

**Statistical and Modelling Techniques to Build Confidence
in the Investigation of Immunology through
Agent-Based Simulation**

Mark Norman Read

Submitted for the degree of Doctor of Philosophy.

University of York,
Department of Computer Science

September 2011

Abstract

For over a century immunological research has made striding advancements that have substantially improved the wellbeing and longevity of the population. There remain, however, many pathogens and diseases to which the immune system and immunologists currently have no answer. Immunological research is as active now in understanding the immune system and its response as it ever has been. Extraordinary technological breakthroughs allow researchers to examine the genes and proteins expressed in individual cells, and the molecules expressed on their cell surfaces. Advancements in imaging techniques allow researchers to observe immune cell interactions at levels of detail never before possible. Yet, despite impressive technological advancements, there is sentiment within immunology that reductionist techniques alone will not be sufficient to explain manner in which system-level behaviours arise from low-level components. Such an understanding is necessary to develop the best strategies for fighting disease. An increasing number of immunologists are combining traditional wet-lab research programs with computational methods that consolidate wet-lab data and attempt to reconstruct the system-level phenomenon observed in the real-world immune system. Such computational techniques facilitate hypothesis formation and exploration, allow for predictive experimentation that is not possible on the real-world, and can guide wet-lab work towards potentially fruitful avenues of research. However, not well appreciated in the literature is the issue of simulation validity: confidence that *in silico* results are representative of the real-world system that simulations attempt to capture. There is considerable uncertainty surrounding many immunological phenomenon, which can complicate the creation of simulations that are themselves highly abstract entities. The results of simulation will not necessarily translate directly into the real-world domain.

This thesis investigates modelling and statistical techniques that establish confidence in *in silico* results being representative of the real-world immune system. A case study in the murine autoimmune disease experimental autoimmune encephalomyelitis (EAE), a model for multiple sclerosis, is undertaken as a means to explore these techniques. The disease is subject to rigorous domain modelling prior to its representation in simulation. Modelling and simulation calibration are performed in close collaboration with a domain expert in EAE. A means of grading simulation executions by the same scale employed in the wet-lab, through examination of an entire mouse, is created. A novel technique is developed whereby the relationship between the accuracy of averaged simulation results and the number of simulation executions sampled in compiling them is established. The completed simulation is explored using a global sensitivity analysis, and a novel robustness analysis technique. These analyses reveal aspects of the simulation that are highly influential to its overall system-level behaviours. The extent to which simulation behaviours are reliant on parameters being assigned particular ranges of values is investigated using robustness analysis. This thesis theorises how these techniques may be considered and applied in the context of domain-specific knowledge to interpret *in silico* results into the original domain. Next, *in silico* experimentation into the nature of EAE is performed, resulting in several predictions concerning the role of particular cells and the spleen in mediating recovery from disease. Lastly, this thesis reflects upon the contribution of these modelling and statistical techniques to provide confidence that simulation results are representative of real-world EAE. A novel approach to development that guides simulations to appropriate levels of abstraction, and demonstrates this to be the case, making extensive use of meta-heuristic search and real-world experimental data is proposed.

Contents

Contents	iii
List of Figures	vii
List of Tables	xii
1 Introduction	1
1.1 Motivation	1
1.1.1 Computational techniques for investigating the immune system .	2
1.1.2 Confidence in simulation	3
1.1.3 The CoSMoS process	5
1.2 Outline of this thesis	7
1.2.1 Thesis structure	8
1.2.2 Thesis contribution	8
1.2.3 Publications relating to this thesis	9
2 Background Immunology and EAE	11
2.1 Introduction to the immune system	11
2.1.1 Organs of the immune system	12
2.1.2 Innate immune system function	13
2.1.3 The adaptive immune system	13
2.1.4 Instigation of the adaptive immune response	18
2.1.5 Terminating the immune response	19
2.1.6 Immune system tolerance	19
2.2 Experimental Autoimmune Encephalomyelitis	20
2.2.1 EAE autoimmunity	21
2.2.2 EAE regulation	23
2.3 Summary	24
3 Modelling and Simulating Immunological Systems	27
3.1 Motivation for modelling and simulating the immune system	27
3.2 Review of modelling and simulation techniques	29
3.2.1 Mathematical approaches	29
3.2.2 Agent-based approaches	30
3.2.3 Hybrid techniques	32
3.3 Integration of wet-lab and <i>in silico</i> techniques	32
3.4 Previous work on modelling EAE	34
3.5 Methodologies and building confidence in simulation	35
3.5.1 Calibration, and the adoption of parameter values	37
3.5.2 Sensitivity analysis	38
3.5.3 Verifying simulation predictions to build confidence	40

3.5.4	Hypothesis evaluation, and acquiring representative results	42
3.5.5	The need for confidence in simulation results	43
3.6	Conclusion	44
4	Domain Model of Experimental Autoimmune Encephalomyelitis	45
4.1	Goal and motivation	45
4.2	Domain model of EAE	46
4.2.1	Relationship between real domain and domain model	47
4.2.2	Modelling <i>perspectives</i>	51
4.2.3	Depicting single-entity dynamics	61
4.2.4	DC and CNS macrophage dynamics	65
4.2.5	Neuron dynamics	68
4.2.6	Cytokine and MBP dynamics	68
4.2.7	Temporal and numerical aspects of EAE	70
4.3	Reflection on the use of UML	71
4.3.1	Reflection on the process of creating a domain model	73
4.3.2	Tackling complexity through perspectives	74
4.3.3	Modelling perspectives as activity diagrams	76
4.3.4	Depicting compounding concurrency with expansion regions	77
4.3.5	Modelling perspectives as sequence diagrams	78
4.3.6	Modelling perspectives as class diagrams	79
4.3.7	Temporal ambiguity on class diagrams	80
4.3.8	Capturing single entity dynamics	81
4.3.9	Summary	83
4.4	Conclusion	84
5	Developing and Calibrating an EAE Simulation Platform	85
5.1	Goal and motivation	85
5.2	Platform model	85
5.2.1	From domain model to platform model	86
5.2.2	Simulation architecture	86
5.2.3	Spatial representation	88
5.2.4	Immunisation mechanism	88
5.2.5	Cellular turnover	90
5.2.6	Periodic migration of DCs	91
5.2.7	Migratory behaviour	91
5.2.8	Probabilistic timing of events	91
5.2.9	Local activation of T cells	93
5.2.10	Polarisation selection by DCs and CD4Th cells	93
5.2.11	Simulation initialisation	93
5.2.12	Cytokine decay	94
5.2.13	T cell specificity	94
5.3	Calibration of simulation platform	94
5.3.1	The role of domain expertise in calibration	96
5.3.2	Calibration process	98
5.3.3	Calibration against multiple data points	98
5.3.4	Calibrated simulation platform dynamics	100
5.3.5	Outstanding discrepancies between simulation platform and <i>in vivo</i> behaviour	100
5.4	Importance of visualisation and state inspection	101
5.5	Conclusion	102

6 Explorative Experimentation	105
6.1 Motivation and goal	105
6.2 Analysis responses	107
6.2.1 <i>In silico</i> EAE Severity Scoring	108
6.3 Compensating for aleatory uncertainty	126
6.4 Determining influence of parameters	130
6.4.1 Experimental procedure	130
6.4.2 Results	132
6.5 Robustness analysis of parameters	140
6.5.1 Experimental procedure	141
6.5.2 Results	142
6.6 Qualifying the significance of simulation results	153
6.7 Conclusion	156
7 Projective Experimentation	159
7.1 Goal and motivation	159
7.2 Elucidation of baseline behaviour	160
7.2.1 Experimental methodology	160
7.2.2 Results	161
7.2.3 Summary	172
7.3 Investigation of CD8Treg mediated regulation	177
7.3.1 Experimental methodology	177
7.3.2 Results	179
7.3.3 Summary	188
7.4 <i>In silico</i> splenectomy	191
7.4.1 Experimental methodology	194
7.4.2 Results	194
7.4.3 Summary	206
7.5 Conclusion	207
8 Discussion	209
8.1 Reflections on establishing confidence in simulation	209
8.1.1 Explicit domain modelling	210
8.1.2 Calibration of simulation mechanics and parameters	210
8.1.3 Demonstrably representative results	211
8.1.4 Interpretation of <i>in silico</i> results in term of the target domain	212
8.1.5 Simulation augmentation and confidence in existing results	213
8.1.6 A meta-heuristic framework for quantifying the accuracy of simulation mechanics	214
8.1.7 Confidence objectivity	216
8.2 Thesis summary, and contribution	218
8.3 Further work	221
8.3.1 Further investigation of EAE	221
8.3.2 Capturing complex systems with the UML	222
8.3.3 Meta-heuristic guided calibration and simulation development	223
8.3.4 Instantiation and further investigation of robustness analysis techniques	225
8.4 Concluding remarks	225
A Primer on UML, and Additional Materials for EAE Domain Model	227
A.1 UML notations	227

A.1.1	Class diagrams	227
A.1.2	Activity diagrams	228
A.1.3	State machine diagrams	229
A.2	Additional domain model diagrams	230
A.2.1	Class diagram of <i>establishment of regulation</i> perspective	230
A.2.2	Class diagram of <i>type 2 deviation of the autoimmune response</i> perspective	232
B	Additional Materials on Calibrating ARTIMMUS	235
B.1	Simulation platform parameters	235
B.2	The calibration of simulation parameters	241
C	Supporting Material for Statistical Techniques	245
C.1	Calculating the <i>A</i> test	245
C.2	Global sensitivity analysis	245
C.3	Robustness sensitivity analysis	264
Glossary of Terms		299
Bibliography		305

List of Figures

1.1	The CoSMoS process for performing simulation-based investigation of complex systems.	5
2.1	The events leading to the activation of lymphocyte cells.	16
2.2	The progression of EAE in individual mice amongst four experimental groups, with five mice in each group.	22
2.3	Abstract depiction of the major cell types involved in EAE autoimmunity and its associated recovery.	22
4.1	An <i>expected behaviours</i> diagram depicting the phenomena observed in the real domain, and the behaviours manifesting from cellular interactions believed to be responsibly for them.	48
4.2	The spatial components of the domain model, and the manner in which the cells of the domain model may migrate between them.	50
4.3	The types of relationship used in activity diagrams.	52
4.4	UML activity diagram depicting the cellular interactions and events that lead to neuronal apoptosis in the CNS following immunisation for EAE.	54
4.5	UML activity diagram depicting the cellular interactions and events that lead to the self-perpetuation of autoimmunity following neuronal apoptosis resulting from immunisation for EAE.	56
4.6	UML class diagram depicting the relationships between entities of the domain model involved in the establishment and perpetuation of the autoimmune response.	57
4.7	UML activity diagram depicting the cellular interactions and events that lead to the instigation and perpetuation of the regulatory immune response.	59
4.8	UML activity diagram depicting the cellular interactions and events that lead to a type 2 deviation of the immune response.	60
4.9	State machine diagram depicting the dynamics of CD4Th cells.	63
4.10	State machine diagram depicting the dynamics of CD4Treg cells.	64
4.11	State machine diagram depicting the dynamics of CD8Treg cells.	65
4.12	State machine diagram depicting the dynamics of dendritic cells.	66
4.13	State machine diagram depicting the dynamics of CNS macrophages.	67
4.14	State machine diagram depicting the dynamics of neurons.	68
4.15	State machine diagram depicting the dynamics of myelin basic protein (MBP).	68
4.16	State machine diagram depicting the dynamics of type 1 cytokine, and its influence on other cells of the domain model.	69
4.17	State machine diagram depicting the dynamics of type 2 cytokine, and its influence on other cells of the domain model.	70
4.18	State machine diagram depicting the dynamics of TNF- α , and its influence on other cells of the domain model.	70

4.19	UML class diagram attempting to capture every entity and relationship within the domain model on a single diagram.	75
4.20	An example sequence diagram showing the interactions between participants A, B and C over time.	76
4.21	Examples of activity diagram expansion regions, and their potential application in depicting compounding concurrency.	78
4.22	An example UML sequence diagram, depicting the messages passed between participants A, B and C.	79
4.23	The use of UML class diagram composition relationships depicting mutually exclusive expression of MHC:peptide complexes on cells.	79
4.24	Select relationships between entities in the EAE domain model.	82
5.1	Class diagram depicting organisation of classes in the simulation platform in terms of inheritance hierarchies and interface implementations. The majority of associations are not shown, in aid of readability.	87
5.2	The lattice grid based spatial representation of ARTIMMUS.	89
5.3	ARTIMMUS's immunization mechanism, and how it is parameterised. . . .	89
5.4	The calculation of probabilities that a migrating cell, such as a T cell, will move either downwards, upwards, or remain at its current vertical level in a particular time step.	92
5.5	Depiction of how the absolute time of events is determined through a normal distribution of possibilities.	93
5.6	Examples of clinical data pertaining to the progression of EAE in mice. . . .	97
5.7	Abstract depiction of how calibration against multiple experimental scenarios can lead to identification of simulation abstractions that better approximate the dynamics of the real-world domain.	99
5.8	Effector T cell population dynamics of the calibrated ARTIMMUS simulation. .	100
5.9	Screenshot of the ARTIMMUS simulation.	102
6.1	The eleven responses derived from simulation behaviour, displayed on graphs depicting baseline T cell dynamics and mean EAE severity over time. .	107
6.2	Example <i>in vivo</i> data on the progression of EAE in individual mice under conditions of physiological recovery, regulation temporarily disabled. . . .	109
6.3	Demonstration of how threshold values for relating the number of neurons apoptosed in an hour of simulation time to EAE severity scores is achieved. .	112
6.4	The assignment of <i>in silico</i> severity scores by thresholding on the rate of neuronal apoptosis, for conditions of physiological recovery and disabled regulation.	113
6.5	Variation in the number of neurons apoptosed per hour of simulation time, and the effect of different sizes of smoothing window when smoothing data. Data is of an example simulation execution under condition of physiological recovery.	115
6.6	EAE progression of 7 mice of the 'B10.PL' strain, having been immunized for EAE. Data obtained from the Kumar lab.	116
6.7	Example EAE progressions for simulations under conditions of physiological recovery and disabled regulation. The mapping of neuronal apoptosis rates to EAE severity was derived using no smoothing of neuronal apoptosis data. The resultant EAE progressions are unrealistic with respect to the frequency with which EAE scores change, as demonstrated by the high powers of low period components in the Fourier transforms.	118

6.8	Example EAE progressions for simulations under conditions of physiological recovery and disabled regulation. The mapping of neuronal apoptosis rates to EAE severity was derived using a smoothing window size of 157 hours. The resultant EAE progressions closely resemble the frequency analysis characteristics of <i>in vivo</i> EAE.	119
6.9	The sum of squared normalised difference measures indicating how well various smoothing window sizes align frequency and magnitude of changes in <i>in silico</i> EAE scores over time with those observed <i>in vivo</i>	120
6.10	The sum of squared normalised differences measure indicating the alignment of threshold values relating neurons apoptosed per hour to EAE severity scores derived under conditions of physiological recovery and disabled regulation for various smoothing window sizes.	120
6.11	The proportion of EAE progressions reaching particular maximum severity scores, for various smoothing window sizes.	122
6.12	The sum of squared difference measures indicating the separation between the proportion of <i>in silico</i> and <i>in vivo</i> individuals reaching particular maximum EAE severity scores, obtained under various smoothing window sizes. Data obtained under conditions of physiological recovery and disabled regulation are shown.	123
6.13	Determining that a smoothing window size of 157 hours best satisfies the five criteria for the most appropriate mapping of neuronal apoptosis rates to <i>in silico</i> EAE severity scores.	124
6.14	The derivation of threshold values for relating rates of neuronal apoptosis to <i>in silico</i> EAE severity scores, and the proportions of simulation reaching particular maximum EAE severity scores, for the final smoothing window size of 157 hours.	125
6.15	The effect of aleatory uncertainty on the results of <i>A</i> test analysis for eleven responses, for various sample sizes. The first set of simulation executions is compared with each of the remaining nineteen in obtaining <i>A</i> test scores.	128
6.16	Maximum <i>A</i> test scores across twenty sets of experimentation, using the same parameter values in all cases, for various responses and sample sizes used in generating distributions.	129
6.17	An example latin hypercube design, obtaining ten samples across two parameters.	131
6.18	Comparison of T cell dynamics when <i>CD4THelper_diff00</i> is set to its default value of 5%, and 100%.	140
6.19	Robustness analysis of the <i>TCell_AICDMean</i> parameter.	143
6.21	The mean severity of EAE experienced under parameter values of 0 and 60 hours for <i>TCell_AICDStdDev</i>	153
6.22	The robustness indices arising from robustness analysis can be contrasted with biologically-plausible ranges of parameter values. This can indicate the degree to which simulation results are explained by uncertainty in the domain, and hence how much confidence one may place in their being representative of the real-world domain rather than underspecified parameter values.	155
6.23	The range of simulation behaviours which may be assumed representative of the domain, rather than the results of uncertainty concerning exact biological figures, may be calculated by constructing a latin hypercube around the ranges of biologically plausible values.	156
7.1	Effector T cell population dynamics.	161

7.2	CD4Th population states of activation.	163
7.3	CD4Treg population states of activation.	163
7.4	CD8Treg population states of activation.	163
7.5	Cumulative count CD4Th cells primed in each compartment.	165
7.6	Cumulative count of CD4Treg cells primed in each compartment.	165
7.7	Cumulative count of CD8Treg cells primed in each compartment.	165
7.8	States of APC activation in the SLO.	167
7.9	States of APC activation in the CLN.	167
7.10	States of APC activation in the Spleen.	168
7.11	States of APC activation in the CNS.	168
7.12	Dendritic cell polarizations in the CLN.	169
7.13	Compartments in which CD4Th1 apoptosis induction by CD8Treg cells takes place.	169
7.14	Neurons apoptosed per hour.	170
7.15	Screenshot of ARTIMMUS, showing the CNS compartment at day 18 post-immunization.	171
7.16	Mean progression of EAE.	171
7.17	Mean distributions of EAE progression. The proportion of simulation runs experiencing each level of EAE and a cumulative representation of this data are shown.	172
7.18	Progression of EAE and the effector T cell population dynamics for 3 example simulation executions, shown one per row, with regulatory activity disabled.	173
7.19	Progression of EAE and the effector T cell population dynamics for 2 further example simulation executions with regulatory activity disabled.	174
7.20	Analyses of number of EAE episodes and remissions, and their durations for physiological recovery and regulation disabled groups.	176
7.21	Median effector T cell population dynamics and mean progression of EAE, derived from <i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i> values of 100%, 5% and 0%.	180
7.22	Neurons apoptosed per hour, and cumulative CD4Th1 cells apoptosed over time, derived using <i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i> values of 100%, 5% and 0%.	181
7.23	Box and whisker plot of <i>CD4Th1</i> at 40 days, and <i>max EAE</i> and <i>EAE at 40 days</i> responses for various parameter values of <i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i>	182
7.24	Median effector T cell population dynamics and mean progression of EAE, derived from <i>Th1Polarization_mhcUnExpressionDelayMean</i> values of 8, 3 and 0 hours.	184
7.25	Neurons apoptosed per hour, and cumulative CD4Th1 cells apoptosed over time, derived using <i>Th1Polarization_mhcUnExpressionDelayMean</i> values of 8, 3 and 0 hours.	185
7.26	Box and whisker plot of <i>CD4Th1</i> @ 40 days, and the mean <i>max EAE</i> and mean <i>EAE at 40 days</i> responses, obtained for various values of <i>Th1Polarization_mhcUnExpressionDelayMean</i>	186
7.27	Cumulative count of the number of effector Th1 cells induced into apoptosis by the CD8Treg population, for various times required to migrate through the circulatory compartment.	189
7.28	Cumulative count of effector Th1 cells induced into apoptosis by the CD8Treg populations, by compartment, for various times required to migrate through the circulatory compartment.	189

7.29	The number of effector CD4Th1 and CD4Th2 cells in the CNS compartment over time, for various times required to migrate through the circulatory compartment.	190
7.30	Cumulative count of CD4Th cells primed by compartment, for various times required for a T cell to migrate through the circulatory system.	190
7.31	Cumulative count of CD4Treg cells primed by compartment, for various times required for a T cell to migrate through the circulatory system.	190
7.32	Cumulative count of CD8Treg cells primed by compartment, for various times required for a T cell to migrate through the circulatory system.	191
7.33	Magnitude of effect measures of the difference in EAE severities amongst control and splenectomy groups at various times following immunization.	196
7.34	Mean progression of EAE under control and splenectomy experiments.	198
7.35	Cumulative counts of the proportion of simulations experiencing each degree of EAE severity, for control and splenectomy experiments.	198
7.36	Long-term cumulative counts of the proportion of simulations experiencing each degree of EAE severity, for control and splenectomy experiments.	198
7.37	Analyses of number of EAE incidences and remissions, and their durations for control and splenectomy experimental groups.	199
7.38	Example EAE progressions of individual splenectomised simulation executions.	202
7.39	The median system wide T cell dynamics, obtained from 500 simulation executions, showing control and splenectomy experiments.	203
7.40	Cumulative count CD4Th cells primed in each compartment, for control and splenectomy experiments.	203
7.41	Cumulative count CD4Treg cells primed in each compartment, for control and splenectomy experiments.	204
7.42	Cumulative count CD8Treg cells primed in each compartment, for control and splenectomy experiments.	204
7.43	States of APC activation in the CNS, for control and splenectomy experiments.	204
7.44	States of APC activation in the CLN, for control and splenectomy experiments.	205
7.45	Polarizations of DCs in the CLN, for control and splenectomy experiments.	205
8.1	A meta-heuristic based technique for quantifying how well a simulation captures its target domain.	217
8.2	Abstract representation of the cells and their interactions involved in the two regulatory pathways of EAE: the induction of apoptosis in CD4Th1 cells by CD8Tregs, and the CD200:CD200R mediated suppression of DC priming capacity by CD8Treg cells.	223
A.1	Example UML class diagram.	228
A.2	Example UML activity diagram.	229
A.3	Example UML state machine diagram.	229
A.4	UML class diagram depicting the relationships between entities of the domain model involved in the instigation and perpetuation of the regulatory immune response.	231
A.5	UML class diagram depicting the relationships between entities of the domain model that lead to the type 2 deviation of the autoimmune response.	233

List of Tables

4.1	Temporal aspects to EAE.	72
6.1	<i>In vivo</i> data showing the maximum scores reached by different experimental groups of mice.	110
6.2	The magnitude of effect size indicated by <i>A</i> test scores.	126
6.3	The total influence of each parameter on simulation responses, measured as a sum of ranks for each response in turn.	133
6.4	Summary of parameter robustness indexes	145
6.5	Robustness indexes for compartment dimensions.	147
6.6	Robustness indexes for initial cell numbers	147
6.7	Robustness indexes for T cell-APC interactions.	147
6.8	Robustness indexes for cytokine secretion and decay.	148
6.9	Robustness indexes for standard deviation parameters.	148
7.1	<i>A</i> test scores indicating the difference between total, and compartmental, regulatory-induced Th1 apoptosis for different migratory times spent by T cells in the circulatory system.	188
7.2	Acute, chronic or relapsing EAE in rats as a function of age or splenectomy. Data replicated from [Ben-Nun <i>et al.</i> 1980], table 2.	193
7.3	The proportions of simulations experiencing each possible maximum EAE severity score, for control and splenectomy groups.	195
7.4	The <i>A</i> test scores indicating magnitude of difference in response distributions between splenectomy and control experimental groups.	200
B.1	The standard parameters of the ARTIMMUS simulation platform.	236
C.1	Partial rank correlation coefficient between parameter values and the <i>CD4Th1 Max</i> response, and the associated p value	247
C.2	Partial rank correlation coefficient between parameter values and the <i>CD4Th1 Max Time</i> response, and the associated p value	248
C.3	Partial rank correlation coefficient between parameter values and the <i>CD4Th2 Max</i> response, and the associated p value	250
C.4	Partial rank correlation coefficient between parameter values and the <i>CD4Th2 Max Time</i> response, and the associated p value	251
C.5	Partial rank correlation coefficient between parameter values and the <i>CD4Treg Max</i> response, and the associated p value	253
C.6	Partial rank correlation coefficient between parameter values and the <i>CD4Treg Max Time</i> response, and the associated p value	254
C.7	Partial rank correlation coefficient between parameter values and the <i>CD8Treg Max</i> response, and the associated p value	256

C.8	Partial rank correlation coefficient between parameter values and the <i>CD8Treg Max Time</i> response, and the associated p value	257
C.9	Partial rank correlation coefficient between parameter values and the <i>CD4Th1 at 40 days</i> response, and the associated p value	259
C.10	Partial rank correlation coefficient between parameter values and the <i>Max EAE</i> response, and the associated p value	260
C.11	Partial rank correlation coefficient between parameter values and the <i>EAE at 40 days</i> response, and the associated p value	262
C.12	Robustness indexes for <i>CD4Th1 Max</i> response	265
C.13	Robustness indexes for <i>CD4Th1 Max Time</i> response	267
C.14	Robustness indexes for <i>CD4Th2 Max</i> response	270
C.15	Robustness indexes for <i>CD4Th2 Max Time</i> response	272
C.16	Robustness indexes for <i>CD4Treg Max</i> response	275
C.17	Robustness indexes for <i>CD4Treg Max Time</i> response	277
C.18	Robustness indexes for <i>CD8Treg Max</i> response	280
C.19	Robustness indexes for <i>CD8Treg Max Time</i> response	283
C.20	Robustness indexes for <i>CD4Th1at40d</i> response	285
C.21	Robustness indexes for <i>Max EAE</i> response	288
C.22	Robustness indexes for <i>EAE at 40d</i> response	290
C.23	Robustness indexes for <i>Max EAE A Test</i> response	293
C.24	Robustness indexes for <i>EAE at 40d A Test</i> response	295

Acknowledgements

Undertaking a PhD and conducting research can be immensely enjoyable, and hugely rewarding. However, it can also be a very difficult and demanding task, especially when it comes to write a thesis. I would not be in the position I find myself now without the support of quite a lot of people, to whom I owe my sincere thanks:

To Jon and Paul for their excellent supervision and advice over the years.

To Vipin for his part in our collaboration, ever patient with a non-immunologist asking questions to which the answers do not yet exist.

To Susan and the NSC group for sending me on countless trips to the wonderful city of La Jolla to conduct my collaboration, and to my colleagues in the CoSMoS project and the NSC group for many engaging discussions.

To Louis, Dan and Tara, without whom I doubt I would have my sanity intact at the end of this thesis.

To my parents, Norman and Marijke, for all they have done to get me to where I am now.

Lastly, and certainly not least, to Helena for years of love, belief, support, patience and understanding. Writing a thesis is very hard, you made it a lot easier.

Author Declaration

The work presented here is the author's own. Some of this work has already been presented in [Read *et al.* 2009b], [Read *et al.* 2009a] and [Read *et al.* 2011]. The experience of conducting this work has formed a case study in [Polack *et al.* 2010]. The simulation developed in this thesis has been extended and used to conduct experimentation outside the scope of thesis, reported in [Williams *et al.* 2011].

Chapter 1

Introduction

This introductory chapter presents the motivation for the work conducted in this thesis, and an overview of how it is guided and structured. Firstly, section 1.1 motivates the investigation of the immune system through computational techniques, and explores the issue of confidence in simulation results. This is followed by section 1.2 which details the aim of the thesis, the research objectives that guide it towards this aim, its content and structure, and lastly its contribution.

1.1 Motivation

The *immune system* is the collective term given to the cells, molecules, processes, tissues and organs that are charged with maintaining the health of an organism. It is responsible for mounting defense against pathogens and for clearing tumorous cells. Whilst the immune system may generally be considered to be effective in fulfilling these functions, it is not perfect, as evidenced by diseases such as leishmaniasis or tuberculosis that often cannot be completely cleared from the body [Manabe & Bishai 2000, Kaye & Scott 2011], the development of life-threatening cancers, or autoimmune diseases where the immune system attacks the host.

Immunology is the study of the immune system. For over a century immunologists have been investigating the operation of the immune system, and how it may be manipulated to improve wellbeing and longevity. The field has lead to the discovery of vaccines that have rendered many potentially lethal pathogens relatively harmless, and countless drugs to improve the immune system's ability to fulfill its function.

Whilst acknowledging the remarkable successes that immunology has had to date, there are still many diseases, such as AIDS or multiple sclerosis, to which the immune system can mount no effective response, and to which the biomedical industry has found no cure. Immunology is as active a field as it ever has been, and recent times have seen phenomenal advances in the technologies used to study the immune system. Researchers are able to detect the molecules being expressed on or secreted by individual cells, the proteins being created within those cells and the genes that are actively transcribed by them. Advances in *in vivo* imaging techniques allow researchers to observe the interactions between individual cells at astonishingly small scales. Geneticists have mapped the genome, and a full mapping of the body's proteins is not far

away [Cohen 2007a]. Collectively, these technologies have generated vast quantities of experimental data.

Yet, for all the unique insight that these technological advancements can deliver, no unified theory of how the immune system operates has emerged [Cohen 2007a]. The top researchers in the field have fundamentally different views regarding the manner in which the immune system operates to direct its immensely destructive ability towards targets that cause harm to the host, but not the host itself¹. As Cohen writes, “The more data we have access to, the more confused we have become” [Cohen 2007a]. The immune system is too complex to be understood through purely reductionist techniques, looking to explain higher-level behaviours by examining in ever more detail its lower-level components. Many complex systems, such as the immune system, exhibit system-level behaviours that cannot be intuitively understood by purely examining their low-level individual components in turn. The immune system, like many biological systems, is rife with the key features of complexity: vast numbers of low-level components, each of which is capable of performing many functions, and each of those many functions can be performed by many types of component [Cohen 2004, Edelman & Gally 2001]. One cannot attribute a single function to a single component, and a wider systems-level overview of how the components interact, support and suppress one another’s function is required to understand its operation.

1.1.1 Computational techniques for investigating the immune system

There is a growing sentiment in immunology that computational methods can aid investigation of the immune system [Germain *et al.* 2011, Cohen 2007a]. The benefits that these methods purport to bring to immunology are echoed by many within the field, and an increasing number of immunological publications report the use of computational methods in their investigations [Kleinsteine 2008]. To summarise, the benefits of a complementary computational and wet-lab approach to immunology are as follows:

- Simulations can capture and consolidate large quantities of data from a wide variety of sources, and provide a system-level overview of what that data represents [Cohen 2007a, Bauer *et al.* 2009, Germain *et al.* 2011].
- Simulations, being computer code, are more amenable to designing, conducting and collecting data from experimentation than the natural system [Forrest & Beauchemin 2007]. Real-world experimentation has physical and ethical limitations that do not computer simulations.
- Having integrated established experimental data, simulation offers a platform for the formulation and evaluation of hypotheses concerning system operation. Such hypotheses can then be tested for consistency against established data [Chakraborty *et al.* 2003, Kirschner & Linderman 2009, Chakraborty & Das 2010].
- Simulation can help in directing wet-lab experimentation, permitting relatively cheap preliminary investigations that can point to interesting aspects of the system under study. Areas where biological knowledge is lacking can be identified,

¹In 2000, the journal *Seminars in Immunology* dedicated an entire issue (volume 12, issue 3) to facilitate this very debate amongst some of the most prominent immunologists in the field. The articles contained in this issue demonstrate the widely diverse theories of immune system function held within the field.

and revealing experiments to discern between competing theories can be designed [Cohen 2007a, Chakraborty & Das 2010, Bauer *et al.* 2009].

The field concerned with the exploration of the immune system through modelling and simulation is referred to as *computational immunology*. A variety of computational methods have been employed in investigating immunological systems. This thesis focuses on the agent-based paradigm [An *et al.* 2009]. Agent-based simulation is characterised by an explicit representation of individual entities in the system, such as cells or molecules. Agent-based simulation captures the various states that particular types of entity in the system may exist in, and the temporal events, interactions and probabilities that promote changes in state. The state that each entity of the system exists in is explicitly maintained in an agent-based system. These systems typically encompass an explicit representation of space and time, within which system components exist, move and interact.

Agent-based simulations allow computational immunologists to perform experimentation at the level of the cell, or the molecule, and observe the system-level results. In the highly controllable environment of a simulation one can engineer and perform experiments that are not possible in the wet-lab, such as disabling only one of many functions that a particular molecule performs, and observe the effects on any aspect of simulation behaviour. Computational immunology represents a *constructionist* approach to immunological investigation, which may form a valuable complement to the more traditional reductionist wet-lab approaches.

1.1.2 Confidence in simulation

Though computational methods hold much potential in complementing traditional wet-lab techniques in the advancement of immunology, this can only be truly accomplished if it can be argued that the simulation is a faithful representation of its target domain. This is especially the case where simulation is used to perform investigations for which the real-world system's behaviour is unknown. In fact, simulations are never a completely faithful representation of their target domains, for reasons explored below: they are highly abstract entities. They can, however, be fit for purpose: one can have confidence that the results of simulation are satisfactorily representative of the real-world domain with respect to the particular *in silico* experimentation being conducted. If simulation is to be used to complement and guide wet-lab experimentation, by facilitating hypothesis exploration and providing predictions of the real-world system, it is necessary to establish confidence that the simulation is fit for purpose. This is not a Boolean all-or-nothing quality. Rather, confidence can be established to varying degrees, since, as noted above, simulation is never an absolutely faithful representation of the real-world domain. This section explores the separation of simulation and real-world domain, and the various aspects of simulation-based investigation in which one must establish confidence.

The immune system is a highly complex system; new cell types and their interactions are continually being identified, exact rates and quantities involved in processes are often unknown² [An *et al.* 2009]. These uncertainties, or simply complete absences of information, can considerably complicate the construction of immunological simulations.

Further to this, simulations are highly abstract representations of their target domains. They typically represent only a small subset of the real-world system's compo-

²For example, [Bevan 2006b, Prlic *et al.* 2007] explore the controversy concerning the role of CD4Th cells in the cytotoxic T cell response.

nents, and at a substantially reduced number. One might hypothesise that the components of the real-world system represented in simulation are the only ones critical to its operation. However, this cannot be known for certain, and in all likelihood those components and interactions of the real system not explicitly represented in simulation do have some influence over its operation. As such, the components of a simulation must compensate for the activities of those that are not represented. In dealing with underspecified system components, and the abstractions made in simulation, simulations require *calibration*: the adjustment of simulation components, interactions and parameters such that simulation dynamics reflect those observed in the real-world system.

Herein lies two aspects of simulation-based investigation where one must provide confidence. Firstly, confidence that the simulation's capture of the cells, processes and interactions present in the real domain is satisfactorily accurate, given the nature of the experimentation to be performed on it. This thesis refers to these aspects of simulation as its *mechanics*. Secondly, confidence that the simulation's *parameters* are assigned appropriate values: that the mechanics have associated with them the correct rates, probabilities, quantities and temporal behaviours. The conceptual separation between simulation mechanics and parameters is made because it is believed that a simulation with inappropriate mechanics can still be parameterised to replicate certain aspects of the real-world system under particular circumstances. Both mechanics and parameters must be appropriate if simulation results are to be representative of the target domain.

Agent-based simulations can be made stochastic in nature: two simulation executions given the exact same input parameters can often reveal different results. This feature of agent-based simulations can be desirable, it is a feature of the real-world system also. It can, however, complicate *in silico* experimentation. Computational immunologists must establish confidence that the results of *in silico* experimentation are representative of the simulation's dynamics and the experimental procedure, rather than stochastic variation. The typical approach here is to sample many simulation executions, and generate averaged results.

The last aspect of establishing confidence in simulation results being representative of the real-world system concerns the interpretation of *in silico* results into the original domain. Owing to the abstract nature of simulation, and the considerable uncertainty concerning many immunological phenomenon, *in silico* results are not directly interpretable with respect to their real-world domain. For example, simulations may represent only a few thousand cells, whereas the real-world immune system contains millions. One must establish the relationship between simulation metrics, and those of the real-world domain. Following this example, if only a few thousand cells are represented *in silico*, what does an experiment that results in the generation of an extra few hundred mean in terms of the original domain? Secondly, one must provide confidence that *in silico* results, such as novel predictions, are not explained merely though underspecified parameter values. Underspecified values may arise from incomplete, ambiguous or contradictory research in the domain. Simulation predictions that hold only under a narrow subset of the biologically plausible parameter values should not necessarily be assumed representative of the real system.

Some of these issues concerning the establishment of confidence that simulation results are representative of the real-world domain have been recognised by the CoSMoS project³, an ongoing research project that seeks to develop support for the investigation of complex systems [Andrews *et al.* 2010]. Their work is reviewed in the following

³The Complex Systems Modelling an Simulation infrastructure (CoSMoS) project, EPSRC grants EP/E053505/1 and EP/049419/1, www.cosmos-research.org, 2007-2011.

section.

1.1.3 The CoSMoS process




Figure 1.1: The CoSMoS process for performing simulation-based investigation of complex systems. The arrows represent the flow of information, not the strict flow of activity. Figure represents an early version of the CoSMoS process, and is obtained from personal communication with members of the CoSMoS project [Andrews *et al.* 2010].

The CoSMoS project is an EPSRC funded project which seeks to build generic tools and techniques to support the modelling, simulation and analysis of complex systems [Andrews *et al.* 2010]. The project is case study driven, drawing on the experience of a range of case studies in analysing and building complex systems, covering a diverse range of domains, such as immunology, sociology and swarm robotics [Polack *et al.* 2010, Owen *et al.* 2010]. The work of the present thesis has been undertaken in parallel with the CoSMoS project, and has served as one such case study.

The CoSMoS project's approach to complex system simulation is captured in the *CoSMoS process* [Andrews *et al.* 2010]. It advocates rigorous modelling activities that underpin simulation-based investigation. The process has remained in development throughout the present thesis, undergoing refinements and additions in reflection of the ongoing case studies that inform it. As such, the version of the process presented here reflects its state around the commencement of this work, and is the version that has been employed in this thesis⁴.

The CoSMoS process is denoted in figure 1.1. It encompasses five artifacts: the domain, the domain model, the platform model, the simulation platform, and the results model. The domain is that which is to be represented in simulation, the remainder are engineered artifacts to support simulation. Each of these is considered in turn.

The *domain* represents the complex system of interest, in this case some aspect of the immune system. Understanding of the domain's nature is achieved through experimentation and observation. This understanding is captured in the *domain model*, which must present a coherent and consistent view of the domain. As such, domain modelling can highlight areas of inconsistent or incomplete domain-specific information, and may lead to the adoption of assumptions. It focusses on scientific understanding, and is free of any potential simulation implementation-specific details. It captures all entities of the system, and how their collective behaviours manifest in system-level

⁴The names of some of the process' artifacts have been updated to their more recent titles, since it is felt that they better reflect the nature of the artifacts. The structure and semantics of the process, however, is unchanged.

phenomenon. The process makes no claims over the sorts of modelling techniques that are best suited to modelling any particular domain.

The *platform model* serves as an implementation-specific specification for the simulation platform, and is derived from the domain model. All explicit representation of system-level behaviours are removed; these system-level ‘emergent’ behaviours should arise from low-level component interactions once the simulation is executed. They should not be directly programmed into the simulation. If simulation execution does not reveal these properties, then either the simulation does not correctly capture the domain model, or the domain model does not capture all the relevant aspects of the domain. Some aspects of the domain may be simplified, for example variable levels of molecular expression by cells may be abstracted into a dichotomous state of expression or not, in which case these simplifications constitute further assumptions made of the domain and must be recorded. Implementation specific additions may be made in the platform model, which could include user interfaces or data logging tools.

The completion of the platform model is followed by the construction of the *simulation platform*, which may be executed and used to perform experimentation. The process makes no claims over what modelling paradigm the simulation must belong to, or which programming language it is implemented in.

The *results model* encapsulates the understanding that arises from experimentation with the simulation platform. It is to the simulation platform what the domain model is to the domain. The results model may be contrasted with the domain model for the purposes of validation. Similar behaviours in the results model to those recorded in the domain model affords confidence that the simulation is a fair representation of the domain being modelled. Further, the results model may reveal results of simulation that constitute predictions made of the original domain.

There are two further points to be made of the CoSMoS process. Any of the artifacts identified above, except the simulation, may be implicit, rather than full physical instantiations, subject to issues of impact and criticality of simulation results. Secondly, the process may be highly iterative in nature. Early transitions through the process may yield simulation platforms that do not adequately capture the domain, in which case further investigation of the domain and amendment of process artifacts may be performed. Should the simulation platform satisfactorily replicate the observations made of the domain, it may be extended to explore novel scientific questions, hence motivating further iterations of the process. In this manner, the CoSMoS process facilitates open ended scientific research.

CoSMoS advocates an interdisciplinary approach to simulation-based investigation of complex systems. An immunologist is unlikely to have the computational and programming skills required to create simulations and perform experimentation on them. Likewise, computer scientists are unlikely to have the biological background necessary to fully comprehend the research conducted into the immune system. In CoSMoS nomenclature, the constructor of a simulation is referred to as the *modeller*, whereas an expert in the domain being investigated is referred to as the *domain expert*.

The modelling activities of the process serve two purposes. Firstly, they serve to guide interaction with the domain expert, providing an explicit agreed upon scope for the simulation based work. Secondly, by explicitly documenting the simulation’s representation of the domain and recording assumptions made in deriving it, simulation-based experimentation is made more transparent; models are more comprehensible than source code. The models aid in establishing confidence that a simulation is an appropriate representation of the domain, the domain expert verifies that the scientific understanding captured in the domain model is appropriate and reasonable. Hence, CoSMoS seeks to address the first aspect of establishing confidence in *in silico* results:

that the simulation's mechanics are a faithful representation of the real-world system.

1.2 Outline of this thesis

Thus far, this chapter has motivated the investigation of the immune system through simulation. It has identified the need to establish confidence that the results of *in silico* experimentation are representative of the real-world system, and detailed four aspects of simulation based investigation where this is must be addressed. These are:

- Confidence that the simulation's capture of cells, processes and interactions of the target domain is satisfactorily accurate, given the investigations performed on it.
- Confidence that the simulation is correctly parameterised.
- Confidence that *in silico* results are representative of the simulation's dynamics, and not merely the result of stochastic variation in the simulation.
- Understanding of the implications of *in silico* results in terms of the original domain; that simulation results are not merely the result of underspecified parameter values owing to incomplete, ambiguous or contradictory research in the domain.

This thesis concerns statistical and modelling techniques to support simulation-based investigation of the immune system, specifically the four aspects of confidence in simulation results identified above. The modelling aspect of this work adheres to the CoSMoS process reviewed above. This scope is reflected in the thesis aim:

To apply and develop statistical and modelling techniques that aid confidence in agent-based simulations of immunology, specifically experimental autoimmune encephalomyelitis.

In order to investigate these techniques this thesis conducts a case study in modelling and simulating a murine autoimmune disease, experimental autoimmune encephalomyelitis (EAE), which has many parallels with multiple sclerosis. EAE was selected because of its considerable complexity; it is felt that if the methods established in this thesis suffice in providing confidence with respect to this case study, then they will hold for many other immunological domains. The modelling and simulation work of this thesis is aligned with a particular wet-lab laboratory, that of Dr. Vipin Kumar at the Torrey Pines Institute for Molecular Studies, in San Diego. Dr. Kumar is an expert in EAE, and assumes the role of the *domain expert* in this thesis. The author assumes the role of *modeller*.

The following research objectives will guide the work of this thesis in achieving its aim:

Obj 1: Explore the role of explicit domain modelling in the EAE case study.

Obj 2: Create an agent-based simulation of EAE.

Obj 3: Investigate and develop techniques for calibrating agent-based simulations.

Obj 4: Perform novel *in silico* experimentation using the agent-based simulation of EAE.

Obj 5: Develop and apply statistical techniques for interpreting *in silico* results in the context of the target domain, EAE.

1.2.1 Thesis structure

This thesis addresses the above research objectives in seven chapters, most of which relate to aspects of the CoSMoS process. Their content, relation to the CoSMoS process, and the contribution to fulfilling the research objectives, are as follows:

Chapter 2 provides background material on immunology and EAE, necessary for understanding the *in silico* experimentation conducted in this thesis. This chapter represents the domain of the CoSMoS process.

Chapter 3 reviews the literature on modelling and simulating the immune system, and provides a critique on how existing works have addressed the question of simulation validity and confidence in *in silico* results. It has no direct relationship to any CoSMoS process artifact.

Chapter 4 details the construction of a domain model of EAE. EAE is modelled using the unified modelling language (UML). Though originally intended for the specification of software systems, it is felt that the various diagrammatic notations encompassed with the UML may be suitable for agent-based modelling of immune system domains. The chapter provides critical reflection on the ability of UML to represent EAE, and hence contributes to research objective 1.

Chapter 5 reports on the construction of a simulation of EAE. It represents the platform model and the simulation artifacts of the CoSMoS process. The chapter details the manner in which the EAE domain model is interpreted as a specification for a simulation, and reports the procedure used to develop the simulation to an appropriate level of abstraction and parameterise it. The chapter addresses research objectives 2 and 3.

Chapter 6 considers techniques for extracting representative data from the simulation, and interpreting it in terms of the original domain. It performs an initial exploration of simulation behaviours through application of a statistical technique that determines the influence of the simulations various parameters on its behaviour. As such, the chapter represents part of the CoSMoS results model, and investigates its relationship to the domain. Contributions to research objectives 4 and 5 are made.

Chapter 7 performs novel *in silico* experimentation into EAE through simulation, hence addressing research objective 4. It also represents the results model of the CoSMoS process.

Chapter 8 provides critical reflection on the work conducted in this thesis, and how it relates to establishing confidence that *in silico* results are representative of the real-world domain. In doing so, the chapter makes further contributes to research objectives 1 and 3. This chapter concludes the thesis.

1.2.2 Thesis contribution

There are two strands of contribution: those made to general techniques for exploring immunology through simulation, and those specific to EAE. These are summarised below, and are expanded upon in chapter 8.

Contributions to the field of computational immunology are as follows:

- Articulation of the necessity to establish confidence in the results of immune system simulations, and identification of the four aspects of simulation-based investigation in which this must be done (chapter 3).
- Assessment of the ability of UML to describe EAE in an agent-based manner, and an approach to addressing and specifying the complexity of systems such as EAE (chapter 4).
- A highly interdisciplinary and collaborative calibration procedure for developing simulations to appropriate levels of abstraction (chapter 5).
- Development of a novel *consistency analysis* technique that establishes the relationship between the accuracy of averaged simulation results, and the number of simulation executions sampled in obtaining them (chapter 6).
- Development of a novel statistical technique, termed a *robustness analysis*, that reveals the extent to which parameters may be perturbed before a statistically significant deviation in simulation behaviour occurs (chapter 6).
- Several proposed methods of application of the robustness analysis whereby the significance of *in silico* results can be qualified in terms of the real-world domain (chapter 6).
- A proposed meta-heuristic search-based technique for assessing whether a simulation appropriately captures the complexity of the target domain. This framework can be employed in guiding simulation development to appropriate levels of abstraction (chapter 8).

In addition, the following contributions have been made to the field of EAE:

- Articulation of the complex nature of EAE, and motivating its investigation through computational modelling and simulation (chapter 2).
- An agent-based simulation of EAE, through which *in silico* experimentation may be conducted (chapter 4).
- Elucidation of important cellular processes and interactions that drive autoimmunity and subsequent regulation-mediated recovery in EAE (chapters 6 and 7).
- Identification of substantial redundancy in the ability of the regulatory network of cells in EAE to regulate the cells that mediate autoimmune behaviour (chapter 7).
- Insight into the role of the spleen in the recovery from EAE (chapter 7).

1.2.3 Publications relating to this thesis

A number of publications have resulted from work conducted towards this thesis. These publications, and their relation to the present thesis are as follows:

[Read *et al.* 2009b] : I am the principle author of this conference extended abstract. The paper expresses the initial findings on the applicability of UML to modelling complex immunological systems, which is expanded upon in chapter 4.

1. INTRODUCTION

[Read et al. 2009a] : I am the principle author of this workshop paper. The paper concerns an early iteration of the EAE domain model, the latest iteration of which appears in chapter 4.

[Read et al. 2011] : I am the principle author of this journal paper, much of which is reproduced in chapter 6.

[Polack et al. 2010] : My experiences in modelling and simulating EAE, and in conducting a cross-disciplinary collaboration with Dr. Kumar have provided a case study for this conference paper. I am not the first author.

[Williams et al. 2011] : The simulation of EAE developed in this thesis has been extended and used for further investigation into EAE. I am not the first author on this conference abstract, and the work reported is not directly related to this thesis.

Chapter 2

Background Immunology and EAE

As outlined in the previous chapter, this thesis conducts a case study in modelling and simulating a murine autoimmune disease, experimental autoimmune encephalomyelitis (EAE). This chapter presents sufficient immunological background to understand EAE, and the modelling and *in silico* experimentation work conducted in the later chapters. Section 2.1 presents a review of background immunology. Immunology is a very wide field, and the section focusses mainly on those aspects that are pertinent to EAE. Section 2.2 reviews EAE, specifically the mouse model used in the Kumar laboratory. Lastly, section 2.3 focusses on the complex nature of the immune system, motivating its exploration through computational modelling and simulation techniques.

2.1 Introduction to the immune system

Immunology is the field of biological science that studies all aspects of the *immune system*, the collective term given to the cells, molecules, processes, tissues and organs that maintain the health of an organism. Whilst this function is frequently interpreted as protection against disease and pathogen, the immune system has been implicated in a much broader range of activities, including the healing of wounds, scar tissue formation, removal of injured or abnormal cells, growing of new blood vessels and the clearance of inter-cellular debris [Cohen 2006]. The review of the immune system presented below is intended only as a general introduction to the immunology required to appreciate the processes underpinning EAE. It is largely based on two immunology text books, [Kindt *et al.* 2007, Janeway *et al.* 2005], and additional literature is cited as appropriate. Readers seeking additional information are encouraged to approach the cited literature. There are many conditions to which the immune system responds, and it is not the intention of this review to cover all such scenarios. Instead, it focusses on the immune system's response to pathogenic invasion, though many of the principles described apply also to other conditions of immune system activation.

The cells of the immune system are termed *leukocytes*. Leukocytes are reactive to the conditions of the environment that they inhabit, which they perceive through the use of *receptors*, complex protein structures bound to the surface of the cell. Stimulation of receptors through binding to the structures to which they are sensitive induces

chemical reactions or DNA transcription events within the cell, altering its behaviour or composition. As the sections that follow will reveal, there is a great deal of inter-cellular communication between leukocytes during the immune response. Communication is carried out through direct cellular contact between cell membrane-bound receptors, but also through soluble protein structures called *cytokines*. The effects of cytokines can be felt over a distance, and may be perceived by multiple cells.

Historically, immunologists have conceptualised the vertebrate immune system as composed of two parts, the *innate* and *adaptive* systems, and whilst recent study increasingly points to the extensive interactions between the two, and blurring of the boundary between them, the separation serves well as a structure in which to explore immunology [Germain 2004]. The primary distinction between the innate and adaptive immune systems relates to how their receptors are generated, and hence the stimuli to which they can react. The cells of the innate immune system are endowed with receptors that have evolved over the timescales of a species to respond to signals of the host and those of pathogenic invaders [Matzinger 1994]. Their receptors recognise heavily conserved structures that are integral to a pathogen's survival; different pathogens of the same class, be it fungi, viruses, bacteria or parasites will often express these structures. The innate immune system response is very fast to react to pathogenic invaders, effectively combatting them immediately upon detection.

The cells of the adaptive immune system are called the *lymphocytes*. The primary separation of lymphocytes from the cells of the innate immune system lies in their receptors. The lymphocyte receptor repertoire, and hence the range of structures they may react to, is established over the lifetime of the host; two individuals of the same species likely have very different lymphocyte receptor repertoires. It is the adaptive immune system that mounts responses to previously unseen pathogenic invaders. Some pathogens, such as viruses, evolve on a much faster time-scale than the species' that they infect, and whilst this affords them some ability to out-manoeuvre the innate immune system's recognition and response, the adaptive immune system can react to these evolutionary pressures much more quickly. It takes around a week for the adaptive immune system to generate a response that is effective in combatting the pathogenic invader.

The conceptual separation of innate and adaptive immune systems should not be taken to indicate that these systems operate independently of one another, there is a great deal of communication and interaction between the cells of both systems. As the sections that follow will highlight, it is the cells of the innate immune system that first prime the adaptive immune response, and the adaptive immune response in turn can co-ordinate the actions of the innate immune system in dealing with pathogens. The remainder of this introduction to immunology is organised as follows. First, section 2.1.1 discusses the organs of the immune system. A summary of the innate immune system is presented in section 2.1.2. Section 2.1.3 summarises the adaptive immune system, discussing its cell types and their receptors, the stringent activation requirements of these cells, and the phenomenon of immunological memory. Section 2.1.4 explains how the adaptive immune system is instigated, and this is followed by a review of how the immune response is terminated in section 2.1.5. Lastly, owing to its relevance in autoimmune diseases such as EAE, section 2.1.6 discusses how the immune system is prevented from attacking the host.

2.1.1 Organs of the immune system

In addition to the cells of the innate and adaptive systems, the immune system comprises several organs specialised in performing particular functions. They are classified

into two groups according to their function. The primary immune (lymphoid) organs are the thymus and the bone marrow, where lymphocytes are generated and mature before entering the circulatory system.

The secondary lymphoid organs comprise the lymph nodes situated throughout the body, and the spleen. They specialise as compartments for the interactions of leukocytes, and lymph. Lymph is the interstitial fluid found in the body's tissues. It slowly drains from the tissues through the lymph nodes and back into the blood. In doing so, the microbes and general debris that exist in the tissues are channelled through the lymph nodes, where leukocytes may react to them. The spleen performs a similar function in filtering the blood.

Immunology literature often references the *periphery*, which is taken to mean the tissues of the body. The periphery is considered separate from the lymphoid organs.

2.1.2 Innate immune system function

The innate immune system comprises a wide variety of cells, differing in the specialised functions that they perform. These include, but are not limited to, dendritic cells (DCs), macrophages, natural killer cells, and granulocytic cells. Some of these cells, such as DCs, naturally reside throughout the body, whereas others remain in the blood stream and migrate to particular locations in the body only when required to do so. Innate immune system cells are induced into performing functions through stimulation of their receptors, or receipt of particular cytokine signals. Together, these cells can perform a variety of functions in response to the perception of harmful pathogens. These include:

- The ingestion and destruction of debris and pathogens, a process termed *phagocytosis*. Cells capable of performing this function are named *phagocytes*.
- The secretion of substances that are harmful to pathogens, such as TNF- α , reactive oxygen species, and nitric oxide.
- The promotion of immune system cell migration towards the sites of bodily damage, a process known as *inflammation*.
- The induction of a controlled cellular death, called *apoptosis*, in viral-infected cells.
- The presentation of materials to the adaptive immune system, thereby initiating the adaptive immune response.

Though these functions are constitutively performed under certain conditions, largely dictated by stimulation of receptors sensitive to highly conserved structures on pathogens or indications of harm to the host's tissues, they are greatly enhanced through the actions of adaptive immune system cells, once the adaptive immune response is instigated.

2.1.3 The adaptive immune system

The adaptive immune system comprises two distinct major lineages of lymphocyte: B cells, which arise from the bone marrow, and T cells which arise from the thymus. There are two forms of adaptive immune response, the *humoral* immune response, which is mediated through B cells, and the *cytotoxic* immune response, which is mediated through a particular lineage of T cell.

The adaptive immune system is responsible for mounting immune response against previously unseen pathogenic invaders. The ability of lymphocytes to recognise such pathogens lies in the manner in which their receptors are generated, through DNA rearrangement. These receptors, and their ability to respond to particular structures, are not directly encoded in the germline; they are generated within the life-time of the host, and as such the repertoire of receptors that one possesses are highly dependent on their immunological and pathogenic history. These receptors are termed the B cell receptor (BCR) for B cells, and the T cell receptor (TCR) for T cells. The term *antigen* is given to any substance that is able to bind to a BCR or TCR in a manner that instigates the adaptive immune response.

Lymphocytes that have not encountered antigen and been activated exist in a *naive* state. They exhibit a highly motile migratory behaviour, moving between the circulatory system, lymph nodes and spleen. This migratory pattern increases the probability of their encountering antigen for which they are specific, should it exist. Should the conditions necessary for a lymphocyte to become activated be realised, it will enter a *proliferative* cycle, during which many more cells of the same kind are generated. At the end of proliferation, a lymphocyte will differentiate into an *effector* cell, capable of mediating the adaptive immune response to pathogenic invasion. Some cells further differentiate into *memory* cells, responsible for mounting more effective responses against previously encountered pathogens. Immunological memory is discussed below in section 2.1.3.3.

The BCR is capable of binding to an extremely wide range of soluble structures (which may, for example, be found on the surface of pathogens), and doing so stimulates the B cell. One effector form of the B cell is the plasma cell, which secretes a soluble form of B cell receptor called an *antibody*. Plasma cells are short lived, but can secrete vast quantities of antibodies during their lifespans. These antibodies will bind to structures on the pathogens that they recognise, and signal to other cells of the immune system, including various cells of the innate immune system, to facilitate their destruction and clearance.

There exist two major sub-populations of T cells, cytotoxic T cells (Tc cells) and T helper (Th) cells. These are characterised by their expression of either CD8 or CD4 molecules¹ respectively, molecules which partially dictate the types of structures that TCRs can bind with. CD4T cells and CD8T cells are referred to throughout this thesis. Whereas BCRs are able to bind a wide variety of structures found on cells, pathogens, molecules and alike, TCRs are only able to bind with peptides presented on specialised molecules called *major histocompatibility complexes* (MHC). The cells that present MHC molecules are called *antigen presenting cells* (APCs). TCR:MHC² interaction, which is integral to the instigation of the adaptive immune response, is discussed in the following section.

CD8Tc cells scan the host's cells for signs of intra-cellular pathogen, such as viruses. Upon recognition of such a pathogen an effector cytotoxic T cell can induce apoptosis in the target cell.

CD4Th cells are required for the initiation of both cytotoxic and humoral immune responses, as discussed below in section 2.1.3.2. There are two major sub-populations of effector CD4Th cell, the CD4Th1 and CD4Th2 cell, both of which arise from naive CD4Th cells. CD4Th1 cells generally promote and support the cytotoxic immune response, whereas CD4Th2 cells are similarly required for the instigation of the humoral

¹ CDx refers to *cluster of differentiation*, an official classification system for cell-surface receptors [Kindt *et al.* 2007].

² X:Y is the notation used throughout this thesis to denote a binding between X and Y.

immune response. CD4Th1 and CD4Th2 cells cross-regulate one another: each population regulates the size and activity of the other through various feedback mechanisms, often involving cytokines. The adaptive immune response will be either predominantly humoral or cytotoxic in nature³. Cross-regulation of CD4Th1 and CD4Th2 cells is achieved through the action of cytokines; CD4Th1 cells secrete a profile of cytokines, termed *type 1* cytokines, that promote the differentiation of activated CD4Th cells into CD4Th1 cells, and suppress the expansion of the CD4Th2 population. In turn, the CD4Th2 population secretes a *type 2* profile of cytokines that have the opposite effect.

2.1.3.1 Specificity, affinity and avidity

A key event in the activation of a lymphocyte is the stimulation of its TCRs or BCRs, achieved through binding with antigen. A lymphocyte is said to be *specific* for a particular antigen if binding of its TCR or BCR receptors with that antigen is capable of activating it (subject to the constraints explored in the following section). Lymphocytes express vast numbers of receptors on their cell surfaces, and their activation typically requires a number of these to bind with antigen. The strength of binding between an antigen and a single receptor is defined as the lymphocyte's *affinity* for the antigen. The term *avidity* describes the binding strength between an antigen and a lymphocyte whilst taking into consideration the fact that many receptors will bind to the antigen.

To illustrate: a T cell with low affinity for an antigen that is expressed in very large quantities on an APC will bind to the APC with equal avidity as a second T cell that has high affinity for a particular MHC:peptide complex, but that is expressed in low quantities.

It follows that any particular lymphocyte can recognise multiple antigen, and any single antigen can be recognised by multiple lymphocytes with distinct receptors [Cohen 2004].

2.1.3.2 Lymphocyte activation

In order to perform their functions in mediating either the humoral or cytotoxic immune responses, or supporting them in the case of CD4Th cells, lymphocytes must become induced into their *effector* states. These immune responses, particularly the cytotoxic response, can cause substantial harm to the host if mis-directed, and hence there exist stringent safeguards governing their generation. CD4Th cells must receive two signals before becoming effector cells; CD8Tc and B cells also require two signals, and further require the support of activated CD4Th cells in inducing their effector functions [Matzinger 1994]. Figure 2.1 depicts the signals required for lymphocyte activation. Note that lymphocyte activation and population expansion is also referred to as *priming*.

CD4Th cells are activated by APCs. Dendritic cells are the most potent activators of naive CD4Th cells. A naive CD4Th cell whose TCRs successfully bind with MHC-II:peptide complexes as displayed by a DC will receive signal 1 and become stimulated. If the DC is expressing co-stimulatory molecules, then these will bind with respective receptors on the CD4Th cell, and it will receive signal 2. At this point the CD4Th cell becomes activated and enters its proliferative cycle, and will eventually differentiate into either an effector or memory CD4Th1 or CD4Th2 cell, the latter being dictated largely by the profile of cytokines being secreted by the DC [Germain 2004, Kapsenberg 2003].

³The humoral immune response can support cytotoxic activity, however this aspect of immunology is beyond the focus of this thesis.




Figure 2.1: The events leading to the activation of lymphocyte cells, adapted from [Kindt *et al.* 2007] and [Janeway *et al.* 2005]. Lymphocytes have stringent requirements for their activation; B and T cells each require the receipt of two signals, termed *signal 1* and *signal 2*, in order to become activated. Transduction of these signals are mediated through the binding of a series of molecules and their corresponding receptors. Activation of B cells requires the activation of CD4Th2 cells, and activation of CD8Tc cells requires the activation of CD4Th1 cells which must license and APC to express the required levels of co-stimulatory molecule.

The activation of a B cell is also contingent on its receipt of two signals. A B cell derives signal 1 when multiple BCRs bind with some antigen. The B cell will ingest the antigen, and process it into peptides, some of which will then be presented as MHC-II:peptide molecules. An effector CD4Th2 cell specific for such MHC-II:peptide complexes will bind to the B cell. An effector CD4Th cell expresses CD40L molecules, and these bind with CD40 on the B cell to induce signal 2 in it. The B cell will then differentiate into either a memory B cell, or a plasma cell.

The activation of CD8Tc cells is also dependent on CD4Th cell function. The TCRs of CD8Tc cells are restricted to recognising peptides presented on MHC-I, rather than MHC-II. Successful binding of a naive CD8Tc cell's TCR receptors with MHC-I:peptide complexes expressed on an APC will deliver signal 1 to the CD8Tc cell. Perhaps owing to the tremendous destructive potential of CD8Tc cells, they require more co-stimulatory stimulation than CD4Th cells. DCs must be induced into expressing the levels of co-stimulatory molecules required to deliver signal 2 to a naive CD8Tc cell, and this is accomplished through the *licensing* of the APC by a CD4Th1 cell. An effector CD4Th1 cell binds with MHC-II:peptide complexes on the APC, and then licenses it through CD40:CD40L interaction. A second means through which CD4Th cell can be required for the activation of CD8Tc cells is through the promotion of non-classical MHC-I molecules⁴ on APCs [Tang *et al.* 2006]. Hence, both the CD8Tc and the CD4Th1 cell must recognise MHC:peptide complexes derived from the same pathogen before the cytotoxic immune response can initiate. Upon receipt of both signals 1 and 2, a CD8Tc cell will enter its proliferative cycle, and eventually differentiate into either an effector CD8Tc cell, or a memory CD8Tc cell.

MHC-I is expressed by most of the host's cells. These molecules specialise in processing samples of a cell's internal contents into peptides and presenting them for the scrutiny of CD8Tc cells. Hence, CD8Tc cells specialise in combatting intra-cellular pathogens such as viruses. MHC-II expression is generally restricted to particular phagocytic cells of the innate immune system, such as DCs and macrophages. These cells process exogenous antigen that they have phagocytosed and present it as MHC-II:peptide complexes. Through a process termed *cross-priming* such APCs are also able to present peptides derived from exogenous antigen on MHC-I complexes [Bevan 2006a]. This a critical process in the instigation of the cytotoxic immune response against intra-cellular pathogens that do not infect DCs or macrophages.

2.1.3.3 Immunological memory

Another phenomenon that the adaptive immune system bestows upon the host is that of immunological memory. The adaptive immune system takes around 10 to 17 days to fully mount a response to a previously un-encountered pathogen. Where a second immunization with the pathogen occurs, the adaptive immune system takes between 2 and 7 days to respond.

This increased rate of response to previously encountered pathogen is mediated through particular forms of effector B and T cells: memory cells. When a B or T cell differentiates into an effector cell, there is some probability that it will become a memory cell. These cells are long-lived, and have less stringent requirements for their activation: no signal 2 is required, and these cells can respond to reduced concentrations of antigen.

It is this immunological memory that vaccines exploit in providing long term protection against particular pathogens. A typical vaccine comprises an attenuated form

⁴There are many different forms of MHC molecule, each of which is capable of presenting distinct repertoires of peptide [Rodgers & Cook 2005].

of the pathogen, providing a target for the development of lymphocyte memory cells, in conjunction with substances called *adjuvants*. The use of attenuated pathogens protects the recipient from exactly that danger which the vaccine seeks to circumvent, however, pathogens that are not capable of harming the host do not necessarily incite a potent immune response [Matzinger 1994]. Adjuvants are substances that strongly stimulate the immune system into action, hence facilitating the perception of an attenuated pathogen as harmful, despite the fact that it is not. With the addition of adjuvant, the immune system will generate a response against the attenuated pathogen, which includes the generation of memory cells.

2.1.4 Instigation of the adaptive immune response

A typical instigation of the adaptive immune response commences with the infiltration of pathogen into a particular site of the host. Pathogenic action causes harm to the local tissues, which alerts the cells of the innate immune system. Dendritic cells, which reside throughout many of the body's tissues as immature cells, are highly sensitive to the presence of pathogenic materials and harm to the body's tissues [Kapsenberg 2003, Matzinger 2007]. Their constitutive phagocytic activity results in their ingestion of the pathogen, and its subsequent presentation as MHC:peptide complexes. DCs that perceive pathogens or harm to tissues become stimulated into up-regulating their co-stimulatory molecule expression. As an aside, this is one role of adjuvant in vaccines, to facilitate the perception of harmful substances by DCs. DCs express a range of receptors, the Toll like receptors (TLRs), that are sensitive to different classes of pathogen (such as viruses, bacteria, parasites and fungi). Depending on which TLRs are stimulated, DCs can be induced into secreting different profiles of cytokine. In this manner, through perception by DCs, the class of pathogen can influence whether a CD4Th response develops into a CD4Th1 or CD4Th2 direction [Germain 2004, Kapsenberg 2003, Sousa 2006]. It has also been suggested that the class of adaptive immune response generated is influenced by signaling of DCs by the tissues, since certain classes of response are highly destructive to certain organs, and administration of the same antigen in different tissues enlists different classes of response [Matzinger 2007, Matzinger 2002].

DCs periodically mature and migrate from the peripheral tissues to the local lymph nodes (or the spleen, depending on the site of pathogenic invasion) where they present antigenic peptides to T cells. Owing to their co-stimulatory molecule expression, any T cells specific for the MHC:peptide complexes that they present will become activated and enter their proliferative cycles. The circulation of lymph from the peripheral tissues, through the lymph nodes, and into the circulatory system results in the transport of some pathogenic material directly into the lymph nodes, permitting the activation of B cells. Antigen transported through the circulation of lymph may also lead to its phagocytosis by lymph-node resident DCs that then present to T cells [Lanzavecchia & Sallusto 2004, Sousa 2006].

Following activation, lymphocytes enter a proliferative cycle, where for several days they repeatedly divide to produce a population of naive daughter cells. These naive cells may be activated through the presence of the same antigenic materials that activated their parents, since their receptors are largely identical⁵. Upon completion of their proliferative cycles, lymphocytes differentiate into effector cells, and leave the lymph nodes.

⁵In the case of T cells, naive progeny resulting from proliferation have exactly the same TCR structures as their parent cells. In the case of B cells, some mutation of the BCR occurs, and those progeny harbouring mutations resulting in greater affinity for antigen are selected for survival.

Innate immune cells such as macrophages at the site of pathogenic invasion secrete, and induce in the local tissues an expression of, a variety of molecules that attract other immune system cells to the sites of damage. This process is called *inflammation*. Effector lymphocytes re-entering the circulatory system after leaving the lymph nodes migrate to the site of pathogenic invasion in response to inflammation. Here they receive additional antigenic stimulus, termed *local activation*, which stimulates them into performing effector functions.

The actions of the adaptive and innate immune systems integrate at the site of infection to combat pathogenic invasion. Lymphocytes can contribute to the inflammatory process, recruiting cells of both the adaptive and innate immune systems. The cytokine milieu can promote the activation of macrophages, inducing in them the secretion of TNF- α , reactive oxygen species, and nitric oxide, substances which are harmful to many pathogens but which can also cause damage to the tissues. DCs are induced into immunogenic phenotypes, up-regulating co-stimulatory molecule expression and migrating to the secondary lymphoid organs to prime further populations of lymphocytes. The cytotoxic immune response can enlist cells of both the innate and adaptive immune systems to induce apoptosis in local tissue cells, such as those infected by viruses. The humoral immune response results in the secretion of antibodies that can mark pathogens for phagocytosis, or tissue cells for apoptosis.

If it persists, the immune response can cause considerable harm to the tissues. The inflammatory process can recruit a huge number of cells that can disrupt tissue organisation and function, and it can promote the generation of scar tissue. Cytotoxic activity can directly destroy substantial portions of tissue, and the substances secreted by innate immune cells can likewise cause widespread disruption of organ function. The immune response is largely supported by the lymphocytes, and as such the following section considers how the immune system is terminated.

2.1.5 Terminating the immune response

An effective immune response will clear the pathogenic invader from the host, after which the large populations of lymphocytes, capable of mediating considerable damage to the host's tissues, are no longer needed. There are two processes by which the adaptive immune system terminates, discussed in turn.

Lymphocytes require persistent stimulation of their BCRs or TCRs in order to survive, and in absence of antigenic stimulation the effector cells generated to combat a pathogen will enter apoptosis upon completion of this task. This form of lymphocytic death is called *death by neglect* [Parijs & Abbas 1998, Matzinger 1994]. It applies throughout the lymphocyte lifecycle: naive cells also require some low level of stimulation in order to persist.

Even in the presence of antigenic stimulation, effector lymphocytes cannot survive indefinitely. Through a process termed *activation induced cell death* (AICD), lymphocytes that are persistently stimulated are induced into apoptosis. This is believed to be a homeostatic mechanism to control the size of effector T cell populations during the immune response [Kabelitz *et al.* 1993].

2.1.6 Immune system tolerance

The DNA rearrangement processes that are responsible for the creation of BCRs and TCRs endow these populations of cells with the ability to recognise an extremely wide range of structures, including those that constitute the host. A long-standing debate in immunology concerns the mechanisms and processes whereby cells of the adaptive

immune system are prevented from targeting the host⁶. It has been argued that all self-reactive B and T cells are eradicated in the thymus and bone marrow, and hence that self-reactive lymphocytes cannot enter the periphery. Under this argument, autoimmunity may be attributed to a failure of these selection and eradication mechanisms. Whilst the role of the thymus and bone marrow in deleting many self-reactive lymphocytes is largely acknowledged, it is now well appreciated that self-reactive lymphocytes do in fact exist in the periphery of healthy individuals; indeed some level of self-reactivity is required to prevent death by neglect in naive lymphocytes. As such, a variety of mechanisms are coming to light whereby the ability of these cells to mount effective immune responses against the body are suppressed.

Self-reactive lymphocytes that escape the bone marrow or thymus can receive an abundance of stimulation in the periphery, where the antigens to which they are reactive are expressed in large quantities. As discussed above, this persistent stimulation of lymphocytes leads to AICD, and the self-reactive lymphocytes will enter apoptosis.

A second mechanism mediating adaptive immune tolerance to the host is through the induction of *anergy* in self-reactive lymphocytes [Parijs & Abbas 1998]. Anergy is defined as unresponsiveness to antigenic stimulation, and also leads to eventual apoptosis. Anergy is induced in lymphocytes that, during activation, receive signal 1 in absence of signal 2. T cells are activated through APCs, predominantly DCs, wherein co-stimulatory molecule expression at levels to deliver signal 2 must be induced. As noted above, direct recognition of pathogen or harm having occurred in the tissues can lead to DC expression of co-stimulatory molecules. Where DCs phagocytose and then present elements of the host to T cells, but where no pathogenic presence or tissue damage has occurred, presentation will be in a *tolerogenic* context⁷, marked by the absence of co-stimulatory molecule expression. T cells binding with such APCs will be anergised. In the case of B cells, the lack of effector CD4Th2 cells capable of providing signal 2 also leads to anergy [Matzinger 1994].

Lastly, recent years have seen the characterisation of a variety of regulatory T (Treg) cells which act to suppress and down-regulate adaptive immune responses, rather than promote them. A wide variety of T cell subsets capable of performing such activities have been identified, of both CD4 and CD8 varieties. These Tregs have been shown to modulate the immune response through a range of mechanisms, including the secretion of cytokines that interfere with T cell activation [Boehmer 2005, Shevach *et al.* 2001], down-regulation of APC ability to prime T cells [Chang *et al.* 2002], promotion of the generation of other Treg subsets [Cortesini *et al.* 2001], and the direct apoptosis of activated T cells [Smith & Kumar 2008, Tang *et al.* 2005]. Some such Tregs are *specific* in their suppression, targetting particular T cells on the basis of the antigen that they recognise, whereas others are *non specific*, secreting cytokines in a manner perceivable by any neighbouring cells.

2.2 Experimental Autoimmune Encephalomyelitis

Experimental autoimmune encephalomyelitis (EAE) is an animal autoimmune disease that arises through the direction of immunity towards the central nervous sys-

⁶In 2000, the journal *Seminars in Immunology* dedicated an entire issue (volume 12, issue 3) to this debate, with some of the most prominent figures in immunology sounding and critiquing the various perspectives on how the immune system targets pathogenic invaders but not the host. For example, see [Zinkernagel 2000, Silverstein & Rose 2000, Medzhitov & Janeway, Jr 2000, Langman & Cohn 2000, Grossman & Paul 2000, Cohen 2000, Bretscher 2000, Anderson & Matzinger 2000].

⁷As opposed to *immunogenic* context.

tem (CNS). It is one of the earliest developed animal models of disease, and has been widely used as a model for studying multiple sclerosis [Baker & Jackson 2007, Baxter 2007, Goverman 2009, Zamvil & Steinman 1990]. EAE is studied in a variety of animals including, but not exclusively, mice, rats, guinea pigs and monkeys. The clinical course of EAE varies with animal and with the immunization protocol used to induce it, and across the different models a wide range of disease phenotypes can be induced, from acute to relapsing to chronic paralysis, of varying degrees of severity [Baker & Jackson 2007, Baxter 2007]. Depending on the protocol used to induce it, a proportion of experimental animals perish from EAE. This thesis, and hence this review of EAE, focusses primarily on the model of EAE used in the Kumar laboratory [Kumar & Sercarz 2001, Kumar 2004], which is studied in mice.

EAE involves direction of immunity towards the myelin sheath, an insulatory material that coats the neurons of the CNS, and which is necessary for their proper function. The disease is typically induced through two methods: the administration of myelin derivatives, such as myelin basic protein (MBP), in conjunction with adjuvants that strongly stimulate the immune system into action; or through the adoptive transfer of myelin specific effector T cells from one animal already induced into disease into another [Zamvil & Steinman 1990].

The susceptibility of experimental animals to autoimmunity following administration of the exact same immunization protocol varies considerably. Some do not present clinical symptoms at all, whilst others perish. The Kumar lab grades the severity of EAE experienced by an experimental mouse on a scale of 0 to 5: 0, no symptoms; 1, flaccid tail; 2, hind limb weakness; 3, hind limb paralysis; 4, whole body paralysis; 5, death [Kumar *et al.* 1996]. Figure 2.2 demonstrates the variation in autoimmune severities experienced by mice undergoing various interventions. In the EAE model employed in the Kumar lab animals frequently experience spontaneous recovery from autoimmune symptoms, as may be seen on the figure. It is this spontaneous recovery that has attracted the interests of immunologists seeking to understand how the same recovery, mediated through cells that are also found in humans, can be induced in sufferers of multiple sclerosis.

Figure 2.3 depicts the major cell types that are involved in the Kumar laboratory's EAE model and its associated recovery. Their role in mediating autoimmunity and recovery therefrom are discussed in sections 2.2.1 and 2.2.2 that follow.

2.2.1 EAE autoimmunity

A typical immunization for EAE entails a sub-cutaneous injection of MBP. This leads to its phagocytosis by dendritic cells, and subsequent display as MHC-II:MBP on the cell surface. The adjuvants that accompany immunization for EAE, CFA and PTx⁸, induce DCs to adopt a highly immunogenic phenotype, making them strong primers of CD4Th1 cell immune responses [Menezes *et al.* 2007].

DCs residing in the periphery, having phagocytosed MBP and been induced into maturity by adjuvant, migrate to the local lymph nodes, and there prime populations of naive MBP-specific CD4Th cells [Goverman 2009]. These cell populations enter their proliferative cycles, and adopt either CD4Th1 or CD4Th2 polarizations. Having been exposed to adjuvant, the DCs promote the adoption of CD4Th1 polarizations, though polarization adoption is probabilistic, and some CD4Th cells will still adopt a CD4Th2

⁸CFA, complete Freund's adjuvant, consists of the inactivated bacterium *mycobacterium tuberculosis*. PTx, pertussis toxin, is a protein produced by the bacterium *mordetella pertussis*. Both these substances are powerful inducers of the immune response.




Figure 2.2: The progression of EAE in individual mice amongst four experimental groups, with five mice in each group. Group A experiences onset and then physiological recovery from disease. Groups B, C and D have received various interventions that interfere with the animal's ability to recover from clinical symptoms. Note the considerable variation in disease progression experienced by mice of the same experimental group, having undergone the exact same immunization procedure. Figure adopted from [Kumar *et al.* 1996].




Figure 2.3: Abstract depiction of the major cell types involved in EAE autoimmunity and its associated recovery. Red arrows indicate interactions leading to autoimmunity, blue arrows indicate regulatory activity that counters the autoimmune response.

phenotype [Ando *et al.* 1989, Menezes *et al.* 2007]. Upon completion of their proliferative cycles, these T cells differentiate into effector T cells, migrate out of the lymph nodes, and rejoin the circulatory system, eventually reaching the CNS compartment.

CD4Th1 and CD4Th2 effector cells are able to cross the blood-brain barrier and gain entry to the CNS. There they derive local activation from APCs presenting MHC-II:MBP. CD4Th1 cells commence the secretion of pro-inflammatory cytokines such as IL-2, and INF- γ [Ando *et al.* 1989, Menezes *et al.* 2007].

These cytokines constitute a pro-inflammatory context to resident antigen presenting cells, such as macrophages and microglia⁹, which become stimulated. Once stimulated, these cells secrete TNF- α , reactive oxygen species (ROS), and nitric oxide (NO), all of which are harmful to neurons. Sufficiently high doses of these substances cause demyelination: the destruction of the insulator myelin sheath that covers the neurons [Hendriks *et al.* 2005, Raivich & Banati 2004, Tambuyzer *et al.* 2009].

Following demyelination, myelin fragments such as MBP are phagocytosed by macrophages, microglia and DCs. Macrophages and microglia express MHC-II:MBP following this phagocytosis, and thereby provide further stimulus for local activation of encephalitogenic¹⁰ CD4Th1 and CD4Th2 infiltrates. The inflammatory conditions now present in the CNS induce these APCs to up-regulate co-stimulatory molecule expression, hence adopting immunogenic phenotypes. The DCs migrate to the cervical lymph nodes (CLN), and there present MHC-II:MBP [Goverman 2009].

This presentation of MBP in the CLN by DCs provides further stimulus for the generation of MBP-specific CD4Th cell populations. Once a CD4Th1 presence is established in the CNS compartment, MBP-presenting DCs will prime further populations of these same cells in the CLN, and hence autoimmunity becomes self-perpetuating, potentially persisting for long after the effects of immunization in the peripheral lymph nodes have expired.

2.2.2 EAE regulation

Many mice induced into EAE spontaneously recover from autoimmune symptoms, even when the mechanisms mediating regulation are interfered with, see figure 2.2. The Kumar lab has identified a network of cells that has a regulatory effect on autoimmunity, and mediates recovery from disease. This regulatory network is indicated on figure 2.3 by blue arrows. The Kumar lab has characterised two forms of Treg that play a significant role in recovery from EAE, CD4Tregs and CD8Tregs. Experiments to deplete or incapacitate these cells results in labored recovery from autoimmune symptoms [Kumar *et al.* 1996, Beeston *et al.* 2010], whereas their artificial premature activation following induction for EAE protects recipients from autoimmunity [Tang *et al.* 2007]. Further, the adoptive transfer¹¹ of activated CD4Treg or CD8Treg cells into mice prior to immunization protects them from subsequent attempts to induce EAE [Tang *et al.* 2006, Kumar 1998, Kumar *et al.* 2001]. This section details how these two Treg cell populations operate within the mouse immune system to mediate recovery from autoimmunity.

Through the natural course of their lifecycles, CD4Th1 cells die of AICD, entering apoptosis and being subsequently phagocytosed by APCs, such as DCs [Kabelitz *et al.* 1993]. The majority of encephalitogenic CD4Th1 cells use the V β 8.2 TCR gene segment to encode their TCRs [Kumar & Sercarz 2001]. DCs derive and present two peptides from two regions of these TCRs: framework 3 (Fr3), and complementarity determining region 1/2. MHC-II:Fr3 complexes are presented by DCs, which prime populations of CD4Tregs [Kumar 1998, Smith *et al.* 2010].

CD4Treg cells secrete IL-2 and INF- γ cytokines necessary for the induction of Qa-1:CDR1/2 expression on DCs [Tang *et al.* 2006], a process known as *licensing*.

⁹Microglia are specialised APCs that reside in the central nervous system.

¹⁰*Encephalitis* relates to inflammation of the tissues of the brain.

¹¹Adoptive transfer experiments take cells from one experimental animal, in this case effector CD4Treg or CD8Treg cells from mice undergoing recovery from EAE, and administer them to another to observe the effect. Here, these cells conferred some protection from attempts to induce EAE following the transfer.

ing. Qa-1 is a form of non-classical MHC-I molecule that presents a substantially smaller repertoire of peptides than classical MHC-I can. Qa-1:CDR1/2 presentation and CD4Treg cytokine secretion leads to the priming of CD8Treg populations¹² [Smith *et al.* 2009, Kumar 2004].

For around 8 hours following differentiation into effector cells, V β 8.2 TCR CD4Th1 cells express Qa-1:CDR1/2 complexes, for which CD8Treg cells are specific. Effector CD8Treg cells are a form of cytotoxic T cell, and upon binding Qa-1:CDR1/2 expressed on a CD4Th1 cell are able to induce its apoptotic death [Beeston *et al.* 2010, Tang *et al.* 2006].

On a system level, this rise in CD8Treg population number leads to a reduction in the CD4Th1 cells that mediate self-perpetuating autoimmunity: recently primed CD4Th1 cells originating from the CLNs are induced into apoptosis before reaching the CNS compartment, and hence the loop of self-perpetuating autoimmunity is broken.

It is only CD4Th1 cells that are subject to regulatory activity [Kumar 2004, Madakamutil *et al.* 2003]. The MBP-specific CD4Th2 cells escape regulation by CD8Treg cells, and as a result of their unhindered expansion the immune response deviates in a type 2 direction [Kumar 2004, Kumar & Sercarz 2001]. In the context of EAE, CD4Th2 cells, and the cytokines they secrete, do not promote demyelination and instead serve to counter the pro-inflammatory encephalitogenic context in the CNS [Kumar 1998].

2.3 Summary

This chapter has served as an introduction to the immune system, and the manner in which the immune response is instigated, regulated and terminated. It has explored the domain of the simulation work conducted in this thesis: the particular EAE mouse model employed by the Kumar laboratory. The perspective of immunology presented here is somewhat simplified, and tailored towards understanding the processes underpinning EAE; the field of immunology extends considerably further than that which is outlined in this chapter. Cutting edge research continues to identify new cells, molecules, pathways and processes that constitute the immune response. The immune system is increasingly revealed to be one of considerable complexity, briefly explored here. It is the need to address this complexity that has partially motivated immunological investigation through *in silico* techniques.

Section 2.1 has only touched upon the myriad of cells that mediate and control the immune response. Many of these cells overlap in the functions that they perform, such as the ability to recognise particular pathogens, though the processes through which they do so differ [Parnes 2004, Cohen 2004]. A particular immune insult will attract the attention of a wide variety of different cell populations. There is extensive communication between the cells of the immune system, which influence the nature of one another's function and role in the immune response. Such communication can be mediated through direct cellular contact, or through cytokines. To date, a vast number of different communication and cytokine molecules have been uncovered¹³.

Cytokines themselves exhibit considerably complexity. One particular cytokine may have a range of different effects on different cells, and several distinct cytokines

¹²CD8Treg cells that are specific for Qa-1 have been implicated in maintaining self-tolerance and countering autoimmune behaviour in a variety of diseases by regulating self-reactive T cells [Lu *et al.* 2006, Kim *et al.* 2010].

¹³Cell-bound molecules are typically assigned numbers of the form 'CDxxx', where *xxx* is an incrementing number following the discovery of new molecules. The immunology textbook Kuby Immunology lists 339 such molecules, and 50 distinct cytokine molecules [Kindt *et al.* 2007].

are capable of inducing similar functional changes on any particular type of cell. They exhibit the properties of synergy and antagonism, where the effect of one cytokine may only be realised in the presence of another, or where one cytokine may counteract the effect of another. Similar principles apply for cell-bound molecules.

The particular form that an immune response acquires is influenced by a wide variety of factors, including the recognition of certain qualities of the pathogen by cells of the innate immune system, the tissues from which the immune response originates, and the various subpopulations of lymphocytes that cross-regulate one another. T cells lie at the heart of cross-regulatory and suppressive activities that shape the class of immune response that occurs. These cells can target one another directly, and can influence the APCs that prime T cell populations. The dendritic cell alone is capable of influencing at least seven distinct T cell responses, each of which can further influence DC behaviour [Sousa 2006]. Matters become even more complicated than this: the evolutionary pressure on pathogens to trigger receptors that promote a class of immune response that is not effective in dealing with them has been highlighted [Matzinger 2002].

In summary, as with many biological systems, every level of the immune system seems riddled with complexity [Cohen 2004, Edelman & Gally 2001]. There are countless processes that mediate positive and negative feedbacks on the particular directions that an immune response might take. From the perspective of autoimmune behaviour, there are simply too many pathways and possibilities for immunologists to have a clear understanding of how autoimmunity manifests, let alone how to effectively extinguish it in a manner that does not leave the host's immune system significantly compromised. The same can apply to a host of other pathogenic invaders that seem impervious to the immune response, whose actions can result in chronic immunity with dire consequences for the host.

It is this complexity that is leading immunologists to explore computational modelling and simulation techniques as a means of complementing their traditional research programs. Identifying how to best manipulate the immune system in combatting disease is likely to require a system-level holistic overview of its function, and it is hoped that constructionist computational methods can assist in providing it [Cohen 2007b, Chakraborty & Das 2010, Germain *et al.* 2011]. Existing applications of computational methods in studying the immune system are the subject of the following chapter. A domain model of EAE is presented in chapter 4.

Chapter 3

Modelling and Simulating Immunological Systems

The previous chapter concluded by highlighting the complex nature of the immune system, and how immunologists are beginning to complement traditional wet-lab research with computational methods. It has been proposed that computational methods can help provide a system-level overview of immune system function, necessary for identifying the best strategies for combatting disease. It is the application of computational approaches to immune system study that forms the focus of this chapter.

The chapter is organised as follows. Section 3.1 motivates the investigation of the immune system through *in silico* techniques. An overview of mathematical and agent-based modelling techniques is presented in section 3.2. Section 3.3 explores the manner in which select computational works have provided a valuable complement to traditional wet-lab techniques. A review of existing modelling work on EAE is presented in section 3.4. The issue of building confidence in simulation results being representative of the target system is articulated in section 3.5. This section draws on representative examples of the literature to provide a critique of best practice in the field. Finally, section 3.6 concludes this chapter.

3.1 Motivation for modelling and simulating the immune system

There is much overlap in the terminology employed in the literature on modelling and simulating the immune system, hence the following terms and definitions are used throughout this thesis. A *model* is assumed to be some abstract description of a target system, theory, or phenomena [Polack *et al.* 2008]. It may be diagrammatic or mathematical in nature. A *simulation* is an instantiation and execution of a model on a computer.

A review of the literature on modelling and simulating the immune system reveals a number of ways in which these techniques are reported to benefit the field of immunology. These are broadly summarised as follows:

- Simulations can be more amenable to designing, conducting and collecting data from intricate and revealing experimentation than the natural system

[Forrest & Beauchemin 2007]. Simulations are simply code, and there are no limitations on how it may be varied in facilitating investigation. On the contrary, there are physical and sometimes ethical limitations to what can currently be performed in the wet-lab. For example, whereas wet-lab technology may offer the possibility of genetic ‘knock out’ experiments, where an experimental animal may be prevented from producing a particular cytokine, simulation allows experimenters to interfere with only one of several possible functions that the same cytokine may perform.

- Simulations can form a common ground for the mass consolidation of experimental data [Chakraborty & Das 2010, Kirschner & Linderman 2009]. It is increasingly appreciated that the advancement of high-throughput wet-lab experimental techniques are generating more data than can be intuitively understood without the aid of integrative tools [Cohen 2007a, Bauer *et al.* 2009, Germain *et al.* 2011]. Simulations can capture and integrate data from a wide variety of labs, using different experimental methods, and provide a system-level overview of what this data represents.
- Computational approaches offer a platform on which to formulate and evaluate competing hypotheses of system operation, permitting experimenters to assess how well different theories match established experimental data [Chakraborty *et al.* 2003, Kirschner & Linderman 2009]. This may, for example, lead to the identification of pathways or interactions not previously appreciated [Chakraborty & Das 2010].
- By means of the points above, simulations can help in directing wet-lab experimentation, allowing investigators to quickly exclude fruitless avenues of research before incurring the considerable cost of wet-lab experimental work. Preliminary investigation *in silico* can point to interesting aspects of the system, can highlight where knowledge is lacking, and can help design revealing experiments to discern between competing theories [Cohen 2007a, Chakraborty & Das 2010, Bauer *et al.* 2009].

A key benefit of modelling and simulation is the insight that may be gained across scales, from molecular to cellular, from cellular to population, from different cellular populations to the level of an individual, and beyond to societies [Cohen & Harel 2007]. The scales can be temporal as well as spatial, where molecular events on a short scale, perhaps at the start of an immune response, may have profound effects on the overall long term progression of the response. Observation across scales is a feature of many modelling works which examine the roles that low-level system components have on overall system behaviour. Most of the literature focusses on manipulating interactions on one scale and observing at the scale above. The value of integrating models across multiple scales has been highlighted, but the technology to do so has not yet been established [Kirschner & Linderman 2009, Kirschner *et al.* 2007]. To give a hypothetical example, application of computational methods have recently isolated a potential target for vaccines against HIV [Dahirel *et al.*], owing to this particular aspect of the virus’s conservation against mutation. It is hypothesised that this particular element of the virus is critical to its function, and hence it cannot be quickly mutated to avoid immune attack, as is the case with much of the rest of the virus’s structure. Before deploying such a vaccine, it would be worthwhile knowing whether this virus could mutate to circumvent the immune system, and if so, what the time scale for this might be. As such, a timetable for vaccine deployment could be established. The multi-scale

computational techniques of the future might be capable of simulating the effectiveness of the vaccine at combatting the virus at the individual level, encompass a notion of how quickly the virus can mutate, integrate this across entire populations, and extrapolate whether such a vaccine is viable as a long term solution to HIV.

Though the field of modelling and simulating the immune system is growing, in many regards it is still in its infancy. The above issue of multi-scale modelling is not addressed in this chapter¹, but another aspect of the field, felt to be absolutely critical to its maturation and establishment as a viable complement to wet-lab research, is that of confidence in simulation [Polack *et al.* 2010]. Simulations are abstractions of their target systems, and much of the target system is unknown, which can be problematic for creating simulations. Given these obstacles, it is vital that researchers build confidence in the community that the results of *in silico* experimentation are genuinely representative of the real domain [Polack 2010, Polack *et al.* 2010]. This issue of confidence in simulation is a principle theme in this thesis, and its recognition in the literature is a central focus of this chapter.

3.2 Review of modelling and simulation techniques

The section begins by reviewing mathematically based approaches to immune system modelling, then computational agent-based approaches, and lastly provides some examples of novel hybrid approaches that balance the tradeoffs between the two.

It should be noted that there is no ‘correct’ or ‘incorrect’ modelling paradigm for investigating the immune system. Rather, the choice of which modelling tools one adopts should be dictated by what is appropriate for the domain being modelled, given the lines of investigation to be carried out. Consideration must be given to the number of system entities that need to be represented, whether stochastic events are significant to the system or whether purely averaged behaviours are appropriate, whether spatial or population heterogeneity is important to the overall system dynamics, and the types of assumptions that one can reasonably make of the target system [Bauer *et al.* 2009].

3.2.1 Mathematical approaches

A mathematical modelling technique that has found frequent application in exploring the immune system is the ordinary differential equation (ODE). Examples include the modelling of influenza infection [Bocharov & Romanyukha 1994], the immune response in tuberculosis [Marino *et al.* 2010], and the competition and co-evolution of different strains of HIV in the presence of the immune response [Bagnoli *et al.* 2006]. The latter was extended to model the effects of three current HIV treatments with respect to dominance switching between HIV strains [Sorathiya *et al.* 2009]. ODEs represent populations of elements as real-numbered variables, and specify through coupled equations the interactions that dictate how population sizes change under one another’s influence.

ODEs are computationally efficient to execute, and are amenable to a wide variety of analytical techniques [Bauer *et al.* 2009]. They can be effective when investigating the emergent dynamics that result from the interactions between large populations of structurally homogeneous entities.

There are however several assumptions inherent in ODE modelling that have prompted modellers to seek alternative representations. ODE models assume a well

¹The reader is referred to [Kirschner & Linderman 2009, Kirschner *et al.* 2007] for discussion on multiscale modelling of the immune system.

mixed space of elements, when in fact spatial considerations *in vivo* are critical to many immunological processes. This is illustrated by an investigation into the spread of a viral infection, where the authors contrast an ODE representation with an equivalent agent-based model [Beauchemin 2006]. In the ODE model, the rate at which infection spreads is related to the ratio of uninfected cells to infected cells, whereas the agent-based model incorporated an explicit notion of space, and allowed only those cells neighbouring infected cells to themselves become infected. The author concludes that the agent-based model better reflects real-world experimental data.

ODE models are suitable for representing large populations of elements, but are susceptible to strange behavioural artifacts when populations are very small. In the real world, the removal of the last member of a population entails a population size of zero, however, typical ODE population decay mechanisms allow populations to approach zero, but never truly reach it, and hence these systems allow for un-realistic rises in population number [Meier-Schellersheim & Mack 1999].

ODEs tend to have fewer parameters than their spatially-explicit counterparts, potentially making them easier to parameterise, and the more appropriate modelling paradigm if wet-lab data was derived from well-mixed sources, such as harvesting cells from the spleen or blood [Bauer *et al.* 2009].

3.2.2 Agent-based approaches

Whereas ODE modelling represents all the entities of a particular variety as a single real-numbered variable, agent based modelling provides an explicit discrete representation of each individual entity, such as a cell, in the system [An *et al.* 2009]. ODEs assume all entities belonging to particular populations to be identical, whereas in the agent-based paradigm agents can be heterogeneous, and maintain their own potentially unique state. The dynamics of agents in the agent-based paradigm are dictated through rules governing the behavioural changes that result from their interactions with one another, and with their environment. Agent behaviours, though governed by these rules, are often stochastic in nature. Hence, distinct agent-based simulation executions can yield different results, despite having the same inputs. For domains where stochasticity or diversity in the history of agents can be influential on the overall system-behaviours, the agent-based approach can be a highly suitable modelling paradigm [Bogle & Dunbar 2010, Germain 2001, Milanesi *et al.* 2009].

Since this paradigm permits the explicit modelling of individual cells and their states, such as molecular expressions or the rules that govern changes therein, agent-based modelling allows for the expression and exploration of hypotheses that are difficult to represent in ODE models [Forrest & Beauchemin 2007]. These may include, for instance, the system-level effects of preventing the expression of a particular molecule, altering its expression duration, or interfering with only one of several of its functions.

Agent-based systems often encompass an explicit representation of space, with the movement of agents through this space dictating their interactions with one another. Such explicit representation permits investigation into the system-level effects that arise from processes inherently dependent on space, such as the structure of cells within lymphoid organs [Efroni *et al.* 2007], or the emergence of highly spatially-structured immune responses to diseases such as tuberculosis [Ray *et al.* 2009]. There are, however, examples of literature employing well-mixed space within the agent-based paradigm [Garrett *et al.* 2007].

A variety of spatial representations may be found in the literature. Lattice grid representations, where space is discretized, are common, and both two dimensional and three dimensional representations have been used [Segovia-Juarez *et al.* 2004,

Bogle & Dunbar 2008]. Continuous spatial representation has also found application in the modelling literature [Jacob *et al.* 2004]. Lattice grid representations do not offer the same granularity as continuous space, but are less computationally demanding; realistic simulation of cells in continuous space requires computationally expensive collision detection, whereas in lattice grid representations, the number of cells that may occupy a grid space can be limited, and determining the cells that lie within one cell's neighbourhood is trivial [Bogle & Dunbar 2010]. Ongoing research is seeking to develop algorithms providing cellular movement in a three dimensional lattice grid that replicates cell motility observed in the real-world [Bogle & Dunbar 2008].

A downside to the agent-based approach is the considerable computational resource required to execute realistic numbers of agents, or at least the numbers required for particular emergent properties to emerge [Bauer *et al.* 2009]. Time in agent-based systems is typically discretised, and in each time step a large number of entities must interact with the environment and each other, and then update their states, which can be computationally demanding. In response to this demand, a number of tools and programming languages are being developed that harness the recent technological advances in multi-core processors and networked clusters of machines for performing large-scale computation. One such example is the occam- π process-oriented programming language, being developed within the CoSMoS project [Welch & Barnes 2005]. Agents are represented as processes that have their own thread of execution; whereas a sequential language such as Java² might be constrained to using a single process to compute all agents at each time step, agents in occam- π may be computed in parallel on multiple processor cores, making use of sophisticated message passing techniques to communicate with one another [Andrews *et al.* 2008b]. Language technologies and program usage patterns to facilitate the execution of occam- π simulations on clusters are being developed, which include, for example, how to implement and manage discrete spatial representations in a manner that minimises the overhead of processes having to communicate across a traditional LAN network [Sampson *et al.* 2009].

Another criticism raised against the agent-based paradigm concerns its documentation [Bauer *et al.* 2009]. Whereas the equations underpinning ODE based systems can be documented in a concise, unambiguous manner, agent-based simulations consist of many lines of computer code written in any number of programming languages. However, Bauer *et al.*'s criticism might be better directed at the manner in which agent-based paradigm is practiced, than at the paradigm itself. Whist discerning the exact dynamics of an agent-based system by unravelling its code is a formidable task, several groups of investigators are making use of diagrammatic modelling languages in expressing the dynamics of their systems. The Harel & Cohen group employ state-charts in the specification of agent-based systems, and generate the simulation code directly from these diagrams [Kam *et al.* 2001]. [Andrews *et al.* 2010] has argued for the underpinning of simulation activities with rigorous modelling activities, not for the purpose of automated code generation, but for arguing that simulations are appropriate representations of their target systems and for transparency of simulation-based scientific investigation.

A purported benefit of the agent-based modelling paradigm over mathematical approaches lies in its intuitive nature. The modelling languages used to express agent-based systems are, for many experimentalists and clinicians, more approachable than ODEs and readily engages them in the modelling process [Bauer *et al.* 2009]. Its perceived parallels with the manner in which immunologists reason about the immune sys-

²The Java programming language specification is maintained by Sun Microsystems, [Gosling *et al.* 2005].

tem has lead to the promotion of the unified modelling language³ (UML) as a tool for the expression of agent-based simulation behaviours [Bersini 2006]. Indeed, Harel statecharts, on which UML state machine diagrams are based, was originally put forward for the modelling of complex systems [Harel 1987], and their similarity with the way that biologists reason about biological systems has been noted [Cohen & Harel 2007]. UML is a collection of diagrammatic notations originally intended for the specification of software systems. It is frequently used to specify systems to be implemented in languages of the object-oriented paradigm. Object-orientation, where system components are represented as discrete entities responsible for maintaining their own state, shares many parallels with the agent-based philosophy.

3.2.3 Hybrid techniques

Hybrid modelling and simulation approaches attempt to incorporate elements of both mathematical and computational paradigms, in an effort to find a balance between computational resource and the level of granularity required for a particular simulation domain.

An illustrative example is found in the stage-structured models of [Chao *et al.* 2004b, Chao *et al.* 2004a], wherein agent life cycles are divided into stages, and all individuals in a particular stage are assumed to be identical. A single integer indicates the number of entities associated with each stage [Chao *et al.* 2004b]. By using discrete, rather than real-valued, variables to represent population sizes, the approach avoids the undesirable artifacts associated with ODEs when populations sizes approach zero. The model does not explicitly represent space, and as such is considerably more computationally lightweight than a typical agent-based system. Probabilities dictating the transition of individuals between stages allows this approach to capture an element of stochasticity.

A popular hybrid approach entails the representation of cells as discrete agents, and lower-level molecules, such as cytokines, as continuous entities. [Ray *et al.* 2009] make use of a two dimensional lattice grid, wherein cells are able to traverse between neighbouring locations. Molecular diffusion, on the other hand, is modelled through partial differential equations. Representing a different hybrid approach, the CyCells simulation employs three dimensional representation of space, in which cells are represented as agents that move around continuous-space, whereas cytokines diffuse between locations on a lattice [Warrender *et al.* 2006].

3.3 Integration of wet-lab and *in silico* techniques

The present section reviews a selection of literature that makes use of modelling and simulation techniques to explore immunology. Special focus is given to the manner in which these works complement real-world experimental techniques in advancing the field of immunology.

For over a decade professors David Harel (a computer scientist) and Irun Cohen (an immunology) have collaborated on modelling and simulating biological systems, in particular the immune system. Their collaboration has culminated in an agent-based technology, called Reactive Animation, for specifying and implementing system behaviours, and providing a visual interface to the system that permits real-time manipulation of the underlying simulation [Efroni *et al.* 2005]. Visualisation is implemented

³www.uml.org, [Rumbaugh *et al.* 2005].

using Adobe Flash⁴, and the simulation engine is specified using Harel statecharts [Harel 1987], which are executed in a software package called Rhapsody.

Reactive animation has been used to simulate the development of T cells in the thymus [Efroni *et al.* 2003], focussing on the molecules and cytokines that dictate the movement of cells around this highly structured organ. The simulation is constructed in a bottom-up fashion, acquiring data from around 300 publications and consolidating it into one simulation. By performing *in silico* knockout experiments, the authors have made several predictions concerning the molecular basis, and the role of competition between cells, in key processes involved in T cell development [Efroni *et al.* 2007]. Two of the predictions to arise from this simulation work have been experimentally verified [Efroni *et al.* 2005].

A second example line of research that makes use of reactive animation technology focuses on the lymph node, again following the group's bottom-up approach to simulation construction and parameterization [Swerdlin *et al.* 2008]. The simulation is used to examine the role of lymph node size and antigen quantity on the number of plasma and memory B cells produced following the generation of the immune response.

The published work on reactive animation places a great deal of emphasis on visualisation, which both represents what is happening in the simulation, and allows the user to manipulate the simulation in real time. Users are able to affect behavioural changes in individual cells, or choose between different hypotheses governing cellular behaviour which influence entire populations. Users can then observe the higher-level effects of different hypotheses governing cellular behaviour, and as such can evaluate the high-level implications of different behavioural hypotheses. The strength of the reactive animation approach lies in its consolidation of large quantities of experimental data, and the ability to visualise it; the authors have argued that visualisation of simulations is key to engaging the minds of wet-lab researchers and gaining an intuition of system operation [Cohen 2007a].

Chao *et al.* employ their stage structured model of CD8Tc responses to antigen to investigate the effects of antigen kinetics on the pool of memory T cells produced [Chao *et al.* 2004a]. The authors conclude that non-replicating antigen (as might be administered in vaccinations) does not produce as wide a variety of antigen specificities in the resultant memory T cell population as does replicating antigen. Whilst the administration of replicating and non-replicating antigen could be accomplished *in vivo*, analysis of the entire repertoire of resultant antigen-specific memory cells could not.

In further investigation, [Chao *et al.* 2005] create a novel computational model of the T cell selection processes that occurs in the thymus that prevents large quantities of potentially self-reactive cells from entering the periphery. Their computational model represents TCRs, MHC molecules and peptides as strings, and simulates the selection process by deleting randomly generated TCRs on the basis of their ability to react to a set of other strings comprising the MHC:peptide complexes of the host. The authors conclude that the range of MHC:foreign-peptide complexes that can activate a T cell is related to the T cell's affinity for the MHC complex, and that this phenomenon may benefit the immune system in efficiently covering the large space of potential peptides with limited T cells. Conclusions are drawn on the basis of simple experiments which would nonetheless be highly challenging to conduct in the wet-lab.

Using a hybrid two-dimensional lattice grid simulation, [Ray *et al.* 2009] build on previous work ([Segovia-Juarez *et al.* 2004]) investigating the immune response to *mycobacterium tuberculosis* (TB), an infectious disease that kills more people per year

⁴Previously Macromedia Flash. It is a multimedia platform for use on web-pages.

than any other. The authors investigate the complex role of the cytokine TNF- α in the immune response, and its success in clearing the bacteria. TNF- α is a highly pleiotropic substance, having a wide range of effects on various cell populations. The authors conduct a series of novel and intricate *in silico* experiments that annul particular effects of this cytokine on particular cell populations, whilst leaving other effects in place. Such experimentation cannot be performed *in vivo*, where at best researchers can knock out the genes encoding TNF- α or its receptors, thereby having wide ranging effects on a large number of cell populations. The authors go on to conclude that it is the bacteria, rather than the cytokine, that is the prime driver of granuloma⁵ structure by performing another *in silico* experiment that cannot be conducted in the wet-lab: it is not possible to separate the effects of bacterial number and TNF- α . However the separation is relatively trivial *in silico*; the authors permit the physiological progression of the immune response, and then artificially fix the number of bacteria in the system at a certain point during the response, thereby separating the relationship between bacterial number and TNF- α secretion.

Simulation has been used to extrapolate beyond the limitations of wet-lab technology used to investigate interactions between T cells and DCs in the lymph node [Linderman *et al.* 2010]. Two-photon microscopy is a technique that allows imaging of living tissue, however it can only observe areas of around a square millimeter, and for time frames of no more than a few hours. The authors extrapolate beyond the limitations of wet-lab technology, to the level of an entire lymph node. They investigate the relationships between the frequency of particular antigen-specific T cells amongst the total population, the percentage of total DCs in the lymph node capable of priming T cells, and the expression levels of MHC:peptide complexes by those DCs. Their findings indicate that caution must be excised when interpreting wet-lab data using two-photon microscopy, where researchers typically artificially raise the number of T cells and DCs in the lymph node during an experiment. The authors go on to reflect upon the selective pressure on pathogens to evolve means of interfering with the processes that lead to T cell priming, and use their simulation to investigate the potential effects of manipulating particular combinations of processes. Once more, these experiments can not be engineered in the real-world system.

A second example of simulation used to extrapolate beyond the limitations of wet-lab experimental techniques arises from the Chakraborty lab, who do so upon wet-lab data in creating an agent-based simulation of T cell-dendritic cell (DC) interactions [Zheng *et al.* 2008, Henrickson *et al.* 2008]. The simulation work is closely tied to experimental work, and the synergy of the two approaches allowed researchers to investigate the nature in which DC density in the lymph node, and the quantity and duration of MHC:peptide presentation on DCs impact T cell motility and activation. Collectively, the authors conclude that there is compelling evidence that T cells integrate signals from multiple DC interactions before entering proliferation.

3.4 Previous work on modelling EAE

The only known work on modelling and simulating EAE to date is that of [Borghans *et al.* 1998]. The authors present a highly abstract ODE model of EAE, encompassing two variables: one representing the number of autoimmune T cells, the

⁵A *granuloma* is a specialised immune structure that forms through the aggregation of immune system cells in response to intra-cellular paracites, with the purpose of containing and combatting the invaders.

other representing the number of regulatory T cells⁶. The ODE model is used to investigate the phenomenon of T cell vaccination (TCV), a form of adoptive transfer intervention wherein activated auto-reactive T cells are taken from one experimental animal and administered into the other. If this is performed with either normal or attenuated auto-reactive T cells, the recipient animal becomes resistant to subsequent attempts to induce autoimmunity in it [Ben-Nun *et al.* 1981].

Two steady states, in terms of T cell population sizes, are identified in this system. The first describes the case of zero auto-reactive cells, and zero regulatory cells. The second reflects a balance reached between slowly proliferating auto-reactive T cells, which fuel the slow proliferation of regulatory T cells, which in turn prevent the further expansion of auto-reactive T cell number. The authors demonstrate that vaccination with normal auto-reactive T cells can bring the system from the default steady state of no cells to the balance state. If the dose of auto-reactive T cells is sufficiently large, then transient autoimmune symptoms may be experienced whilst the regulatory T cell population expands to sufficient number to reduce the auto-reactive T cell population. Hereafter, TCV with large numbers of normal auto-reactive T cells causes perturbation of the system, but not sufficiently so that autoimmune symptoms arise; the system is immune to further attempts to induce autoimmunity.

Using their model, [Borghans *et al.* 1998] predict that TCV with attenuated autoimmune cells will convey only transient protection from autoimmunity. This is because TCV induces regulatory T cells, but does not produce additional auto-reactive T cells since attenuated T cells cannot proliferate. TCV with normal auto-reactive T cells, however, switches the system to the balanced steady state, conferring long term protection against autoimmunity. The authors extend this prediction as a potentially verifiable wet-lab experiment.

The authors go on to demonstrate how the remitting and relapsing form of autoimmunity observed in many human cases of autoimmune disease may be explained through auto-reactive and regulatory T cells failing to reach the balanced steady state. Instead, these cell populations oscillate in number: expansion of auto-reactive T cells leads to expansion of regulatory T cells, which then aggressively reduce auto-reactive T cell number, which starves the regulatory T cell population of the simulation required to maintain their number, which then permits the further expansion of auto-reactive T cells, and so the cycle continues. This behaviour could be realised in their ODE model by particular manipulation of parameters, namely, the ability of regulatory T cells to abrogate the auto-reactive T cell population was increased.

3.5 Methodologies and building confidence in simulation

“Because extrapolation of results to situations outside the range of experiment is fraught with error, supplementation of experimental results with the in silico simulations is useful.” [Linderman *et al.* 2010]

This quote is taken from [Linderman *et al.* 2010], referring to their simulation-borne finding that artificially raising the number of T cells and DCs when performing two-photon microscopy experiments in the wet-lab can substantially perturb the observed system away from its physiological behaviour. However, it is suggested in this thesis that exploring natural systems through simulation is just as big an extrapolation, if not more so, than the wet-lab convention that the authors refer to. This section explores the

⁶CD4Treg and CD8Treg cells have been abstracted into a single entity. No other cells in the system are explicitly represented.

issue of establishing confidence in the results produced by immune system simulations, and the manner in which the literature has addressed it.

The benefits that modelling and simulation techniques can bring to the field of immunology are argued in a wide variety of prominent journal publications, for example [Forrest & Beauchemin 2007, Germain *et al.* 2011, Chakraborty & Das 2010]. An increasing number of journal publications report the use of computational methods along side traditional techniques in deriving their results [Kleinsteiner 2008]. Yet despite this, very few authors address, or even acknowledge, the issue of correctly interpreting simulation results in terms of the original domain. This is an important issue that requires address within the field, considering that an incorrect or inappropriate simulation might hold as much potential to mislead as it does to clarify.

Models and simulations of the immune system are, by necessity, abstract representations of the target systems. It is computationally and conceptually intractable to represent every aspect of the real world immune system, too little is understood of much of the field. Typically, a small subset of the host's cells are represented in a simulation, and at a substantially reduced number than that which engage in the real system's immune response. Entities in a simulation may be labelled 'T cells' or 'APCs', but they in fact represent a much wider set of real-world cells: is the 'T' cell a Th1 cell, a Th2 cell, a Tc cell, or any number of other T cell subsets that are continually discovered, and at which stage of its life cycle? The entities which comprise simulations do not represent real-world cells in full, but rather an amalgamation of real-world concepts, collectively abstracted to a level that is often arbitrarily determined. These abstractions are intended to be representative of the real-world phenomena being observed, and in doing so, the abstract entities must compensate for the activities of real-world components that are not explicitly represented in the simulation. It follows that parameter values do not translate directly from one domain to another; an immunologist may establish the lifespan of an effector CD8Tc cell, but this may not hold across all T cell subtypes. Furthermore, if the immunological observation was made, for example, *in vitro* by extracting such cells from an experimental animal, and placing them in the presence of APCs to prevent death by neglect, then the figures may still fail to exactly represent the lifespan of CD8Tc cells *in vivo* where they are subject to an abundance of regulatory mechanisms not present in the petri-dish. At best, parameter values acquired in the real immunological domain should be considered guidelines for their corresponding values in simulation, not exact figures.

These factors all impact on how representative a simulation is of its target domain. If a simulation is being used to gain insight into an immunological phenomenon, then the conceptual separation between the simulation and the target domain must be considered when interpreting and communicating results [Polack 2010]. Results from simulation must be appropriately framed in the real domain, avoiding conclusions about the real-system beyond what the simulation can truly attest to. The issue of establishing confidence in simulation results is an open question, subject to ongoing research in the wider community, and consideration in this thesis. It is also not a Boolean property; confidence in simulation cannot be established absolutely, it is intimately related to the simulation's purpose, and the experimentation performed upon it. A particular simulation at a particular level of abstraction may be suitable for investigating particular issues, but not others. Confidence must be established that a simulation's results are genuinely representative of the system that it simulates, given the nature of the problem that it has been used to investigate.

This section reviews the literature on agent-based simulation of the immune system, with a particular focus on the implications of current practice with respect to establishing confidence in simulation results. Section 3.5.1 considers the issue of pa-

parameterisation, and calibration of simulations against real-world systems. The use of sensitivity analysis in qualifying the significance of underspecified parameters is considered in section 3.5.2, which also touches upon the use of these techniques in exploring simulation behaviours. A popular approach to building confidence in simulation is to demonstrate the consistency of *in silico* results with those of real-world experimentation, which is the focus of section 3.5.3. Section 3.5.4 applies only to stochastic systems, and examines the use of statistical methods to gain representative results from simulation data. Finally, section 3.5.5 provides an example wherein two simulations of the same disease have provided contradictory information, re-affirming the present issue of confidence in simulation, and consolidating the sections that precede it.

3.5.1 Calibration, and the adoption of parameter values

As explored in this section, a widely adopted approach to simulation parameterisation is to adopt parameters directly from the literature, typically being taken from a variety of sources; some *in vivo*, some *in vitro*, and often from experimentation conducted in different labs using different systems or experimental models (for example, human versus mouse). Agent-based simulations in particular tend to encompass many parameters, the biological values of which are not always known. As such, modellers frequently have to derive these values through other means, typically opting to estimate the values themselves.

An alternative approach to pure estimation is to calibrate the unknown parameters. In the context of simulation, *calibration* is the activity of adjusting parameter values⁷ in order to align simulation behaviour with that observed in the target system. As such, parameters for which the corresponding biological data does not exist can be assigned appropriate values, with some justification for their assignment. Furthermore, the arguments above concerning the abstractive nature of simulations would suggest that all parameters should be subject to calibration, even if their corresponding biological values are well known; these parameters can be adjusted to compensate for the activities of elements of the real-world system that do not find explicit representation in the simulation. If there is a lack of consensus in the immunological data, then calibration can compensate for this disparity when determining parameter values; the value adopted is that which best aligns simulation behaviour with that of the real-world system, given the abstractions made in the simulation.

Despite the advantages that calibration would seem to bring to the parameterisation of simulations, it is an activity that is often not acknowledged or reported as having been performed in the literature. For example, the seminal works performed by the Cohen & Harel laboratory, in which data is consolidated from a very wide variety of sources⁸, do not report to have performed any calibration activities, despite stating that literature estimates for particular parameters vary by several orders of magnitude⁹ [Efroni *et al.* 2003, Efroni *et al.* 2005, Swerdlin *et al.* 2008, Efroni *et al.* 2007].

Where the literature does indicate that calibration has been performed, the details of which target data the simulation was calibrated against, or even the calibration procedure itself, are often not reported. To illustrate, [Ray *et al.* 2009] state that calibration was used in determining the values assigned to some parameter values, but do

⁷The calibration technique employed and reported in chapter 5 of this thesis extends beyond adjustment of parameter values, and considers adjustment of simulation mechanics also. Such use of calibration is not reported in any literature reviewed here.

⁸300 papers are reported to have been consulted in their simulation of T cell development in the thymus [Efroni *et al.* 2005].

⁹[Swerdlin *et al.* 2008] cite literature that reports the frequency of B cells specific for a particular antigen to lie in the range of 1 in 10k to 1 in 1000k.

not indicate which, or the calibration procedure followed, or the data against which calibration was performed.

A very recent ODE model of the immune response to tuberculosis infection contradicts these general trends [Marino *et al.* 2010]. Here, the authors publish the full details of a sophisticated model fitting algorithm, and the data against which their model was calibrated. The model itself is tied into the experimental work of one particular immunology lab, where a series of wet-lab experiments were performed to yield data against which the mathematical model could be calibrated. Employing data from only a single lab, thereby ensuring a high level of consistency, is again contrary to the general trend in the literature of adopting data from a wide variety of sources.

An interesting case in parameter value adoption is found in the agent-based model of influenza infection, published by [Beauchemin *et al.* 2005]. The authors report that all but two of their simulation's parameter values were adopted directly from another publication, an ODE model of the same disease ([Bocharov & Romanukha 1994]), which is stated to be more complex in representing a much wider variety of cells than the authors' present simulation. The authors of the ODE model report to have derived these parameter values through estimation followed by calibration. Hence, Beauchemin *et al.* directly adopt parameter values obtained through calibration of a different system, representing a different set of cells, and which makes fundamentally different assumptions concerning spatial representation¹⁰. We suggest that these parameters will not necessarily translate directly across to a new representation, from ODE to agent-based. It is noteworthy that, although Beauchemin *et al.* consider their model to be a sufficient analogue to that of the original ODE model, they report notable differences in behaviour between the two. A more rigorous process of calibration can aid in translating parameter values from one domain into another, compensating for differences in abstraction, and providing a justification for the values that are adopted. A further example of immunological simulation work that adopts parameters from other simulations can be found in [Swerdlin *et al.* 2008].

It has been argued in this section that calibration can aid in parameterising simulations, allowing values to be adopted across domains and adjusted to compensate for abstraction. It is especially useful in assigning suitable values to parameters for which the biological counterpart is not known. However, returning to the notion that one may hold varying degrees of confidence in simulation results, assignment of values to parameters for which the biological analogue is completely unknown is less compelling than the case where some guidelines may be extracted from the literature. The following section highlights the value of sensitivity analysis in addressing this issue. Sensitivity analysis techniques can establish which simulation parameters are influential with respect to its overall behaviour, and which are relatively inconsequential. If an underspecified parameter is found to have no substantial influence on simulation behaviour, then it can be argued that there is little cause for concern. If this is not the case, then the effects of underspecified parameter values aught to be considered when interpreting and relaying results of *in silico* experimentation.

3.5.2 Sensitivity analysis

Sensitivity analysis (SA) is an umbrella term for a collection of statistical techniques that attribute variation in a system's inputs to variations in its outputs. The techniques

¹⁰Coincidentally, it is Beauchemin who later conducts the comparison of ODE and agent-based simulations of the same disease in arriving at the conclusion that different spatial representations can have important impacts on the overall results of simulation, reported above in section 3.2.2 [Beauchemin 2006].

have implications for parameterisation of a simulation, and its exploration. They operate by varying inputs and correlating the resultant variations in output. *One-at-a-time* (OAT) SA techniques perturb only a single system input at a time, whereas in *global* analyses all parameters are perturbed simultaneously. If the extent of one simulation parameter's influence is dependent on the value held by another, then global techniques are more likely to reflect this fact than OAT methods [Saltelli *et al.* 2000].

SA techniques form a powerful basis for exploration of a simulation. Where a parameter is found to be highly influential, it may be concluded that the mechanisms or interactions that it parameterizes are important in the simulation's operation. These results can be highly relevant when, for instance, using simulation to identify potential targets in the system for combatting diseases or developing vaccines [Marino *et al.* 2008].

The techniques are also highly relevant to calibration and exploration of simulations of the immune system. The application of SA to a simulation's parameters can reveal how influential particular parameters are with respect to its higher level behaviours. This has implications for simulation parameterization, particularly where simulation parameter values cannot be extracted from the literature, and where estimations or calibrations are required to assign them appropriate values. If the real-world value corresponding to a particular simulation parameter is not known, and if SA reveals that this parameter is of little consequence to overall simulation behaviour, then assigning that parameter an estimated or arbitrary value may not impact the accuracy of the simulation. Conversely, if a parameter is shown to be highly influential to simulation output, yet the corresponding biological values cannot be ascertained to a narrow band of values, then a great deal of caution must be exercised when interpreting the results of said simulation. Although none of the literature surveyed has directly acknowledged this relationship, it is hinted at in the theoretical work of [Marino *et al.* 2008], who note that reducing uncertainty in parameter values can lead to stronger predictions.

SA techniques are well established in particular fields, but have found little application in the domain of immunological simulations. They are more readily applied to deterministic than they are to stochastic systems. Stochastic simulations, reviewed in much of this chapter, can require considerable computational resource when subjected to these analyses. The issue lies in their stochasticity: a deterministic system will always yield the same results given the same parameter inputs, however this is not the case for stochastic systems, where many executions are required to obtain representative averaged results. Thorough and comprehensive sensitivity analysis can require substantial exploration of a system's parameter space, and if each point in parameter space requires a large number of executions to obtain representative results, then the total computational resource required can quickly escalate by orders of magnitude.

The vast majority of immunological simulation works reviewed in compiling this chapter do not employ any SA techniques in addressing the issues raised above, an observation shared by [Marino *et al.* 2008]. There are, however, some exceptions, reviewed here.

In their agent-based investigation of influenza infection, [Beauchemin *et al.* 2005] conduct an informal SA of their simulation. Simulation performance is investigated whilst varying each parameter individually of the others, constituting an OAT approach. For each parameter in turn one value above and one below the default baseline value are investigated. These values relate to "their biologically allowed range" [Beauchemin *et al.* 2005]. The simulation's behaviours in each case are plotted and presented alongside one another. Though this analysis gives an intuitive appreciation of which aspects of the simulation are influential on its overall performance, there is no formal comparison between parameter sensitivities.

A second publication to employ a similar informal sensitivity analysis is

[Zheng *et al.* 2008]. Here the authors investigate the quantitative difference in simulation behaviour that arises from assigning parameters alternative values in an OAT approach. The values selected are seemingly arbitrarily determined. The focus of the analysis, in addition to providing some informal exploration of the simulation, is to demonstrate that simulation results from which the authors draw conclusions are stable under these variations in parameter values.

A more structured and comprehensive sensitivity analysis is applied to the work of [Ray *et al.* 2009], who employ a global SA technique in separating and understanding how several distinct effects that the cytokine TNF- α has on different cell populations contributes to the progression of the immune response to TB. A purely conventional use of this technique would attribute correlation coefficients to the influence of individual parameters to system outputs, however in the authors' novel application of this technique, the correlations with system output at various points during simulation time are determined. In this manner the authors identify that certain parameters, and hence mechanisms in the system, have effects that drastically reverse during the course of the immune response. In a more recent ODE modelling publication, authors based at the same lab again use SA to explore the role of various immune factors in granuloma development in response to TB [Marino *et al.* 2010]. However, in neither case do the authors consider parameter sensitivities in the context of the biological certainty underpinning the values they are assigned when interpreting simulation results into the original domain.

A highly novel method for using SA to explore behaviours across different scales represented in the simulation is provided by [Marino *et al.* 2008]. The authors suggest applications where, for example, one might examine how the number of T cells leaving the lymph node affects the bacterial load present in the lung. Hence, the authors highlight how the inputs to SA need not be simulation parameters, but may be emergent properties within the simulation.

3.5.3 Verifying simulation predictions to build confidence

A popular approach to instilling confidence in simulation results in the literature is to demonstrate that the simulation's behaviour is consistent with that of the target system under different experimental conditions. There are numerous examples to be found in the literature where a model or simulation has been constructed, after which its behaviour has been contrasted with that observed in real-world systems, a selection of examples of which are reviewed below. As noted above, confidence in simulation is not a Boolean all-or-nothing property. It is related to the use of the simulation; a particular simulation may be accurate with respect to certain classes of experiment, but not others. Demonstrating simulation consistency with real-world data can increase one's confidence in the system, but the practice is not without its limitations.

Firstly, the contrast between *in silico* results and real-world data is often subjectively made by the modeller, using no statistical techniques to quantify exactly how well the two data sets are aligned. Such techniques do exist, for example the sum of squared differences measure can be used to quantify the proximity between two data sets. Such a measure has been used outside of computational immunology in assessing the alignment of various models of stem cell dynamics with *in vivo* data [Wath & Liò 2008]. The measure returned by this analysis is highly dependent on the particular data sets being used, hence making comparisons between different data sets or models difficult. However, conventions of best practice and acceptable deviations could be established within the field. More applicable examples may be found in the form of statistical tests

such as the t test, or its non-parametric counterparts¹¹, which can quantify the difference between two distributions and indicate whether it is significant or not. Not only is the lack of application of objective measures of deviation between *in silico* and real-world data commonplace within the literature, there are examples in the literature where the authors explicitly note areas of deviation between simulation results and those of the real world, for example [Beauchemin *et al.* 2005, Warrender *et al.* 2006]. Without objective measures to quantify whether deviations are statistically or scientifically significant or not, it is difficult to assess their implications.

Secondly, immunology is a field that is rife with inconsistent or contradictory data, arising from different labs using different experimental models or species of animal, different experimental techniques, measured by different individuals with their own subjective dispositions. We argue that comparison of *in silico* results against data derived from multiple different sources is less reliable than when with data from the same source or the same lab. Example works that argue the consistency of *in silico* results with wet-lab results from a variety of labs include [Chao *et al.* 2004b, Swerdlin *et al.* 2008, Linderman *et al.* 2010]. Given the irregular nature of data across the immunological literature, a criticism that may be raised against post-simulation development selection of real-world data to contrast *in silico* results against is that the authors might be accused of having “cherry picked” the experiments that they select. This approach may be strengthened by either explicitly citing the experiments against which a simulation is to be contrasted *before* it is developed, or validating against experiments that are conducted in a single laboratory. In the latter case, the simulation can be tied to the work of a single lab. Two examples of this latter approach may be found in the literature. In a relatively unique example, [Marino *et al.* 2010] perform extensive wet-lab experimentation in order to acquire data necessary for parameterization of their model of the immune response to TB infection. Though the authors do not cite the benefit of acquiring data from a single consistent source as their motivation, the strength of their approach is noted here. A second example of tying simulation work in with wet-lab experimentation may be found in [Zheng *et al.* 2008], where simulation behaviour is frequently contrasted with wet-lab results from a single strand of work (published in [Henrickson *et al.* 2008]), and where the final conclusions of the research are based on results from both *in vivo* and *in silico* study.

A stronger approach to arguing that a simulation is an accurate representation of the target domain is by generating *in silico* predictions that are then verified in the wet-lab. There can be no accusation of cherry-picking experiments: the predictions are inherently linked to the simulation. This approach is advocated as best practice in a recent review paper [Bauer *et al.* 2009], and has been accomplished by [Efroni *et al.* 2007, Efroni *et al.* 2005] in their work on T cell maturation in the thymus, where two of the predictions arising from their simulation have since been validated. Once more, however, it is noted here that there are criticisms to be raised against this approach. Wet-lab experiments can be time-consuming and expensive to conduct. They can require the sacrifice of experimental animals, which in turn require that experiments be ethically certified. Perhaps most importantly, they require access to wet-lab researchers willing to carry out these experiments. The value of this approach in providing confidence in a simulation is acknowledged, however one might consider what happens if the prediction does not hold, given the effort required in per-

¹¹The Vargha-Delaney A test is a non-parametric effect magnitude test, which is introduced in chapter 6 [Vargha & Delaney 2000]. A review of basic statistics can be found in [Kranzler & Moursund 1999].

forming wet-lab experimentation. The modellers would be compelled to rethink and redesign their systems, after which further predictions may be generated, and wet-lab researchers may be re-appraised. It is not hard to conceive that experimentalists may grow weary of this process, should it persist. Though this approach certainly has merit in conducting simulation-based research, it is suggested that a rigorous, comprehensive simulation development methodology that serves to deliver confidence in simulations *as they are constructed* would be most valuable to the field. Such approaches, however, do not exist. The CoSMoS project is actively engaged in developing methodologies to assist in this manner.

3.5.4 Hypothesis evaluation, and acquiring representative results

This final section examining aspects of instilling confidence in *in silico* results is applicable only to stochastic simulations. As noted above, subsequent executions of stochastic simulations given the same input parameters can differ in the results that they generate. When performing *in silico* experimentation, one hopes to derive results that are genuinely representative of the system, rather than the result of stochasticity. In order to achieve this, many simulation executions are performed, and the results are averaged. However, the literature on immunological simulations contains no formal procedure or convention used to establish the number of simulation executions required to provide an acceptable level of accuracy in results.

The simulation-based research that originates from the Kirschner laboratory readily reports the use of averaging techniques applied over multiple simulation executions in deriving their results [Ray *et al.* 2009, Linderman *et al.* 2010]. In their work on simulating tuberculosis, [Ray *et al.* 2009] collect either 10 or 15 samples, dependent on the particular experiment being performed. When calculating simulation behaviours at various points in parameter space, conducted during their sensitivity analysis, the authors collect only 4 simulation executions per parameter set¹². In their work on lymphocyte priming dynamics in the lymph node, [Linderman *et al.* 2010] collect 10 samples. It is implicit that the authors of these studies consider this level of sampling sufficient to reduce experimental error for their purposes, however no procedure by which this has been determined is presented, and intuitively, the numbers of samples collected seem small.

The Cohen & Harel group's simulation work, which makes use of their visualisation-centric 'reactive animation' approach, reports no use of statistics to compile data across multiple runs [Efroni *et al.* 2007, Efroni *et al.* 2005, Efroni *et al.* 2003, Swerdlin *et al.* 2008]. The hypotheses that these simulations investigate are expressed at a coarse granularity of molecular expression; most cells either express particular molecules, or they do not, and for a few cells a single intermediate level of molecular expression is also represented. Their literature leads one to conclude that hypothesis evaluation is performed subjectively by examination of simulation visualisations. If deviation in simulation behaviour under the different hypotheses investigated in these works is very large, and consistently so, then this approach can suffice for the authors' needs. However, were one to investigate effects at finer granularities, varying molecular expression levels or durations by smaller quantities, then it is suggested here that a more rigorous and objective statistical framework would be required.

¹²[Marino *et al.* 2008] note that, for this particular global sensitivity analysis technique, a large number of samples in parameter space can compensate for "moderate" stochasticity in the absence of any averaging at samples in parameter space, but no relationships are given between degree of stochasticity and sample size.

3.5.5 The need for confidence in simulation results

Assessing the accuracy of a model or simulation's representation of the immune system is a critical issue that must be addressed in conducting *in silico* experimentation. It has been proposed here that simulations have as much potential to mislead as they do to clarify and inform. A specific example of this can be found in the works of [Ray *et al.* 2009] and [Marino *et al.* 2010]. Both examples, conducted at the same laboratory, model the immune response to tuberculosis. Ray *et al.* make use of a spatially explicit two dimensional agent-based representation, and model a small section of the lung, whereas Marino *et al.* make use of a large ODE system that models both the lung and a lymph node. Ray *et al.* demonstrate a major role for bacterial number in determining the size of granuloma, whereas Marino *et al.* find that bacterial number is relatively inconsequential. The inconsistency is noted in the later work of [Marino *et al.* 2010], who highlight the “importance of using different mathematical approaches to explore a specific biological question to uncover the effects of different process that are present in the system”. Had only one of these works been conducted, researchers may have accepted the potentially incorrect finding, and, if following the purported benefits that modelling and simulation bring to immunological research, designed further experimentation around this fact. We suggest that the discrepancy in the findings points to the pressing requirement for reliable and comprehensive validation techniques in creating immune system simulations. Could better methodologies for the construction and exploitation of simulation, and rigorous statistical comparison against established wet-lab data point to inappropriate abstractions or inadequate parameterization before *in silico* results are published to the wider community? At the very least, could such techniques indicate the degree to which simulation-based results are able to accurately represent the target domain?

To illustrate, [Marino *et al.* 2010] make an unusual assumption in not explicitly representing bacterial load in their ODE system. Rather, time series data of bacterial load in mice has been acquired from wet-lab experimentation, and this data is used as an input to the ODE system. Hence, the ability of the immune response to combat infection is not explicitly represented: the simulated immune system actions are not reflected in bacterial number. The authors have used sensitivity analysis (SA) to explore the influence of various system components on overall immune response, and have derived a method to include the time series data in the SA¹³. However, where SA alters a parameter, and hence the influence of a particular immune system component, the result on the immune response's ability to combat the pathogen is not realised; pathogen dynamics continue as they would had the change not been made. As such, the feedbacks operating through the immune system's response to changing bacterial number will not be invoked, and the true wider implications of the parameter perturbation will not be observed. The authors do not comment on the implications of this assumption, but it seems likely that it is significant. If so, then the application of discriminatory validation techniques might highlight this as an inappropriate assumption; it is believed here that the establishment of such validation techniques is essential for the progression of the simulation-assisted immunological research.

Citing various *in silico* investigations of T cell signalling pathways, [Chakraborty & Das 2010] divides computational studies into type I and type II simulations. The former describes simulations which aim to provide quantitative insight into the system, whereas the second focuses on qualitative contributions,

¹³It is this inclusion that has permitted their observation that bacterial number is relatively inconsequential to the number of cells recruited into the immune response, a corollary for granuloma size.

specifically, identifying novel pathways or formulating hypotheses to explain experimental observations. This thesis argues that, with the issues of simulation abstraction and underspecified data presented above, quantitative simulations are in even greater need of firm validation techniques than the qualitative variety. If one hopes to extract actual hard numbers rather than trends from simulation, then it *must* be a demonstrably accurate surrogate for the real-world system.

Though the conceptual separation has not been examined above, the issue of building confidence in a simulation's accurate representation of a target system may be thought of in two parts. Firstly, the mechanics of the simulation must be sufficiently representative of the target system to replicate its results, ideally under a wide variety of interventions or experiments. This relates to the entities in the target system, how they are represented or abstracted, and how they interact with one another. The second consideration is parameter values, the values that annotate the mechanisms of the simulation. For instance, the simulation's mechanics might dictate that cell B can kill cell A, and the parameters then dictate the probability of this happening, or how long is required for it to do so. Calibration and sensitivity analysis present potential strategies for instilling confidence at the level of simulation parameters. The question of a simulation's mechanics is more difficult to address. The difficulty in identifying an appropriate level of abstraction, which elements of the domain to include in a simulation, is noted by [Forrest & Beauchemin 2007]. There exists no guidance in the literature reviewed here with respect to this aspect of computational immunology. The approach advocated by [Andrews *et al.* 2010], reviewed in section 1.1.3 and which this thesis adheres to, is to approach computational immunology in an inter-disciplinary manner. Through close collaboration, modellers and domain experts can work together to ensure that a simulation is a sensible and justified representation of the target system, and that no obviously inappropriate assumptions are made. This aspect of confidence in simulation is revisited in the discussion chapter, chapter 8, of this thesis.

The modelling and simulation of EAE reported in this thesis is of the agent-based paradigm. As such, it permits the investigation of hypotheses at the level of molecular expression, and examination of spatial behaviours. Many of the above issues of building confidence in simulation are addressed. The simulation work is aligned with a murine model of EAE as used in one particular laboratory. As such parameter values are derived from the experience and knowledge of Dr. Kumar, who heads the lab.

3.6 Conclusion

This chapter has reviewed the existing literature on studying the immune system through the use of computational modelling and simulation techniques. Section 3.2 examines ODE and agent-based modelling paradigms, the two most popular computational methods employed in the literature. Select examples from the literature that illustrate exactly how computational methods have complemented wet-lab research have been explored in section 3.3. The only known work on modelling EAE is reviewed in section 3.4. Section 3.5 explores the separation between models and simulations, and the real-world immune system. It motivates the need to establish confidence that simulation results are representative of their target domains, and examines the manner in which the literature addresses this issue. Shortcomings are noted, and the work of this thesis in addressing them is motivated.

The next chapter details the construction of a domain model of EAE, the first stage in the CoSMoS process.

Chapter 4

Domain Model of Experimental Autoimmune Encephalomyelitis

The first stage of the CoSMoS process is the creation of a *domain model* that captures a comprehensive understanding of the domain [Andrews *et al.* 2010]. The domain is EAE, as detailed in chapter 2. This chapter serves two purposes: section 4.2 presents the latest iteration of the EAE domain model, largely expressed using the unified modelling language (UML)¹ [Rumbaugh *et al.* 2005]; and section 4.3 assesses the ability of UML to capture this complex immunological domain. Firstly, section 4.1 expands on the motivation for creating and maintaining a domain model, and the use of UML in doing so. Section 4.4 concludes the chapter. The work in this chapter contributes to the thesis aim by directly addressing research objective 1, exploring the role of domain modelling in the EAE case study.

4.1 Goal and motivation

The present chapter serves two purposes, the presentation of a domain model of EAE expressed using the unified modelling language (UML) [Rumbaugh *et al.* 2005], and an assessment of the strengths and weaknesses of UML in expressing this domain. The domain model presented in this chapter represents a key component in the CoSMoS process. It serves two purposes: to facilitate the exploration of the domain prior to the construction of a simulation, and to present a comprehensive and transparent understanding of the domain that underpins subsequent simulation-based experimental work.

The nature of the immunological domain dictates that there exist aspects of EAE that are unknown, or contested, within the literature. The rigorous modelling process of creating a domain model highlights areas of inconsistency or underspecification, which must be resolved through making assumptions of the domain. The model presents a medium through which the modeller can interact with the domain expert, facilitating and directing discussion, and it constitutes a mutually agreed scope of the domain that the simulation will represent. The domain model is validated by the domain expert, which helps ensure that domain-specific understanding that underpins *in silico*

¹www.uml.org

experimentation is reasonable. The demonstration that a simulation is an adequately faithful representation of the original domain is an open question, and is addressed here through validation of the domain model by the domain expert.

The second purpose of the domain model is in the facilitation of scientific transparency. In order for one to have confidence that a simulation is a faithful representation of the real-world domain that it purports to capture, its origins must be made explicit. The manner in which various aspects of the real-world domain have been captured within the simulation, and the assumptions made in doing so, are recorded. The domain model is a key component in linking simulation work to the real-world domain. The platform model, presented in the following chapter, provides additional details pertaining to the manner in which the domain model has been implemented.

The domain modelling process is iterative and ongoing, involving both modeller and the domain expert, reflecting the changes made to the simulation in facilitating open-ended investigation of the domain. The present domain model represents its most recent iteration.

The second half of this chapter provides an assessment of the applicability of the unified modelling language (UML) to representing a complex immunological domain such as EAE. The UML is a collection of diagrammatic modelling notations designed to aid in the specification and construction of software systems. Its notations allow for a wide range of specification scopes, from high-level system overviews to low-level focusses on particular elements of the system [Fowler 2004, Rumbaugh *et al.* 2005]. It can represent both static and dynamic views of a system. Static views depict the relationships that entities hold with one another, whilst dynamic views express the collaborations between system entities and the changes to their internal states that influence their external behaviours.

The multiple scopes for specifying a system encompassed within the UML has made it a popular modelling tool, and it finds application outside of the software domain within which it was conceived [Fowler 2004]. There have been numerous applications of UML to modelling biological systems, for example, [Garnett *et al.* 2008, Bersini 2006, Sadot *et al.* 2008, Andrews *et al.* 2008a, Efroni *et al.* 2003, Kam *et al.* 2001]²

The modelling of EAE presented here has provided insight into the strengths and weakness of the UML in representing various aspects of this complicated immunological disease. These observations form principles and practices that can inform others seeking to model immunological systems. Whilst other graphical notations for expressing biological systems exist, none have been accepted as a universal standard by the wider computational biology community [Le Novère *et al.* 2009]. It is this lack of standardisation, and the ambiguities and lack of universal applicability of existing notation, that has motivated the creation of the systems biology graphical notation (SBGN) which seeks to achieve universal acceptance [Le Novère *et al.* 2009]. Despite this, the SBGN is primarily targeted at biochemical systems.

Note that graphical notations for expressing biological systems do exist, motivated by the lack of a standard universally accepted biological notation the the systems biology graphical notation has been created [Le Novère *et al.* 2009].

4.2 Domain model of EAE

This section presents a domain model of EAE, largely expressed using the unified modelling language (UML). It has been developed based on the literature surveyed in

²These authors used state charts [Harel 1987] as their modelling medium. State charts are the basis for the state machine diagrams of UML.

chapter 2 and personal communication with the domain expert. The model is presented in a top-down manner, comprising three layers that depict the system at different levels of abstraction. They are:

1. A system-level overview of the domain model. This highly abstract layer details how the cells of the domain model interact in order to produce system-level behaviours, and how these behaviours are believed to correspond to phenomena observed in the real-world domain. This modelling stage does not make use of any UML notation.
2. Modelling of system *perspectives*: decompositions of EAE's complex onset and recovery. This modelling layer describes in greater detail the cell-level events and interactions that together constitute system-wide behaviours marking each stage of disease and recovery. These models are expressed using UML activity and class diagrams.
3. Cell-level dynamics. This layer represents the lowest-level modelling in the domain model, detailing the dynamics of individual entities in the system. UML state machine diagrams have been used in expressing these models.

Section 4.2.1 presents the top level of modelling. It delineates the scope of the domain model, and relates it to the real-world domain. Section 4.2.2 presents four *perspectives* of EAE. Lastly, cell-level dynamics are presented in section 4.2.3. Diversions from the standard UML notation are indicated in the sections that make use of them, and the standard UML notations are described in appendix section A.1 for reference.

4.2.1 Relationship between real domain and domain model

Figure 4.1 delineates the system of interest, and is termed an *expected behaviours* diagram. It does not conform to any UML notation. The diagram abstractly indicates those entities of the real domain that are represented in the domain model, how the interactions between them result in system-wide behaviours, and how these system-wide behaviours are believed to constitute the phenomena observed in the real domain. Hence, the figure indicates how the system described in the domain model relates to the real domain.

There exists substantial quantities of literature on EAE, and various aspects of this complex disease are independently studied by a wide variety of labs. It is computationally intractable to represent every aspect of the real-world system in models and simulations, and as such a subset is investigated. It is not claimed that other aspects of the system not represented in the domain model are irrelevant, only that those aspects of the system that are represented are believed to be *sufficient* for the manifestation of the phenomena observed in the real domain. Boxes annotated with '<<expected>>' tags represent emergent phenomena that are expected to manifest at the large-scale from low-level interactions of entities within the system. The entities themselves, and an abstract indication of their interactions with one another, are represented at the bottom of the diagram. The remainder of this section details how these domain model entities relate to cells and molecules found in the real domain. The section provides a high-level overview of how the emergent phenomena indicated on the figure manifest from the interactions of domain model entities. The emergent phenomena and the domain model entities themselves are elaborated upon in subsequent sections.

Experimentation carried out within Kumar's lab has lead to observations that categorise into three distinct phenomena. EAE as resulting from sub-cutaneous immunisation with myelin basic protein (MBP), complete Freund's adjuvant (CFA), and




Figure 4.1: An *expected behaviours* diagram depicting the phenomena observed in the real domain, and the behaviours manifesting from cellular interactions believed to be responsible for them.

pertussis toxin (PTx), leads to damage in the central nervous system (CNS). This leads to paralysis in the subject. Following the induction of EAE, the majority of experimental animals experience physiological recovery from paralysis; no experimental intervention is administered in facilitating recovery. Lastly, mice having undergone recovery from autoimmunity are resistant to subsequent attempts to induce paralysis with similar immunisation.

It is hypothesised that paralysis following immunisation with MBP, CFA and PTx is the result of T cell autoimmunity targeting myelin as expressed by cells in the CNS. Myelin is expressed by a variety of cells, and the domain model abstracts these into a single cell titled the *neuron*. It is believed that MBP is phagocytosed by dendritic cells (DCs) which prompts them to prime MBP-specific T cells. Multiple T cell sub-species may comprise these populations, including CD4Th1, CD4Th2 and CD4Th17³ cells. The domain model abstracts the actions of CD4Th1 and CD4Th17 cells into a single entity, termed the *CD4Th1* cell. CD4Th1 and CD4Th2 cells suppress the priming of one another's populations, largely through the secretion of type 1 and type 2 cytokines. There are numerous types of cytokine secreted by each cell type, and their effects are complex. They have been collectively abstracted into two domain model entities, *type 1 cytokine*, and *type 2 cytokine*.

Whilst both CD4Th1 and CD4Th2 cells gain access to the CNS, CD4Th2 cells do not contribute to paralysis. Upon infiltrating the CNS, CD4Th1 cells induce the activation of macrophages and microglia that reside there. Microglia ordinarily reside in the CNS, and the inflammatory context that follows the infiltration of CD4Th1 cells to the CNS results in the recruitment of macrophages to the site. These various APCs have been abstracted in the domain model, and are represented as a single cell type termed the *CNS macrophage*. The activation of CNS macrophages leads to the secretion of pro-inflammatory molecules such as TNF- α , nitric oxide, and reactive oxygen species. The domain model abstracts these molecules into a single entity titled *TNF- α* . TNF- α is harmful to neurons, leading to demyelination: the stripping of myelin from neurons. Neuronal myelin is not explicitly represented in the domain model, which instead represents damage to neurons as their entering apoptosis.

The spontaneous recovery from paralysis observed *in vivo* is believed to be the result of CD8Treg cells inducing apoptosis in MBP-specific CD4Th1 cells. The physiological turnover of CD4Th1 cells leads to their phagocytosis by DCs, and the subsequent priming of CD4Treg populations by those DCs. CD4Treg cells license DCs to prime CD8Treg populations; it has been established that the absence of CD4Treg cells in mice induced into EAE results in increased paralysis and delayed recovery from autoimmunity [Kumar *et al.* 1996]. It is believed that the role of these CD4Treg cells is to induce the expression of Qa-1 in DCs, required for the priming of CD8Treg populations. These CD8Tregs have been categorised as CD8 $\alpha\alpha$ TCR $\alpha\beta$ T regulatory cells by the Kumar laboratory, and are simply termed *CD8Treg* cells in this domain model.

The cell-level events that lead to protection against subsequent attempts to induce EAE are currently unknown. This domain model makes no claim concerning the manifestation of this phenomenon, and it is not explicitly represented.

The inter-cellular interactions outlined above, and indicated on figure 4.1, are done so at a very abstract level of detail. The following sections examine how these cellular behaviours manifest in system-wide emergent behaviours, termed *perspectives*, in greater detail. EAE paralysis and its associated recovery are characterised by consider-

³The role of CD4Th17 cells in EAE has not been explicitly investigated by the Kumar lab, but it is acknowledged that they likely do mediate autoimmune behaviour in their model. For reference on CD4Th17 in EAE, see [Zepp *et al.* 2011].

able complexity, and to facilitate exploration of how the cellular interactions contribute to the expected behaviours outlined in figure 4.1 each behaviour is divided into two *perspectives*. The resultant four *perspectives* describe the collective consequences of cellular interactions, and may be considered as four stages of EAE and its recovery. They are:

1. The *initial establishment of autoimmunity* in the CNS following immunisation.
2. The *self-perpetuation of autoimmunity*.
3. The *establishment of regulation* that results in the apoptosis of CD4Th1 cells.
4. The *type 2 deviation of autoimmune response* that results from regulatory activity and ultimately leads to the termination of both autoimmune and regulatory immune responses.

First, however, the spatial abstraction of a real-world experimental animal is detailed.

4.2.1.1 Spatial representation in domain model




Figure 4.2: The spatial components of the domain model, and the manner in which the cells of the domain model may migrate between them.

The present domain model abstracts the physical space of a real-world experimental animal into six compartments, depicted in figure 4.2. The manner in which cells are able to migrate from compartment to compartment is indicated.

The *periphery* represents the sub-cutaneous region of the experimental animal where immunisation with myelin basic protein (MBP), complete Freund's adjuvant (CFA) and pertussis toxin (PTx) are administered for the induction of EAE. Immunogenic type 1 polarised DCs expressing MHC-II:MBP peptides migrate from the periphery to the draining lymph nodes upon maturation. No other cells are considered to enter the periphery.

The *SLO* (secondary lymphoid organ) compartment is an abstraction of the lymph nodes that drain the region of the mouse where immunisation for EAE is administered.

In addition to the DCs that migrate to the SLO following immunisation, this compartment is considered to contain a small number of ordinary DCs. Any T cell that is not in either an effector or apoptotic state can migrate into this compartment from the circulatory system, and any T cell in any state may migrate from it back into the circulatory system.

The *CNS* compartment represents the central nervous system of an experimental animal. It contains neurons, CNS macrophages, and DCs. Effector CD4Th1 and CD4Th2 cells are able to migrate into this compartment from the circulatory system, however CD4Treg and CD8Treg cells are not. Neurons and CNS macrophages permanently reside in this compartment, and DCs migrate from the CNS to the CLN upon maturation.

The cervical lymph nodes (CLNs) of an experimental animal are lymph nodes in the neck that drain the CNS. These are represented in the domain model by the *CLN* compartment. It contains a small number of permanently residential DCs, and the mature DCs that migrate from the CNS compartment. Apoptotic CD4Th1 and CD4Th2 cells are able to migrate from the CNS compartment into the CLN. Any T cell that is not in either an effector or apoptotic state is able to migrate into the CLN from the circulatory system. Any T cell, in any state, is able to migrate from the CLN into the circulatory system.

The *spleen* compartment represents the spleen of an experimental animal. It contains a large number of permanently residential DCs. Any T cell in any state may enter or leave this compartment via the circulatory system.

The *circulatory system* compartment is the domain model's representation of the circulatory system of an experimental animal. Only T cells may reside in this compartment. It is through this compartment that T cells are able to migrate to the other compartments represented in the domain model, with the exception of a direct connection from the CNS to the CLN and from the periphery to the SLO.

4.2.2 Modelling perspectives

The following sections depict the manifestation of EAE *perspectives*, decompositions of the disease into four stages covering its onset and recovery, from cellular-level events and interactions. This is accomplished through the use of UML activity and class diagrams. Only standard UML notation is employed in class diagrams, however several additions have been made to the standard UML activity diagram notation, as described below. As explored more fully in section 4.3.3 below, these additions have arisen owing to the inability of standard UML constructs to satisfactorily represent important concepts in the domain.

Figure 4.3a depicts the standard state transition of a UML activity diagram. Activity diagrams that highlight the order in which activities are undertaken by entities in the simulation for some outcome to be realized make frequent use of this sequential relationship. It should however be noted, when interpreting the activity diagrams in this domain model, that the entity responsible for undertaking some action does not *necessarily* become inactive afterwards, as might be suggested by the sequential relationship arrow; it may continue to interact with other entities in the system. For example, a cell A may instill some effect in cell B, which may in turn lead to other activities taking place. However, cell A has not necessarily terminated, it may go on to perform other functions that are not depicted on the diagram, but that still contribute to the perspective being modelled.

A number of concepts relevant to the domain being modelled could not be expressed using standard UML activity diagram notations. As such, additional notation has been




Figure 4.3: The types of relationship used in activity diagrams.

created to represent these concepts. These are described below.

Figure 4.3b denotes a *propagation* relationship, conceived to demonstrate that an entity may perform an action that has some consequence elsewhere, but without wishing to imply that this activity stops. This relationship is indicated by an arrow originating from a line that is perpendicular to the border of the activity, but does not lie in contact with it. In the example, activity A leads to activity B. The occurrence of activity B has consequences represented in activities C and D, which may be considered as new paths of execution. However, the activity B does not terminate. Consequences C and D may each occur any number of times whilst activity B is undertaken. New occurrences of C and D cease once the transit from B to E is made. For example, activity B might represent a cell in a proliferating state, in which case action C could represent the creation of a daughter cell. That daughter cell would go on to interact with other entities in the system, independently of the proliferating cell. Several daughter cells may be created as a result of proliferation, before the proliferative activity ceases.

Figure 4.3c denotes a relationship that can be *interrupted* or down-regulated. Activity A leads to activity B, however activity C has the consequence of either fully or partially preventing that transit. For example, activity A could represent an activity undertaken by a population of cells, and that activity could lead to some consequence in another population of cells. Activity C may represent the secretion of some molecule that interferes with the ability of the population represented by A having the stated consequence on the population represented by B. The interference is not necessarily

absolute, it may be partial. Continuing the example, it may be that as a result of C, A has a diminished ability to affect B, rather than that the relationship is completely prevented.

Figure 4.3d depicts a *contributory* relationship. Activity A leads to a decision, B, after which either activity D or E will commence. The nature of that decision is influenced by activity C. C does instigate the decision (A leads to the decision), it only influences it when it occurs. In the example, the contributory influence of C on B is shown as being propagative, meaning that C may continue to influence many occurrences B, until a sequential transition away from C occurs.

The activities depicted on the activity diagrams that follow are largely cellular state changes and interactions that take place in particular compartments of the domain model. Dotted lines, representing activity diagram ‘swim lanes’, have been used to segregate groups of activities in accordance to the compartments in which they take place. This is demonstrated in figure 4.3e, in which activity A takes place in compartment X, and gives rise to activity B taking place in compartment Y.

Activities on the activity diagrams that follow are typically expressed at the level of a single cell. However, there are in all cases many such cells undergoing similar processes, and this is indicated in the text that accompanies the diagrams.

Note that some of the *perspectives* depicted are cyclic in nature, and hence do not contain end states. Start states have been given for cyclic diagrams in order to facilitate exploration of the diagram in the text.

The following sections model each of the four perspectives in turn. Section 4.2.2.1 depicts the *initial establishment of autoimmunity* in the CNS following immunisation. This leads to the *self-perpetuation of autoimmunity*, modelled in section 4.2.2.2. The manner in which the autoimmune response incites the regulatory response is described in section 4.2.2.3. Lastly, recovery is marked by a *type 2 deviation of the autoimmune response*, and this is depicted in section 4.2.2.4.

4.2.2.1 Initial establishment of autoimmunity

This section describes the onset of EAE autoimmunity, detailing the events that lead from immunisation to the apoptosis of neurons in the CNS. These events are depicted as an activity diagram in figure 4.4. Note that whilst the figure depicts activities at the single-cell level, there are in all cases many such cells undergoing similar processes.

EAE is induced in experimental animals through the sub-cutaneous administration of myelin basic protein (MBP), complete Freund’s adjuvant (CFA), and pertussis toxin (PTx). MBP is phagocytosed by DCs resident in the periphery where immunisation occurs, and this leads to their expression of MHC-II:MBP complexes. CFA and PTx are powerful immunopotentiators, and induce a type 1 polarisation in DCs, prompting them to upregulate co-stimulatory molecules. Following maturation, each DC will migrate to the draining lymph nodes, represented in the domain model as the *secondary lymphoid organ* (SLO) compartment.

Once in the SLO compartment, type 1 polarised DCs secrete type 1 cytokine. They come into contact with MBP-specific CD4Th cells, which bind to the MHC-II:MBP complexes they express. This binding delivers signal 1 to the T cells. The type 1 cytokines secreted by the DC prompts the CD4Th cells to primarily adopt type 1 polarisations, leading to their differentiation into CD4Th1 cells. Note that the adoption of either type 1 or type 2 polarisations by CD4Th cells is probabilistic, and not absolute. Some proportion of the CD4Th population will differentiate into CD4Th2 cells, despite the secretion of type 1 cytokine by the DCs on which they prime. CD4Th2 cells are not believed to be of any significant consequence at this early stage of disease onset,




Figure 4.4: UML activity diagram depicting the cellular interactions and events that lead to neuronal apoptosis in the CNS following immunisation for EAE.

and as such their actions are not explicitly detailed on the diagram. However, their behaviours are largely identical to that of CD4Th1 cells, as discussed in more detail in sections 4.2.2.4 and 4.2.3.1 below.

The CD4Th1 cells bind with the co-stimulatory molecules expressed by the DCs on which they prime, and derive signal 2. Receipt of both signals 1 and 2 leads a T cell to enter its proliferative cycle. During this cycle a CD4Th1 cell will divide and

produce a naive daughter cell, a process termed *spawning* in this domain model. Whilst in its proliferative cycle, a single CD4Th1 cell may produce multiple daughter cells. The majority of these naive daughter cells will immediately bind the MHC-II:MBP complexes expressed by the priming DC, and follow a similar sequence of events as the parent cell. Those that do not begin priming on the same DC as their parents assume migratory behaviour.

Once the proliferative cycle is complete, a CD4Th1 cell will differentiate into an effector T cell and will resume migratory behaviour, leaving the SLO compartment. These cells will eventually migrate into the CNS compartment. The CNS compartment is populated with CNS macrophages, some proportion of which express MHC-II:MBP complexes. This expression is the result of physiological turnover of neurons in the CNS, which the CNS macrophages residing there will phagocytose. The infiltrating CD4Th1 cells bind with MHC-II:MBP as expressed on these CNS macrophages, and become locally activated. Local activation prompts CD4Th1 cells to secrete type 1 cytokines, which stimulate CNS macrophages to secrete TNF- α . In sufficient concentration, TNF- α is harmful to neurons, and leads to their apoptosis.

4.2.2.2 Self-perpetuation of autoimmunity

Following the initial establishment of autoimmunity, a series of events lead to its self-perpetuation. The key events and interactions between cells that constitute this behaviour are presented as an activity diagram in figure 4.5. Due to the self-perpetuating nature of this autoimmunity the diagram has no end state.

Apoptotic neurons in the CNS compartment are phagocytosed by resident DCs, and by CNS macrophages. Note that any single apoptotic neuron can only be phagocytosed by either a DC or a CNS macrophage; the use of the fork relationship here represents the requirement for both DCs and CNS macrophages to phagocytose neurons in order for autoimmunity to self-perpetuate. Both DCs and CNS macrophages present MHC-II:MBP complexes as a result of phagocytosis. The type 1 cytokines secreted by CD4Th1 infiltrates in the CNS induce type 1 polarisation in DCs, prompting their expression of co-stimulatory molecules.

Upon maturation, CNS-resident DCs migrate to the CLN, where they present MHC-II:MBP to naive and proliferating CD4Th populations, delivering signal 1 to them. Their type 1 polarisation leads DCs to secrete type 1 cytokine. As a result of this, proliferating CD4Th cells preferentially adopt type 1 polarisations and differentiate into CD4Th1 cells. Again, this preferential adoption of type 1 polarisations is not absolute, even in the presence of DCs secreting type 1 cytokine some CD4Th cells will adopt CD4Th2 polarisations. CD4Th2 cells play no role in the self-perpetuation of the autoimmune response, and as such are not indicated on the diagram. The co-stimulatory molecules expressed by the DCs on which CD4Th cells prime deliver signal 2 to them, inducing these T cells to enter their proliferation cycles. Proliferating CD4Th cells spawn naive daughter cells, many of which are immediately primed on the same DC as their parent cells. Upon completion of their proliferative cycles, CD4Th1 cells differentiate into effector cells and resume migratory behaviour. This leads them to infiltrate the CNS compartment.

Effector CD4Th1 cells entering the CNS compartment are locally activated by CNS macrophages expressing MHC-II:MBP. Following local activation, a CD4Th1 cell secretes type 1 cytokine, which in turn stimulates CNS macrophages to secrete TNF- α . In sufficient concentration, TNF- α is harmful to neurons, and results in their apoptotic death. These apoptotic neurons are phagocytosed by CNS-resident DCs and CNS macrophages. Hence, autoimmunity self-perpetuates.




Figure 4.5: UML activity diagram depicting the cellular interactions and events that lead to the self-perpetuation of autoimmunity following neuronal apoptosis resulting from immunisation for EAE.

Figure 4.6 depicts a UML class diagram of the cells and molecules involved in both the instigation and self-perpetuation of autoimmunity. Immunisation with MBP, CFA, and PTx results in MBP being phagocytosed by DCs that are induced into immunogenic⁴ phenotypes. This interaction is modelled such that a single instance of MBP, CFA and PTx can only have these effects on a single DC, but that any single DC may engage in this relationship with zero or more instances of MBP, CFA and PTx. DCs also derive MBP from the phagocytosis of apoptotic neurons. A single DC or CNS macrophage may phagocytose zero or more neurons, and a neuron can be

⁴Expressing MHC:peptide complexes in conjunction with co-stimulatory molecules.




Figure 4.6: UML class diagram depicting the relationships between entities of the domain model involved in the establishment and perpetuation of the autoimmune response.

phagocytosed by at most one DC, or one CNS macrophage, but not both. DCs secrete type 1 cytokines if they are type 1 polarised, however they do not if they are type 2 polarised. As such, any single DC may secrete zero to many type 1 cytokine molecules. A single instance of a type 1 cytokine molecule is secreted by either a DC or a CD4Th1 cell, but not both. In addition to CFA and PTx, type 1 cytokine can induce a type 1 polarisation in a DC. Many type 1 cytokine molecules are required for this induction, and since this interaction does not destroy the molecule, a single type 1 molecule may engage in this relationship with several distinct DCs before it decays, however it may only do so with a single DC at a time.

DCs that have phagocytosed MBP may present MHC-II:MBP complexes upon maturation. Such DCs are able to prime many CD4Th cells, and as such are capable of presenting many MHC-II:MBP complexes. Any particular DC may present between zero and many co-stimulatory molecules, depending on whether expression has been induced. Perception of either CFA and PTx, or sufficient type 1 cytokine prompts this expression of co-stimulatory molecules. A co-stimulatory molecule is expressed by exactly one DC. Co-stimulatory molecules deliver signal 2 to CD4Th cells, and prime them to proliferate. However, a CD4Th cell is not necessarily a recipient of signal 2 in its lifespan. Co-stimulatory molecules have been modelled such that a single instance is sufficient to deliver signal 2 to a T cell. A CD4Th cell is either a recipient of signal

2, or it is not, and it cannot receive the signal more than once. Likewise, MHC-II:MBP is modelled in such a manner that a single instance is sufficient to deliver signal 1 to a CD4Th cell. A single instance of MHC-II:MBP is also sufficient to locally activate a CD4Th cell. Signal 1 can be delivered to a CD4Th1 cell only once, whereas local activation may occur multiple times. An MHC-II:MBP molecule is able to engage in these relationships with only one CD4Th cell at a time. However, it may participate in many such events with many different CD4Th cells over the course of its existence, since CD4Th cells that differentiate into effector cells migrate away from the DC and allow other cells to bind with it.

Upon receipt of signal 2, a CD4Th cell enters its proliferative cycle, during which time it may spawn between zero and many daughter cells. A daughter cell has only a single parent cell which spawned it. If a CD4Th cell perceives sufficient type 1 cytokine upon receipt of signal 2, then it adopts a CD4Th1 polarisation. Many type 1 cytokine molecules are required for this induction, and since the perception of cytokine molecules does not result in their destruction, any single molecule may engage in multiple such relationships. CD4Th1 cells secrete many type 1 cytokine molecules upon being locally activated. Simultaneous perception of sufficient numbers of type 1 cytokine molecules can stimulate a CNS macrophage into secreting TNF- α . Only CNS macrophages secrete this cytokine in this domain model, and as such a TNF- α molecule must have been secreted by exactly one CNS macrophage. Simultaneous perception of sufficient TNF- α molecules induces apoptosis in a neuron.

4.2.2.3 Establishment of regulation

The physiological lifecycle of T cells eventually leads all T cells to enter apoptosis, including MBP-specific CD4Th cells. It is the apoptosis of CD4Th cells that leads to the instigation of the regulatory T cell response that ameliorates autoimmune activity, as depicted on figure 4.7. Following their apoptosis, these CD4Th cells are phagocytosed by DCs. There is no constraint over which compartment a CD4Th cell resides in upon entering apoptosis, and they may be phagocytosed by any DC. CD4Th cells contain Fr3 and CDR1/2 peptides, which are components of the cell's T cell receptor (TCR). Upon phagocytosis of a CD4Th cell, a DC processes these peptides and is able to present them on MHC complexes. A pro-inflammatory cytokine milieu induces co-stimulatory molecule expression on a DC.

If the DC resided in the spleen during its immature state, then it remains there following maturation. Likewise CLN and SLO resident DCs do not migrate elsewhere upon maturation. If the DC resided in the CNS, then it migrates to the CLN. These migratory details are not indicated on the figure in the interest of reducing its complexity. Once mature, a DC will express MHC-II:Fr3, which leads to the delivery of signal 1 to naive CD4Treg cells that encounter it. The co-stimulatory molecules expressed by the DC deliver signal 2 to CD4Treg cells, causing them to enter their proliferative cycles. During this cycle any single CD4Treg may produce a number of naive CD4Treg daughter cells, which will either begin priming on the same DC as the parent cell, or resume migratory activity. Upon completion of its proliferative cycle, a CD4Treg will differentiate into an effector cell. As an effector cell, binding with MHC-II:Fr3 on a DC allows the CD4Treg to license the DC for Qa-1 molecule expression. Hence, the DC on which the CD4Treg was priming becomes licensed, however an effector CD4Treg may license any MHC-II:Fr3 expressing DC.

Once licensed for Qa-1 expression, DCs that have phagocytosed CD4Th cells are able to express Qa-1:CDR1/2 complexes. Naive CD8Treg cells that encounter the DC may derive signal 1 through binding with these complexes. Signal 2 is de-




Figure 4.7: UML activity diagram depicting the cellular interactions and events that lead to the instigation and perpetuation of the regulatory immune response.

rived through binding with co-stimulatory molecules expressed on the DC. As with CD4Tregs, CD8Tregs enter their proliferative cycles upon receipt of both signals 1 and 2, and produce naive daughter CD8Treg cells during this process. Upon completion of the proliferative cycle, a CD8Treg will differentiate into an effector cell and resume migratory activity.

CD4Th1 cells express Qa-1:CDR1/2 for around 8 hours following their differentiation into effector cells. During this period, an effector CD8Treg may bind with the CD4Th1 cell, and induce it into apoptosis. Apoptotic CD4Th1 cells, like any apoptotic




Figure 4.8: UML activity diagram depicting the cellular interactions and events that lead to a type 2 deviation of the immune response.

CD4Th cell, may be phagocytosed by a DC which derives Fr3 and CDR1/2 peptides from it. Both CD4Treg and CD8Treg cells secrete type 1 cytokine following local activation, and this contributes to the pro-inflammatory cytokine milieu that induces DCs into expressing co-stimulatory molecules necessary for the priming of further Treg populations. Hence, the regulatory immune response perpetuates, as is reflected in the lack of a terminating state on figure 4.7.

A class diagram facilitated discussion on the relationships between system entities responsible for the establishment of regulation may be found in appendix section A.2.1.

4.2.2.4 Type 2 deviation of the autoimmune response

This section describes the last of the four stages of EAE and its associated recovery: the deviation of the autoimmune response in a type 2 direction. Figure 4.8 depicts an activity diagram showing the cellular interactions and events that lead to the type 2 deviation. Once more the actions depicted here are cyclic, type 2 deviation does not

occur as a single atomic action within the system, but emerges as a gradual shift in behaviours spanning multiple populations of cells. However, the autoimmune response does eventually terminate, and hence an end state is indicated.

Immunisation for EAE induces a heavily type 1 response, with CFA and PTx prompting DCs to adopt a type 1 polarisation which causes priming CD4Th cells to preferentially adopt a CD4Th1 polarisation. As noted above, a DC's influence on the polarisation adopted by priming CD4Th cells is not absolute; although the majority of CD4Th cells primed on type 1 DCs adopt CD4Th1 polarisations, some still differentiate into CD4Th2 cells. Both CD4Th1 and CD4Th2 cells migrate to the CNS compartment, where they secrete type 1 and type 2 cytokines respectively. In the early stages of EAE, CD4Th1 cells outnumber their CD4Th2 counterparts, and as such the CNS cytokine milieu is composed primarily of type 1 cytokine. In sufficient concentration, type 1 cytokine induces apoptosis in neurons.

DCs in the CNS phagocytose apoptotic neurons, and upon maturation adopt either a type 1 or type 2 polarisation, depending on the balance of type 1 and type 2 cytokines in their local vicinity. Likewise, by default a DC will not express co-stimulatory molecules, these are induced in the presence of sufficient type 1 cytokine. A DC that expresses both MHC:peptide complexes and co-stimulatory molecules is immunogenic, and is able to prime T cell populations. Absence of co-stimulatory molecules renders a DC tolerogenic, in which case it induces anergy in T cells, preventing their proliferation and differentiation. Immunogenic DCs are either type 1 or type 2 polarising, which influences the polarisation adopted by the CD4Th cells that they prime.

The physiological turnover of CD4Th cells leads to their phagocytosis by DCs, and the subsequent priming of CD4Treg and CD8Treg populations. During the \sim 8 hours immediately following their differentiation into effector cells, CD4Th1 cells are susceptible to regulation, wherein CD8Tregs induce apoptosis in CD4Th1 cells. At a population level this regulatory action serves to reduce the number of CD4Th1 cells that infiltrate the CNS compartment. CD8Treg cells do not regulate CD4Th2 cells, which continue to enter the CNS and secrete type 2 cytokine. The reduction in CD4Th1 cells results in the cytokine milieu in the CNS shifting towards a type 2 dominance. DCs able to present MBP will adopt type 2 polarisations, and prime predominantly CD4Th2 cells. Once more, this adoption of type 2 polarisations by the CD4Th population is not absolute, and some CD4Th1 cells will continue to be primed, however many of these will be subject to regulation. As the quantity of CD4Th1 cells infiltrating the CNS reduces, so too does the quantity of type 1 cytokine being secreted there. This results in a reduction in neuronal apoptosis, and hence a reduction in MBP-presenting DCs migrating to the CLN. DCs migrating from the CNS to the CLN will fail to express co-stimulatory molecules, hence becoming tolerogenic and unable to prime T cell populations.

As the number of CD4Th cells primed in the system reduces, so too does the number of DCs presenting MHC-II:Fr3, and consequently Qa-1:CDR1/2. The regulatory immune response terminates alongside the autoimmune response, and experimental animals recover from paralysis.

Appendix section A.2.2 presents a UML class diagram-lead discussion on the relationships between system entities responsible for the type 2 deviation of the immune response.

4.2.3 Depicting single-entity dynamics

The following sections detail the complete dynamics for the various physical entities included in this domain model, at the single-entity level. Dynamics are expressed

using standard UML state machine diagram notation. In facilitating the expression of transition guards, the following symbols have been used:

‘&’: logical conjunction. $(A \& B)$ evaluates to true when both A and B are true, and false otherwise.

‘|’: logical disjunction. $(A | B)$ evaluates to true when either A or B are true, and false otherwise.

‘ $\delta(\text{condition})$ ’: used to indicate probabilistic events. Will evaluate to true with some probability.

‘ $\lambda(\text{condition})$ ’: used to indicate temporal events. Will evaluate to true when some period of time has elapsed.

Section 4.2.3.1 examines the single-cell dynamics of the various T cells represented in the domain model. Section 4.2.4 explores the dynamics of dendritic cells and CNS macrophages, and neurons are the subject of section 4.2.5. Lastly, section 4.2.6 focusses on the dynamics of MBP and the various cytokines represented in the domain model.

4.2.3.1 T cell dynamics


Figures 4.9, 4.10 and 4.11 depict state machine diagrams of CD4Th, CD4Treg and CD8Treg cells respectively. The three T cell types represented in the domain model share many similarities. They commence their life cycles in a naive state, in the circulatory system, and assume migratory behaviour. All naive T cells can migrate into the SLO, the CLN, and the spleen from the circulatory system, and vice versa.

Migratory behaviour continues until a T cell encounters a DC presenting MHC:peptide complexes for which it is specific. Thereupon the T cell will cease migratory behaviour and bind to the APC. It derives signal 1, leading it to enter a partially activated state. In the case of CD4Th cells, a polarisation is adopted at this point. These cells become either CD4Th1 or CD4Th2 cells, depending on the local cytokine milieu, largely influenced by whether or not the DC secretes type 1 cytokine. A predominantly type 1 cytokine milieu increases the probability that a CD4Th1 polarisation be adopted, whereas lack of type 1 cytokine promotes CD4Th2 polarisation.


If the DC to which the T cell binds expresses co-stimulatory molecules, then the T cell derives signal 2, and enters a proliferative state. Whilst in either a naive or partially activated state, a T cell may survive some period of time without receiving signal 1 and then signal 2 before dying of neglect. This is represented on the state machine diagrams as $\lambda(\text{neglect})$.

T cells remain in a proliferative state for some period of time, represented as $\lambda(\text{effector})$ before they differentiate into effector cells. Whilst proliferating, a T cell divides and creates naive daughter cells of the same type, a process termed *spawning* in this domain model. The time required for this to occur is represented as $\lambda(\text{prolif})$. Daughter T cells are created in the same location as their parents, indicated by the H state in the state machine diagrams.

Effector T cells detach from the DCs on which they prime, and resume migratory behaviour. If primed in either the CLN, SLO or spleen compartments, these T cells re-enter the circulatory system. Effector T cells lack the adhesion molecules required for them to migrate back into the CLN or SLO compartments, however they may re-enter the spleen. Effector CD4Th1 and CD4Th2 cells are able to migrate into the CNS compartment, where they remain until they enter apoptosis, upon which they may migrate into the CLN.




(a) Dynamics of CD4Th cells.




(b) Decomposition of the lymphoid organ states: the SLO, CLN and spleen.

Figure 4.9: State machine diagram depicting the dynamics of CD4Th cells.

Effector T cells require local activation before they may secrete cytokines. It requires some time following differentiation into an effector cell before a T cell may be locally activated through TCR:MHC:peptide interaction. TCR:MHC:peptide interaction prior to the elapse of this time, represented by $\lambda(localAct)$, may allow a T cell to perform some effector function, but not cytokine secretion. In the case of CD4Treg cells, TCR:MHC-II:Fr3 binding at any point during their effector cycles results in the CD4Treg cell attempting to license the DC with which it is bound. For CD8Treg cells, TCR:Qa-1:CDR1/2 binding with a CD4Th1 cell results in the induction of apoptosis in the CD4Th1 cell. After $\lambda(localAct)$ has elapsed, TCR:MHC:peptide binding induces cytokine secretion in effector T cells. CD4Th1, CD4Treg and CD8Treg cells secrete



(a) Dynamics of CD4Treg cells.




(b) Decomposition of the lymphoid organ states: the SLO, CLN and spleen.


Figure 4.10: State machine diagram depicting the dynamics of CD4Treg cells.

type 1 cytokine, whereas CD4Th2 cells secrete type 2 cytokine.

As with naive and partially activated T cells, effector T cells must receive stimulation from TCR:MHC:peptide binding regularly to avoid entering apoptosis through neglect. For those cells that reach locally activated states, effector lifespan is still limited, and sufficient stimulation of this form leads to activation induced cell death (*AICD*). Apoptotic T cells are eventually phagocytosed by APCs.



(a) Dynamics of CD8Treg cells.



(b) Decomposition of the lymphoid organ states: the SLO, CLN and spleen.

Figure 4.11: State machine diagram depicting the dynamics of CD8Treg cells.

4.2.4 DC and CNS macrophage dynamics

Dendritic cells are responsible for priming T cell populations. Their dynamics are depicted on figure 4.12. Dendritic cells begin their existence in an immature state, and mature some time later, represented by $\lambda(maturation)$.

In an immature state a DC is highly phagocytic, and the capacity to perform phagocytosis is reduced following the maturation of the cell. DCs may reside in five compartments in the domain model: the CNS, CLN, SLO, periphery, and the spleen. If the DC is in any of the CLN, SLO or spleen during their immature lifespan, then they do not migrate to away from these compartments upon maturation. However,




Figure 4.12: State machine diagram depicting the dynamics of dendritic cells.

upon maturation DCs residing in the CNS or the periphery migrate to the CLN and the SLO respectively.

For a DC to present MHC:peptide complexes, it must be capable of presenting both the MHC and peptide components. The peptide components considered in this domain model are MBP, Fr3 and CDR1/2. MBP is derived from phagocytosis of a neuron or from direct phagocytosis of MBP following immunisation. Fr3 and CDR1/2 are derived from the TCRs of CD4Th cells. In addition to phagocytosis of such cells, capacity to present peptides is probabilistic: only a very small proportion of phagocytosis events lead to the derivation of presentable peptides. This is represented as $\delta(\text{derive presentable peptides})$ on the diagram.

A DC is able to express MHC-II molecules upon maturation. Qa-1 molecules can




Figure 4.13: State machine diagram depicting the dynamics of CNS macrophages.

only be expressed once the DC has been licensed to do so by a CD4Treg cell. By default, a mature DC does not express co-stimulatory molecules upon maturation. Expression is induced by the perception of a sufficient concentration of type 1 cytokine, and may be induced through licensing by either a CD4Th cell or a CD4Treg cell. Co-stimulatory molecule expression may be induced at any point following maturation.

Upon maturation a DC adopts either a type 1 or type 2 polarisation. Type 1 polarised DCs secrete type 1 cytokine. The decision of which polarisation to adopt is entirely dependent on the balance of type 1 versus type 2 cytokine in the location of the DC upon maturation.

DCs expire once their lifespan is exceeded, represented as $\lambda(expire)$.

CNS macrophages exist only in the central nervous system (CNS). The only MHC:peptide complex that they present is MHC-II:MBP, and in a similar manner to DCs, the presentation requires the phagocytosis of a neuron, and is probabilistic. Some small proportion of CNS macrophages express MBP immediately, represented by δ (*basal expression*). This is to reflect the fact that the physiological turnover of neurons, which is not in itself represented in the domain model, will result in their phagocytosis by CNS macrophages, and the presentation of MHC-II:MBP complexes.

CNS macrophages exist in immature and mature states. Whilst immature they are more phagocytic than when mature. Maturation occurs some time into their lifespan, represented by $\lambda(\text{maturation})$, but may also be induced through perception of a sufficient concentration of type 1 cytokine. Perception of sufficient concentration of type 1 cytokine induces TNF- α secretion in CNS macrophages. Both immature and mature CNS macrophages are able to express MHC-II molecules.

CNS macrophages do not exist indefinitely, and expire after some period of time,

represented by $\lambda(expire)$.

4.2.5 Neuron dynamics




Figure 4.14: State machine diagram depicting the dynamics of neurons.

Neurons reside exclusively in the CNS compartment. Their dynamics are depicted on figure 4.14. All neurons express MBP. They are alive until perception of sufficient concentration of TNF- α induces their apoptotic death. Apoptotic neurons are eventually phagocytosed by an APC.

4.2.6 Cytokine and MBP dynamics




Figure 4.15: State machine diagram depicting the dynamics of myelin basic protein (MBP).

Myelin basic protein (MBP) is the substance to which encephalitogenic T cells are specific; it is integral to autoimmune activity. Type 1 and type 2 cytokines influence a variety of cellular behaviours that dictate the progression of autoimmunity and recovery. Although these entities of the domain model do not strictly carry state, state machine diagrams have been used to depict the system from their perspectives.

MBP dynamics are depicted on figure 4.15. It is introduced to the system through two means: it is manufactured and expressed by neurons, or it is injected into the




Figure 4.16: State machine diagram depicting the dynamics of type 1 cytokine, and its influence on other cells of the domain model.

system by an experimenter. Neurons that enter apoptosis are eventually phagocytosed by an APC. Likewise, MBP residing in extra-cellular space as a result of immunisation is eventually phagocytosed by an APC.

Once internal to the APC, there is a probability that MBP will be successfully processed for presentation as MHC-II:MBP complexes. This is represented as δ (*derive presentable peptides*). If MBP is successfully presented on MHC, then the MBP molecule is destroyed when the APC expires, if it is not, then it is considered to be destroyed immediately; it is of no further consequence to the system.

Type 1 cytokine, depicted in figure 4.16, is secreted by a variety of cells: CD4Th1 cells, CD4Treg cells, CD8Treg cells, and dendritic cells (DCs). Once secreted it exists in extra-cellular space for some period of time before it decays. Whilst in extra-cellular space, a type 1 cytokine molecule may be perceived by a variety of cells, and may influence their subsequent behaviours. Note that the induction of behavioural changes in cells requires the simultaneous perception of multiple cytokine molecules. Perception of a cytokine molecule by a cell does not destroy the molecule, and it continues to exist in extra-cellular space after dis-engaging from the cell.

Perception of sufficient type 1 cytokines by a CD4Th cell can result in it adopting a CD4Th1 polarisation, however this adoption is preferential, not absolute. The probability of a CD4Th1 polarisation being adopted is represented as δ (*type 1*). If type 1 cytokine is perceived as a sufficiently dominant constituent of the local cytokine milieu, then a maturing DC will adopt a type 1 polarisation. Type 1 cytokine perception also induces co-stimulatory molecule expression in mature DCs. Lastly, sufficient concentrations of type 1 cytokine induce TNF- α secretion in CNS macrophages.




Figure 4.17: State machine diagram depicting the dynamics of type 2 cytokine, and its influence on other cells of the domain model.




Figure 4.18: State machine diagram depicting the dynamics of TNF- α , and its influence on other cells of the domain model.

The dynamics of type 2 cytokine, and its influence on other cells of the domain model, are depicted on figure 4.17. The dynamics of type 2 cytokine are identical to those of type 1 cytokine, with the exception that it is secreted by CD4Th2 cells only. Its influences, however, differ. It has two influences over cells of the domain model. When perceived in sufficient concentration, it induces a preference for CD4Th cells to adopt type 2 polarisations, again this preference is not absolute. It may also induce type 2 polarisation in DCs if perceived as a sufficiently dominant constituent of the local cytokine milieau.

TNF- α , figure 4.18, is secreted only by CNS macrophages and exists in the CNS compartment. When perceived in sufficient concentration it induces apoptosis in neurons.

4.2.7 Temporal and numerical aspects of EAE

The three layers of modelling outlined above describe the cellular interactions and events that together constitute EAE onset and recovery, but they do not parameterise

these dynamics. Table 4.1 outlines various temporal aspects to EAE. It depicts the major population-level events that characterise EAE, and the times at which these are observed to occur *in vivo*. Also indicated are the times associated with certain cell-level events. The following items describe probabilistic and numerical details concerning the cells involved in EAE. These figures are derived through interaction with the domain expert.

- There exists around $\frac{1}{10}$ the density of DCs in the CNS as exist in the SLO and CLN compartments.
- At the peak of the regulatory immune response, there exist around $\frac{2}{3}$ the number of CD8Tregs as CD4Tregs.
- There exist around $\frac{1}{7}$ the number of CNS macrophages in the CNS as neurons.
- A DC is around 7 times the size of a T cell.
- CD4Th1 population expansion following immunisation is around 10 to 50 fold.
- Increase of DCs in the SLO compartment following immunisation is around 2 to 4 fold.
- The probability that the phagocytosis of a cell leads to the successful derivation and presentation of peptides is only a few percent.

There are a great many other aspects to the present domain model that require parameterisation, however these details are not known in the domain and can hence not be specified at this time. They are the subject of calibration activities as reported in the following chapter.

4.3 Reflection on the use of UML

The present section draws upon the experience of creating the EAE domain model in providing a commentary on the strengths and weaknesses of UML in modelling this complex immunological domain.

The commentary commences with section 4.3.1, which examines the benefit of a top-down approach to creating and documenting the domain model. Section 4.3.2 motivates the address of complexity through the derivation of *perspectives* of the domain's higher-level behaviours, and decomposing them into manageable sub-behaviours that can be modelled in turn. Section 4.3.3 reflects on the suitability of representing perspectives using activity diagrams, and highlights issues of representing the compounding concurrency of ever increasing numbers of cells that engage in activities that lead to the emergence of system-wide behaviours. The use of expanding regions, an activity diagram notation, in expressing this concurrency is covered in section 4.3.4. UML sequence diagrams represent another diagrammatic notation for expressing dynamic interactions between entities, and their suitability for depicting the manifestation of high-level behaviours from cell-level interactions is considered in section 4.3.5. Section 4.3.6 considers the value of modelling perspectives using class diagrams, but notes that an entity's history is important when considering the relationships that it may or may not engage in. In continuation of this, section 4.3.7 considers how the various temporal factors related to a particular relationship between two entities in the domain can lead to ambiguities when that relationship is expressed on a class diagram. Section 4.3.8 considers the expression of single-entity dynamics using state machine diagrams.

Time	Event
0 days	immunisation with MBP, CFA and PTx in the periphery
3-5 days	Detectable proliferation of CD4Th1 and CD4Th2 cells in the SLO
5-7 days	Detectable infiltration of CD4Th1 and CD4Th2 cells in the CNS
10-15 days	Visible paralysis of mouse
10 days	Detectable proliferation of CD4Treg and CD8Treg cells in CLN
30 days	CD4Th cells no longer found in CNS
30-40 days	Recovery from EAE

Event	Time
Duration of Qa-1:CDR1/2 expression on CD4Th1 cells following differentiation into effector cells	8 hours
Time required for spawning of daughter CD4Th1, CD4Treg and CD8Treg cells	1 day
Time required for spawning of daughter CD4Th2 cell	1.5 day
Time T cell spent in proliferative state	5 days
AICD death occurs in CD4Th1, CD4Treg and CD8Treg effector cells	5 days
AICD death occurs in CD4Th2 effector cells	8 days
Time a naive T cell spends in lymph node	12 hours
Time a naive T cell spends in the spleen	5 hours
Time a naive T cell spends in the circulatory system	30 min

(a) Times at which population-level events occur in EAE. The times given are estimates, not exact figures.

(b) The times associated with cell-level events in EAE. The times given are estimates, not exact figures.

Table 4.1: Temporal aspects to EAE.

4.3.1 Reflection on the process of creating a domain model

EAE is a highly complex, stochastic and dynamic disease. The system-wide behaviours such as paralysis and recovery manifest through large-scale low-level interactions amongst a wide variety of cells spanning several spatial compartments. System-wide behaviours emerge from the balance of population-level dynamics, and as a result of the many feedbacks that exist between cell populations: shifts in disease severity can be both swift or prolonged. The complex nature of this disease renders the task of deriving and presenting a domain model non-trivial. The top-down modelling approach reported above was arrived at following a variety of attempts to capture different aspects of the domain using different diagrammatic notations. The top-down approach represents the most successful manner of both exploring the domain in deriving the model, and presenting it to third parties.

The top-down manner in which the EAE domain model above is presented commences with a high-level system overview before expanding upon particular behaviours and dynamics. The purpose of constructing and maintaining a domain model is two fold: to conceptually explore the domain prior to engaging in simulation construction, and to present a comprehensive and transparent perspective of the domain that underpins the *in silico* experimentation that follows, such that it be open to scientific scrutiny. The benefit of adopting a top-down approach to modelling with respect to each purpose is considered in turn.

Domain modelling is an iterative task, and the model of EAE presented above represents the current iteration. Construction of the first iterations of the domain model, based upon domain-specific literature and interaction with the domain expert was performed in quasi-top-down manner; diagrams were iteratively amended as inconsistencies arose and additional information from the domain was brought to light. Nonetheless, a focus on top-down exploration of the system is advocated. The observation of phenomena in the real-world system, such as mice developing paralysis and undergoing spontaneous recovery, readily motivates one to construct hypotheses concerning the manifestation of these phenomena, and eventually design experiments aimed at confirming or falsifying them. When modelling, appreciation of hypotheses concerning the manifestation of system-level emergent behaviours arising from cellular activities and interactions serves to inform and guide the specification of single-entity level dynamics.

A top-down presentation of a domain model is considered a more intuitive approach to presenting a comprehensive and transparent perspective of the domain than a bottom-up presentation. EAE is a highly complex disease, arising from a large range of interactions between a multitude of cells spanning several spatial compartments. Inviting the reader to conceptualise system-wide behaviours from single-entity dynamics by reporting them first is believed to be considerably less intuitive than expanding upon system-wide behaviours by presenting the model in progressively increasing detail.

The top-down approach employed in constructing and presenting the EAE domain model comprises three levels of modelling:

1. System-level overview of the domain model components, how their interactions integrate to provide high-level behaviours, and how these high-level behaviours are believed to contribute to phenomena observable in the real-world domain. This phase indicates which aspects of the domain are to be incorporated into the model, since only a subset of the entire real-world domain is represented.
2. Decomposition of disease progression into various *perspectives*. System-level patterns in EAE disease are identified, termed *perspectives*, and their hypothesised

manifestation from cellular-level interactions are explored. This modelling level entails tackling the complexity of EAE by breaking it into more manageable subsets of behaviour.

3. Specification of single-entity level dynamics. The physical low-level entities of the system, such as cells and molecules, are considered as individuals and their behaviours specified.

Level 1 in the EAE domain model was graphically modelled in figure 4.1. It provides third parties with an immediate overview of the domain being modelled, the cells that are hypothesised to play a significant role in the real-world phenomena and the system-level behaviours that result from their interactions. The bottom layer of this diagram expands upon the behaviours abstractly indicated in the top layer, which indicates how they tie together to constitute the system being modelled. Note that this diagram does not conform to any UML notation.

The second level of modelling broaches the complexity of EAE onset and recovery by deriving several *perspectives* of disease progression. Experience indicates that the disease is too complex to capture on a single diagram, and this decomposition is necessary to relate different aspects of the disease to the cellular interactions that comprise them. The perspectives derived in modelling EAE relate to four stages of disease progression. Each of these perspectives has been modelled using UML activity and class diagrams.

The third level of modelling represents the highly stochastic and dynamic low-level entities of the domain. The cellular events and interactions that lead to the perspectives indicated in the second level of modelling are hypothesised to be the dominant behaviours exhibited by cells of the domain; owing to stochasticity, there may be great variation in the behaviours that are actually exhibited by individual cells. The full range of behaviours that cells and molecules can exhibit are captured through use of UML state machine diagrams in specifying entity dynamics.

The sections that follow examine the strengths and weaknesses of UML in representing particular aspects of EAE. The expected behaviours diagram of level 1 is not considered, since it does not adhere to any UML formalism.

4.3.2 Tackling complexity through perspectives

A primary purpose of the domain model is to convey understanding, and as such there is a clear motivation to avoid excess complexity on diagrams. The need to separate aspects of EAE disease progression into subsets of behaviour arose following attempts to capture the entirety of the disease's manifestation from single-entity interactions on a single diagram. This could not be accomplished as the full multi-dimensionality of EAE cannot be captured on a single two-dimensional diagram.

Figure 4.19 represents an attempt to capture every entity in the domain model as a class diagram. Relationship cardinalities have been omitted, yet despite this the diagram is marked with considerable complexity; there are 16 entities represented on the diagram sharing a total of 35 relationships between them. Each entity is connected to a mean average of $\sim 4\frac{1}{4}$ others, and the DC entity alone is connected to 11 other entities. Furthermore, many of these relationships are subject to constraints which have not been graphically depicted. For example, a particular type 1 cytokine molecule instance can only have resulted from secretion by a single cell: either a DC, CD4Th1, CD4Treg or CD8Treg. Information of this nature can only be captured as constraints, of which there are sufficiently many that depicting them all on a single diagram would render it incomprehensible.




Figure 4.19: UML class diagram attempting to capture every entity and relationship within the domain model on a single diagram.

The use of activity diagrams to describe the manifestation of system-level and emergent behaviours arose after attempting to capture how single-entity interactions lead to these behaviours as class diagrams. Figure 4.20 depicts an attempt to model the manifestation of autoimmunity as a class diagram. Depicting emergent properties of the domain as classes on class diagrams has previously been demonstrated by [Garnett *et al.* 2008], where the emergent property was the route through which molecules flow within the structure of a plant. In the present domain of EAE this approach of depicting an emergent phenomena such as ‘autoimmunity’ as a class on a class diagram is deemed inappropriate. An ‘autoimmunity’ cannot be instantiated as a physical entity, and as such it is unclear what cardinalities should be assigned to the ‘autoimmunity’ end of relationships.

Emergent behaviours within EAE, such as the establishment of autoimmunity in the CNS or the manner in which regulation results in type 2 deviation of the autoimmune response, are instead modelled as high-level behaviours, and termed *perspectives*. The manner in which perspectives manifest through low-level cellular interactions are depicted on activity diagrams. The EAE domain model details four perspectives, which




Figure 4.20: An example sequence diagram showing the interactions between participants A, B and C over time.

are derived as various stages of disease onset and recovery. The *expected behaviours* diagram of figure 4.1 contains only two high-level behaviours, covering autoimmunity harming neurons and regulation ameliorating autoimmunity. Each of these ‘expected behaviours’ is divided into two perspectives; this is in the interest of maintaining the clarity of the expected behaviours diagram: indicating each of the four perspectives and how the cells of the system relate to each of them would result in an excessively cluttered diagram. It is another example of the top-down manner in which EAE is modelled and presented.

4.3.3 Modelling perspectives as activity diagrams

Each perspective has been modelled as a UML activity diagram, which treats the perspective as a higher-level behavioural system property and details how these behaviours are believed to manifest through cell-level events and interactions.

UML activity diagrams permit any abstract concept to be represented as an activity, and are thereby highly amenable to depicting the order in which cellular events and interactions occur in constituting a higher-level behaviour. However, activity diagrams do not convey the large scale concurrency and stochasticity that exists in the real domain. The cellular events and interactions depicted on perspective activity diagrams are mostly expressed at the single-cell level, however for the perspective being modelled to occur in the real domain requires that a large number of cells engage in the activities being expressed. Furthermore, activity diagrams should not be interpreted as implying that all cells exhibit the behaviours indicated, cells are highly stochastic entities, and different individuals of the same type may experience vastly different paths through their possible range of dynamics. Manifestation of a particular perspective requires that sufficiently many, but not necessarily all, cells follow the activities indicated on the diagrams. Lastly, the cells that do engage in the activities depicted on the diagrams do not necessarily do so simultaneously, at each point in time there may be many populations of cells undertaking each of the activities depicted on the diagrams.

Activity diagrams imply a sequential transfer of execution that does not exist in the real domain where, as noted above, cells act concurrently and will continue to perform functions after the concept described by an activity has occurred. Some of these concurrent activities are particularly important to the perspective being modelled, and as such have lead to the development of the *propagating* relationship, depicted in figure 4.3. Similarly, EAE disease and recovery is mediated by feedback loops between networks of cells. The stimulus for the regulatory immune response is related to the intensity of the autoimmune response, and since the regulatory response ameliorates the autoimmune response, these two system-wide behaviours are coupled. Changes in the intensities of these two immune responses arise from population-level dynamics, which may be cyclic in nature, and are hence difficult to depict on activity diagrams.

The desire to depict these relationships, the ability of one activity to interfere with the succession of a second activity to a third, and the influence that one activity has on the particular outcome of another, motivated the *interruptible* and *contributory* relationships depicted on figure 4.3.

The concurrencies and feedbacks existing throughout the populations of cells involved in EAE, when expressed using activity diagrams, have resulted in forks that do not necessarily have corresponding joins. Under the strict ‘token passing’ interpretation of execution through these diagrams, many would never truly terminate. Whilst the corresponding immune responses in the real-world often do terminate eventually, they may correspond to a huge number of instantiations of all the activities depicted on the diagram, and hence this lack of termination is not entirely a mis-representation. The cyclic nature of the perspectives being modelled has lead to activity diagrams that contain no end state, and could be considered to have no start state, though these have been provided to facilitate exploration.

Whilst some activities depicted on activity diagrams may occur in any spatial compartment, there are others that occur in only certain locations. Activity diagram ‘swim lanes’, dotted lines that segregate activities on the diagram, are ordinarily used to indicate responsibility for an activity. In the present domain model these dotted lines have proven informative in indicating the spatial compartments in which particular events and interactions take place.

4.3.4 Depicting compounding concurrency with expansion regions

The use of activity diagram expansion regions in depicting the compounding concurrency in cell populations has been considered. Expansion regions are a notation used in activity diagrams to indicate multiple invocations of some activity or activities. The region marks an area of the diagram where actions occur once for each item in a collection comprising the region’s input. Inputs and outputs are denoted using small adjoining boxes.

Figure 4.21a shows an example expansion region, where the output of A leads to multiple instances of B leading to C. All invocations of B and C must complete before activity D is undertaken. A shorthand compact notation for a single activity being invoked multiple times is shown in figure 4.21b, where activity B is executed multiple times before C is executed.

Figure 4.21c depicts three different applications of expansion region concepts to a simple hypothetical perspective from the EAE domain, where immunisation leads to the maturation of many DCs, each of which go on to prime many T cells. The first example, (I), shows the encapsulation of expansion regions, reflecting the fact that each of many DCs primes many T cells. When applied to a larger and more realistically scoped perspective however, entailing many more compounding concurrent activities, this notation has the potential to add significant complexity to the diagram. Furthermore, with cyclic paths in perspectives, it is not clear where the termination of the expansion region should lie.

The second example, (II), makes use of the compact notation. The output collections have been omitted in an attempt to indicate that the activities that follow an expansion region do not have to wait for all the invocations of the region to complete. However, whilst the number of outputs from a region does not have to equal the number of inputs, this discrepancy indicates that the region acts as a filter, with some activities being dropped [Fowler 2004]. This is not the intended interpretation. As such, the third example, (III), makes use of multiple outputs from a region containing only a single input, to indicate that many invocations may follow the activity.




Figure 4.21: Examples of activity diagram expansion regions, and their potential application in depicting compounding concurrency.

Although the third example appears to most closely and match the intended concept, it fails to do so satisfactorily. The undesired implication of a sequential transfer of control remains, that T cells cannot commence being primed until all DCs have matured as a result of immunisation. It is more likely that these activities are ongoing and overlapping. Expansion regions are not currently used in the EAE domain model, however their potential is noted and further exploration of their semantics is warranted. However, this lies outside the scope of the current thesis.

4.3.5 Modelling perspectives as sequence diagrams

UML sequence diagrams offer another diagrammatic notation for expressing how system entities collaborate in some dynamic behaviour. A simple example sequence diagram is presented in figure 4.22. Like activity diagrams, sequence diagrams allow for the depiction of concurrent and alternative paths of execution (not shown).

The possibility of modelling perspectives using sequence diagrams in place of activity diagrams has been considered, however this was decided against. The depiction of time linearly along only one dimension would prevent the expression of cyclic groups of events, which appear in all of the perspectives modelled as activity diagrams. Fur-




Figure 4.22: An example UML sequence diagram, depicting the messages passed between participants A, B and C.




Figure 4.23: The use of UML class diagram composition relationships depicting mutually exclusive expression of MHC:peptide complexes on cells.

thermore the vertical notation depicting ‘lifelines’ during which messages are passed between participants implies that participants pass control to one another and wait until other lifelines have completed before they are themselves able to continue. This is in contradiction with biological reality, where entities are entirely concurrent and do not block one another from being able to interact with one another. Although activity diagram sequence relationships also imply a transfer of execution, it is felt that the sequence diagram notation, and its implication of participants being ‘blocked’, is even less appropriate.

4.3.6 Modelling perspectives as class diagrams

Activity diagrams are effective in indicating cellular-level events and the orders in which they occur for higher-level behaviours to manifest. The activities depicted on such diagrams are typically expressed at the single-cell level, despite the fact that populations of cells are required to perform these actions in order for system-level properties to emerge. Activity diagrams offer no means of specifying the quantities of entities that engage in the particular activities depicted on the diagram, and as such class diagrams have been created in an effort to provide this information.

There is some redundancy between the information conveyed between class diagrams of EAE, as evidenced by the lack of a class diagram for the first perspective, the manifestation of autoimmunity. Once drawn, the information on this diagram was sufficiently similar to that of the self-perpetuation of autoimmunity perspective that the two were combined into a single diagram. Some concepts, which are expressed on multiple diagrams, are common throughout the domain model. Examples include the fact that a cytokine can only be secreted from one source, or that a particular MHC:peptide complex instance is expressed by only one of the several cell types able to express it. Such constraints have been expressed in the text accompanying the class diagrams; the issue of rendering a class diagram excessively complex through the depiction of large

numbers of constraints was noted above.

The mutually exclusive nature in which cell-surface molecule instances such as MHC-II:MBP can be expressed by only one of the multiple cell types capable of expressing them is effectively captured using composition relationships, as demonstrated on figure 4.23. These relationships have ‘no sharing’ semantics that entail only one composition relationship may be exercised by a particular entity instance at a time [Fowler 2004]. Furthermore, the destruction of the ‘owner’ instance entails the destruction of all ‘owned’ instances; the apoptosis of a cell includes the destruction of all MHC:peptide complexes it may be expressing. However, the same does not hold for cell-cytokine secretion relationships, where a cytokine instance can only be secreted by one of several cell types, but the apoptosis of the cell does not entail the destruction of the cytokine instance.

The generalization relationship has been used in only one context in the perspective class diagrams, denoting that both CD4Th1 and CD4Th2 cells are examples of CD4Th cells. Conceptually, generalization may be considered to have no place on a domain model class diagram. CoSMoS domain models are not intended to provide software specification, where generalization is a fundamental aspect of quality coding practice. At the biological level, cells and molecules have no notion of generalization, however these concepts can prove useful when attempting to abstract behaviours into a domain model. They can serve to reduce the number of associations on diagrams. In the EAE domain model generalization has been used to indicate that CD4Th1 and CD4Th2 cells arise following differentiation of CD4Th cells, but that both contain Fr3 and CDR1/2 peptides necessary for the priming of Treg populations (e.g. figure A.4).

EAE is a highly dynamic and stochastic system, particular instances of cells may take very different paths through the dynamics depicted on their state machine diagrams. Cells may engage and disengage with many other cells through the course of their lifespans. Some might do this many times, others might never participate in a particular interaction that they are capable of. These dynamic aspects of EAE are not well represented on class diagrams, where the notation may hinder rather than support comprehension of the domain being modelled. Class diagrams are intended to represent static relationships between entities, yet cells, molecules and their interactions are inherently dynamic. This can lead to ambiguity when interpreting class diagrams, as detailed in the following section.

4.3.7 Temporal ambiguity on class diagrams

When interpreting an association and its cardinality on a class diagram, which together specify the nature of the relationship between two entities, the time frame over which the relationship holds is not necessarily apparent. A cardinality is used to indicate the number of instances that engage in a particular relationship. However, there are multiple temporal aspects of a relationship that are relevant to the EAE domain, and these cannot all be simultaneously graphically depicted. Hence, interpretation of a cardinality can be ambiguous. To illustrate, figure 4.24a shows the relationship between a type 1 cytokine molecule and a CNS macrophage cell. A type 1 molecule ‘stimulates’ a CNS macrophage, and both ends of the relationship are annotated with ‘0..*’ cardinalities. The following items illustrate the relationship between a type 1 molecule and a CNS macrophage, and how different aspects of this relationship result in the requirement for different cardinalities.

- Over the course of its lifespan, a CNS macrophage may perceive any number of type 1 cytokines, depending on the state of the simulation. Some CNS

macrophage instances may perceive none, if their lifetime corresponds to a period of no immune activity in the CNS. Others may perceive a great many over their lifespans. This suggests that cardinality A should be ‘0..*’.

- For a CNS macrophage to become stimulated, it must simultaneously perceive a number of type 1 cytokines exceeding some threshold. This suggests cardinality A should be ‘*’. However, this does not reflect that fact that not all CNS macrophages will be stimulated during their lifetimes. A cardinality of ‘0..*’ could be interpreted as reflecting the optional nature of this relationship, but it does not capture the notion that stimulation requires simultaneous perception of a threshold number of cytokine molecules.
- Perception of a cytokine molecule by a cell is not considered to destroy it, following perception a molecule may disengage from a cell. Hence, any particular type 1 cytokine may engage in stimulating many CNS macrophages over the course of its existence, which suggests a ‘0..*’ cardinality for B. However, a cytokine molecule can only be perceived by a single cell at any single point in time, suggesting a ‘0..1’ cardinality.


A further example of ambiguity concerning the timeframe over which relationship cardinalities apply may be found on figure 4.24b, which depicts the relationship that a CD4Th entity has with itself, describing the manner in which cells proliferate. It is tempting to interpret cardinalities as the number of instances that engage in the relationship at any single point in time. Under this interpretation, this relationship would indicate that a cell may provide any number of daughter cells simultaneously, for example that it may divide 5 ways to produce 5 daughter cells. However, this is biologically incorrect, a daughter cell is produced through cell division, and only one such daughter cell may be produced at a time. The relationship is intended to describe the fact that a CD4Th cell has only a single parent cell, but that a parent cell may spawn any number of daughter cells during its proliferative cycle.

EAE is a highly dynamic and stochastic domain, cells engage and disengage frequently. Some instances of entities may engage in particular relationships many times, others not at all. There are many temporal aspects of the relationship that are relevant to the domain, and UML class diagrams present no means to reflect this. Each end of a relationship is permitted only one cardinality. There are no guidelines on the temporal domain over which a cardinality applies, and further, in EAE this is specific to the particular relationship. Modelling dynamic and stochastic systems such as EAE using a notation intended for specifying static relationships can lead to ambiguities that confuse interpretation of the domain, rather than supporting it.


4.3.8 Capturing single entity dynamics

UML state machine diagrams have been used to specify low-level behavioural dynamics, and have been constructed for all cell types in the domain model, and for molecules that influence system dynamics.

Many of the biological cells of this domain model, such as the T cells or the DC, have complex multi-dimensional dynamics. Further, these dimensions are not necessarily completely orthogonal. These features of cellular dynamics can complicate the construction of state machine diagrams that describe them; expressing high-dimensional partially-orthogonal information on a two-dimensional diagram is challenging. For example, figure 4.9 captures the dynamics of CD4Th cells as a state machine diagram.



(a) The relationship between type 1 cytokines and CNS macrophages. The cardinalities have been labelled ‘A’ and ‘B’ to facilitate discussion in the text.



(b) The relationship that a CD4Th cell has in producing additional instances of itself

Figure 4.24: Select relationships between entities in the EAE domain model.

The locations in which CD4Th cells may reside are depicted as a mutually exclusive set of states that are orthogonal to the rest of the cell’s dynamics, such as its states of maturation. However, these sets of states are not entirely mutually exclusive: state transitions resulting from TCR:MHC-II:MBP binding can only occur when the T cell resides in the SLO, CLN, spleen or CNS compartments; there are no APCs in the circulatory system. The diagram is not incorrect, since cross-reference with CNS macrophage and DC state machine diagrams (figures 4.13 and 4.12) reveals that these APCs cannot reside in the circulatory system. Depicting these constraints diagrammatically on the CD4Th cell state machine diagram would increase its complexity and hinder its comprehensibility, which undermines the goal in presenting a transparent and informative domain model. UML state machine diagrams allow for transitions to be guarded, and this notation has been employed in dealing with partial-orthogonality. For example, figure 4.13 depicts the conditions for maturation in a CNS macrophage. The transition from ‘highly phagocytic’ to ‘reduced phagocytic capacity’ is guarded by ‘maturation’, being allowed to occur only when a CNS macrophage matures.

State machine diagrams are effective in capturing dichotomous and categorical concepts, where states do not overlap. Molecular expression in the present domain model has been represented in this manner, a DC either expresses MHC-II:MBP or it does not. However, were the domain model to be more representative of the underlying biology, MHC expression levels might be modelled as a variable quantity. There is evidence to suggest that DCs express low levels of MHC-II whilst immature [Kindt *et al.* 2007], and that expression is up-regulated following maturation. Naive T cells require more TCR:MHC bindings with an APC to derive signal 1 than effector cells require to derive local activation. Continuous-domained aspects such as levels of molecule expression or rates of cytokine secretion cannot be captured using UML state machine diagram notation, and would require equations or textual accompaniment.

Two other concepts have proven highly relevant to modelling EAE, yet UML state machine diagrams do not provide a notation to express them. First is the notion that a transition may occur only after some period of time has elapsed. This is highly relevant to the stages of a cell’s lifecycle, and nearly every cell’s state machine diagram contains such a notion. This has been represented by a guard containing a statement of the form $\lambda(condition)$, used to indicate the passage of sufficient time. Secondly, EAE is a stochastic system, and some transitions occur probabilistically. Such transitions have been modelled as guards containing statements of the form $\delta(condition)$. It is envisaged that both temporal and probabilistic conditions are highly relevant concepts in many biological and complex domains, and modelling such domains requires that they be representable.

Several of the state machine diagrams in this domain model depict single states that have no transitions to alternative states, and that are orthogonal to other aspects of the diagram. An example of this may be found on figure 4.13, where a CNS macrophage is always capable of expressing MHC-II. It is unconventional for states on state machine diagrams to exist in isolation such as this, however it is relevant and informative with respect to the present domain.

The EAE domain model contains several examples of domain aspects modelled as states that are not necessarily the internal state of the entity being modelled. It may be argued that the spatial compartment in which a cell resides is not part of its internal state, yet it has proven informative to indicate on cellular state machine diagrams which compartments the cell may reside in, and the conditions necessary for it to migrate elsewhere. A similar approach has been employed in [Garnett *et al.* 2008]. More extreme examples are comprised of state machine diagrams depicting the influence of cytokines on cells of the system; figures 4.16, 4.17 and 4.18 pertain to type 1 cytokine, type 2 cytokine, and TNF- α . The main focus of these diagrams has been to indicate the conditions necessary for their perception to induce behavioural changes in the various cells of the system. The influence a molecule has on another entity of the domain is not part of its internal state. Yet it is felt to be informative to consolidate this information onto a single diagram.

4.3.9 Summary

There exist aspects of EAE that are both well captured, and not well captured using UML notations. It is the complex, dynamic, and high partially-orthogonal dimensional aspects of the disease that UML does not represent well. Unsatisfactory attempts to capture the entirety of the domain within single diagrams lead to the approach of defining *perspectives* of the domain; decomposing its high-level complexity into more manageable components to be modelled individually.

Perspectives have been modelled using UML activity diagrams. This notation was chosen for its ability to represent any abstract event as an activity, and link them together in order. However, the notation does not capture the large-scale concurrency or stochasticity of the real domain. *In vivo* there exist many populations of cells undergoing different activities depicted on the diagrams at many points in time; there is no sequential transfer of control as suggested on the activity diagram, a single cell may interact with many others at the same time, and may continue to do so after it has instigated an event in another cell.

Sequence diagrams were considered for the modelling of perspectives, but their even stronger implication of transfer of control, and ‘blocking’ of cell activity whilst some other activity is undertaken lead to a preference for activity diagrams. In addition, the cyclic nature in which nearly all high-level system behaviours manifest from the activities of low-level entities could not be depicted on sequence diagrams, owing to their monotonic depiction of time down the vertical axis.

The concurrent and cyclic nature of several perspectives, when depicted on activity diagrams, has lead to fork relationships that do not have corresponding joins. Under strict ‘token passing’ interpretation of activity diagrams many of these diagrams do not terminate, and some do not have end states. Whilst unorthodox, this is not an inaccurate reflection of the real domain.

There were three relational concepts that could not be depicted using standard activity diagram notation, but were considered sufficiently fundamental to the domain that a bespoke notation was created. The first is the notion that entity A may conduct an activity that simultaneously leads to others being instigated in other entities,

without A's activity terminating. Next is the notion that entity A undertakes an activity that partially disrupts a second activity from inducing a third. Lastly, the notion that one activity actively influences a decision resulting from a second activity without inducing it, or itself terminating.

Class diagrams are considered to be of limited benefit in depicting EAE, owing to its highly dynamic and stochastic nature. Class diagrams can become highly connected, and require numerous constraints to accurately reflect the real domain, all of which add to the diagram's complexity and hinder its comprehension. Furthermore, an entity's ability to engage in some relationship frequently depends on its history of interactions with other entities. In such dynamic systems there are various temporal aspects to a relationship between two entities, and these cannot all be depicted, leading to ambiguities when interpreting diagrams.

UML state machine diagrams have proven effective in depicting single-entity dynamics. Once more, it is high and partially-orthogonal dimensionality that raises issues. However, these have been satisfactorily overcome through use of guards. It is noted that many transitions depend on probabilistic or temporal conditions, and a notation was devised to represent these aspects.

The *expected behaviours* diagram reflects the most abstract level of modelling in the EAE domain model. It depicts the phenomena observed in the real domain, and describes how, at a very abstract level, domain model entities interact to produce system-wide behaviours believed to be sufficient for the manifestation of real-world phenomena. The diagram does not conform to any UML notation, since none could be found that provides for the specification of such concepts.

4.4 Conclusion

As stated in section 4.1, the goal of this chapter has been twofold: to present the latest iteration of the EAE domain model, and to provide an assessment of UML's ability to express this complex immunological domain. It has contributed to the thesis aim by directly addressing research objective 1: to explore the role of domain modelling in the EAE case study. Note that further contribution to this objective is made in chapter 8, which reflects upon the confidence that domain modelling brings to simulation-based work.

Section 4.2 presents the EAE domain model. Diagrammatic modelling is performed at three layers of abstraction, and is largely expressed using UML notations. The model is presented in a top down manner, commencing with a highly abstract depiction of how domain model cells interact to produce system-wide behaviours believed to significantly contribute to the phenomenon of paralysis and subsequent recovery as observed *in vivo*. EAE's complexity has been tackled by presenting four *perspectives* of the disease, which correspond with four stages of onset and recovery. Activity and class diagrams describe how domain model cells and molecules are involved in manifesting each of these perspectives. Lastly, the lowest level of modelling entails the specification of cell and molecule dynamics through the use of state machine diagrams.

Construction of the EAE domain model has afforded insight into the ability of UML to represent this complex immunological domain, as reported in section 4.3. UML is effective in describing many aspects of the disease, however the qualities of the domain that are difficult to express are large scale concurrency, stochasticity, and partially-orthogonal high dimensionality.

The next chapter describes the creation of ARTIMMUS, a simulation of EAE that is based on the domain model expressed here.

Chapter 5

Developing and Calibrating an EAE Simulation Platform

This chapter describes the derivation of the platform model from the domain model of the previous chapter. The platform model serves as a specification for a simulation of EAE, named *ARTIMMUS*: “artificial murine multiple sclerosis.” The platform model is presented in section 5.2. Next, section 5.3 reports on the novel procedure used to calibrate ARTIMMUS, and guide its development to an appropriate representation of real-world EAE. Section 5.4 briefly reflects upon the role of visualisation in the development of ARTIMMUS, providing insight into simulation’s mechanics at each iteration and supporting assessment of whether they are an adequate representation of the target domain. Lastly, section 5.5 concludes the chapter. First, however, section 5.1 provides a motivation for the work carried out here.

5.1 Goal and motivation

This chapter reports the specification and calibration of ARTIMMUS, a simulation platform on which *in silico* investigation into EAE can be performed. The creation of the simulation platform provides a case study for investigating the calibration of immunological simulations, specifically how the calibration process might contribute to establishing confidence in a simulation’s faithful representation of the target domain. Hence, the chapter specifically addresses research objectives 2 and 3: to create an agent-based simulation of EAE, and to investigate and develop techniques for calibrating agent-based simulations.

5.2 Platform model

The present section constitutes the platform model of ARTIMMUS. Much of the model is implicit. Rather than reproduce large quantities of the domain model, section 5.2.1 instead indicates the manner in which the domain model is interpreted into the platform model, detailing which aspects are considered to constitute part of the platform model, and which are not.

The remaining sections describe how specific aspects of the domain model are implemented. These sections make explicit the assumptions that are adopted where the domain model is not able to fully specify a particular behaviour. Section 5.2.2 details the implementation language of ARTIMMUS, and shows how the computer code is organised. Section 5.2.3 describes the abstraction of physical space. The simulation’s representation of immunisation for EAE is described in section 5.2.4. The turnover of cells is described in section 5.2.5. Section 5.2.6 denotes how APCs in ARTIMMUS undergo periodic transitions from immature to mature states. The representation of blood flow in the simulation, which dictates the migratory behaviour of many motile cells is described in section 5.2.7. Section 5.2.8 indicates how normal distributions are used to represent the period of time that elapses between cellular state transitions. The implementation of local activation of T cells is described in section 5.2.9. Section 5.2.10 reports on the conditions that dictate the adoption of polarisations by DCs and CD4Th cells. The manner in which the simulation is initialised, such that simulation artifacts arising from cells starting in identical states and locations are minimised, is detailed in section 5.2.11. The mechanism representing the decay of cytokine molecules is described in section 5.2.12. Lastly, the implementation of T cell specificity is covered by section 5.2.13. Many of these sections reference ARTIMMUS parameters, of which a full listing is provided in appendix section B.1. Note that parameters are named according to their function; for example, the names of all parameters relating to T cells begin with the text “*TCell_*.” Parameter names are indicated in italics.

5.2.1 From domain model to platform model

The EAE domain model of the preceding chapter explores EAE instigation and recovery at several levels of abstraction: a highly abstract system-level overview of the cells believed to be responsible for disease onset and recovery; a decomposition of disease onset and recovery into four system-level behaviours, indicating how these manifest from cell-level interactions; and the full dynamical behaviours of individual cells in the system.

The individual cell-level behaviours are expressed as state machine diagrams, and it is these that form the specifications for the simulation platform: cellular behaviours are coded on the basis of state machine diagrams. The other diagrams, which cover system-level behaviours, describe the expected result of the large scale simulation of individual cells. These diagrams are in no way used as specification for code in the simulation platform; these behaviours should emerge as a result of cellular interactions represented in the state machine diagrams.

5.2.2 Simulation architecture

The ARTIMMUS simulation platform is developed in Java, and within the MASON¹ simulation framework [Balan *et al.* 2003, Luke *et al.* 2004]. Java belongs to the object oriented programming paradigm, wherein logical entities in the program are explicitly represented and are responsible for maintaining their own state. As such, the object oriented paradigm maps well onto the concept of agents in agent-based modelling. MASON provides facilities for spatial representation, visualisation, and inspection of state within the simulation. It provides a simulation engine that manages the execution of simulation agents. Time is discretized into *time steps*; every time step the state of every cell and compartment in the simulation is updated.

¹Multi-Agent Simulator of Neighborhoods (MASON),
<http://www.cs.gmu.edu/~eclab/projects/mason/>.




Figure 5.1: Class diagram depicting organisation of classes in the simulation platform in terms of inheritance hierarchies and interface implementations. The majority of associations are not shown, in aid of readability.

Figure 5.1 depicts the organisation of classes in the simulation platform, focussing on inheritance and implementation of interfaces; the majority of associations are not shown, as these may be derived from the domain model presented in the previous chapter. The simulation platform makes extensive use of abstract classes, inheritance, and interface implementation. The simulation ultimately consists of cells all of which inherit from `Cell_Impl`, and spatial compartments in which they exist, which inherit from `Compartiment2D_Impl`. Both these abstract classes implement `Steppable`, the MASON interface required for execution within the simulation engine. The class `EAESimulation` is the driver of the simulation, and is responsible for its initialisation.

There are many behavioural aspects common to multiple cell types. Such behaviours, such as maintaining the state of a T cell, are coded as far up the hierarchy tree as possible. This coding practice reduces redundant code, and leaves a single point of code maintenance should it ever need amendment. Where cells are concerned, all classes encoding cellular behaviour are abstract, with the exception of those that may actually be instantiated to directly represent a cell of the domain model. A similar programming pattern is used for coding the simulation's spatial compartments. These share many commonalities, which are implemented in `Compartiment2D_Impl` and inher-

ited in the concrete subclasses. In either case, abstract classes are often used to provide default behaviours which are overridden by specific exceptions where necessary.

The simulation platform makes use of interfaces with respect to the MHC:peptide complexes that cells may express. This serves to reduce the complexity of code. For example, a CD8Treg is specific for Qa-1:CDR1/2, which may be expressed by either a DC or a CD4Th cell. From the perspective of the `CD8Treg` class, it need only establish whether an instance of a cell implements the `Qa_1_CDR12` interface to know whether they may interact through this complex, rather than querying whether the cell is an instance of the `CD4THelper` or `Dendritic Cell` types.

The `Molecule` abstract class is the superclass for all the cytokines (`Type1`, `Type2` and `TNF- α`) and peptides (`Fr3`, `CDR1/2` and `MBP`) of the system. Cytokines are not represented as individual agents in the simulation. Rather, their concentrations in space are maintained by `Compartment2D_Impl` and represented as real numbers; this is discussed in further detail in section 5.2.3. The concrete cytokine classes are programmed using the singleton pattern, and as such references to these singletons are used when cells either secrete cytokines into the compartment or request their concentration. These singleton instances are also referred to by cells when being phagocytosed by APCs; the APCs need not know what form of cell is being phagocytosed to establish which peptides they contain, these references to these singleton peptides are passed directly to the APC by the cell being phagocytosed at the time.

5.2.3 Spatial representation

Physical space is explicitly represented within ARTIMMUS through use of a two dimensional lattice grid, graphically depicted in figure 5.2. The lattice grid consists of four layers, one for each type of entity explicitly represented in the simulation. There is one layer to maintain the location of cells, and one for each of the cytokine types: type 1 cytokine, type 2 cytokine, and $\text{TNF-}\alpha$.

Cells occupy grid spaces in the lattice, and may move between them. There are two sizes of cell considered in the simulation: T cells, and all other cells. T cells are assumed to be $\frac{1}{7}$ the size of all other cells. A single grid space is equivalent to the area of a single larger cell. As such, a soft upper limit of 7 T cells may occupy a single grid space, however for all other cells only a single cell may occupy a given grid space.

Cytokines diffuse around the grid, in accordance to algorithm 1, which is adapted from [Andrews & Timmis 2006]. Diffusion is performed at each time step in the simulation, and is implemented as an atomic operation. At each time step the majority of cytokine contents of a particular grid space are randomly diffused to one of its eight neighbours. Cytokine concentrations are represented as real numbers, however diffusion operates only on whole numbers, moving quantities of multiples of 1.0 to neighbours at a time. Concentrations less than 1.0 remain in the same grid space.

With respect to cells, lattice grids are toroidal only in the horizontal dimension. Cells migrating into a compartment are placed at the top of the lattice grid, those moving off the bottom of the grid are considered to have migrated elsewhere. With respect to cytokines, grid spaces are both vertically and horizontally toroidal.

5.2.4 Immunisation mechanism

Immunisation for EAE is accomplished *in vivo* through the administration of MBP, PTx and CFA. These immunisation substances do not find explicit representation within the simulation, which instead represents immunisation through the appearance of MBP-presenting immunogenic type 1 polarising DCs in the SLO compartment. Hence, the




Figure 5.2: The lattice grid based spatial representation of ARTIMMUS. Each grid space in the lattice may contain cells and molecules, such as cytokines, with bespoke layers representing the spatial occupancy of each within the lattice. The ability of cells to move and for molecules to diffuse between grid spaces is indicated.




Figure 5.3: ARTIMMUS's immunization mechanism, and how it is parameterised. The label “*Simulation_immunizationLinear*” has been omitted from parameter names.

periphery compartment of the domain model is not represented in the platform model, and is not implemented in ARTIMMUS.

The immunisation mechanism, depicted in figure 5.3, is parameterised through 4 parameters: *Simulation_immunizationDC0*, *Simulation_immunizationLinearFreq*, *Simulation_immunizationLinearGradient*, and *Simulation_immunizationLinearInitial*. The last specifies the number of immunisation DCs placed into the SLO compartment at time zero, as a one-off event. The remainder parameterise a linearly reducing number of DCs that are added to the SLO periodically. The period is defined by *Simulation_immunizationLinearFreq*. *Simulation_immunizationDC0* and *Simulation_immunizationLinearGradient* describe the level of DCs inserted at time zero, and the rate of linear decay. Every *Simulation_immunizationLinearFreq* hours, the value described by these two parameters, given the current simulation time, is rounded to the nearest whole number of DCs which are then placed in the SLO.

Algorithm 1: Cytokine diffusion algorithm. Designed as an atomic operation, all of the cytokine molecules present in a grid space at the start of the operation are distributed amongst its neighbours. Algorithm has been adapted from [Andrews & Timmis 2006].

```
input: grid, lattice grid containing cytokine concentrations
output: grid, lattice grid containing cytokine concentrations following diffusion
newgrid ← grid.clone
for x ← 1 to grid.width do
    for y ← 1 to grid.height do
        quantity ← newgrid[x][y]
        if quantity ≠ 0 then
            share ← ⌊quantity ÷ 8⌋
            grid[x][y] ← quantity // reduce by ‘quantity’
            grid[x - 1][y - 1] ←+ share // increase by ‘share’
            grid[x][y - 1] ←+ share
            grid[x + 1][y - 1] ←+ share
            grid[x - 1][y] ←+ share
            grid[x + 1][y] ←+ share
            grid[x - 1][y + 1] ←+ share
            grid[x][y + 1] ←+ share
            grid[x + 1][y + 1] ←+ share
            remainder ← quantity - (8 × share)
            repeat remainder times
                // select random gridspace from the neighbourhood
                xi ← (random(3) - 1) + x
                yi ← (random(3) - 1) + y
                grid[xi][yi] ←+ 1
```

5.2.5 Cellular turnover

Neurons enter apoptosis upon exposure to a sufficient concentration of TNF- α in a single time step. Upon phagocytosis by an APC, a new neuron is placed at the exact same location as that which preceded it.

The basal size of T cell populations are homeostatically maintained, through a mechanism that acts independently of any proliferative activities, which reflects the activities of the thymus in creating new T cells. The probability that a naive T cell is inserted into the simulation at each time step is calculated as the initial population size of the particular T cell type, divided by the mean lifespan of a naive T cell, multiplied by the proportion of an hour that a single time step represents. T cells created in this manner are inserted into a random location in the circulatory system.

Like neurons, DCs and CNS macrophages are replaced by new cells upon entering apoptosis. The immature cells that replace such apoptotic cells are placed at random locations within the compartments where they originated from.

5.2.6 Periodic migration of DCs

There exist two forms of DC in the simulation: those capable of migratory behaviour, and those that are not. Non-migratory DCs are those placed in the spleen, SLO and CLN compartments upon simulation initialisation. Migratory DCs exist in the CNS as immature cells, and move into the CLN compartment upon maturation.

This migratory behaviour is periodic, DCs remain in the CNS compartment as immotile immature cells for some predetermined period of time, randomly determined for each individual DC, before commencing migration to the CLN compartment. Once in the CLN migratory DCs continue to follow the blood flow for some randomly determined period of time before coming to a rest somewhere in the compartment.

5.2.7 Migratory behaviour

Compartments in ARTIMMUS are represented as lattice grids. In each time step a motile² cell may move to any of the eight grid spaces surrounding the one in which it current resides, or it may remain stationary. Cell movements are calculated by considering horizontal and vertical movements independently. In the horizontal plane, a cell has equal probabilities of moving left, right, or remaining stationary.


Cells migrating into a compartment are placed at the top of the lattice grid. Conversely, traversing beyond the bottom of the lattice grid is interpreted as the cell leaving the compartment. The probabilities of vertical movement, being a move down, up, or remaining vertically stationary, are dictated by the time required for a T cell to migrate through the compartment. The times required for migration through a compartment are compartment-specific, and are represented by parameters of the form *CLN.timeToCrossOrgan*, many of which hold literature defined values. Figure 5.4 indicates how the probabilities of a cell performing a particular vertical movement are derived, based on a compartment's height, and how many time steps a cell will generally remain in a compartment for.

5.2.8 Probabilistic timing of events

Many of the cellular state transitions described in the domain model are implemented to occur after some probabilistically determined period of time. For example, on average DCs remain immature for a particular period of time, however individual DCs differ in the exact durations they spend in this state. Such state transitions are implemented through use of normal distributions, as depicted in figure 5.5. When a cell enters a particular state, event A, a probability distribution is used to select an absolute time in the future for the transition to the following state, event B. Hence, such distributions describe the duration of time cells spend in particular states. They are parameterised by a mean and standard deviation³. Note that, since a normal distribution is used to describe the durations of time spend in particular states, it is possible that event B will be selected to have occurred before event A. This is guarded against by converting all events occurring in the past into events occurring 1 hour after event A.

²Not all cells are motile. T cells are motile only when not bound to an APC, neurons are never motile, and APCs are motile only in certain states.

³The parameters of ARTIMMUS that describe standard deviations actually describe twice the standard deviation. Hence, their values indicate that ~95% of transitions must occur within a particular period of time surrounding the mean.



height of compartment in grid spaces = ϕ

time to cross compartment = ω

time slice = τ

time steps in which cell must cross compartment = $\delta = \frac{\omega}{\tau}$

proportion of total time steps that result in a net downwards movement = $\alpha = \frac{\phi}{\delta}$

There are three constraints on the probabilities of moving down (x), up (z), or remaining stationary (y): all possible movements must sum to 1.0; any upwards movement must be reflected in the probability of downwards movement, in order to traverse the compartment in α time steps; and the probability of remaining stationary is the mean of probabilities of moving up and down. This last constraint is an assumption made to make the equations solvable. The constraints are mathematically expressed as follows:

$$x + y + z = 1.0 \quad (5.1)$$

$$x = z + \alpha \quad (5.2)$$

$$y = \left(\frac{x + z}{2} \right) \quad (5.3)$$

Derivation of how to calculate z , y , and x :

$$z = x - \alpha \quad (5.4)$$

$$y = \left(\frac{x + z}{2} \right) = \left(\frac{x + x - \alpha}{2} \right) = \left(\frac{2x - \alpha}{2} \right) \quad (5.5)$$

$$\begin{aligned} x + y + z &= 1.0 \\ x + \left(\frac{2x - \alpha}{2} \right) + (x - \alpha) &= 1 \\ 2x + (2x - \alpha) + 2x - 2\alpha &= 2 \\ 6x - 3\alpha &= 2 \\ x &= \left(\frac{2 - 3\alpha}{6} \right) \\ x &= \left(\frac{1 - 1.5\alpha}{3} \right) \end{aligned} \quad (5.6)$$

Figure 5.4: The calculation of probabilities that a migrating cell, such as a T cell, will move either downwards, upwards, or remain at its current vertical level in a particular time step. Probabilities are represented within the range 0.0 to 1.0. The probabilities of moving down, up, or remaining vertically stationary are represented by x , z and y respectively. The calculations are derived through consideration of a compartment's height, and how long a cell has to traverse this distance.




Figure 5.5: Depiction of how the absolute time of events is determined through a normal distribution of possibilities. When event A occurs, the absolute time at which event B is to take place is derived from a normal distribution of relative times in the future. Simulation parameters specifying the mean and $2 \times$ the standard deviation describe the distribution. Event B cannot take place in the past, hence absolute time for B that are less than the time of event A are converted into future times.

5.2.9 Local activation of T cells

The domain model stipulates that effector T cells cannot secrete cytokines until they have been locally activated. The intended meaning of local activation is that the T cell has left the location at which it was primed and is re-stimulated at the site of the immune response. In order to prevent T cells in ARTIMMUS from being locally activated by the APCs on which they prime, a time limitation has been placed on their ability to be locally activated. T cells can only derive local activation some time after differentiation into effector cells, parameterised by *TCell_timeLocalActivationDelay*.

5.2.10 Polarisation selection by DCs and CD4Th cells

Upon maturation DCs become either type 1 or type 2 polarising. Type 1 polarising DCs secrete type 1 cytokine that induces preferential selection of CD4Th1 polarisations by CD4Th cells. The polarisation selected by a DC is dependent on the local cytokine milieu at the time of maturation. If the total proportion of type 2 cytokine in the cytokine milieu is less than the figure represented by the parameter *DendriticCell-cytokineType2PolarizationRatio* then a type 1 polarisation is adopted. Otherwise a type 2 polarisation is adopted, including the case that there is a complete absence of any type 1 or type 2 cytokine in the milieu.

CD4Th cells adopt either a CD4Th1 or CD4Th2 polarisation upon entering their proliferative cycles. One more, the decision of which polarisation to adopt is based on the local cytokine milieu in the grid space where the cell resides. If the proportion of type 1 cytokine in the milieu is greater than or equal to 80%⁴, then a type 1 polarisation is preferentially adopted. Otherwise, including the case that there exist no cytokines in the milieu, a type 2 polarisation is preferred. The adoption of a polarisation is probabilistic, with the degree of ‘preference’ in each case parameterised through *CD4THelper_diff00* and *CD4THelper_diff08*.

5.2.11 Simulation initialisation

Simulation initialisation, at time zero, is done in a manner that minimises simulation artifacts arising from cells starting in identical states and locations. DCs and CNS macrophages are created to be part way through their lifespans as immature cells; this prevents large quantities of these cells entering maturity at unrealistically similar times.

⁴Represented at 0.8, percentages are scaled onto the range between 0.0 and 1.0.

The initial populations of naive T cells are placed in randomly determined compartments, excluding the CNS. Some proportion of migratory DCs are placed in the CNS as immature cells, whereas the remainder are placed in the CLN compartment as mature cells. The exact proportion is dictated by the proportion of a DCs entire lifespan that is spent in an immature state.

5.2.12 Cytokine decay

In addition to diffusing around compartmental space, cytokines also decay over time. This has been implemented in discrete time as follows:

$$C_{t+\tau} = C_t \left(\frac{1}{2} \right)^{\frac{\tau}{\lambda}} \quad (5.7)$$

Where τ is the duration of time represented by a time step, in hours (e.g., 0.125), λ is the decay rate, also in hours, and C_t is the concentration of cytokine in a particular grid space at time t . Cytokine concentrations, represented as real numbers, subject to this implementation of decay will never reach a concentration of exactly 0.0. As such, at each time step, the concentration of cytokine in a particular grid space is set to 0.0 if it is less than the value indicated by the parameter `Molecule_decayThreshold`. This thresholding adds realism to the simulation; cytokines secreted by cells *in vivo* do not have infinitely far reaching influence.

5.2.13 T cell specificity

Upon contact of a T cell with a cell expressing MHC:peptide for which it is specific, a binding between the two cells is probabilistically determined by parameters. This reflects the range of specificities that *in vivo* T cells have for particular MHC:peptide complexes. A successful binding will lead to further interactions between the cells. In the case of T cells induced into their proliferative cycles, the T cells become immotile, bound to the APC for the duration of their proliferative cycle. In contrast, an unsuccessful binding does not lead to cellular interaction.

The determination of whether a binding is successful is conducted when the two cells first contact one another, as indicated by their occupancy of either neighbouring or the same grid spaces, and is probabilistic in nature. All T cells that are not created through proliferation are assigned specificities upon creation, represented as a real number randomly selected to lie between the values indicated by the simulation parameters `TCell_specificityLowerLimit` and `TCell_specificityUpperLimit`. These two parameters obey the constraint:

$$0.0 \leq \text{TCell_specificityLowerLimit} \leq \text{TCell_specificityUpperLimit} \leq 1.0$$

Daughter cells arising from proliferation inherit the specificities of their parent cells. This specificity represents the probability that contact with a neighbouring cell for which a T cell is specific will result in a successful binding.

5.3 Calibration of simulation platform

The present section describes the calibration procedure used to align the simulation with *in vivo* progression and recovery of EAE. The parameter values that result from

calibration, and the manner in which each parameter was subject to the following procedure, are denoted in appendix section B.1.

Calibration is an essential stage in the construction of a simulation platform for the following reasons, with each discussed in turn in the text that follows:

1. The domain model, and by necessity the simulation platform, are *abstractions* of the cellular and molecular interactions that govern real-world EAE. It is not known in advance whether the abstractions and assumptions that constitute a particular iteration of the domain model will reproduce real-world behaviours when simulated.
2. There are a great many simulation parameters for which the corresponding *in vivo* values are either not known or are not known exactly.
3. Since the domain model and simulation platforms are abstract representations of real-world EAE, *in vivo* parameter values, if known, will not necessarily translate directly into the simulation platform.

The domain model presented in chapter 4, and the platform model outlined above, are the end result of an iterative development process; the first iteration and implementation of the domain model as a simulation platform did not realise the system-level dynamics observed *in vivo*. Through undertaking the present calibration process, the abstract representation of cellular and molecular interactions and dynamics are brought to a point that can satisfactorily reproduce disease onset and recovery as observed *in vivo*.

With regards the second point, the parameter values of the simulation platform are arrived at through three means. Some values are based on domain expertise or literature. Where no domain-specific guidance exists for the value a parameter should adopt, arbitrary values are assumed. Finally, many but not all parameters are subject to calibration, their values altered to best align simulation behaviour with that observed *in vivo*.

Point 3 above states that real-world parameters, where known, will not necessarily translate directly into simulation parameters. The real-world elements of the domain that find explicit representation within the simulation platform must compensate for the effects of those elements that do not. It is hypothesised that those elements not represented are not critical to disease progression and recovery, however this is not known absolutely; they may have some effect. Consider the type 1 and type 2 cytokines of the domain model. These are abstractions of a myriad of real-world cytokines, abstracted on the basis of their promotion of either a type 1 or type 2 T cell response. Were the secretion and decay rates of all these real-world cytokines to be known with great certainty, none would translate directly into the simulation platform. Calibration aligns simulation behaviour with that observed *in vivo*, and in doing so the effects of the missing elements are to some degree compensated for.

The result of calibration is the amendment of the simulation's mechanics and/or the tuning of its parameter values. The early stages of calibration focus primarily on altering the abstractions and assumptions underpinning the cellular and molecular dynamics and interactions; where these simulation mechanics are incorrect, the deviation of simulation behaviour from the real-world EAE are very large. Once the simulation's mechanics are relatively stable and its dynamics more closely match the target system, calibration focusses on simulation parameters to make relatively minor changes to the simulation's system-level behaviours.


The end result of calibration is a simulation platform that encompasses mechanics and parameter values that satisfactorily reproduce the dynamics of real-world EAE. This state of the simulation is referred to as the *baseline* in subsequent chapters, and forms a basis against which experimentation may be contrasted and interpreted.

The following sections describe the *collaborative calibration procedure*, a novel calibration procedure that makes extensive use of domain expertise and calibration against multiple real-world experimental scenarios. Its application in developing ARTIMMUS to a satisfactory level of abstraction is also described. Section 5.3.1 motivates calibration against a domain expert's expertise. The calibration procedure itself is described in section 5.3.2. Section 5.3.3 expands on the criteria for calibration, motivating the use of multiple real-world experimental scenarios as opposed to just one. The calibrated simulation's dynamics are presented in section 5.3.4, and section 5.3.5 explores areas of this baseline behaviour that are recognised as being inconsistent with that of the domain.

5.3.1 The role of domain expertise in calibration


The *collaborative calibration procedure* consults the domain expert with respect to how well the simulation is aligned with *in vivo* behaviour. This is to mitigate the significant uncertainty and variance in data from the real domain, which does not exist in a format that the simulation may be directly calibrated against. Figure 5.6 highlights the type of data typically derived from EAE experiments *in vivo*. Figure 5.6a shows the mean severity of EAE, measured on a 5 point scale [Kumar *et al.* 1996], experienced by groups of 6-8 mice of different strains following the same induction of EAE administered in each group. There is considerable variation in EAE progression experienced by each strain of mouse. Figure 5.6b shows the severity of EAE experienced by groups of 5 mice of the same strain, with each group undergoing a different intervention. It can be seen that within groups of exactly the same experimental animal undergoing exactly the same intervention, there is considerable variation in the severities of EAE experienced. Figure 5.6c indicates the number of CD4Th1 cells found in different bodily compartments at various times following immunization in several experimental animals. It can be seen that there is once again significant spread in the data, with several examples where there are no data points lying on the calculated mean values. Acquiring data such as this can require the sacrifice of an experimental animal, and since genetically identical animals undergoing exactly the same EAE induction can experience significantly different responses (figure 5.6b), it is difficult to compile a representative progression of EAE in terms of individual cell population number.

Data of this nature can be challenging to calibrate a simulation against. The simulation platform can be executed many hundreds or thousands of times in order to obtain highly representative averaged values for a particular metric of interest. The same fidelity of data is not available in the immunological literature, where the number of samples obtained in acquiring an average rarely exceeds ten. Whereas a computational simulation can readily provide exact numbers of a particular cell type over time, the data of figure 5.6c only *indicates* their number, there is not an easily established exact mapping between the metric used and the actual number of cells observed by the instrument. For this reason the calibration process consults with the domain expert, who is uniquely placed to interpret how data arising from the simulation may map to that which he has encountered of the real-world domain. In this manner the simulation platform is calibrated against the domain expert's understanding of the domain, some of which is captured in the domain model, rather than directly against any particular immunological data.



(a) The mean EAE severity experienced by groups of 6-8 mice of different strains following the same induction of EAE administered in each group. Adapted from [Madakamutil *et al.* 2008].

(b) EAE severity experienced by each of 5 mice under different interventions (boxes A-C), over time (in days). Taken from [Kumar *et al.* 1996]. The interventions involve injecting mice experiencing EAE with varying quantities of Treg cells.



(c) An indication of the number of CD4Th1 cells (EI = expansion index on the y-axis) residing in a selection of organs at various times following the induction of EAE from [Menezes *et al.* 2007]. Organs examined were: lymph nodes (LN), spleen (Sp), blood (Bl), lungs (Lu), bone marrow (BM), liver (Li), mesenteric lymph nodes (mLN), thymus (Thy).

Figure 5.6: Examples of clinical data pertaining to the progression of EAE in mice. EAE severity is scored on a scale of 0 to 5, “1, flaccid tail; 2, hind limb weakness; 3, hind limb paralysis; 4, whole body paralysis; 5, death.” [Kumar *et al.* 1996].

5.3.2 Calibration process

The process through which the simulation platform is calibrated is as follows.

1. The simulation is executed with a “best guess” set of parameters and representative median results collected, thereby avoiding assumptions concerning normality of distributions.
2. The domain expert identifies aspects of simulation dynamics that deviate from his perspective of the real system. Both domain expert and modeller discuss to identify the source of the deviation in the simulation. Both parties’ input are invaluable: the domain expert brings a wealth of domain-specific understanding, whilst the modeller, having built the simulation, has a detailed understanding and intuition of how this information has been abstracted and how the simulation operates.
3. Potential avenues of model and simulation amendment and development are identified, and each one is independently integrated into the simulation in turn. In each case the simulation is then executed to once more obtain representative median results.
4. Subsequent interactions between domain expert and modeller re-examine the results and decide upon which amendments are to be permanently adopted into the simulation. As such, the calibration process is iterative.

Throughout this calibration process, the simulation’s behaviour is examined under two disparate experimental circumstances: that of normal EAE progression and recovery following induction of EAE; and the progression of EAE with the regulatory network disabled. The motivation for calibration against multiple experiments is explored in the following section. This calibration procedure performs *qualitative* alignment of simulation against *in vivo* dynamics; quantitative comparison is not possible for reasons highlighted above: simulation is highly abstractive, encompasses far fewer cells than exist in an experimental animal, and the existing *in vivo* data on EAE is unsuitable for direct comparison with simulation data.

As noted above, the calibration process is undertaken as a collaborative effort between the modeller (Mark Read) and the domain expert (Dr. Vipin Kumar) in person at TPIMS⁵. The calibration process calls for a great deal of collaboration between the two parties, who must work closely to identify potential sources of misalignment between the simulation platform and the target system. Given very different backgrounds of the two parties, it is felt that this activity is best performed in person, face to face.

5.3.3 Calibration against multiple data points

One use of simulation, as reported in the following chapters, is in performing predictive experimentation: investigating experimental scenarios that are not possible in the real-world. This is a situation where confidence in a simulation’s faithful representation of the domain is important; the real-world system’s behaviour under these circumstances is not known. An important feature of the present calibration procedure is its calibration against multiple experimental scenarios, such as knock-out experiments. Calibration against multiple experimental conditions instead of only one helps avoid

⁵Torrey Pines Institute for Molecular Studies (TPIMS), San Diego, USA. The institute at which Dr. Kumar and his lab are based.




Figure 5.7: Abstract depiction of how calibration against multiple experimental scenarios can lead to identification of simulation abstractions that better approximate the dynamics of the real-world domain. A simulation with inappropriate mechanics (e.g., s_1) may be fitted to a single datapoint (x_2), but will fail to faithfully replicate real-world dynamics (r) under other experimental scenarios. Simulation calibration against multiple data points (x_1 and x_2) could reveal a simulation to be an inappropriate abstraction of the real-world domain, and lead to further development (s_2). This increases confidence in predictive *in silico* results under experimental scenarios that cannot be conducted in the wet-lab (shaded area, and x_3).

the fitting an incorrect model to a single data point. This is an important feature of the calibration process, and is illustrated figure 5.7.

The space of different experimental scenarios is represented along the x axis, and the system's behaviour under those scenarios is depicted on the y axis. The function $r(x)$ represents the behaviour of the real-world system under scenario x , which originates from the interactions and dynamics of the many cells and molecules that comprise the real-world system. Scenarios x_1 and x_2 are established experiments, where the real-world system's behaviour is known. The grey box around scenario x_3 , on the other hand, represents a series of experiments that are problematic to perform in the real-world, and which is to be investigated *in silico*.

The function s_1 represents a highly abstract and naive simulation representation of the real system. It has been calibrated against the scenario x_2 only. Function s_2 represents another simulation, whose alignment with real-world system behaviour has been demonstrated under both scenarios x_1 and x_2 . The figure demonstrates that an inappropriately naive model can be fitted to a single data point, however its behaviour at other data points can vary considerably from that of the real-world system. It is conceivable that the parameters of an incorrect model can be manipulated such that its behaviour corresponds to any *single* datapoint. Simulation s_2 may have been derived when attempting to further match s_1 to the scenario x_1 , and realising that s_1 was too naive to properly replicate the rich behaviours of the real-world system; some element of r was inappropriately represented or abstracted, and hence s_2 is created. Simulation s_2 is still not an exact analogue of the real-system, but it is demonstrably better than s_1 .

Following this logic, the present EAE simulation platform is calibrated against the domain expert's expertise under two experimental scenarios: that of physiological recovery from autoimmunity following the induction of EAE, and the progression of EAE in mice where the ability of the regulatory network to regulate CD4Th1 cells has been revoked⁶. The demonstrable ability of the simulation to replicate multiple real-world

⁶This is achievable in the real domain through a number of different interventions

experiments affords confidence in the results of subsequent predictive experimentation.

Criteria against which the simulation was evaluated for the first scenario are the population level events described in the domain model, section 4.2.7 and table 4.1a. In addition, the domain expert indicated that the effector CD4Treg population peak should not outnumber the CD4Th1 population peak.

The criteria for the second scenario are similar to the first, with the exception that the CD4Th1 population will not generally have been abrogated by 50 days; this is however stochastic, the CD4Th1 population may reduce to zero in absence of regulation in some cases, but not in the general case.

5.3.4 Calibrated simulation platform dynamics

Upon the completion of the calibration process, the simulation platform is considered to satisfactorily reproduce the *in vivo* dynamics of the target system under two experimental set ups: the onset and physiological recovery of EAE, and the prolonged autoimmune symptoms experienced when the regulatory network has been abrogated. The simulation platform's dynamics under these two experimental circumstances are shown in figure 5.8, which depicts the number of effector T cells throughout the system over time. Note that the 'physiological regulation' scenario is considered to be the *baseline* behaviour of the simulation.




Figure 5.8: Effector T cell population dynamics of the calibrated ARTIMMUS simulation.

5.3.5 Outstanding discrepancies between simulation platform and *in vivo* behaviour

Although the domain expert has deemed the simulation's behaviour to be sufficiently representative of the real-world to commence experimentation, there remains one aspect of simulation behaviour that deviates from the *in vivo* system. It is the time at which effector Treg cells are first generated in the system. The domain expert believes that these cells appear a few days earlier *in vivo* than occurs in the simulation platform. The present section explores three potential causes of this discrepancy.

Firstly, T cells in the simulation traverse and explore spatial compartments in a random fashion, albeit in a manner that results in their migration out of the compartments.

[Kumar *et al.* 1996, Beeston *et al.* 2010]. It is achieved in the simulation through setting the parameter *CD8Treg_cd8TregToCD4ThelperSpecificityDropOff* to 0.0, hence revoking the ability of a CD8Treg to induce apoptosis in a Qa-1 expressing CD4Th1 cell.

ment within the required time. *In vivo*, mature DCs form a highly organised network of cells that specialise in presenting antigen to T cells in a highly efficient manner, greatly enhancing the likelihood that T cells will interact with DCs carrying antigen for which they are specific [Lindquist *et al.* 2004].

Secondly, DC maturation and migration in the simulation is entirely periodic. An immature DC residing in the CNS will mature only when a pre-determined period of time has elapsed; during this time it is unsusceptible to any cytokine or environmental influence. However, *in vivo*, inflammatory conditions can promote the early maturation and migration of DCs [Kindt *et al.* 2007, Janeway *et al.* 2005]. It is possible that this accelerated maturation promotes earlier presentation of MHC:peptide complexes required for Treg priming *in vivo* than in the simulation.

Finally, the simulation platform represents substantially fewer T cells than exist in the real-world system. The presentation of MHC:peptide complexes required for initial Treg priming results from physiological turnover of CD4Th cells. Since the CD4Th population is relatively small during the initial stages of the simulation, naive CD4Th cells resulting from proliferation will readily gain access to APCs and enter their proliferative cycles. Hence, these cells will enter apoptosis only at the end of their full life cycles. Transitions between T cell states are probabilistic; the period of time that a cell remains in a particular state is extracted from a normal distribution. It is possible, though relatively unlikely, that a CD4Th cell will spend a smaller than average duration of time in each of its lifecycle states prior to entering apoptosis. However, in a system containing many orders of magnitude more cells than currently represented in the simulation, a large number of CD4Th cells will transit through their life cycles very quickly, and hence provide the peptides required to prime Treg populations earlier than currently occurs in the simulation platform. Hence, the computational restrictions that prevent the simulation from containing realistic numbers of cells can affect the onset of processes within the system that require threshold quantities of cells to engage in some low probability behaviour.

5.4 Importance of visualisation and state inspection

The domain model presented in the preceding chapter was arrived at in an iterative manner, partially through simulation-based inspection of the system-level behaviours that arise from the cellular dynamics incorporated in each iteration. The iterative development of the simulation platform, and the models that underpin it, has greatly benefited from visualisation of the simulation. The creation of graphical user interfaces can be time consuming activity; a key motivation for adopting the MASON framework is the visualisation capability that it provides. A full screenshot of the completed simulation platform is provided in figure 5.9. The facilities provided by MASON allow the user to double click on a grid space and inspect and manipulate the state of individual cells and cytokine concentrations contained therein. These facilities have allowed close inspection of the system-level behaviours that arise from the cellular-level dynamics implemented in the simulation platform.

Since it is not known what system-level behaviours will arise from a particular set of cell-level dynamics upon execution, hence the need for iterative development, it can be unclear as to whether an unexpected system-level behaviour is a genuine result of a particular set of cell-level dynamics, or a coding error. Visualisation, and particularly inspection of cell states, has aided significantly in establishing which is the case. This concurs with the value attributed to visualisation in “debugging” simulation code, expressed by [Efroni *et al.* 2003]. Furthermore, in concurrence with Cohen’s

third requirement for useful models of biological systems, to “engage the mind of the experimentalist with understandable, visual representations”, visually observing the interactions between cells and the system-level behaviours that result has greatly aided understanding of the system being modelled.




Figure 5.9: Screenshot of the ARTIMMUS simulation. The top right window is being used to inspect the state of a particular cell in a grid space in the CLN compartment.

5.5 Conclusion

This chapter has reported the specification and construction of ARTIMMUS, the *artificial murine multiple sclerosis* simulation platform. Section 5.2 presents the platform model. Derived from the domain model, the platform model serves as a specification for the simulation platform. Explicit representations of system-level of emergent properties are removed from the domain model, since these must emerge from the simulation’s low level dynamics rather than be programmed directly into it. Any implementation-specific amendments and additions are reported in the platform model.

The domain model of the previous chapter and the platform model presented here represent the current iteration of the CoSMoS process; the initial iterations of the CoSMoS process failed to produce a simulation that adequately reproduced the dynamics of the real-world system. Simulation development and evaluation against the real-world domain is driven by the collaborative calibration procedure outlined in section 5.3. ARTIMMUS is calibrated against the model of EAE employed by a single specific laboratory. As such, rather than adopting parameters values from disparate sources, as is the approach of many immunological simulation works reviewed in chapter 3, guidelines on parameter values are obtained directly from the domain expert. The role of the domain expert in the calibration process is to mitigate the considerable variation found in immunological literature. Furthermore, in circumstances where a particular aspect of the immune system is simply unknown in immunology the domain expert

can provide guidance and intuition as to the best course of action. The collaborative calibration procedure is novel in its simultaneous calibration of the simulation against multiple real-world experimental scenarios. As explored in section 5.3.3, it is felt that this feature of the procedure helps avoid fitting an incorrect simulation against only a single data point.

Lastly, section 5.4 reflects upon the role of simulation visualisation in exploring simulation dynamics, and assessing whether they are representative of its mechanics or arise through coding errors.

This chapter has addressed research objectives 2 and 3: to create an agent-based simulation of EAE, and investigate and develop techniques for calibrating agent-based simulations. The following chapters report *in silico* experimentation performed on ARTIMMUS, and explore statistical techniques for interpreting *in silico* results in the context of EAE.

Chapter 6

Explorative Experimentation

The previous chapter reported the specification and calibration of ARTIMMUS, a simulation platform for the investigation of EAE. The present chapter is concerned with the statistical exploration of the dynamics comprising ARTIMMUS, and the interpretation of results of *in silico* exploration.

Section 6.1 presents the motivation and goals of this chapter. Section 6.2 establishes several simulation *responses*, metrics of simulation behaviour required for the analyses explored in this chapter. The section goes on to develop a system for assigning similar autoimmunity severity scores to ARTIMMUS as are used in wet-lab experimentation. In terms of the CoSMoS process, this system constitutes a mapping between the results model and the domain model. A technique for establishing the relationship between the accuracy of averaged simulation results, and the number of simulation executions sampled in generating them is derived in section 6.3. This is referred to as the *consistency analysis technique*. Section 6.4 describes the application of a global sensitivity analysis, based on latin hypercube sampling, to ARTIMMUS. This analysis explores simulation mechanics, and how the various pathways and features of the simulation constitute its overall behaviour. Section 6.5 derives a novel robustness analysis technique that is used to explore the fragility of simulation behaviour with respect to those simulation parameters that were assigned arbitrary values, as described in the previous chapter. The theory underpinning section 6.5 is expanded upon in section 6.6 to propose a method for employing robustness analysis jointly with sensitivity analysis to qualify the significance of simulation-derived results in terms of the original domain. This chapter is concluded in section 6.7.

6.1 Motivation and goal

This chapter is concerned with the exploration of simulation through statistical techniques, and the interpretation of *in silico* results into the original domain. The chapter addresses research objectives 4 and 5: to perform novel *in silico* experimentation using ARTIMMUS, and to develop and apply statistical techniques for interpreting *in silico* results in the context of EAE.

Chapters 1 and 3 have explored how simulations are abstract representations of their target domains, and how *in silico* results cannot necessarily be directly interpreted into the original domain. Rather, the link between simulation and the real-world must be

established. Real-world experimentation in the Kumar lab makes use of a scale for grading the severity of autoimmunity experienced in an experimental animal on a 6 point scale, reviewed in section 2.2. This chapter make use of statistical techniques in developing a system to grade simulation results on a similar scale. As such, *in silico* results may be considered and discussed in similar terms to those of the wet-lab. Note that results will still not necessarily translate directly across, but the grading mechanism does to some extent bridge the gap.

Agent-based simulations such as ARTIMMUS are stochastic in nature, and a preliminary consideration for interpreting *in silico* results in terms of the real domain is ensuring that they are representative of the simulation itself. This is handled by the novel consistency analysis technique described in this chapter. It is used to establish the number of simulation executions required to minimise the effect of simulation stochasticity on averaged simulation results to an acceptable level.

The chapter makes use of a sensitivity analysis technique in exploring the influence and criticality of cells, pathways and parameters to ARTIMMUS's overall dynamics. Broadly speaking, such methods relate variation in a system's output to variation in its inputs. Global sensitivity analyses operate by simultaneously perturbing system inputs and measuring how the simulation's outputs respond. The pathways, interactions and mechanics that comprise the simulation are annotated by parameter values. For example, CNS macrophages can apoptose neurons through the secretion of TNF- α , but the exact secretion rates and sensitivities to this cytokine are dictated by parameter values. Parameters found to be highly influential on simulation behaviour represent critical pathways and components within the system. Conversely, those found to be relatively uninfluential indicate the either redundancies in the system, or pathways and components that are of no consequence. Global sensitivity analyses that reveal this information can not be performed in real-world experimental animals, and as such they constitute novel *in silico* experimentation.

The last strand of work reported in this chapter concerns establishing the significance of simulation-based results in terms of the original domain. This relates to the certainty one has in the values that simulation parameters are assigned. Chapter 5 details how simulation parameters are devised: some are informed by domain-specific knowledge, some are iteratively tuned in order to obtain results that match *in vivo* behaviour, and some are arbitrarily assigned. If simulation behaviours and predictions critically hinge on parameter values that were assigned arbitrary values, then the predictions may not hold for the real-world domain, and might instead be artifacts of arbitrarily assigned parameter values.

To explore this possibility, this chapter develops and employs a novel robustness analysis to investigate the ranges of parameter values for which simulation behaviour is statistically consistent. If simulation behaviour changes significantly when arbitrarily assigned parameters are perturbed by a small amount, then choosing different values for these parameters may have a substantial impact on simulation results. Such a discovery does not necessarily undermine activities to explore a system through simulation, but it is important to convey such criticalities when presenting simulation results; other members of the scientific community may have reason to query or contest assumptions and parameter values used. If the biologically plausible ranges of parameter values are known, then these robustness and sensitivity analysis techniques may be used to identify the simulation behaviours at the extremities of ranges. These represent the range of behaviours that established data can attest to, and experimentation that results in behaviours exceeding such ranges have a stronger case for being argued as significant in terms of the real domain. The way in which sensitivity and robustness analyses can accomplish this are explored in this chapter.




Figure 6.1: The eleven responses derived from simulation behaviour, displayed on graphs depicting baseline T cell dynamics and mean EAE severity over time. These are: (a), the maximum effector T cell population sizes, the times at which they occur, and the number of CD4Th1 cells at 40 days; (b) maximum EAE severity, and EAE severity at 40 days.

Firstly, however, sensitivity and robustness analyses require that system *responses*, metrics of simulation behaviour, be defined. The results of sensitivity and robustness analyses can only be considered in terms of the system responses. The chapter commences by deriving nine responses based on effector T cell dynamics. A further two responses are derived from assigning EAE autoimmunity severity scores to simulation behaviour.

6.2 Analysis responses

Sensitivity and robustness analysis techniques make use of system *responses* in their analyses. A response is some aspect of system behaviour that may be measured. Analysis results express the relationship between system parameters and system responses.

A number of responses have been adopted in analysing ARTIMMUS. Nine are derived directly from the four T cell populations that exist within the simulation. These are: the maximum number of effector CD4Th1, CD4Th2, CD4Treg and CD8Treg cells reached within the simulation at any point during its execution; the time at which each of those four peaks are attained; and the number of effector CD4Th1 cells residing within the simulation at 40 days post immunisation. These responses are graphically depicted in figure 6.1a. This choice in responses allows analyses to reveal how parameters influence different aspects of autoimmune and regulatory behaviour independently. The strength of each immune response is indicated by the peak number of T cell populations, and the speed of onset is represented by the time at which they occur. The number of effector CD4Th1 cells residing in the system at 40 days reveals the extent to which autoimmunity has been curtailed.

A further two responses result from assigning EAE severity scores to simulation executions. The significance of an *in vivo* experimental intervention is expressed through changes in autoimmunity severity experienced by groups of mice. To better integrate simulation-derived results with those obtained *in vivo*, a means of assigning autoimmunity severity scores to simulation executions has been devised, and is discussed in section 6.2.1.

Two additional responses are based on this *in silico* EAE severity scoring system, depicted in figure 6.1b. They are the severity of EAE at 40 days post-immunization,

and the maximum severity experienced at any point in time. These responses permit analysis of how simulation parameters influence the severity of autoimmunity in the simulation, and the ability of the regulatory immune response to combat it. Further, they align with similar measures employed *in vivo*.

6.2.1 *In silico* EAE Severity Scoring

This section describes the derivation of a system to grade the severity of autoimmunity experienced by simulation executions using the same 0 to 5 scale as employed in real-world experimentation. This scoring permits a better interpretation of simulation-derived results into the real-world domain.

In vivo, the severity of autoimmunity experienced by an experimental animal is scored on a scale of 0 to 5, describing the extent of paralysis in the mouse, as described in section 2.2. The significance of an intervention’s effect is measured in terms of this scale, typically differentiating the maximum scores reached by individuals in control and experimental groups, or the length of time for which EAE symptoms persist.



This disease score is measured by observing an entire mouse; it results from a wide range of complex interactions between the mouse’s immune system and its central nervous system. The number of immune cells, their levels of activity, the degree of inflammation, and the resilience of the mouse’s central nervous system to the effects of immune system cells all contribute to the overall severity of paralysis.

Figure 6.2a shows the *in vivo* progression of EAE in groups of five mice. The mice depicted in window A experience physiological recovery, through the action of regulatory T cells. The mice in window C are similarly induced into EAE, however their CD4Treg populations have been depleted until around 30 days. It is difficult to differentiate the progressions of individual mice, as such figures 6.2b and 6.2c provide clarification. The progressions of EAE differ significantly between individual mice. EAE symptoms in control mice, window A, are generally less severe than in those in which CD4Treg cells have been temporarily depleted, window C.

Further evidence of the variation of autoimmunity experienced amongst individual mice, and the influence of regulation in mediating recovery, is provided by [Beeston *et al.* 2010]. The authors investigate the manner through which CD8Treg cells induce apoptosis in encephalitogenic CD4Th1 cells, and conclude that regulation is critically dependent on the perforin pathway¹ [Beeston *et al.* 2010]. The authors investigate the potential for pre-immunization vaccination with Fr3 and CDR1/2 peptides, derived from encephalitogenic CD4Th1 cells, to reduce autoimmune symptoms in mice. The control groups of mice receive sham vaccinations, whereas experimental groups receive the peptides. This is performed in both wild-type (perforin *+/+*) and perforin *-/-* (knockout) mice, comprising a total of four experimental groups. The maximum EAE scores experienced by individuals in each of these four groups are reported, and are replicated in table 6.1a. The proportion of mice that reach a maximum of each possible score can be calculated, as shown in table 6.1b.

In Beeston *et al.*’s perforin-knockout mice, immunisation for EAE results in a \sim 40% mortality rate, as opposed to \sim 10% in control mice. 78% of control mice experience maximum EAE severities of level 3 or less, whereas only 22% of mice with disabled regulatory capacity experience severities of level 3 or less. Whereas [Kumar *et al.* 1996]’s experimental procedure temporarily disables regulation through depletion of CD4Treg populations for 3 weeks, the regulatory pathway in perforin-knockout mice is permanently obstructed. As a result, the severity of autoimmune symptoms are increased.

¹The perforin pathway is one of many through which cells can induce apoptosis in one another.

(a) Data obtained from [Kumar *et al.* 1996] figure 3, windows A and C.

(b) Clarification of window A.

(c) Clarification of window C.

Figure 6.2: Example *in vivo* data on the progression of EAE in individual mice under conditions of physiological recovery, regulation temporarily disabled. Data taken from [Kumar *et al.* 1996] figure 3, windows A and C. Figures 6.2b and 6.2c provide clarification of the data; hence lines are slightly staggered.

A faithful *in silico* EAE severity scoring system should reproduce the variation in the autoimmunity progressions of individual mice, and the difference in progressions resulting from intact or obstructed regulatory capacity. The remainder of this section details how data from [Beeston *et al.* 2010] and [Kumar *et al.* 1996] is used to calibrate an *in silico* EAE severity grading system meeting these requirements. The grading system operates by relating the rate of neuronal apoptosis in the simulation to EAE severity scores.

6.2.1.1 Relating rate of neuronal apoptosis to EAE severity

The *in silico* EAE severity scoring system operates by assigning scores based on the rate of neuronal apoptosis, since this is a direct result of autoimmune activity. The relationship between rates of neuronal apoptosis and the corresponding EAE severity score is based on threshold values for the number of neurons apoptosed in an hour of simulated time.

Threshold values are calculated through calibration against *in vivo* data, obtained from [Beeston *et al.* 2010] as described above, in which mice undergoing two experimental procedures are investigated. These are: the physiological recovery from EAE in

6. EXPLORATIVE EXPERIMENTATION

Mice	Treatment	Maximum score for each mouse
Perforin +/+	control	(5, 4, 3, 3, 2, 2, 2, 1, 1)
	vaccination	(3, 1, 1, 1, 0, 0, 0, 0, 0, 0)
Perforin -/-	control	(5, 5, 5, 4, 4, 4, 3)
	vaccination	(5, 5, 5, 4, 4, 4, 3, 3, 1)

(a) Wildtype (perforin +/+) and perforin -/- (knockout) mice vaccinated with either CD4Th1-derived peptides or control substance, and induced into EAE. Data is replicated from table 2 of [Beeston *et al.* 2010].

Score	PR (%)	RD 1 (%)	RD 2 (%)	RD 1 + 2 (%)
5	11.1	42.9	36.4	38.9
4	11.1	42.9	36.4	38.9
3	22.2	14.3	18.2	16.7
2	33.3	0	0	0
1	22.2	0	9.1	5.6
0	0	0	0	0

(b) The proportion of mice that reach maximum EAE severity scores. PR, physiological recovery, perforin +/+ control. RD 1, regulation disabled, perforin -/- control. RD 2, perforin -/- vaccination. RD 1 + 2, combined groups of RD 1 and RD 2; these groups may be combined since perforin -/- mice are unable to regulate encephalitogenic CD4Th1 cells, regardless of whether they received control or vaccination substances.

Table 6.1: *In vivo* data showing the maximum scores reached by different experimental groups of mice.

wild-type mice, where regulatory action is intact; and mice experiencing exaggerated autoimmune symptoms since the ability of CD8Treg cells to regulate the autoimmune response has been permanently disabled. Beeston *et al.*'s data is reproduced in table 6.1a, and the proportions of mice experiencing particular maximum EAE severities following immunisation are reported in table 6.1b.

These proportions are used to calibrate threshold values relating the rate of neuronal apoptosis to *in silico* EAE severity scores. Threshold values are derived from physiological recovery data and disabled regulation data independently, and are then combined. The physiological recovery data adopted is the 'PR' column in table 6.1b. The regulation disabled data adopted is the 'RD 1 + 2' column in the same table, which is derived from combining the perforin -/- vaccination and control groups; since perforin -/- mice are unable to regulate the autoimmune response the vaccination is of no consequence, and the control and vaccination groups may be combined to provide more data samples.

Representative simulation behaviours reflecting physiological recovery and regulation disabled experimental conditions are obtained in the form of 1000 simulation executions for each. Physiological recovery behaviour is obtained using the standard baseline simulation parameters. Conditions of disabled regulation are achieved through setting the simulation parameter `CD8Treg_cd8TregToCD4HelperSpecificityDropOff` to 0.0, which completely prevents CD8Tregs from inducing apoptosis in CD4Th1 cells. The maximum number of neurons apoptosed in any hour is extracted for each simulation execution. Expressed as a cumulative distribution function, this data indicates the proportion of simulations that reach *at most* a particular rate of neuronal apoptosis. The solid blue and red lines on figures 6.3a and 6.3b represent the cumulative distribution functions for 1000 simulation executions under conditions of physiological recovery and disabled regulation respectively. Under physiological recovery, simulation neuronal

apoptosis rates always approach zero around day 35, with peak rates occurring around day 15. However, in the case of disabled regulation neuronal apoptosis rates do not return to zero, and high peak values may occur much later following immunisation. As such, cumulative distribution plots are based neuronal apoptosis rates observed within the first 40 days of simulation time for physiological recovery, and the first 60 days for regulation disabled. Figures 6.4a and 6.4b show the number of neurons killed per hour in example simulation executions under each experimental condition.


The proportion of *in vivo* mice that experience particular maximum severities of autoimmunity are used to define the thresholds relating rates of neuronal apoptosis to each severity score. This is demonstrated in figure 6.3. For example, column 'PR' in table 6.1b indicates that 11% of mice reach a maximum score of 5. Figure 6.3a indicates that only 11% of simulation runs reach a maximum rate of neuronal apoptosis of 111 neurons or more per hour. As such, 111 neurons per hour is mapped to level 5.

As may be seen in figures 6.3a and 6.3b, the threshold values obtained for physiological recovery do not exactly reflect those obtained under disabled regulation. The final threshold value adopted for each severity score is the median average of the values obtained under each data set. This is demonstrated in figures 6.3c and 6.3d.


Figure 6.4 depicts the assignment of EAE severity scores to two examples simulation executions, derived under conditions of physiological recovery and disabled regulation. EAE severity scores are assigned to mice *in vivo* every one or two days. *In vivo*, the severity of autoimmunity does not change sufficiently within 24 hours to warrant a more frequent sampling. In order to align *in silico* autoimmunity grading with wet-lab practice, the rate of neuronal apoptosis is sampled once every 24 hours, and severity scores assigned based on these samples. The sample rate of neuronal apoptosis when deriving cumulative distribution functions from which threshold values are extracted is also once every 24 hours. Hence, although both examples exceed the level 5 threshold, neither are assigned this score since the breach does not fall on one of the daily sample points. Note that where such a breach does occur, level 5 is assigned for the remainder of observation time, regardless of the neuronal apoptosis data from that point onwards; level 5 represents death.

6.2.1.2 Smoothing neuronal apoptosis data

Figures 6.4a and 6.4b depict the number of neurons apoptosed each hour for example simulation executions. There is considerable variation in the number of neurons apoptosed from one hour to the next, and one day to the next. As a result, *in silico* EAE severities scores change score frequently, and with large magnitudes. This is not the case with the progression of EAE as observed *in vivo*, figure 6.2, where scores remain stable for several days at a time. Creating an accurate reflection of *in vivo* EAE progression in ARTIMMUS requires that the trend of neurons apoptosed per hour be extracted from this data, with the high frequency components removed. This is achieved through smoothing the data using a sliding window filter. Sliding window filters take as a parameter a window size. They operate by adjusting each value in the data set to the median of all values in the window, with the window being centered on the value being adjusted. The effect of applying the sliding window filter with a variety of smoothing window sizes is depicted for an example simulation execution in figure 6.5




(a) Defining thresholds relating neurons apoptosed in an hour to EAE severity scores for physiological recovery data. The solid blue line represents the cumulative distribution function.



(b) Defining thresholds relating neurons apoptosed in an hour to EAE severity scores for regulation disabled data. The solid red line represents the cumulative distribution function.

Score	Neurons apoptosed / hour		
	Physiological recovery	Regulation disabled	Final
5	111	106	108.5
4	104	96	100
3	98	89	93.5
2	87	89	88
1	65	75	70

(c) Threshold values for physiological recovery, regulation disabled, and the final values adopted based on their median.



(d) Final thresholds, based on the median values for thresholds defined using data physiological recovery and regulation disabled, figures 6.3a and 6.3b respectively.

Figure 6.3: Demonstration of how threshold values for relating the number of neurons apoptosed in an hour of simulation time to EAE severity scores is achieved. Figure 6.3a shows how thresholds are selected based on the proportion of simulation runs that should attain particular maximum EAE scores for physiological recovery data. Figure 6.3b shows thresholds for regulation disabled data. The median for each of these thresholds form the final thresholds used in grading *in silico* EAE severities.




Figure 6.4: The assignment of *in silico* severity scores by thresholding on the rate of neuronal apoptosis, for conditions of physiological recovery and disabled regulation. EAE severities are assigned by sampling rates of neuronal apoptosis every 24 hours, hence level 5 is not assigned, despite the corresponding threshold being exceeded in several hours.

The sliding window filter is applied to neuronal apoptosis time series data both when deriving threshold values through the activities described above, and also to simulation data before applying thresholds to grade *in silico* autoimmune severities.

In addition to removing the high frequency variation in the neuronal apoptosis data, smoothing in this manner alters the cumulative distribution functions for the maximum number of neurons apoptosed per hour. As such, the thresholds relating rates of neuronal apoptosis to EAE severity scores are altered. Further, the alterations to cumulative distribution functions obtained under conditions of physiological recovery and disabled regulation differ, since the rate of neuronal apoptosis under physiological recovery returns to zero, whereas under disabled regulation it does not. As such, some choices in smoothing window size better align the thresholds obtained under these two data sets than others. The choice of window size influences the final thresholds adopted in grading simulation executions, and different smoothing functions affect the final proportions of simulation runs that reach maximum severity scores. These final proportions for *in silico* EAE progressions may be compared to the initial *in vivo* specifications. The choice in smoothing function is, hence, non-trivial. These observations together form the following five criteria used to establish the optimal choice in smoothing function, the exact measures for each are explained in more detail in the sections that follows:

1. The frequency and magnitude of changes in *in silico* EAE scores over time must reflect that of *in vivo* data, both for conditions of physiological recovery and disabled regulation.
2. The threshold values relating rates of neuronal apoptosis to each EAE severity, obtained under conditions of physiological recovery and disabled regulation, must lie in close proximity to one another.
3. The proportion of *in silico* simulation executions reaching maximum EAE scores must resemble the *in vivo* specification, both for conditions of physiological recovery and disabled regulation.

The degree to which different choices of smoothing function satisfy each of these criteria are expressed in terms of sum of squared differences measures, each of which is described in detail in the sections that follow. The final choice in smoothing function is based on these measures, and is detailed in section 6.2.1.6. A variety of smoothing window sizes are investigated, ranging from 1 hour (no smoothing) to 481 (smoothing based on 10 days either side of every data item).

6.2.1.3 Measuring frequency and magnitude of changes in *in silico* EAE progression

Figure 6.7 demonstrates that when no smoothing is applied to neuronal apoptosis data when deriving and applying threshold values indicating autoimmune severities, the resultant EAE progressions contain high frequency fluctuations. *In vivo* mice do not typically lapse into and out of EAE symptoms every day, and do not typically experience day-to-day changes spanning several severity scores. The alignment between frequency and magnitude of severity score changes for *in silico* and *in vivo* data is measured using the Fourier transform of the respective data sets [Bracewell 1999].

The Fourier transform is a mathematical operation that decomposes a signal, in this case EAE severity scores, into its constituent frequencies. It expresses a signal in terms of sine waves, indicating the magnitude and phase of each. The magnitude




Figure 6.5: Variation in the number of neurons apoptosed per hour of simulation time, and the effect of different sizes of smoothing window when smoothing data. Data is of an example simulation execution under condition of physiological recovery.

of each component sine wave, with phase information removed, is termed the *power*. Application of the Fourier transform to *in vivo* EAE progression data sets gives an indication of how frequently *in vivo* severity scores change, and how big these changes are. Similar application of the Fourier transform to *in silico* EAE progressions permits comparison with *in vivo* data.

The Fourier transform has been applied to the two *in vivo* EAE progression data sets outlined in [Kumar *et al.* 1996], discussed above and depicted in figure 6.2. These datasets are referred to as ‘Kumar Fig3A’ and ‘Kumar Fig3C’. The Fourier transform has also been applied to an additional data set obtained from the Kumar lab, termed the ‘B10.PL’ dataset, depicted in figure 6.6. In this experiment mice have been immunized in a manner that does not entail complete recovery by 30 days. It is included here as an additional example of the frequency with which *in vivo* EAE scores change. These Fourier transform data allow comparison between *in vivo* and *in silico* EAE progressions, and are shown alongside similarly derived simulation data on the figures that follow. The Fourier transforms of *in vivo* data shown on such figures are the medians, taken across the individual mice comprising the experimental group. In all cases, the Fourier analysis is applied to the first 30 days of *in vivo* and *in silico* data. This is to permit a fair comparison between *in silico* and *in vivo* data. The ‘Kumar Fig3A’ data represents an experiment considered equivalent to the condition of physiological recovery, and the ‘Kumar Fig3C’ dataset is considered equivalent to that of disabled regulation².

The high frequencies and magnitudes with which *in silico* EAE progressions change score when no smoothing is applied to the neuronal apoptosis data is depicted for two examples in figure 6.7. The transforms of the simulation data are shown in the context of those for the three *in vivo* data sets. The Fourier transform of the example in figure 6.7a reveals larger high frequency powers and smaller low frequency powers than

²Whilst the procedure employed to disable regulatory action in this dataset is not permanent, its effects are believed to extend beyond 30 days.




Figure 6.6: EAE progression of 7 mice of the ‘B10.PL’ strain, having been immunized for EAE. Data obtained from the Kumar lab.

the *in vivo* data sets. Figure 6.7b also reveals unrealistically large high frequency powers.

Figure 6.8 depicts EAE progressions and corresponding Fourier transforms of the same two simulation executions, but using a different mapping of neuronal apoptosis rates to EAE severities. Here the mapping is derived and applied using a smoothing window size of 157 hours. As such, simulation and *in vivo* Fourier transforms are more closely aligned.

The above examples demonstrate that the sliding widow sizes influence the frequency and magnitude of changes *in silico* EAE scores. A variety of smoothing window sizes have been examined, and their ability to align *in silico* EAE progressions with *in vivo* data measured through the following procedure.

1. Given a smoothing window size, the threshold values relating neuronal apoptosis rates to severity scores are calculated.
2. These thresholds are used to calculate EAE progressions for each of 1000 simulation runs. The Fourier transform of each EAE progression is taken, and median values for the powers across all 1000 runs are calculated. This is performed independently for conditions of physiological recovery and disabled regulation.
3. Averaged *in silico* Fourier transforms are contrasted with their *in vivo* counterparts by calculating sum of squared difference measures. *In silico* physiological recovery data is contrasted with the Kumar Fig3A data set, whilst disabled regulation data is contrasted with the Kumar Fig3C data set.
 - a) For each frequency component in turn, the powers of the *in silico* and *in vivo* transforms are identified. Both are then divided by whichever is larger, hence transforming absolute numbers into normalised numbers.
 - b) The difference between these normalised figures is then calculated, and squared.

- c) The measure of difference between the *in silico* and *in vivo* data set is obtained by summing these squared normalised differences across all frequency components. The normalisation ensures that the difference between each frequency component is treated equally, rather than being biased towards those that have larger absolute power values.

The sum of squared normalised difference measures obtained for various smoothing window sizes are shown on figure 6.9.


6.2.1.4 Measuring alignment of thresholds derived under physiological recovery and disabled regulation

As noted previously, the size of window used when smoothing neuronal apoptosis data influences the cumulative distribution functions that describe the proportion of simulation executions that reach particular maximum numbers of neurons apoptosed in an hour. As such, the thresholds derived from identifying the proportion of simulation runs required to reach a particular level of EAE severity change with window size. The cumulative distribution functions derived under conditions of physiological recovery and disabled regulation do not change in the same manner, and as such the alignment of threshold values derived under each of these conditions changes with the choice of smoothing window size.


The alignment between the threshold values obtained under each condition is measured using a sum of squared normalised differences measure, which informs the eventual choice in smoothing window size. Given a particular smoothing window size, the physiological recovery and regulation disabled data sets are smoothed, and the threshold values relating the number of neurons apoptosed in an hour to EAE severity scores for each data set are derived, as described above in section 6.2.1.1. For each severity score in turn, the two threshold values for that severity are divided by whichever is larger, hence expressing one threshold as a proportion of the other³. The difference between these normalised threshold values is calculated and then squared. The sum of these squares for all 5 severity levels forms the eventual score for how well threshold values are aligned given a particular smoothing window size.

Figure 6.10 shows how different smoothing window sizes align threshold values, as defined through the sum of squared normalised difference measures described above.

³As with the previous sum of squared difference measure, differences here are normalised to ensure that each component of the sum of squares measure is treated equally, rather than biasing towards larger severity scores which will inherently be associated with larger figures for neuronal apoptosis rates.




(a) Example EAE progression in simulation of physiological recovery, and corresponding Fourier transform.




(b) Example EAE progression in simulation with regulation disabled, and corresponding Fourier Transform

Figure 6.7: Example EAE progressions for simulations under conditions of physiological recovery and disabled regulation. The mapping of neuronal apoptosis rates to EAE severity was derived using no smoothing of neuronal apoptosis data. The resultant EAE progressions are unrealistic with respect to the frequency with which EAE scores change, as demonstrated by the high powers of low period components in the Fourier transforms.



(a) Example EAE progression in simulation of physiological recovery, and corresponding Fourier transform.



(b) Example EAE progression in simulation with regulation disabled, and corresponding Fourier Transform

Figure 6.8: Example EAE progressions for simulations under conditions of physiological recovery and disabled regulation. The mapping of neuronal apoptosis rates to EAE severity was derived using a smoothing window size of 157 hours. The resultant EAE progressions closely resemble the frequency analysis characteristics of *in vivo* EAE.




Figure 6.9: The sum of squared normalised difference measures indicating how well various smoothing window sizes align frequency and magnitude of changes in *in silico* EAE scores over time with those observed *in vivo*.




Figure 6.10: The sum of squared normalised differences measure indicating the alignment of threshold values relating neurons apoptosed per hour to EAE severity scores derived under conditions of physiological recovery and disabled regulation for various smoothing window sizes.

6.2.1.5 Measuring alignment of proportions of simulation and *in vivo* data reaching maximum severity scores

The final thresholds adopted in relating rates of neuronal apoptosis to EAE severity scores are derived from thresholds calculated under conditions of physiological recovery and disabled regulation, as described above. Median values are taken, since the values obtained under each condition do not align exactly. Hence, in adopting median threshold values, the proportions of *in silico* executions that reach particular maximum severity scores under each circumstance no longer exactly match their *in vivo* counterparts. The discrepancy between actual *in silico* and *in vivo* proportions depends on where the discrepancies in threshold values lies on the cumulative distribution functions used to derive them. This may be seen in figure 6.3, by observing how threshold values move between graphs (a) and (b), and (c). The size of the discrepancies between proportions of *in silico* and *in vivo* executions reaching maximum severity scores must be considered when selecting a smoothing window size. Discrepancies are measured for both physiological recovery and disabled regulation conditions, using the following procedure:


1. Given a smoothing window size, the final thresholds used to grade simulation EAE severities are calculated.
2. The same smoothing window size is used to smooth the neuronal apoptosis time series data from 1000 simulation executions.
3. Each simulation execution is graded using the final thresholds, and the maximum score attained in each is extracted. Hence, the proportion of simulation executions reaching each maximum severity is calculated.
4. For each EAE severity score (0 to 5) the difference between the proportion of the population experiencing that maximum score *in vivo* and *in silico* is calculated and squared. The squared differences for all severities are then summed to derive the final measure.

The procedure is performed twice, using both physiological recovery and disable regulation simulation data. Figure 6.11 depicts these discrepancies between *in silico* and *in vivo* proportions for each condition, for various smoothing window sizes. Figure 6.12 shows the sum of squared differences measures for various smoothing window sizes.


6.2.1.6 Determining an appropriate size of smoothing window

The most appropriate smoothing window size to be adopted in mapping neuronal apoptosis rates to EAE severity scores is determined through the five sum of squared differences metrics described in the above sections. These metrics describe how well various smoothing window sizes deliver EAE progressions that change with similar frequency and magnitudes as observed *in vivo*, figure 6.9; align the threshold values obtained under conditions of physiological recovery and disabled regulation, figure 6.10; and deliver similar proportions of *in silico* and *in vivo* individuals reaching particular maximum severity scores, figure 6.12.

In establishing the most appropriate smoothing window size, data derived through each of these five metrics have been normalised in turn, dividing the sum of squared difference value for each smoothing window size by the maximum value for any smoothing window size. The normalised data from each of these five measures is plotted on



(a) Physiological recovery.



(b) Regulation disabled.

Figure 6.11: The proportion of EAE progressions reaching particular maximum severity scores, for various smoothing window sizes. Values for *in silico* data are depicted as solid lines, whereas values for *in vivo* data are dotted lines. Data derived under conditions of physiological recovery and disabled regulation are shown on separate graphs. Note that where the proportions for multiple *in vivo* severity scores are equal, see table 6.1b, lines have been separated slightly for clarity.






Figure 6.12: The sum of squared difference measures indicating the separation between the proportion of *in silico* and *in vivo* individuals reaching particular maximum EAE severity scores, obtained under various smoothing window sizes. Data obtained under conditions of physiological recovery and disabled regulation are shown.

figure 6.13a. The normalised data for each metric are summed for each smoothing window size, as depicted on figure 6.13b, and the smallest sum of sum of squared differences value is associated with a smoothing window size of 157 hours. EAE severity scores assigned to *in silico* data for the remainder of this thesis are done so through application of threshold values derived using this smoothing window size, and applied to simulation neuronal time series data having been smoothed in a similar manner. The threshold values and their derivation are shown in figure 6.14.




(a) Normalised sum of squared difference measures for the five metrics used to identify how well various smoothing window sizes satisfy each of the five criteria



(b) The sum of the five metrics shown in figure 6.13a. The smallest value is associated with a smoothing window size of 157 hours.

Figure 6.13: Determining that a smoothing window size of 157 hours best satisfies the five criteria for the most appropriate mapping of neuronal apoptosis rates to *in silico* EAE severity scores.




(a) Defining thresholds relating neurons apoptosed in an hour to EAE severity scores for physiological recovery data. The solid blue line represents the cumulative distribution function.

(b) Defining thresholds relating neurons apoptosed in an hour to EAE severity scores for regulation disabled data. The solid red line represents the cumulative distribution function.

Score	Neurons apoptosed / hour			Proportions (<i>in vivo</i> <i>in silico</i>)	
	PR	RD	Final	PR(%)	RD(%)
5	79.29	77.98	78.64	11.1 14.5	38.9 33.8
4	76.79	72.95	74.96	11.1 20.0	38.9 29.5
3	73.82	69.48	71.65	22.2 25.0	16.7 22.1
2	69.17	69.48	69.32	33.3 17.0	0.0 9.7
1	49.99	63.32	56.66	22.2 22.8	5.6 4.9
0				0.0 0.2	0.0 0.0

(c) Threshold values for physiological recovery (PR), regulation disabled (RD), and the final values adopted based on their median. The proportions of mice reaching particular maximum EAE severity scores, and the proportions of simulations reaching similar scores using final threshold values are also shown.



(d) Final thresholds, based on the median values for thresholds defined using data physiological recovery and regulation disabled, figures 6.14a and 6.14b respectively.

Figure 6.14: The derivation of threshold values for relating rates of neuronal apoptosis to *in silico* EAE severity scores, and the proportions of simulation reaching particular maximum EAE severity scores, for the final smoothing window size of 157 hours.

6.3 Compensating for aleatory uncertainty

Sensitivity and robustness analyses are performed on a system in order to elucidate the relationships between its inputs and its outputs. Stochastic systems, however, introduce noise (termed *aleatory uncertainty*) [Helton 2008]) into these analyses. Two simulation executions, with different seeds, can produce vastly different results given the same input parameters. The results obtained from *in silico* experimentation are influenced by both the nature of the system, and the aleatory uncertainty that arises within it. Employing averaged results from many executions of a stochastic system can reduce the influence of aleatory uncertainty on analysis results, making them more representative of a simulation’s dynamics, and the experimentation performed upon it.

Hence, there exists a tradeoff in which a balance must be struck. Stochastic-agent based simulations, such as ARTIMMUS, can be computationally expensive to execute. The number of simulation executions that are sampled must represent a balance between minimising the effect of aleatory uncertainty in simulation results, and the computational expense that can be afforded. The following procedure is employed to elucidate the relationship between these two conflicting criteria, and hence inform the choice of the number of simulation executions employed in conducting *in silico* experimentation. It is termed the *consistency analysis* procedure.

The procedure is based around the Vargha-Delaney *A* test, a non-parameteric effect magnitude test that indicates the scientific significance of the difference between two populations of samples [Vargha & Delaney 2000]. It compares two populations of samples, A and B, and returns a value between 0.0 and 1.0 that indicates the probability that a randomly chosen sample from population A is larger than a randomly chosen sample from population B. In this manner it may be used to assess the difference between two populations of samples. A value of 0.5 indicates no difference, where values above 0.71 or below 0.29 indicates a “large” difference in the distributions [Vargha & Delaney 2000]. Table 6.2 details how *A* test scores relate to various magnitudes of difference between two populations. The matlab code used to generate *A* test scores is provided in appendix section C.1.

Difference	Large	Medium	Small	None	Small	Medium	Large
<i>A</i> test score	<0.29	<0.36	<0.44	0.50	>0.56	>0.64	>0.71

Table 6.2: The magnitude of effect size indicated by *A* test scores [Vargha & Delaney 2000].

The procedure operates by assessing the scientific significance of stochasticity on distributions of simulation response values, and how this may be reduced by increasing the sample size of simulation executions used in gaining representative indications of simulation behaviour. The *A* test is used to compare response distributions obtained using the exact same parameters, and hence any difference observed is attributed purely to stochasticity in the simulation. The procedure is as follows.

For each distribution sample size n to be investigated, twenty sets of n simulation executions are derived, using the baseline simulation parameter values. Since no parameter values are changed, the distributions from which response samples (simulation executions) are being drawn remains static. The resulting twenty sets of samples will more closely resemble one another as n increases. With respect to the ARTIMMUS simulation, eleven distributions are obtained for each of the twenty sets of experimentation, based on the eleven responses described above in section 6.2. The distributions arising from the first of the twenty sets of simulation runs is compared to corresponding distributions from the remaining nineteen sets using the *A* test. A score of 0.5 is always

indicated for the first set of experimentation, since it this data is being compared to itself. Plotting each of the eleven A test scores against the twenty sets of experimentation indicates the magnitude of effect that aleatory uncertainty has on simulation results for different sample sizes.

Figure 6.15 demonstrates how sample sizes of 5, 50 and 500 influence the effect of aleatory uncertainty on A test scores. Increasing the number of samples used when generating response distributions decreases the differences in those distributions that arise because of aleatory uncertainty. The maximum A test scores across these 19 measures can be plotted against the sample size n , as in figure 6.16. Results show that at least 500 samples are required to consistently reduce the effect of aleatory uncertainty to a less-than-small effect magnitude size across all responses. Unless otherwise stated, the remaining experimentation reported in this thesis employs a sample size of 500 simulation runs when deriving representative results. This figure represents an acceptable computational expense, and satisfactorily reduces the effect of aleatory uncertainty on analysis results.





Figure 6.15: The effect of aleatory uncertainty on the results of A test analysis for eleven responses, for various sample sizes. The first set of simulation executions is compared with each of the remaining nineteen in obtaining A test scores.

(a) Maximum A test scores obtained across twenty sets of experimentation.

Response	Sample Size							
	5	10	50	100	200	350	500	1000
<i>CD4Th1 Max</i>	0.80	0.70	0.66	0.58	0.57	0.55*	0.53*	0.52*
<i>CD4Th1 Max Time</i>	0.96	0.68	0.60	0.60	0.53*	0.55*	0.54*	0.53*
<i>CD4Th2 Max</i>	0.88	0.81	0.62	0.59	0.56*	0.57	0.53*	0.53*
<i>CD4Th2 Max Time</i>	1.00	0.67	0.57	0.57	0.55*	0.55*	0.53*	0.53*
<i>CD4Treg Max</i>	0.76	0.65	0.55*	0.55*	0.56*	0.54*	0.54*	0.52*
<i>CD4Treg Max Time</i>	0.92	0.69	0.59	0.62	0.55*	0.56*	0.54*	0.52*
<i>CD8Treg Max</i>	0.92	0.88	0.60	0.56*	0.55*	0.54*	0.53*	0.52*
<i>CD8Treg Max Time</i>	0.92	0.65	0.61	0.60	0.53*	0.55*	0.54*	0.52*
<i>CD4Th1 @ 40d</i>	0.70	0.60	0.55*	0.57	0.53*	0.53*	0.52*	0.51*
<i>Max EAE</i>	0.70	0.77	0.61	0.56*	0.54*	0.54*	0.53*	0.52*
<i>EAE @ 40d</i>	0.70	0.60	0.56*	0.55*	0.53*	0.52*	0.53*	0.51*

(b) Maximum A test scores, rounded to two significant figures. Scores indicating an effect magnitude of less than small are indicated with *.Figure 6.16: Maximum A test scores across twenty sets of experimentation, using the same parameter values in all cases, for various responses and sample sizes used in generating distributions.

6.4 Determining influence of parameters

This section describes the application of a global sensitivity analysis technique to the ARTIMMUS simulation. The sensitivity analysis reveals the sensitivity of simulation responses to perturbation of its parameters. The pathways and activities that comprise the simulation are annotated by parameter values that provide the specifics of these mechanics. These include, for example, the periods of time that cells spend in particular states, the rates of cytokine secretion and cellular sensitivities, and the probabilities that particular events lead to particular actions. By identifying the relative contributions of these parameters to simulation behaviour, sensitivity analysis affords insight into the relative influences and interactions of the mechanics of the simulation. It highlights simulation components that are critical to system operation, and redundancies in the system.

The global sensitivity analysis employed here leads an exploration of ARTIMMUS components, gaining insight into how the system operates. The experimental procedure employed is described in section 6.4.1. The results of this analysis are detailed in section 6.4.2.

6.4.1 Experimental procedure

The sensitivity analysis employed here is *global* in that all parameters analysed are perturbed simultaneously. This provides a more thorough exploration of simulation space than *one at a time* alternatives where each parameter is perturbed independently of others. Global techniques show compound effects where a parameter's influence depends on the exact values held by others. The sensitivity analysis employed here is based on latin hypercube design [McKay *et al.* 1979], which represents an efficient sampling of parameter space using a minimal number of samples.

A latin hypercube design is used to obtain 500 samples of parameter space around the baseline. Figure 6.17 depicts an example latin hypercube design which obtains ten samples across two parameters. The region indicated between upper and lower limits for each parameter to be included in the design is divided into n regions, where n is the number of samples to be obtained by the design. In the example of figure 6.17, and the design used in analyzing the EAE simulation, the regions are uniformly distributed, though this need not be the case if ensuring a high number of samples in a particular area of a parameter's domain is desired. The latin hypercube procedure ensures that each of the n regions for every parameter is sampled exactly once.

The range over which parameters are perturbed is similar for each parameter, and is defined as a proportion of the parameters' baseline values. The domain over which each parameter is perturbed was chosen to be between 10% above and 10% below that parameter's baseline value. When creating a latin hypercube design, each parameter perturbation is calculated as a real valued number lying in this domain. However, the type of the parameter being perturbed dictates that further manipulation may be required. A full list of ARTIMMUS simulation parameters, including their types and default values, may be found in appendix section B.1.

The EAE simulation parameters analysed here are of three types: real valued, probabilistic, and natural number valued. For real valued parameters, the values returned by the latin hypercube design may be directly adopted.

Probabilistic parameters may hold any real value between 0.0 and 1.0. If the baseline value of this parameter lies such that 10% above or below breaches one of these thresholds, then the perturbation range is shifted such that the size of the range is maintained, but falls between 0.0 and 1.0. For example, *CD8Treg_cd8TregToCD4T-*




Figure 6.17: An example latin hypercube design, obtaining ten samples across two parameters.

`helperSpecificityDropOff`, which dictates the probability of a CD8Treg cell apoptosing an effector Qa-1 expressing CD4Th1 cell upon binding, is ordinarily set to 1.0. The perturbation range over which the latin hypercube design will assign values hence falls between 0.9 and 1.1. Owing to the probabilistic domain of this variable, the perturbation range would be adjusted to fall within the boundaries of 0.8 to 1.0.

Simulation parameters that take as values only natural numbers, such as `Simulation_numCNS`, which dictates the number of neurons in the CNS compartment, must be rounded to the nearest natural number before the simulation may be executed. Hence, the perturbation range for these parameters will be approximately 10% above and below the baseline value, but may not fall exactly on those bounds. For parameters which may take only even values, values are rounded to the nearest even natural number.

The latin hypercube design employed here creates 500 samples of simulation parameter space. Due to the stochastic nature of the simulation, acquiring representative simulation responses for each sample requires obtaining averaged simulation results. 500 simulation executions are attempted at each sample point in the latin hypercube design, in accordance with the analysis of section 6.3 above. However, it is possible that the latin hypercube design selects points in parameter space at which simulation execution is computationally intractable. There exists finite space within the compartments of the simulation, were the latin hypercube design to select parameters such that T cells proliferate (divide) quickly for extended periods of time, going on to exist as effector cells for extended periods of time, it is possible that simulation execution will stall as there are more cells than space to contain them. In such extreme cases where it is not possible to obtain a full set of 500 simulation executions for a particular sample point in parameter space, that sample is excluded from the analysis.

For all the remaining sample points, response distributions are calculated from the 500 simulation executions, based on the responses outlined in section 6.2 above. For the present analysis, averaged results are obtained from each response distribution. This is the median value for all T cell related responses, but mean values are used for the two EAE based responses. The T cell responses may adopt a very large range of values, however the two EAE based responses can assume only values of 0 to 5. Although the underlying distributions are known to be non-normal, employing mean averaged results permits a greater number of response values to be assumed, and hence

facilitates a higher fidelity analysis of the influence that a particular parameter has on EAE severity scores. Using mean averaged results is also consistent with wet-lab practice.

The analysis proceeds by considering each parameter in turn. The 500 points in parameter space are ordered in accordance to their value for the particular parameter of interest. The correlations between averaged response values and the value for the parameter of interest are calculated. Partial rank correlation coefficients are employed here.

Partial rank correlation coefficients (PRCC) give a measure of rank correlation between two variables, the simulation parameter and response, with the linear effects of other simulation parameters on the response removed [Marino *et al.* 2008]. As such, the effect of an imperfect latin hypercube design, where there exists correlation between two or more parameter values, are somewhat mitigated.

The p values that describe the likelihood of observing a particular correlation purely through chance are also recorded. It should be noted that the results of the present analysis indicate influences of parameter values only for that volume of parameter space around the baseline, within which the simulation's parameters are perturbed. Whilst some continuity of these influences may be assumed, it is possible that there exist other areas of parameter space where the significance of simulation parameters differ considerably.

6.4.2 Results

The sensitivities of the simulation with respect to each response and parameter are recorded in tables C.1 to C.11 in appendix section C.2. These tables report the PRCC, associated p values, and ranks of influence of simulation parameters with respect to each of the eleven responses. Table 6.3 summarises the ranks of each parameter with respect to each response. An indication of global importance of each parameter is achieved through summing all the individual ranks to create a total.

Table 6.3: The total influence of each parameter on simulation responses, measured as a sum of ranks for each response in turn. Responses are indicated as follows: 1M, $CD4Th1$ Max; 1MT, $CD4Th1$ Max Time; 2M, $CD4Th2$ Max; 2MT, $CD4Th2$ Max Time; 4M, $CD4Treg$ Max; 4MT, $CD4Treg$ Max Time; 8M, $CD8Treg$ Max; 8MT, $CD8Treg$ Max Time; Th40, $CD4Th1$ at 40 Days; ME, Max EAE; E40, EAE at 40 Days. Ranks marked with either (+) or (-) represent P values ≤ 0.01 , and indicate whether the correlation of parameter value with response was positive or negative.

Rank	Parameter Name	1M	1MT	2M	2MT	4M	4MT	8M	8MT	Th40	ME	E40	Total
1	CNSCell_apoptosisTNFaThreshold	2(-)	3(-)	3(-)	2(-)	3(-)	2(-)	3(-)	2(-)	2(-)	1(-)	2(-)	25
2	CNSMacrophage_tnfaSecretedPerHourWhenStimulated	3(+)	2(+)	2(+)	1(+)	4(+)	3(+)	4(+)	3(+)	1(+)	2(+)	1(+)	26
3	Molecule_molecularHalflife	6(+)	4(+)	4(+)	3(+)	7(+)	4(+)	5(+)	4(+)	4(+)	3(+)	3(+)	47
4	CD4THelper_diff08	1(+)	6(-)	1(-)	7(+)	5(+)	17(-)	6(+)	17(-)	6(+)	8(+)	8	82
5	CNS_height	14(-)	9(-)	8(-)	4(-)	15(-)	7(-)	12(-)	7(-)	7(-)	4(-)	4(-)	91
7	TCell_proliferationMean	7(-)	1(+)	9(-)	6(+)	1(-)	1(+)	2(-)	1(+)	5(+)	23	54	110
7	APC_timeOfDeathMean	5(+)	7(-)	5(+)	14	2(+)	6(-)	1(+)	6(-)	48	9(+)	7(+)	110
8	CNS_width	17(-)	12(-)	10(-)	5(-)	17(-)	9(-)	15(-)	10(-)	17	5(-)	5(-)	122
9	Simulation_immunizationLinearDC0	4(+)	8(-)	6(+)	12(+)	6(+)	16(-)	7(+)	12(-)	37	10(+)	10	128
10	Simulation_numCNSMacrophage	18(+)	16(+)	16(+)	11(+)	23(+)	15(+)	17(+)	15(+)	21	6(+)	6(+)	164
11	TCell_AICDMean	8(+)	13(+)	7(+)	10(+)	10(+)	14(+)	9(+)	16(+)	63	7(+)	9	166
12	Simulation_immunizationLinearFreq	9(-)	20	13(-)	22	14(-)	11(+)	14(-)	11(+)	50	11	23	198
13	TCell_becomeEffectorMean	13(-)	5(-)	12(-)	13(-)	9(+)	5(-)	35	5(-)	35	56	20	208
14	Simulation_immunizationLinearGradient	10(+)	10(+)	17(+)	9(+)	11(+)	22	10(+)	40	24	31	42	226
15	TCell_cellsPerGridspace	12(+)	19	14(+)	18	12(+)	20	11(+)	21	16	42	46	231
16	TCell_specificityUpperLimit	26	27	55	38	27	13(-)	20(+)	13(-)	11(-)	43	23	296
17	Simulation_numDCCNS	33	28	30	21	24	24	26	24	19	64	31	324
18	APC_probabilityPhagocytosisToPeptide	15(+)	30	11(+)	8(+)	8(+)	68	8(+)	60	57	49	12	326
19	Circulation_timeToCrossOrgan	24	39	24	69	20(-)	10(+)	21(-)	8(+)	46	19	48	328
20	CLN_timeToCrossOrgan	60	35	32	36	25	28	23	22	26	25	18	330
21	Th2Polarization_proliferationMean	16(-)	14(+)	19	53	18(-)	53	16(-)	42	44	35	33	343
22	Simulation_numCNS	37	31	44	24	59	29	52	29	22	15	11	353
23	TCell_timeLocalActivationInducedEffectorFunctionFor	11(+)	24	15(+)	40	16(+)	21	18(+)	19	72	55	63	354
24	CNSMacrophage_type1RequiredForActivation	30	21	28	16	21(-)	64	22(-)	62	63	13	15	355
25	TCell_apoptosisNaiveMean	26	16(+)	44	54	59	8(+)	51	9(+)	66	18	13	364
26	CD4THelper_diff00	59	64	26	17	28	40	28	34	39	21	14	370
27	TCell_becomeEffectorStdDev	34	11(+)	46	52	19(-)	12(+)	67	14(+)	68	23	43	389
28	DendriticCell_phagocytosisProbabilityImmature	28	57	21	33	22(+)	46	19(+)	43	41	51	30	391
29	CNSMacrophage_phagocytosisProbabilityImmature	56	41	39	64	31	30	31	27	20	26	28	393
30	TCell_specificityLowerLimit	21	17(-)	22	55	35	18(-)	24	18(-)	70	58	64	402
31	TCell_AICDStdDev	28	46	39	15	38	58	39	68	25	28	24	408
32	APC_immatureDurationMean	19(-)	61	36	46	13(-)	37	13(-)	35	58	68	32	418
33	TCell_apoptosisPartialMaturityMean	32	50	40	30	48	44	36	37	18	20	65	420
34	Th1Polarization_mhcUnExpressionDelayMean	20	23	63	42	33	52	59	60	3(-)	14	55	424
35	CD8Treg_type1SecretedPerHourWhenActivated	31	38	34	19	34	43	42	53	29	62	49	434
36	DendriticCell_type1RequiredForActivation	22	63	27	23	72	61	45	56	41	16	16	442
37	CD4Treg_type1SecretedPerHourWhenActivated	35	52	52	33	29	42	32	49	72	25	26	447
38	TCell_apoptosisNaiveStdDev	46	47	33	39	70	41	54	44	13	49	21	457
39	DendriticCell_cytokineType2PolarizationRatio	29	62	18(-)	59	47	27	42	32	32	61	58	467
40	Spleen_width	68	55	57	49	53	26	38	23	23	42	35	469
41	Simulation_numCD4Th	43	23	49	56	49	34	51	33	65	39	36	478

Continued on Next Page...

Table 6.3 – Continued

Rank	Parameter Name	1M	1MT	2M	2MT	4M	4MT	8M	8MT	Th40	ME	E40	Total
42	CNSMacrophage_basalMBPExpressionProbability	53	33	41	43	41	57	33	55	49	39	38	482
43	Th2Polarization_type2SecretedPerHourWhenActivated	49	66	20	46	71	36	64	30	31	44	27	484
44	DendriticCell_type1SecretedPerHourImmunized	63	46	60	20	39	50	71	70	14	33	19	485
45	Simulation_numCD8Treg	66	43	60	42	43	51	27	52	9(–)	40	53	486
46	Th1Polarization_type1SecretedPerHourWhenActivated	43	49	31	60	69	31	47	38	10(+) 37	37	72	487
47	SLO_width	41	25	70	25	56	32	56	28	61	45	53	492
49	DendriticCell_phagocytosisProbabilityMature	70	56	67	29	26	69	25	65	34	17	35	493
49	TCell_timeLocalActivationDelay	45	65	69	50	65	23	37	25	55	30	29	493
50	APC_immatureDurationStdDev	40	18(+) 25	25	27	37	67	30	66	67	71	50	498
51	Th2Polarization_proliferationStdDev	57	60	54	31	54	25	69	31	39	35	47	502
52	SLO_height	47	29	62	48	45	63	42	61	42	32	40	511
53	APC_timeOfDeathStdDev	36	58	45	29	42	60	46	64	52	70	17	519
54	Th1Polarization_mhcUnExpressionDelayStdDev	23	54	50	26	46	65	43	51	28	66	68	520
55	CNS_timeToCrossOrgan	39	67	71	66	37	45	29	45	54	28	41	522
56	DendriticCellMigrates_lengthOfTimeMovingFollowingMigration	72	54	23	61	66	35	72	27	30	29	62	531
57	Simulation_numCD4Treg	65	49	56	34	32	39	57	40	64	53	53	542
58	TCell_apoptosisPartialMaturityStdDev	52	26	37	38	51	71	49	72	37	46	67	546
59	CLN_width	69	36	65	66	40	39	44	41	27	60	60	547
60	TCell_proliferationStdDev	51	42	72	69	30	49	34	47	34	65	60	553
61	Molecule_decayThreshold	67	40	35	70	62	19	64	20	61	72	44	554
62	Spleen_timeToCrossOrgan	58	34	51	72	44	33	62	36	56	47	69	562
63	SLO_timeToCrossOrgan	61	69	42	35	50	55	66	46	69	37	39	569
64	Circulation_width	64	32	68	58	61	55	53	58	8(+) 52	52	62	571
65	CNSMacrophage_phagocytosisProbabilityMature	38	68	60	62	67	72	68	71	46	12	26	590
66	Simulation_numDC	62	71	53	48	52	70	59	69	15	51	45	595
67	CD8Treg_cd8TregToCD4ThelperSpecificityDropOff	44	44	29	51	59	62	60	57	53	69	71	599
68	Circulation_height	54	37	47	72	55	48	49	50	52	67	70	601
69	CLN_height	48	52	66	67	65	59	62	63	12(+) 58	58	57	609
70	Spleen_height	51	72	48	64	65	47	70	48	43	59	56	623
71	Simulation_numDCSpleen	71	70	62	57	60	56	55	54	47	63	38	633
72	Simulation_immuneActivationLinearInitial	55	59	65	44	68	66	66	68	59	55	66	671

6.4.2.1 Neuronal apoptosis

The driving stimulus of both autoimmune and regulatory activity within the simulation is the apoptosis of neurons. Apoptotic neurons are phagocytosed by CNS-resident dendritic cells (DCs), which migrate to the cervical lymph node (CLN) to further prime populations of encephalitogenic CD4Th1 cells. The physiological turn over of CD4Th1 cells leads to their phagocytosis and consequently to the priming of CD4 and CD8Treg cells.

Consistent with this driver of simulation behaviour, the results of the global sensitivity analysis indicate that the most influential parameters in the simulation are those directly related to the death of neurons in the CNS, see table 6.3. The first and second most influential parameters in the simulation are *CNSCell_apoptosisTNFaThreshold* and *CNSMacrophage_tnfSecretedPerHourWhenStimulated*, which respectively dictate the threshold concentration of TNF- α which induces apoptosis in neurons and the quantity of TNF- α secreted per hour by activated macrophages in the CNS. *Molecule_molecularHalfLife* controls the rate of cytokine decay in the simulation. When set to high values, this parameter dictates that TNF- α will exist for longer, and diffuse further, hence increasing neuronal apoptosis. When these parameters are adjusted to minimise damage to neurons, both autoimmune and hence regulatory activity decrease in intensity and duration; peaks of T cell populations are reduced and occur sooner. Results concerning *Simulation_numCNSMacrophage*, which specifies the number of CNS macrophages in the CNS compartment, concur with this interpretation. Exhibiting a significant influence on all simulation responses, additional CNS macrophages lead to additional TNF- α secretion during autoimmune activity, which leads to additional neuronal apoptosis.

CNS_height and *CNS_width* are found to be more influential than any of the other parameters dictating compartment dimensions, which are relatively inconsequential to simulation behaviour. These two parameters dictate how densely populated the CNS compartment is with CD4Th cells, CNS macrophages and neurons. A more densely populated compartment will promote greater neuronal apoptosis, which promotes greater autoimmune and hence regulatory activity.

The *TCell_proliferationMean* parameter dictates the mean length of time required for a proliferating T cell to spawn a daughter cell. It is the most influential parameter in dictating the maximum population size of CD4Treg cells, second most influential with respect maximum number of CD8Treg cells, and the most influential regarding the times at which population sizes are reached for CD4Treg, CD8Treg and CD4Th1 cells. When this parameter is increased independently of *TCell_becomeEffectorMean*, which dictates the duration of a T cell's proliferative activity before differentiating into an effector cell, the maximum population sizes for these T cells is reduced. When set to higher values, this parameter dictates that fewer effector CD4Th1 cells enter the CNS at earlier stages, hence taking longer to cause neuronal apoptosis that results in a strong stimulus for continued autoimmune activity. Likewise, fewer effector CD4Th1 cells in the system at early stages of disease progression results in less CD4Th1 phagocytosis by DCs that would eventually prime regulatory T cell populations. Increasing the length of time required for proliferation has the effect of prolonging autoimmune and regulatory phases of disease.

The parameters discussed thus far have a direct influence on the rate of neuronal apoptosis in the CNS compartment. *Th1Polarization_type1SecretedPerHourWhenActivated* and *CNSMacrophage_type1RequiredForActivation* represent 'second order' parameters in this respect, dictating the quantity of type 1 cytokine secreted in the CNS compartment, and the concentration of this cytokine required to induce activation and

TNF- α secretion in CNS macrophages. *CNSMacrophage_type1RequiredForActivation* is found to be significant (p value < 0.01) only to the *CD4Treg max* and *CD8Treg max* responses, whereas *Th1Polarization_type1SecretedPerHourWhenActivated* is found to significantly influence only the *CD4Th1 at 40 days* response. These results are somewhat surprising, and suggest a great redundancy in the ability of the CD4Th1 population to activate CNS macrophages, likely owing to the very large number of these cells that infiltrate the CNS compartment.

Simulation_numCNS controls the number of neurons in the CNS compartment, and hence has a direct relevance to the number of neurons that can be apoptosed through autoimmune activity. It is surprising that this parameter has no significant correlation with any response, and invites the following interpretation. Dendritic cells in the simulation are not motile until they mature, at which point their phagocytic capacity is reduced from 100% to 30%, and they migrate to the CLN following a relatively straight path. Activated CNS macrophages also become motile upon activation, and share a 30% phagocytic capacity. However their movement is random in all directions. Coupled with the fact that CNS macrophages reside within the CNS compartment for the entire duration of their mature lifecycle, and that there are 75 macrophages versus 40 dendritic cells within the CNS, it is highly likely that macrophages phagocytose significantly more apoptotic neurons than dendritic cells do. Whereas additional peptide-presenting DCs can prime vast numbers of T cells, the additional MBP-presenting activated CNS macrophages likely do not contribute significantly to the immune response, other than to provide marginally quicker local activation to CD4Th infiltrates. Hence, a 10% perturbation of *Simulation_numCNS* does not necessarily translate to significant deviations in simulation behaviour.

Related to this parameter, *Simulation_numDCCNS*, which dictates the number of dendritic cells in the CNS, is not found to have any significant influence over any of the responses. Under default parameter values, the maximum number of peptide-presenting immunogenic DCs in the CLN compartment (at peak immune activity) is around 30, despite there being around 92 CNS-originating DCs present in the compartment at any point in time. The majority of DCs in the CLN do not present peptides. At $\pm 10\%$ of *Simulation_numDCCNS*'s default value, this ratio translates into 27 and 33 peptide-presenting DCs in the CLN at peak immune activity. Not all DCs present both CD4Th-derived and MBP peptides together, hence the change in number of peptide presenting DCs from the default value of 30 is shared between Treg and CD4Th populations. The lack of significant influence by *Simulation_numDCCNS* over simulation responses suggests that these differences in peptide presentation by the DC population do not translate to significant differences in T cell population dynamics or neuronal apoptosis.

6.4.2.2 Immunisation mechanism

The simulation's immunisation mechanism, whereby EAE is induced in the simulation, is parameterised through four parameters: *Simulation_immuneLinearInitial*, *Simulation_immuneLinearDC0*, *Simulation_immuneLinearGradient* and *Simulation_immuneLinearFreq*. Upon immunisation, an initial number of immunized MBP-presenting DCs are placed within the SLO compartment at time zero. This number is specified by *Simulation_immuneLinearInitial*. Thereafter, immunized MBP-presenting DCs are added to the SLO compartment periodically, as defined by *Simulation_immuneLinearFreq*. The number inserted reduces linearly over time, described by the parameters *Simulation_immuneLinearDC0* and *Simulation_immuneLinearGradient* (as described in section 5.2.4).

Simulation_immunizationLinearDC0, *Simulation_immunizationLinearGradient* and *Simulation_immunizationLinearFreq* are found to exhibit statistically significant influence over simulation responses. Increasing either *Simulation_immunizationLinearDC0* or *Simulation_immunizationLinearGradient* (which is, by default, set to a negative value) results in more immunising DCs being inserted into the SLO. This can be considered to reflect increasing the doses of MBP, CFA and PTx used to immunize mice *in vivo*. The result is that the maximum number of CD4Th1 cells reached increases, as does the maximum number of CD4Th2 cells. This in turn provides additional stimulation to the regulatory immune response, with maximum number of CD4Treg and CD8Treg cells increasing. *Simulation_immunizationLinearDC0* has the effect of increasing the onset of immunity for CD4Th1, CD4Treg and CD8Treg cells. This is not the case for *Simulation_immunizationLinearGradient* which delays the peak of immune activity, and has no significant influence over this response for CD4Treg and CD8Treg cells. This difference may be explained if the same proportional increase in *Simulation_immunizationLinearGradient* results in more immunisation DCs being introduced later during the simulation than the same increase in *Simulation_immunizationLinearDC0*.

Simulation_immunizationLinearFreq has a pattern of response almost exactly opposite to that of *Simulation_immunizationLinearDC0* and *Simulation_immunizationLinearGradient*. Increasing this parameter has the effect of reducing the number of immunisation DCs administered to the simulation, but maintaining the length of time over which this may occur. Reduced stimulation for CD4Th priming results in reduced stimulation for CD4Treg and CD8Treg populations, resulting in reduced maximum population size, and reduced intensity of response onset.

It is of note that *Simulation_immunizationLinearInitial* is the least influential parameter in the simulation.

6.4.2.3 Regulatory capacity

A surprising result is the apparent lack of influence exhibited by the parameters annotating regulatory action on simulation responses. *CD8Treg_cd8TregToCD4Thelper-SpecificityDropOff* is ranked as the sixth least significant parameter in the simulation. This parameter dictates the probability that a binding of an effector CD8Treg to an effector CD4Th1 will result in the induction of apoptosis in the CD4Th1 cell. As such this parameter can dictate how effective the regulatory network is in the system. A related parameter is *Th1Polarization_mhcUnExpressionDelayMean*, which dictates the mean length of time that Qa-1 is expressed by CD4Th1 cells following their differentiation into effector cells. CD4Th1 cells are only susceptible to regulatory action whilst expressing Qa-1. No statistical significance is found for *CD8Treg_cd8TregToCD4ThelperSpecificityDropOff* with respect to any response, and *Th1Polarization_mhcUnExpressionDelayMean* is significant to only one; it is the third most influential parameter in relation to the *CD4Th1 at 40d* response. These results motivate further investigation into the robust nature of the regulatory network, conducted in the following chapter.

6.4.2.4 Type 1 / type 2 balance

The balance between type 1 and type 2 immune response is largely dictated through cytokine profiles, both in the CNS where DCs are activated and adopt a polarisation, and in the area surrounding immunogenic DCs which directly influence the polarisation that primed CD4Th cells adopt. Type 1 cytokine secretion by CD4Treg and CD8Treg cells exhibits no significant influence on any response. Given that these cells do not

enter the CNS compartment, and that they are not likely to be in the vicinity of CLN-resident DCs that prime CD4Th cells by the time that they are locally activated for cytokine secretion, these results are to be expected.

However, cytokine secretion by CD4Th1 and CD4Th2 cells, which do enter the CNS, are also found to be largely insignificant to simulation behaviour. *Th1Polarization_type1SecretedPerHourWhenActivated*, which dictates the rate of type 1 cytokine secretion by CD4Th1 cells is significant with respect to only the *CD4Th1 at 40 days* response. *Th2Polarization_type2SecretedPerHourWhenActivated*, the rate of type 2 cytokine secretion by CD4Th2 cells does not exhibit a statistically significant influence on any response. *DendriticCell_cytokineType2PolarizationRatio*, which dictates the ratio of type 2 cytokine to type 1 cytokine required for a DC to adopt a type 2 polarisation, exhibits statistically significant influence on the *CD4Th2 max* response only. With respect to this response the parameter is ranked as the 18th most significant parameter.

These results collectively suggest that simulation behaviour does not hinge on a delicate balance of cytokine profiles, rather that cytokines secretion rates and half lives dictate that in any scenario where cytokine balance affects the polarisation adopted by a cell, one type of cytokine exists in significantly greater quantity than the other.

This conclusion is further supported by results for parameters *CNSMacrophage-type1RequiredForActivation* and *DendriticCell_type1RequiredForActivation*, which respectively dictate the concentration of type 1 cytokine required to activate CNS macrophages and induce the upregulation of co-stimulatory molecules in DCs. *CNSMacrophage-type1RequiredForActivation* is shown significant with respect to only two responses, in which it is ranked 21st and 22nd. *DendriticCell_type1RequiredForActivation* is not significant with respect to any responses. These results indicate that during an immune response, APCs come into close proximity with sufficiently many type 1 secreting cells such that they perceive type 1 cytokines in significantly higher concentrations than the range dictated by $\pm 10\%$ of their default activation thresholds.

6.4.2.5 Peptide derivation and presentation by APCs, and T cell priming

Various parameters pertaining to peptide phagocytosis and presentation, and priming of T cells are now considered. Peptides are derived from apoptotic cells, as phagocytosed by DCs and CNS macrophages. These cells are more phagocytic when immature than mature. Parameters specifying the probability that contact between an APC and an apoptotic cell will result in the phagocytosis of the latter by the former are *DendriticCell_phagocytosisProbabilityImmature*, *DendriticCell_phagocytosisProbabilityMature*, *CNSMacrophage_phagocytosisProbabilityImmature* and *CNSMacrophage_phagocytosisProbabilityMature*. Only *DendriticCell_phagocytosisProbabilityImmature* exhibits a statistically significant influence on simulation responses, and does so on only *CD4Treg max*, where it is ranked as the 22nd most influential parameter, and *CD8Treg max*, where it is ranked 19th. These results indicate a robustness in the stimulation's operation; an apoptotic cell comes into contact with sufficient phagocytic APCs such that perturbing the phagocytic capacity of APCs by 10% does not significantly alter the simulation's behaviour. Furthermore, for a phagocytosis event to influence simulation behaviour, the phagocytosing APC must derive and present peptides for the priming of T cell populations. However, the probability of a phagocytosis event leading to peptide derivation is only 2%. As discussed next, the parameter dictating this probability is considerably more influential on simulation behaviour.

The probability that a phagocytosis event will lead to the derivation of peptides that are presented on MHC is identical for both DCs and CNS macrophages, and is

specified through *APC_probabilityPhagocytosisToPeptide*. This parameter is found to be significant with respect the maximum number of CD4Th1, CD4Th2, CD4Treg and CD8Treg cells attained through simulation execution, and the time at which the CD4Th2 population peaks. As discussed in the following chapter, around a third of mature APCs present peptides. Coupled with the fact that the default value for *APC_probabilityPhagocytosisToPeptide* is only 2%, this suggests that a very high proportion of APCs phagocytose at least one cell. Adjusting this parameter by 10% is hence highly influential on simulation behaviour; a single immunogenic APC can produce a large number of T cells. It is proposed that it is the considerable redundancy in the regulatory network which prevents this parameter also significantly influencing EAE-severity related parameters, and the number of CD4Th1 cells at 40 days.

APC_immatureDurationMean specifies the duration of time that an APC remains in an immature state for. It is found to be relatively insignificant, influencing *CD4Th1 max*, *CD4Treg max* and *CD8Treg max* responses, but ranked below the top ten in all cases. Increasing this parameter is found to decrease the maximum number of T cells attained during simulation execution. The overall lifespan of an APC is dictated by the length of time that it remains immature, *APC_immatureDurationMean*, and the length of time that it remains in a mature state before entering apoptosis, *APC_timeOfDeathMean*. Excluding DCs arising from the simulation's immunisation mechanism, the number of APCs, both DCs and CNS macrophages, in the simulation at any point in time is static; an APC entering apoptosis is immediately replaced by another immature APC. Hence, increasing *APC_immatureDurationMean* will generally decrease the total number of APCs presenting peptides at any point in time, and hence reduce the peak number of T cells generated during immune responses. This result suggests that increasing the number of mature APCs in the simulation at any point in time has a greater potential to prime larger T cell populations than increasing the time these APCs spend in a highly phagocytic state, and hence the probability that they will derive peptides through which T cell priming may be accomplished.

A parameter that exhibits significant influence over many of the simulation's responses is *APC_timeOfDeathMean*. It dictates the mean period of time that APCs remain in a mature state before entering apoptosis. As such, this parameter determines how long proliferative T cells can derive stimulation required to produce daughter cells for. Since daughter cells typically prime on the same APCs as their parent cells, there is an exponential relationship between an APC's mature lifespan and the size of the T cell population that it induces.

TCell_specificityLowerLimit and *TCell_specificityUpperLimit* dictate the boundaries of the range from which T cell specificities may be selected. T cells with low specificity are less likely to derive effective TCR signaling from contact with MHC:peptide complexes for which they are specific. It is interesting to note that *TCell_specificityUpperLimit* is more influential, being higher ranked in nearly all responses, than *TCell_specificityLowerLimit*.

6.4.2.6 T cell dynamics

CD4THelper_diff08 and *CD4THelper_diff00* specify the probability of a CD4Th cell adopting a type 1 polarisation, as opposed to type 2, when the proportion of type 1 cytokine comprising the local cytokine milieu is above or below 80% respectively. The former parameter is shown to be highly influential in the simulation, exhibiting statistically significant effects in all but one response, *EAE at 40 days*. In contrast, *CD4THelper_diff00* exhibits no statistically significant influence over any response. Figure 6.18 shows the system-wide T cell dynamics when *CD4THelper_diff00* is set to its

baseline value of 5%, and 100%, wherein DCs that would ordinarily strongly favour CD4Th differentiation into CD4Th2 cells now exclusively induce CD4Th1 cells. There is very little difference in the CD4Th1 dynamics that result from this large parametric change, suggesting that the recovery from autoimmunity in the simulation is significantly driven by regulatory action, wherein the CD4Th1 population is apoptosed by CD8Tregs. In contrast type 2 deviation, where additional CD4Th2 cells in the CNS compartment induce type 2 polarising DCs in place of type 1 polarising, is a result of the recovery process, rather than a driver of it.




Figure 6.18: Comparison of T cell dynamics when *CD4THelper_diff00* is set to its default value of 5%, and 100%.

TCell_AICDMean exerts significant influence over nearly all of the simulation’s responses. It specifies the mean lifespan of effector T cells, prior to their entering apoptosis. Increasing this parameter increases the maximum population sizes of all T cells, and the times at which they occur. Manipulating effector T cell lifespan will influence the maximum number of CD4Th1 cells that reside within the CNS during autoimmunity. As such, additional type 1 cytokine is secreted, which in turn increases stimulation for CNS macrophage activation and the secretion of neuron-harming TNF- α . In this manner, this parameter exerts significant influence over the *max EAE* response also.

6.5 Robustness analysis of parameters

The ARTIMMUS simulation’s parameter values are derived through a variety of methods. Some are based on domain-specific knowledge, others are arbitrarily chosen, and some are tuned in order to deliver simulation behaviour that reflects that observed *in vivo*. This section describes the derivation and application of a robustness analysis that ascertains the range of values that each parameter may adopt before a significant change in simulation behaviour occurs.

The analysis indicates where simulation behaviours critically depend on certain parameters holding values in a very small range of their possible domains. Parameter values may be known with some certainty, if they have been derived directly from domain-specific knowledge, or have undergone calibration to ensure that simulation behaviour reflects that observed *in vivo*. If, however, simulation behaviour is found to critically depend on parameters that have been assigned arbitrary values, as may be the case where domain knowledge cannot specify an appropriate value, then simulation-based results may differ if another arbitrarily determined value were to be assigned.

This information is important to consider when interpreting simulation-derived results and qualifying their significance.

The following section details the experimental procedure of the robustness analysis. Thereafter it is employed in ascertaining ARTIMMUS's robustness with respect to parameters that were assigned arbitrary values, and the implications of these results are considered.

6.5.1 Experimental procedure

The robustness analysis technique devised here is a *one-at-a-time* technique, in that each parameter is perturbed independently of the others, which remain at their baseline values. The technique employs two indications of when a significant change in simulation behaviour has occurred, discussed below. Each parameter is perturbed away from its baseline value, and the boundary at which a significant deviation in simulation behaviour, measured through the responses identified in section 6.2, is recorded. In this manner, the fragility of the simulation with respect to parameter perturbation may be ascertained; for each parameter, the range of values either side of the default value that do not cause a significant change in simulation behaviour are identified. A complete robustness analysis of the *TCell_AICDMean* parameter is depicted in figure 6.19 to illustrate the experimental procedure described in this section.

To facilitate comparison of simulation fragility with respect to various parameters, a *robustness index* measure has been devised. The boundaries on either side of the baseline value at which a significant deviation in simulation behaviour occurs are identified. Only the nearest boundaries to the baseline value are considered, and there are at most two such boundaries, termed the *upper* and *lower* boundaries. The distance between each of these boundaries and the baseline parameter value is expressed as a percentage of the parameter's baseline value, comprising the *upper* and *lower* indices. If no significant deviation in simulation behaviour is observed then the boundary and corresponding percentage index is assigned the value *not-a-number* (NaN). Where a significant deviation in simulation behaviour does occur, but the baseline parameter value is 0, the corresponding percentage index is assigned the value infinity (Inf). Values for upper and lower boundaries that are not NaN are accompanied by either (+) or (-), which respectively indicate a significant increase or decrease in response value. For clarity, NaN scores are indicated by ‘.’ in the tables of this thesis.

The range of values over which a parameter is perturbed is determined by its type, and what is considered reasonable. Some parameters represent probabilities, or proportions. Such parameters may take values in the range of 0.0 to 1.0. The full range of values is typically explored, with the notable exception of *TCell_specificityUpperLimit* and *TCell_specificityLowerLimit*, where the latter may not exceed the value held by the former, and vice versa. Many pairs of simulation parameters, which are real numbers, represent normal distributions from which the periods of time that individual cells remain in particular states are drawn. Such distributions are defined through a parameter that describes the mean, and another that describes the standard deviation of the distribution. The parameters describing the standard deviations are perturbed to values ranging from 0 to the value held by the mean⁴. The parameters describing the mean are perturbed to values considered reasonable, or until a significant deviation in simulation behaviour is observed. The robustness analysis employs two means of determining when a change in simulation response behaviour may be considered significant.

⁴Recall from section 5.2.8 that for parameters describing standard deviations, these normal distributions actually describe two times the standard deviation, covering 95% of the distribution.

The first makes use of the Vargha-Delaney A test [Vargha & Delaney 2000], introduced above in section 6.3, to indicate when a perturbation has resulted in a scientifically significant change in simulation behaviour. The use of a non-parametric test avoids the assumption that the underlying distributions from which samples are drawn are normally distributed. An effect magnitude test is employed to determine scientific significance in place of statistical significance since a sufficiently large sample size always reveals statistically significance, unless the variable of interest has no effect.

The A test is used to indicate significant deviations in simulation behaviour for 11 responses, outlined in section 6.2.1. These are: the nine responses derived from T cell population dynamics, the maximum EAE score experienced during simulation execution, and the EAE score at 40 days. For each parameter perturbation 500 simulation executions are conducted, and the 11 responses are calculated for each individual execution. For each response in turn, the 500 samples derived from a particular parameter perturbation are compared using the A test to another 500 samples derived from simulation executions using baseline parameter values. A test scores above 0.71 or below 0.29 indicate a “large” difference between two distributions [Vargha & Delaney 2000]. These values are assumed to indicate when scientifically significant changes in simulation behaviour have resulted from parameter perturbation.


The other means of identifying significant changes in simulation behaviour is built around assigning EAE severity scores to simulation executions, and considers both the maximum EAE severity experienced and the EAE severity remaining at 40 days. Whilst using the A test to indicate significant deviations in simulation behaviour is statistically grounded, it offers no indication of the biological significance of the change. The number of T cells in the simulation is a vast reduction of the number in an experimental animal, and there is no direct way of relating changes in terms of simulation cell number to the physical effects that are observed *in vivo*. Interpretation of *in vivo* experimentation is underpinned by the assignment of EAE severity scores to individual mice, from which group averages are compiled. The simulation’s EAE severity scoring mechanism, detailed in section 6.2.1 above, permits a similar approach to interpreting simulation derived results. The domain expert has indicated that a change of ± 1.0 in a group’s mean EAE score may be considered significant. Averages are compiled from 500 simulation runs, in accordance to the results of the consistency analysis in section 6.3.

Though the distribution of EAE scores is non-normal, as scores below 0 or above 5 are not possible, mean averages are used in preference to medians in order to align the present analysis with *in vivo* practice, and to facilitate greater fidelity when contrasting group scores.

6.5.2 Results

This section details the application of the robustness analysis to ARTIMMUS simulation parameters, with a focus on those that are arbitrarily assigned their baseline values.

Simulation parameter values are determined through either domain-specific knowledge, iterative calibration to deliver simulation behaviours that are consistent with those observed *in vivo*, or are arbitrarily assigned since no guidance on appropriate values exists. Parameters assigned through iterative calibration are built upon those that are arbitrarily assigned. Finding that simulation behaviours critically depend on these parameters holding values from a very small range has implications for the interpretation of simulation results; were different arbitrary values to be chosen for such parameters, simulation behaviours, and hence results, might differ significantly. This



(c) Summary of robustness indices, lower and upper boundaries and indices for all responses. RI, robustness index; LI, lower index; UI, upper index; LB, lower boundary; UB, upper boundary. For clarity, NaN is indicated by ‘.’.

Figure 6.19: Robustness analysis of the *TCell_AICDMean* parameter.

does not necessarily undermine activities to explore a system through simulation, since it cannot be ruled out that other simulation parameters cannot be re-calibrated to re-align simulation behaviour with that observed *in vivo*. It is, however, important to communicate this simulation criticality when presenting the results of simulation-based experimentation.

At the time of writing the author knows of no reason to question the values that the simulation’s parameters have been assigned. However, were reasons to question simulation parameter values to arise in the future, this analysis can indicate where implications on simulation results exist. The results of this analysis can indicate the implications on simulation results were reasons to doubt simulation parameters to arise in the future. Although the focus of this section is on arbitrarily determined parameter values, a summary of robustness indices for all simulation parameters with respect to all simulation responses is presented in table 6.4, and appendix section C.3 provides a full robustness analysis, including robustness boundaries and indices, with respect to each response.

It may be seen from table 6.4 that some simulation responses are highly fragile with respect to parametric perturbation, with several examples of perturbations of less than 2% being sufficient to result in significant deviation in simulation behaviours. There are many parameters for which perturbations cause significant deviation in behaviours for all simulation responses.

The findings of this section reveal that, in comparison, the simulation is relatively robust with respect to perturbation of its arbitrarily assigned parameters. Most are able to affect significant deviation in simulation behaviour for less than half the responses,

but require perturbations of around 40% or more to do so. Noteworthy exceptions are *CNS_height*, *CNS_width*, and *CNSMacrophage_tnfaSecretedPerHourWhenStimulated*. Small perturbations of less than 10% result in significant behavioural deviations in nearly all responses. These results point to the criticality of the rate of neuronal apoptosis in the system, to which all three parameters directly contribute.

Table 6.4: Summary of parameter robustness indexes, ordered by total rank. Responses are indicated as follows: 1M, *CD4Th1 Max*; 1MT, *CD4Th1 Max Time*; 2M, *CD4Th2 Max*; 2MT, *CD4Th2 Max Time*; 4M, *CD4Treg Max*; 4MT, *CD4Treg Max Time*; 8M, *CD8Treg Max*; 8MT, *CD8Treg Max Time*; Th40, *CD4Th1 at 40 Days*; MEA, *Max EAE*; E40A, *EAE at 40 Days*. Significant deviation is indicated through the *A* test. ME and E40 represent *Max EAE* and *EAE at 40 Days*, with significant deviations in response behaviour defined as a change of at least 1.0 in the mean EAE score. Not-a-number values, representing no significant deviation in behaviour, are marked with a period for clarity. Response columns show the smallest of the two robustness indexes for each parameter-response combination. The ‘total’ is the sum of ranks for each parameter across all responses, with small response indexes being ranked highest.

Rank	Parameter Name	1M	1MT	2M	2MT	4M	4MT	8M	8MT	Th40	MEA	E40A	ME	E40	Total
1	CNSMacrophage.tnfaSecretedPerHourWhenStimulated	5.64	11.97	2.99	5.74	4.84	9.01	6.67	9.22	16.26	1.55	1.86	1.41	0.89	28
2	CNSCell_apoptosisTNFaThreshold	5.26	13.78	3.27	5.20	5.27	11.88	6.63	13.03	14.70	1.45	2.04	1.36	0.97	30
3	CNS.height	6.53	14.79	4.02	6.72	5.82	10.99	7.40	12.67	17.45	3.66	3.83	3.31	1.82	56
4	Molecule_molecularHalflife	7.44	16.73	4.17	7.76	6.79	12.80	8.73	13.93	25.29	2.22	2.58	2.02	1.23	64
5	Simulation_numCNSMacrophage	12.12	23.71	6.79	13.13	13.18	18.65	13.97	19.00	42.02	3.41	3.60	3.08	1.71	106
6	APC.timeOfDeathMean	7.30	26.84	7.62	35.80	4.27	42.36	5.79	47.02	26.88	11.79	24.68	11.87	11.76	115
7	Simulation_immuneActivationLinearDC0	6.67	16.71	9.02	43.58	8.72	22.22	12.88	21.38	43.09	13.94	24.95	15.90	14.13	124
8	CNS.width	9.27	23.01	4.66	7.02	7.20	16.40	8.12	17.10	.	3.90	3.93	3.25	1.87	150
9	Simulation_immuneActivationLinearGradient	16.92	22.76	15.16	44.10	12.95	36.71	16.24	39.82	54.15	33.83	50.24	34.83	37.75	184
10	Simulation_numCNS	43.41	56.83	27.59	38.11	34.59	45.09	36.73	46.67	.	6.92	11.51	7.55	5.41	253
11	Simulation_numDC CNS	39.70	51.70	23.84	39.09	16.67	63.58	22.29	57.64	112.10	65.01	138.00	71.18	55.11	262
12	Th1Polarization_type1SecretedPerHourWhenActivated	38.04	62.49	38.01	44.88	35.32	53.98	39.64	56.07	477.10	26.07	141.30	27.92	48.21	263
13	TCell_proliferationMean	9.28	9.18	9.81	.	3.54	5.58	5.88	5.91	15.59	18.24	.	19.17	.	270
14	TCell_cellsPerGridspace	8.98	47.96	10.41	33.17	8.56	39.20	9.61	42.67	37.56	22.03	.	24.78	.	272
15	TCell_becomeEffectorMean	18.22	24.94	20.64	47.50	9.53	14.98	17.27	14.94	29.03	28.88	.	29.29	.	286
16	Simulation_immuneActivationLinearFreq	10.31	.	16.13	36.11	26.81	33.46	35.57	37.20	.	23.16	33.20	20.72	16.67	290
17	TCell_AICDMean	8.81	.	8.54	24.62	12.10	.	13.92	.	99.74	9.80	17.71	10.61	10.26	322
18	Simulation_numCD4Th	48.10	41.43	54.33	74.73	45.91	42.25	57.20	41.97	79.81	62.24	.	64.00	243.80	330
19	TCell_timeLocalActivationInducedEffectorFunctionFor	14.59	65.05	19.02	49.56	17.56	68.72	21.95	53.73	.	22.52	.	23.26	31.35	336
20	CNSMacrophage_type1RequiredForActivation	58.00	127.50	50.32	75.85	71.60	108.30	80.11	106.30	96.71	34.70	58.82	36.37	35.39	341
21	Circulation.height	61.06	46.95	36.99	182.60	32.55	43.99	39.21	45.58	29.27	46.44	.	51.30	.	361
22	CD4THelper_diff08	5.66	47.23	2.69	92.33	8.84	90.82	10.18	92.70	.	16.58	.	16.93	.	365
23	APC_probabilityPhagocytosisToPeptide	32.53	65.77	16.92	25.85	9.78	60.98	12.01	74.93	49.54	441
24	SLO_width	67.65	44.32	56.01	162.80	48.82	44.10	137.80	44.75	.	119.50	.	130.90	.	485
25	DendriticCell_type1SecretedPerHourImmunized	88.89	.	60.43	1834.00	84.73	91.87	86.42	91.50	.	93.41	889.50	93.30	604.80	487
26	Simulation_numCD8Treg	94.96	76.89	.	37.63	73.86	24.09	73.60	53.27	.	99.01	.	97.46	498	
27	DendriticCell_phagocytosisProbabilityImmature	76.65	91.47	53.31	70.32	54.63	92.77	61.30	92.95	.	92.48	.	93.67	.	517
28	TCell_timeLocalActivationDelay	64.22	75.41	58.56	325.50	88.58	57.11	81.86	59.87	.	93.46	.	96.19	.	518
29	Simulation_numCD4Treg	.	83.07	.	.	30.89	47.23	55.04	48.05	91.56	.	98.98	.	97.94	530
30	TCell_apoptosisNaiveMean	38.20	53.26	49.23	.	50.61	54.93	.	63.73	.	65.05	.	68.47	.	555
31	CNSMacrophage_basalMBPExpressionProbability	76.59	95.18	78.32	96.19	82.58	.	85.40	.	.	83.73	.	83.49	285.70	565
32	CNSMacrophage_phagocytosisProbabilityMature	.	.	78.38	90.65	83.66	24.24	56.55	28.15	25.25	588
33	Circulation.timeToCrossOrgan	82.13	.	59.66	.	44.98	78.82	57.61	76.69	.	69.59	.	70.98	.	591
34	DendriticCell_cytokineType2PolarizationRatio	81.25	83.80	49.16	.	77.75	82.52	79.53	82.23	484.50	606
35	Molecule_decayThreshold	752.70	.	235.20	.	664.70	3561.00	760.10	3721.00	87.05	1317.00	.	1283.00	.	628
36	CLN_width	198.40	57.31	51.53	.	48.45	39.93	59.73	665

Continued on Next Page...

Table 6.4 – Continued

Table 6.5: Robustness indexes for parameters pertaining to compartmental dimensions. Responses are indicated as follows: 1M, *CD4Th1 Max*; 1MT, *CD4Th1 Max Time*; 2M, *CD4Th2 Max*; 2MT, *CD4Th2 Max Time*; 4M, *CD4Treg Max*; 4MT, *CD4Treg Max Time*; 8M, *CD8Treg Max*; 8MT, *CD8Treg Max Time*; Th40, *CD4Th1 at 40 Days*; MEA, *Max EAE*; E40A, *EAE at 40 Days*. Significant deviation is indicated through the *A* test. ME and E40 represent *Max EAE* and *EAE at 40 Days*, with significant deviations in response behaviour defined as a change of at least 1.0 in the mean EAE score.

Parameter Name	1M	1MT	2M	2MT	4M	4MT	8M	8MT	Th40	MEA	E40A	ME	E40
<i>CLN_height</i>	44.60	59.28	122.30	59.79	45.61
<i>CLN_width</i>	198.40	57.31	51.53	.	44.12	.	53.26	48.73
<i>CNS_height</i>	6.53	14.79	4.02	6.72	5.81	12.57	7.44	11.92	17.45	3.66	3.83	3.31	1.82
<i>CNS_width</i>	9.27	23.01	4.66	7.02	7.25	17.88	8.35	20.97	.	3.90	3.93	3.25	1.87
<i>Circulation_height</i>	61.06	46.95	36.99	182.60	34.30	43.86	39.66	45.19	29.27	46.44	.	51.30	.
<i>Circulation_width</i>	126.50
<i>SLO_height</i>	82.59	131.00	.	160.60	140.00	.	141.20	.
<i>SLO_width</i>	67.65	44.32	56.01	162.80	47.34	45.93	141.30	45.71	.	119.50	.	130.90	.
<i>Spleen_height</i>	.	.	160.00	.	160.50	57.50	165.10	58.37	140.70
<i>Spleen_width</i>	50.39	35.94	73.76	34.69

Table 6.6: Robustness indexes for parameters specifying initial cell numbers. Responses indicated as in table 6.5.

Parameter Name	1M	1MT	2M	2MT	4M	4MT	8M	8MT	Th40	MEA	E40A	ME	E40
<i>Simulation_numCD4Th</i>	48.10	41.43	54.33	74.73	47.15	43.88	58.67	43.05	79.81	62.24	.	64.00	243.80
<i>Simulation_numCD8Treg</i>	94.96	76.89	.	.	39.77	75.10	23.54	78.07	53.27	.	99.01	.	97.46
<i>Simulation_numCD4Treg</i>	.	83.07	.	.	30.90	48.09	56.81	47.05	91.56	.	98.98	.	97.94
<i>Simulation_numDCSpleen</i>	51.47	.	51.16	.	82.67
<i>Simulation_numDC</i>	335.60	.	456.60

Table 6.7: Robustness indexes for parameters pertaining to T cell-APC interactions. Responses indicated as in table 6.5

Parameter Name	1M	1MT	2M	2MT	4M	4MT	8M	8MT	Th40	MEA	E40A	ME	E40
<i>CNSMacrophage_phagocytosisProbabilityImmature</i>	87.70	46.66	.	52.72	.
<i>CNSMacrophage_phagocytosisProbabilityMature</i>	.	.	78.38	90.65	83.48	24.24	56.55	28.15	25.25
<i>DendriticCell_phagocytosisProbabilityImmature</i>	76.65	91.47	53.31	70.32	53.52	93.54	60.78	93.36	.	92.48	.	93.67	.
<i>DendriticCell_phagocytosisProbabilityMature</i>	86.09
<i>TCell_specificityLowerLimit</i>	.	86.34	.	.	93.33	60.70	74.95	62.75
<i>TCell_specificityUpperLimit</i>	42.14	.	43.10

Table 6.8: Robustness indexes for cytokine secretion and decay. Responses indicated as in table 6.5.

Parameter Name	1M	1MT	2M	2MT	4M	4MT	8M	8MT	Th40	MEA	E40A	ME	E40
CNSMacrophage.tnfaSecretedPerHourWhenStimulated	5.64	11.97	2.99	5.74	4.79	9.29	6.71	9.68	16.26	1.55	1.86	1.41	0.89
DendriticCell.type1SecretedPerHourImmunized	88.89	.	60.43	1834.00	85.23	.	87.52	.	.	93.41	889.50	93.30	604.80
CD4Treg_type1SecretedPerHourWhenActivated
CD8Treg_type1SecretedPerHourWhenActivated
Th1Polarization_type1SecretedPerHourWhenActivated	38.04	62.49	38.01	44.88	35.75	57.56	40.50	55.19	477.10	26.07	141.30	27.92	48.21
Th2Polarization_type2SecretedPerHourWhenActivated	.	.	87.03	94.89	384.30	587.30	455.00	540.30	97.59
Molecule_decayThreshold	752.70	.	235.20	.	673.80	.	808.60	.	87.05	1317.00	.	1283.00	.

Table 6.9: Robustness indexes for parameters specifying standard deviations of timing distributions. Responses indicated as in table 6.5.

6.5.2.1 Compartmental dimensions

Table 6.5 shows the robustness indices for parameters specifying compartmental dimensions. Altering the heights and widths of compartments, independently of any other parameter, adjusts the density of cells in the simulation.

In the case of the CNS compartment this has a marked effect on many of the simulation's responses. The importance of the rate of neuronal apoptosis on simulation behaviour has been established above in section 6.4.2.1: apoptotic neurons incite self-perpetuating autoimmunity which in turn stimulates the expansion of regulatory T cells. Altering *CNS_height* and *CNS_width* adjusts the density of cells in the CNS compartment. When these parameters are decreased, the same quantity of TNF- α , secreted by the same number of CNS macrophages, reaches more neurons. There exists less physical space over which to dissipate, increasing its concentration. As such the rate at which neurons enter apoptosis increases. The simulation is highly sensitive to perturbation of these two parameters, with perturbations of less than 5% resulting in significant behavioural changes in numerous responses.

The other compartmental dimension parameters generally require perturbations of at least a third of the default value in order to attain significant behavioural changes, and in many cases no significant change is affected.

A common trend in compartments where T cell priming takes place, the CLN, SLO and spleen, is for perturbation of a compartment's width to have a greater effect on a simulation response than the same perturbation of its height. This is believed to be a simulation artifact, resulting from the manner in which cells move within the compartment. Cells migrate through a compartment from top to bottom. Hence, a migratory T cell will always explore the full height of a compartment. The same is not true for a compartment's width, where the probability of a T cell moving left, right or remaining stationary are the same. On average, a T cell will traverse down through a compartment in a straight line, exploring relatively little of the width. Reducing a compartment's width means that DCs are packed horizontally closer together, and this increases the probability that T cells will encounter DCs expressing MHC:peptide complexes for which they are specific.

The same phenomenon can explain why the simulation is significantly less robust towards perturbation of *Circulation_height* than *Circulation_width*. Decreasing the width of the circulatory compartment, where significant CD4Th1 apoptosis by CD8Treg cells occurs, increases the probability that cells will interact with one another. This results in increased CD4Th1 regulation of CD8Treg cells, which in turn increases the degree of stimulation for Treg priming at that point in time. Hence, the peaks of Treg population sizes are increased.

In summary, these results indicate that altering the density of cells in the CNS is critical to the behaviour of the simulation, but less so for other compartments.

6.5.2.2 Seeding of simulation with cells

Table 6.6 summarises simulation robustness with respect to those parameters that dictate the initial number of cells in the simulation, and that were arbitrarily assigned baseline values. The initial number of T cells in the simulation, dictated by *Simulation_numCD4Th*, *Simulation_numCD4Treg* and *Simulation_numCD8Treg*, was arbitrarily chosen, though the ratio between Treg cells and CD4Th1 cells was informed by the domain expert. Robustness analysis finds that perturbations of around 40% or more are typically required to affect significant changes in behaviour.

Simulation_numDC and *Simulation_numDCSpleen* were also assigned arbitrary values, though the ratio between them was informed by the domain expert. The former is unable to affect significant changes in simulation behaviour unless perturbed by several times the default value. Coupled with the lack of significant influence portrayed by this parameter, table 6.3, these results do not cause concern. The latter requires perturbation of 50%, and this affects only three responses, pertaining to the maximum number of effector CD4Treg and CD8Tregs attained, and the number of CD4Th1 cells remaining at 40 days.

These parameters influence the initial density of cells in the simulation. In the case of T cells, reducing the density prolongs the typical length of time required for immune responses to instigate, since the probabilities of T cells encountering immunogenic APCs are reduced. It is noted that the simulation does not encompass any notion of chemokine⁵ secretion by immunogenic APCs, which has the effect of attracting T cells to immunogenic APCs. At present immune response instigation relies on random movement of T cells through secondary lymphoid compartments where APCs reside. Were the density of MBP-specific CD4Th cells, CD4Treg or CD8Treg cells in the simulation to be found inconsistent with *in vivo* findings and cause concern, implementation of this chemokine attractant mechanism may compensate for high numbers of T cells. The same applies were the number of DCs in the spleen found to be artificially high.

In summary, at present, the robustness of the simulation with respect to these parameters does not cause concern.

6.5.2.3 T cell-APC interaction

The parameters *TCell_specificityUpperLimit*, *TCell_specificityLowerLimit*, *DendriticCell_phagocytosisProbabilityImmature*, *DendriticCell_phagocytosisProbabilityMature*, *CNSMacrophage_phagocytosisProbabilityImmature* and *CNSMacrophage_phagocytosisProbabilityMature* have somewhat arbitrarily defined parameter values. Table 6.7 summarises the robustness indices of these parameters.

The simulation is relatively robust with respect perturbation of the APC phagocytosis parameters, in most cases no significant deviation in behaviour occurs, and where it does, perturbations of over 50% are typically required. These robustness indices are not considered problematic. It is noteworthy that *CNSMacrophage_phagocytosisProbabilityMature* is able to increase the mean level of EAE experienced at 40 days by 1.0 when increased by 25% of its default value. The lack of any such deviation in behaviour for *CNSMacrophage_phagocytosisProbabilityImmature* indicates that phagocytosis, and hence peptide derivation, by mature CNS macrophages is important for maintaining neuronal apoptosis. Figure 6.20 shows the effector T cell dynamics and the mean progression of EAE over time for *CNSMacrophage_phagocytosisProbabilityMature* values of 20% and 40%. It may be observed that whilst the difference in T cell population dynamics appears minor, the EAE severities differ substantially. Mature APCs are highly motile, and are more likely to contact apoptotic neurons, which are stationary. Whilst increasing this parameter by 20% may reduce the phagocytosis of apoptotic CD4Th1 cells and neurons by dendritic cells, the lack of substantial difference in T cell priming patterns suggests that this effect is minor. Rather than explain the increased EAE severity through the increased priming of CD4Th1 cells, it is likely that increased phagocytic activity in mature CNS macrophages is resulting in increased MHC-II:MBP expression by these cells, and hence recent CD4Th1 infiltrates in the CNS are able to derive local activation more quickly. Earlier local activation of

⁵Chemokines are molecules, similar to cytokines, that influence cellular movements.

CD4Th1 cells will result in, at a population level, increased type 1 cytokine secretion, which prompts more widespread TNF- α secretion by CNS macrophages, and hence increased neuronal apoptosis.

Phagocytic activity of mature CNS macrophages is hence shown to impact neuronal apoptosis in the CNS. However, should future experimentation indicate that this parameter is incorrectly set, the implications for ARTIMMUS are not overly severe. This parameter's range of effects does not extend to T cell population dynamics, rather, it is limited to the rate of neuronal apoptosis. As such, the *in silico* EAE severity scoring mechanism would likely require re-calibration.




Figure 6.20: Changes in the phagocytic potential of mature CNS macrophages has a marked effect on the mean progression of EAE, although relatively little effect is seen in actual T cell dynamics. This is believed to be due to the time taken for effector CD4Th1 cells to derive local activation in the CNS: increased CNS macrophage phagocytic activity leads to increased expression of MHC-II:MBP.

The simulation also displays robustness with respect to perturbation of T cell specificity related parameters. *TCell_specificityUpperLimit* perturbation is found to be significant with respect to only two responses, when perturbed by over 40%. *TCell_specificityLowerLimit* affects significant changes only when perturbed by over 60% of its default value.

6.5.2.4 Cytokine secretion and decay

Table 6.8 summarises simulation robustness with respect to perturbation of parameters specifying cellular cytokine secretion and cytokine decay. Perturbation of *CD4Treg*-

type1SecretedPerHourWhenActivated and *CD8Treg_type1SecretedPerHourWhenActivated*, which were assigned arbitrary values significantly less than *Th1Polarization_type1SecretedPerHourWhenActivated*, is shown to be inconsequential to simulation behaviour. *Th2Polarization_type2SecretedPerHourWhenActivated* is shown to significantly alter behaviour when perturbed by at least 85%.

Perturbation of both *Th1Polarization_type1SecretedPerHourWhenActivated* and *CNSMacrophage_tnfaSecretedPerHourWhenStimulated* incites significant deviation in simulation behaviour for all responses. In the former case a perturbation of at least 25% is required, in the latter 1% is sufficient to increase the level of EAE severity experienced at 40 days by 1.0. As established in section 6.4.2.1, *CNSMacrophage_tnfaSecretedPerHourWhenStimulated* is directly related to the rate of neuronal apoptosis, which in turn drives both autoimmune and hence regulatory responses. *Th1Polarization_type1SecretedPerHourWhenActivated* may be considered a ‘second order’ parameter in influencing the rate of neuronal apoptosis, type 1 cytokine being required for the activation of CNS macrophages. These results are significant, and point to the importance of balancing parameters that lead to neuronal apoptosis such that the behaviours observed *in vivo* are accurately replicated in the simulation. Together, parameters describing the secretion rates of cytokines, cell sensitivities to cytokines, and *Molecule_molecularHalfLife* dictate the physical reach of cytokine effects. Should it be found that the two arbitrarily determined cytokine secretion parameters discussed here contribute to an unrealistic diffusion of cytokine effect in the simulation, then many other parameters will need recalibration to ensure that this discrepancy is corrected. These include the other parameters directly involved in neuronal apoptosis: *CNSCell_apoptosisTNFaThreshold*, *APC_probabilityPhagocytosisToPeptide*, *Simulation_numCNS*, and *Simulation_numCNSMacrophage*.

Molecule_decayThreshold represents the minimum concentration of cytokine that may exist in a grid space before it is considered to be zero; without this feature the simulation’s decay mechanism would allow cytokine concentrations to approach but never reach zero. This parameter was assigned an arbitrary value intended to be many orders of magnitude smaller than the smallest quantity of cytokine secreted by any cell. Parameter perturbation is found to significantly alter simulation behaviour for several responses, however a perturbation of at least 85% is required, and in most cases the required perturbation is several hundred percent. It may be concluded that simulation behaviour is not critically defined by the exact value held by this parameter.

6.5.2.5 Standard deviation related parameters.

The duration of time that cells spend in particular states are typically drawn from a normal distribution, specified by a mean and standard deviation. The mean values are based on domain-specific knowledge, or are tuned in order to align simulation behaviour with that observed *in vivo*. Parameters specifying standard deviations, listed in table 6.9 have been arbitrarily assigned values that are far from zero but less than those held by the corresponding mean parameter.

Table 6.9 summarises robustness indices for arbitrarily assigned standard deviation parameters. This analysis indicates that perturbation of these parameters is largely inconsequential to simulation behaviour. *TCell_AICDStdDev* is found to exert influence on the *max EAE* response when perturbed by around 30% or more. Figure 6.21 depicts the mean severity of EAE experienced under *TCell_AICDStdDev* parameter values of 0 and 60 hours, and confirms this parameter’s marked effect on the maximum severity of EAE reached. The result may be explained through the requirement for a CD4Th1 cell to reach the CNS and become locally activated before contributing to autoimmune

activity. To illustrate with a hypothetical scenario, were $TCell_AICDStdDev$ to be set to zero, and the average time required for a CD4Th1 to migrate into the CNS compartment and become activated be 40 hours, then the average time that each CD4Th1 cell would perform effector function for would be 20 hours. Increasing $TCell_AICDStdDev$ to 60 hours, many cells would remain in effector state for significantly longer than 60 hours, and many would enter apoptosis prior to reaching the CNS.

For any one CD4Th1 cell, increasing values for $TCell_AICDStdDev$ that further reduce its effector life-span below the 40 hours required to reach the CNS does not further decrease its contribution to autoimmunity; it makes no contribution. However, the other extreme is that a cell experiencing a similarly extended lifespan *would* provide additional contribution to autoimmune activity. Since some proportion of a CD4Th1 cell's effector lifespan is spent migrating to the CNS and awaiting local activation, increases in autoimmune contribution resulting from increases in $TCell_AICDStdDev$ outweigh the reduced contribution of cells that apoptose before being locally activated in the CNS.

Should the default value of 56 hours for $TCell_AICDStdDev$ be deemed inappropriate, the *in silico* EAE severity scoring mechanism would require re-calibration to reflect the fact that the rate of neuronal apoptosis is reasonably dependent on this parameter.




Figure 6.21: The mean severity of EAE experienced under parameter values of 0 and 60 hours for $TCell_AICDStdDev$.

6.6 Qualifying the significance of simulation results

The previous section has examined simulation robustness with respect to perturbation of parameters that were assigned arbitrary values, since no guidance exists on appropriate values that they might otherwise adopt. Finding that simulation behaviours critically depend on these parameters holding values from a very small range of their possible domain has implications on the interpretation of simulation results; were these parameters assigned other arbitrary values, simulation behaviours may differ significantly. Reporting these criticalities is important when communicating *in silico* results, other researchers may have reason to doubt the values adopted in simulation. The robustness analysis is an initial step towards qualifying the significance of simulation-derived results in the real domain, and the confidence that may be placed in them. The present section expands on this theory, explaining how the robustness and sensitivity analyses may be combined with domain-specific knowledge concerning biologically plau-

sible ranges of parameter values to further this goal. Three methods of application are explored below.

Firstly, the boundaries of biologically plausible parameter ranges can be contrasted with robustness boundaries, derived through the robustness analysis, to ascertain whether a particular simulation behaviour may be explained by potentially inaccurate parameter values, or whether it is genuinely representative of the mechanics upon which the simulation is built. This is illustrated in figure 6.22. If it is found that simulation behaviours significantly deviate when simulation parameters are perturbed to other biologically-plausible values within their domain, then less confidence should be placed in those results being representative of the original domain. Should this occur, the significance of the discrepancy, and hence the degree of confidence assumed in the results, may be ascertained through assessing that parameter's influence over simulation behaviour using a global sensitivity analysis. If the parameter is found to have little influence over the simulation's behaviour, then the robustness boundaries lying within the biologically-plausible boundaries is not as significant as would be the case for a very influential parameter.

Secondly, the robustness analysis technique need not take the simulation's calibrated baseline parameter values as the point from which parameters are perturbed. It may instead be set to a point in parameter space that reflects a prediction of the system. In this case the application of the robustness analysis technique would constitute breaking the prediction, after which the conditions required to break the prediction can be contrasted with domain-specific knowledge. For example, *in silico* experimentation might investigate the nature of a regulatory pathway by perturbing a parameter specifying the efficacy of its operation. This might reveal a switching in simulation behaviour only when the efficacy is reduced to a certain point; that the regulatory pathway continues to operate correctly until efficacy is reduced to a particular level. Assuming this level as the baseline value for that parameter, and then performing a robustness analysis would estimate⁶ the area of parameter space under which this prediction holds. This may then be contrasted with domain-specific knowledge to qualify the certainty of the prediction: the degree to which it is an artifact of assuming not-entirely biologically plausible parameter values, or a genuine representation of the original domain.

Lastly, an alternative approach to qualifying when simulation results may be assumed genuinely representative of the original domain is to ascertain the simulation behaviours that occur when using the most extreme values of biologically-plausible parameter ranges. These simulation behaviours outline the range of behaviours

Lastly, an alternative approach to qualifying when simulation results may be assumed genuinely representative of the original domain is as follows. It relies on ascertaining the extremes of simulation behaviour that occur when using any combination of parameter values from within their biologically-plausible ranges. These extremes represent the range of simulation behaviours that can be attributed to biological uncertainty (uncertainty concerning exact figures in the real domain). By contrasting the results of *in silico* experimentation with these extremes, one can determine whether results are genuinely representative of the real-world domain, or may simply be explained by the inability to fully specify parameter values. As illustrated in figure 6.23, this strategy could be implemented using latin hypercube sampling, where the range of values from which samples in parameter space are extracted are set to biologically plausible analogues. Representative simulation behaviours may then be obtained at

⁶Application of the above robustness analysis in this context would not constitute exactly identifying the area of parameter space, since it is a *one at a time* analysis; compound effects of moving two parameters away from their baseline values would not be revealed.

each of these points in parameter space, and the extremities of simulation behaviour extracted therefrom. If there exist compounding effects between several parameters, for example if simulation behaviours change substantially when two parameters hold values within their biologically plausible ranges but not near the extremities, then these may be highlighted through use of latin hypercube sampling.

The approaches to qualifying the significance of *in silico* results in terms of the real-world domain proposed here are theoretical only. Their application and further exploration is considered outside the scope of this thesis.




Figure 6.22: The robustness indices arising from robustness analysis can be contrasted with biologically-plausible ranges of parameter values. This can indicate the degree to which simulation results are explained by uncertainty in the domain, and hence how much confidence one may place in their being representative of the real-world domain rather than underspecified parameter values.




Figure 6.23: The range of simulation behaviours which may be assumed representative of the domain, rather than the results of uncertainty concerning exact biological figures, may be calculated by constructing a latin hypercube around the ranges of biologically plausible values. Simulation behaviour at each point in parameter space is ascertained, and the range of simulation behaviours arising from biological uncertainty extracted. Results of *in silico* experimentation falling outside this range may be assumed more representative of the real-world domain, rather than the result of underspecified parameter values.

6.7 Conclusion

This chapter has explored the importance of cells, pathways and parameters in ARTIMMUS through statistical techniques. Further, it has investigated how statistical techniques may be employed to aid interpretation of *in silico* results in terms of the original domain. As such, the chapter has addressed research objectives 4 and 5: to perform novel *in silico* experimentation using ARTIMMUS, and to develop and apply statistical techniques for interpreting *in silico* results in the context of EAE.

Section 6.2 establishes metrics of simulation behaviour necessary for the application of sensitivity and uncertainty analyses. Nine of the metrics relate directly to system-level T cell dynamics, the remaining two concern a means to grade simulation executions using the same six point scale employed in the wet-lab. The mapping of this six point grading system onto simulation executions better integrates simulation and

real-world results. A novel *consistency analysis technique* that establishes the relationship between the accuracy of averaged simulation results and the number of simulation executions sampled in deriving them is presented in section 6.3. This technique allows researchers to balance computational requirements with the precision required in their analyses. Section 6.4 reports the application of a global sensitivity analysis, based on latin hypercube sampling, in assessing the influences of ARTIMMUS’s components and pathways on its overall behaviour. The most influential parameters in the simulation are those that pertain directly to neuronal apoptosis; this phenomenon drives the autoimmune response, which in turn provides stimulation for the regulatory immune response. A novel robustness sensitivity analysis is presented in section 6.5. It establishes the range of parameter values over which simulation behaviour is statistically consistent. The analysis is used to assess the extent to which ARTIMMUS’s behaviour critically depends on the parameters assigned arbitrary values during calibration. Should simulation behaviours critically depend on these parameters holding very specific values, then caution must be exercised in drawing conclusions from the results, had different arbitrary values been assigned the simulation’s behaviour might differ substantially. Section 6.6 expands upon this theory by considering a variety of ways in which sensitivity and robustness analyses can be applied and considered in the context of domain-specific knowledge to qualify the significance of *in silico* results in the target domain.

The application of the global sensitivity analysis reveals aspects of the simulation that are influential on its overall behaviour. Such an analysis could not be performed in the wet-lab, and as such it constitutes novel *in silico* experimentation. Several insightful predictions have arisen from this analysis, summarised here:

- The relatively inconsequential effect of perturbing parameters pertaining to the regulatory pathway, the ability of CD8Treg cells to apoptose CD4Th1 cells and the duration for which such cells express Qa-1 for, leads to the prediction that there is considerable redundancy in the regulatory pathway’s ability to ameliorate autoimmune behaviour. This is investigated further in the following chapter.
- Results of the analysis indicate that there is not a delicate balance between type 1 and type 2 cytokine in the grid spaces where they influence cellular behaviours. Rather, the relatively inconsequential effect of perturbing cytokine secretion and sensitivities leads to the prediction that in any grid space where cytokines lead to polarisation related decisions, one cytokine substantially outnumbers the other in concentration.
- The type 2 deviation of the autoimmune response, observed during recovery, is more a result of regulatory action than a driver of it. The parameters dictating the probabilities of CD4Th cells adopting particular polarisations were found to be of little consequence when perturbed. This was further investigated by altering parameters such that the conditions that ordinarily prompt strong preference for CD4Th2 polarisations instead exclusively lead to adoption of type 1 polarisations revealed very little change in the ability of the system to recover from autoimmunity.

The following chapter conducts further *in silico* experimentation into EAE.

Chapter 7

Projective Experimentation

The previous chapter has explored the role of various statistical techniques in interpreting *in silico* results. The application of global sensitivity analysis has lead to novel insight into EAE. This chapter reports further novel *in silico* experimentation into EAE, performed on ARTIMMUS.

Firstly, section 7.1 motivates this chapter. There are three stands of experimentation reported here. The first, section 7.2, examines the cellular dynamics underpinning autoimmunity and recovery both in the presence and absence of regulatory activity, the two conditions used in calibrating ARTIMMUS. Secondly, section 7.3 examines various aspects of regulation that lead to recovery. Lastly, section 7.4 reports the use of ARTIMMUS in gaining insight into the role of the spleen in EAE. Lastly, section 7.5 concludes this chapter.

7.1 Goal and motivation

This chapter directly addresses research objective 4: perform novel *in silico* experimentation using the agent-based simulation of EAE. In addition to providing contribution to the field of EAE, performing *in silico* experimentation provides context in which to consider the extent to which the modelling and statistical techniques explored in this thesis provide confidence in simulation results being representative of the real-world domain.

There are three themes of experimentation reported in this chapter. Section 7.2 investigates the behaviour of ARTIMMUS under conditions of physiological recovery and prolonged autoimmunity, achieved through disabling the ability of CD8Treg cells to regulate CD4Th1 cells. These two complementary conditions reflect those used in the calibration of the simulation. This experimentation examines the states of T cell activation over time, the locations in which T cell populations are primed, the states of APC activation in each compartment over time, the polarizations of DCs in the CLN during the course of the immune response, the locations in which regulation of the CD4Th1 population takes place, the rate of neuronal apoptosis and resultant progression of EAE severity over time, and lastly an analysis of the relapsing nature of EAE when regulation is disabled.

The second theme of experimentation, presented in section 7.3, examines the regulatory pathway by which effector CD8Treg cells are able to induce apoptosis in Qa-1

expressing CD4Th1 cells in greater detail. This regulatory pathway has a critical role in recovery from EAE, and the *in silico* experiments conducted here cannot be performed in the real-world domain. Three sets of experimentation are presented. The first examines the effect of reducing the efficacy of the regulatory pathway, reducing the probability that a successful binding between an effector CD8Treg cell and a Qa-1 expressing CD4Th1 cell will lead to the former inducing apoptosis in the latter. The second considers the effect of altering the mean duration of time for which effector CD4Th1 cells express Qa-1, Qa-1 expression being a requisite for regulation. The last experiment considers the possibility that section 7.2's finding that the majority of CD4Th1 cells are induced into apoptosis in the circulatory system is an artifact of an underspecified simulation parameter, rather than being genuinely representative of the underlying simulation mechanics.

The last theme of experimentation, presented in section 7.4, employs ARTIMMUS in gaining insight into a set of *in vivo* experimentation which demonstrated that splenectomy (removal of the spleen) in rats subsequently induced into EAE is associated with increased autoimmune severity. A splenectomy experiment is engineered into the simulation, and the manner in which this reduces regulatory capacity of Treg populations is examined.

7.2 Elucidation of baseline behaviour

The present section details the use of the ARTIMMUS simulation in gaining insight into the nature of EAE. Two experimental scenarios are considered, that of physiological CD8Treg mediated recovery, and prolonged autoimmunity through the abrogation of the regulatory pathway. These two scenarios were the targets of calibration activities reported in chapter 5; the simulation and its parameters were amended in a manner that aligned simulation behaviour with that observed *in vivo* under these two scenarios. Both these scenarios are considered here as the simulation's behaviour in the presence and absence of regulation affords greater insight into its overall dynamics.

Section 7.2.1 details the experimental methodology used to acquire the data that is then presented and analysed in section 7.2.2. The following aspects of simulation behaviour are examined: the dynamics of each T cell population over time, with consideration of their states of activation; the locations of T cell priming, and the numbers of each T cell sub-type that are primed in each; the states of APC activation in each of the simulation's compartments; the polarizations of DCs in the CLN compartment; the locations in which CD8Treg mediated apoptosis of effector CD4Th1 cells takes place; the rates of neuronal apoptosis over time, and the resultant dynamics of EAE severity; and lastly the relapsing and remitting nature of EAE in simulations with regulatory function disabled is analysed. These results are summarised in section 7.2.3.

7.2.1 Experimental methodology

The condition of physiological recovery is achieved in the simulation through using baseline parameter values, detailed in section B.1. This experimental scenario represents the default behaviour of the simulation.

The prolonged autoimmunity scenario is achieved through setting the *CD8Treg-cd8TregToCD4ThelperSpecificityDropOff* parameter to 0.0. This parameter dictates the probability that a successful binding between a CD8Treg and a Qa-1 expressing effector CD4Th1 cell will lead to the apoptosis of the latter as induced by the former. By setting this parameter to 0, CD8Treg cells are unable to regulate the CD4Th1 population.

In acquiring representative simulation behaviours, 500 simulation executions have been obtained for each of these two experimental scenarios. Representative behaviour over time is then compiled by calculating, for any phenomenon of interest, the median value across 500 simulation runs for that phenomenon at each point in time.

7.2.2 Results

Figure 7.1 depicts effector T cell population dynamics under conditions of physiological recovery and prolonged autoimmunity. These graphs are provided for reference in the sections that follow.




Figure 7.1: Effector T cell population dynamics.

7.2.2.1 States of T cell activation

This section examines the states of activation of the CD4Th, CD4Treg and CD8Treg populations over time. Figure 7.2 depicts the CD4Th population, figure 7.3 the CD4Treg population, and figure 7.4 the CD8Treg population.

In all T cell populations, proliferating cells outnumber those in any other state. The increase in naive T cells that lags behind the increase in proliferating cells suggests that naive T cells resulting from proliferation cannot always gain immediate antigenic stimulus from the same DC that their parents are primed by. The delay between naive and proliferating T cell increases indicates that this occurs during the heavy T cell priming stage of the immune response. This suggests that during the immune response dendritic cells (DCs) become *spatially saturated*. Each DC is surrounded by eight grid spaces that T cells may occupy. This allows up to 56 T cells to physically interact with each DC. T cells spend on average 60 hours in a proliferative state, in which they produce a daughter naive T cell every \sim 20 hours. The mean lifespan of a mature DC is 110 hours. Discounting time spent migrating (for DCs originating from the CNS), and assuming that a single T cell begins priming on the DC immediately following its maturation, this allows for 5 complete generations of T cells to be primed on the average APC. A single T cell will, after 5 generations of proliferation, differentiation and subsequent migration, produce \sim 28 proliferating T cells that are bound to the DC. However, these hypothetical conditions better approximate the system at the start of an immune response. When a priming DC enters apoptosis the proliferating cells with which it was bound will resume migratory behaviour, and if they have not already differentiated into effector cells they may bind another immunogenic DC to resume

proliferative behaviour. As the immune response progresses, the number of proliferating T cells in the system increases, and the likelihood that DCs become spatially saturated increases, since many more than one naive or proliferating T cell may bind with it.

A common trend across all the T cell populations is the lack of partially activated T cells. This suggests, and is supported by figures 7.8, 7.9, 7.11 and 7.10 which are discussed below, that there are very few tolerogenic APCs present through the course of the immune response.

7.2.2.2 Locations of T cell priming

Figures 7.5, 7.6 and 7.7 show cumulative counts of the number of CD4Th, CD4Treg and CD8Treg cells primed in each compartment, under conditions of both physiological recovery and prolonged autoimmunity.

With regulatory function intact, the number of CD4Th cells that are primed in the SLO as a result of immunization almost equals the number primed in the CLN compartment. CD4Th priming in the CLN is the result of autoimmune activity in the CNS, wherein dendritic cells (DCs) obtain MBP which is then presented on MHC-II:MBP complexes through the phagocytosis of apoptotic neurons before migrating into the CLN compartment. Priming in the CLN begins around day 9 following immunization, and SLO-based priming of the CD4Th population ceases around day 17. Hence, the priming of the CD4Th population shifts from the SLO to the CLN during this time. The rate of priming in the SLO is greater than in the CLN, this is likely because the number of immunogenic DCs in the SLO resulting from immunization is generally higher than the number of MBP-presenting immunogenic DCs in the CLN at any point in time. This is supported by figures 7.8 and 7.9, which depict the number of immature, tolerogenic and immunogenic APCs in the SLO and CLN compartments, and is discussed below. Under prolonged autoimmune conditions, shown in figure 7.5b, the priming of CD4Th cells continues in the CLN until day 200, at which time observation ends.

Priming of the CD4Treg population, figure 7.6, takes place in the SLO, the CLN and the spleen. In comparison with the CLN and the spleen, very little priming takes place in the SLO. The number of CD4Tregs primed in the spleen is just over double the number in the CLN. There are multiple factors that contribute to these compartmental priming profiles, considered below.

Priming in the SLO is likely the result of apoptotic CD4Th cells being phagocytosed by the 10 resident DCs that permanently exist there; the DCs placed there as a result of immunization have reduced to very few in number (figure 7.8, discussed below) by the time that the little priming of the CD4Treg population that takes place in the SLO occurs. This can explain why relatively little CD4Treg priming occurs in the SLO.

Apoptotic T cells cannot migrate from the circulatory compartment into any of the CLN, SLO, or CNS compartments. Any apoptotic CD4Th cells that reach the circulatory system must be phagocytosed in the spleen. CD4Th cells may enter apoptosis from naive, partially activated, or effector states of activation, given sufficient lack of TCR stimulation. Figures 7.6 and 7.7 show that the majority of CD4Treg and CD8Treg cells are primed in the spleen. These data suggest that the majority of apoptotic CD4Th cells fail to be phagocytosed in the CLN or the SLO, or by DCs in the CNS, and are eventually phagocytosed in the spleen. Note that these data do not indicate how many CD4Th cells are phagocytosed by CNS macrophages in the CNS, which would not result in priming of Treg populations.

Between days 15 and 35 the CD8Treg population apoptoses \sim 3000 CD4Th1 cells (figure 7.13, discussed below). With the regulatory pathway intact, the maximum




Figure 7.2: CD4Th population states of activation.




Figure 7.3: CD4Treg population states of activation.




Figure 7.4: CD8Treg population states of activation.

number of CD4Th1 cells in the system reaches a peak of around 1000 at day 15, and then reduces to around 200 in the following 5 days (figure 7.1). This is several times less than exists at 20 days with regulatory function disabled, and suggests that a considerable proportion of CD4Th1 cells are apoptosed by the CD8Treg population before reaching the CNS. Figure 7.13, discussed below, confirms the majority of this regulatory activity takes place in the circulatory system. Hence, these CD8Treg-apoptosed Th1 cells can only be phagocytosed in the spleen. The restriction that these apoptotic cells be phagocytosed in the spleen explains why abrogating regulatory function increases the proportion of total Treg priming that occurs in the CLN.

The compartmental priming profile of CD8Tregs, figure 7.7, is qualitatively very similar to that of CD4Tregs. The notable exception is that whilst the number of CD4Tregs being primed in the CLN is between 40% and 50% of those being primed in the spleen, here the ratios are 20% with regulation intact, and 35% with regulatory activity disabled. This phenomenon may be due to the requirement for a DC to be licensed for Qa-1 expression by an effector CD4Treg before it may begin priming CD8Tregs, and the fact that CLN based DCs must first migrate out of the CNS compartment and then settle in the CLN compartment before any priming can take place. A DC exists in a mature state for around 110 hours before entering apoptosis; in the case of CNS-originating DCs some of this time is spent migrating between compartments. This is not the case for splenic DCs, which may begin priming immediately. There exists an exponential relationship between the time that a T cell population spends priming on a DC, and the number of T cells that are eventually produced. The coupling of the delay in CLN DCs migrating between compartments with the time taken for them to become licensed for Qa-1 may account for the reduced proportion of CD8Treg cells that arise from priming in the CLN when compared to the spleen.

7.2.2.3 States of APC activation

This section considers the states of activation and different contexts in which APCs present peptides in the various compartments of the simulation. Figure 7.8 shows the SLO, figure 7.9 the CLN, figure 7.10 the spleen, and figure 7.11 the CNS. In each case four APC states are considered: the total number of APCs in the compartment, the number of immature APCs, the number of tolerogenic APCs, and lastly the number of immunogenic APCs. Note that for an APC to be considered either tolerogenic or immunogenic it must be presenting peptides. Mature APCs that do not present any peptides are not considered.

The SLO compartment ordinarily contains 10 resident DCs that are replaced by immature DCs when they expire. Figure 7.8 shows how immunization increases the total number of DCs to over 50. These ‘immunization’ DCs represent DCs that migrate to the lymph nodes from the periphery following immunization with MBP, CFA and PTx. The plateau of immunogenic DCs around 18 in number between days 10 and 13 is the result of a temporary steady state being reached between the rate at which DCs are placed in the SLO and the rate at which they expire. This is an artifact of the immunization mechanism: immunogenic DCs are placed in the SLO every 6 hours, and the number that are added slowly decreases over time, according to the parameters *Simulation_immunizationLinearDC0* and *Simulation_immunizationLinearGradient*. These two parameters describe a linear decay in the level of DCs that are added over time, the actual number that appear every 6 hours is a rounding of this level to the nearest natural number. For the first 5 days following immunization 2 DCs are added every 6 hours, thereafter this number reduces to 1 DC until around 12 days, after which the influx of immunogenic DCs ceases.




Figure 7.5: Cumulative count CD4Th cells primed in each compartment.




Figure 7.6: Cumulative count of CD4Treg cells primed in each compartment.




Figure 7.7: Cumulative count of CD8Treg cells primed in each compartment.

The CNS compartment contains around 40 immature DCs and exactly 75 CNS macrophages at any point in time. The initial fluctuations and gradual increase in total APC number are due to the system finding a steady state between the number of mature CNS-originating DCs in the CLN and the number of immature DCs in the CNS. These cells exist as a closed system; a mature CNS-originating DC entering apoptosis in the CLN is replaced by an immature DC in the CNS. Upon simulation initiation, the numbers of DCs that should be placed in the CNS and in the CLN are estimated based on the assumption that all CNS-originating DCs in the CLN are mature and all those DCs in the CNS are immature. In reality this is not the case, some CNS-resident APCs are mature and in the process of migrating towards the CLN. Hence, the initial number of CNS-resident DCs is underestimated, and the number in the CLN overestimated. This results in a minor correction as simulation execution progresses.

At peak autoimmune behaviour, around 50 of the 120 APCs in the CNS are immunogenic. These cells are capable of providing local activation for effector CD4Th infiltrates. Discounting the ~40 immature DCs, this figure suggests that around 65% of CNS macrophages are immunogenic and presenting MBP¹. A small quantity of CNS APCs adopt a tolerogenic phenotype during heavy immune activity. Since no tolerogenic DCs appear in the CLN (figure 7.9), these cells may be migratory DCs that perceive sufficient type 1 cytokine so as to adopt an immunogenic phenotype en route to the CLN, or they may be CNS macrophages. In either case the tolerogenic phenotype will not influence the CD4Th infiltrates; these effector cells seek only MHC-II:MBP stimulation, the lack of co-stimulatory molecules expressed on tolerogenic APCs only influences T cells at time of priming.

As with the SLO, the CLN compartment contains 10 resident DCs, each of which is replaced by an immature DC when it perishes. As depicted in figure 7.9, there are around 90 additional DCs that originate from the CNS, migrating to the CLN upon maturation. At peak autoimmunity around 30 DCs exist in an immunogenic state. Notably, at no point during simulation execution, under conditions of either physiological regulation or prolonged autoimmunity, do any cells in the CLN exist in a tolerogenic state. This suggests that, even in peak or prolonged autoimmune conditions, only 30% of migratory DCs are able to prime and immunize T cell populations.

Figure 7.10 shows the states of splenic DCs over time. There are 100 DCs residing in the spleen at all times. During peak immune activity around 25 DCs are immunogenic. Preceding the establishment of immunogenic DCs is a short period during which at most 10 DCs exist in a tolerogenic state. This is likely due to the lack of type 1 cytokine being secreted in the spleen. The earliest T cells to become primed into effectors are CD4Th cells, and these will not generally secrete type 1 cytokine until they are locally activated in the CNS, though the stochastic nature of the simulation dictates that some may derive local activation in the compartments in which they are primed, should they reside there for sufficient time. CD4Treg and CD8Treg effector cells also secrete type 1 cytokine, and it is likely that these cells are the major inducers of immunogenic presentation contexts in splenic DCs.

It is of note that there exists no period of tolerogenic APC activation at the end of either the autoimmune (days 20 to 30) or regulatory (days 40 to 50) immune responses, in any of the compartments. Tolerogenic APCs induce anergy in priming T cells, and it is theoretically possible that this aids in curtailing the immune response. However, the present data suggests that this is not the case in ARTIMMUS.

¹For an APC to be considered either immunogenic or tolerogenic it must express either Fr3, CDR1/2 or MBP peptides. Since CD4Treg and CD8Treg cells do not infiltrate the CNS, the ability of CNS macrophages to express Fr3 or CDR1/2 is not represented in the simulation. Hence for them to be considered immunogenic or tolerogenic, they must be expressing MBP.




Figure 7.8: States of APC activation in the SLO.




Figure 7.9: States of APC activation in the CLN.

7.2.2.4 Polarizations of DCs in the CLN

Figure 7.12 depicts the number of type 1 and type 2 polarized immunogenic DCs in the CLN over time, for conditions of physiological recovery and prolonged autoimmunity. The vast majority of immunogenic DCs in the CLN originate from the CNS compartment, hence these graphs indicate how the cytokine milieu resulting from immune action in the CNS influences DC priming behaviour. The first immunogenic DCs to migrate out of the CNS do so at around 7 days post-immunization. The majority of these cells are type 1 polarized, reaching a peak of around 19 cells at day 17.

With regulatory activity intact, the number of type 1 polarized DCs reduces to 0 by day 28. Interestingly, there is a marked type 2 polarized DC presence throughout this period, reaching a peak, and becoming the majority polarization type in the CLN, at around day 22. The presence of type 2 polarized DCs extends beyond the duration of type 1 DCs, and eventually returns to 0 by day 37. This majority of type 2 DCs from 22 to 37 days explains the type 2 deviation in the CD4Th immune response seen to peak at day 30 on figure 7.1a.

Under conditions of prolonged autoimmunity, the number of type 1 polarized DCs continually outnumbers type 2 polarized DCs. Of note, the maximum number of type 2 polarized DCs decreases from 16 to 11 under prolonged autoimmunity. The presence of type 2 polarized DCs remains throughout the 200 days of observation, consistently




Figure 7.10: States of APC activation in the Spleen.




Figure 7.11: States of APC activation in the CNS.

representing a third of the total immunized DC population in the CLN.

7.2.2.5 Locations of CD8Treg mediated CD4Th1 apoptosis

Figure 7.13 depicts a cumulative count of the number of CD4Th1 cells induced into apoptosis by CD8Tregs in each of the simulation's compartments, in simulations undergoing physiological recovery. The CNS is not included, since CD8Tregs cannot gain entry to this compartment. It may be observed that although CD8Treg mediated apoptosis of CD4Th1 cells takes place in the CLN, SLO, spleen and circulatory system, the vast majority of apoptosis induction takes place in the circulatory system.

7.2.2.6 Neuronal apoptosis and EAE severity

It was established in the previous chapter (section 6.4.2.1) that neuronal apoptosis is the driving stimulus of autoimmune behaviour, and hence the regulatory immune response that counteracts it. The present section examines neuronal apoptosis and the resultant severities of EAE in the simulation under conditions of physiological recovery and disabled regulation.

Figure 7.14 depicts the number of neurons apoptosed per hour of simulated time. The results presented here are median values taken across 500 simulation executions.




Figure 7.12: Dendritic cell polarizations in the CLN.




Figure 7.13: Compartments in which CD4Th1 apoptosis induction by CD8Treg cells takes place.

Despite this large sample size, the variation in number of cells apoptosed from one hour to the next is striking. The smoothed rates of neuronal apoptosis, from which EAE severities are calculated, are indicated on the diagrams. Under conditions of physiological recovery the smoothed number of neurons apoptosed per hour reaches a peak ~ 70 around day 15. With regulatory action disabled, this peak is raised slightly, and may be the result of the slightly increased maximum number of effector CD4Th1 cells reached over time, shown in figure 7.1. Whereas regulatory action reduces the rate of neuronal apoptosis to 0 by day 40, disabled regulation sees a relatively steady level of around 60 to 50 neurons being apoptosed every hour. Cross reference of figures 7.1b and 7.14b reveals that around 1000 effector CD4Th1 cells cause neuronal apoptosis at a rate of 70 neurons per hour, and around 400 effector CD4Th1 cells reduces this number to only 50 neurons per hour. This suggests a nonlinear relationship between the number of effector CD4Th1 cells in the system and the rate of neuronal apoptosis, with a diminishing increase in neuronal apoptosis resulting from linear increase in CD4Th1 cell number. This can be explained through the spatial occupancy of the CD4Th populations in the CNS. Figure 7.15 shows a screenshot of the ARTIMMUS CNS compartment at day 18 following immunization, the time of peak autoimmunity. It is representative of typical simulation behaviour. It can be seen that CD4Th1 and CD4Th2 cells reside at significantly higher densities at the top of the compartment than at the bottom. It is likely that additional CD4Th1 cells present in the top of the

CNS, where many others already reside, will not contribute to significant additional neuronal apoptosis.

The mean progression of EAE severity, for both physiological recovery and prolonged autoimmunity, is depicted in figure 7.16. The distributions of EAE severities over time amongst 500 simulation executions are depicted in figure 7.17.

Physiological recovery sees a mean peak in EAE severity of around 2.5. 86% of simulation executions experience full recovery, whilst the remaining 14% perish; this is to be expected, given the manner in which the EAE severity scoring mechanism was calibrated in the previous chapter. It is notable that nearly all simulation executions experience at least level 1 EAE at peak of autoimmunity. Around half reach a maximum EAE severity of 2.

With the regulatory pathway disabled, the severities of EAE are generally increased. The peak mean EAE severity is raised to nearly 3.5. Over 35% of simulations perish, and over 50% of executions experience at least level 3 EAE severity at peak autoimmunity. At 200 days post-immunization just over 50% of simulations experience no clinical symptoms. Around 10% remain on level 1.

Figures 7.18 and 7.19 together show the progression of EAE and effector T cell population dynamics for five example simulation executions with regulatory activity disabled. The executions are indicative of the variations that exist between individual simulation executions. In the first example of figure 7.18, the T cell dynamics reduce considerably around 150 days, however, once a simulation's neuronal apoptosis rate passes level 5 EAE severity the simulation is deemed to have perished. The second example demonstrates that the most severe episode of autoimmunity may occur long after initial immunization. In this case immunization induces level 2 autoimmunity within the first 30 days, but level 4 is reached at around 130 days. In absence of regulatory capacity, the autoimmune response does not tend to cease, though it may reduce considerably. This is depicted in the third example of figure 7.18, where all T cell subtype between days 100 and 200 lie below 200 in number, sufficiently low for no clinical symptoms to emerge after day 80. Yet the immune responses do not cease completely. The first example of figure 7.19 demonstrates how there may exist long remissions from autoimmune episodes, over 50 days in length. However, since the T cell responses do not cease, further autoimmunity will frequently present later. The last example of figure 7.19 shows a simulation where autoimmune symptoms present for a large portion of the 200 days of observation, but never strong enough to induce death.




Figure 7.14: Neurons apoptosed per hour.




Figure 7.15: Screenshot of ARTIMMUS, showing the CNS compartment at day 18 post-immunization. Red cells are Th1 cells; white, Th2; blue, apoptotic Th cells; gray, neurons; purple, DCs; pink, CNS macrophages.




Figure 7.16: Mean progression of EAE.

7.2.2.7 Analysis of relapsing autoimmunity

The previous section established that simulations undergoing Treg-mediated physiological recovery from autoimmunity experience complete recovery. In contrast, with regulatory action disabled, simulations experience prolonged autoimmune symptoms over the 200 days of observation. The present section examines the nature of this prolonged autoimmunity in greater detail.

Figure 7.20 presents an analysis of the relapsing nature of EAE experienced when regulatory function is disabled, and contrasts this with the progression of EAE with regulation intact.

The mortality rate of mice undergoing physiological recovery is 14%, and abrogation of the regulatory pathway increases this to 36%. In simulations lacking regulatory function the vast majority of death from EAE occurs in the first onset of autoimmune symptoms, as seen by the convergence of the duration of single incidence of EAE towards the mortality rate of 36.4% on figure 7.20c. For simulations that do not perish during their first episode of autoimmune symptoms, very few do not experience relapsing autoimmunity. 5 to 6 relapses of autoimmunity are common, with over 10% of simulation runs experiencing each of these numbers. 6.6% of simulations experience over 10 relapses in the 200 day observation period. The initial episode of autoimmunity is extended considerably when regulatory function is disabled. In presence of regulatory action EAE symptoms last no longer than 13 days, unless the simulation has




Figure 7.17: Mean distributions of EAE progression. Figures 7.17a and 7.17b show the proportion of simulation runs experiencing each level of EAE. Figures 7.17c and 7.17d show a cumulative representation of this data, the proportion of simulation runs experiencing *at least* the specified EAE severity.

perished. The most common durations are between 9 and 11 days in length. Disabling the regulatory pathway increases the median duration from 10 days to 20, excluding simulations that perish. Relapses following this initial episode are shorter, the median reducing to around 7 days or less.

In simulations with regulatory function disabled, the median duration of remissions is around 10 days. This does not change noticeably with subsequent remissions, with the exception that the 6th remission and greater tend to be shorter.

7.2.3 Summary

This section has detailed the use and analysis of the ARTIMMUS simulation in gaining insight into the nature of EAE. Two experimental conditions have been examined, that of physiological CD8Treg mediated recovery from autoimmunity, and prolonged autoimmunity achieved through disabling the ability of CD8Treg cells to apoptose Qa-1 expressing effector CD4Th1 cells. These two conditions were used in the calibration of ARTIMMUS, as detailed in chapter 5. The examination of both these scenarios provides deeper insight into the cellular dynamics and interactions underpinning EAE and its recovery. The results are summarised as follows.

- Section 7.2.2.1 examined the dynamics of the four T cell populations, in terms of




Figure 7.18: Progression of EAE and the effector T cell population dynamics for 3 example simulation executions, shown one per row, with regulatory activity disabled.




Figure 7.19: Progression of EAE and the effector T cell population dynamics for 2 further example simulation executions with regulatory activity disabled.


their states of activation. The data suggested that APCs may become spatially saturated, reaching the upper limit on the number of T cells that are able to gain physical access to each dendritic cell (DC).

- The spleen was revealed as a major source of CD4Treg and CD8Treg priming in section 7.2.2.2, which analysed the compartments in which T cell priming takes place.
- Section 7.2.2.3 examines the states of APC activation in the simulation’s compartments. It establishes that, in peak immune activity, 65% of CNS macrophages, 30% of CNS-originating DCs in the CLN, and 25% of DCs in the spleen are immunogenic, being both mature and capable of priming T cell populations. Further, it was revealed that there exists no period of significant tolerogenic APC activity in any of the simulation’s compartments at the end of the autoimmune and regulatory immune responses, suggesting that induction of anergy in T cell populations, as mediated through tolerogenic APCs, is not an integral component in terminating the immune responses.
- Section 7.2.2.4 establishes that type 2 polarized DCs exist in the CLN throughout the autoimmune response, and in presence of regulatory activity become the majority DC type towards the end of T cell priming activity in this compartment. In the absence of regulatory activity, there is a consistent presence of type 2 polarized DCs, but at consistently lower levels than type 1 polarized DCs.


- The locations in which CD8Treg mediated apoptosis of the CD4Th1 cell population takes place was examined in section 7.2.2.5. It was found that the vast majority of this regulation takes place in the circulatory compartment.
- Section 7.2.2.6 investigates the rate of neuronal apoptosis and the resultant EAE severities experienced by simulations. The nonlinear relationship between the number of effector CD4Th1 cells in the CNS compartment and the rate of neuronal apoptosis was revealed; there is a diminishing increase in neuronal apoptosis with linearly increasing numbers of CD4Th1 infiltrates. Under conditions of physiological recovery, half of all simulations experience at most level 2 autoimmunity. In absence of regulatory capacity, over 50% of simulations reach at least level 3 autoimmunity at peak autoimmune activity. The abrogation of regulatory capacity raised the mortality rate of EAE from 14% to 36%.
- Lastly, section 7.2.2.7 examines the relapsing nature of EAE in absence of regulation. It is revealed that the vast majority of death resulting from autoimmunity occurs in the first episode of autoimmune symptoms. Those simulations that do not perish during this first episode almost exclusively go on to experience remissions and then relapses in autoimmune symptoms. Relapses are significantly shorter in duration than the initial episode of autoimmunity following immunization. Abrogation of the regulatory pathway significantly increases the duration of the first episode of EAE symptoms, from a median of 10 days to 20. s

EAE episodes	Phys. rec. (%)	Dis. reg. (%)
Mortality rate	14.2	36.4
0 EAE episodes	0.2	0.0
1 EAE episode	99.6	36.8
2 EAE episodes	0.2	2.2
3 EAE episodes	0.0	6.2
4 EAE episodes	0.0	7.4
5 EAE episodes	0.0	10.6
6 EAE episodes	0.0	10.2
7 EAE episodes	0.0	8.4
8 EAE episodes	0.0	6.6
9 EAE episodes	0.0	5.0
≥ 10 EAE episodes	0.0	6.6


(a) The mortality rate and total number of autoimmune episodes experienced by physiological recovery (Phys. rec.) and disabled regulation (Dis. reg.) groups.



(b) Cumulative distribution plot of the duration of EAE episodes for physiological recovery group.



(c) Cumulative distribution plot of the duration of each EAE episode for regulation disabled group.



(d) Cumulative distribution plot of the duration of remission from EAE for regulation disabled group. Note that remissions are numbered following the first incidence of EAE; the period between immunization and the initial onset of autoimmune symptoms is not shown. Further, by definition, a remission must end in another episode of EAE. This is to exclude those that are terminated by the end of observation, which skew the data.

Figure 7.20: Analyses of number of EAE episodes and remissions, and their durations for physiological recovery and regulation disabled groups.

7.3 Investigation of CD8Treg mediated regulation

The induction of apoptosis in effector CD4Th1 cells by CD8Tregs is the central mechanism in mediating recovery of EAE [Beeston *et al.* 2010, Kumar *et al.* 1996]. As demonstrated in the previous investigation, its abrogation results in substantially increased autoimmune severity. This section details experimentation carried out into the nature of this regulatory pathway, examining the robustness of the pathway to perturbation, and thereby identifying the points at which switching behaviours take place. Since this regulatory pathway is critical to expediting recovery from EAE, understanding its points of failure is an important component in appreciating how prolonged autoimmunity might manifest, and identifying potential medical interventions. The *in silico* experiments conducted here are impossible to perform in the real-world domain.

Two aspects of the regulatory mechanism are investigated here: the efficacy of CD8Treg cells in inducing apoptosis in target effector CD4Th1 cells, and the mean duration of time that effector CD4Th1 cells express Qa-1 for.

By default, a successful binding between an effector CD8Treg and a Qa-1 expressing CD4Th1 cell will always result in the CD8Treg apoptosis of the CD4Th1 cell, reflecting the fact that the simulation's *CD8Treg_cd8TregToCD4ThelperSpecificityDropOff* parameter is set to 1.0 (100%). This parameter value assumes that CD8Treg cells may bind with equal avidity to a Qa-1 expressing CD4Th1 cell as they would a Qa-1:CDR1/2 expressing DC. There exists no domain specific knowledge to underpin this assumption, and as such the effects of alternative values are investigated in section 7.3.2.1. These investigations may be interpreted in terms of the real-world domain, projecting how interference with the regulatory pathway's efficacy influences the recovery from EAE.

CD4Th1 cells express Qa-1 for around 8 hours following their differentiation into effector cells. Only during this time are they susceptible to regulation by the CD8Treg population. Section 7.3.2.2 analyses how reducing or extending this period of time effects the simulation's behaviour.

Additionally, a further investigation into the locations in which CD4Th1 cells are apopposed by the CD8Treg population is conducted. Section 7.2.2.5, above, deduced that the majority of such regulation takes place in the circulatory system. It is however possible that this figure is highly influenced by the length of time that T cells remain in the circulatory system. By default this value is 5 hours, a significant portion of the time that CD4Th1 cells express Qa-1 for following differentiation into effector cells. As detailed in chapter 5, this figure is arbitrarily defined. Domain specific knowledge dictates that T cells spend around 30 min in the circulatory system before migrating to another compartment, however since only a small subset of all possible locations a T cell may migrate to are represented in the simulation, it was felt that assigning only 30 minutes to this parameter would result in T cells migrating into compartments directly related to EAE at an abnormally high rate. Section 7.3.2.3 examines the locations in which CD8Treg mediated CD4Th1 apoptosis takes place whilst varying the duration of time required for T cells to migrate through the circulatory system. Thereby, this experimentation explores the extent to which the finding that this regulation occurs primarily in the circulatory system relies on an arbitrarily assigned parameter value.

7.3.1 Experimental methodology

Experimentation into the efficacy of the regulatory pathway, the probability that a successful binding between a CD8Treg and a Qa-1 expressing CD4Th1 cell will lead to the apoptosis of the CD4Th1 cells, is conducted through perturbation of the *CD8Treg_cd8TregToCD4ThelperSpecificityDropOff* simulation parameter. By default this

parameter is set to 1.0, representing a probability of 100%. It is assigned a series of values between 0.0 and 1.0, with 500 simulation executions being conducted at each value. For each set of 500 executions, the distribution of the number of CD4Th1 cells residing in the system at 40 days, the mean maximum EAE score attained throughout simulation executions, and the mean level of EAE remaining at 40 days are extracted. These metrics are contrasted with similar metrics obtained using the default parameter value. In doing so, the point at which the efficacy of the regulatory pathway is sufficiently reduced so as to constitute a significant change in simulation behaviour is identified. For the number of CD4Th1 cells residing in the system at 40 days, this is ascertaining through use of the *A* test [Vargha & Delaney 2000], assuming ‘large’ differences of ≥ 0.71 and ≤ 0.29 to be scientifically significant. For the two EAE severity metrics, changes of ± 1.0 in mean score are assumed to be significant, as indicated by the domain expert.

Experimentation into the mean duration of Qa-1 expression by effector CD4Th1 cells is achieved through perturbation of the *Th1Polarization_mhcUnExpressionDelay-Mean* simulation parameter. By default this parameter is assigned a value of 8 hours, and this is perturbed to a series of values ranging from 0 to 24 hours. As with the above experimentation into the efficacy of the regulatory pathway, the points at which perturbation constitutes significant changes in simulation behaviour are calculated using metrics of: the number of CD4Th1 cells residing in the system at 40 days, the maximum mean EAE score attained at any point during simulation, and the level of EAE remaining at 40 days. Determination of significant changes in simulation behaviour is as above. Metrics are extracted from 500 simulation executions at each parameter value.

The last set of experimentation concerns the locations in which CD8Treg mediated apoptosis of Qa-1 expressing CD4Th1 cells takes place, and how these locations depend on the length of time that T cells require to migrate through the circulatory system. This migratory time is adjusted through perturbation of the simulation’s *Circulation-timeToCrossOrgan* parameter. By default this parameter is set to 5 hours, and through this experimentation it is reassigned values ranging from the default value to 30 minutes. The dimensions of the circulatory compartment have been altered to facilitate this experimentation. Ordinarily, the compartment is 40 grid spaces high, and 62 wide. T cells are able to move a single grid space in any direction once every $7\frac{1}{2}$ minutes. Hence, at 40 grid spaces high, the minimum time required for a T cell to migrate through the circulatory system is 5 hours. The dimensions of the circulatory system have been adjusted to 4 grid spaces in height, and 620 in width. These revised dimensions allow T cells to migrate through the circulatory system in at least 30 min, and maintain an area identical to that of the default circulatory system dimensions. Once more the points at which significant changes in simulation behaviours occur are deduced. This is performed for five aspects of simulation behaviour: the total number of CD4Th1 cells induced into apoptosis by the CD8Treg population at 50 days, and the total number of CD4Th1 cells induced into apoptosis in each of the circulatory, CLN, SLO, and spleen compartments at 50 days. 500 simulation executions are performed for each parameter value, and 500 values for each of these metrics are extracted accordingly. These are contrasted to similar metrics derived using the default parameter value using the *A* test. Once more, ‘large’ differences, indicated as ≥ 0.71 and ≤ 0.29 , are assumed to be scientifically significant.

7.3.2 Results

Results are presented here in three subsections, examining the efficacy of the regulatory pathway, the duration of Qa-1 expression by effector CD4Th1 cells, and lastly challenging the finding that the majority of CD4Th1 cells are induced into apoptosis by the CD8Treg population in the circulatory system.

7.3.2.1 Efficacy of the regulatory pathway

Results indicate that there is considerable redundancy in the ability of the CD8Treg population to induce apoptosis in Qa-1 expressing CD4Th1 cells. Figure 7.21 depicts the effector T cell dynamics and mean progression of EAE using *CD8Treg_cd8TregToCD4HelperSpecificityDropOff* values of 100%, 5% and 0%. The regulatory pathway is able to substantially reduce the number of CD4Th1 cells in the system when set to just 5% efficacy. Figure 7.22 shows the number of neurons apoptosed per hour, and a cumulative count of CD4Th1 cells apoptosed over time. It supports the observation that 5% efficacy substantially reduces CD4Th1 population size by 40 days; with this twenty-fold reduction in regulatory efficacy resulting in only one eighth less the total number of CD4Th1 cells apoptosed at 40 days.

Figure 7.23a summarises the number of CD4Th1 cells remaining at 40 days of simulation execution, which serves as an indication of how effective the regulatory pathway is in reducing autoimmune T cell activity. Box and whiskers are coloured black to indicate the default parameter value (100%), and red or blue to indicate that a scientifically significant change either has or has not occurred, as measured by the *A* test. The efficacy of the regulatory pathway may be reduced to 30% without causing a significant change in the number of CD4Th1 cells remaining at 40 days.

Figure 7.23b shows how the mean maximum EAE score and mean EAE score at 40 days are influenced by altering *CD8Treg_cd8TregToCD4HelperSpecificityDropOff* between 0 and 100%. The domain expert has indicated that changes of ± 1.0 in mean EAE severity may be considered significant. Although no significant change in behaviour for the maximum EAE metric is observed under this parametric manipulation, a significant increase in EAE experienced at 40 days occurs when *CD8Treg_cd8TregToCD4HelperSpecificityDropOff* is reduced to 0%.

The *CD8Treg_cd8TregToCD4HelperSpecificityDropOff* parameter is a simulation abstraction of the interaction between CD8Treg and CD4Th1 cells, and mediated through a variety of surface molecules, including MHC:TCR and adhesion molecules. The present results suggest that the typical binding affinity between these cells may be considerably less than the default level of 100%, and that as little as 2% is sufficient to maintain effective regulation of the CD4Th1 population, measured in terms of change in EAE severity at 40 days.




Figure 7.21: Median effector T cell population dynamics and mean progression of EAE, derived from *CD8Treg_cd8TregToCD4ThelperSpecificityDropOff* values of 100%, 5% and 0%.






Figure 7.22: Neurons apoptosed per hour, and cumulative CD4Th1 cells apoptosed over time, derived using *CD8Treg_cd8TregToCD4ThelperSpecificityDropOff* values of 100%, 5% and 0%.



(a) Box and whisker plot of the number of CD4Th1 effector cells residing in the system at 40 days following immunization, for various efficacies of regulation. Efficacy of regulation is defined as the probability that a binding between an effector CD8Treg and a Qa-1 expressing CD4Th1 cell will lead to the former inducing apoptosis in the latter. Boxes coloured blue indicates no significant change with respect to the default value of 100%, whereas red indicates that a significant change has occurred.



(b) The mean maximum EAE severity score reached at any point, and the mean EAE severity score experienced at 40 days, obtained for various efficacies of regulation. Significant changes are assumed as ± 1.0 from the mean severity score for the default efficacy of 100%

Figure 7.23: Box and whisker plot of *CD4Th1 at 40 days*, and *max EAE* and *EAE at 40 days* responses for various parameter values of *CD8Treg_cd8TregToCD4ThelperSpecificityDropOff*.

7.3.2.2 Duration of Qa-1 expression by CD4Th1 cells

CD4Th1 cells are susceptible to being regulated only whilst expressing Qa-1 molecules, which they do for a mean duration of 8 hours following their differentiation into effector cells. The effect of manipulating this duration is investigated *in silico*. A CD4Th1 cell differentiating into an effector cell selects a duration for which it will express Qa-1 from a normal distribution, the mean and $2 \times$ the standard deviation of which are specified by the *Th1Polarization_mhcUnExpressionDelayMean* and *Th1Polarization_mhcUnExpressionDelayStdDev* simulation parameters respectively.

Simulation behaviours using *Th1Polarization_mhcUnExpressionDelayMean* values between 0 and 24 hours have been obtained. Figure 7.24 shows the effector T cell dynamics and mean progression of EAE for parameter values of 8 (the default), 3, and 0 hours. Figure 7.25 shows the number of neurons apoptosed per hour, and a cumulative count of the number CD4Th1 cells apoptosed over time, for the same parameter values. The effect of this parametric adjustment on the number of CD4Th1 cells remaining at 40 days, the maximum mean EAE score attained during simulation execution and the mean EAE severity score at 40 days is summarised in figure 7.26. It is observed that increasing the duration of Qa-1 expression has no significant effect on the capacity of the CD8Treg population to regulate the CD4Th1 population. Figure 7.26a indicates that reducing the mean duration of Qa-1 expression to 6 hours causes a significant increase, as measured through the *A* test, in the number of effector CD4Th1 cells remaining in the simulation at 40 days. The graph suggests that this effect is emphasized at 4 hours of expression or less. Assuming a change of ± 1.0 in mean EAE progression to be indicative of significance deviation in simulation behaviour, as specified by the domain expert, figure 7.26b shows that no significant change in the maximum EAE level attained during simulation behaviour results from adjusting this parameter between 0 and 24 hours. A reduction from the default value of 8 hours to 0 hours does however constitute a significant increase in the severity of EAE experienced at 40 days.




Figure 7.24: Median effector T cell population dynamics and mean progression of EAE, derived from *Th1Polarization_mhcUnExpressionDelayMean* values of 8, 3 and 0 hours.






Figure 7.25: Neurons apoptosed per hour, and cumulative CD4Th1 cells apoptosed over time, derived using *Th1Polarization_mhcUnExpressionDelayMean* values of 8, 3 and 0 hours.



(a) Box and whisker plot of the number of CD4Th1 effector cells residing in the system at 40 days following immunization, for various mean durations of Qa-1 expression by such cells. Boxes coloured blue indicate no significant change with respect to the default value of 8 hours, whereas red indicates that a significant change has occurred.



(b) The mean maximum EAE severity score reached at any point, and the mean EAE severity score experienced at 40 days, for various durations of Qa-1 expression by effector CD4Th1 cells. Significant changes are assumed as ± 1.0 from the mean severity score for the default duration of 8 hours.

Figure 7.26: Box and whisker plot of *CD4Th1 @ 40 days*, and the mean *max EAE* and mean *EAE at 40 days* responses, obtained for various values of *Th1Polarization_mhcUnExpression-DelayMean*.

7.3.2.3 Challenging the location of CD4Th1 regulation

Section 7.2.2.5 established that the majority of CD4Th1 apoptosis induction the CD8Treg population occurs in the circulatory system. In contrast, very little occurs in the spleen or the CLN. As detailed in section B.1, the length of time that T cells spend migrating through the circulatory system is set to 5 hours. Whereas migratory times for the SLO, CLN and spleen are based on domain specific knowledge, the length of time that T cells spend in the circulatory system is arbitrarily defined. [Kindt *et al.* 2007] indicates that T cells spend 30 min in the circulatory system before migrating into another compartment, however, since the EAE simulation represents only compartments that are integral to EAE, adopting a figure of 30 min would substantially bias T cell migrations towards compartments directly involved in EAE, hence the figure of 5 hours was adopted.

Effector CD4Th1 cells are only susceptible to regulation whilst they express Qa-1, which they do for a mean period of 8 hours following differentiation into effector cells. It is possible that the dominance of the circulatory system as a site for regulatory action is an artifact of the parameter value that dictates migratory times for this compartment. Since the parameter value is not underpinned by domain specific knowledge, the projection that this compartment is the main site of regulatory action should be challenged before being accepted as representative of the underlying system. The present section investigates how the locations of CD4Th1 apoptosis induction by CD8Treg cells changes with differing times required for T cells to migrate through the circulatory system. A range of values is investigated, from 30 min to the baseline value of 5 hours. The effect on the total number of CD4Th1 cells induced into apoptosis in each compartment are reported in table 7.1.

It is observed that altering the migratory time of the circulatory system within the range of 1 and 5 hours has no significant effect on the total number of CD4Th1 cells induced into apoptosis by the CD8Treg population, as measured by the *A* test. Reducing the migratory time to 30 min causes a significant increase the total number of CD4Th1 cells apoptosed throughout the system. Figure 7.27 illustrates how this aspect of simulation behaviour for the extreme values of 30 min and 5 hours.

Figure 7.28 depicts cumulative counts of the number of Th1 cells apoptosed by the CD8Treg population in each of the simulations compartments, for different migratory times through the circulatory system. Here the effect of altering migration times is substantial; a value of 5 hours results in the circulatory system being the predominant site of regulatory activity, however when reduced to 30 min regulation occurs almost equally in each of the circulatory, CLN and spleen compartments. The significance of this difference is reported in table 7.1: reducing migratory duration to 3 hours significantly increases the total number of CD4Th1 cells apoptosed in the spleen; a reduction to $2\frac{1}{2}$ hours induces significant increases in the CLN and SLO compartments; and lastly, a reduction to $1\frac{1}{2}$ hours significantly reduces the number of CD4Th1 cells apoptosed in the circulatory system. It is of note that the figures for the circulatory system are the most robust to alteration of this parameter, and this most likely reflects the considerable variation in this data amongst the 500 simulation executions from which it is extracted.

The reduction in regulatory action occurring in the circulatory system and the increase occurring in the spleen, resulting from reduced migratory times for the circulatory compartment, may be expected: effector CD4Th1 cells will reach the spleen and CNS compartments more rapidly and spend a greater proportion of the time for which they express Qa-1 there. However, the increase in regulatory action occurring in the CLN is surprising; effector T cells cannot migrate into this compartment from

Response	Time required to migrate through circulatory system (hours)							
	$\frac{1}{2}$	1	$1\frac{1}{2}$	2	$2\frac{1}{2}$	3	4	5
Cumulative	0.71*	0.67	0.64	0.62	0.60	0.58	0.53	0.50
Circulatory	0.00*	0.05*	0.18*	0.31	0.38	0.43	0.46	0.50
CLN	0.96*	0.91*	0.84*	0.77*	0.72*	0.68	0.57	0.50
SLO	0.97*	0.92*	0.86*	0.78*	0.74*	0.68	0.56	0.50
Spleen	1.00*	1.00*	1.00*	0.99*	0.96*	0.90*	0.69	0.50

Table 7.1: *A* test scores indicating the difference between total, and compartmental, regulatory-induced Th1 apoptosis for different migratory times spent by T cells in the circulatory system. In each case distributions for the indicated times are compared with that of 5 hours. Scores marked * are assumed significant, comprising a ‘large’ magnitude of effect.

the circulatory system. Regulation occurring there must result from the interaction between recently-differentiated effector CD4Th1 and CD8Treg cells. The reduced migratory time of the circulatory system results in increased effector CD4Th1 infiltrates in the CNS, as illustrated in figure 7.29. This in turn will increase the rate of neuronal apoptosis, and the number of CD4Th cells entering apoptosis in the CNS, which in turn increases the number of DCs migrating from the CNS to the CLN that present Fr3, CDR1/2 and MBP peptides. These DCs will prime greater numbers of CD4Th, CD4Treg and CD8Treg populations, as shown in figures 7.30, 7.31 and 7.32. The increased numbers of T cells priming in the CLN may increase the incidences of effector Qa-1 expressing CD4Th1 cells interacting, and being apoptosed by, effector CD8Treg cells. Figures 7.30, 7.31 and 7.32 also show increased T cell priming in the SLO compartment, and hence the increase in CD4Th1 regulation occurring in this compartment may be explained through similar means, increased incidences of effector Qa-1 expressing CD4Th1 cells interacting with effector CD8Treg cells.

7.3.3 Summary

This section has reported *in silico* experimentation into the regulatory pathway by which CD8Treg cells induce apoptosis in the Qa-1 expressing effector CD4Th1 population, conducted using the ARTIMMUS simulation. The experimentation has examined the effect of altering the efficacy of the regulatory pathway at the cellular level: the probability that a successful binding of a CD8Treg cell to Qa-1 expressing effector CD4Th1 cell will result in the apoptosis of the latter as induced by the former; and the mean duration of Qa-1 expression by effector CD4Th1 cells. The experimentation carried out makes use of both the *A* test and domain expert informed metrics to discern when significant changes to simulation behaviour have occurred. Simulation behaviours under study relate to the ability of the regulatory pathway to effectively reduce autoimmune activity, as measured through the number of effector CD4Th1 cells residing in the system at 40 days, and both the maximum severity of EAE experienced during simulation execution and that remaining at 40 days.

These experiments have exposed the considerable redundancy that exists within this regulatory pathway. The efficacy of the pathway could be reduced from the default value of 100% to 0% before a domain expert-indicated significant change in simulation behaviour takes place. Increasing the duration of Qa-1 expression by effector CD4Th1 cells has no significant effect on recovery from autoimmunity. This duration could be reduced from the default value of 8 hours to 0 before domain expert-indicated significant increases in the mean level of EAE experienced at 40 days takes place.

It is noteworthy that in neither of these experiments was the maximum mean EAE




Figure 7.27: Cumulative count of the number of effector Th1 cells induced into apoptosis by the CD8Treg population, for various times required to migrate through the circulatory compartment.




Figure 7.28: Cumulative count of effector Th1 cells induced into apoptosis by the CD8Treg populations, by compartment, for various times required to migrate through the circulatory compartment.




Figure 7.29: The number of effector CD4Th1 and CD4Th2 cells in the CNS compartment over time, for various times required to migrate through the circulatory compartment.




Figure 7.30: Cumulative count of CD4Th cells primed by compartment, for various times required for a T cell to migrate through the circulatory system.




Figure 7.31: Cumulative count of CD4Treg cells primed by compartment, for various times required for a T cell to migrate through the circulatory system.




Figure 7.32: Cumulative count of CD8Treg cells primed by compartment, for various times required for a T cell to migrate through the circulatory system.

score amongst 500 simulations significantly altered. It may be seen in figure 7.1b, wherein the regulatory pathway has been completely disabled, that there is a substantial reduction in effector CD4Th1 cell number following peak autoimmune activity. These results indicate that regulatory function does not engage rapidly enough to significantly alter the maximum level of EAE symptoms experienced by simulations.

The last set of experimentation carried out in this section challenges the finding that the majority of CD4Th1 apoptosis induction by the CD8Treg population occurs in the circulatory system, first reported in section 7.2.2.5. The duration of time required for a T cell to migrate through the circulatory system is 5 hours, a substantial proportion of the total time for which Qa-1 is expressed by effector Th1 cells. This figure is not underpinned by domain specific knowledge, and as such the projection's validity is challenged. It is found that, although the total number of effector CD4Th1 cells induced into apoptosis by the CD8Treg population is largely independent of this parameter, the proportion of CD4Th1 cells that are apoptosed in the circulatory system is not. Reduction of the time required for a T cell to migrate through the circulatory system is found to reduce the number of CD4Th1 cells regulated in the circulatory system, and increase the number regulated in the CLN, spleen and SLO compartments. This result stresses the importance of challenging design assumptions when interpreting simulation-derived projections of real world behaviour. ARTIMMUS cannot be used to project the proportion of CD8Treg induced CD4Th1 apoptosis occurring in the real system unless the time that T cells require to migrate through the circulatory system is better specified, or the alternative compartments into which T cells may migrate is better represented.

7.4 *In silico* splenectomy

The present section details the employment of ARTIMMUS in exploring the role of the spleen in recovery from EAE. The important role of the spleen in the recovery from EAE is highlighted in [Ben-Nun *et al.* 1980], who perform experimentation with rats. Further *in silico* investigation provides insight into the processes underpinning its role in EAE. A splenectomy experiment is engineered into ARTIMMUS, and through contrasting these results with a control (non-splenectomised) group and changes to EAE progression, T cell dynamics, regulatory capacity of Treg populations, and states

of activation of APC populations are examined.

The influential, but complex role of the spleen in the recovery from EAE is highlighted by [Ben-Nun *et al.* 1980], who perform splenectomies on rats of varying ages, between 1 and 15 months, and induce EAE within a month thereof. Three experiments are performed, with an experimental and two control groups in each case. The experimental groups are splenectomised at $1\frac{1}{2}$, 3 or 12 months of age. Control groups are not splenectomised, though the ‘sham’ group does undergo surgery. Both experimental and control groups are then induced into EAE one month later, at $2\frac{1}{2}$, 4 or 13 months of age. The authors’ data is reproduced in table 7.2.

Rats induced into EAE at $2\frac{1}{2}$ months experience acute EAE, and all but the ~25% that die experience recovery by day 17, regardless of whether or not they were splenectomised a month before EAE induction. 2 out of 8 mice perish in the control group, in contrast with 3 out of 8 in the splenectomy group.

The 4 month control group experiences a ~10% mortality rate, which is raised to 20% in the experimental group splenectomised one month earlier. All control mice that did not perish experience full recovery by day 17. In contrast, all the splenectomised mice experience chronic symptoms until observation ends at 32 days. The authors define chronic symptoms as unremitting paralysis from its onset until the end of observation at 32 days.

The mortality rate in the 13 month group is similar regardless of splenectomy, being around 40%. A few control mice experience recovery around day 20, but all then experience relapsing symptoms at day 24. The remaining control mice experience chronic symptoms throughout experimental observation. Of the 60% of splenectomised mice that did not die, all experience chronic symptoms.

Ben-Nun *et al.*’s experiments demonstrate the increasing severity of EAE with increasing age at time of induction. In all cases, splenectomy results in increased EAE severity. The authors note the role of regulatory action on encephalitogenic T cell responses in the recovery from EAE. They suggest that the spleen is involved in this suppression. However, though they go on to suggest that “the populations of cells mediating recovery are not always resident in or dependent on the spleen”, with reference to how splenectomy does not always interfere with recovery.

Data from ARTIMMUS, section 7.2 above, indicates that significant priming of Treg populations occurs in the spleen, being around 70% of all CD4Treg cells and 75% of all CD8Treg cells, figures 7.6a and 7.6b respectively. The majority of the remaining Treg populations are primed in the CLN compartment. Figure 7.5a shows that no CD4Th cell priming occurs in the spleen. These simulation-derived results suggest that the role of the spleen in EAE is as a major site for Treg priming. This is investigated further through the engineering of a splenectomy experiment within ARTIMMUS.

Age (months)	Treatment	Acute EAE			Recovery		Chronic EAE Incidence	Relapse		Recovery from relapse Incidence	Grade of pathology
		Day of onset	Total	Incidence Mortality	Day	Incidence		Day	Incidence		
$2\frac{1}{2}$	None	11	8/8	2/8	17	6/6	0/6	—	0/6	—	++
	Sham	12	8/8	2/8	16	6/6	0/6	—	0/6	—	++
	Splenectomy	11	8/8	3/8	17	5/5	0/5	—	0/5	—	++
4	None	12	8/8	1/8	17	7/7	0/7	—	0/7	—	++
	Sham	12	10/10	1/10	17	9/9	0/9	—	0/9	—	++
	Splenectomy	11	10/10	2/10	—	0/8	8/8	—	—	—	+++
12	None	11	10/10	3/10	20	2/7	5/7	24	2/2	0/2	++++
	Sham	12	10/10	4/10	19	2/6	4/6	24	2/2	0/2	++++
	Splenectomy	11	10/10	4/10	—	0/6	6/6	—	—	—	++++

Table 7.2: Acute, chronic or relapsing EAE in rats as a function of age or splenectomy. Data replicated from [Ben-Nun *et al.* 1980], table 2. The indicated age in months is time of immunization for EAE, treatment was performed 1 month prior to immunization. ‘Sham’ indicates surgery, but no splenectomy was performed. Chronic EAE is defined as unremitting paralysis from onset until termination of the experiment at 32 days post immunization. Pathology is graded on a scale of + to +++, and is based on post-experimental examination of the animal’s central nervous system.

7.4.1 Experimental methodology

In order to engineer a splenectomy experiment in the ARTIMMUS simulation, modifications have been made to the spleen compartment. The splenectomised spleen compartment contains no dendritic cells, and has no capacity for cells to reside within it; with the exception of apoptotic CD4Th cells, all cells entering the compartment exit it immediately. In the case of apoptotic CD4Th cells, these cells are removed from the simulation. Ordinarily, a large number of CD4Th cells are phagocytosed in the spleen, as evidenced by the number of DCs that present to Treg cells in this compartment. Apoptotic CD4Th cells cannot enter the CNS or the SLO, and unless migrating from the CNS compartment, cannot enter CLN. Hence, the number of apoptotic CD4Th cells residing in the circulatory system will continually increase, since they cannot be phagocytosed. The lack of removal of these cells from the simulation will impact simulation performance, interfering with other cells that attempt to migrate through the circulatory compartment. In order to prevent this simulation artifact, the splenectomised spleen removes apoptotic CD4Th cells from the simulation when they attempt to migrate through it.

In splenectomy experiments, the splenectomy spleen compartment replaces the standard compartment in ARTIMMUS, and is connected to the other simulation compartments as before.

To facilitate analysis of the effect that splenectomy has on simulation behaviour, two experimental groups of simulation executions have been compiled, termed the *control* and *splenectomy* groups. Both groups comprise 500 simulation executions using unique seeds. All simulation executions in each group are obtained using the same parameter values, being the baseline values outlined in section B.1. The control group simulations are derived from standard ARTIMMUS executions, whereas the splenectomy group simulations are obtained by substituting ARTIMMUS's spleen compartment with the splenectomy variety described above.

Where comparisons between the control and splenectomy groups are made using the *A* test [Vargha & Delaney 2000], scores comprising 'medium' and 'large' effect magnitudes are indicated. Medium effects are defined as being ≥ 0.64 or ≤ 0.36 , large effects are defined as ≥ 0.71 or ≤ 0.29 . Both medium and large effects are assumed to be statistically significant.

7.4.2 Results

It is important to note that ARTIMMUS was calibrated against mouse data from the Kumar lab, whereas the experiments performed by Ben-Nun *et al.* are conducted on rats. The metric used to grade EAE severity by Ben-Nun differs from that of the Kumar lab, against which the simulation's own EAE grading system was calibrated. Ben-Nun *et al.*'s metric directly examines the damage to the CNS once the experiment has been terminated, whereas that used by the Kumar lab is based on the degree of paralysis observed. Hence, direct quantitative comparison between *in silico* and *in vivo* results are not possible. Ben-Nun *et al.* do, however, report mortality rates and period of autoimmune symptoms, noting days of onset and recovery, and incidence of relapsing and chronic EAE in a manner that facilitates qualitative comparison with the results gained from ARTIMMUS.

Results are presented in the following manner. Section 7.4.2.1 considers the effect that splenectomy has on the initial onset of EAE autoimmune symptoms, and qualitatively contrasts these results with those of Ben-Nun *et al.* [Ben-Nun *et al.* 1980]. Section 7.4.2.2 observes simulation behaviour for 200 days and reports on the relapsing

nature of autoimmunity following splenectomy, in contrast to the mono-phasic disease of the control group. Section 7.4.2.3 describes the effect of splenectomy on the simulation’s T cell populations, noting the significant reduction in effector CD4Treg and CD8Treg population sizes. Section 7.4.2.4 reports on how this reduction in effector Treg population size leads to a delayed ability of splenectomised mice to completely eradicate effector CD4Th1 populations, which in turn leads to relapsing autoimmune symptoms. Lastly, section 7.4.2.5 considers the effect that splenectomy has on the APC populations in the CLN and the CNS.

7.4.2.1 Splenectomy impact on initial EAE episode

Figure 7.34 depicts the mean progression of EAE among the 500 simulations in each of the control and splenectomy groups. The mean maximum EAE level is marginally increased from 2.5 in the control group to 2.75 in the splenectomy group. Splenectomy increases the level of EAE at 40 days, once the severity of EAE has plateaued, from under 0.7 to 1.0. Table 7.3 indicates the proportions of simulations that reach particular maximum severity scores over time, for control and splenectomy groups. Splenectomy generally increases the maximum severity of autoimmunity experienced by simulations, with the proportions of simulations experiencing levels 0 to 2 decreasing, and levels 3 to 5 increasing.

Figure 7.35 shows the proportion of each experimental group that experiences each severity of EAE over the first 50 days following immunization. The peak of autoimmune symptoms occurs around day 15 in both the control and splenectomy groups. Contrasting figures 7.35a and 7.35b reveals that splenectomy tends to slow recovery from EAE. The control group experiences complete recovery by day 23 whilst the splenectomy group takes until day 30, after which a very small proportion of simulations continue to exhibit level 1 autoimmunity.

Figure 7.33a reports the magnitude of effect that splenectomy has on EAE progression at various points in time by using the *A* test [Vargha & Delaney 2000] to contrast distributions of EAE scores between control and splenectomy groups.

Score	Control (%)	Splenectomy (%)
5	14.2	19.4
4	22.0	26.4
3	24.2	26.4
2	15.8	13.6
1	23.6	14.2
0	0.2	0.0


Table 7.3: The proportions of simulations experiencing each possible maximum EAE severity score, for control and splenectomy groups.

The *max EAE* response is time-independent, it contrasts the distributions formed by the maximum severity of EAE experienced by each simulation at any point in time. As noted above, splenectomy is observed to increase the severity of EAE at peak autoimmune behaviour, however the *A* test confirms that this increase is not significant. The observation that splenectomy delays recovery from EAE is also confirmed by the *A* test, and shown to be a significant effect between days 19 and 23.

It is interesting to note that despite the mortality rate increasing from 14% to 19% under splenectomy, this difference is not considered significant by applying the *A* test to EAE severity at day 40, table 7.33a.

Response	<i>A</i> test score
<i>Max EAE</i>	0.57
<i>EAE at 16d</i>	0.55
<i>EAE at 18d</i>	0.62
<i>EAE at 20d</i>	0.67 ⁺
<i>EAE at 25d</i>	0.55
<i>EAE at 30d</i>	0.53
<i>EAE at 40d</i>	0.53

(a) *A* test scores of the difference between EAE severities experienced by the control and splenectomy group at select points in time. Scores marked * are indicative of ‘large’ differences between response distributions, being ≤ 0.29 or ≥ 0.71 . Those marked + are indicative of ‘medium’ differences, ≤ 0.36 or ≥ 0.64 .



(b) Plot of *A* test scores between EAE severities in the control and splenectomy groups over time. The ‘medium’ and ‘large’ effect magnitude scores are indicated.

Figure 7.33: Magnitude of effect measures of the difference in EAE severities amongst control and splenectomy groups at various times following immunization.

Two further measures of significant changes in simulation behaviour have been indicated by the domain expert; changes of ± 1.0 in the maximum mean progression of EAE, and the level of EAE at 40 days may be assumed to be significant. As may be seen in figure 7.34, the splenectomy procedure again fails to constitute significant differences in terms of these metrics.

Although the metric used by Ben-Nun *et al.* to grade EAE severity differs from that of the Kumar lab and ARTIMMUS, there are several parallels between the results from ARTIMMUS and those from Ben-Nun *et al.* Ben-Nun *et al.* found that splenectomised rats tended to experience relapses of autoimmune symptoms, or experienced chronic symptoms until the termination of their experiments at day 32. Recovery from EAE symptoms in ARTIMMUS is also delayed following splenectomy, though almost all simulations have experienced full remission by day 32. Ben-Nun *et al.* note that the mortality rate of rats following splenectomy was, in all cases, increased following splenectomy. This too is consistent with findings from ARTIMMUS, where splenectomy increases the mortality rate from 14% in the control group to 19% in the splenectomy group.

7.4.2.2 Splenectomy induces long term relapses in autoimmunity

Figure 7.36 shows the proportion of control and splenectomy simulations experiencing each grade of EAE severity over 200 days. Figure 7.37 presents an analysis of the relapses and remissions of into and out of autoimmune symptoms amongst control and

splenectomy groups. The proportion of simulations that experience varying numbers of relapses are presented. Cumulative distribution plots show the distribution of durations for each relapse and remission amongst experimental populations.

Over 99.5% of simulations in the control group experience a single phase of EAE autoimmunity. 14% of these simulations perish, whilst the remaining simulations recover by day 23 and (with the exception of one 1 of 500) experience no relapses over the 200 day observation period. In contrast, a small proportion of the splenectomy group continues to experience level 1 autoimmune symptoms throughout the 200 day observation period. Figure 7.37c establishes that proportion is constituted through relapsing autoimmune symptoms amongst the population, as opposed to single simulations that fail to return to EAE severity 0; the longest single period of initial autoimmune symptoms, in the 81% of simulations that do not perish, is around 20 days.

In contrast with the control group, only 81% of splenectomised simulations experience a single period of clinical symptoms. The remaining 19% of splenectomised simulations all experience relapsing EAE. 13% of simulations experience a single relapse following the first episode of autoimmune symptoms, and 5% experience a second.

Splenectomy increases the duration of clinical symptoms. In the control group EAE lasts as little as 5 days and at most 13, provided the simulation has not perished; periods of 11 to 13 days are the most common. In contrast, the first bout of EAE in splenectomised simulations is rarely less than 9 days, and may be as much as 17, with a median of 12 days duration, excluding those that perish. Successive relapses in splenectomised simulations tend to be shorter in duration, though 80% of any relapse is less than 7 days in duration. Relapsing symptoms do not reach level 5 severity; of the 19% of simulations that perish following splenectomy all do so in their first episode of EAE.

Figure 7.37d depicts cumulative distribution plots of the proportion of splenectomy simulations that experience particular durations of remission before autoimmune symptoms reoccur. This is shown for successive remissions. The data indicates that the first remission is the longest, and that subsequent remissions tend to be shorter. The median duration of the first remission is over 70 days. Half of mice that experience a second incidence of EAE do so after 12 days of remission.

Collectively, this data indicates that the chance of a splenectomised simulation experiencing a relapse of autoimmune symptoms decreases following each relapse, as evidenced by the decreasing proportion of the population that experience increasing numbers of EAE incidence. However, where subsequent relapses do occur, they occur more quickly and last for decreasing periods of time. Figure 7.38 depicts the progression of EAE in four example splenectomy simulation executions. These examples are indicative of the relapsing nature of EAE following splenectomy, mono-phasic EAE examples are not shown. The examples demonstrate that the first episode tends to be the most severe. Thereafter simulations may enter extended periods of remission before clinical symptoms re-present. Relapses are very rarely more severe than level 1, and do not tend to last as long as the initial episode.

7.4.2.3 Impact on effector T cell populations

The nine responses pertaining to effector T cell dynamics, outlined in section 6.2, have been calculated for control and splenectomy groups. Table 7.4 summarises the magnitude of effect that splenectomy has on each of these responses, using the A test to contrast response distributions between control and splenectomy experimental groups.

The results indicate that splenectomy has the biggest effect measurable by the A test on the maximum number of effector CD4Treg and effector CD8Treg cells, which




Figure 7.34: Mean progression of EAE under control and splenectomy experiments.




Figure 7.35: Cumulative counts of the proportion of simulations experiencing each degree of EAE severity, for control and splenectomy experiments.





Figure 7.36: Long-term cumulative counts of the proportion of simulations experiencing each degree of EAE severity, for control and splenectomy experiments.

EAE incidence	Control (%)	Splenectomy (%)
Mortality rate	14.2	19.4
0 EAE episodes	0.2	0.0
1 EAE episode	99.6	80.6
2 EAE episodes	0.2	12.6
3 EAE episodes	0.0	4.8
4 EAE episodes	0.0	1.6
5 EAE episodes	0.0	0.2
≥ 6 EAE episodes	0.0	0.2

(a) The mortality rate and total number of autoimmune occurrences experienced by control and splenectomised simulation groups.



(d) Cumulative distribution plot of the duration of remission from EAE for splenectomy group. Note that remissions are numbered following the first incidence of EAE; the period between immunization and the initial onset of autoimmune symptoms is not shown. Further, by definition, a remission must end in another incidence of EAE. This is to exclude those that are terminated by the end of observation, which skew the data.

Figure 7.37: Analyses of number of EAE incidences and remissions, and their durations for control and splenectomy experimental groups.

experience considerable reductions in size. As indicated by figures 7.41 and 7.42, the spleen is a major site of Treg priming. Its removal explains this reduction in CD4Treg and CD8Treg number at peak autoimmunity.

The number of effector CD4Th1 cells remaining at 40 days is also altered by splenectomy, the *A* test once more indicating the biggest magnitude of effect possible. These effects may be observed in figure 7.39, which depicts the median effector T cell dynamics. Whereas effector CD4Th1 cell number is reduced to zero by day 40 in the control group, splenectomy sees the effector CD4Th1 population persist at under 100 cells in size for over 150 days of observation.

The intensity of the CD4Th1 and CD4Th2 immune responses, as measured by the maximum number of these cells reached at any point in time, are not significantly affected, though a ‘medium’ effect is found in the delay at which the CD4Th1 response peaks. The times at which CD4Treg and CD8Treg immune responses peak are not significantly affected by splenectomy.

Response	<i>A</i> test score
<i>CD4Th1 Max</i>	0.58
<i>CD4Th1 Max Time</i>	0.64 ⁺
<i>CD4Th2 Max</i>	0.44
<i>CD4Th2 Max Time</i>	0.49
<i>CD4Treg Max</i>	0.00*
<i>CD4Treg Max Time</i>	0.38
<i>CD8Treg Max</i>	0.00*
<i>CD8Treg Max Time</i>	0.40
<i>CD4Th1 at 40d</i>	0.99*

Table 7.4: The *A* test scores indicating magnitude of difference in response distributions between splenectomy and control experimental groups. Scores marked * are indicative of ‘large’ differences between response distributions, being ≤ 0.29 or ≥ 0.71 . Those marked + are indicative of ‘medium’ differences, ≤ 0.36 or ≥ 0.64 .

7.4.2.4 Splenectomy reduces regulatory capacity of Treg populations

As demonstrated in figure 7.40, priming of the CD4Th population ceases around day 40 in the control group, however in the splenectomy group priming continues, albeit at a reduced rate, until observation ends at 200 days. Hence, splenectomy reduces the capacity for regulation to counter autoimmune behaviour to such a degree that complete and permanent cessation of autoimmune T cell responses is often not possible. Figure 7.36b demonstrates that a small proportion of the simulation population experience EAE grade 1 severity at any point in time for the full 200 days of observation. As described above, these incidences of clinical symptoms tend to present in relapsing forms.

It has been established in section 6.4.2.1 of the previous chapter that the stimulus for Treg priming is the severity of the autoimmune response. In absence of the spleen, a primary site of Treg priming (figures 7.41 and 7.42) the regulatory immune responses are unable to reach a sufficient size to completely eradicate autoimmune behaviour. Regulation is effective in reducing the size of the CD4Th1 population, however when this population becomes small the capacity of DCs to prime Treg populations reduces, and the CD8Treg population dies out before it is able to completely abrogate the CD4Th1 population. The lack of CD8Tregs allows the CD4Th1 population to once more expand, and in the period taken for DC populations to further prime Treg cells

in response, clinical symptoms re-present. Whilst this pattern may repeat for extended periods of time, figure 7.39c suggests that the CD4Th1 immune response is abrogated completely by day 170 in over 50% of simulations (figure 7.39c depicts the median simulation behaviour).

7.4.2.5 Influence on APC populations

Figure 7.43 shows the states of APC activation in the CNS over 200 days of observation for both control and splenectomy groups. In the control group, the number of immunogenic APCs reaches a peak of 50 at the height of the autoimmune response. This figure is similar in the splenectomy group, however the number of immunogenic APCs does not return to basal levels, instead holding relatively steady at over 20. The peak of autoimmune activity sees around 1000 effector CD4Th1 cells in the system. The observation that 1000 effector CD4Th1 cells induce immunogenic phenotypes in 50 CNS-resident APCs, and that only \sim 100 effector CD4Th1 cells (between days 50 and 100 in splenectomy group) can induce immunogenic phenotypes in \sim 30 APCs, supports the conclusion that there exists a non-linear relationship between CD4Th1 population size and the degree of autoimmunity experienced. This diminishing return on increasing CDTh1 cell number was previously established in section 7.2.2.6.

Figure 7.44, the states of APC activation in the CLN, demonstrates that fewer than 10 immunogenic DCs in the CLN, out of a total number of \sim 95, is sufficient to maintain the CD4Th1 population at levels capable of inducing relapsing clinical EAE symptoms, but often insufficient to prime the number of Treg cells required to abrogate autoimmunity.

It has been established, in section 7.2.2.4 and shown in figure 7.45a, that the end of the CD4Th1 immune response is marked by a transition of the majority immunogenic DCs in the CLN adopting a type 2 polarization in place of type 1. Figure 7.45a reveals that this majority also occurs in splenectomised mice, between days 25 and 55. However, this majority is not sufficient to result in a type 2 deviation of CD4Th cells capable of out-competing the effect of the CD4Th1 population; the CD4Th2 cells do not outnumber the CD4Th1 cells, the number of type 1 polarized DCs in the CLN increases to equal the number of type 2 polarized, and the autoimmune response does not terminate.




Figure 7.38: Example EAE progressions of individual splenectomised simulation executions.




Figure 7.39: The median system wide T cell dynamics, obtained from 500 simulation executions, showing control and splenectomy experiments.




Figure 7.40: Cumulative count CD4Th cells primed in each compartment, for control and splenectomy experiments.




Figure 7.41: Cumulative count CD4Treg cells primed in each compartment, for control and splenectomy experiments.




Figure 7.42: Cumulative count CD8Treg cells primed in each compartment, for control and splenectomy experiments.




Figure 7.43: States of APC activation in the CNS, for control and splenectomy experiments.




Figure 7.44: States of APC activation in the CLN, for control and splenectomy experiments.




Figure 7.45: Polarizations of DCs in the CLN, for control and splenectomy experiments.

7.4.3 Summary

The section has described the use of the ARTIMMUS simulation in investigating the results of a splenectomy experiment. The role of the spleen in EAE has been investigated *in vivo* by [Ben-Nun *et al.* 1980]. Ben-Nun *et al.*'s experiments on rats exposed the spleen's role in aiding recovery from EAE; its removal results in more severe EAE. A similar experiment, the removal of the spleen before the induction of EAE, has been engineered into ARTIMMUS with the aim of gaining insight into its role in EAE.

It is important to note that ARTIMMUS was calibrated against mouse data from the Kumar lab, whereas the experiments performed by Ben-Nun *et al.* were on rats. The metric used to grade EAE severity by Ben-Nun *et al.* differs from that of the Kumar lab, against which the simulation's own EAE grading system was calibrated. Ben-Nun's *et al.*'s metric directly examines the damage to the CNS once the experiment has been terminated, whereas that used by the Kumar lab is based on the degree of paralysis observed. Ben-Nun *et al.* do, however, report mortality rates and period of autoimmune symptoms, noting days of onset, and recovery, and incidence of relapsing and chronic EAE in a manner that facilitates qualitative comparison with the results gained from ARTIMMUS.

The simulation's results are consistent with those observed *in vivo* in that splenectomy results in more severe EAE symptoms. Following splenectomy, autoimmune conditions present more severely and last longer. Splenectomy in ARTIMMUS had considerable effect on the long term recovery from EAE. In absence of splenectomy all mice that do not perish experience complete recovery with no relapsing symptoms over the 200 days of observation. In contrast, splenectomy results in 20% of all mice experiencing at least one relapse of clinical symptoms.

Results from ARTIMMUS demonstrated that the majority of Treg priming occurs in the spleen, and that its removal results in significantly reduced CD4Treg and CD8Treg numbers. This reduced Treg population is able to regulate the CD4Th1 population to the point that the difference in EAE severities between splenectomy and control groups at 40 days is not statistically significant, but is often not able to completely abrogate the CD4Th1 population. This results in the remitting autoimmune symptoms noted above.

It is noteworthy that the effect of splenectomy on the levels of EAE experienced at peak immunity and at various points following immunization was not deemed 'large' through use of the *A* test, though 'medium' effects did occur. The present experimentation has not investigated whether greater magnitudes of effect might be reached with different strengths of immunization for EAE. The immunization protocols used to induce EAE differ between labs, and between experiments conducted by the same lab. An experiment intended to investigate the ability of an intervention to reduce the severity of symptoms might induce more severe EAE than an experiment intended to investigate the potential of an intervention to increase symptoms. The experimentation reported in this section has investigated splenectomy under a single immunization protocol. Section 7.4.2.4 suggested that the inability of splenectomised simulations to completely abrogate the CD4Th1 population was due to the reduced Treg immune responses. The reduced CD8Treg number in splenectomised simulations are able to substantially reduce CD4Th1 cell number, but not before the reduction in CD4Th1 cells results in reduced Treg priming capacity, which once more permits the expansion of the CD4Th1 cell population. Further experimentation with ARTIMMUS could examine the ability of splenectomised simulations to effectively abrogate the CD4Th1 population as a function of immunization strength; a stronger immunization for EAE might result in enhanced autoimmune activity that the reduced Treg populations are

less capable of regulating, hence resulting in increasingly severe and relapsing autoimmunity. Alternatively, it is possible that under a stronger immunization the increased CD4Th1 number might provide additional priming capacity for Treg populations that are able to abrogate the autoimmune response after a single episode of autoimmune symptoms. Such experimentation would provide greater insight into the interaction between autoimmune and regulatory immune responses, and the role of the spleen in recovery from EAE.

7.5 Conclusion

This chapter has reported three strands of *in silico* experimentation into EAE. Section 7.2 has examined the cellular dynamics that underpin autoimmunity and regulation in EAE. It also explored the relapsing nature of EAE in simulations where the regulatory pathway is disabled.

Section 7.3 focuses on various aspects of the regulatory pathway through which CD8Treg cells apoptose CD4Th1 cells. The section challenges a finding of the previous section, that the majority of CD4Th1 apoptosis through regulation occurs in the circulatory system. It is found that this finding critically depends on an arbitrarily assigned parameter value, and as such the finding cannot be accepted as representative of real-world EAE. This demonstrates the importance of appreciating the effect of arbitrarily assigned parameter values. *In silico* predictions that rest on arbitrarily assigned parameter values should be challenged.

Lastly, section 7.4 reports on using ARTIMMUS to gain insight into the role of the spleen, motivated by wet-lab experimentation conducted in the 1980s that revealed splenectomised rats to experience more severe EAE than their control group counterparts. The section predicts that the spleen is a major source of Treg priming, and that its removal results in reduced numbers of Tregs which cannot combat autoimmunity with the same efficiency, though most mice do still recover. Direct objective comparison between the wet-lab data and *in silico* results is not possible, the wet-lab metrics of autoimmunity severity differ from those of the Kumar lab that ARTIMMUS employs. However, the two sets of results do seem consistent with one another, and demonstrate the generality of ARTIMMUS: it is able to replicate phenomena from another model, against which it was not calibrated. This finding serves to increase confidence that simulation results are representative of real-world EAE.

This chapter has addressed research objective 4: to perform novel *in silico* experimentation using the agent-based simulation of EAE. It provides contributions to the field of EAE, and provides a context for exploring how statistical and modelling techniques provide confidence in simulation results being representative of real-world EAE. This is the subject of the following and final chapter of this thesis.

Chapter 8

Discussion

The present chapter concludes the work of this thesis, and commences in section 8.1 with a reflection upon the extent to which the modelling and statistical techniques applied and developed are considered to provide confidence in simulation results being representative of the real-world domain. The last contribution of this thesis is made: a novel meta-heuristic search based technique that is proposed to guide simulation development to appropriate levels of abstraction, and demonstrate this to have been the case. Next, section 8.2 returns to the thesis aim and the research objectives, and provides a chapter-by-chapter summary of content and contributions. Section 8.3 details four strands of further work that have arisen from this thesis. Lastly, section 8.4 presents the concluding remarks of this thesis.

8.1 Reflections on establishing confidence in simulation

This thesis is motivated by the need to establish confidence that *in silico* results are representative of the immune system domains that they intend to capture. The introductory chapter considered four aspects of agent-based *in silico* experimentation in which confidence must be provided. These are:

- Confidence that the simulation’s capture of the cells, processes and interactions of the target domain is satisfactorily accurate, given the investigations that are to be performed on it. These aspects of the simulation have been referred to as the simulation’s *mechanics*.
- Confidence that the simulation is correctly parameterised; that the mechanics have associated with them the correct rates, probabilities, quantities and temporal behaviours. These aspects have been referred to as the simulation’s *parameters*.
- Confidence that *in silico* results are representative of the simulation’s dynamics, and not merely the result of stochasticity in the simulation.
- Understanding of the implications of *in silico* results in terms of the original domain: how do simulation metrics map to the real-world system, and are interesting simulation results merely the result of underspecified parameter values?

The aim of this thesis has been to address these issues of simulation confidence through modelling and statistics, by undertaking a case study in modelling and simulating a complex murine autoimmune disease, experimental autoimmune encephalomyelitis (EAE). This is reflected in the thesis aim:

To apply and develop statistical and modelling techniques for agent-based simulations of immunology, specifically experimental autoimmune encephalomyelitis.

The present section reflects on how the modelling and statistical techniques explored in this thesis establish confidence that the results arising from *in silico* experimentation with ARTIMMUS are representative of real-world EAE. It proposes a novel method for quantifying how well a simulation captures its target domain, and explaining how this may be integrated into the CoSMoS process to demonstrably guide simulation development towards a faithful representation of the target system. The concluding section explores how the various techniques examined in this thesis may add objectivity to the establishment of confidence in simulation.

8.1.1 Explicit domain modelling

The CoSMoS process, reviewed in the introductory chapter, underpins simulation-based experimentation with explicit domain modelling [Andrews *et al.* 2010]. The process promotes simulation-based investigation of complex systems as an inherently interdisciplinary activity, and as such the present case study in EAE has been conducted in collaboration with Dr. Kumar, an immunologist and expert in the disease. The CoSMoS process may be considered to address the first of four aspects of simulation confidence detailed above: establishing confidence that the simulation's mechanics are a faithful representation of the target domain. The domain modelling work conducted in chapter 4 has been highly beneficial in addressing the complexity of EAE and gaining a consistent perspective of its dynamics. The creation of a full specification of the domain has highlighted areas of ambiguity and underspecification in immunology, and the activity has been valuable in guiding interaction with the domain expert. Confidence in the simulation's faithful representation of the domain is achieved through domain expert verification that the dynamics underpinning it are realistic and appropriate. Though time-consuming, it is felt that explicit domain modelling has been highly beneficial to exploring EAE through simulation.

8.1.2 Calibration of simulation mechanics and parameters

The CoSMoS process is inherently iterative; as further scientific questions of the domain are addressed through simulation, the artifacts of the process are updated. The CoSMoS process acknowledges that the first pass of the process may not lead to a simulation whose dynamics are consistent with the target domain, and that this discrepancy can motivate further iterations of the process [Andrews *et al.* 2010]. Indeed, this has been the case in the present study of EAE. However, the process offers no guidance on deciding whether a simulation satisfactorily captures a target-domain, or how to rectify inconsistencies should they be found. Chapter 5 reports the *collaborative calibration procedure*, by which this was achieved through close collaboration between the domain expert and the modeller. The procedure explores the space of both simulation mechanics and parameters in developing the simulation to a point where it satisfactorily reflects the dynamics of real-world EAE under two experimental scenarios. Once more

it is felt that substantial input from the domain expert during development provides confidence that the simulation is appropriately grounded in the real-world domain.

A strength of the procedure is its calibration against multiple real-world experimental *scenarios*, real-world interventions such as adoptive transfer experiments, genetic knockout experiments, or administering different doses of immunogen. The more scenarios in which a simulation can replicate the behaviour of the real-world, the greater the confidence that may be assumed in its faithful representation of the target system, and the greater the certainty attributed to its results under scenarios where real-world behaviours are not known. Chapter 3's review of the existing literature on modelling and simulating the immune system revealed that most simulations are calibrated against only one experimental scenario, and parameter values are adopted from a wide variety of sources. As explained in chapter 5, calibration against only a single experimental scenario does not protect against fitting an incorrect model to only a single data point. The literature reviewed in chapter 3 often demonstrate that simulation dynamics are consistent with wet-lab experimental data, but also point to discrepancies between the two. We argue that the calibration procedure developed in this thesis is the stronger approach: firstly, discrepancies between *in silico* and wet-lab data can motivate further simulation development; secondly, simulation mechanics and parameters are set to best replicate multiple scenarios, rather than biasing towards any one. This procedure addresses the first two aspects of establishing confidence in simulation: confidence that both simulation mechanics and parameters are appropriate representations of the real-world domain.

The ARTIMMUS simulation was calibrated against two experimental scenarios, and the experience of having done so has revealed how many parameter combinations and mechanistic abstractions can lead to alignment of the simulation with real-world EAE under only one of the two scenarios. It is felt that this collaborative calibration procedure greatly improves confidence that the simulation is representative of real-world EAE. However, it is also felt that ARTIMMUS could benefit from calibration against more than two experimental scenarios, and this is reflected in the further work section below. It should also be noted that a criticism made in chapter 3 of the subjective manner in which much of the existing literature on computational immunology contrasts *in silico* results with wet-lab data applies also to the present calibration of ARTIMMUS. Calibration received substantial input from the domain expert, who made a subjective assessment of whether the dynamics of ARTIMMUS matched their conception of the *in vivo* system. This approach was motivated by the absence of domain-specific data against which ARTIMMUS could be directly contrasted and calibrated. It is, however, believed that this subjective assessment is better made by a domain expert than a modeller.

8.1.3 Demonstrably representative results

The *in silico* experimentation that has been conducted using ARTIMMUS has at times involved subtle changes in molecular expression, the effects of which have ranged from large to small. ARTIMMUS is a stochastic simulation, and as such the results of simulation executions given the same input parameters can differ. Interpretation of *in silico* results requires that one understands the extent to which variations in simulation behaviour can be attributed to the experimental procedure, as opposed to simulation stochasticity. To this end, chapter 6 developed a *consistency analysis technique* to establish the relationship between the level of accuracy reflected in averaged results, and the number of simulation executions sampled in generating them. It allows computational immunologists to ascertain the number of simulation executions required to

correctly interpret data at the level of accuracy they desire, hence instilling confidence that *in silico* results are representative of simulation dynamics, not stochastic variation. As such, it relates directly to the third aspect of instilling confidence in the results of simulation-based experimentation: that results are representative of simulation dynamics and the current experimental procedure, not merely stochastic variation.

As a non-parametric effect magnitude test, the *A* test which underpins this procedure can also be used in hypothesis evaluation: computational immunologists can decide in advance the magnitude of difference they consider significant, and then evaluate the results in the context of hypotheses [Vargha & Delaney 2000]. This consistency analysis technique can indicate the minimum effect magnitude that can be accepted as being representative of the simulation, given the number of simulation executions sampled in generating results.

In summary, this procedure can effectively establish confidence in *in silico* results being representative of simulation dynamics under experimental procedures, especially where the *A* test is used to interpret the significance of results.

8.1.4 Interpretation of *in silico* results in term of the target domain

Of the four aspects of establishing confidence that *in silico* results reflect the target-system, identified above, the last is perhaps the least well appreciated in the literature; it is not considered in any of the literature surveyed. The results that arise from simulation-based experimentation reflect the dynamics of the simulation, and not necessarily the target-domain. It is important to appreciate the separation of simulation from target domain when making claims of immunology that arise from simulation. Results of simulation will not necessarily translate directly into the immunological domain. Simulations represent only a small subset of the entire immune system, with many cells and physical compartments being omitted. Where cell types are represented, their number rarely exceed a few thousand, whereas the real-world system encompasses perhaps millions. Interpretation of *in silico* results into the real-world domain requires the mapping between simulation and target system to be established.

In the case of ARTIMMUS, this mapping has concerned the six point scale system used to grade the severity of autoimmunity in wet-lab mice. The metric relies on human observation of the degree of paralysis experienced by the animal. ARTIMMUS does not represent the vast majority of the mouse, and hence there is no obvious analogue by which to grade the autoimmune severity in the simulation. This has been addressed by engineering two experimental scenarios into the simulation, and calibrating a mechanism whereby the number of autoimmune CD4Th1 cells in the CNS compartment is related to the severity of autoimmunity experienced by real-world experimental animals undergoing the same experiments. This mapping has greatly facilitated the appreciation of *in silico* results in the same terms as employed in the wet-lab. The *in vivo* data against which the mapping was calibrated was relatively sparse: 9 data points for one experiment, and 18 for the other, which were used to form distributions of maximum EAE severity scores across the group. It is felt that the calibrated mapping is as accurate as could be expected given the sparsity of data, and hence the confidence that one may have in this *in silico* autoimmunity severity scoring system being representative of the real-world system could not be improved upon without more wet-lab data.

A second aspect of interpreting *in silico* results into the real-world domain that has been considered in this thesis stems from the substantial variation in, or complete absence of, immunological data required to inform and parameterise the simulation. The values reported for a particular immunological phenomenon, used to guide simula-

tion parameterisation, can vary by orders of magnitude. The effect of this variation or absence on simulation behaviour must be considered prior to interpreting simulation results back into immunology: an insightful or critical *in silico* prediction may only hold for a small range of parameter values, smaller than the range of possible values reported in the immunological literature. If this is the case, then caution much be exercised when interpreting results.

The present thesis has addressed this issue through the development of a novel *robustness analysis*. The analysis may be used to identify the range of simulation parameter values for which simulation behaviour is scientifically consistent, and the points at which scientifically significant deviations in behaviour occur. The analysis has been used in this thesis to determine such values for parameters assigned arbitrary values. Many simulation parameters have been informed through domain expertise, or have been determined through calibration, giving some confidence in their values. However, some parameters have no direct biological analogue, and as such there is no guidance for their parameterisation. In ARTIMMUS these parameters were assigned arbitrary values such that calibration could commence. The analysis reveals the extent upon which the simulation's default behaviour requires these parameters to hold a particular range of values. It is important for computational immunologists to consider these results, and ensure that they are comfortable with them. Furthermore, the results of this analysis should be communicated to the academic community, alongside *in silico* predictions. Researchers in the wider community can then examine these criticalities and assess whether or not they consider them to be important or significant. If there is cause for concern, then less weight may be held in the results of the experimentation.

Several methods of application for this robustness analysis have been proposed in chapter 6. These methods allow researchers to establish the range of simulation behaviors that are explained through underspecified parameter values: variation in domain-specific knowledge. *In silico* results exceeding these ranges may be assumed more representative of the real-world domain.

The robustness analysis, and the methods of application proposed, help to ground *in silico* results in the real domain. They establish the degree to which simulation-derived experimental results are supported by established domain knowledge, and hence how much confidence may be placed in them. Full integration of *in silico* and wet-lab techniques requires that simulation results be appropriately interpreted into the real domain, and the present techniques offer a starting point by which this may be accomplished. With respect to ARTIMMUS, the simplest form of the robustness analysis has been applied, and criticalities of simulation behaviour upon parameter values are known. However, the extended methods of application require more domain-specific data and literature to be gathered, and this is considered beyond the scope of the thesis. It is felt that this work is highly worthwhile, since the potential it has to provide confidence in *in silico* results is thought to be substantial. As such, it is reflected in the further work section below.

8.1.5 Simulation augmentation and confidence in existing results

It has been maintained throughout this thesis that simulation validity is not an all-or-nothing quality, and is best thought of in terms of confidence that the results of *in silico* experimentation are representative of the real-world domain. The techniques explored in this thesis can be applied to either increase that confidence, or demonstrate *in silico* results are not representative. The application of the CoSMoS process and the interdisciplinary collaborative calibration procedure in developing ARTIMMUS do provide confidence in the results of *in silico* experimentation. The splenectomy experiments

of section 7.4 demonstrate that ARTIMMUS produces results consistent with those of a different lab's EAE model, against which it was not calibrated. This demonstrates some level of generality in the simulation, and again serves to improve confidence in its results. However, it is important to consider what may happen to existing simulation results were additions to be made to ARTIMMUS.

One motivation for performing iterations of the CoSMoS process is to address subsequent scientific questions of a domain. As explored more fully in the further work section below, a second regulatory pathway has recently been identified by the Kumar laboratory whereby CD8Treg cells down-regulate the immune response in general. These cells are thought to interact with DCs through CD200:CD200R molecular signalling and suppress their ability to prime T cell populations. Very little is currently known about the influence of this pathway in EAE, and hence its investigation *in silico* is left as further work. However, consider the hypothetical scenario where this pathway has substantial influence in EAE. DCs prime autoimmune CD4Th1 cells, and as such the suppression of this priming ability could help mediate recovery from autoimmunity. However, DCs also prime CD4Th2, CD4Treg and CD8Treg cells, which are involved in the abrogation of autoimmunity. Hence, the CD8Treg cells may regulate themselves, which could serve to enhance autoimmune activity. The effects of this pathway are difficult to predict, but could be substantial. Its implementation in ARTIMMUS might lead to a significant reduction in CD8Treg priming. One of the *in silico* predictions of this thesis is that there exists considerable redundancy in the ability of the CD8Treg population to apoptose autoimmune CD4Th1 cells, thereby mediating recovery from autoimmunity. If the CD200 regulatory pathway were to substantially reduce CD8Treg number, then the pathway may no longer exhibit such redundancy. The *in silico* prediction would have been incorrect, yet the application of the CoSMoS and calibration processes would not have revealed this to be the case *a priori*. In this hypothetical scenario, the CD200 pathway is important and had been investigated. However, had it not been investigated, perhaps because the Kumar laboratory did not perform the experiments to discover it, the *in silico* prediction may have lead to publications which then directed further wet-lab work, all based on incorrect assumptions.

The CoSMoS process and the collaborative calibration procedure can instill confidence that a simulation is an accurate representation of the target system, but they do not *demonstrate* this to be the case, and there is no means of quantifying to what extent it is a faithful surrogate of real-world EAE. The following section proposes an extension of the collaborative calibration procedure which could potentially quantify how well a simulation captures the complexity of the target domain. It can make use of extensive experimental data in doing so. This extension can be employed during development to guide a simulation to an appropriate level of abstraction.

8.1.6 A meta-heuristic framework for quantifying the accuracy of simulation mechanics

The demonstration that a simulation can replicate the results observed in the real-world system under a variety of different scenarios can increase confidence that the simulation has correctly captured the target domain. These *scenarios* represent different experiments, for example: genetic knock-out experiments, administration of different doses of virus, or adoptive transfer experiments. The present section extends this idea, and proposes a framework that makes use of meta-heuristic search techniques in establishing the extent to which a simulation appropriately captures the complexity of the target.

Meta-heuristic search algorithms operate over a space of potential solutions to a problem, and attempt to identify the best quality solution [Luke 2009]. The quality of a

particular solution can be ascertained through application of a fitness function, supplied to the algorithm *a priori*. These algorithms are stochastic in nature, and operate by maintaining either one or many candidate solutions to the problem. Candidate solutions are stochastically perturbed in an effort to improve their quality, and as such these algorithms are iterative in nature. Meta-heuristic search algorithms are not guaranteed to find the global optimum solution to a problem, and will often converge on sub-optimal solutions. They are however well suited to application in very large problem spaces, and where the fitness landscape of solutions is highly non-linear or noisy¹.

The framework builds on the concept that the same experiment performed on the same system will reveal, given allowances for stochasticity, the same results. A simulation intended as a surrogate for another target system will produce similar results if it correctly captures the complexities underlying the target system; it is these complexities that define the way the target system responds to different scenarios. If the complexity is incorrectly captured, then the same set of results cannot be achieved without manipulating some aspect of the simulation's representation in order to compensate.

The proposed framework is captured diagrammatically in figure 8.1, which makes use of only three real-world experiments, but there is no theoretical limit on the number that are actually used in practice. Application of the framework begins by identifying a series of experiments that have been performed on the target system, and engineering each of these into the simulation. Meta-heuristic search is then applied to each experiment *in turn*, attempting to tune simulation parameters to align simulation behaviour under each particular experiment with the corresponding results obtained in the target system. Hence, if N real world experiments have been identified, meta-heuristic search will be performed N times, and will ultimately yield N sets of optimised parameter values, which may be considered as N points in the simulation's parameter space. The degree to which the simulation captures the complexity of the target system is revealed in the proximity of these points in parameter space to one another. If they lie close together, then the simulation may be deemed a faithful representation of the target system. If, on the other hand, the N points in parameter space lie far apart, then the opposite may be concluded. This is because some critical aspect of the target system's complexity has been incorrectly abstracted, and as such the meta-heuristic process must substantially adjust parameter values to compensate.

To reinforce this concept, if the target system was not the real-world system, but in fact the simulation itself, then the meta-heuristic search would always return parameter values that lie in close proximity to one another; the simulation is a perfect representation of itself. It is the distance between points in parameter space at the end of this optimisation process that reveals how well one system captures another.

The framework is not intended as a replacement for the CoSMoS process, rather it constitutes a valuable complement to the simulation development process. The first iterations of the CoSMoS process are often spent developing a simulation to a state that is considered appropriate for the questions that are to be asked of it, however the CoSMoS process does not indicate how this might be ascertained [Andrews *et al.* 2010]. The present framework provides supporting evidence for this decision. The input of the domain expert in grounding simulation work within the real domain cannot be understated, however the present framework provides the objective measure that this

¹This refers to problems where the quality of nearby solutions in problem space may vary considerably, making the application of numerical or analytical optimisation techniques that follow the gradients of neighbouring solution qualities to arrive at optimal solutions problematic.

has been accomplished: that the simulation demonstrably captures the complexities of the real-world system. It is envisaged that the framework be applied at the end of each iteration of simulation development, at which point it can indicate whether or not further development is required. Although determination of which amendments are to be made to a simulation that requires further development are best left to the collaboration between the domain expert and the modeller, it is possible that the results of the present framework can provide insight into which aspects of the simulation are in need of re-consideration. By examining the nature of the real world experiments that could not be replicated without substantial parameter alteration, and consideration of which parameters were being perturbed to compensate for inadequate replication of the target system, insight may be gained into those aspects of the target system and simulation that are not well aligned.

The practical implementation of this framework is not without its challenges, particularly with reference to the tendency of meta-heuristic search to find local optima rather than the global, however discussion thereof is deferred to the further work sections below.

8.1.7 Confidence objectivity

There is another important contribution to establishing confidence in *in silico* results that this framework could provide: objectivity. This thesis has noted that confidence is not an all-or-nothing quality, that one may hold varying degrees of confidence that a simulation is representative of the real system for particular types of experimentation. However, as evidenced in this discussion, confidence is often subjective and difficult to quantify. Of the four aspects of establishing confidence in simulation results explored in this thesis, one can be argued in a highly objective manner: confidence that averaged simulation results are not simply the result of stochasticity.

The consistency analysis technique for establishing the link between accuracy of averaged simulation results and number of simulation samples obtained in generating them is firmly grounded in objective statistical methods. The robustness analysis technique, which quantifies the range of simulation parameter values for which simulation behaviour is scientifically consistent, is also highly objective, though its interpretation is problem specific and may be subjective as a result. The methods of application by which sensitivity and robustness analysis techniques can be applied and considered in the context of domain-specific knowledge to qualify the significance of simulation-based results in terms of the target domain are somewhat objective. The sensitivity and robustness analysis techniques are based on objective statistical methods, however they rely on one having a firm idea of the biologically plausible ranges of parameter values. As presented in this thesis, these methods are theoretical and have not been instantiated. Further research into how these techniques may be employed in scenarios where biologically plausible parameter ranges are known only as probabilities or loose estimates is left as further work, discussed below.

With respect to establishing confidence in a simulation's mechanics, or its parameters, the CoSMoS process does not indicate when a simulation is 'good enough', since this is dependent on the particular study being conducted. The collaborative calibration process explored in this thesis employs the domain expert's opinion in ascertaining when this is the case, an approach which is also subjective. The literature review of existing modelling and simulation works in immunology explored how researchers contrast *in silico* behaviours with established wet-lab data in order to argue the accuracy of the simulation, however comparisons are again subjective. The present meta-heuristic search approach to guiding simulation development to appropriate levels of abstraction




Figure 8.1: A meta-heuristic based technique for quantifying how well a simulation captures its target domain. Established real-world experimental procedures in the target system are identified, and each is engineered into the simulation. Meta-heuristic search is applied to each *in silico* experiment in turn, optimising simulation parameters with respect to the dynamics observed in the real-world system. The resultant optimised parameter sets, of which there is one for each experimental procedure, are then considered. Should these points in parameter space lie in close proximity to one another, then the simulation faithfully captures the complexities of the target domain. Should they lie far apart, then it does not: some critical aspect of the real-world system is inappropriately abstracted in the simulation, and the meta-heuristic search process is making considerable adjustments to parameter values to compensate. Note that only three experiments are depicted here for clarity; the larger the number of experiments, the stronger the technique.

offers an opportunity to argue confidence in simulation mechanics and parameters in more objective terms. The quality of a simulation's representation of its target domain is encoded in the spatial spread of optimised points in parameter space. Whilst it would be challenging to define a general principle for how much spread is acceptable, since it is dependent on the number and nature of parameters in particular simulations, measuring how successive iterations of any one simulation's development reduce this spread could indicate when to stop development: when the improvement in spread is not longer statistically significant.

8.2 Thesis summary, and contribution

The present section returns to the aims and objectives of this thesis, and provides a summary of how the objectives were met, and the contributions that were made in doing so.

To apply and develop statistical and modelling techniques for agent-based simulations of immunology, specifically experimental autoimmune encephalomyelitis.

Five research objectives were identified to guide the work of this thesis in addressing its aim. These were as follows:

Obj 1: Explore the the role of explicit domain modelling in the EAE case study.

Obj 2: Create an agent-based simulation of EAE.

Obj 3: Investigate and develop techniques for calibrating agent-based simulations.

Obj 4: Perform novel *in silico* experimentation using the agent-based simulation of EAE.

Obj 5: Develop and apply statistical techniques for interpreting *in silico* results in the context of the target domain, EAE.

The remainder of this section provides a chapter-by-chapter summary of this thesis, indicating where research objectives were met, and listing the contributions of this thesis.

Chapter 1: The motivation for exploring immunological phenomenon through computational methods is provided. The issue of establishing confidence that the results of *in silico* experimentation are representative of the real-world domain is introduced. The CoSMoS process, which proposes to underpin simulation-based investigation of complex systems with explicit and rigorous modelling activities, is reviewed. Lastly, the thesis aim and research objectives that guide the current work in exploring modelling and simulation techniques to support *in silico* experimentation are detailed, and an overview of the thesis is presented.

Chapter 2: A review of immunology at the level of detail required to comprehend EAE is provided. The EAE mouse model employed by the Kumar lab is described. The complexity of immunology is used to motivate the application of computational methods in its exploration. This chapter forms the domain of the CoSMoS process. **Contributions:**

- Articulation of the complex nature of EAE, and motivating its investigation through computational modelling and simulation.

Chapter 3: The literature on exploring immunology through modelling and simulation is reviewed. Focus is given to the manner in which computational methods have complemented more traditional wet-lab techniques. The question of simulation accuracy is considered, and the approaches through which the literature addresses it are reviewed.

Contributions:

- Making explicit the need to establish confidence in the results of *in silico* experimentation being representative of their target simulation, and the identification of four aspects of simulation-based experimentation in which this must be done: simulation mechanics, simulation parameters, appropriately accurate simulation results, and the interpretation of simulation results into the original domain.

Chapter 4: The chapter reports the construction of a domain model of EAE, expressed using the UML. A detailed assessment of UML's strengths and weaknesses in specifying this complex autoimmune disease is presented. This chapter addresses research objective 1.

Contributions:

- A complete and comprehensive domain model of EAE, no such agent-based model existed prior to this thesis.
- A methodology for creating and presenting complex immunological systems such as EAE using UML.
- A detailed analysis of UML's ability to capture particular immune system aspects present in EAE.

Chapter 5: The construction and calibration of the EAE simulation platform, ARTIMMUS, is presented. The manner in which the domain model is interpreted as a specification for the simulation is detailed, and implementation-specific amendments and additions are reported. In terms of the CoSMoS process, this chapter details the transition from domain model to platform model. The chapter reports the novel and highly interdisciplinary collaborative calibration procedure used to develop and parameterise the simulation to a satisfactory representation of real-world EAE. This chapter addresses research objectives 2 and 3.

Contributions:

- A highly interdisciplinary and collaborative calibration procedure for developing simulations to appropriate levels of abstraction, and the articulation of the importance of calibrating immunological simulations against multiple data points.
- The first agent-based simulation of EAE known to exist, which may be used to perform *in silico* experimentation into the nature of the disease.

Chapter 6: This chapter concerns the development and application of statistical methods to support *in silico* experimentation in ARTIMMUS. In terms of the CoSMoS process it represents part of the results model, the manner in which it is validated against the domain model, and how predictions are interpreted into the domain. Firstly, the chapter derives a means of grading simulation executions in

terms of the 6 point scale employed in wet-lab research. This greatly facilitates interpretation of *in silico* results, and helps to contextualise them in terms of the real domain. The *A* test is used in creating and applying a novel procedure for establishing the relationship between the accuracy of averaged simulation results, and the number of simulation executions sampled in deriving them, named the *consistency analysis technique*. The results of this analysis are used to determine how many simulation executions are employed in deriving averaged simulation results in the remainder of this thesis. A global sensitivity analysis, based on latin hypercube sampling, is performed on ARTIMMUS. The results provide insight into the relative importance of different cells, molecules, and their interactions to the autoimmunity and regulatory aspects of EAE. The *A* test is once more employed in the creation of a novel robustness analysis, capable of determining the range of parameter values over which simulation behaviour is scientifically consistent. This technique is applied to the parameters of ARTIMMUS, and focus is given to those that were assigned arbitrary values. The implications of these results being inappropriately assigned their values are considered. Lastly, the chapter considers how the implications of *in silico* results can be established in terms of the real-world domain. To this end, several approaches in which sensitivity and robustness analyses can be applied and considered in the context of domain-specific knowledge concerning the biologically plausible ranges of parameter values are proposed. Research objectives 4 and 5 are addressed.

Contributions:

- A novel technique for establishing the relationship between the accuracy of averaged simulation results, and the number of simulation executions sampled in deriving them.
- The prediction that the type 2 deviation of the autoimmune response observed during the recovery from EAE is a result of regulation of CD4Th1 cells by CD8Tregs, rather than an active driver of recovery.
- A novel robustness analysis technique, capable of determining the boundaries of simulation parameters at which scientifically significant deviations in simulation behaviour occur.
- Several theoretical methods of application whereby this robustness analysis technique, sensitivity analysis and domain specific knowledge can be used to qualify the significance of *in silico* results in terms of the target domain.

Chapter 7: The ARTIMMUS simulation platform is used to perform novel experimentation into EAE, hence it further represents the CoSMoS results model. Firstly, two experimental scenarios, being physiological recovery from autoimmunity and hampered recovery through abrogation of the regulatory pathway, are examined in detail to reveal the cellular dynamics which give rise to these system-level phenomenon. Secondly, the CD8Treg mediated regulatory pathway is investigated *in silico*, examining the effect of perturbing its efficacy. Lastly, ARTIMMUS is employed in gaining insight into the role of the spleen in EAE. These results are contrasted with those of a different laboratory that performed splenectomy experiments on rats. The results are consistent, and demonstrate the generality of ARTIMMUS: it may be used to gain insight into other EAE models than that of the Kumar laboratory. This chapter addresses research objective 4.

Contributions:

- *In silico* elucidation of the system-level cellular and compartmental dynamics that give rise to autoimmunity and its recovery.
- Prediction that there exists considerable redundancy in the regulatory capacity of the CD8Treg population to counter autoimmunity through the apoptosis of CD4Th1 cells.
- Prediction that the primary role of the spleen in EAE models is as a site of substantial Treg priming, and that hindered recovery from autoimmunity in splenectomised mice is due to the substantially reduced numbers of Tregs generated within the system.

Chapter 8: The conclusive chapter of this thesis. The contribution that explicit domain modelling, calibration and statistical methods make to establishing confidence in *in silico* results being representative of real-world EAE are discussed. An extension to the collaborative calibration procedure of chapter 5 that makes use of established real-world experimental data in quantifying how well simulation captures the target domain is proposed. The contributions of this thesis are summarised, avenues of further work are presented, and final concluding remarks are made. This chapter makes additional contribution to research objective 3.

Contributions:

- A theoretical technique for employing meta-heuristic search algorithms in quantifying how faithful a simulation's representation of the target domain is, and the manner in which this can be used to develop simulations to appropriate levels of abstraction.

From the above summary of this thesis it may be concluded that each of the five research objectives identified to guide the thesis in achieving its aim have been met. As such, it is concluded that this thesis has accomplished its aim.

8.3 Further work

The contributions of this thesis have revealed several avenues of potential further work, which are explored in turn in the present section. Four broad strands of future work are identified, relating to: the *in silico* investigation of EAE, section 8.3.1; the modelling of complex systems using the UML, section 8.3.2; the use of meta-heuristic search in guiding simulation development and calibration, section 8.3.3, and lastly the instantiation and further investigation of robustness analysis methods, section 8.3.4.

8.3.1 Further investigation of EAE

This thesis has made use of ARTIMMUS, a simulation platform for the investigation of EAE, in providing novel investigation and insight into the nature of the disease.

One prediction to arise from this *in silico* experimentation is that type 2 deviation of the autoimmune response is a reactive result of regulatory activity by CD8Treg cells, rather than a driving force in the abrogation of autoimmune behaviour. When DCs that would ordinarily prime CD4Th2 cells are altered to prime only CD4Th1 cells, hence substantially revoking the ability of type 2 deviation to oppose autoimmune activity, the changes to T cell system dynamics were marginal. Two more predictions arising from ARTIMMUS are that there exists great redundancy in the ability of the regulatory network to control autoimmune behaviour, and that the reason splenectomised mice experience hampered recovery from autoimmune symptoms is due to its

priming of large populations of regulatory T cells: its removal substantially reduces the number of Treg cells generated in response to autoimmunity. As highlighted in chapter 3, one means to promote confidence in simulation results is to verify *in silico* predictions experimentally in the real-world system. This has not been performed as clinical immunological investigation is outside the scope of this thesis, and relies on the domain expert currently having the resources to allocate to such experimentation. Wet-lab investigation into the hypotheses arising from ARTIMMUS is left as further work.

In consultation with the domain expert, the following aspect of the real-world domain has been highlighted as a point of further work. Recent wet-lab experimentation has revealed that effector CD8Treg cells express the CD200 molecule. Its corresponding receptor, CD200R, is expressed on dendritic cells (DCs). CD200:CD200R interaction leads to the down-regulation of DC priming ability. Hence, the CD200 pathway represents a second regulatory pathway that may influence the onset and recovery from autoimmunity in EAE. Figure 8.2 depicts the two regulatory pathways in EAE, the induction of apoptosis in autoimmune-inducing CD4Th1 cells by CD8Treg cells already represented in ARTIMMUS, and the suppression of DC priming capacity through CD200:CD200R signalling by CD8Tregs. As may be appreciated from the diagram, it is difficult to predict the exact impact of the CD200 pathway: reduced DC priming ability can lead to the priming of fewer CD4Th1 cells, responsible for mediating autoimmune behaviour, but it may also lead to the priming of fewer Treg cells responsible for opposing autoimmune behaviour. It is proposed that the incorporation of this pathway into ARTIMMUS, and the appropriate re-calibration of the simulation, can lead to novel insight into the role of this pathway in EAE, and its potential as a target for enhancing regulatory activity in mice.

Lastly, should this further work on investigating EAE through ARTIMMUS be performed, the simulation will require additional development and re-parameterisation. It is suggested that this be accompanied by calibration against additional experimental scenarios. Although current calibration against two experimental scenarios is believed to have guided the simulation towards a more faithful representation of the domain than might have been achieved by considering only one, more scenarios would provide further confidence in the simulation's results.

8.3.2 Capturing complex systems with the UML

Chapter 4 has explored the applicability of the UML to the specification of EAE, a complex autoimmune disease. This modelling case study has revealed usage patterns by which the high-dimensional complexity of EAE can be decomposed and described on two-dimensional diagrams. It has highlighted the inability of standard UML notation to express certain aspects of the domain, such as mass-concurrency which escalates as the disease progresses, the ambiguity in attempting to capture both an entity's individual history and the number of entities that may engage in a particular relationship at a system-level, and stochastic and time-dependent behaviours.

Further work in modelling other immunological or complex system domains can reveal whether these findings relate only to the present case study in EAE, or are broader issues in modelling complex systems in general. If found to be general issues, then efforts can be made to define more formally the semantics underpinning the syntactic notations and the patterns of usage identified and created in chapter 4. These contributions can then form useful guidelines and support for others attempting to model complex systems.




Figure 8.2: Abstract representation of the cells and their interactions involved in the two regulatory pathways of EAE: the induction of apoptosis in CD4Th1 cells by CD8Tregs, and the CD200:CD200R mediated suppression of DC priming capacity by CD8Treg cells. Red lines indicate promotion of autoimmune behaviour, whereas blue lines indicate promotion of recovery from autoimmunity. Only DCs and T cells are shown.

8.3.3 Meta-heuristic guided calibration and simulation development

The calibration procedure identified in chapter 5 makes extensive use of the domain expert's expertise to parameterise the simulation. In this manner, the simulation was calibrated and aligned with the work of one particular lab, rather than taking parameter values from potentially unrelated sources and literature. However, the procedure calls for a collaborative decision between modeller and domain expert to be made with respect to altering either simulation parameters or amending the simulation's mechanics. The ARTIMMUS simulation platform contains around 60 parameters which were largely perturbed by hand and assessed by eye, admittedly guided by modelling- and domain-expertise, in aligning simulation behaviour with that of the real-world system. It is felt that this process could benefit from automation and more objective assessment. It is impractical to investigate large numbers of parameter perturbations manually, and it is conceivable that certain simulation abstractions were unnecessarily considered inappropriate on the basis of inability to satisfactorily replicate real-world dynamics as a result.

Meta-heuristic search techniques, introduced above in section 8.1.6, can aid parameterisation by more fully exploring parameter space. In fact, a natural consequence of the application of the meta-heuristic guided development procedure described above is that the points in parameter space necessary to replicate a variety of real-world behaviours should lie in proximity to one another. Hence, the framework can not only indicate when a simulation's development represents an appropriate abstraction of the target domain, but will automatically deliver suitable parameter values as well. These points in parameter space may not correspond exactly, in which case one could either adopt one of them, select a point in parameter space that lies at their centre, or apply a larger multi-objective meta-heuristic search to identify a parameter set that optimises

the simulation with respect to all the real-world behaviours simultaneously.

A challenge for this approach, and of the meta-heuristic guided development framework in general, is the lack of suitable data available in the literature against which to calibrate the simulation. This was noted in chapter 5, and was a motivation for the incorporation of the domain expert in the process. However, if the domain expert is able to subjectively assess whether an observed simulation behaviour matches their intuition of the real-world system, one would hope that they can also describe the criteria against which they are making this assessment. This could form the basis for the fitness evaluation functions required by meta-heuristic algorithms, and whilst still somewhat subjective, it is more objective than relying purely on intuition through subsequent iterations of the calibration procedure.

Hence, further work is proposed to investigate the integration of the aforementioned meta-heuristic guided development with the calibration approach of chapter 5, and with the CoSMoS process. The meta-heuristic guided development process can better explore parameter space, and more objectively indicate whether the simulation appropriately captures the real-world domain. Where this is found to not be the case, collaborative consultation between modeller and domain expert can lead to suggestions for how to best amend the simulation, and as indicated in section 8.1.6 above. It is possible that the meta-heuristic guided development process can suggest areas on which to focus. This leads to an iteration in the CoSMoS process; the domain is explored, domain model diagrams are amended, the simulation platform is altered and finally re-assessed.

A second strand of further work concerns the application of the meta-heuristic guided development process itself. The ability of this process to indicate whether the simulation appropriately abstracts the target domain rests on the ability of meta-heuristic search to locate points in parameter space that best align simulation behaviour with that of the real-world. However, as noted in section 8.1.6, meta-heuristic search techniques do not necessarily reveal the global optimum solution, and can converge on local optimum. This can mislead the overall process: a simulation that does appropriately abstract the target-domain may be incorrectly assessed as not doing so because the search processes converge on local optima that lie far apart. There are a number of potential solutions to this problem. Firstly, certain meta-heuristic search algorithms, such as differential evolution, naturally permit the specification of upper and lower boundaries in which to constrain the search process [Storn & Price 1995]. These boundaries could be used to constrain search to biologically-plausible parameter ranges. Secondly, certain meta-heuristic search algorithms, such as genetic algorithms, maintain a population of candidate solutions [Luke 2009]. These could be seeded with the results of previous optimisations over other experimental scenarios. Hence, the search process is constrained to find local optima that are at least biologically plausible, and encouraged to converge on the same optima, if appropriate. Lastly, the degree to which a simulation faithfully captures the complexity of the real-world domain is encoded in the proximity of points in parameter space resulting from the optimisation process. It seems unlikely that these points will all lie exactly atop one another. Investigation will be required to ascertain the relationship between a satisfactorily faithful representation, and the proximity of these points to one another. This thesis has proposed the meta-heuristic search guided development process, but has not applied it nor investigated the technologies necessary to achieve it. These avenues of research are left as further work.

8.3.4 Instantiation and further investigation of robustness analysis techniques

Chapter 6 saw the creation of a novel robustness analysis technique, capable of indicating the extent to which parameter values can be perturbed before a scientifically significant deviation in simulation behaviour occurs. It was used to identify these parameter boundaries in the ARTIMMUS simulation platform. Section 6.6 then proposed three applications of this technique that could qualify the significance of simulation derived results in terms of the original domain; separating those results arising through underspecified parameter values from those that are genuinely representative of the domain. The application of these techniques to the case study in EAE was considered outside the scope of this thesis. The application of these techniques to a case study where more data on biologically plausible parameter values are known, and hence an evaluation of their contribution to integrating *in silico* results with the simulation's target domain, is left as further work.

A second avenue of further work regarding the application of the robustness analysis is as follows. The three approaches of application for the robustness analysis identified in section 6.6 depend upon one having a firm idea of biologically plausible ranges for parameter values. In many domains this may not be the case, immunologists may be able to quantify a rough probability distribution describing the likely value for a parameter, but might not be willing to provide firm boundaries for its likely value. The application of these robustness analysis techniques in the presence of these distributions of values is worthy of investigation, since it would make the approach more applicable to a wider range of problem domains.

8.4 Concluding remarks

The field of immunology, and the bio-medical industries in general, have made tremendous advancements over the last 150 years. The value of these fields is well appreciated by society, delivering improved well-being and longer lives.

Research methods used in immunology have, likewise, made tremendous advancements. Modern assays can now examine molecular expression by cells and DNA transcription events, leading to data concerning the genes, proteins and events influencing susceptibility and development of disease. *In vivo* imaging techniques can examine cellular interactions at incredibly small scales. However, there is a growing sentiment within immunology that these low level reductionist approaches, and vast quantities of data that they produce, alone cannot lead to a coherent understanding of how low-level system components and interactions lead to disease and recovery [Cohen & Harel 2007, Germain *et al.* 2011].

Immunologists are increasingly looking to computational modelling and simulation methods to aid in integrating, animating and reasoning about the data that they collect from the wet-lab. These methods help in providing an oversight of a system, allow for the formulation and evaluation of hypotheses in the context of established data, and identifying potentially fruitful avenues of further research.

However, a review of the growing literature on modelling and simulating immunological systems reveals that the question of simulation validity is not well appreciated in the field. Computational models and simulations are highly abstract representations of the systems that they intend to capture. They are man-made artifacts, and results of *in silico* experimentation are not necessarily representative of the real-world system. One example where two simulations of the same disease created in the same laboratory have provided directly contradictory results has been highlighted in section 3.5.5 of this

thesis, it is not unreasonable to assume that there may exist many more. For computational techniques to be genuinely insightful the means to determine that a simulation is a faithful representation of its target domain, or at least ascertain the extent to which *in silico* results may be assumed representative, must be established.

The CoSMoS project has identified this need, and is actively researching methods by which to address it. Their CoSMoS process underpins simulation-based investigation with rigorous and explicit domain modelling, promoting close collaboration with domain experts as a means of verifying that the work conducted *in silico* is a fair representation of the real-world system. Further, these models make simulation-based investigation transparent to the wider community, allowing others to assess the assumptions underpinning a simulation, and hence how much trust to hold in its results.

It is the position of this thesis that whilst domain modelling is a valuable and necessary means of arguing confidence in simulation based work, it alone is not sufficient. If a simulation is to be accepted as a faithful surrogate of some target system, then it must be demonstrated as such. The implications of underspecified biological values must be understood in the simulation domain, and results from the simulation domain must be appropriately interpreted in the real-world domain. This thesis has investigated statistical methods that build towards this goal.

The benefits that computational methods can bring to immunological research are compelling. Far sighted researchers are proposing ‘personalised medicine’, computational artifacts that operate at the level of the individual. Hence, an individual’s specific immunological history and current symptoms can be considered *in silico* when selecting the best strategy for treatment for that individual [Hood *et al.* 2004, An *et al.* 2009]. Simulating specific individuals is a formidable goal, and one imagines that it will someday be realised. However, the challenges of simulation validity articulated in this thesis must be met before this can be accomplished. Simulation can have as much potential to mislead as it does to inform. If it is to become an integral component the immunological research of the future, then we must have confidence that its results are insightful and trustworthy.

Appendix A

Primer on UML, and Additional Materials for EAE Domain Model

The present appendix chapter provides additional materials on modelling EAE. It accompanies chapter 4, the domain model of experimental autoimmune encephalomyelitis. Section A.1 provides an introduction to the UML’s class, activity and state machine diagrams, which are used extensively in creating the domain model of EAE. Section A.2 provides additional diagrams that form part of the EAE domain model.

A.1 UML notations

The domain model of EAE, presented in chapter 4 makes use of UML class diagrams, activity diagrams and state machine diagrams. The present section provides an introduction to these UML diagrammatic notations, exploring through example the standard syntax and semantics that are used in the EAE domain model. Departures from standard UML notation in modelling EAE are described in the sections where they are used. For additional information, the reader is referred to [Fowler 2004].

Section A.1.1 describes UML class diagrams, section A.1.2 explores activity diagrams, and lastly section A.1.3 describes state machine diagram notation.

A.1.1 Class diagrams

Figure A.1 shows an example UML class diagram, demonstrating the various concepts and relationships used in the present thesis. Classes represent logical entities in the domain being modelled. These entities may be instantiated, constituting individuals which may carry their own state. The relationships that instances of particular classes have with one another are depicted by lines connecting classes together. The class to which an instance belongs is referred to as its ‘type’.

Relationships may be named, as demonstrated by role name R . Role names may be placed at both ends of a relationships, but are optional. It is customary to provide at least one role name. In the present example, instances of A share relationship $R1$ with instances of class B . The number of instances that engage in a particular relationship




Figure A.1: Example UML class diagram.

may be indicated through ‘multiplicities’, which are again optional. In the present example, every single instance of type *A* is associated with many instances of type *B*, indicated by the symbol $*$, whereas every single instance of type *B* is associated with exactly one instance of type *A*. The relationship between classes *A* and *B* is an example of an ‘association’ relationship; the two classes are related but one does not ‘own’ the other.

A second type or relationship is termed ‘composition’, and is demonstrated by the relationships held between class *B* and classes *C* and *D*. They are used to indicate that instances of some class are fundamental components of instances of some other class. Composition relationships have some restrictions not applicable to standard association relationships. Firstly, a particular instance may engage in only a single composition. Hence, an instance of *B* may be a fundamental component of either an instance of *C* or an instance of *D*, but not both. Secondly, composition entails that the destruction of an instance entails the destruction of all its components. If an instance of *B* is a component of an instance of *C*, and the instance of *C* is destroyed, so too is the instance of *B*. Composition relationships may also be given role names and multiplicities, though these are again optional. In the example, an instance of *C* may be composed of either a single instance, or no instances *B*. On the other hand, instances of *D* may be composed of any number of instances of *B*.

A third type of relationship is termed ‘generalization’. These relationships may have role names, but do not have multiplicities. Generalization is used to indicate that some class is everything that another class is, but may have additional features. For example, classes *E* and *F* have all the features that class *D* does, including the composition relationship with *B*. In addition, class *E* shares relationship *R2* with class *G*, whereas class *F* does not. From an instance of class *B*’s perspective, it does not care whether the instance of *D* that it may be associated with is really an instance of *D*, or *E* or *F*.

A.1.2 Activity diagrams

UML activity diagrams, an example of which is depicted in figure A.2, offer a means to specify the flow of events. ‘Activities’ are indicated in rounded-corner rectangles, and may describe any abstract notion. One activity typically leads to another, as indicated by the arrow transitions. The solid circle indicates the initial node, it is from here that interpretation of the activity diagram commences. The solid circle enclosed within another circle indicates the final node, at which point interpretation terminates.

In the example, the first activity to occur is *A*. This leads to activity *B*. Activities may be undertaken by entities or parties, and this responsibility is indicated by dotted




Figure A.2: Example UML activity diagram.

lines names ‘swim lanes’. In the present example, X is responsible for undertaking activity A , whilst Y is responsible for undertaking all other activities indicated.

Activity diagrams allow for activities to be undertaken in parallel. Completion of activity B leads to activities C and D , which are undertaken in parallel, as indicated by the solid black bar, termed a ‘fork’ relationship. When both C and D complete, the ‘join’ relationship leads to the triggering of E . Both C and D must complete before E commences.

E represents a decision relationship. If condition Z holds true, then E leads to the final state. Otherwise activity F is undertaken, which in turn also leads to the final state.

A.1.3 State machine diagrams




Figure A.3: Example UML state machine diagram.

The dynamics of a logical entity in the model may be depicted using a UML state machine diagram. These diagrams depict the states that such an entity may exist in. State machine diagrams are typically drawn for the classes in the domain model (see

section A.1.1 above). Instances of a class have their own state, as described using a state machine diagram, and two instances of the same class need not be in identical states.

Figure A.3 depicts an example state machine diagram. The solid black circle represents the start state, and the transition from it indicates that upon creation an instance exists in state *B*. States may be encapsulated: state *A* is a composite state, whilst the instance exists in states *B*, *C* or *D*, it is still in state *A*.

Transitions between states are indicated by arrows. They may be labelled with a name, as with the transition from *D* to *E*. Transitions may be guarded, to indicate that some condition must hold true before the transition may occur. The transition from *B* to *C* is guarded by condition *G1*. Furthermore, transitions may entail that some action be undertaken, which may have consequences elsewhere in the system. On transitions, actions are indicated by the ‘/’ symbol, and in the example the transition from *B* to *C* is accompanied by action *Act1*. Names, guards and actions are all optional.

Whilst in a particular state, an entity may also perform actions. This is indicated by the syntax ‘do/’. Whilst in state *C*, an instance in this example performs action *Act2*.

Sequences of states may be orthogonal to one another, and this is indicated by dotted lines. In the example, state *D* is orthogonal to *B* and *C*. An instance in state *A* is also in state *D*, regardless of whether it is in state *B* or *C*.

Any transition leaving a composite state requires that all other sub-states also be left. In the example, undertaking transition *T1* entails that the entity is no longer in state *D*, and undertaking transition *T2* entails that the entity is no longer in state *C*. If *T1* is taken, than the entity does not enter *E*, whilst it does if *T2* is taken. *T1* leads to the final state, in which case the instance is destroyed.

Lastly, states may be decomposed into more complex states, as with state *D* in the example. Decomposition is indicated a by symbol comprising two small circles connected by a line. The expansion of a state is depicted on a separate composite state in which the name appears in italics. In the example, an entity entering state *D* also enters state *F*. It may transit to *G* if condition *G2* is true, and may transit back to *F* by taking transition *T3*. All the while the instance is also in state *D*.

A.2 Additional domain model diagrams

The present section details additional diagrams of the EAE domain model, presented in chapter 4. Section A.2.1 presents a class diagram of the *establishment of regulation* perspective, and section A.2.2 details a class diagram of the *type 2 deviation of the autoimmune response* perspective.

A.2.1 Class diagram of *establishment of regulation* perspective

Figure A.4 depicts a class diagram describing the static relationships between entities involved in the instigation and perpetuation of the regulatory immune response. The response is instigated through the phagocytosis of CD4Th (including both CD4Th1 and CD4Th2) cells by DCs. Any particular DC may phagocytose between zero and many CD4Th cells over the course of its lifespan. A CD4Th cell is phagocytosed by exactly one DC. This prompts the DC to express MHC-II:Fr3. Qa-1:CDR1/2 complexes are expressed only after a DC has been licensed by a CD4Treg, as discussed below. Expression is only possible on DCs that have phagocytosed a least one CD4Th1 cell, in which case many such complexes are expressed. Otherwise none are. A MHC-II:Fr3 complex is expressed by exactly one DC, and the same holds for Qa-1:CDR1/2 cells.




Figure A.4: UML class diagram depicting the relationships between entities of the domain model involved in the instigation and perpetuation of the regulatory immune response.

An MHC-II:Fr3 complex can bind with at most one CD4Treg cell at a time, but over the course of its existence may bind with many different CD4Treg cells. Such binding events deliver signal 1 to CD4Treg cells. A CD4Treg can receive signal 1 at most once. The same relationships hold for Qa-1:CDR1/2 complexes and CD8Treg cells.

Having received signal 1, both CD4Treg and CD8Treg cells can receive signal 2 if the DC to which they are bound expresses co-stimulatory molecules. A DC must be induced into expressing co-stimulatory molecules, and this only occurs upon the perception of a sufficient concentration of type 1 cytokines. A DC is only induced into co-stimulatory molecule expression at most once, after which these molecules are expressed for the remainder of its mature lifespan. Co-stimulatory molecule expression is induced by the simultaneous perception of sufficient concentration of type 1 cytokine, once a DC is mature. Perception by cells does not destroy cytokine molecules, and a particular type 1 cytokine may be perceived by any number of DCs during its existence. When and if expression is induced, a DC will express many co-stimulatory molecules. Co-stimulatory molecules are modelled in such a manner that a single instance is sufficient to deliver signal 2 to a T cell upon binding. As with signal 1, T cells can only receive signal 2 at most once, and they may not receive it at all during their lifespans.

Upon receipt of both signals 1 and 2, both CD4Tregs and CD8Tregs enter their proliferative cycles. During this time they may spawn any number of naive daughter T cells, though only one cell may be spawned at a time. A T cell has only one parent

T cell.

Proliferating T cells differentiate into effector cells once their proliferative cycle completes. An effector CD4Treg cell licenses a DC for Qa-1 expression upon binding with MHC-II:Fr3 complexes as expressed by that DC. This binding also locally activates the CD4Treg. A DC can be licensed for Qa-1 expression at most once, after which time it may express Qa-1 molecules for the remainder of its lifespan as a mature cell. This does not however prevent any number of effector CD4Tregs binding with a particular DC, and hence deriving local activation. A CD4Treg may license any number of DCs.

Both CD8Treg and CD4Treg cells secrete type 1 cytokines upon being locally activated. They may be locally activated any number of times after their differentiation into effector cells. If locally activated, these cells secrete many type 1 cytokines. A particular type 1 cytokine molecule can only be secreted from one source, in this case either a CD4Treg or a CD8Treg (other sources not indicated on the diagram).

Effector CD8Treg cells may bind with Qa-1:CDR1/2 complexes as expressed by CD4Th1 cells, a binding from which they derive local activation. Following this binding, a CD8Treg cell will induce apoptosis in the CD4Th1 cell. A CD4Th1 cell can be induced into apoptosis by a CD8Treg at most once, and a CD8Treg cell can induce apoptosis in any number of CD4Th1 cells.

A.2.2 Class diagram of *type 2 deviation of the autoimmune response* perspective

Figure A.5 denotes a class diagram of the cells and molecules involved in type 2 deviation, and indicates relationships between these entities, and the numbers of entities that engage in such relationships over time. The central component of type 2 deviation of the autoimmune response is the DC, which primes T cells populations. Any particular T cell primes only on a single DC at a time. In the majority of cases a T cell entering proliferation will complete this cycle whilst bound to the same DC, however should the DC expire, a proliferating T cell will continue migratory behaviour until either differentiating into an effector cell, or binding to another DC expressing MHC:peptide complexes for which it is specific. If a DC is not presenting MHC:peptide complexes then it cannot prime T cells, otherwise it may prime many.

The expression of MHC:peptide complexes requires that a DC have phagocytosed either a neuron for expression of MBP peptides, or a CD4Th cell for the expression of Fr3 and CDR1/2 peptides. Any particular apoptotic cell must be phagocytosed by exactly one APC (CNS macrophages are not indicated on figure A.5, since they are not directly related to type 2 deviation of the immune response). The phagocytosis of Treg cells is not indicated on the diagram since they contain no peptides pertaining to EAE.

A pre-requisite for priming of T cells is that the DC express co-stimulatory molecules, a single instance of which is modelled here as being sufficient to induce signal 2 in a single T cell at a time. Co-stimulatory molecule expression is induced in DCs through the perception of type 1 cytokine. Many such cytokine molecules are required for this induction. A type 1 cytokine molecule is secreted by only one cell, which is either a DC, a locally activated effector CD4Th1, CD4Treg (not shown on figure A.5) or CD8Treg cell.

An immunogenic DC is either type 1 or type 2 polarized, and this influences the polarization adopted by CD4Th cells that it primes. Type 1 polarized DCs secrete type 1 cytokine, whereas type 2 polarized DCs do not secrete any cytokines. The polarization that a DC adopts is based on the balance of type 1 versus type 2 cytokine in its immediate vicinity upon maturation. Cytokines molecules are not destroyed




Figure A.5: UML class diagram depicting the relationships between entities of the domain model that lead to the type 2 deviation of the autoimmune response.

following their perception by cells, and as such they may influence anywhere between zero and many cells.

Neurons are indirectly induced into apoptosis by CD4Th1 cells, through the secretion of type 1 cytokine and stimulation of CNS macrophages. Owing to this indirect relationship, a single neuron may be induced into apoptosis through the actions of at least one CD4Th1 cell, and a single CD4Th1 cell is responsible for inducing apoptosis in any number of neurons.

Appendix B

Additional Materials on Calibrating ARTIMMUS

The present appendix chapter provides additional materials on ARTIMMUS and its calibration. Section B.1 provides a full listing of the parameters in ARTIMMUS, noting their name, type, default calibrated values, and function. Section B.2 provides details on how each parameter was calibrated.

B.1 Simulation platform parameters

The following table, table B.1, is a comprehensive list of all the parameters to the ARTIMMUS simulation platform. It lists their name, type, default value following calibration, and a note describing its function. Parameters are grouped according to their function; for example, the names of all parameters relating to T cells begin with the text “*TCell_*.” Some parameters are overridden by others, for example, *TCell_ProliferationMean* applies to all T cells except for CD4Th2 cells, for which the parameter value is provided by *Th2Polarization_proliferationMean*. The prefixes that parameter names are assigned are related to the abstract and concrete classes in the simulation platform’s design, indeed it is in these classes that the parameters are held in the simulation platform. As such, parameters relating to abstract classes may have influence over several concrete cell classes, as indicated by the class diagram of figure 5.1. Hence, parameters starting with the text “*APC_*” apply to both CNS macrophages and DCs.

Table B.1: The standard parameters of the ARTIMMUS simulation platform. The table details their types, default values (following calibration), and what they parameterise. Note that to aid transparency parameter names relate to their function, and are presented in a hierarchical manner. For example, the first two parameters in the table are *APC.immatureDurationMean* and *APC.immatureDurationStdDev*. \mathbb{R} refers to the real numbers, \mathbb{N} refers to the natural numbers. Percentages are indicated in the range 0.0 to 1.0, indicated by the notation (% / 100). $2 \times \mathbb{N}$ indicates even numbers.

Parameter details	Description
APC	
<i>immatureDurationMean</i> type: $0.0 \leq \mathbb{R}$ default: 48.0 hours	The mean of the normal distribution of times from which the duration that an APC spends in an immature state before maturing is selected from.
<i>immatureDurationStdDev</i> type: $0.0 \leq \mathbb{R}$ default: 24.0 hours	$2 \times$ the standard deviation of the normal distribution of times from which the duration that an APC spends in an immature state before maturing is selected from.
<i>timeOfDeathMean</i> type: $0.0 \leq \mathbb{R}$ default: 110.0 hours	The mean of the normal distribution of times from which an APCs mature lifespan is selected from.
<i>timeOfDeathStdDev</i> type: $0.0 \leq \mathbb{R}$ default: 48.0 hours	$2 \times$ the standard deviation of the normal distribution of times from which an APCs mature lifespan is selected from.
<i>probabilityPhagocytosisToPeptide</i> type: $0.0 \leq \mathbb{R} \leq 1.0$ default: 0.02 (% / 100)	The probability (scaled between 0.0 and 1.0) that the phagocytosis of a cell by an APC will lead to the derivation of presentable peptides.
CD4THelper	
<i>diff00</i> type: $0.0 \leq \mathbb{R} \leq 1.0$ default: 0.05 %/100	The probability (scaled between 0.0 and 1.0) that a CD4Th cell will adopt a type 1 polarization if type 1 cytokine comprises less than 80% of the cytokine milieu in the grid space where the cell resides.
<i>diff08</i> type: $0.0 \leq \mathbb{R} \leq 1.0$ default: 0.85 (%/100)	The probability (scaled between 0.0 and 1.0) that a proliferating CD4Th cell will adopt a type 1 polarization if type 1 cytokine comprises 80% or more of the cytokine milieu in the grid space where the cell resides.
CD4Treg	
<i>type1SecretedPerHourWhenActivated</i> type: $0.0 \leq \mathbb{R}$ default: 10.0 molecules	The quantity of type 1 cytokine secreted per hour by locally activated effector CD4Treg cells.
CD8Treg	
<i>type1SecretedPerHourWhenActivated</i> type: $0.0 \leq \mathbb{R}$ default: 10.0 molecules	The quantity of type 1 cytokine secreted per hour by locally activated effector CD8Treg cells.
<i>cd8TregToCD4ThelperSpecificityDropOff</i> type: $0.0 \leq \mathbb{R} \leq 1.0$ default: 1.0 (%/100)	The probability (scaled between 0.0 and 1.0) that a successful binding between an effector CD8Treg and a Qa-1 expressing CD4Th cell will lead to the former inducing the apoptosis of the latter.
CNSCell	
<i>apoptosisTNFaThreshold</i> type: $0.0 \leq \mathbb{R}$ default: 5.0 molecules	The threshold quantity of TNF- α above which a neuron occupying the same grid space is induced into apoptosis.

Continued on Next Page...

Table B.1 – Continued

Parameter details	Description
<i>Circulation</i>	
<i>width</i> type: $2 \times \mathbb{N}$ default: 62 grid spaces	The width of the circulation compartment in terms of grid spaces.
<i>height</i> type: $2 \times \mathbb{N}$ default: 40 grid spaces	The height of the circulation compartment in terms of grid spaces.
<i>timeToCrossOrgan</i> type: $0.0 \leq \mathbb{R}$ default: 5.0 hours	The average time required for a T cell to cross the circulation compartment.
<i>CLN</i>	
<i>width</i> type: $2 \times \mathbb{N}$ default: 50 grid spaces	The width of the CLN compartment in terms of grid spaces.
<i>height</i> type: $2 \times \mathbb{N}$ default: 50 grid spaces	The height of the CLN compartment in terms of grid spaces.
<i>timeToCrossOrgan</i> type: $0.0 \leq \mathbb{R}$ default: 12.0 hours	The average time required for a naive T cell to migrate through the CNS compartment.
<i>CNS</i>	
<i>width</i> type: $2 \times \mathbb{N}$ default: 50 grid spaces	The width of the CNS compartment in terms of grid spaces.
<i>height</i> type: $2 \times \mathbb{N}$ default: 50 grid spaces	The height of the CNS compartment in terms of grid spaces.
<i>timeToCrossOrgan</i> type: $0.0 \leq \mathbb{R}$ default: 20.0 hours	The average time for an apoptotic T cell to migrate through the CNS compartment.
<i>CNSMacrophage</i>	
<i>basalMBPExpressionProbability</i> type: $0.0 \leq \mathbb{R} \leq 1.0$ default: 0.2 (%) / 100)	The probability that a newly created CNS macrophage will express MHC-II:MBP complexes.
<i>type1RequiredForActivation</i> type: $0.0 \leq \mathbb{R}$ default: 2.5 molecules	The threshold quantity of type 1 cytokine above which a CNS macrophage occupying the same grid space is induced into TNF- α secretion.
<i>tnfaSecretedPerHourWhenStimulated</i> type: $0.0 \leq \mathbb{R}$ default: 100.0 molecules	The quantity of TNF- α secreted per hour by an activated CNS macrophage.
<i>phagocytosisProbabilityImmature</i> type: $0.0 \leq \mathbb{R} \leq 1.0$ default: 0.7 (%) / 100)	The probability that contact between an immature CNS macrophage and an apoptotic cell will result in the former phagocytosing the latter.
<i>phagocytosisProbabilityMature</i> type: $0.0 \leq \mathbb{R} \leq 1.0$ default: 0.3 (%) / 100)	The probability that contact between a mature CNS macrophage and an apoptotic cell will result in the former phagocytosis the latter.
<i>DendriticCell</i>	
<i>type1RequiredForActivation</i> type: $0.0 \leq \mathbb{R}$ default: 2.0 molecules	The minimum quantity of type 1 cytokine required for the induction of co-stimulatory molecule expression in a mature DC occupying

Continued on Next Page...

B. ADDITIONAL MATERIALS ON CALIBRATING ARTIMMUS

Table B.1 – Continued

Parameter details	Description
	the same grid space.
<i>phagocytosisProbabilityImmature</i> type: $0.0 \leq R \leq 1.0$ default: 1.0 (%) / 100)	The probability that contact between an immature DC and an apoptotic cell will result in the former phagocytosing the latter.
<i>phagocytosisProbabilityMature</i> type: $0.0 \leq R \leq 1.0$ default: 0.3 (%) / 100)	The probability that contact between a mature DC and an apoptotic cell will result in the former phagocytosing the latter.
<i>type1SecretedPerHourImmunized</i> type: $0.0 \leq R$ default: 10.0 molecules	The quantity of type 1 cytokine secreted per hour of an immunogenic DC.
<i>cytokineType2PolarizationRatio</i> type: $0.0 \leq R \leq 1.0$ default: 0.17 (%) / 100)	The minimum proportion of type 2 cytokine comprising the cytokine milieu required for the adoption of a type 2 polarization in a DC occupying the same grid space.
<i>DendriticCellMigrates</i>	
<i>lengthOfTimeMovingFollowingMigration</i> type: $0.0 \leq R$ default: 3.5 hours	The length of time that a migratory DC is motile for after migrating into a lymphoid compartment from its original compartment.
<i>Molecule</i>	
<i>molecularHalfLife</i> type: $0.0 \leq R$ default: 30 min	The half life of secreted cytokines in hours.
<i>decayThreshold</i> type: $0.0 \leq R$ default: 0.01	The threshold concentration of cytokine in a grid space below which the concentration is set to 0.0.
<i>Simulation</i>	
<i>immunizationLinearFreq</i> type: $0.0 \leq R$ default: 6.0 hours	The periodicity with which immunization DCs are inserted into the SLO compartment, whilst this process persists.
<i>immunizationLinearInitial</i> type: \mathbb{N} default: 14 cells	The number of immunization DCs placed in the SLO compartment at time zero (time at which immunization for EAE is administered).
<i>immunizationLinearDC0</i> type: $0.0 \leq R$ default: 2.0 cells	The number of immunization DCs placed in the SLO compartment as a result of immunization for EAE decreases linearly over time. This parameter dictates the number at time zero.
<i>immunizationLinearGradient</i> type: $R \leq 0.0$ default: -0.005 cells/hour	The number of immunization DCs placed in the SLO compartment as a result for immunization for EAE decreases linearly over time. This parameter dictates the gradient of linear decrease.
<i>numCD4Th</i> type: \mathbb{N} default: 40 cells	The basal number of naive CD4Th cells in the simulation.
<i>numCD4Treg</i> type: \mathbb{N} default: 30 cells	The basal number of naive CD4Treg cells in the simulation.
<i>numCD8Treg</i>	The basal number of naive CD8Treg cells in

Continued on Next Page...

Table B.1 – Continued

Parameter details	Description
<code>type: IN</code> <code>default: 30 cells</code>	the simulation.
<code>numCNS</code> <code>type: IN</code> <code>default: 500 cells</code>	The number of neurons in the simulation.
<code>numCNSMacrophage</code> <code>type: IN</code> <code>default: 75 cells</code>	The number of CNS macrophages in the simulation.
<code>numDC</code> <code>type: IN</code> <code>default: 10 cells</code>	The number of DCs permanently residing in the CLN and SLO compartments in the simulation.
<code>numDCCNS</code> <code>type: IN</code> <code>default: 40 cells</code>	The number of migratory DCs in the CNS compartment at the start of the simulation.
<code>numDCSpleen</code> <code>type: IN</code> <code>default: 100 cells</code>	The number of DCs permanently residing in the spleen compartment.
<i>SLO</i>	
<code>width</code> <code>type: 2 × IN</code> <code>default: 50 grid spaces</code>	The width of the SLO compartment in terms of grid spaces.
<code>height</code> <code>type:</code> <code>default: 50 grid spaces</code>	The height of the SLO compartment in terms of grid spaces.
<code>timeToCrossOrgan</code> <code>type: 0.0 ≤ R</code> <code>default: 12.0 hours</code>	The average time required for a naive T cell to migrate through the SLO compartment.
<i>Spleen</i>	
<code>width</code> <code>type: 2 × IN</code> <code>default: 62 grid spaces</code>	The width of the SLO compartment in terms of grid spaces.
<code>height</code> <code>type: 2 × IN</code> <code>default: 40 grid spaces</code>	The height of the SLO compartment in terms of grid spaces.
<code>timeToCrossOrgan</code> <code>type: 0.0 ≤ R</code> <code>default: 5.0 hours</code>	The average time required for a naive T cell to migrate through the Spleen compartment.
<i>TCell</i>	
<code>apoptosisNaiveMean</code> <code>type: 0.0 ≤ R</code> <code>default: 30.0 hours</code>	The mean of the normal distribution of times from which the durations that naive T cells can survive without receiving MHC:peptide stimulation are selected from.
<code>apoptosisNaiveStdDev</code> <code>type: 0.0 ≤ R</code> <code>default: 17.0 hours</code>	$2 \times$ the standard deviation of the normal distribution of times from which the durations that naive T cells can survive without receiving MHC:peptide stimulation are selected from.
<code>apoptosisPartialMaturityMean</code> <code>type: 0.0 ≤ R</code> <code>default: 12.0 hours</code>	The mean of the normal distribution of times from which the durations that partially activated T cells can survive without receiving MHC:peptide stimulation are selected from.
<code>apoptosisPartialMaturityStdDev</code> <code>type: 0.0 ≤ R</code> <code>default: 6.0 hours</code>	$2 \times$ the standard deviation of the normal distribution of times from which the durations that partially activated T cells can survive

Continued on Next Page...

Table B.1 – Continued

Parameter details	Description
	without receiving MHC:peptide stimulation are selected from.
<i>AICDMean</i> type: $0.0 \leq R$ default: 60.0 hours	The mean of the normal distribution of times from which the durations that effector T cells can survive for, given frequent local activation, are selected from.
<i>AICDStdDev</i> type: $0.0 \leq R$ default: 56.0 hours	$2 \times$ the standard deviation of the normal distribution of times from which the durations that effector T cells can survive for, given frequent local activation, are selected from.
<i>becomeEffectorMean</i> type: $0.0 \leq R$ default: 60.0 hours	The mean of the normal distribution of times from which the durations that T cells spend in proliferating states are selected from.
<i>becomeEffectorStdDev</i> type: $0.0 \leq R$ default: 56.0 hours	$2 \times$ the standard deviation of the normal distribution of times from which the durations that T cells spend in proliferating states are selected from.
<i>proliferationMean</i> type: $0.0 \leq R$ default: 19.2 hours	The mean of the normal distribution of times from which the durations required for proliferating T cells to produce naive daughter cells are selected from. This parameter does not apply to CD4Th2 cells.
<i>proliferationStdDev</i> type: $0.0 \leq R$ default: 9.6 hours	$2 \times$ the standard deviation of the normal distribution of times from which the durations required for proliferating T cells to produce naive daughter cells are selected from. This parameter does not apply to CD4Th2 cells.
<i>cellsPerGridspace</i> type: $1 \leq N$ default: 7 cells	The maximum number of T cells that fit into a single simulation grid space.
<i>specificityUpperLimit</i> type: $0.0 \leq R \leq 1.0$ default: 0.9 %/100	The upper limit of specificities that T cells may have for their cognate MHC:peptide complexes.
<i>specificityLowerLimit</i> type: $0.0 \leq R \leq 1.0$ default: 0.5 %/100	The lower limit of specificities that T cells may have for their cognate MHC:peptide complexes.
<i>timeLocalActivationInducedEffector... ...FunctionFor</i> type: $0.0 \leq R$ default: 48.0 hours	The length of time that effector T cells survive for following each local activation event, before entering apoptosis through neglect.
<i>timeLocalActivationDelay</i> type: $0.0 \leq R$ default: 10.0 hours	The length of time that must elapse following differentiation into effector cells before a T cell is susceptible to local activation.
<i>Th1Polarization</i>	
<i>mhcUnExpressionDelayMean</i> type: $0.0 \leq R$ default: 8.0 hours	The mean of the normal distribution of times from which the durations of Qa-1 expression by effector CD4Th1 cells are selected from.
<i>mhcUnexpressionDelayStdDev</i> type: $0 \leq R$ default: 2.0 hours	$2 \times$ the standard deviation of the normal distribution of times from which the durations of Qa-1 expression by effector CD4Th1 cells are selected from.
<i>type1SecretedPerHourWhenActivated</i> type: $0.0 \leq R$	The quantity of type 1 cytokine secreted per hour by locally activated effector CD4Th1

Continued on Next Page...

Table B.1 – Continued

Parameter details	Description
default: 100.0 molecules	cells.
<i>Th2Polarization</i>	
<p><i>proliferationMean</i> type: $0.0 \leq \mathbb{R}$ default: 28.8 hours</p>	The mean of the normal distribution of times from which the durations required for proliferating CD4Th2 cells to produce naive daughter cells are selected from.
<p><i>proliferationStdDev</i> type: $0.0 \leq \mathbb{R}$ default: 19.2 hours</p>	$2 \times$ the standard deviation of the normal distribution of times from which the durations required for proliferating CD4Th2 cells to produce naive daughter cells are selected from.
<p><i>type2SecretedPerHourWhenActivated</i> type: $0.0 \leq \mathbb{R}$ default: 100.0 molecules</p>	The quantity of type 2 cytokine secreted per hour by locally activated effector CD4Th2 cells.

B.2 The calibration of simulation parameters

This section gives an overview of how the parameters of the simulation, listed above in table B.1, were assigned their default parameter values. Some parameters were assigned arbitrary values, whereas others are based on domain-specific knowledge or literature. The values of some parameters were adjusted in order to align simulation behaviour with that observed *in vivo*. An overview of the calibration procedure is provided in section 5.3.2. The present section provides more detail on which parameters were calibrated, and how. It is not, however, intended as a comprehensive report of the order in which parameters or simulation behaviours were addressed, or the manner in which their exact values were altered.

The compartmental dimensions, parameterised by *Circulation_height*, *Circulation_width*, *CLN_height*, *CLN_width*, *CNS_height*, *CNS_width*, *SLO_height*, *SLO_width*, *Spleen_height* and *Spleen_width* are all assigned arbitrary values of 50, with the following exception. The height and width of the circulation and spleen compartments are set to permit a cell to migrate through the compartment within the required time, yet maintain an area close to that of other parameters, being 2500 (50×50) grid spaces.

The time for a T cell to migrate through each compartment is parameterised by the following: *Circulation_timeToCrossOrgan*, *CLN_timeToCrossOrgan*, *CNS_timeToCrossOrgan*, *SLO_timeToCrossOrgan* and *Spleen_timeToCrossOrgan*. The values for the CLN, SLO and spleen are derived from immunological literature [Kindt *et al.* 2007]. The value for the CNS was arbitrarily picked to be greater than the migration time for any other compartment. Although immunological literature indicates that the migration time for the circulatory compartment should be 30 minutes, it is felt that a figure of 30 min substantially biases T cell migrations towards compartments directly involved in EAE, since the simulation platform represents only compartments that are integral to EAE. Hence the migration time is arbitrarily increased to 5 hours.

The parameters *Simulation_numCD4Th*, *Simulation_numCD4Treg* and *Simulation_numCD8Treg* dictate the basal number of T cells in the simulation. Their figures are arbitrarily picked, though the ratios between them are based on the domain expert's input. *Simulation_numDC* is assigned an arbitrarily small value, whilst *Simulation_numDCSpleen* is assigned an arbitrary a value significantly larger than *Simulation_numDC*.

Immunization for EAE in the simulation is parameterised through four parameters: *Simulation_immunizationLinearInitial*, *Simulation_immunizationLinearDC0*, *Simulation_immunizationLinearGradient* and *Simulation_immunizationLinearFreq*. *Simulation_immunizationLinearInitial* is arbitrarily picked. The remaining three are set such that the peak of the CD4Th1 population expansion occurs at approximately the correct time, determined with help of the domain expert. It was also set such that the population of CD4Th1 cells produced through immunization is sufficient to instigate self-perpetuating autoimmunity in the CNS after all the immunization DCs have expired.

TCell_specificityUpperLimit and *TCell_specificityLowerLimit* are arbitrarily selected to give a reasonably wide range of specificities. All the T cells in the simulation are intended to be specific to a particular MHC:peptide complex, and as such *TCell_specificityLowerLimit* is not set to lie near the value 0.0, since this would effectively constitute a T cell with no specificity for its intended MHC:peptide complex.

DendriticCell_phagocytosisProbabilityImmature, *DendriticCell_phagocytosisProbabilityMature*, *CNSMacrophage_phagocytosisProbabilityImmature*, *CNSMacrophage_phagocytosisProbabilityMature* are assigned arbitrary values, but ensuring that mature DCs are the most phagocytic type of APC in the simulation, and that mature APCs are more phagocytic than immature APCs.

The following parameters all pertain to the induction of apoptosis in neurons, and are calibrated as a group: *Simulation_numCNSMacrophage*, *numDCCNS*, *Simulation_numCNS*, *DendriticCell_cytokineType2PolarizationRatio*, *CNSMacrophage_type1RequiredForActivation* and *APC_probabilityPhagocytosisToPeptide*. They are calibrated on the basis of their ability to maintain a self-perpetuating autoimmunity in the CNS, yet permit the regulatory network to abrogate autoimmune behaviour. The domain expert has indicated that *APC_probabilityPhagocytosisToPeptide* should hold a value of a few percent, and this is used as its starting value in calibration. *Simulation_numCNS* is selected such that there remain sufficient space in the CNS compartment for T cells, DCs and CNS macrophages to migrate around the space without excessive hinderance. The domain expert has indicated that there should exist around 7 neurons for every CNS macrophage; this is reflected in the domain model. *DendriticCell_type1RequiredForActivation* is set to be more sensitive than *CNSMacrophage_type1RequiredForActivation*.

CNSCell_apoptosisTNFaThreshold is tuned by examining the concentrations of TNF- α in simulation gridspaces over time. From this graph it is possible to identify peaks in concentration which occur when activated CNS macrophages come into close proximity with neurons. The threshold was selected near the top of these peaks, indicating that close proximity to a TNF- α secreting CNS macrophage is required to induce apoptosis in a neuron.

Molecule_molecularHalflife is assigned its value based on an estimate by the domain expert. *Molecule_decayThreshold* is assigned an arbitrary value substantially smaller than the quantity of cytokines secreted by any cell in a single time step of the simulation, and intended such that cytokines can not have infinitely far reaching influence.

The following parameters all relate to the standard deviations of distributions from which the durations of time that cells spend in particular states are selected from: *APC_timeOfDeathStdDev*, *APC_immatureDurationStdDev*, *TCell_proliferationStdDev*, *TCell_apoptosisNaiveStdDev*, *TCell_apoptosisPartialMaturityStdDev*, *TCell_becomeEffectorStdDev*, *TCell_AICDStdDev*, *Th2Polarization_proliferationStdDev*. These parameters are assigned arbitrarily determined values close to, but not exceeding, the corresponding mean parameter for the particular distribution that they parameterise.

The value of *CNSMacrophage_basalMBPExpressionProbability* is fine tuned by identifying the percentage required for autoimmune behaviour to become established in the CNS, and the turnover of CD4Th1 cells in absence of regulatory activity to self-perpetuate once immunization DCs have expired. When set too low, the autoimmune response does not gain the critical intensity required for self-perpetuation after the MHC-II:MBP presenting immunogenic DCs resulting from immunization for EAE expire. The value of *TCell_timeLocalActivationDelay* is assigned on the basis of empirical observation of simulation behaviour. It is set sufficiently high such that T cells do not generally become locally activated in the same compartment as where they were primed. The value of *TCell_timeLocalActivationInducedEffectorFunctionFor* is assigned with guidance from the domain expert.

The values for the parameters *CNSMacrophage_tnfaSecretedPerHourWhenStimulated*, *DendriticCell_type1SecretedPerHourImmunized*, *Th1Polarization_type1SecretedPerHourWhenActivated* and *Th2Polarization_type2SecretedPerHourWhenActivated* are arbitrarily picked. *CD4Treg_type1SecretedPerHourWhenActivated* and *CD8Treg_type1SecretedPerHourWhenActivated* are arbitrarily picked to be significantly less than type 1 secreted by CD4Th1 cells.

APC_timeOfDeathMean was informed by domain expert. *APC_immatureDurationMean* was picked in order to achieve a level of presentation able to sustain immune responses.

The following parameters all pertain to the durations of time that T cells remain in particular stages of their lifecycles for: *TCell_apoptosisNaiveMean*, *TCell_apoptosisPartialMaturityMean*, *TCell_proliferationMean*, *Th2Polarization_proliferationMean*, *TCell_AICDMean* and *TCell_becomeEffectorMean*. Their initial values are all based on consultation with the domain expert, but are then subject to fine-tuning in order to adjust the height of T cell population expansions, and the times at which these occur. This is performed by simultaneously adjusting all parameters by the same proportion.

CD4THelper_diff00 and *CD4THelper_diff08* are selected such the balance of CD4Th1 to CD4Th2 cells during the immune response are perceived as correct by the domain expert. Furthermore, these parameters are set to reflect the strong cytokine stimulus required for the adoption of type 1 polarization by CD4Th cells, which will more readily adopt a type 2 polarization [Kindt *et al.* 2007].

The value of *TCell_cellsPerGridspace* is informed by domain expert, based on the diameter of a T cell in comparison to a dendritic cell.

The values of *Th1Polarization_mhcUnExpressionDelayMean* and *Th1Polarization_mhcUnExpressionDelayStdDev* were informed by domain expert.

CD8Treg_cd8TregToCD4THelperSpecificityDropOff is arbitrarily set to 100%, since the domain expert had no reason to believe that interaction between a CD8Treg and Qa-1 expressing effector CD4Th1 cell would constitute any weaker a binding than a CD8Treg with a Qa-1:CDR1/2 expressing DC.

DendriticCellMigrates_lengthOfTimeMovingFollowingMigration is selected such that migratory DCs originating from the CNS on average move around half way through the CLN compartment before becoming immotile.

Appendix C

Supporting Material for Statistical Techniques

The present appendix chapter provides additional material and data on the *consistency analysis* technique of section 6.3, and on applying the global sensitivity and robustness analysis techniques of sections 6.4 and 6.5 to ARTIMMUS.

Section C.1 provides the code used to generate *A* test scores. Data pertaining to global sensitivity analysis is presented in section C.2, and that pertaining to the robustness analysis is presented in section C.3.

C.1 Calculating the *A* test

The following matlab¹ code is used to generate Vargha-Delaney *A* test scores in this paper [Vargha & Delaney 2000]. It was kindly provided through personal communication with Prof. Susan Stepney and Simon Poulding at the University of York. Note that it is only valid if the two distributions being compared contain the exact same number of samples.

```
function A = Atest(X, Y)
[p,h,st] = ranksum(X,Y,'alpha',0.05);
N = size(X,1); M = size(Y,1);
A = (st.ranksum/N - (N+1)/2)/M;
```

C.2 Global sensitivity analysis

This section details all results of performing a latin hypercube based sensitivity analysis on the EAE simulation, as described in section 6.4. Data for each response is presented in turn, in the following order: *CD4Th1 Max*, *CD4Th1 Max Time*, *CD4Th2 Max*, *CD4Th2 Max Time*, *CD4Treg Max*, *CD4Treg Max Time*, *CD8Treg Max*, *CD8Treg Max Time*, *CD4Th1 at 40 days*, *Maximum EAE*, and *EAE at 40 days*. Data for each response proceeds with a table that details the partial rank correlation coefficient for each parameter and the particular response, and its associated p value. Graphs of

¹<http://www.mathworks.co.uk/>.

parameter values against response values are drawn for all parameters for which the associated p value was less than or equal to 0.1.

Note that the partial rank correlation coefficient removes the effect of other parameter changes in its calculation, however these changes are not removed from the response-parameter graphs. As such, the partial rank correlation coefficients reported do not bear a direct relation to the correlations that may be observed on the graphs.

Table C.1: Partial rank correlation coefficient between parameter values and the *CD4Th1 Max* response, and the associated p value.

Response Name	PRCC	P Value	Rank
<i>CD4THelper_diff08</i>	0.82	6.9e-105	1
<i>CNSCell_apoptosisTNFaThreshold</i>	-0.77	4.6e-83	2
<i>CNSMacrophage_tnfaSecretedPerHourWhenStimulated</i>	0.77	2.2e-82	3
<i>Simulation_immunizationLinearDC0</i>	0.76	8.1e-81	4
<i>APC_timeOfDeathMean</i>	0.76	2.0e-78	5
<i>Molecule_molecularHalflife</i>	0.65	1.5e-50	6
<i>TCell_proliferationMean</i>	-0.63	2.3e-47	7
<i>TCell_AICDMean</i>	0.62	1.5e-46	8
<i>Simulation_immunizationLinearFreq</i>	-0.61	6.4e-44	9
<i>Simulation_immunizationLinearGradient</i>	0.42	7.5e-19	10
<i>TCell_timeLocalActivationInducedEffectorFunctionFor</i>	0.39	6.1e-17	11
<i>TCell_cellsPerGridspace</i>	0.38	1.8e-15	12
<i>TCell_becomeEffectorMean</i>	-0.27	1.4e-08	13
<i>CNS_height</i>	-0.26	3.8e-08	14
<i>APC_probabilityPhagocytosisToPeptide</i>	0.21	1.3e-05	15
<i>Th2Polarization_proliferationMean</i>	-0.18	0.00024	16
<i>CNS_width</i>	-0.18	0.00026	17
<i>Simulation_numCNSMacrophage</i>	0.16	0.00077	18
<i>APC_immatureDurationMean</i>	-0.16	0.0012	19
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	-0.093	0.057	20
<i>TCell_specificityLowerLimit</i>	0.087	0.076	21
<i>DendriticCell_type1RequiredForActivation</i>	0.085	0.082	22
<i>Th1Polarization_mhcUnExpressionDelayStdDev</i>	-0.081	0.097	23
<i>Circulation_timeToCrossOrgan</i>	-0.081	0.1	24
<i>TCell_apoptosisNaiveMean</i>	0.078	0.11	26
<i>TCell_specificityUpperLimit</i>	0.078	0.11	26
<i>DendriticCell_phagocytosisProbabilityImmature</i>	0.077	0.12	28
<i>TCell_AICDStdDev</i>	-0.077	0.12	28
<i>DendriticCell_cytokineType2PolarizationRatio</i>	-0.072	0.14	29
<i>CNSMacrophage_type1RequiredForActivation</i>	-0.071	0.15	30
<i>CD8Treg_type1SecretedPerHourWhenActivated</i>	0.063	0.2	31
<i>TCell_apoptosisPartialMaturityMean</i>	-0.061	0.22	32
<i>Simulation_numDCCNS</i>	0.059	0.23	33
<i>TCell_becomeEffectorStdDev</i>	-0.058	0.23	34
<i>CD4Treg_type1SecretedPerHourWhenActivated</i>	-0.058	0.24	35
<i>APC_timeOfDeathStdDev</i>	-0.055	0.26	36
<i>Simulation_numCNS</i>	0.046	0.34	37
<i>CNSMacrophage_phagocytosisProbabilityMature</i>	0.044	0.37	38
<i>CNS_timeToCrossOrgan</i>	0.042	0.39	39
<i>APC_immatureDurationStdDev</i>	-0.039	0.43	40
<i>SLO_width</i>	-0.038	0.44	41
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	0.032	0.51	43
<i>Simulation_numCD4Th</i>	0.032	0.51	43
<i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i>	-0.029	0.56	44
<i>TCell_timeLocalActivationDelay</i>	0.028	0.57	45
<i>TCell_apoptosisNaiveStdDev</i>	-0.027	0.58	46
<i>SLO_height</i>	-0.026	0.59	47
<i>CLN_height</i>	0.025	0.62	48
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	0.024	0.63	49
<i>Spleen_height</i>	0.023	0.64	51
<i>TCell_proliferationStdDev</i>	-0.023	0.64	51

Continued on Next Page...

Table C.1 – Continued

Response Name	PRCC	P Value	Rank
<i>TCell_apoptosisPartialMaturityStdDev</i>	-0.022	0.65	52
<i>CNSMacrophage_basalMBPExpressionProbability</i>	0.022	0.66	53
<i>Circulation_height</i>	-0.021	0.66	54
<i>Simulation_immunizationLinearInitial</i>	0.02	0.68	55
<i>CNSMacrophage_phagocytosisProbabilityImmature</i>	0.02	0.69	56
<i>Th2Polarization_proliferationStdDev</i>	0.018	0.71	57
<i>Spleen_timeToCrossOrgan</i>	-0.017	0.73	58
<i>CD4THelper_diff00</i>	-0.016	0.74	59
<i>CLN_timeToCrossOrgan</i>	0.015	0.75	60
<i>SLO_timeToCrossOrgan</i>	0.015	0.76	61
<i>Simulation_numDC</i>	0.013	0.79	62
<i>DendriticCell_type1SecretedPerHourImmunized</i>	0.012	0.8	63
<i>Circulation_width</i>	-0.01	0.84	64
<i>Simulation_numCD4Treg</i>	-0.0049	0.92	65
<i>Simulation_numCD8Treg</i>	-0.0042	0.93	66
<i>Molecule_decayThreshold</i>	0.004	0.93	67
<i>Spleen_width</i>	0.0034	0.94	68
<i>CLN_width</i>	-0.0032	0.95	69
<i>DendriticCell_phagocytosisProbabilityMature</i>	0.0025	0.96	70
<i>Simulation_numDCSpleen</i>	-0.0021	0.97	71
<i>DendriticCellMigrates_lengthOfTimeMovingFollowingMigration</i>	0.00056	0.99	72

 Table C.2: Partial rank correlation coefficient between parameter values and the *CD4Th1 Max Time* response, and the associated p value.

Response Name	PRCC	P Value	Rank
<i>TCell_proliferationMean</i>	0.83	7.5e-109	1
<i>CNSMacrophage_tnfaSecretedPerHourWhenStimulated</i>	0.75	1.4e-76	2
<i>CNSCell_apoptosisTNFaThreshold</i>	-0.73	2.9e-71	3
<i>Molecule_molecularHalflife</i>	0.59	8.7e-41	4
<i>TCell_becomeEffectorMean</i>	-0.48	1.2e-25	5
<i>CD4THelper_diff08</i>	-0.43	3.4e-20	6
<i>APC_timeOfDeathMean</i>	-0.43	5.7e-20	7
<i>Simulation_immunizationLinearDC0</i>	-0.41	1.7e-18	8
<i>CNS_height</i>	-0.3	3.4e-10	9
<i>Simulation_immunizationLinearGradient</i>	0.24	7.2e-07	10
<i>TCell_becomeEffectorStdDev</i>	0.24	8.6e-07	11
<i>CNS_width</i>	-0.22	3.7e-06	12
<i>TCell_AICDMean</i>	0.19	6.9e-05	13
<i>Th2Polarization_proliferationMean</i>	0.17	0.00041	14
<i>Simulation_numCNSMacrophage</i>	0.15	0.0016	16
<i>TCell_apoptosisNaiveMean</i>	0.15	0.0016	16
<i>TCell_specificityLowerLimit</i>	-0.13	0.0062	17
<i>APC_immatureDurationStdDev</i>	0.13	0.0073	18
<i>TCell_cellsPerGridspace</i>	0.12	0.011	19
<i>Simulation_immunizationLinearFreq</i>	0.12	0.012	20
<i>CNSMacrophage_type1RequiredForActivation</i>	-0.12	0.015	21
<i>Simulation_numCD4Th</i>	-0.11	0.021	23
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	-0.11	0.021	23
<i>TCell_timeLocalActivationInducedEffectorFunctionFor</i>	0.11	0.024	24

Continued on Next Page...

Table C.2 – Continued

Response Name	PRCC	P Value	Rank
<i>SLO_width</i>	0.091	0.062	25
<i>TCell_apoptosisPartialMaturityStdDev</i>	0.085	0.084	26
<i>TCell_specificityUpperLimit</i>	-0.079	0.11	27
<i>Simulation_numDCCNS</i>	0.076	0.12	28
<i>SLO_height</i>	0.075	0.13	29
<i>APC_probabilityPhagocytosisToPeptide</i>	0.073	0.14	30
<i>Simulation_numCNS</i>	0.072	0.14	31
<i>Circulation_width</i>	0.071	0.15	32
<i>CNSMacrophage_basalMBPExpressionProbability</i>	0.064	0.19	33
<i>Spleen_timeToCrossOrgan</i>	0.062	0.21	34
<i>CLN_timeToCrossOrgan</i>	0.058	0.24	35
<i>CLN_width</i>	0.054	0.27	36
<i>Circulation_height</i>	0.053	0.28	37
<i>CD8Treg_type1SecretedPerHourWhenActivated</i>	-0.048	0.33	38
<i>Circulation_timeToCrossOrgan</i>	0.047	0.33	39
<i>Molecule_decayThreshold</i>	-0.045	0.35	40
<i>CNSMacrophage_phagocytosisProbabilityImmature</i>	-0.041	0.4	41
<i>TCell_proliferationStdDev</i>	0.038	0.43	42
<i>Simulation_numCD8Treg</i>	-0.035	0.47	43
<i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i>	-0.033	0.5	44
<i>DendriticCell_type1SecretedPerHourImmunized</i>	0.033	0.51	46
<i>TCell_AICDStdDev</i>	0.033	0.51	46
<i>TCell_apoptosisNaiveStdDev</i>	0.031	0.53	47
<i>Simulation_numCD4Treg</i>	-0.029	0.55	49
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	0.029	0.55	49
<i>TCell_apoptosisPartialMaturityMean</i>	0.027	0.58	50
<i>CD4Treg_type1SecretedPerHourWhenActivated</i>	0.025	0.61	52
<i>CLN_height</i>	0.025	0.61	52
<i>Th1Polarization_mhcUnExpressionDelayStdDev</i>	0.024	0.63	54
<i>DendriticCellMigrates_lengthOfTimeMovingFollowingMigration</i>	-0.024	0.63	54
<i>Spleen_width</i>	0.023	0.64	55
<i>DendriticCell_phagocytosisProbabilityMature</i>	0.02	0.69	56
<i>DendriticCell_phagocytosisProbabilityImmature</i>	-0.018	0.71	57
<i>APC_timeOfDeathStdDev</i>	-0.017	0.72	58
<i>Simulation_immunizationLinearInitial</i>	-0.017	0.73	59
<i>Th2Polarization_proliferationStdDev</i>	0.016	0.75	60
<i>APC_immatureDurationMean</i>	-0.015	0.76	61
<i>DendriticCell_cytokineType2PolarizationRatio</i>	0.014	0.78	62
<i>DendriticCell_type1RequiredForActivation</i>	-0.013	0.79	63
<i>CD4THelper_diff00</i>	0.012	0.81	64
<i>TCell_timeLocalActivationDelay</i>	-0.011	0.82	65
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	0.01	0.83	66
<i>CNS_timeToCrossOrgan</i>	-0.0095	0.85	67
<i>CNSMacrophage_phagocytosisProbabilityMature</i>	-0.0088	0.86	68
<i>SLO_timeToCrossOrgan</i>	-0.0071	0.89	69
<i>Simulation_numDCSpleen</i>	-0.0046	0.93	70
<i>Simulation_numDC</i>	0.0028	0.95	71
<i>Spleen_height</i>	0.0018	0.97	72

C. SUPPORTING MATERIAL FOR STATISTICAL TECHNIQUES

Table C.3: Partial rank correlation coefficient between parameter values and the *CD4Th2 Max* response, and the associated p value.

Response Name	PRCC	P Value	Rank
<i>CD4THelper_diff08</i>	-0.86	3.6e-125	1
<i>CNSMacrophage_tnfaSecretedPerHourWhenStimulated</i>	0.78	7.7e-86	2
<i>CNSCell_apoptosisTNFaThreshold</i>	-0.77	4.1e-83	3
<i>Molecule_molecularHalflife</i>	0.63	6.3e-48	4
<i>APC_timeOfDeathMean</i>	0.43	2.6e-20	5
<i>Simulation_immunizationLinearDC0</i>	0.4	7.7e-18	6
<i>TCell_AICDMean</i>	0.36	1.9e-14	7
<i>CNS_height</i>	-0.34	6.9e-13	8
<i>TCell_proliferationMean</i>	-0.24	7.7e-07	9
<i>CNS_width</i>	-0.23	2.6e-06	10
<i>APC_probabilityPhagocytosisToPeptide</i>	0.22	5.2e-06	11
<i>TCell_becomeEffectorMean</i>	-0.21	9.6e-06	12
<i>Simulation_immunizationLinearFreq</i>	-0.2	4.7e-05	13
<i>TCell_cellsPerGridspace</i>	0.19	6.0e-05	14
<i>TCell_timeLocalActivationInducedEffectorFunctionFor</i>	0.16	0.00077	15
<i>Simulation_numCNSMacrophage</i>	0.16	0.00094	16
<i>Simulation_immunizationLinearGradient</i>	0.15	0.0026	17
<i>DendriticCell_cytokineType2PolarizationRatio</i>	-0.13	0.0081	18
<i>Th2Polarization_proliferationMean</i>	-0.12	0.016	19
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	0.12	0.018	20
<i>DendriticCell_phagocytosisProbabilityImmature</i>	0.1	0.039	21
<i>TCell_specificityLowerLimit</i>	0.094	0.054	22
<i>DendriticCellMigrates_lengthOfTimeMovingFollowingMigration</i>	0.092	0.06	23
<i>Circulation_timeToCrossOrgan</i>	-0.087	0.076	24
<i>APC_immatureDurationStdDev</i>	-0.081	0.1	25
<i>CD4THelper_diff00</i>	-0.08	0.1	26
<i>DendriticCell_type1RequiredForActivation</i>	-0.075	0.13	27
<i>CNSMacrophage_type1RequiredForActivation</i>	-0.074	0.13	28
<i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i>	0.068	0.16	29
<i>Simulation_numDCCNS</i>	0.067	0.17	30
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	-0.065	0.18	31
<i>CLN_timeToCrossOrgan</i>	-0.062	0.2	32
<i>TCell_apoptosisNaiveStdDev</i>	0.062	0.21	33
<i>CD8Treg_type1SecretedPerHourWhenActivated</i>	0.057	0.25	34
<i>Molecule_decayThreshold</i>	0.054	0.27	35
<i>APC_immatureDurationMean</i>	-0.047	0.34	36
<i>TCell_apoptosisPartialMaturityStdDev</i>	-0.046	0.35	37
<i>TCell_AICDStdDev</i>	0.044	0.37	39
<i>CNSMacrophage_phagocytosisProbabilityImmature</i>	0.044	0.37	39
<i>TCell_apoptosisPartialMaturityMean</i>	-0.042	0.4	40
<i>CNSMacrophage_basalMBPExpressionProbability</i>	0.041	0.4	41
<i>SLO_timeToCrossOrgan</i>	0.039	0.43	42
<i>TCell_apoptosisNaiveMean</i>	0.038	0.44	44
<i>Simulation_numCNS</i>	0.038	0.44	44
<i>APC_timeOfDeathStdDev</i>	-0.036	0.46	45
<i>TCell_becomeEffectorStdDev</i>	-0.036	0.47	46
<i>Circulation_height</i>	-0.033	0.5	47
<i>Spleen_height</i>	-0.032	0.51	48
<i>Simulation_numCD4Th</i>	0.031	0.53	49
<i>Th1Polarization_mhcUnExpressionDelayStdDev</i>	-0.03	0.54	50
<i>Spleen_timeToCrossOrgan</i>	-0.027	0.58	51

Continued on Next Page...

Table C.3 – Continued

Response Name	PRCC	P Value	Rank
<i>CD4Treg_type1SecretedPerHourWhenActivated</i>	-0.021	0.66	52
<i>Simulation_numDC</i>	-0.018	0.71	53
<i>Th2Polarization_proliferationStdDev</i>	0.017	0.73	54
<i>TCell_specificityUpperLimit</i>	0.016	0.74	55
<i>Simulation_numCD4Treg</i>	0.014	0.77	56
<i>Spleen_width</i>	-0.014	0.78	57
<i>Simulation_numCD8Treg</i>	0.013	0.79	60
<i>DendriticCell_type1SecretedPerHourImmunized</i>	0.013	0.79	60
<i>CNSMacrophage_phagocytosisProbabilityMature</i>	-0.013	0.79	60
<i>Simulation_numDCSpleen</i>	0.012	0.8	62
<i>SLO_height</i>	-0.012	0.8	62
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	-0.011	0.83	63
<i>CLN_width</i>	0.0069	0.89	65
<i>Simulation_immunizationLinearInitial</i>	0.0069	0.89	65
<i>CLN_height</i>	0.0035	0.94	66
<i>DendriticCell_phagocytosisProbabilityMature</i>	0.0025	0.96	67
<i>Circulation_width</i>	-0.0022	0.96	68
<i>TCell_timeLocalActivationDelay</i>	-0.0019	0.97	69
<i>SLO_width</i>	0.00057	0.99	70
<i>CNS_timeToCrossOrgan</i>	-0.00048	0.99	71
<i>TCell_proliferationStdDev</i>	-0.00021	1.0	72

Table C.4: Partial rank correlation coefficient between parameter values and the *CD4Th2 Max Time* response, and the associated p value.

Response Name	PRCC	P Value	Rank
<i>CNSMacrophage_tnfaSecretedPerHourWhenStimulated</i>	0.81	3.1e-97	1
<i>CNSCell_apoptosisTNFaThreshold</i>	-0.8	6.8e-96	2
<i>Molecule_molecularHalflife</i>	0.67	5.0e-56	3
<i>CNS_height</i>	-0.41	5.7e-18	4
<i>CNS_width</i>	-0.31	4.7e-11	5
<i>TCell_proliferationMean</i>	0.3	2.5e-10	6
<i>CD4THelper_diff08</i>	0.27	1.9e-08	7
<i>APC_probabilityPhagocytosisToPeptide</i>	0.24	1.1e-06	8
<i>Simulation_immunizationLinearGradient</i>	0.2	2.8e-05	9
<i>TCell_AICDMean</i>	0.2	4.1e-05	10
<i>Simulation_numCNSMacrophage</i>	0.19	7.2e-05	11
<i>Simulation_immunizationLinearDC0</i>	0.18	0.00021	12
<i>TCell_becomeEffectorMean</i>	-0.15	0.0027	13
<i>APC_timeOfDeathMean</i>	0.12	0.011	14
<i>TCell_AICDStdDev</i>	0.12	0.014	15
<i>CNSMacrophage_type1RequiredForActivation</i>	-0.11	0.028	16
<i>CD4THelper_diff00</i>	-0.086	0.079	17
<i>TCell_cellsPerGridspace</i>	0.08	0.1	18
<i>CD8Treg_type1SecretedPerHourWhenActivated</i>	-0.078	0.11	19
<i>DendriticCell_type1SecretedPerHourImmunized</i>	0.074	0.13	20
<i>Simulation_numDCCNS</i>	0.065	0.19	21
<i>Simulation_immunizationLinearFreq</i>	0.061	0.21	22
<i>DendriticCell_type1RequiredForActivation</i>	-0.058	0.23	23
<i>Simulation_numCNS</i>	0.058	0.24	24

Continued on Next Page...

Table C.4 – Continued

Response Name	PRCC	P Value	Rank
<i>SLO_width</i>	0.057	0.24	25
<i>Th1Polarization_mhcUnExpressionDelayStdDev</i>	0.055	0.26	26
<i>APC_immatureDurationStdDev</i>	-0.049	0.32	27
<i>APC_timeOfDeathStdDev</i>	-0.048	0.33	29
<i>DendriticCell_phagocytosisProbabilityMature</i>	-0.048	0.33	29
<i>TCell_apoptosisPartialMaturityMean</i>	-0.044	0.37	30
<i>Th2Polarization_proliferationStdDev</i>	-0.041	0.41	31
<i>CD4Treg_type1SecretedPerHourWhenActivated</i>	-0.04	0.42	33
<i>DendriticCell_phagocytosisProbabilityImmature</i>	0.04	0.42	33
<i>Simulation_numCD4Treg</i>	0.036	0.46	34
<i>SLO_timeToCrossOrgan</i>	0.035	0.48	35
<i>CLN_timeToCrossOrgan</i>	0.034	0.48	36
<i>TCell_specificityUpperLimit</i>	-0.033	0.5	38
<i>TCell_apoptosisPartialMaturityStdDev</i>	-0.033	0.5	38
<i>TCell_apoptosisNaiveStdDev</i>	-0.032	0.52	39
<i>TCell_timeLocalActivationInducedEffectorFunctionFor</i>	-0.031	0.53	40
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	-0.029	0.55	42
<i>Simulation_numCD8Treg</i>	-0.029	0.55	42
<i>CNSMacrophage_basalMBPExpressionProbability</i>	0.027	0.59	43
<i>Simulation_immuneActivationLinearInitial</i>	0.026	0.6	44
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	0.025	0.61	46
<i>APC_immatureDurationMean</i>	0.025	0.61	46
<i>SLO_height</i>	-0.023	0.64	48
<i>Simulation_numDC</i>	0.023	0.64	48
<i>Spleen_width</i>	-0.021	0.68	49
<i>TCell_timeLocalActivationDelay</i>	-0.02	0.69	50
<i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i>	-0.019	0.7	51
<i>TCell_becomeEffectorStdDev</i>	0.014	0.78	52
<i>Th2Polarization_proliferationMean</i>	0.013	0.8	53
<i>TCell_apoptosisNaiveMean</i>	-0.012	0.81	54
<i>TCell_specificityLowerLimit</i>	0.008	0.87	55
<i>Simulation_numCD4Th</i>	0.0076	0.88	56
<i>Simulation_numDCSpleen</i>	-0.0074	0.88	57
<i>Circulation_width</i>	0.0071	0.89	58
<i>DendriticCell_cytokineType2PolarizationRatio</i>	-0.0068	0.89	59
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	-0.0066	0.89	60
<i>DendriticCellMigrates_lengthOfTimeMovingFollowingMigration</i>	-0.006	0.9	61
<i>CNSMacrophage_phagocytosisProbabilityMature</i>	-0.0055	0.91	62
<i>CNSMacrophage_phagocytosisProbabilityImmature</i>	0.0046	0.93	64
<i>Spleen_height</i>	-0.0046	0.93	64
<i>CNS_timeToCrossOrgan</i>	0.0045	0.93	66
<i>CLN_width</i>	0.0045	0.93	66
<i>CLN_height</i>	-0.0033	0.95	67
<i>Circulation_timeToCrossOrgan</i>	-0.0032	0.95	69
<i>TCell_proliferationStdDev</i>	-0.0032	0.95	69
<i>Molecule_decayThreshold</i>	0.0026	0.96	70
<i>Circulation_height</i>	-0.0017	0.97	72
<i>Spleen_timeToCrossOrgan</i>	-0.0017	0.97	72

Table C.5: Partial rank correlation coefficient between parameter values and the *CD4Treg Max* response, and the associated p value.

Response Name	PRCC	P Value	Rank
<i>TCell_proliferationMean</i>	-0.91	4.9e-157	1
<i>APC_timeOfDeathMean</i>	0.87	3.7e-130	2
<i>CNSCell_apoptosisTNFaThreshold</i>	-0.78	2.0e-88	3
<i>CNSMacrophage_tnfaSecretedPerHourWhenStimulated</i>	0.78	2.1e-85	4
<i>CD4THelper_diff08</i>	0.69	1.9e-59	5
<i>Simulation_immunizationLinearDC0</i>	0.67	2.9e-55	6
<i>Molecule_molecularHalflife</i>	0.66	2.7e-54	7
<i>APC_probabilityPhagocytosisToPeptide</i>	0.59	2.6e-41	8
<i>TCell_becomeEffectorMean</i>	0.51	1.6e-29	9
<i>TCell_AICDMean</i>	0.45	2.5e-22	10
<i>Simulation_immunizationLinearGradient</i>	0.44	7.5e-21	11
<i>TCell_cellsPerGridspace</i>	0.35	1.3e-13	12
<i>APC_immatureDurationMean</i>	-0.29	1.8e-09	13
<i>Simulation_immunizationLinearFreq</i>	-0.27	1.6e-08	14
<i>CNS_height</i>	-0.24	5.2e-07	15
<i>TCell_timeLocalActivationInducedEffectorFunctionFor</i>	0.23	2.7e-06	16
<i>CNS_width</i>	-0.18	0.0003	17
<i>Th2Polarization_proliferationMean</i>	-0.17	0.00058	18
<i>TCell_becomeEffectorStdDev</i>	-0.16	0.00075	19
<i>Circulation_timeToCrossOrgan</i>	-0.15	0.0022	20
<i>CNSMacrophage_type1RequiredForActivation</i>	-0.14	0.003	21
<i>DendriticCell_phagocytosisProbabilityImmature</i>	0.14	0.0052	22
<i>Simulation_numCNSMacrophage</i>	0.13	0.01	23
<i>Simulation_numDCCNS</i>	0.12	0.013	24
<i>CLN_timeToCrossOrgan</i>	0.12	0.016	25
<i>DendriticCell_phagocytosisProbabilityMature</i>	-0.09	0.067	26
<i>TCell_specificityUpperLimit</i>	0.089	0.069	27
<i>CD4THelper_diff00</i>	-0.084	0.087	28
<i>CD4Treg_type1SecretedPerHourWhenActivated</i>	-0.08	0.1	29
<i>TCell_proliferationStdDev</i>	0.076	0.12	30
<i>CNSMacrophage_phagocytosisProbabilityImmature</i>	-0.069	0.16	31
<i>Simulation_numCD4Treg</i>	0.068	0.17	32
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	-0.066	0.18	33
<i>CD8Treg_type1SecretedPerHourWhenActivated</i>	0.064	0.19	34
<i>TCell_specificityLowerLimit</i>	0.063	0.2	35
<i>CNS_timeToCrossOrgan</i>	0.061	0.22	37
<i>APC_immatureDurationStdDev</i>	-0.061	0.22	37
<i>TCell_AICDStdDev</i>	-0.058	0.24	38
<i>DendriticCell_type1SecretedPerHourImmunized</i>	0.056	0.26	39
<i>CLN_width</i>	-0.055	0.26	40
<i>CNSMacrophage_basalMBPExpressionProbability</i>	0.055	0.27	41
<i>APC_timeOfDeathStdDev</i>	0.054	0.27	42
<i>Simulation_numCD8Treg</i>	-0.053	0.28	43
<i>Spleen_timeToCrossOrgan</i>	-0.051	0.3	44
<i>SLO_height</i>	-0.046	0.35	45
<i>Th1Polarization_mhcUnExpressionDelayStdDev</i>	-0.045	0.36	46
<i>DendriticCell_cytokineType2PolarizationRatio</i>	-0.042	0.39	47
<i>TCell_apoptosisPartialMaturityMean</i>	-0.04	0.41	48
<i>Simulation_numCD4Th</i>	0.04	0.42	49
<i>SLO_timeToCrossOrgan</i>	0.039	0.43	50
<i>TCell_apoptosisPartialMaturityStdDev</i>	-0.029	0.56	51

Continued on Next Page...

Table C.5 – Continued

Response Name	PRCC	P Value	Rank
<i>Simulation_numDC</i>	0.028	0.57	52
<i>Spleen_width</i>	-0.026	0.59	53
<i>Th2Polarization_proliferationStdDev</i>	-0.025	0.6	54
<i>Circulation_height</i>	-0.024	0.62	55
<i>SLO_width</i>	-0.023	0.64	56
<i>TCell_apoptosisNaiveMean</i>	0.021	0.67	59
<i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i>	-0.021	0.67	59
<i>Simulation_numCNS</i>	0.021	0.67	59
<i>Simulation_numDCSpleen</i>	0.019	0.69	60
<i>Circulation_width</i>	-0.019	0.7	61
<i>Molecule_decayThreshold</i>	-0.017	0.72	62
<i>CLN_height</i>	-0.016	0.74	65
<i>Spleen_height</i>	0.016	0.74	65
<i>TCell_timeLocalActivationDelay</i>	0.016	0.74	65
<i>DendriticCellMigrates_lengthOfTimeMovingFollowingMigration</i>	-0.014	0.77	66
<i>CNSMacrophage_phagocytosisProbabilityMature</i>	0.013	0.79	67
<i>Simulation_immuneActivationLinearInitial</i>	0.011	0.82	68
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	0.01	0.83	69
<i>TCell_apoptosisNaiveStdDev</i>	0.0096	0.84	70
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	0.0066	0.89	71
<i>DendriticCell_type1RequiredForActivation</i>	0.0036	0.94	72

Table C.6: Partial rank correlation coefficient between parameter values and the *CD4Treg Max Time* response, and the associated p value.

Response Name	PRCC	P Value	Rank
<i>TCell_proliferationMean</i>	0.92	2.6e-172	1
<i>CNSCell_apoptosisTNFaThreshold</i>	-0.78	1.3e-87	2
<i>CNSMacrophage_tnfSecretedPerHourWhenStimulated</i>	0.78	6.9e-86	3
<i>Molecule_molecularHalflife</i>	0.68	2.6e-58	4
<i>TCell_becomeEffectorMean</i>	-0.51	9.1e-29	5
<i>APC_timeOfDeathMean</i>	-0.39	6.4e-17	6
<i>CNS_height</i>	-0.28	5.2e-09	7
<i>TCell_apoptosisNaiveMean</i>	0.25	1.9e-07	8
<i>CNS_width</i>	-0.23	1.4e-06	9
<i>Circulation_timeToCrossOrgan</i>	0.23	3.3e-06	10
<i>Simulation_immuneActivationLinearFreq</i>	0.22	5.0e-06	11
<i>TCell_becomeEffectorStdDev</i>	0.2	3.6e-05	12
<i>TCell_specificityUpperLimit</i>	-0.2	5.6e-05	13
<i>TCell_AICDMean</i>	0.2	5.9e-05	14
<i>Simulation_numCNSMacrophage</i>	0.18	0.00015	15
<i>Simulation_immuneActivationLinearDC0</i>	-0.16	0.00074	16
<i>CD4THelper_diff08</i>	-0.15	0.002	17
<i>TCell_specificityLowerLimit</i>	-0.15	0.0025	18
<i>Molecule_decayThreshold</i>	-0.11	0.019	19
<i>TCell_cellsPerGridspace</i>	0.11	0.022	20
<i>TCell_timeLocalActivationInducedEffectorFunctionFor</i>	0.097	0.047	21
<i>Simulation_immuneActivationLinearGradient</i>	0.093	0.059	22
<i>TCell_timeLocalActivationDelay</i>	0.089	0.068	23
<i>Simulation_numDCCNS</i>	0.086	0.079	24

Continued on Next Page...

Table C.6 – Continued

Response Name	PRCC	P Value	Rank
<i>Th2Polarization_proliferationStdDev</i>	-0.085	0.084	25
<i>Spleen_width</i>	0.083	0.09	26
<i>DendriticCell_cytokineType2PolarizationRatio</i>	0.081	0.1	27
<i>CLN_timeToCrossOrgan</i>	0.078	0.11	28
<i>Simulation_numCNS</i>	0.077	0.11	29
<i>CNSMacrophage_phagocytosisProbabilityImmature</i>	-0.076	0.12	30
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	0.074	0.13	31
<i>SLO_width</i>	0.073	0.14	32
<i>Spleen_timeToCrossOrgan</i>	0.069	0.16	33
<i>Simulation_numCD4Th</i>	-0.068	0.16	34
<i>DendriticCellMigrates_lengthOfTimeMovingFollowingMigration</i>	-0.067	0.17	35
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	-0.066	0.18	36
<i>APC_immatureDurationMean</i>	0.063	0.2	37
<i>Simulation_numCD4Treg</i>	-0.055	0.26	39
<i>CLN_width</i>	0.055	0.26	39
<i>CD4THelper_diff00</i>	-0.054	0.27	40
<i>TCell_apoptosisNaiveStdDev</i>	0.044	0.37	41
<i>CD4Treg_type1SecretedPerHourWhenActivated</i>	-0.043	0.38	42
<i>CD8Treg_type1SecretedPerHourWhenActivated</i>	0.042	0.39	43
<i>TCell_apoptosisPartialMaturityMean</i>	0.042	0.4	44
<i>CNS_timeToCrossOrgan</i>	0.041	0.4	45
<i>DendriticCell_phagocytosisProbabilityImmature</i>	0.041	0.41	46
<i>Spleen_height</i>	0.04	0.41	47
<i>Circulation_height</i>	0.039	0.43	48
<i>TCell_proliferationStdDev</i>	-0.038	0.44	49
<i>DendriticCell_type1SecretedPerHourImmunized</i>	0.036	0.46	50
<i>Simulation_numCD8Treg</i>	-0.035	0.47	51
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	-0.032	0.51	52
<i>Th2Polarization_proliferationMean</i>	0.028	0.56	53
<i>Circulation_width</i>	0.026	0.6	55
<i>SLO_timeToCrossOrgan</i>	0.026	0.6	55
<i>Simulation_numDCSpleen</i>	-0.025	0.61	56
<i>CNSMacrophage_basalMBPExpressionProbability</i>	0.021	0.66	57
<i>TCell_AICDStdDev</i>	0.02	0.68	58
<i>CLN_height</i>	0.019	0.7	59
<i>APC_timeOfDeathStdDev</i>	0.018	0.72	60
<i>DendriticCell_type1RequiredForActivation</i>	0.016	0.75	61
<i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i>	0.015	0.75	62
<i>SLO_height</i>	0.015	0.76	63
<i>CNSMacrophage_type1RequiredForActivation</i>	-0.014	0.77	64
<i>Th1Polarization_mhcUnExpressionDelayStdDev</i>	0.013	0.79	65
<i>Simulation_immuneActivationLinearInitial</i>	-0.011	0.82	66
<i>APC_immatureDurationStdDev</i>	0.0096	0.85	67
<i>APC_probabilityPhagocytosisToPeptide</i>	0.0063	0.9	68
<i>DendriticCell_phagocytosisProbabilityMature</i>	0.0057	0.91	69
<i>Simulation_numDC</i>	-0.0054	0.91	70
<i>TCell_apoptosisPartialMaturityStdDev</i>	-0.0044	0.93	71
<i>CNSMacrophage_phagocytosisProbabilityMature</i>	-0.00021	1.0	72

C. SUPPORTING MATERIAL FOR STATISTICAL TECHNIQUES

Table C.7: Partial rank correlation coefficient between parameter values and the *CD8Treg Max* response, and the associated p value.

Response Name	PRCC	P Value	Rank
APC_timeOfDeathMean	0.85	2.0e-119	1
TCell_proliferationMean	-0.83	1.7e-106	2
CNSCell_apoptosisTNFaThreshold	-0.8	5.9e-95	3
CNSMacrophage_tnfaSecretedPerHourWhenStimulated	0.79	1.2e-91	4
Molecule_molecularHalflife	0.69	2.0e-60	5
CD4THelper_diff08	0.67	6.7e-55	6
Simulation_immunizationLinearDC0	0.64	1.4e-48	7
APC_probabilityPhagocytosisToPeptide	0.59	5.3e-40	8
TCell_AICDMean	0.5	4.1e-28	9
Simulation_immunizationLinearGradient	0.43	6.9e-20	10
TCell_cellsPerGridspace	0.39	1.7e-16	11
CNS_height	-0.28	7.8e-09	12
APC_immatureDurationMean	-0.25	2.1e-07	13
Simulation_immunizationLinearFreq	-0.23	1.7e-06	14
CNS_width	-0.21	1.9e-05	15
Th2Polarization_proliferationMean	-0.17	0.00049	16
Simulation_numCNSMacrophage	0.15	0.0019	17
TCell_timeLocalActivationInducedEffectorFunctionFor	0.15	0.0025	18
DendriticCell_phagocytosisProbabilityImmature	0.14	0.004	19
TCell_specificityUpperLimit	0.14	0.0041	20
Circulation_timeToCrossOrgan	-0.14	0.0052	21
CNSMacrophage_type1RequiredForActivation	-0.13	0.0071	22
CLN_timeToCrossOrgan	0.12	0.017	23
TCell_specificityLowerLimit	0.11	0.019	24
DendriticCell_phagocytosisProbabilityMature	-0.11	0.021	25
Simulation_numDCCNS	0.11	0.029	26
Simulation_numCD8Treg	0.089	0.071	27
CD4THelper_diff00	-0.083	0.09	28
CNS_timeToCrossOrgan	0.082	0.095	29
APC_immatureDurationStdDev	-0.081	0.096	30
CNSMacrophage_phagocytosisProbabilityImmature	-0.081	0.099	31
CD4Treg_type1SecretedPerHourWhenActivated	-0.071	0.15	32
CNSMacrophage_basalMBPExpressionProbability	0.07	0.15	33
TCell_proliferationStdDev	0.068	0.17	34
TCell_becomeEffectorMean	-0.067	0.17	35
TCell_apoptosisPartialMaturityMean	-0.064	0.19	36
TCell_timeLocalActivationDelay	0.063	0.2	37
Spleen_width	-0.058	0.24	38
TCell_AICDStdDev	-0.057	0.24	39
CD8Treg_type1SecretedPerHourWhenActivated	0.05	0.31	42
DendriticCell_cytokineType2PolarizationRatio	-0.05	0.31	42
SLO_height	-0.05	0.31	42
Th1Polarization_mhcUnExpressionDelayStdDev	-0.049	0.32	43
CLN_width	-0.047	0.34	44
DendriticCell_type1RequiredForActivation	0.038	0.44	45
APC_timeOfDeathStdDev	0.037	0.45	46
Th1Polarization_type1SecretedPerHourWhenActivated	0.036	0.46	47
TCell_apoptosisPartialMaturityStdDev	-0.035	0.47	49
Circulation_height	-0.035	0.47	49
TCell_apoptosisNaiveMean	-0.035	0.48	51
Simulation_numCD4Th	0.035	0.48	51

Continued on Next Page...

Table C.7 – Continued

Response Name	PRCC	P Value	Rank
Simulation_numCNS	0.032	0.51	52
Circulation_width	-0.025	0.61	53
TCell_apoptosisNaiveStdDev	0.022	0.65	54
Simulation_numDCSpleen	0.021	0.67	55
SLO_width	-0.02	0.69	56
Simulation_numCD4Treg	-0.019	0.69	57
Th1Polarization_mhcUnExpressionDelayMean	-0.019	0.7	59
Simulation_numDC	0.019	0.7	59
CD8Treg_cd8TregToCD4ThelperSpecificityDropOff	-0.016	0.75	60
Spleen_timeToCrossOrgan	-0.015	0.76	62
CLN_height	-0.015	0.76	62
Th2Polarization_type2SecretedPerHourWhenActivated	0.012	0.8	64
Molecule_decayThreshold	-0.012	0.8	64
SLO_timeToCrossOrgan	-0.01	0.84	66
Simulation_immunizationLinearInitial	0.01	0.84	66
TCell_becomeEffectorStdDev	0.0097	0.84	67
CNSMacrophage_phagocytosisProbabilityMature	0.0096	0.85	68
Th2Polarization_proliferationStdDev	-0.0092	0.85	69
Spleen_height	0.0087	0.86	70
DendriticCell_type1SecretedPerHourImmunized	0.0037	0.94	71
DendriticCellMigrates_lengthOfTimeMovingFollowingMigration	-0.0012	0.98	72

Table C.8: Partial rank correlation coefficient between parameter values and the *CD8Treg Max Time* response, and the associated p value.

Response Name	PRCC	P Value	Rank
TCell_proliferationMean	0.93	3.7e-183	1
CNSCell_apoptosisTNFaThreshold	-0.77	5.5e-84	2
CNSMacrophage_tnfaSecretedPerHourWhenStimulated	0.76	3.2e-81	3
Molecule_molecularHalflife	0.67	2.0e-55	4
TCell_becomeEffectorMean	-0.57	5.0e-37	5
APC_timeOfDeathMean	-0.45	5.0e-22	6
CNS_height	-0.26	5.0e-08	7
Circulation_timeToCrossOrgan	0.24	1.1e-06	8
TCell_apoptosisNaiveMean	0.23	2.6e-06	9
CNS_width	-0.22	4.3e-06	10
Simulation_immunizationLinearFreq	0.22	7.9e-06	11
Simulation_immunizationLinearDC0	-0.21	1.6e-05	12
TCell_specificityUpperLimit	-0.21	2.3e-05	13
TCell_becomeEffectorStdDev	0.2	3.5e-05	14
Simulation_numCNSMacrophage	0.17	0.00034	15
TCell_AICDMean	0.17	0.00042	16
CD4THelper_diff08	-0.17	0.00068	17
TCell_specificityLowerLimit	-0.16	0.0015	18
TCell_timeLocalActivationInducedEffectorFunctionFor	0.12	0.014	19
Molecule_decayThreshold	-0.11	0.025	20
TCell_cellsPerGridspace	0.1	0.034	21
CLN_timeToCrossOrgan	0.1	0.04	22
Spleen_width	0.089	0.068	23
Simulation_numDCCNS	0.083	0.09	24

Continued on Next Page...

Table C.8 – Continued

Response Name	PRCC	P Value	Rank
<i>TCell_timeLocalActivationDelay</i>	0.083	0.092	25
<i>CNSMacrophage_phagocytosisProbabilityImmature</i>	-0.08	0.1	27
<i>DendriticCellMigrates_lengthOfTimeMovingFollowingMigration</i>	-0.08	0.1	27
<i>SLO_width</i>	0.077	0.12	28
<i>Simulation_numCNS</i>	0.076	0.12	29
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	-0.074	0.13	30
<i>Th2Polarization_proliferationStdDev</i>	-0.073	0.14	31
<i>DendriticCell_cytokineType2PolarizationRatio</i>	0.07	0.15	32
<i>Simulation_numCD4Th</i>	-0.068	0.16	33
<i>CD4THelper_diff00</i>	-0.068	0.17	34
<i>APC_immatureDurationMean</i>	0.067	0.17	35
<i>Spleen_timeToCrossOrgan</i>	0.063	0.2	36
<i>TCell_apoptosisPartialMaturityMean</i>	0.061	0.21	37
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	0.061	0.22	38
<i>Simulation_numCD4Treg</i>	-0.055	0.26	40
<i>Simulation_immuneLinearGradient</i>	0.055	0.26	40
<i>CLN_width</i>	0.053	0.28	41
<i>Th2Polarization_proliferationMean</i>	0.05	0.3	42
<i>DendriticCell_phagocytosisProbabilityImmature</i>	0.047	0.34	43
<i>TCell_apoptosisNaiveStdDev</i>	0.046	0.35	44
<i>CNS_timeToCrossOrgan</i>	0.045	0.36	45
<i>SLO_timeToCrossOrgan</i>	0.044	0.37	46
<i>TCell_proliferationStdDev</i>	-0.042	0.39	47
<i>Spleen_height</i>	0.041	0.4	48
<i>CD4Treg_type1SecretedPerHourWhenActivated</i>	-0.039	0.42	49
<i>Circulation_height</i>	0.038	0.44	50
<i>Th1Polarization_mhcUnExpressionDelayStdDev</i>	0.033	0.5	51
<i>Simulation_numCD8Treg</i>	-0.032	0.52	52
<i>CD8Treg_type1SecretedPerHourWhenActivated</i>	0.029	0.56	53
<i>Simulation_numDCSpleen</i>	-0.027	0.58	54
<i>CNSMacrophage_basalMBPExpressionProbability</i>	0.023	0.63	55
<i>DendriticCell_type1RequiredForActivation</i>	0.023	0.64	56
<i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i>	0.02	0.68	57
<i>Circulation_width</i>	0.02	0.69	58
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	-0.019	0.69	60
<i>APC_probabilityPhagocytosisToPeptide</i>	-0.019	0.69	60
<i>SLO_height</i>	0.018	0.71	61
<i>CNSMacrophage_type1RequiredForActivation</i>	-0.017	0.74	62
<i>CLN_height</i>	0.015	0.76	63
<i>APC_timeOfDeathStdDev</i>	0.01	0.83	64
<i>DendriticCell_phagocytosisProbabilityMature</i>	0.0095	0.85	65
<i>APC_immatureDurationStdDev</i>	0.0094	0.85	66
<i>TCell_AICDStdDev</i>	0.0088	0.86	68
<i>Simulation_immuneLinearInitial</i>	-0.0088	0.86	68
<i>Simulation_numDC</i>	-0.0071	0.89	69
<i>DendriticCell_type1SecretedPerHourImmunized</i>	-0.003	0.95	70
<i>CNSMacrophage_phagocytosisProbabilityMature</i>	0.0028	0.95	71
<i>TCell_apoptosisPartialMaturityStdDev</i>	-0.0012	0.98	72

Table C.9: Partial rank correlation coefficient between parameter values and the *CD4Th1 at 40 days* response, and the associated p value.

Response Name	PRCC	P Value	Rank
<i>CNSMacrophage_tnfaSecretedPerHourWhenStimulated</i>	0.45	6.0e-22	1
<i>CNSCell_apoptosisTNFaThreshold</i>	-0.36	3.2e-14	2
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	-0.34	9.6e-13	3
<i>Molecule_molecularHalflife</i>	0.3	2.2e-10	4
<i>TCell_proliferationMean</i>	0.25	1.3e-07	5
<i>CD4THelper_diff08</i>	0.24	7.3e-07	6
<i>CNS_height</i>	-0.23	3.1e-06	7
<i>Circulation_width</i>	0.18	0.0003	8
<i>Simulation_numCD8Treg</i>	-0.15	0.002	9
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	0.15	0.003	10
<i>TCell_specificityUpperLimit</i>	-0.13	0.0065	11
<i>CLN_height</i>	0.13	0.0092	12
<i>TCell_apoptosisNaiveStdDev</i>	0.12	0.011	13
<i>DendriticCell_type1SecretedPerHourImmunized</i>	0.12	0.013	14
<i>Simulation_numDC</i>	0.12	0.014	15
<i>TCell_cellsPerGridspace</i>	-0.11	0.022	16
<i>CNS_width</i>	-0.11	0.024	17
<i>TCell_apoptosisPartialMaturityMean</i>	-0.11	0.029	18
<i>Simulation_numDCCNS</i>	0.1	0.042	19
<i>CNSMacrophage_phagocytosisProbabilityImmature</i>	-0.098	0.046	20
<i>Simulation_numCNSMacrophage</i>	0.095	0.052	21
<i>Simulation_numCNS</i>	0.09	0.067	22
<i>Spleen_width</i>	0.084	0.088	23
<i>Simulation_immunizationLinearGradient</i>	-0.077	0.11	24
<i>TCell_AICDStdDev</i>	0.071	0.15	25
<i>CLN_timeToCrossOrgan</i>	0.065	0.18	26
<i>CLN_width</i>	-0.063	0.2	27
<i>Th1Polarization_mhcUnExpressionDelayStdDev</i>	-0.062	0.21	28
<i>CD8Treg_type1SecretedPerHourWhenActivated</i>	0.056	0.26	29
<i>DendriticCellMigrates_lengthOfTimeMovingFollowingMigration</i>	-0.055	0.26	30
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	-0.054	0.27	31
<i>DendriticCell_cytokineType2PolarizationRatio</i>	-0.051	0.3	32
<i>TCell_proliferationStdDev</i>	0.05	0.31	34
<i>DendriticCell_phagocytosisProbabilityMature</i>	0.05	0.31	34
<i>TCell_becomeEffectorMean</i>	0.045	0.35	35
<i>Simulation_immunizationLinearDC0</i>	-0.045	0.36	37
<i>TCell_apoptosisPartialMaturityStdDev</i>	-0.045	0.36	37
<i>Th2Polarization_proliferationStdDev</i>	0.042	0.39	39
<i>CD4THelper_diff00</i>	0.042	0.39	39
<i>DendriticCell_type1RequiredForActivation</i>	0.04	0.41	41
<i>DendriticCell_phagocytosisProbabilityImmature</i>	0.04	0.41	41
<i>SLO_height</i>	0.04	0.42	42
<i>Spleen_height</i>	-0.039	0.43	43
<i>Th2Polarization_proliferationMean</i>	-0.037	0.45	44
<i>Circulation_timeToCrossOrgan</i>	-0.036	0.46	46
<i>CNSMacrophage_phagocytosisProbabilityMature</i>	0.036	0.46	46
<i>Simulation_numDCSpleen</i>	-0.035	0.48	47
<i>APC_timeOfDeathMean</i>	0.031	0.53	48
<i>CNSMacrophage_basalMBPExpressionProbability</i>	0.029	0.55	49
<i>Simulation_immunizationLinearFreq</i>	-0.027	0.58	50
<i>Circulation_height</i>	0.026	0.59	52

Continued on Next Page...

C. SUPPORTING MATERIAL FOR STATISTICAL TECHNIQUES

Table C.9 – Continued

Response Name	PRCC	P Value	Rank
<i>APC_timeOfDeathStdDev</i>	-0.026	0.59	52
<i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i>	-0.026	0.6	53
<i>CNS_timeToCrossOrgan</i>	0.022	0.66	54
<i>TCell_timeLocalActivationDelay</i>	0.021	0.67	55
<i>Spleen_timeToCrossOrgan</i>	-0.018	0.71	56
<i>APC_probabilityPhagocytosisToPeptide</i>	-0.015	0.76	57
<i>APC_immatureDurationMean</i>	-0.014	0.77	58
<i>Simulation_immuneActivationLinearInitial</i>	0.012	0.8	59
<i>SLO_width</i>	-0.012	0.81	61
<i>Molecule_decayThreshold</i>	0.012	0.81	61
<i>CNSMacrophage_type1RequiredForActivation</i>	-0.011	0.82	63
<i>TCell_AICDMean</i>	0.011	0.82	63
<i>Simulation_numCD4Treg</i>	-0.01	0.83	64
<i>Simulation_numCD4Th</i>	-0.0086	0.86	65
<i>TCell_apoptosisNaiveMean</i>	-0.0068	0.89	66
<i>APC_immatureDurationStdDev</i>	-0.0042	0.93	67
<i>TCell_becomeEffectorStdDev</i>	-0.0034	0.94	68
<i>SLO_timeToCrossOrgan</i>	-0.0018	0.97	69
<i>TCell_specificityLowerLimit</i>	-0.0016	0.97	70
<i>CD4Treg_type1SecretedPerHourWhenActivated</i>	0.0014	0.98	72
<i>TCell_timeLocalActivationInducedEffectorFunctionFor</i>	-0.0014	0.98	72

Table C.10: Partial rank correlation coefficient between parameter values and the *Max EAE* response, and the associated p value.

Response Name	PRCC	P Value	Rank
<i>CNSCell_apoptosisTNFaThreshold</i>	-0.88	3.7e-134	1
<i>CNSMacrophage_tnfSecretedPerHourWhenStimulated</i>	0.87	7.1e-133	2
<i>Molecule_molecularHalfLife</i>	0.79	3.3e-91	3
<i>CNS_height</i>	-0.4	4.2e-17	4
<i>CNS_width</i>	-0.32	1.3e-11	5
<i>Simulation_numCNSMacrophage</i>	0.28	3.7e-09	6
<i>TCell_AICDMean</i>	0.23	2.1e-06	7
<i>CD4THelper_diff00</i>	0.19	0.00014	8
<i>APC_timeOfDeathMean</i>	0.17	0.00053	9
<i>Simulation_immuneActivationLinearDC0</i>	0.15	0.0016	10
<i>Simulation_immuneActivationLinearFreq</i>	-0.11	0.025	11
<i>CNSMacrophage_phagocytosisProbabilityMature</i>	0.1	0.034	12
<i>CNSMacrophage_type1RequiredForActivation</i>	-0.1	0.038	13
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	0.095	0.051	14
<i>Simulation_numCNS</i>	0.089	0.07	15
<i>DendriticCell_type1RequiredForActivation</i>	-0.077	0.11	16
<i>DendriticCell_phagocytosisProbabilityMature</i>	-0.077	0.12	17
<i>TCell_apoptosisNaiveMean</i>	-0.073	0.14	18
<i>Circulation_timeToCrossOrgan</i>	-0.068	0.17	19
<i>TCell_apoptosisPartialMaturityMean</i>	0.062	0.21	20
<i>CD4THelper_diff00</i>	-0.061	0.22	21
<i>TCell_becomeEffectorStdDev</i>	0.057	0.25	23
<i>TCell_proliferationMean</i>	-0.057	0.25	23
<i>CLN_timeToCrossOrgan</i>	0.054	0.27	25

Continued on Next Page...

Table C.10 – Continued

Response Name	PRCC	P Value	Rank
<i>CD4Treg_type1SecretedPerHourWhenActivated</i>	-0.054	0.27	25
<i>CNSMacrophage_phagocytosisProbabilityImmature</i>	0.047	0.34	26
<i>CNS_timeToCrossOrgan</i>	-0.043	0.38	28
<i>TCell_AICDStdDev</i>	0.043	0.38	28
<i>DendriticCellMigrates_lengthOfTimeMovingFollowingMigration</i>	0.041	0.4	29
<i>TCell_timeLocalActivationDelay</i>	-0.04	0.41	30
<i>Simulation_immuneActivationLinearGradient</i>	0.039	0.43	31
<i>SLO_height</i>	-0.036	0.46	32
<i>DendriticCell_type1SecretedPerHourImmunized</i>	0.035	0.48	33
<i>Th2Polarization_proliferationStdDev</i>	-0.034	0.48	35
<i>Th2Polarization_proliferationMean</i>	-0.034	0.48	35
<i>SLO_timeToCrossOrgan</i>	0.034	0.49	37
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	0.034	0.49	37
<i>Simulation_numCD4Th</i>	0.033	0.5	39
<i>CNSMacrophage_basalMBPExpressionProbability</i>	0.033	0.5	39
<i>Simulation_numCD8Treg</i>	-0.032	0.51	40
<i>TCell_cellsPerGridspace</i>	0.029	0.55	42
<i>Spleen_width</i>	-0.029	0.55	42
<i>TCell_specificityUpperLimit</i>	0.028	0.57	43
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	-0.027	0.58	44
<i>SLO_width</i>	0.027	0.59	45
<i>TCell_apoptosisPartialMaturityStdDev</i>	-0.026	0.6	46
<i>Spleen_timeToCrossOrgan</i>	0.025	0.6	47
<i>APC_probabilityPhagocytosisToPeptide</i>	-0.025	0.61	49
<i>TCell_apoptosisNaiveStdDev</i>	-0.025	0.61	49
<i>DendriticCell_phagocytosisProbabilityImmature</i>	-0.023	0.64	51
<i>Simulation_numDC</i>	0.023	0.64	51
<i>Circulation_width</i>	0.02	0.68	52
<i>Simulation_numCD4Treg</i>	0.018	0.72	53
<i>TCell_timeLocalActivationInducedEffectorFunctionFor</i>	0.016	0.74	55
<i>Simulation_immuneActivationLinearInitial</i>	0.016	0.74	55
<i>TCell_becomeEffectorMean</i>	-0.015	0.76	56
<i>CLN_height</i>	0.014	0.78	58
<i>TCell_specificityLowerLimit</i>	-0.014	0.78	58
<i>Spleen_height</i>	0.013	0.8	59
<i>CLN_width</i>	-0.012	0.8	60
<i>DendriticCell_cytokineType2PolarizationRatio</i>	0.011	0.82	61
<i>CD8Treg_type1SecretedPerHourWhenActivated</i>	-0.011	0.83	62
<i>Simulation_numDCSpleen</i>	-0.0096	0.84	63
<i>Simulation_numDCCNS</i>	0.0069	0.89	64
<i>TCell_proliferationStdDev</i>	-0.0063	0.9	65
<i>Th1Polarization_mhcUnExpressionDelayStdDev</i>	0.0059	0.9	66
<i>Circulation_height</i>	0.0058	0.91	67
<i>APC_immatureDurationMean</i>	0.0054	0.91	68
<i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i>	-0.0013	0.98	69
<i>APC_timeOfDeathStdDev</i>	-0.0011	0.98	70
<i>APC_immatureDurationStdDev</i>	0.001	0.98	71
<i>Molecule_decayThreshold</i>	-9.1e-05	1.0	72

C. SUPPORTING MATERIAL FOR STATISTICAL TECHNIQUES

Table C.11: Partial rank correlation coefficient between parameter values and the *EAE at 40 days* response, and the associated p value.

Response Name	PRCC	P Value	Rank
CNSMacrophage_tnf _a SecretedPerHourWhenStimulated	0.77	2.2e-82	1
CNSCell_apoptosisTNFaThreshold	-0.77	6.5e-82	2
Molecule_molecularHalflife	0.65	8.3e-52	3
CNS_height	-0.37	7.9e-15	4
CNS_width	-0.32	4.2e-11	5
Simulation_numCNSMacrophage	0.28	5.1e-09	6
APC_timeOfDeathMean	0.17	0.00055	7
CD4THelper_diff08	0.12	0.013	8
TCell_AICDMean	0.12	0.016	9
Simulation_immuneActivationLinearDC0	0.11	0.023	10
Simulation_numCNS	0.1	0.034	11
APC_probabilityPhagocytosisToPeptide	0.095	0.052	12
TCell_apoptosisNaiveMean	-0.093	0.059	13
CD4THelper_diff00	-0.081	0.099	14
CNSMacrophage_type1RequiredForActivation	-0.081	0.1	15
DendriticCell_type1RequiredForActivation	-0.076	0.12	16
APC_timeOfDeathStdDev	-0.075	0.13	17
CLN_timeToCrossOrgan	0.073	0.14	18
DendriticCell_type1SecretedPerHourImmunized	0.068	0.16	19
TCell_becomeEffectorMean	-0.066	0.18	20
TCell_apoptosisNaiveStdDev	-0.065	0.18	21
Simulation_immuneActivationLinearFreq	-0.061	0.21	23
TCell_specificityUpperLimit	0.061	0.21	23
TCell_AICDStdDev	0.056	0.26	24
CNSMacrophage_phagocytosisProbabilityMature	0.055	0.26	26
CD4Treg_type1SecretedPerHourWhenActivated	-0.055	0.26	26
Th2Polarization_type2SecretedPerHourWhenActivated	-0.05	0.3	27
CNSMacrophage_phagocytosisProbabilityImmature	0.048	0.33	28
TCell_timeLocalActivationDelay	-0.044	0.37	29
DendriticCell_phagocytosisProbabilityImmature	-0.043	0.38	30
Simulation_numDCCNS	0.037	0.46	31
APC_immatureDurationMean	-0.036	0.46	32
Th2Polarization_proliferationMean	-0.036	0.47	33
Spleen_width	-0.031	0.53	35
DendriticCell_phagocytosisProbabilityMature	-0.031	0.53	35
Simulation_numCD4Th	0.03	0.54	36
Simulation_numDCSpleen	-0.029	0.56	38
CNSMacrophage_basalMBPExpressionProbability	-0.029	0.56	38
SLO_timeToCrossOrgan	0.028	0.57	39
SLO_height	-0.027	0.58	40
CNS_timeToCrossOrgan	0.027	0.59	41
Simulation_immuneActivationLinearGradient	-0.026	0.6	42
TCell_becomeEffectorStdDev	0.025	0.62	43
Molecule_decayThreshold	0.024	0.63	44
Simulation_numDC	0.023	0.64	45
TCell_cellsPerGridspace	0.021	0.67	46
Th2Polarization_proliferationStdDev	-0.018	0.71	47
Circulation_timeToCrossOrgan	-0.018	0.72	48
CD8Treg_type1SecretedPerHourWhenActivated	0.017	0.73	49
APC_immatureDurationStdDev	0.015	0.76	50
SLO_width	0.012	0.8	53

Continued on Next Page...

Table C.11 – Continued

Response Name	PRCC	P Value	Rank
<i>Simulation_numCD8Treg</i>	-0.012	0.8	53
<i>Simulation_numCD4Treg</i>	0.012	0.8	53
<i>TCell_proliferationMean</i>	0.012	0.81	54
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	0.011	0.83	55
<i>Spleen_height</i>	-0.01	0.83	56
<i>CLN_height</i>	-0.0072	0.88	57
<i>DendriticCell_cytokineType2PolarizationRatio</i>	0.0068	0.89	58
<i>CLN_width</i>	-0.0063	0.9	60
<i>TCell_proliferationStdDev</i>	0.0063	0.9	60
<i>Circulation_width</i>	0.0057	0.91	62
<i>DendriticCellMigrates_lengthOfTimeMovingFollowingMigration</i>	0.0057	0.91	62
<i>TCell_timeLocalActivationInducedEffectorFunctionFor</i>	0.0037	0.94	63
<i>TCell_specificityLowerLimit</i>	0.0034	0.94	64
<i>TCell_apoptosisPartialMaturityMean</i>	0.003	0.95	65
<i>Simulation_immuneActivationLinearInitial</i>	0.0025	0.96	66
<i>TCell_apoptosisPartialMaturityStdDev</i>	0.0024	0.96	67
<i>Th1Polarization_mhcUnExpressionDelayStdDev</i>	-0.0023	0.96	68
<i>Spleen_timeToCrossOrgan</i>	-0.0017	0.97	69
<i>Circulation_height</i>	-0.0011	0.98	70
<i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i>	-0.00051	0.99	71
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	0.00043	0.99	72

C.3 Robustness sensitivity analysis

This section presents the full results of applying the robustness analysis of section 6.5 to the ARTIMMUS simulation. Data is presented for each response is presented in turn, in the following order: *CD4Th1 Max*, *CD4Th1 Max Time*, *CD4Th2 Max*, *CD4Th2 Max Time*, *CD4Treg Max*, *CD4Treg Max Time*, *CD8Treg Max*, *CD8Treg Max Time*, *CD4Th1 at 40 days*, *Maximum EAE*, and *EAE at 40 days*. Two further responses are presented hereafter, *Maximum EAE A test* and *EAE at 40 days A test*, which employ the A test on *Maximum EAE* and *EAE at 40 days* responses to determine when a scientifically significance change in simulation behaviour has taken place. Data is presented in tables, with each table presenting the simulation's parameters, the robustness index, lower and upper indexes, the lower and upper boundaries at which significant deviation in simulation behaviours take place, the default value for each parameter. Parameters are ranked according to their robustness indexes. For a full explanation of these measures refer to section 6.5.1.

Table C.12: Robustness indexes for parameters with respect to the *CD4 Th1 Max* response. RI, robustness index; LI, lower index; UI, upper index; LB, lower boundary; DV, default value; UB, upper boundary. Results indicating no significant deviation in behaviour are marked with a period.

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
CNSCell_apoptosisTNFaThreshold	5.263	5.263	7.752	4.737(+)	5	5.388(−)	1
CNSMacrophage_tnfaSecretedPerHourWhenStimulated	5.638	7.201	5.638	92.8(−)	100	105.6(+)	2
CD4THelper_diff08	5.66	6.334	5.66	0.7962(−)	0.85	0.8981(+)	3
CNS_height	6.53	6.53	7.432	46.73(+)	50	53.72(−)	4
Simulation_immunizationLinearDC0	6.672	7.838	6.672	1.843(−)	2	2.133(+)	5
APC_timeOfDeathMean	7.3	7.3	9.124	102(−)	110	120(+)	6
Molecule_molecularHalflife	7.439	9.284	7.439	0.4536(−)	0.5	0.5372(+)	7
TCell_AICDMean	8.806	8.806	11.16	54.72(−)	60	66.69(+)	8
TCell_cellsPerGridspace	8.98	8.98	.	6.371(−)	7	.	9
CNS_width	9.273	9.273	12.56	45.36(+)	50	56.28(−)	10
TCell_proliferationMean	9.282	11.97	9.282	16.9(+)	19.2	20.98(−)	11
Simulation_immunizationLinearFreq	10.31	10.31	11.67	5.382(+)	6	6.7(−)	12
Simulation_numCNSMacrophage	12.12	14.29	12.12	64.28(−)	75	84.09(+)	13
TCell_timeLocalActivationInducedEffectorFunctionFor	14.59	14.59	23.67	40.99(−)	48	59.36(+)	14
Simulation_immunizationLinearGradient	16.92	16.92	21.62	-0.005846(−)	-0.005	-0.003919(+)	15
TCell_becomeEffectorMean	18.22	26.21	18.22	44.27(−)	60	70.93(−)	16
APC_probabilityPhagocytosisToPeptide	32.53	32.53	64.08	0.01349(−)	0.02	0.03282(+)	17
Th1Polarization_type1SecretedPerHourWhenActivated	38.04	38.04	509.8	61.96(−)	100	609.8(+)	18
TCell_apoptosisNaiveMean	38.2	38.2	104.4	18.54(−)	30	61.32(+)	19
Simulation_numDCCNS	39.7	40.88	39.7	23.65(−)	40	55.88(+)	20
Th2Polarization_proliferationMean	41.66	41.66	.	16.8(+)	28.8	.	21
Simulation_numCNS	43.41	43.41	.	283(−)	500	.	22
Simulation_numCD4Th	48.1	48.1	83.79	20.76(−)	40	73.51(+)	23
APC_immatureDurationMean	54.25	54.25	98.42	21.96(+)	48	95.24(−)	24
CNSMacrophage_type1RequiredForActivation	58	76.5	58	0.5875(+)	2.5	3.95(−)	25
Circulation_height	61.06	.	61.06	.	50	80.53(−)	26
TCell_timeLocalActivationDelay	64.22	64.22	258.4	3.578(−)	10	35.84(−)	27
SLO_width	67.65	.	67.65	.	50	83.82(−)	28

Continued on Next Page...

Table C.12 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
CNSMacrophage_basalMBPExpressionProbability	76.59	76.59	175.6	0.04683(–)	0.2	0.5511(+)	29
DendriticCell_phagocytosisProbabilityImmature	76.65	76.65	.	0.2335(–)	1	.	30
DendriticCell_cytokineType2PolarizationRatio	81.25	81.25	.	0.03188(–)	0.17	.	31
Circulation_timeToCrossOrgan	82.13	.	82.13	.	5	9.107(–)	32
SLO_height	82.59	.	82.59	.	50	91.29(–)	33
DendriticCell_type1SecretedPerHourImmunized	88.89	88.89	860.5	1.111(–)	10	96.05(+)	34
Th1Polarization_mhcUnExpressionDelayMean	93.54	93.54	.	0.5165(+)	8	.	35
Simulation_numCD8Treg	94.96	94.96	.	1.513(+)	30	.	36
CD8Treg_cd8TregToCD4HelperSpecificityDropOff	99.8	99.8	.	0.001967(+)	1	.	37
DendriticCell_type1RequiredForActivation	99.93	99.93	.	0.001395(–)	2	.	38
CLN_width	198.4	.	198.4	.	50	149.2(–)	39
Molecule_decayThreshold	752.7	.	752.7	.	0.01	0.08527(–)	40
APC_immatureDurationStdDev	24	.	72
APC_timeOfDeathStdDev	48	.	72
CD4THelper_diff00	0.05	.	72
CD4Treg_type1SecretedPerHourWhenActivated	10	.	72
CD8Treg_type1SecretedPerHourWhenActivated	10	.	72
CLN_height	50	.	72
CLN_timeToCrossOrgan	12	.	72
CNS_timeToCrossOrgan	20	.	72
CNSMacrophage_phagocytosisProbabilityImmature	0.7	.	72
CNSMacrophage_phagocytosisProbabilityMature	0.3	.	72
Circulation_width	50	.	72
DendriticCellMigrates_lengthOfTimeMovingFollowingMigration	3.5	.	72
DendriticCell_phagocytosisProbabilityMature	0.3	.	72
SLO_timeToCrossOrgan	12	.	72
Simulation_immuneActivationLinearInitial	14	.	72
Simulation_numCD4Treg	30	.	72
Simulation_numDC	10	.	72
Simulation_numDCSpleen	100	.	72
Spleen_height	50	.	72

Continued on Next Page...

Table C.12 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>Spleen_timeToCrossOrgan</i>	5	.	72
<i>Spleen_width</i>	50	.	72
<i>TCell_AICDStdDev</i>	56	.	72
<i>TCell_apoptosisNaiveStdDev</i>	17	.	72
<i>TCell_apoptosisPartialMaturityMean</i>	12	.	72
<i>TCell_apoptosisPartialMaturityStdDev</i>	6	.	72
<i>TCell_becomeEffectorStdDev</i>	56	.	72
<i>TCell_proliferationStdDev</i>	9.6	.	72
<i>TCell_specificityLowerLimit</i>	0.5	.	72
<i>TCell_specificityUpperLimit</i>	0.9	.	72
<i>Th1Polarization_mhcUnExpressionDelayStdDev</i>	2	.	72
<i>Th2Polarization_proliferationStdDev</i>	19.2	.	72
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	100	.	72

Table C.13: Robustness indexes for parameters with respect to the *CD4Th1 Max Time* response. RI, robustness index; LI, lower index; UI, upper index; LB, lower boundary; DV, default value; UB, upper boundary. Results indicating no significant deviation in behaviour are marked with a period.

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>TCell_proliferationMean</i>	9.177	9.177	11.47	17.44(–)	19.2	21.4(+)	1
<i>CNSMacrophage_tnfaSecretedPerHourWhenStimulated</i>	11.97	11.97	19.41	88.03(–)	100	119.4(+)	2
<i>CNSCell_apoptosisTNFaThreshold</i>	13.78	14.78	13.78	4.261(+)	5	5.689(–)	3
<i>CNS_height</i>	14.79	16.28	14.79	41.86(+)	50	57.4(–)	4
<i>Simulation_immunizationLinearDC0</i>	16.71	17.84	16.71	1.643(–)	2	2.334(–)	5
<i>Molecule_molecularHalflife</i>	16.73	16.73	31.33	0.4164(–)	0.5	0.6566(+)	6
<i>Simulation_immunizationLinearGradient</i>	22.76	22.76	38.57	-0.006138(–)	-0.005	-0.003071(–)	7
<i>CNS_width</i>	23.01	.	23.01	.	50	61.51(–)	8
<i>Simulation_numCNSMacrophage</i>	23.71	23.71	29.6	57.22(–)	75	97.2(+)	9
<i>TCell_becomeEffectorMean</i>	24.94	24.94	45.57	45.04(+)	60	87.34(–)	10

Continued on Next Page...

Table C.13 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>APC_timeOfDeathMean</i>	26.84	.	26.84	.	110	139.5(–)	11
<i>Simulation_numCD4Th</i>	41.43	41.43	55.51	23.43(+)	40	62.2(–)	12
<i>Th2Polarization_proliferationMean</i>	42.96	42.96	.	16.43(–)	28.8	.	13
<i>SLO_width</i>	44.32	44.32	51.27	27.84(–)	50	75.64(+)	14
<i>Circulation_height</i>	46.95	46.95	.	26.53(–)	50	.	15
<i>CD4THelper_diff08</i>	47.23	47.23	.	0.4485(+)	0.85	.	16
<i>TCell.cellsPerGridspace</i>	47.96	47.96	.	3.642(–)	7	.	17
<i>Simulation_numDCCNS</i>	51.7	51.7	.	19.32(–)	40	.	18
<i>TCell_apoptosisNaiveMean</i>	53.26	53.26	213.3	14.02(–)	30	93.98(+)	19
<i>Simulation_numCNS</i>	56.83	56.83	.	215.9(–)	500	.	20
<i>CLN_width</i>	57.31	57.31	.	21.34(–)	50	.	21
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	62.49	62.49	.	37.51(–)	100	.	22
<i>TCell.timeLocalActivationInducedEffectorFunctionFor</i>	65.05	65.05	.	16.78(–)	48	.	23
<i>APC_probabilityPhagocytosisToPeptide</i>	65.77	65.77	.	0.006846(–)	0.02	.	24
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	69.2	69.2	.	2.464(+)	8	.	25
<i>TCell.timeLocalActivationDelay</i>	75.41	75.41	364	2.459(–)	10	46.4(–)	26
<i>Simulation_numCD8Treg</i>	76.89	76.89	249.6	6.933(+)	30	104.9(–)	27
<i>Simulation_numCD4Treg</i>	83.07	83.07	496.2	5.079(+)	30	178.9(–)	28
<i>DendriticCell_cytokineType2PolarizationRatio</i>	83.8	83.8	.	0.02755(–)	0.17	.	29
<i>TCell_specificityLowerLimit</i>	86.34	86.34	.	0.06832(+)	0.5	.	30
<i>DendriticCell_phagocytosisProbabilityImmature</i>	91.47	91.47	.	0.08532(–)	1	.	31
<i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i>	92.76	92.76	.	0.07239(+)	1	.	32
<i>CNSMacrophage_basalMBPExpressionProbability</i>	95.18	95.18	.	0.009648(–)	0.2	.	33
<i>DendriticCell_type1RequiredForActivation</i>	99.77	99.77	.	0.00461(–)	2	.	34
<i>CNSMacrophage_type1RequiredForActivation</i>	127.5	.	127.5	.	2.5	5.687(–)	35
<i>Simulation_immuneActivationLinearInitial</i>	937.5	.	937.5	.	14	145.3(–)	36
<i>TCell_AICDMean</i>	60	.	72
<i>Simulation_immuneActivationLinearFreq</i>	6	.	72
<i>APC_immatureDurationMean</i>	48	.	72
<i>Circulation_timeToCrossOrgan</i>	5	.	72
<i>SLO_height</i>	50	.	72

Continued on Next Page...

Table C.13 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
DendriticCell_type1SecretedPerHourImmunized	10	.	72
Molecule_decayThreshold	0.01	.	72
APC_immatureDurationStdDev	24	.	72
APC_timeOfDeathStdDev	48	.	72
CD4THelper_diff00	0.05	.	72
CD4Treg_type1SecretedPerHourWhenActivated	10	.	72
CD8Treg_type1SecretedPerHourWhenActivated	10	.	72
CLN_height	50	.	72
CLN_timeToCrossOrgan	12	.	72
CNS_timeToCrossOrgan	20	.	72
CNSMacrophage_phagocytosisProbabilityImmature	0.7	.	72
CNSMacrophage_phagocytosisProbabilityMature	0.3	.	72
Circulation_width	50	.	72
DendriticCellMigrates_lengthOfTimeMovingFollowingMigration	3.5	.	72
DendriticCell_phagocytosisProbabilityMature	0.3	.	72
SLO_timeToCrossOrgan	12	.	72
Simulation_numDC	10	.	72
Simulation_numDCSpleen	100	.	72
Spleen_height	50	.	72
Spleen_timeToCrossOrgan	5	.	72
Spleen_width	50	.	72
TCell_AICDStdDev	56	.	72
TCell_apoptosisNaiveStdDev	17	.	72
TCell_apoptosisPartialMaturityMean	12	.	72
TCell_apoptosisPartialMaturityStdDev	6	.	72
TCell_becomeEffectortStdDev	56	.	72
TCell_proliferationStdDev	9.6	.	72
TCell_specificityUpperLimit	0.9	.	72
Th1Polarization_mhcUnExpressionDelayStdDev	2	.	72
Th2Polarization_proliferationStdDev	19.2	.	72
Th2Polarization_type2SecretedPerHourWhenActivated	100	.	72

Table C.14: Robustness indexes for parameters with respect to the *CD4 Th2 Max* response. RI, robustness index; LI, lower index; UI, upper index; LB, lower boundary; DV, default value; UB, upper boundary. Results indicating no significant deviation in behaviour are marked with a period.

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>CD4THelper.diff08</i>	2.688	3.205	2.688	0.8228(+)	0.85	0.8729(−)	1
<i>CNSMacrophage_tnfaSecretedPerHourWhenStimulated</i>	2.992	2.992	3.343	97.01(−)	100	103.3(+)	2
<i>CNSCell_apoptosisTNFaThreshold</i>	3.267	3.267	4.1	4.837(+)	5	5.205(−)	3
<i>CNS_height</i>	4.018	4.018	4.854	47.99(+)	50	52.43(−)	4
<i>Molecule_molecularHalflife</i>	4.171	4.171	5.513	0.4791(−)	0.5	0.5276(+)	5
<i>CNS_width</i>	4.657	4.657	6.711	47.67(+)	50	53.36(−)	6
<i>Simulation_numCNSMacrophage</i>	6.788	7.189	6.788	69.61(−)	75	80.09(+)	7
<i>APC_timeOfDeathMean</i>	7.617	7.617	9.935	101.6(−)	110	120.9(+)	8
<i>TCell_AICDMean</i>	8.544	8.544	10.33	54.87(−)	60	66.2(+)	9
<i>Simulation_immuneActivationLinearDC0</i>	9.022	9.022	10.91	1.82(−)	2	2.218(+)	10
<i>TCell_proliferationMean</i>	9.809	11.93	9.809	16.91(+)	19.2	21.08(−)	11
<i>TCell_cellsPerGridspace</i>	10.41	10.41	.	6.271(−)	7	.	12
<i>Simulation_immuneActivationLinearGradient</i>	15.16	17.62	15.16	-0.005881(−)	-0.005	-0.004242(+)	13
<i>Simulation_immuneActivationLinearFreq</i>	16.13	16.13	37.34	5.032(+)	6	8.24(−)	14
<i>APC_probabilityPhagocytosisToPeptide</i>	16.92	16.92	19.48	0.01662(−)	0.02	0.0239(+)	15
<i>TCell_timeLocalActivationInducedEffectorFunctionFor</i>	19.02	19.02	65.13	38.87(−)	48	79.26(+)	16
<i>TCell_becomeEffectorMean</i>	20.64	29.22	20.64	42.47(−)	60	72.38(−)	17
<i>Simulation_numDCCNS</i>	23.84	26.32	23.84	29.47(−)	40	49.53(+)	18
<i>Simulation_numCNS</i>	27.59	27.59	93.45	362(−)	500	967.3(−)	19
<i>Circulation_height</i>	36.99	.	36.99	.	50	68.49(−)	20
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	38.01	38.01	.	61.99(−)	100	.	21
<i>Th2Polarization_proliferationMean</i>	45.21	45.21	.	15.78(+)	28.8	.	22
<i>DendriticCell_cytokineType2PolarizationRatio</i>	49.16	49.16	.	0.08643(+)	0.17	.	23
<i>TCell_apoptosisNaiveMean</i>	49.23	49.23	156.7	15.23(−)	30	77.02(+)	24
<i>CNSMacrophage_type1RequiredForActivation</i>	50.32	54.99	50.32	1.125(+)	2.5	3.758(−)	25
<i>CLN_width</i>	51.53	51.53	.	24.23(+)	50	.	26

Continued on Next Page...

Table C.14 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
DendriticCell_phagocytosisProbabilityImmature	53.31	53.31	.	0.4669(–)	1	.	27
APC_immatureDurationMean	53.55	53.55	.	22.29(+)	48	.	28
Simulation_numCD4Th	54.33	54.33	103.4	18.27(–)	40	81.36(+)	29
SLO_width	56.01	56.01	97.02	22(+)	50	98.51(–)	30
TCell_timeLocalActivationDelay	58.56	58.56	212	4.144(+)	10	31.2(–)	31
Circulation_timeToCrossOrgan	59.66	.	59.66	.	5	7.983(–)	32
DendriticCell_type1SecretedPerHourImmunized	60.43	60.43	.	3.957(+)	10	.	33
CNSMacrophage_basalMBPExpressionProbability	78.32	78.32	132.1	0.04335(–)	0.2	0.4643(+)	34
CNSMacrophage_phagocytosisProbabilityMature	78.38	78.38	.	0.06485(+)	0.3	.	35
Th2Polarization_type2SecretedPerHourWhenActivated	87.03	92.12	87.03	7.879(–)	100	187(+)	36
Spleen_height	160	.	160	.	50	130(–)	37
Molecule_decayThreshold	235.2	.	235.2	.	0.01	0.03352(+)	38
CD4THelper_diff00	1345	.	1345	.	0.05	0.7226(–)	39
Th1Polarization_mhcUnExpressionDelayMean	8	.	72
Simulation_numCD8Treg	30	.	72
Simulation_numCD4Treg	30	.	72
TCell_specificityLowerLimit	0.5	.	72
CD8Treg_cd8TregToCD4THelperSpecificityDropOff	1	.	72
DendriticCell_type1RequiredForActivation	2	.	72
Simulation_immuneActivationLinearInitial	14	.	72
SLO_height	50	.	72
APC_immatureDurationStdDev	24	.	72
APC_timeOfDeathStdDev	48	.	72
CD4Treg_type1SecretedPerHourWhenActivated	10	.	72
CD8Treg_type1SecretedPerHourWhenActivated	10	.	72
CLN_height	50	.	72
CLN_timeToCrossOrgan	12	.	72
CNS_timeToCrossOrgan	20	.	72
CNSMacrophage_phagocytosisProbabilityImmature	0.7	.	72
Circulation_width	50	.	72
DendriticCellMigrates_lengthOfTimeMovingFollowingMigration	3.5	.	72

Continued on Next Page...

Table C.14 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
DendriticCell_phagocytosisProbabilityMature	0.3	.	72
SLO_timeToCrossOrgan	12	.	72
Simulation_numDC	10	.	72
Simulation_numDCSpleen	100	.	72
Spleen_timeToCrossOrgan	5	.	72
Spleen_width	50	.	72
TCell_AICDStdDev	56	.	72
TCell_apoptosisNaiveStdDev	17	.	72
TCell_apoptosisPartialMaturityMean	12	.	72
TCell_apoptosisPartialMaturityStdDev	6	.	72
TCell_becomeEffectorStdDev	56	.	72
TCell_proliferationStdDev	9.6	.	72
TCell_specificityUpperLimit	0.9	.	72
Th1Polarization_mhcUnExpressionDelayStdDev	2	.	72
Th2Polarization_proliferationStdDev	19.2	.	72

Table C.15: Robustness indexes for parameters with respect to the *CD4Th2 Max Time* response. RI, robustness index; LI, lower index; UI, upper index; LB, lower boundary; DV, default value; UB, upper boundary. Results indicating no significant deviation in behaviour are marked with a period.

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
CNSCell_apoptosisTNFaThreshold	5.2	5.2	6.986	4.74(+)	5	5.349(–)	1
CNSMacrophage_tnfSecretedPerHourWhenStimulated	5.742	5.742	6.124	94.26(–)	100	106.1(+)	2
CNS_height	6.717	7.156	6.717	46.42(+)	50	53.36(–)	3
CNS_width	7.021	7.021	11.46	46.49(+)	50	55.73(–)	4
Molecule_molecularHalflife	7.759	7.759	8.434	0.4612(–)	0.5	0.5422(+)	5
Simulation_numCNSMacrophage	13.13	13.13	13.49	65.15(–)	75	85.12(+)	6
TCell_AICDMean	24.62	24.62	32.68	45.23(–)	60	79.61(+)	7
APC_probabilityPhagocytosisToPeptide	25.85	25.85	58.4	0.01483(–)	0.02	0.03168(+)	8

Continued on Next Page...

Table C.15 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>TCell</i> _cellsPerGridspace	33.17	33.17	.	4.678(–)	7	.	9
<i>APC</i> _timeOfDeathMean	35.8	35.8	.	70.62(–)	110	.	10
<i>Simulation</i> _immunizationLinearFreq	36.11	36.11	.	3.834(–)	6	.	11
<i>Simulation</i> _numCNS	38.11	38.11	96.73	309.4(–)	500	983.6(–)	12
<i>Simulation</i> _numDCCNS	39.09	39.09	57.43	24.37(–)	40	62.97(+)	13
<i>Simulation</i> _immunizationLinearDC0	43.58	43.58	.	1.128(+)	2	.	14
<i>Simulation</i> _immunizationLinearGradient	44.1	.	44.1	.	-0.005	-0.002795(+)	15
<i>Th1Polarization</i> _type1SecretedPerHourWhenActivated	44.88	44.88	.	55.12(–)	100	.	16
<i>TCell</i> _becomeEffectorMean	47.5	47.5	.	31.5(–)	60	.	17
<i>TCell</i> _timeLocalActivationInducedEffectorFunctionFor	49.56	49.56	.	24.21(–)	48	.	18
<i>DendriticCell</i> _phagocytosisProbabilityImmature	70.32	70.32	.	0.2968(–)	1	.	19
<i>Simulation</i> _numCD4Th	74.73	74.73	.	10.11(+)	40	.	20
<i>CNSMacrophage</i> _type1RequiredForActivation	75.85	86.19	75.85	0.3452(+)	2.5	4.396(–)	21
<i>CNSMacrophage</i> _phagocytosisProbabilityMature	90.65	90.65	.	0.02806(+)	0.3	.	22
<i>CD4THelper</i> _diff08	92.33	92.33	.	0.06518(–)	0.85	.	23
<i>Th2Polarization</i> _type2SecretedPerHourWhenActivated	94.89	94.89	.	5.113(–)	100	.	24
<i>CNSMacrophage</i> _basalMBPExpressionProbability	96.19	96.19	387.3	0.007622(–)	0.2	0.9745(+)	25
<i>SLO</i> _width	162.8	.	162.8	.	50	131.4(+)	26
<i>Circulation</i> _height	182.6	.	182.6	.	50	141.3(–)	27
<i>TCell</i> _timeLocalActivationDelay	325.5	.	325.5	.	10	42.55(–)	28
<i>DendriticCell</i> _type1SecretedPerHourImmunized	1834	.	1834	.	10	193.4(–)	29
<i>TCell</i> _proliferationMean	19.2	.	72
<i>Th2Polarization</i> _proliferationMean	28.8	.	72
<i>DendriticCell</i> _cytokineType2PolarizationRatio	0.17	.	72
<i>TCell</i> _apoptosisNaiveMean	30	.	72
<i>CLN</i> _width	50	.	72
<i>APC</i> _immatureDurationMean	48	.	72
<i>Circulation</i> _timeToCrossOrgan	5	.	72
<i>Spleen</i> _height	50	.	72
<i>Molecule</i> _decayThreshold	0.01	.	72
<i>CD4THelper</i> _diff00	0.05	.	72

Continued on Next Page...

Table C.15 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
Th1Polarization_mhcUnExpressionDelayMean	8	.	72
Simulation_numCD8Treg	30	.	72
Simulation_numCD4Treg	30	.	72
TCell_specificityLowerLimit	0.5	.	72
CD8Treg_cd8TregToCD4HelperSpecificityDropOff	1	.	72
DendriticCell_type1RequiredForActivation	2	.	72
Simulation_immuneActivationLinearInitial	14	.	72
SLO_height	50	.	72
APC_immatureDurationStdDev	24	.	72
APC_timeOfDeathStdDev	48	.	72
CD4Treg_type1SecretedPerHourWhenActivated	10	.	72
CD8Treg_type1SecretedPerHourWhenActivated	10	.	72
CLN_height	50	.	72
CLN_timeToCrossOrgan	12	.	72
CNS_timeToCrossOrgan	20	.	72
CNSMacrophage_phagocytosisProbabilityImmature	0.7	.	72
Circulation_width	50	.	72
DendriticCellMigrates_lengthOfTimeMovingFollowingMigration	3.5	.	72
DendriticCell_phagocytosisProbabilityMature	0.3	.	72
SLO_timeToCrossOrgan	12	.	72
Simulation_numDC	10	.	72
Simulation_numDCSpleen	100	.	72
Spleen_timeToCrossOrgan	5	.	72
Spleen_width	50	.	72
TCell_AICDStdDev	56	.	72
TCell_apoptosisNaiveStdDev	17	.	72
TCell_apoptosisPartialMaturityMean	12	.	72
TCell_apoptosisPartialMaturityStdDev	6	.	72
TCell_becomeEffectorStdDev	56	.	72
TCell_proliferationStdDev	9.6	.	72
TCell_specificityUpperLimit	0.9	.	72

Continued on Next Page...

Table C.15 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
Th1Polarization_mhcUnExpressionDelayStdDev	2	.	72
Th2Polarization_proliferationStdDev	19.2	.	72

Table C.16: Robustness indexes for parameters with respect to the *CD4 Treg Max* response. RI, robustness index; LI, lower index; UI, upper index; LB, lower boundary; DV, default value; UB, upper boundary. Results indicating no significant deviation in behaviour are marked with a period.

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
TCell_proliferationMean	3.539	3.539	3.799	18.52(+)	19.2	19.93(−)	1
APC_timeOfDeathMean	4.273	4.273	4.392	105.3(−)	110	114.8(+)	2
CNSMacrophage_tnfaSecretedPerHourWhenStimulated	4.842	6.025	4.842	93.97(−)	100	104.8(+)	3
CNSCell_apoptosisTNFaThreshold	5.274	5.274	6.125	4.736(+)	5	5.306(−)	4
CNS_height	5.822	5.822	6.523	47.09(+)	50	53.26(−)	5
Molecule_molecularHalflife	6.794	6.933	6.794	0.4653(−)	0.5	0.534(+)	6
CNS_width	7.202	7.202	9.32	46.4(+)	50	54.66(−)	7
TCell_cellsPerGridspace	8.562	8.562	.	6.401(−)	7	.	8
Simulation_immunizationLinearDC0	8.725	8.725	9.881	1.825(−)	2	2.198(+)	9
CD4THelper_diff08	8.844	8.844	9.845	0.7748(−)	0.85	0.9337(+)	10
TCell_becomeEffectorMean	9.525	9.525	21.13	54.28(−)	60	72.68(+)	11
APC_probabilityPhagocytosisToPeptide	9.778	9.778	10.68	0.01804(−)	0.02	0.02214(+)	12
TCell_AICDMean	12.1	12.1	13.02	52.74(−)	60	67.81(+)	13
Simulation_immunizationLinearGradient	12.95	18.51	12.95	-0.005926(−)	-0.005	-0.004353(+)	14
Simulation_numCNSMacrophage	13.18	13.25	13.18	65.06(−)	75	84.89(+)	15
Simulation_numDCCNS	16.67	16.67	20.75	33.33(−)	40	48.3(+)	16
TCell_timeLocalActivationInducedEffectorFunctionFor	17.56	17.56	48.79	39.57(−)	48	71.42(+)	17
Simulation_immunizationLinearFreq	26.81	26.81	30.71	4.392(+)	6	7.843(−)	18
Simulation_numCD4Treg	30.89	30.89	32.79	20.73(−)	30	39.84(+)	19
APC_immatureDurationMean	32.42	32.42	37.5	32.44(+)	48	66(−)	20

Continued on Next Page...

Table C.16 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>Circulation_height</i>	32.55	.	32.55	.	50	66.28(–)	21
<i>Th2Polarization_proliferationMean</i>	34.15	34.15	.	18.96(+) 327.1(–)	28.8	.	22
<i>Simulation_numCNS</i>	34.59	34.59	72.22	500	861.1(–)	23	
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	35.32	35.32	693.8	100	793.8(+) 64.68(–)	24	
<i>Simulation_numCD8Treg</i>	37.63	37.63	44.89	18.71(+) 0.5328(–)	30	43.47(–)	25
<i>TCell_specificityUpperLimit</i>	40.8	40.8	.	0.9	.	.	26
<i>Circulation_timeToCrossOrgan</i>	44.98	.	44.98	5	7.249(–)	27	
<i>CLN_height</i>	45.65	45.65	74.06	27.18(+) 50	87.03(–)	28	
<i>Simulation_numCD4Th</i>	45.91	45.91	84.8	21.64(–) 26.39(–)	40	73.92(+) 50	29
<i>Spleen_width</i>	47.22	47.22	141.1	25.77(+) 25.59(+) 50	120.6(–)	30	
<i>CLN_width</i>	48.45	48.45	49.1	50	74.55(–)	31	
<i>SLO_width</i>	48.82	48.82	99.03	50	99.51(–)	32	
<i>Simulation_numDCSpleen</i>	50.17	50.17	.	49.83(–) 100	.	.	33
<i>TCell_apoptosisNaiveMean</i>	50.61	50.61	.	14.82(–) 30	.	.	34
<i>TCell_becomeEffectorStdDev</i>	52.67	52.67	.	26.51(+) 56	.	.	35
<i>DendriticCell_phagocytosisProbabilityImmature</i>	54.63	54.63	.	0.4537(–) 1	.	.	36
<i>CNSMacrophage_type1RequiredForActivation</i>	71.6	92.85	71.6	0.1787(+) 0.1782(–)	2.5	4.29(–) 0.17	37
<i>TCell_proliferationStdDev</i>	75.77	.	75.77	.	9.6	16.87(+) 0.17	38
<i>DendriticCell_cytokineType2PolarizationRatio</i>	77.75	77.75	.	0.03782(–) 0.03485(–)	0.17	.	39
<i>CNSMacrophage_basalMBPExpressionProbability</i>	82.58	82.58	355.4	0.2	0.9108(+) 0.5	.	40
<i>CNSMacrophage_phagocytosisProbabilityMature</i>	83.66	83.66	.	0.04901(+) 0.05732(–)	0.3	.	41
<i>DendriticCell_type1SecretedPerHourImmunized</i>	84.73	84.73	627	1.527(–) 0.08541(+) 0.05732(–)	10	72.7(+) 0.7 0.5	42
<i>CNSMacrophage_phagocytosisProbabilityImmature</i>	87.8	87.8	.	0.08541(+) 0.05732(–)	0.7	.	43
<i>TCell_specificityLowerLimit</i>	88.54	88.54	.	0.05732(–)	0.5	.	44
<i>TCell_timeLocalActivationDelay</i>	88.58	88.58	219.6	1.142(–) 0.023(+) 0.05732(–)	10	31.96(–) 2 0.01 0.05	45
<i>DendriticCell_type1RequiredForActivation</i>	98.85	98.85	643.9	0.023(+) 0.05732(–)	2	14.88(–) 50 10 0.01 0.05	46
<i>Spleen_height</i>	157.9	.	157.9	.	50	129(–) 10 44.21(+) 0.05	47
<i>Simulation_numDC</i>	342.1	.	342.1	.	10	44.21(+) 0.05	48
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	369.1	.	369.1	.	100	469.1(–) 0.01 0.05	49
<i>Molecule_decayThreshold</i>	664.7	.	664.7	.	0.01	0.07647(–) 0.05	50
<i>CD4THelper_diff00</i>	72

Continued on Next Page...

Table C.16 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
Th1Polarization_mhcUnExpressionDelayMean	8	.	72
CD8Treg_cd8TregToCD4THelperSpecificityDropOff	1	.	72
Simulation_immunizationLinearInitial	14	.	72
SLO_height	50	.	72
APC_immatureDurationStdDev	24	.	72
APC_timeOfDeathStdDev	48	.	72
CD4Treg_type1SecretedPerHourWhenActivated	10	.	72
CD8Treg_type1SecretedPerHourWhenActivated	10	.	72
CLN_timeToCrossOrgan	12	.	72
CNS_timeToCrossOrgan	20	.	72
Circulation_width	50	.	72
DendriticCellMigrates_lengthOfTimeMovingFollowingMigration	3.5	.	72
DendriticCell_phagocytosisProbabilityMature	0.3	.	72
SLO_timeToCrossOrgan	12	.	72
Spleen_timeToCrossOrgan	5	.	72
TCell_AICDStdDev	56	.	72
TCell_apoptosisNaiveStdDev	17	.	72
TCell_apoptosisPartialMaturityMean	12	.	72
TCell_apoptosisPartialMaturityStdDev	6	.	72
Th1Polarization_mhcUnExpressionDelayStdDev	2	.	72
Th2Polarization_proliferationStdDev	19.2	.	72

Table C.17: Robustness indexes for parameters with respect to the *CD4 Treg Max Time* response. RI, robustness index; LI, lower index; UI, upper index; LB, lower boundary; DV, default value; UB, upper boundary. Results indicating no significant deviation in behaviour are marked with a period.

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
TCell_proliferationMean	5.584	5.584	5.836	18.13(–)	19.2	20.32(+)	1
CNSMacrophage_tnfaSecretedPerHourWhenStimulated	9.006	9.006	14.99	90.99(–)	100	115(+)	2

Continued on Next Page...

Table C.17 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>CNS_height</i>	10.99	17.94	10.99	41.03(+)	50	55.49(–)	3
<i>CNSCell_apoptosisTNFaThreshold</i>	11.88	13.28	11.88	4.336(+)	5	5.594(–)	4
<i>Molecule_molecularHalflife</i>	12.8	12.8	19.96	0.436(–)	0.5	0.5998(+)	5
<i>TCell_becomeEffectorMean</i>	14.98	14.98	.	51.01(+)	60	.	6
<i>CNS_width</i>	16.4	16.4	16.84	41.8(+)	50	58.42(–)	7
<i>Simulation_numCNSMacrophage</i>	18.65	18.65	22.78	61.01(–)	75	92.09(+)	8
<i>Simulation_immunizationLinearDC0</i>	22.22	22.22	.	1.556(+)	2	.	9
<i>Spleen_width</i>	32.11	32.11	63.56	33.94(–)	50	81.78(+)	10
<i>Simulation_immunizationLinearFreq</i>	33.46	33.46	.	3.992(–)	6	.	11
<i>Simulation_immunizationLinearGradient</i>	36.71	60.12	36.71	-0.008006(+)	-0.005	-0.003164(+)	12
<i>TCell_cellsPerGridspace</i>	39.2	39.2	.	4.256(–)	7	.	13
<i>CLN_width</i>	39.93	39.93	.	30.03(–)	50	.	14
<i>Simulation_numCD4Th</i>	42.25	42.25	77.48	23.1(+)	40	70.99(–)	15
<i>APC_timeOfDeathMean</i>	42.36	42.36	.	63.4(–)	110	.	16
<i>Circulation_height</i>	43.99	49.09	43.99	25.45(–)	50	72(+)	17
<i>SLO_width</i>	44.1	44.1	62.38	27.95(–)	50	81.19(+)	18
<i>Simulation_numCNS</i>	45.09	45.09	.	274.6(–)	500	.	19
<i>Simulation_numCD4Treg</i>	47.23	47.23	69.81	15.83(+)	30	50.94(–)	20
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	53.98	53.98	.	46.02(–)	100	.	21
<i>TCell_apoptosisNaiveMean</i>	54.93	54.93	71.98	13.52(–)	30	51.59(+)	22
<i>TCell_timeLocalActivationDelay</i>	57.11	57.11	.	4.289(–)	10	.	23
<i>TCell_specificityLowerLimit</i>	60.03	60.03	.	0.1999(+)	0.5	.	24
<i>APC_probabilityPhagocytosisToPeptide</i>	60.98	60.98	.	0.007804(–)	0.02	.	25
<i>Simulation_numDCCNS</i>	63.58	63.58	89.9	14.57(–)	40	75.96(+)	26
<i>TCell_timeLocalActivationInducedEffectorFunctionFor</i>	68.72	68.72	.	15.02(–)	48	.	27
<i>Simulation_numCD8Treg</i>	73.86	73.86	114.8	7.841(+)	30	64.44(–)	28
<i>Circulation_timeToCrossOrgan</i>	78.82	.	78.82	.	5	8.941(+)	29
<i>DendriticCell_cytokineType2PolarizationRatio</i>	82.52	82.52	.	0.02972(–)	0.17	.	30
<i>TCell_proliferationStdDev</i>	86.54	.	86.54	.	9.6	17.91(–)	31
<i>CD4THelper_diff08</i>	90.82	90.82	.	0.078(–)	0.85	.	32
<i>DendriticCell_type1SecretedPerHourImmunized</i>	91.87	91.87	.	0.8128(–)	10	.	33

Continued on Next Page...

Table C.17 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
DendriticCell_phagocytosisProbabilityImmature	92.77	92.77	.	0.0723(–)	1	.	34
DendriticCell_type1RequiredForActivation	97.67	97.67	578.8	0.04653(–)	2	13.58(+)	35
CNSMacrophage_type1RequiredForActivation	108.3	.	108.3	.	2.5	5.207(–)	36
Spleen_height	131.6	.	131.6	.	50	115.8(+)	37
Th2Polarization_type2SecretedPerHourWhenActivated	495.3	.	495.3	.	100	595.3(–)	38
Molecule_decayThreshold	3561	.	3561	.	0.01	0.3661(–)	39
TCell_AICDMean	60	.	72
APC_immatureDurationMean	48	.	72
Th2Polarization_proliferationMean	28.8	.	72
TCell_specificityUpperLimit	0.9	.	72
CLN_height	50	.	72
Simulation_numDCSpleen	100	.	72
TCell_becomeEffectorStdDev	56	.	72
CNSMacrophage_basalMBPExpressionProbability	0.2	.	72
CNSMacrophage_phagocytosisProbabilityMature	0.3	.	72
CNSMacrophage_phagocytosisProbabilityImmature	0.7	.	72
Simulation_numDC	10	.	72
CD4THelper_diff00	0.05	.	72
Th1Polarization_mhcUnExpressionDelayMean	8	.	72
CD8Treg_cd8TregToCD4THelperSpecificityDropOff	1	.	72
Simulation_immunizationLinearInitial	14	.	72
SLO_height	50	.	72
APC_immatureDurationStdDev	24	.	72
APC_timeOfDeathStdDev	48	.	72
CD4Treg_type1SecretedPerHourWhenActivated	10	.	72
CD8Treg_type1SecretedPerHourWhenActivated	10	.	72
CLN_timeToCrossOrgan	12	.	72
CNS_timeToCrossOrgan	20	.	72
Circulation_width	50	.	72
DendriticCellMigrates_lengthOfTimeMovingFollowingMigration	3.5	.	72

Continued on Next Page...

Table C.17 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
DendriticCell_phagocytosisProbabilityMature	0.3	.	72
SLO_timeToCrossOrgan	12	.	72
Spleen_timeToCrossOrgan	5	.	72
TCell_AICDStdDev	56	.	72
TCell_apoptosisNaiveStdDev	17	.	72
TCell_apoptosisPartialMaturityMean	12	.	72
TCell_apoptosisPartialMaturityStdDev	6	.	72
Th1Polarization_mhcUnExpressionDelayStdDev	2	.	72
Th2Polarization_proliferationStdDev	19.2	.	72

Table C.18: Robustness indexes for parameters with respect to the *CD8Treg Max* response. RI, robustness index; LI, lower index; UI, upper index; LB, lower boundary; DV, default value; UB, upper boundary. Results indicating no significant deviation in behaviour are marked with a period.

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
APC_timeOfDeathMean	5.794	5.794	6.062	103.6(–)	110	116.7(+)	1
TCell_proliferationMean	5.881	8.698	5.881	17.53(+)	19.2	20.33(–)	2
CNSCell_apoptosisTNFaThreshold	6.628	6.628	7.897	4.669(+)	5	5.395(–)	3
CNSMacrophage_tnfSecretedPerHourWhenStimulated	6.674	6.752	6.674	93.25(–)	100	106.7(+)	4
CNS_height	7.4	7.4	7.626	46.3(+)	50	53.81(–)	5
CNS_width	8.118	8.118	10.5	45.94(+)	50	55.25(–)	6
Molecule_molecularHalflife	8.731	8.731	9.419	0.4563(–)	0.5	0.5471(+)	7
TCell_cellsPerGridspace	9.613	9.613	.	6.327(–)	7	.	8
CD4THelper_diff08	10.18	10.18	11.72	0.7635(–)	0.85	0.9497(+)	9
APC_probabilityPhagocytosisToPeptide	12.01	12.01	14.29	0.0176(–)	0.02	0.02286(+)	10
Simulation_immuneActivationLinearDC0	12.88	12.99	12.88	1.74(–)	2	2.258(+)	11
TCell_AICDMean	13.92	13.92	14.77	51.65(–)	60	68.86(+)	12
Simulation_numCNSMacrophage	13.97	13.97	15.22	64.52(–)	75	86.42(+)	13
Simulation_immuneActivationLinearGradient	16.24	22.65	16.24	-0.006133(–)	-0.005	-0.004188(+)	14

Continued on Next Page...

Table C.18 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>TCell.becomeEffectorMean</i>	17.27	17.27	23.7	49.64(–)	60	74.22(–)	15
<i>TCell.timeLocalActivationInducedEffectorFunctionFor</i>	21.95	21.95	.	37.46(–)	48	.	16
<i>Simulation_numDCCNS</i>	22.29	22.29	36.09	31.08(–)	40	54.43(+)	17
<i>Simulation_numCD8Treg</i>	24.09	24.09	27.19	22.77(–)	30	38.16(+)	18
<i>Simulation_immuneActivationLinearFreq</i>	35.57	35.57	60.17	3.866(+)	6	9.61(–)	19
<i>Th2Polarization_proliferationMean</i>	36.36	36.36	.	18.33(+)	28.8	.	20
<i>Simulation_numCNS</i>	36.73	36.73	84.14	316.3(–)	500	920.7(–)	21
<i>Circulation_height</i>	39.21	.	39.21	.	50	69.6(–)	22
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	39.64	39.64	.	60.36(–)	100	.	23
<i>TCell_specificityUpperLimit</i>	42.35	42.35	.	0.5189(–)	0.9	.	24
<i>APC_immatureDurationMean</i>	50.52	50.52	55.64	23.75(+)	48	74.71(–)	25
<i>Simulation_numDCSpleen</i>	50.76	50.76	.	49.24(–)	100	.	26
<i>Simulation_numCD4Treg</i>	55.04	59.46	55.04	12.16(+)	30	46.51(–)	27
<i>Simulation_numCD4Th</i>	57.2	57.2	121.2	17.12(–)	40	88.46(+)	28
<i>Circulation_timeToCrossOrgan</i>	57.61	.	57.61	.	5	7.881(–)	29
<i>CLN_width</i>	59.73	59.73	63.65	20.14(+)	50	81.83(–)	30
<i>DendriticCell_phagocytosisProbabilityImmature</i>	61.3	61.3	.	0.387(–)	1	.	31
<i>TCell_specificityLowerLimit</i>	73.45	73.45	.	0.1327(–)	0.5	.	32
<i>DendriticCell_cytokineType2PolarizationRatio</i>	79.53	79.53	.	0.03481(–)	0.17	.	33
<i>CNSMacrophage_type1RequiredForActivation</i>	80.11	.	80.11	.	2.5	4.503(–)	34
<i>Spleen_width</i>	81.03	.	81.03	.	50	90.52(–)	35
<i>TCell_timeLocalActivationDelay</i>	81.86	81.86	275.6	1.814(–)	10	37.56(–)	36
<i>CNSMacrophage_basalMBPExpressionProbability</i>	85.4	85.4	.	0.02919(–)	0.2	.	37
<i>DendriticCell_type1SecretedPerHourImmunized</i>	86.42	86.42	.	1.358(–)	10	.	38
<i>DendriticCell_phagocytosisProbabilityMature</i>	89.19	89.19	.	0.03243(+)	0.3	.	39
<i>CLN_height</i>	114.8	.	114.8	.	50	107.4(–)	40
<i>SLO_height</i>	131.4	.	131.4	.	50	115.7(–)	41
<i>SLO_width</i>	137.8	.	137.8	.	50	118.9(–)	42
<i>Spleen_height</i>	163.1	.	163.1	.	50	131.5(–)	43
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	409.8	.	409.8	.	100	509.8(–)	44
<i>Simulation_numDC</i>	467	.	467	.	10	56.7(+)	45

Continued on Next Page...

Table C.18 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
Molecule_decayThreshold	760.1	.	760.1	.	0.01	0.08601(–)	46
TCell_apoptosisNaiveMean	30	.	72
TCell_proliferationStdDev	9.6	.	72
DendriticCell_type1RequiredForActivation	2	.	72
TCell_becomeEffectorStdDev	56	.	72
CNSMacrophage_phagocytosisProbabilityMature	0.3	.	72
CNSMacrophage_phagocytosisProbabilityImmature	0.7	.	72
CD4THelper_diff00	0.05	.	72
Th1Polarization_mhcUnExpressionDelayMean	8	.	72
CD8Treg_cd8TregToCD4THelperSpecificityDropOff	1	.	72
Simulation_immunizationLinearInitial	14	.	72
APC_immatureDurationStdDev	24	.	72
APC_timeOfDeathStdDev	48	.	72
CD4Treg_type1SecretedPerHourWhenActivated	10	.	72
CD8Treg_type1SecretedPerHourWhenActivated	10	.	72
CLN_timeToCrossOrgan	12	.	72
CNS_timeToCrossOrgan	20	.	72
Circulation_width	50	.	72
DendriticCellMigrates_lengthOfTimeMovingFollowingMigration	3.5	.	72
SLO_timeToCrossOrgan	12	.	72
Spleen_timeToCrossOrgan	5	.	72
TCell_AICDStdDev	56	.	72
TCell_apoptosisNaiveStdDev	17	.	72
TCell_apoptosisPartialMaturityMean	12	.	72
TCell_apoptosisPartialMaturityStdDev	6	.	72
Th1Polarization_mhcUnExpressionDelayStdDev	2	.	72
Th2Polarization_proliferationStdDev	19.2	.	72

Table C.19: Robustness indexes for parameters with respect to the *CD8Treg Max Time* response. RI, robustness index; LI, lower index; UI, upper index; LB, lower boundary; DV, default value; UB, upper boundary. Results indicating no significant deviation in behaviour are marked with a period.

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>TCell_proliferationMean</i>	5.911	5.911	6.219	18.07(−)	19.2	20.39(+)	1
<i>CNSMacrophage_tnfaSecretedPerHourWhenStimulated</i>	9.218	9.218	20.23	90.78(−)	100	120.2(+)	2
<i>CNS_height</i>	12.67	.	12.67	.	50	56.34(−)	3
<i>CNSCell_apoptosisTNFaThreshold</i>	13.03	13.03	13.8	4.348(+)	5	5.69(−)	4
<i>Molecule_molecularHalflife</i>	13.93	13.93	28.95	0.4304(−)	0.5	0.6448(+)	5
<i>TCell_becomeEffectorMean</i>	14.94	14.94	.	51.04(+)	60	.	6
<i>CNS_width</i>	17.1	17.1	20.9	41.45(+)	50	60.45(−)	7
<i>Simulation_numCNSMacrophage</i>	19	19	27.91	60.75(−)	75	95.93(+)	8
<i>Simulation_immunizationLinearDC0</i>	21.38	21.38	.	1.572(+)	2	.	9
<i>Spleen_width</i>	33.11	33.11	70.29	33.45(−)	50	85.14(+)	10
<i>Simulation_immunizationLinearFreq</i>	37.2	37.2	.	3.768(−)	6	.	11
<i>Simulation_immunizationLinearGradient</i>	39.82	57.97	39.82	-0.007899(+)	-0.005	-0.003009(+)	12
<i>Simulation_numCD4Th</i>	41.97	41.97	73.05	23.21(+)	40	69.22(−)	13
<i>TCell_cellsPerGridspace</i>	42.67	42.67	.	4.013(−)	7	.	14
<i>SLO_width</i>	44.75	44.75	68.55	27.62(−)	50	84.28(+)	15
<i>Circulation_height</i>	45.58	50.81	45.58	24.59(−)	50	72.79(+)	16
<i>Simulation_numCNS</i>	46.67	46.67	.	266.7(−)	500	.	17
<i>APC_timeOfDeathMean</i>	47.02	47.02	.	58.28(−)	110	.	18
<i>Simulation_numCD4Treg</i>	48.05	48.05	69.65	15.59(+)	30	50.9(−)	19
<i>TCell_timeLocalActivationInducedEffectorFunctionFor</i>	53.73	53.73	.	22.21(−)	48	.	20
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	56.07	56.07	.	43.93(−)	100	.	21
<i>Simulation_numDCCNS</i>	57.64	57.64	99.66	16.94(−)	40	79.87(+)	22
<i>Spleen_height</i>	57.75	57.75	131.7	21.13(−)	50	115.9(+)	23
<i>TCell_timeLocalActivationDelay</i>	59.87	59.87	398.1	4.013(−)	10	49.81(−)	24
<i>TCell_specificityLowerLimit</i>	61.09	61.09	.	0.1946(+)	0.5	.	25
<i>TCell_apoptosisNaiveMean</i>	63.73	63.73	77.7	10.88(−)	30	53.31(+)	26
<i>Simulation_numCD8Treg</i>	73.6	73.6	118.9	7.921(+)	30	65.67(−)	27
<i>APC_probabilityPhagocytosisToPeptide</i>	74.93	74.93	.	0.005014(−)	0.02	.	28

Continued on Next Page...

Table C.19 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>Circulation_timeToCrossOrgan</i>	76.69	.	76.69	.	5	8.834(+)	29
<i>DendriticCell_cytokineType2PolarizationRatio</i>	82.23	82.23	.	0.03021(–)	0.17	.	30
<i>TCell_proliferationStdDev</i>	90.68	.	90.68	.	9.6	18.31(–)	31
<i>DendriticCell_type1SecretedPerHourImmunized</i>	91.5	91.5	.	0.8497(–)	10	.	32
<i>CD4THelper_diff08</i>	92.7	92.7	.	0.06202(–)	0.85	.	33
<i>DendriticCell_phagocytosisProbabilityImmature</i>	92.95	92.95	.	0.07054(–)	1	.	34
<i>DendriticCell_type1RequiredForActivation</i>	97.33	97.33	579.6	0.05345(–)	2	13.59(+)	35
<i>CNSMacrophage_type1RequiredForActivation</i>	106.3	.	106.3	.	2.5	5.158(–)	36
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	545.8	.	545.8	.	100	645.8(–)	37
<i>Molecule_decayThreshold</i>	3721	.	3721	.	0.01	0.3821(–)	38
<i>TCell_AICDMean</i>	60	.	72
<i>Th2Polarization_proliferationMean</i>	28.8	.	72
<i>TCell_specificityUpperLimit</i>	0.9	.	72
<i>APC_immatureDurationMean</i>	48	.	72
<i>Simulation_numDCSpleen</i>	100	.	72
<i>CLN_width</i>	50	.	72
<i>CNSMacrophage_basalMBPExpressionProbability</i>	0.2	.	72
<i>DendriticCell_phagocytosisProbabilityMature</i>	0.3	.	72
<i>CLN_height</i>	50	.	72
<i>SLO_height</i>	50	.	72
<i>Simulation_numDC</i>	10	.	72
<i>TCell_becomeEffectorStdDev</i>	56	.	72
<i>CNSMacrophage_phagocytosisProbabilityMature</i>	0.3	.	72
<i>CNSMacrophage_phagocytosisProbabilityImmature</i>	0.7	.	72
<i>CD4THelper_diff00</i>	0.05	.	72
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	8	.	72
<i>CD8Treg_cd8TregToCD4THelperSpecificityDropOff</i>	1	.	72
<i>Simulation_immuneActivationInitial</i>	14	.	72
<i>APC_immatureDurationStdDev</i>	24	.	72
<i>APC_timeOfDeathStdDev</i>	48	.	72
<i>CD4Treg_type1SecretedPerHourWhenActivated</i>	10	.	72

Continued on Next Page...

Table C.19 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>CD8Treg_type1SecretedPerHourWhenActivated</i>	10	.	72
<i>CLN_timeToCrossOrgan</i>	12	.	72
<i>CNS_timeToCrossOrgan</i>	20	.	72
<i>Circulation_width</i>	50	.	72
<i>DendriticCellMigrates_lengthOfTimeMovingFollowingMigration</i>	3.5	.	72
<i>SLO_timeToCrossOrgan</i>	12	.	72
<i>Spleen_timeToCrossOrgan</i>	5	.	72
<i>TCell_AICDStdDev</i>	56	.	72
<i>TCell_apoptosisNaiveStdDev</i>	17	.	72
<i>TCell_apoptosisPartialMaturityMean</i>	12	.	72
<i>TCell_apoptosisPartialMaturityStdDev</i>	6	.	72
<i>Th1Polarization_mhcUnExpressionDelayStdDev</i>	2	.	72
<i>Th2Polarization_proliferationStdDev</i>	19.2	.	72

Table C.20: Robustness indexes for parameters with respect to the *CD4Th1at40d* response. RI, robustness index; LI, lower index; UI, upper index; LB, lower boundary; DV, default value; UB, upper boundary. Results indicating no significant deviation in behaviour are marked with a period.

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>CNSCell_apoptosisTNFaThreshold</i>	14.7	14.7	.	4.265(+)	5	.	1
<i>TCell_proliferationMean</i>	15.59	.	15.59	.	19.2	22.19(+)	2
<i>CNSMacrophage_tnfaSecretedPerHourWhenStimulated</i>	16.26	.	16.26	.	100	116.3(+)	3
<i>CNS_height</i>	17.45	17.45	.	41.27(+)	50	.	4
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	17.57	17.57	.	6.595(+)	8	.	5
<i>Molecule_molecularHalflife</i>	25.29	.	25.29	.	0.5	0.6265(+)	6
<i>APC_timeOfDeathMean</i>	26.88	26.88	.	80.43(+)	110	.	7
<i>TCell_becomeEffectorMean</i>	29.03	29.03	48.65	42.58(+)	60	89.19(+)	8
<i>Circulation_height</i>	29.27	.	29.27	.	50	64.63(+)	9
<i>TCell_cellsPerGridspace</i>	37.56	37.56	.	4.371(+)	7	.	10

Continued on Next Page...

Table C.20 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
Simulation_numCNSMacrophage	42.02	.	42.02	.	75	106.5(+)	11
Simulation_immuneActivationLinearDC0	43.09	43.09	.	1.138(+)	2	.	12
CLN_height	45.61	.	45.61	.	50	72.81(+)	13
APC_probabilityPhagocytosisToPeptide	49.54	49.54	.	0.01009(+)	0.02	.	14
Simulation_numCD8Treg	53.27	53.27	.	14.02(+)	30	.	15
Simulation_immuneActivationLinearGradient	54.15	.	54.15	.	-0.005	-0.002292(+)	16
APC_immatureDurationMean	57.96	57.96	.	20.18(+)	48	.	17
CD8Treg_cd8TregToCD4ThelperSpecificityDropOff	72.78	72.78	.	0.2722(+)	1	.	18
Simulation_numCD4Th	79.81	79.81	.	8.075(+)	40	.	19
Simulation_numDCSpleen	82.67	82.67	.	17.33(+)	100	.	20
Molecule_decayThreshold	87.05	87.05	.	0.001295(+)	0.01	.	21
Simulation_numCD4Treg	91.56	91.56	375.3	2.532(+)	30	142.6(+)	22
CNSMacrophage_type1RequiredForActivation	96.71	96.71	.	0.08234(+)	2.5	.	23
Th2Polarization_type2SecretedPerHourWhenActivated	97.59	97.59	.	2.41(+)	100	.	24
TCell_AICDMean	99.74	.	99.74	.	60	119.8(+)	25
Simulation_numDCCNS	112.1	.	112.1	.	40	84.84(+)	26
Circulation_width	126.5	.	126.5	.	50	113.3(+)	27
Spleen_height	140.7	.	140.7	.	50	120.3(+)	28
SLO_height	160.6	.	160.6	.	50	130.3(+)	29
Th1Polarization_type1SecretedPerHourWhenActivated	477.1	.	477.1	.	100	577.1(+)	30
DendriticCell_cytokineType2PolarizationRatio	484.5	.	484.5	.	0.17	0.9936(+)	31
CD4THelper_diff00	1060	.	1060	.	0.05	0.5798(+)	32
CNS_width	50	.	72
Spleen_width	50	.	72
Simulation_immuneActivationLinearFreq	6	.	72
SLO_width	50	.	72
Simulation_numCNS	500	.	72
TCell_timeLocalActivationInducedEffectorFunctionFor	48	.	72
TCell_timeLocalActivationDelay	10	.	72
TCell_specificityLowerLimit	0.5	.	72
TCell_apoptosisNaiveMean	30	.	72

Continued on Next Page...

Table C.20 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
Circulation_timeToCrossOrgan	5	.	72
TCell_proliferationStdDev	9.6	.	72
DendriticCell_type1SecretedPerHourImmunized	10	.	72
CD4THelper_diff08	0.85	.	72
DendriticCell_phagocytosisProbabilityImmature	1	.	72
DendriticCell_type1RequiredForActivation	2	.	72
Th2Polarization_proliferationMean	28.8	.	72
TCell_specificityUpperLimit	0.9	.	72
CLN_width	50	.	72
CNSMacrophage_basalMBPExpressionProbability	0.2	.	72
DendriticCell_phagocytosisProbabilityMature	0.3	.	72
Simulation_numDC	10	.	72
TCell_becomeEffectorStdDev	56	.	72
CNSMacrophage_phagocytosisProbabilityMature	0.3	.	72
CNSMacrophage_phagocytosisProbabilityImmature	0.7	.	72
Simulation_immunizationLinearInitial	14	.	72
APC_immatureDurationStdDev	24	.	72
APC_timeOfDeathStdDev	48	.	72
CD4Treg_type1SecretedPerHourWhenActivated	10	.	72
CD8Treg_type1SecretedPerHourWhenActivated	10	.	72
CLN_timeToCrossOrgan	12	.	72
CNS_timeToCrossOrgan	20	.	72
DendriticCellMigrates_lengthOfTimeMovingFollowingMigration	3.5	.	72
SLO_timeToCrossOrgan	12	.	72
Spleen_timeToCrossOrgan	5	.	72
TCell_AICDStdDev	56	.	72
TCell_apoptosisNaiveStdDev	17	.	72
TCell_apoptosisPartialMaturityMean	12	.	72
TCell_apoptosisPartialMaturityStdDev	6	.	72
Th1Polarization_mhcUnExpressionDelayStdDev	2	.	72
Th2Polarization_proliferationStdDev	19.2	.	72

Table C.21: Robustness indexes for parameters with respect to the *Max EAE* response. RI, robustness index; LI, lower index; UI, upper index; LB, lower boundary; DV, default value; UB, upper boundary. Results indicating no significant deviation in behaviour are marked with a period.

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>CNSCell_apoptosisTNFaThreshold</i>	1.361	1.361	1.597	4.932(+)	5	5.08(−)	1
<i>CNSMacrophage_tnfaSecretedPerHourWhenStimulated</i>	1.412	1.846	1.412	98.15(−)	100	101.4(+)	2
<i>Molecule_molecularHalflife</i>	2.022	2.736	2.022	0.4863(−)	0.5	0.5101(+)	3
<i>Simulation_numCNSMacrophage</i>	3.081	3.741	3.081	72.19(−)	75	77.31(+)	4
<i>CNS_width</i>	3.247	3.247	4.32	48.38(+)	50	52.16(−)	5
<i>CNS_height</i>	3.306	3.306	3.478	48.35(+)	50	51.74(−)	6
<i>Simulation_numCNS</i>	7.553	7.553	8.636	462.2(−)	500	543.2(+)	7
<i>TCell_AICDMean</i>	10.61	10.61	12.25	53.63(−)	60	67.35(+)	8
<i>APC_timeOfDeathMean</i>	11.87	11.87	14.68	96.95(−)	110	126.1(+)	9
<i>Simulation_immunizationLinearDC0</i>	15.9	15.99	15.9	1.68(−)	2	2.318(+)	10
<i>CD4THelper_diff08</i>	16.93	16.93	.	0.7061(−)	0.85	.	11
<i>TCell_proliferationMean</i>	19.17	.	19.17	.	19.2	22.88(−)	12
<i>Simulation_immunizationLinearFreq</i>	20.72	20.72	36.19	4.757(+)	6	8.171(−)	13
<i>TCell_timeLocalActivationInducedEffectorFunctionFor</i>	23.26	23.26	107.8	36.84(−)	48	99.73(+)	14
<i>TCell_cellsPerGridspace</i>	24.78	24.78	.	5.265(−)	7	.	15
<i>Th1Polarization_type1SecretedPerHourWhenActivated</i>	27.92	27.92	58.19	72.08(−)	100	158.2(+)	16
<i>CNSMacrophage_phagocytosisProbabilityMature</i>	28.15	28.15	32.94	0.2155(−)	0.3	0.3988(+)	17
<i>TCell_becomeEffectorMean</i>	29.29	29.29	59.28	42.43(−)	60	95.57(−)	18
<i>Simulation_immunizationLinearGradient</i>	34.83	34.83	41.15	-0.006741(−)	-0.005	-0.002943(+)	19
<i>CNSMacrophage_type1RequiredForActivation</i>	36.37	42.72	36.37	1.432(+)	2.5	3.409(−)	20
<i>TCell_AICDStdDev</i>	37.63	37.63	.	34.93(−)	56	.	21
<i>Circulation_height</i>	51.3	.	51.3	.	50	75.65(−)	22
<i>CNSMacrophage_phagocytosisProbabilityImmature</i>	52.72	52.72	.	0.331(−)	0.7	.	23
<i>Simulation_numCD4Th</i>	64	64	252	14.4(−)	40	140.8(+)	24
<i>TCell_apoptosisNaiveMean</i>	68.47	68.47	.	9.458(−)	30	.	25
<i>Circulation_timeToCrossOrgan</i>	70.98	.	70.98	.	5	8.549(−)	26

Continued on Next Page...

Table C.21 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>Simulation_numDCCNS</i>	71.18	71.18	75.32	11.53(–)	40	70.13(+)	27
<i>CNSMacrophage_basalMBPExpressionProbability</i>	83.49	83.49	244.6	0.03301(–)	0.2	0.6891(+)	28
<i>DendriticCell_type1SecretedPerHourImmunized</i>	93.3	93.3	594.6	0.6696(–)	10	69.46(+)	29
<i>DendriticCell_phagocytosisProbabilityImmature</i>	93.67	93.67	.	0.06329(–)	1	.	30
<i>TCell_timeLocalActivationDelay</i>	96.19	96.19	273.9	0.381(–)	10	37.39(–)	31
<i>SLO_width</i>	130.9	.	130.9	.	50	115.4(–)	32
<i>SLO_height</i>	141.2	.	141.2	.	50	120.6(–)	33
<i>Molecule_decayThreshold</i>	1283	.	1283	.	0.01	0.1383(–)	34
<i>Simulation_numCD4Treg</i>	30	.	72
<i>Simulation_numCD8Treg</i>	30	.	72
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	8	.	72
<i>CLN_height</i>	50	.	72
<i>APC_probabilityPhagocytosisToPeptide</i>	0.02	.	72
<i>APC_immatureDurationMean</i>	48	.	72
<i>CD8Treg_cd8TregToCD4THelperSpecificityDropOff</i>	1	.	72
<i>Simulation_numDCSpleen</i>	100	.	72
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	100	.	72
<i>Circulation_width</i>	50	.	72
<i>Spleen_height</i>	50	.	72
<i>DendriticCell_cytokineType2PolarizationRatio</i>	0.17	.	72
<i>CD4THelper_diff00</i>	0.05	.	72
<i>Spleen_width</i>	50	.	72
<i>TCell_specificityLowerLimit</i>	0.5	.	72
<i>TCell_proliferationStdDev</i>	9.6	.	72
<i>DendriticCell_type1RequiredForActivation</i>	2	.	72
<i>Th2Polarization_proliferationMean</i>	28.8	.	72
<i>TCell_specificityUpperLimit</i>	0.9	.	72
<i>CLN_width</i>	50	.	72
<i>DendriticCell_phagocytosisProbabilityMature</i>	0.3	.	72
<i>Simulation_numDC</i>	10	.	72
<i>TCell_becomeEffectorStdDev</i>	56	.	72

Continued on Next Page...

Table C.21 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
Simulation_immuneActivationLinearInitial	14	.	72
APC_immuneDurationStdDev	24	.	72
APC_timeOfDeathStdDev	48	.	72
CD4Treg_type1SecretedPerHourWhenActivated	10	.	72
CD8Treg_type1SecretedPerHourWhenActivated	10	.	72
CLN_timeToCrossOrgan	12	.	72
CNS_timeToCrossOrgan	20	.	72
DendriticCellMigrates_lengthOfTimeMovingFollowingMigration	3.5	.	72
SLO_timeToCrossOrgan	12	.	72
Spleen_timeToCrossOrgan	5	.	72
TCell_apoptosisNaiveStdDev	17	.	72
TCell_apoptosisPartialMaturityMean	12	.	72
TCell_apoptosisPartialMaturityStdDev	6	.	72
Th1Polarization_mhcUnExpressionDelayStdDev	2	.	72
Th2Polarization_proliferationStdDev	19.2	.	72

Table C.22: Robustness indexes for parameters with respect to the *EAE at 40d* response. RI, robustness index; LI, lower index; UI, upper index; LB, lower boundary; DV, default value; UB, upper boundary. Results indicating no significant deviation in behaviour are marked with a period.

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
CNSMacrophage_tnfaSecretedPerHourWhenStimulated	0.8865	.	0.8865	.	100	100.9(+)	1
CNSCell_apoptosisTNFaThreshold	0.9662	0.9662	.	4.952(+)	5	.	2
Molecule_molecularHalflife	1.227	.	1.227	.	0.5	0.5061(+)	3
Simulation_numCNSMacrophage	1.714	.	1.714	.	75	76.29(+)	4
CNS_height	1.822	1.822	.	49.09(+)	50	.	5
CNS_width	1.874	1.874	.	49.06(+)	50	.	6
Simulation_numCNS	5.405	.	5.405	.	500	527(+)	7
TCell_AICDMean	10.26	.	10.26	.	60	66.15(+)	8

Continued on Next Page...

Table C.22 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
APC_timeOfDeathMean	11.76	.	11.76	.	110	122.9(+)	9
Simulation_immuneActivationLinearDC0	14.13	.	14.13	.	2	2.283(+)	10
Simulation_immuneActivationLinearFreq	16.67	16.67	.	5(+)	6	.	11
CNSMacrophage_phagocytosisProbabilityMature	25.25	.	25.25	.	0.3	0.3758(+)	12
TCell_timeLocalActivationInducedEffectorFunctionFor	31.35	.	31.35	.	48	63.05(+)	13
CNSMacrophage_type1RequiredForActivation	35.39	35.39	.	1.615(+)	2.5	.	14
Simulation_immuneActivationLinearGradient	37.75	.	37.75	.	-0.005	-0.003113(+)	15
Th1Polarization_type1SecretedPerHourWhenActivated	48.21	.	48.21	.	100	148.2(+)	16
Simulation_numDCCNS	55.11	.	55.11	.	40	62.05(+)	17
Simulation_numCD8Treg	97.46	97.46	.	0.7634(+)	30	.	18
Th1Polarization_mhcUnExpressionDelayMean	97.7	97.7	.	0.1842(+)	8	.	19
Simulation_numCD4Treg	97.94	97.94	.	0.6187(+)	30	.	20
CD8Treg_cd8TregToCD4HelperSpecificityDropOff	99.82	99.82	.	0.001759(+)	1	.	21
Simulation_numCD4Th	243.8	.	243.8	.	40	137.5(+)	22
CNSMacrophage_basalMBPExpressionProbability	285.7	.	285.7	.	0.2	0.7714(+)	23
DendriticCell_type1SecretedPerHourImmunized	604.8	.	604.8	.	10	70.48(+)	24
CD4THelper_diff08	0.85	.	72
TCell_proliferationMean	19.2	.	72
TCell_cellsPerGridspace	7	.	72
TCell_becomeEffectorMean	60	.	72
TCell_AICDStdDev	56	.	72
Circulation_height	50	.	72
CNSMacrophage_phagocytosisProbabilityImmature	0.7	.	72
TCell_apoptosisNaiveMean	30	.	72
Circulation_timeToCrossOrgan	5	.	72
DendriticCell_phagocytosisProbabilityImmature	1	.	72
TCell_timeLocalActivationDelay	10	.	72
SLO_width	50	.	72
SLO_height	50	.	72
Molecule_decayThreshold	0.01	.	72
CLN_height	50	.	72

Continued on Next Page...

Table C.22 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
APC_probabilityPhagocytosisToPeptide	0.02	.	72
APC_immatureDurationMean	48	.	72
Simulation_numDCSpleen	100	.	72
Th2Polarization_type2SecretedPerHourWhenActivated	100	.	72
Circulation_width	50	.	72
Spleen_height	50	.	72
DendriticCell_cytokineType2PolarizationRatio	0.17	.	72
CD4THelper_diff00	0.05	.	72
Spleen_width	50	.	72
TCell_specificityLowerLimit	0.5	.	72
TCell_proliferationStdDev	9.6	.	72
DendriticCell_type1RequiredForActivation	2	.	72
Th2Polarization_proliferationMean	28.8	.	72
TCell_specificityUpperLimit	0.9	.	72
CLN_width	50	.	72
DendriticCell_phagocytosisProbabilityMature	0.3	.	72
Simulation_numDC	10	.	72
TCell_becomeEffectorStdDev	56	.	72
Simulation_immuneActivationLinearInitial	14	.	72
APC_immatureDurationStdDev	24	.	72
APC_timeOfDeathStdDev	48	.	72
CD4Treg_type1SecretedPerHourWhenActivated	10	.	72
CD8Treg_type1SecretedPerHourWhenActivated	10	.	72
CLN_timeToCrossOrgan	12	.	72
CNS_timeToCrossOrgan	20	.	72
DendriticCellMigrates_lengthOfTimeMovingFollowingMigration	3.5	.	72
SLO_timeToCrossOrgan	12	.	72
Spleen_timeToCrossOrgan	5	.	72
TCell_apoptosisNaiveStdDev	17	.	72
TCell_apoptosisPartialMaturityMean	12	.	72
TCell_apoptosisPartialMaturityStdDev	6	.	72

Continued on Next Page...

Table C.22 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
Th1Polarization_mhcUnExpressionDelayStdDev	2	.	72
Th2Polarization_proliferationStdDev	19.2	.	72

Table C.23: Robustness indexes for parameters with respect to the *Max EAE A Test* response. RI, robustness index; LI, lower index; UI, upper index; LB, lower boundary; DV, default value; UB, upper boundary. Results indicating no significant deviation in behaviour are marked with a period.

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
CNSCell_apoptosisTNFaThreshold	1.452	1.452	1.472	4.927(+)	5	5.074(–)	1
CNSMacrophage_tnfaSecretedPerHourWhenStimulated	1.549	1.604	1.549	98.4(–)	100	101.5(+)	2
Molecule_molecularHalflife	2.22	2.352	2.22	0.4882(–)	0.5	0.5111(+)	3
Simulation_numCNSMacrophage	3.411	3.456	3.411	72.41(–)	75	77.56(+)	4
CNS_height	3.664	3.827	3.664	48.09(+)	50	51.83(–)	5
CNS_width	3.896	3.896	4.093	48.05(+)	50	52.05(–)	6
Simulation_numCNS	6.915	6.915	8.834	465.4(–)	500	544.2(+)	7
TCell_AICDMean	9.797	9.797	13.17	54.12(–)	60	67.9(+)	8
APC_timeOfDeathMean	11.79	11.79	16.35	97.03(–)	110	128(+)	9
Simulation_immunizationLinearDC0	13.94	13.94	17.24	1.721(–)	2	2.345(+)	10
CD4THelper_diff08	16.58	16.58	.	0.7091(–)	0.85	.	11
TCell_proliferationMean	18.24	.	18.24	.	19.2	22.7(–)	12
TCell_cellsPerGridspace	22.03	22.03	.	5.458(–)	7	.	13
TCell_timeLocalActivationInducedEffectorFunctionFor	22.52	22.52	98.96	37.19(–)	48	95.5(+)	14
Simulation_immunizationLinearFreq	23.16	23.16	35.75	4.611(+)	6	8.145(–)	15
CNSMacrophage_phagocytosisProbabilityMature	24.24	24.24	34.52	0.2273(–)	0.3	0.4036(+)	16
Th1Polarization_type1SecretedPerHourWhenActivated	26.07	26.07	61.19	73.93(–)	100	161.2(+)	17
TCell_becomeEffectorMean	28.88	28.88	57.06	42.67(–)	60	94.24(–)	18
TCell_AICDStdDev	31.36	31.36	.	38.44(–)	56	.	19
Simulation_immunizationLinearGradient	33.83	33.83	42.12	-0.006691(–)	-0.005	-0.002894(+)	20

Continued on Next Page...

Table C.23 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
CNSMacrophage_type1RequiredForActivation	34.7	42.87	34.7	1.428(+)	2.5	3.367(–)	21
Circulation_height	46.44	.	46.44	.	50	73.22(–)	22
CNSMacrophage_phagocytosisProbabilityImmature	46.66	46.66	.	0.3734(–)	0.7	.	23
Simulation_numCD4Th	62.24	62.24	267.2	15.1(–)	40	146.9(+)	24
Simulation_numDCCNS	65.01	65.01	80.51	14(–)	40	72.21(+)	25
TCell_apoptosisNaiveMean	65.05	65.05	.	10.48(–)	30	.	26
Circulation_timeToCrossOrgan	69.59	.	69.59	.	5	8.48(–)	27
CNSMacrophage_basalMBPExpressionProbability	83.73	83.73	264.4	0.03254(–)	0.2	0.7289(+)	28
DendriticCell_phagocytosisProbabilityImmature	92.48	92.48	.	0.07523(–)	1	.	29
DendriticCell_type1SecretedPerHourImmunized	93.41	93.41	620.2	0.6591(–)	10	72.02(+)	30
TCell_timeLocalActivationDelay	93.46	93.46	265.5	0.6538(–)	10	36.55(–)	31
SLO_width	119.5	.	119.5	.	50	109.7(–)	32
SLO_height	140	.	140	.	50	120(–)	33
Molecule_decayThreshold	1317	.	1317	.	0.01	0.1417(–)	34
Th1Polarization_mhcUnExpressionDelayMean	8	.	72
CLN_height	50	.	72
APC_probabilityPhagocytosisToPeptide	0.02	.	72
Simulation_numCD8Treg	30	.	72
APC_immatureDurationMean	48	.	72
CD8Treg_cd8TregToCD4ThelperSpecificityDropOff	1	.	72
Simulation_numDCSpleen	100	.	72
Simulation_numCD4Treg	30	.	72
Th2Polarization_type2SecretedPerHourWhenActivated	100	.	72
Circulation_width	50	.	72
Spleen_height	50	.	72
DendriticCell_cytokineType2PolarizationRatio	0.17	.	72
CD4THelper_diff00	0.05	.	72
Spleen_width	50	.	72
TCell_specificityLowerLimit	0.5	.	72
TCell_proliferationStdDev	9.6	.	72
DendriticCell_type1RequiredForActivation	2	.	72

Continued on Next Page...

Table C.23 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
Th2Polarization_proliferationMean	28.8	.	72
TCell_specificityUpperLimit	0.9	.	72
CLN_width	50	.	72
DendriticCell_phagocytosisProbabilityMature	0.3	.	72
Simulation_numDC	10	.	72
TCell_becomeEffectorStdDev	56	.	72
Simulation_immuneActivationLinearInitial	14	.	72
APC_immatureDurationStdDev	24	.	72
APC_timeOfDeathStdDev	48	.	72
CD4Treg_type1SecretedPerHourWhenActivated	10	.	72
CD8Treg_type1SecretedPerHourWhenActivated	10	.	72
CLN_timeToCrossOrgan	12	.	72
CNS_timeToCrossOrgan	20	.	72
DendriticCellMigrates_lengthOfTimeMovingFollowingMigration	3.5	.	72
SLO_timeToCrossOrgan	12	.	72
Spleen_timeToCrossOrgan	5	.	72
TCell_apoptosisNaiveStdDev	17	.	72
TCell_apoptosisPartialMaturityMean	12	.	72
TCell_apoptosisPartialMaturityStdDev	6	.	72
Th1Polarization_mhcUnExpressionDelayStdDev	2	.	72
Th2Polarization_proliferationStdDev	19.2	.	72

Table C.24: Robustness indexes for parameters with respect to the *EAE at 40d A Test* response. RI, robustness index; LI, lower index; UI, upper index; LB, lower boundary; DV, default value; UB, upper boundary. Results indicating no significant deviation in behaviour are marked with a period.

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
CNSMacrophage_tnfaSecretedPerHourWhenStimulated	1.862	.	1.862	.	100	101.9(+)	1
CNSCell_apoptosisTNFaThreshold	2.041	2.041	.	4.898(+)	5	.	2

Continued on Next Page...

Table C.24 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
Molecule_molecularHalflife	2.577	.	2.577	.	0.5	0.5129(+)	3
Simulation_numCNSMacrophage	3.599	.	3.599	.	75	77.7(+)	4
CNS_height	3.827	3.827	.	48.09(+)	50	.	5
CNS_width	3.934	3.934	.	48.03(+)	50	.	6
Simulation_numCNS	11.51	.	11.51	.	500	557.5(+)	7
TCell_AICDMean	17.71	.	17.71	.	60	70.62(+)	8
APC_timeOfDeathMean	24.68	.	24.68	.	110	137.1(+)	9
Simulation_immunizationLinearDC0	24.95	.	24.95	.	2	2.499(+)	10
Simulation_immunizationLinearFreq	33.2	33.2	.	4.008(+)	6	.	11
Simulation_immunizationLinearGradient	50.24	.	50.24	.	-0.005	-0.002488(+)	12
CNSMacrophage_phagocytosisProbabilityMature	56.55	.	56.55	.	0.3	0.4696(+)	13
CNSMacrophage_type1RequiredForActivation	58.82	58.82	.	1.029(+)	2.5	.	14
Simulation_numCD4Treg	98.98	98.98	.	0.3068(+)	30	.	15
Simulation_numCD8Treg	99.01	99.01	.	0.2977(+)	30	.	16
Simulation_numDCCNS	138	.	138	.	40	95.22(+)	17
Th1Polarization_type1SecretedPerHourWhenActivated	141.3	.	141.3	.	100	241.3(+)	18
DendriticCell_type1SecretedPerHourImmunized	889.5	.	889.5	.	10	98.95(+)	19
CD4THelper_diff08	0.85	.	72
TCell_proliferationMean	19.2	.	72
TCell_cellsPerGridspace	7	.	72
TCell_timeLocalActivationInducedEffectorFunctionFor	48	.	72
TCell_becomeEffectorMean	60	.	72
TCell_AICDStdDev	56	.	72
Circulation_height	50	.	72
CNSMacrophage_phagocytosisProbabilityImmature	0.7	.	72
Simulation_numCD4Th	40	.	72
TCell_apoptosisNaiveMean	30	.	72
Circulation_timeToCrossOrgan	5	.	72
CNSMacrophage_basalMBPExpressionProbability	0.2	.	72
DendriticCell_phagocytosisProbabilityImmature	1	.	72
TCell_timeLocalActivationDelay	10	.	72

Continued on Next Page...

Table C.24 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>SLO_width</i>	50	.	72
<i>SLO_height</i>	50	.	72
<i>Molecule_decayThreshold</i>	0.01	.	72
<i>Th1Polarization_mhcUnExpressionDelayMean</i>	8	.	72
<i>CLN_height</i>	50	.	72
<i>APC_probabilityPhagocytosisToPeptide</i>	0.02	.	72
<i>APC_immatureDurationMean</i>	48	.	72
<i>CD8Treg_cd8TregToCD4ThelperSpecificityDropOff</i>	1	.	72
<i>Simulation_numDCSpleen</i>	100	.	72
<i>Th2Polarization_type2SecretedPerHourWhenActivated</i>	100	.	72
<i>Circulation_width</i>	50	.	72
<i>Spleen_height</i>	50	.	72
<i>DendriticCell_cytokineType2PolarizationRatio</i>	0.17	.	72
<i>CD4THelper_diff00</i>	0.05	.	72
<i>Spleen_width</i>	50	.	72
<i>TCell_specificityLowerLimit</i>	0.5	.	72
<i>TCell_proliferationStdDev</i>	9.6	.	72
<i>DendriticCell_type1RequiredForActivation</i>	2	.	72
<i>Th2Polarization_proliferationMean</i>	28.8	.	72
<i>TCell_specificityUpperLimit</i>	0.9	.	72
<i>CLN_width</i>	50	.	72
<i>DendriticCell_phagocytosisProbabilityMature</i>	0.3	.	72
<i>Simulation_numDC</i>	10	.	72
<i>TCell_becomeEffectorStdDev</i>	56	.	72
<i>Simulation_immunizationLinearInitial</i>	14	.	72
<i>APC_immatureDurationStdDev</i>	24	.	72
<i>APC_timeOfDeathStdDev</i>	48	.	72
<i>CD4Treg_type1SecretedPerHourWhenActivated</i>	10	.	72
<i>CD8Treg_type1SecretedPerHourWhenActivated</i>	10	.	72
<i>CLN_timeToCrossOrgan</i>	12	.	72

Continued on Next Page...

Table C.24 – Continued

Parameter Name	RI (%)	LI (%)	UI (%)	LB	DV	UB	Rank
<i>CNS_timeToCrossOrgan</i>	20	.	72
<i>DendriticCellMigrates_lengthOfTimeMovingFollowingMigration</i>	3.5	.	72
<i>SLO_timeToCrossOrgan</i>	12	.	72
<i>Spleen_timeToCrossOrgan</i>	5	.	72
<i>TCell_apoptosisNaiveStdDev</i>	17	.	72
<i>TCell_apoptosisPartialMaturityMean</i>	12	.	72
<i>TCell_apoptosisPartialMaturityStdDev</i>	6	.	72
<i>Th1Polarization_mhcUnExpressionDelayStdDev</i>	2	.	72
<i>Th2Polarization_proliferationStdDev</i>	19.2	.	72

Glossary of Terms

Adaptive immune system. The *specific* arm of the immune system, capable of recognising structures and components of individual pathogens and responding accordingly. The adaptive immune system is responsible for long lasting immunity against a pathogen.

Adjuvant. A substance that stimulates the immune response to antigen, but which does not itself confer immunity.

Anergy. A state of un-responsiveness in lymphocytes. This is triggered if a lymphocyte receives signal 1 in absence of subsequent signal 2.

Antibody. A soluble form of the B cell receptor, secreted by effector B cells.

Antigen. Any substance that is recognised through the TCR or T cells or the BCR (B cell receptor) receptors.

APC. Antigen Presenting Cell. Any cell that is able to present MHC-peptide complexes to T cells. Note that *professional* APCs present MHC-II, whereas most cells of the body do not.

Apoptosis. A controlled cellular death.

Autoimmunity. The phenomenon whereby the immune system targets elements of the host, manifesting as a disease in severe cases.

B cell. A cell of the adaptive immune system, originating from the bone marrow. Effector B cells secrete a soluble form of their receptor called the antibody.

Circulatory system. An organ system that transports cells and substances between the organs and tissues of the body. ARTIMMUS contains a circulatory system compartment, an abstraction for the circulatory system in the body.

CD4T cell. A T cell expressing the CD4 receptor. These are typically T helper cells.

CD8T cell. A T cell expressing the CD8 receptor. These are typically cytotoxic T cells.

CD4Th1. A polarisation of CD4THelper cell, which expresses pro-inflammatory type 1 cytokines. These cells typically promote the cytotoxic immune response. In ARTIMMUS the CD4Th1 cell is an abstraction of EAE causing MBP-specific type 1 effector T helper cells.

CD4Th2. A polarisation of CD4THelper cell, which expresses type 2 cytokines. These cells typically promote the humoral immune response. In ARTIMMUS the CD4Th2 cell is an abstraction of MBP-specific type 2 effector T cells, which do not cause EAE disease.

CD4Th cell. A type of T cell, expressing CD4 receptors, and hence a T helper cell (see entry on T helper cell).

CD4THelper. In ARTIMMUS the term CD4THelper cell can refer to any MBP-specific T helper cell, regardless of state of activation or polarisation.

CD4Treg. A CD4 receptor expressing regulatory T cell. These cells are T helper cells. In ARTIMMUS they are specific for MHC-II:Fr3 complexes, and responsible for licensing dendritic cells for the expression of Qa-1 complexes.

CD8Treg. A CD8 receptor expressing regulatory T cell. These cells are cytotoxic T cells. In ARTIMMUS, they are specific for Qa-1:CDR1/2 complexes.

CDR1/2. Complementarity Determining Region 1/2. The complementarity determining region is part of the TCR receptor that interacts with antigen, and determines specificity. In ARTIMMUS, CDR1/2 refers to a peptide derived from the TCRs of CD4THelper cells, which binds with Qa-1 to form Qa-1:CDR1/2, for which CD8Treg cells are specific.

Chemokine. A subset of cytokines that influence the migration of cells around the body.

CLN. Cervical Lymph Node. These lymph nodes are found in the neck, and drain the tissues of the brain. In ARTIMMUS, the CLN compartment drains the CNS compartment.

CNS. Central Nervous System. Part of the nervous system, containing a large quantity of neurons. ARTIMMUS contains a CNS compartment as an abstraction of the brain and spinal cord.

CNS cell. An abstract representation in ARTIMMUS of neurons and associated cells in the CNS compartment. CNS cells are induced into apoptosis by sufficiently high concentrations of TNF- α .

CNS macrophage. An abstraction used in ARTIMMUS to represent microglia and macrophages found in the CNS compartment. These cells can locally activate infiltrating CD4Th cells, and secrete TNF- α upon stimulation.

Complete Freund's adjuvant. A type of adjuvant. Often shortened to CFA.

Co-stimulatory molecule. A receptor expressed by professional APCs that provide an essential signal in the activation of T cells. This signalling pathway is often referred to as *signal 2*.

Cytokine. Cell signalling protein molecules, secreted by cells.

Cytotoxic immune response. A type of adaptive immune response, mediated through cytotoxic T cells.

Cytotoxic T cell. A type of T cell, expressing the CD8 receptor. Cytotoxic T cells are specific for MHC-I:peptide² complexes. Effector cytotoxic T cells can induce apoptosis in cells expressing MHC:peptide complexes for which they are specific, though this is not considered to include the dendritic cell in this thesis.

Disease. An abnormal condition affecting the body, to the detriment of its function.

Dendritic cell. A cell of the innate immune system. Dendritic cells are highly phagocytic, are able to express MHC-I, MHC-II and Qa-1 complexes, and are the main primers of naive T cells. They often exist in the peripheral tissues as immature cells, migrating to the lymph nodes upon maturation. They influence whether naive T helper cells are polarised in a type 1 or type 2 direction upon priming.

EAE. Experimental Autoimmune Encephalomyelitis. An autoimmune disease whereby the immune system targets the tissues of the central nervous system. It is used as an animal model for multiple sclerosis.

Effector T cell. A T cell existing in a state able to carry out effector functions such as cytotoxic activity or cytokine secretion. T cells reach the effector stage of their life cycles in the lymph nodes or spleen in which they were primed, leave these compartments and migrate through the tissues, typically towards to sites of inflammation.

Encephalitogenic. Tending to cause inflammation of the brain (encephalitis).

²This includes Qa-1:peptide.

Fr3. Framework Region 3. This is a part of the TCR receptor, not responsible for determining specificity. In ARTIMMUS, Fr3 refers to a peptide derived from the TCRs of CD4THelper cells. It binds with MHC-II to form MHC-II:Fr3, for which CD4Treg cells are specific.

Humoral immune response. A type of adaptive immune response, mediated through antibodies secreted by B cells.

Immune System. A collective term given to the cells, molecules, processes and organs that maintain the health of an organism.

Immunity. The biological state whereby the immune system is able to protect the host from a particular infection or disease.

Immunize. To cause an immune response to a particular antigen.

Inflammation. The promotion of immune cell migration towards the sites of bodily damage.

Innate immune system. The *non-specific* arm of the immune system, responding to pathogens in a generic manner. Unlike the adaptive immune system, the innate immune system does not confer long-lasting immunity against a pathogen. It plays a critical role in the activation of T cells.

In silico. A latin term referring to experimentation carried out using computers, typically as simulations.

In vivo. A latin term referring to experimentation carried out using a whole, living organism. In this thesis the term also encompasses *in vitro* studies.

In vitro. A latin term referring to experimentation carried out in a test tube or petri dish.

Licensing. The phenomenon whereby a dendritic cell must first interact with an effector CD4Th cell before it is able to prime CD8T cells.

Lymph node. An immune system organ, small and spherical in shape. Lymph nodes are situated throughout the body, and drain the interstitial fluid (lymph) found in the body's tissues. They specialise in providing a compartment where cells of the innate and adaptive immune system interact, and where the adaptive immune response originates.

Lymphocyte. A cell of the adaptive immune system.

Macrophage. A cell of the innate immune system. Macrophages are capable of phagocytosis, express MHC-I and MHC-II complexes, and secrete a range of cytokines.

MBP. Myelin Basic Protein. A derivative of myelin. In ARTIMMUS MBP represents a peptide that is derived by APCs through the phagocytosis of CNS cells. It binds with MHC-II to form MHC-II:MBP complexes for which CD4THelper cells are specific.

MHC. Major Histocompatibility Complex. A receptor that binds with TCRs, subject to TCR specificity. It is expressed by APCs. MHC can be loaded with peptides, hence presenting peptides to T cells. MHC comes in two major forms, MHC-I to which CD8 T cells bind, and MHC-II to which CD4 T cells bind.

MHC-II:Fr3. In ARTIMMUS, this complex comprises an Fr3 peptide presented by an MHC-II molecule. CD4Treg cells are specific for this complex.

MHC-II:MBP. In ARTIMMUS, this complex comprises an MBP peptide presented by an MHC-II molecule. CD4THelper cells are specific for this complex.

Microglia. Specialised macrophages that reside in the CNS compartment.

Model. An abstract depiction of a target system, theory or phenomenon, either mathematical or diagrammatic in nature.

Multiple sclerosis. An autoimmune disease in which the immune system targets the central nervous system.

Murine. Pertaining to rats and mice.

Myelin. A protein comprising the insulator sheath that surrounds the neurons, necessary for their function.

Naive T cell. A T cell that has yet to bind with antigen for which it is specific.

Neuron. Cells of the nervous system, responsible for transmitting electrical signals around the body.

Pathogen. A disease causing agent. For example, a virus, parasite or bacterium.

Peptide. A string of amino-acids. Their relevance to the present immunology is that peptides are derived by cells from endogenous and, in the case of phagocytes, exogenous material and presented on MHC complexes.

Pertussis toxin. An adjuvant, also referred to as *PTx*.

Phagocyte. A cell capable of phagocytosis.

Phagocytosis. A cellular process of engulfing foreign objects, for example inter-cellular debris or bacteria, through the cell membrane.

PRCC. Partial Rank Correlation Coefficient. A measure of correlation between two random variables, with the effects of a set of other controlling variables removed. The measure employs ranks, rather than raw data values, and as such is better able to discern associations where non-linear relationships between variables exist.

Proliferation. Cell division in adaptive immune system cells. A lymphocyte that has been activated through receipt of signals 1 and 2 enters the proliferative phase of its lifecycle and produces naive daughter lymphocytes.

Qa-1. A form of non-classical MHC-I molecule that presents a substantially smaller repertoire of peptides than classical MHC-I can. In ARTIMMUS, Qa-1 presents CDR1/2 peptides as Qa-1:CDR1/2 complexes for which CD8Treg cells are specific.

Qa-1:CDR1/2. In ARTIMMUS, this complex comprises a CDR1/2 peptide presented by a Qa-1 molecule. CD8Treg cells are specific for these complexes.

Receptor. A molecule found on the surface of a cell that receives signals from receptors found on other cells, or soluble factors such as cytokines.

Regulatory T cell. A type of T cell that suppresses the activities of other T cells, either by moderating their proliferation, suppressing their effector activities, or directly inducing apoptosis in them.

Sensitivity analysis. An umbrella term given to a statistical procedure that quantifies the effects that a system's inputs has on its outputs. In the present thesis the system is a simulation, the inputs are its parameters, and its outputs are particular metrics of behaviour termed *responses*.

Signal 1. Necessary for the activation of a lymphocyte, signal 1 refers to the successful binding of the lymphocyte's receptor (be it a TCR or BCR) with antigen.

Signal 2. Necessary for the activation of a lymphocyte, signal 2 refers to receipt of co-stimulatory molecule signalling by the lymphocyte.

Simulation. An instantiation and execution of a model on a computer.

SLO. Secondary Lymphoid Organ. These organs specialise in the interaction of innate and adaptive immune system cells, examples include the spleen and the lymph nodes. In ARTIMMUS the SLO is a compartment representing a single lymph node at the site of immunization for EAE.

Specificity. Refers to the antigen(s) to which a lymphocyte is able to successfully bind.

Spleen. An immune system organ. The spleen performs a similar function to the lymph node, but drains the blood rather than the interstitial fluids of the tissues.

T cell. A cell of the adaptive immune system, the T cell originates from the thymus. It comes in two forms, the T helper cell and the cytotoxic T cell. T cells express TCR receptors, and interact with MHC:peptide complexes.

T helper cell. A form of T cell that plays a critical role in the generation of the humoral and cytotoxic adaptive immune responses. These cells express the CD4 molecule, and bind with MHC-II:peptide complexes. They come in two forms: CD4Th1 and CD4Th2 cells (see entries above).

TCR. T Cell Receptor. Every T cell expresses large quantities of these receptors. TCRs may differ between T cells, but are identical on any particular T cell. The TCR determines the specificity of a T cell, that antigen to which it is responsive. TCRs interact with MHC molecules, as expressed on APCs.

TNF- α . Tumor Necrosis Factor α . A type of cytokine, secreted by CNS macrophages, and harmful to neurons above a threshold concentration.

Type 1 cytokine. An abstraction used in ARTIMMUS that represents cytokines secreted by CD4Th1 cells and dendritic cells, favouring their expansion and suppressing the proliferation of CD4Th2 cells.

Type 2 cytokine. An abstraction used in ARTIMMUS that represents cytokines secreted by CD4Th2 cells, favouring their expansion and suppressing the proliferation of CD4Th1 cells.

Vaccine. A preparation, often comprising adjuvant and attenuated micro-organism as antigen, that aims to provide long lasting immunity to a particular disease or pathogen.

Bibliography

[An *et al.* 2009] Gary An, Qi Mi, Joyeeta Dutta-Moscato, and Yorah Vodovotz. Agent-based models in translational systems biology. *Wiley Interdisciplinary Reviews Systems Biology and Medicine*, 1(2):159–171, 2009.

[Anderson & Matzinger 2000] Colin C. Anderson and Polly Matzinger. Danger: the view from the bottom of the cliff. *Seminars in Immunology*, 12:231–238, 2000.

[Ando *et al.* 1989] Dale G. Ando, Julie Clayton, Dwight Kono, James L. Urban, and Eli Sercarz. Encephalitogenic T cells in the B10.PL model of experimental allergic encephalomyelitis (EAE) are of the Th-1 lymphokine subtype. *Cellular immunology*, 124(1):132–143, 1989.

[Andrews & Timmis 2006] Paul S. Andrews and Jon Timmis. A computational model of degeneracy in a lymph node. In [Bersini & Carneiro 2006], pages 164–177.

[Andrews *et al.* 2008a] Paul S. Andrews, Fiona Polack, Adam T. Sampson, Jon Timmis, Lisa Scott, and Mark Coles. Simulating biology: towards understanding what the simulation shows. In [Stepney *et al.* 2008], pages 93–124.

[Andrews *et al.* 2008b] Paul S. Andrews, Adam T. Sampson, John Markus Bjørndalen, Susan Stepney, Jon Timmis, Douglas N. Warren, and Peter H. Welch. Investigating patterns for the process-oriented modelling and simulation of space in complex systems. In S. Bullock, J. Noble, R. Watson, and M. A. Bedau, editors, *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, pages 17–24. MIT Press, Cambridge, MA, 2008.

[Andrews *et al.* 2010] Paul S. Andrews, Fiona A. C. Polack, Adam T. Sampson, Susan Stepney, and Jon Timmis. The CoSMoS Process Version 0.1: A Process for the Modelling and Simulation of Complex Systems. Technical Report YCS-2010-453, Department of Computer Science, the University of York, 2010.

[Bagnoli *et al.* 2006] Franco Bagnoli, Pietro Liò, and Luca Sguanci. Modeling viral coevolution: HIV multi-clonal persistence and competition dynamics. *Physica A*, 366:333–346, 2006.

[Baker & Jackson 2007] David Baker and Samuel J. Jackson. Models of multiple sclerosis. *Advances in Clinical Neuroscience & Rehabilitation*, 6(6):10–12, 2007.

[Balan *et al.* 2003] Gabriel Catalin Balan, Claudio Cioffi-Revilla, Sean Luke, Liviu Panait, and Sean Paus. MASON: A Java Multi-Agent Simulation Library. In *Proceedings of the Agent 2003 Conference*, 2003.

[Bauer *et al.* 2009] Amy L. Bauer, Catherine AA Beauchemin, and Alan S. Perelson. Agent based modeling of host-pathogen systems: the successes and challenges. *Information Sciences*, 179(10):1379–1389, 2009.

[Baxter 2007] Alan G. Baxter. The origin and application of experimental autoimmune encephalomyelitis. *Nature Reviews Immunology*, 7(11):904–912, 2007.

[Beauchemin *et al.* 2005] Catherine Beauchemin, John Samuel, and Jack Tuszyński. A simple cellular automaton model for influenza A viral infections. *Journal of Theoretical Biology*, 232(2):223–234, 2005.

[Beauchemin 2006] Catherine Beauchemin. Probing the effects of the well-mixed assumption on viral infection dynamics. *Journal of Theoretical Biology*, 242(2):464–477, 2006.

[Beeston *et al.* 2010] Tara Beeston, Trevor R. F. Smith, Igor Maricic, Xiaolei Tang, and Vipin Kumar. Involvement of INF- γ and perforin, but not Fas/FasL interactions in regulatory T cell-mediated suppression of experimental autoimmune encephalomyelitis. *Journal of Neuroimmunology*, 229(1):91–97, 2010.

[Ben-Nun *et al.* 1980] Avraham Ben-Nun, Yacov Ron, and Irun R. Cohen. Spontaneous remission of autoimmune encephalomyelitis is inhibited by splenectomy, thymectomy or ageing. *Nature*, 288(5789):389–390, 1980.

[Ben-Nun *et al.* 1981] Avraham Ben-Nun, Hartmut Wekerle, and Irun R. Cohen. Vaccination against autoimmune encephalomyelitis with T-lymphocyte line cells reactive against myelin basic protein. *Nature*, 292(5818):60–61, 1981.

[Bersini & Carneiro 2006] Hugues Bersini and Jorge Carneiro, editors. *Artificial Immune Systems: 5th International Conference, ICARIS 2006, Oeiras, Portugal, September 4-6, 2006, Proceedings*, volume 4163 of *Lecture Notes in Computer Science*. Springer, 2006.

[Bersini 2006] Hugues Bersini. Immune system modeling: The OO way. In [Bersini & Carneiro 2006], pages 150–163.

[Bevan 2006a] Michael J. Bevan. Cross-priming. *Nature Immunology*, 7(4):363–365, 2006.

[Bevan 2006b] Michael J. Bevan. Helping the cd8(+) response. *Nature Reviews Immunology*, 4(8):595–602, 2006.

[Bocharov & Romanyukha 1994] G. A. Bocharov and A. A. Romanyukha. Mathematical model of antiviral immune response. III. influenza a virus infection. *Journal of Theoretical Biology*, 167(4):323–360, 1994.

[Boehmer 2005] Harald von Boehmer. Mechanisms of suppression by suppressor T cells. *Nature Immunology*, 6(4):338–344, 2005.

[Bogle & Dunbar 2008] Gib Bogle and Rod R. Dunbar. Simulating T-cell motility in the lymph node paracortex with a packed lattice geometry. *Immunology and Cell Biology*, 86(8):676–687, 2008.

[Bogle & Dunbar 2010] Gib Bogle and Rod R. Dunbar. T cell responses in lymph nodes. *Wiley Interdisciplinary Reviews. Systems Biology and Medicine*, 2(1):107–116, 2010.

[Borghans *et al.* 1998] José A. M. Borghans, Rob J. de Boer, Eli Sercarz, and Vipin Kumar. T cell vaccination in experimental autoimmune encephalomyelitis: A mathematical model. *The Journal of Immunology*, 161(3):1087–1093, 1998.

[Bracewell 1999] Ronald N. Bracewell. *The Fourier transform and its applications*. McGraw-Hill Higher Education, 3 edition, 1999.

[Bretscher 2000] Peter Bretscher. Contemporary models for peripheral tolerance and the classical ‘historical postulate’. *Seminars in Immunology*, 12:221–229, 2000.

[Chakraborty & Das 2010] Arup K. Chakraborty and Jayajit Das. Pairing computation with experimentation: a powerful coupling for understanding t cell signalling. *Nature Reviews Immunology*, 10(1):59–71, 2010.

[Chakraborty *et al.* 2003] Arup K. Chakraborty, Michael L. Dustin, and Andrey S. Shaw. In silico models for cellular and molecular immunology: successes, promises and challenges. *Nature Immunology*, 4(10):933–936, 2003.

[Chang *et al.* 2002] C. C. Chang, R. Ciubotariu, J. S. Manavalan, J. Yuan, A. I. Colovai, F. Piazza, S. Lederman, M. Colonna, R. Cortesini, R. Dalla-Favera, and N. Suciu-Foca. Tolerization of dendritic cells by T_s cells: the crucial role of inhibitory receptors ILT3 and ILT4. *Nature Immunology*, 3(3):237–243, 2002.

[Chao *et al.* 2004a] Dennis L. Chao, Miles P. Davenport, Stephanie Forrest, and Alan S. Perelson. Modelling the impact of antigen kinetics on T-cell activation and response. *Immunology and Cell Biology*, 82(1):55–61, 2004.

[Chao *et al.* 2004b] Dennis L. Chao, Miles P. Davenport, Stephanie Forrest, and Alan S. Perelson. A stochastic model of cytotoxic T cell responses. *Journal of Theoretical Biology*, 228(2):227–240, 2004.

[Chao *et al.* 2005] Dennis L. Chao, Miles P. Davenport, Stephanie Forrest, and Alan S. Perelson. The effects of thymic selection on the range of t cell cross-reactivity. *European Journal of Immunology*, 35(12):3452–3459, 2005.

[Cohen & Harel 2007] Irun R. Cohen and David Harel. Explaining a complex living system: dynamics, multi-scaling and emergence. *Journal of the Royal Society Interface*, 4(13):175–182, 2007.

[Cohen 2000] Irun R. Cohen. Discrimination and dialogue in the immune system. *Seminars in Immunology*, 12:215–219, 2000.

[Cohen 2004] Irun R. Cohen. *Tending Adam’s Garden : Evolving the Cognitive Immune Self*. Elsevier Academic Press, August 2004.

[Cohen 2006] Irun R. Cohen. Immune system computation and the immunological homunculus. In Oscar Nierstrasz, Jon Whittle, David Harel, and Gianna Reggio, editors, *MoDELS*, volume 4199 of *Lecture Notes in Computer Science*, pages 499–512. Springer, 2006.

[Cohen 2007a] Irun R. Cohen. Modeling immune behavior for experimentalists. *Immunological Reviews*, 216(1):232–236, 2007.

[Cohen 2007b] Irun R. Cohen. Real and artificial immune systems: computing the state of the body. *Nature Reviews Immunology*, 7, 2007.

[Cortesini *et al.* 2001] Raffaello Cortesini, Joël LeMaoult, Rodica Ciubotariu, and Nicole Suciu Foca Cortesini. CD8 $^+$ CD28 $^-$ T suppressor cells and the induction of antigen-specific antigen-presenting cell-mediated suppression of Th reactivity. *Immunological Reviews*, 182:201–206, 2001.

[Dahirel *et al.*] Vincent Dahirel, Karthik Shekhar, Florencia Pereyra, Toshiyuki Miura, Mikita Artyomov, Shiv Talsania, Todd M. Allen, Marcus Altfeld, Mary Carrington, Darrell J. Irvine, Bruce D. Walker, and Arup K. Chakraborty. Coordinate linkage of HIV evolution reveals regions of immunological vulnerability. *Proceedings of the National Academy of Sciences*, published online ahead of print, accessed online at <http://dx.doi.org/10.1073/pnas.1105315108> on 22nd August 2011.

[Edelman & Gally 2001] Gerald M. Edelman and Joseph A. Gally. Degeneracy and complexity in biological systems. *Proceedings of the National Academy of Sciences*, 98(24):13763–13768, 2001.

[Efroni *et al.* 2003] Sol Efroni, David Harel, and Irun R. Cohen. Toward rigorous comprehension of biological complexity: Modeling, execution, and visualization of thymic T-cell maturation. *Genome Research*, 13:2485–2497, 2003.

[Efroni *et al.* 2005] Sol Efroni, David Harel, and Irun R. Cohen. Reactive animation: Realistic modeling of complex dynamic systems. *IEEE Computer*, 38(1):38–47, 2005.

[Efroni *et al.* 2007] Sol Efroni, David Harel, and Irun R. Cohen. Emergent dynamics of thymocyte development and lineage determination. *PLoS Computational Biology*, 3(1):0127–0136, 2007.

[Forrest & Beauchemin 2007] Stephanie Forrest and Catherine Beauchemin. Computer immunology. *Immunological Reviews*, 216(1):176–197, 2007.

[Fowler 2004] Martin Fowler. *UML Distilled*. Addison-Wesley, 3rd edition, 2004.

[Garnett *et al.* 2008] Philip Garnett, Susan Stepney, and Ottoline Leyser. Towards an executable model of auxin transport canalisation. [Stepney *et al.* 2008], pages 63–91.

[Garrett *et al.* 2007] Simon Garrett, Martin Robbins, Joanne Walker, William Wilson, and Uwe Aickelin. *In Silico Immunology*, chapter 5 : Modelling immunological memory. Springer, 2007.

[Germain *et al.* 2011] Ronald N. Germain, Martin Meier-Schellersheim, Aleksandra Nitza Lazar, and Iain D. Fraser. Systems biology in immunology: a computational modeling perspective. *Annual Review of Immunology*, 29(1):527–585, 2011.

[Germain 2001] Ronald N. Germain. The art of the probable: System control in the adaptive immune system. *Science*, 293(5528):240–245, 2001.

[Germain 2004] Ronald N. Germain. An innately interesting decade of research in immunology. *Nature Medicine*, 10(12):1307–1320, 2004.

[Gosling *et al.* 2005] James Gosling, Bill Joy, Guy Steele, and Gilad Bracha. *The Java language specification*. Addison Wesley, 3 edition, 2005.

[Goverman 2009] Joan Goverman. Autoimmune T cell responses in the central nervous system. *Nature Reviews Immunology*, 9:393–407, 2009.

[Grossman & Paul 2000] Zvi Grossman and William E. Paul. Self-tolerance: Context dependent tuning of T cell antigen recognition. *Seminars in Immunology*, 12:197–203, 2000.

[Harel 1987] David Harel. Statecharts: A visual formalism for complex systems. *Science of Computer Programming*, 8(3):231–274, June 1987.

[Helton 2008] J. C. Helton. Uncertainty and sensitivity analysis for models of complex systems. In Timothy J. Barth, Michael Griebel, David E. Keyes, Risto M. Nieminen, Dirk Roose, and Tamar Schlick, editors, *Computational Methods in Transport: Verification and Validation*, Lecture Notes in Computational Science and Engineering, pages 207–228. Springer Berlin Heidelberg, 2008.

[Hendriks *et al.* 2005] Jerome J. A. Hendriks, Charlotte E. Teunissen, Helga E. de Vries, and Christine D. Dijkstra. Macrophages and neurodegeneration. *Brain Research Reviews*, 48(2):185–195, 2005.

[Henrickson *et al.* 2008] Sarah E. Henrickson, Thorsten R. Mempel, Irina B. Mazo, Bai Liu, Maxim N. Artyomov, Huan Zheng, Antonio Peixoto, Michael P. Flynn, Balimkiz Senman, Tobias Junt, Hing C. Wong, Arup K. Chakraborty, and Ulrich H. von Andrian. T cell sensing of antigen dose governs interactive behavior with dendritic cells and sets a threshold for T cell activation. *Nature Immunology*, 9(3):282–291, 2008.

[Hood *et al.* 2004] Leroy Hood, James R. Heath, Michael E. Phelps, and Biaoyang Lin. Systems biology and new technologies enable predictive and preventative medicine. *Science*, 306(5696):640–643, 2004.

[Jacob *et al.* 2004] Christian Jacob, Julius Litorco, and Leo Lee. Immunity through swarms: Agent-based simulations of the human immune system. In Giuseppe Nicosia, Vincenzo Cutello, Peter J. Bentley, and Jon Timmis, editors, *ICARIS*, volume 3239 of *Lecture Notes in Computer Science*, pages 400–412. Springer, 2004.

[Janeway *et al.* 2005] Charles Janeway, Paul Travers, Mark Walport, and Mark Shlomchik. *Immunobiology*. Garland Science Publishing, 6 edition, 2005.

[Kabelitz *et al.* 1993] D. Kabelitz, T. Pohl, and K. Pechhold. Activation-induced cell death (apoptosis) of mature peripheral T lymphocytes. *Immunology Today*, 14(7):338–339, 1993.

[Kam *et al.* 2001] Na’aman Kam, Irun R. Cohen, and David Harel. The immune system as a reactive system: Modeling T cell activation with statecharts. In *Proceedings of Visual Languages and Formal Methods (VLFM’01), part of IEEE Symposium on Human-Centric Computing*, pages 15–22, 2001.

[Kapsenberg 2003] Martien L. Kapsenberg. Dendritic cell control of pathogen-driven T cell polarization. *Nature Reviews Immunology*, 3:984–992, 2003.

[Kaye & Scott 2011] Paul Kaye and Phillip Scott. Leishmaniasis: complexity at the host-pathogen interface. *Nature Reviews Microbiology*, 9:604–615, 2011.

[Kim *et al.* 2010] Hye-Jung J. Kim, Bert Verbinnen, Xiaolei Tang, Linrong Lu, and Harvey Cantor. Inhibition of follicular T-helper cells by CD8(+) regulatory T cells is essential for self tolerance. *Nature*, 467(7313):328–332, 2010.

[Kindt *et al.* 2007] Thomas J. Kindt, Richard A. Goldsby, and Barbara A. Osbourne. *Kuby Immunology*. W. H. Freeman and Company, 6th edition, 2007.

[Kirschner & Linderman 2009] Denise E. Kirschner and Jennifer J. Linderman. Mathematical and computational approaches can complement experimental studies of host-pathogen interactions. *Cellular Microbiology*, 11(4):531–539, 2009.

[Kirschner *et al.* 2007] Denise E. Kirschner, Stewart T. Chang, Thomas W. Riggs, Nicolas Perry, and Jennifer J. Linderman. Toward a multiscale model of antigen presentation in immunity. *Immunological Reviews*, 216:93–118, 2007.

[Kleinsteine 2008] Steven H. Kleinsteine. Getting started in computational immunology. *PLoS Computational Biology*, 4(8):e1000128+, 2008.

[Kranzler & Moursund 1999] Gerald Kranzler and Janet Moursund. *Statistics for the Terrified*. Prentice Hall, 2 edition, 1999.

[Kumar & Sercarz 2001] Vipin Kumar and Eli Sercarz. An integrative model of regulation centered on recognition of TCR peptide/MHC complexes. *Immunological Reviews*, 182:113–121, 2001.

[Kumar *et al.* 1996] Vipin Kumar, Karl Stellrecht, and Eli Sercarz. Inactivation of T cell receptor peptide-specific CD4 regulatory T cells induces chronic experimental autoimmune encephalomyelitis. *Journal of Experimental Medicine*, 184:1609–1617, 1996.

[Kumar *et al.* 2001] Vipin Kumar, Jeannie Maglione, Jayant Thatte, Brian Pederson, Eli Sercarz, , and E. Sally Ward. Induction of a type 1 regulatory CD4 T cell response following V β 8.2 DNA vaccination results in immune deviation and protection from experimental autoimmune encephalomyelitis. *International Immunology*, 13(6):835–841, 2001.

[Kumar 1998] Vipin Kumar. TCR peptide-reactive T cells and peripheral tolerance to myelin basic protein. *Research in Immunology*, 149(9):827–834, 1998.

[Kumar 2004] Vipin Kumar. Homeostatic control of immunity by TCR peptide-specific Tregs. *The Journal of Clinical Investigation*, 114(9):1222–1226, 2004.

[Langman & Cohn 2000] R. E. Langman and M. Cohn. A minimal model of the self-nonself discrimination: A return to the basics. *Seminars in Immunology*, 12:189–195, 2000.

[Lanzavecchia & Sallusto 2004] Antonio Lanzavecchia and Federica Sallusto. Lead and follow: The dance of the dendritic cell and T cell. *Nature Immunology*, 5(12):1201–1202, 2004.

[Le Novère *et al.* 2009] Nicolas Le Novère, Michael Hucka, Huaiyu Mi, Stuart Moodie, Falk Schreiber, Anatoly Sorokin, Emek Demir, Katja Wegner, Mirit I. Aladjem, Sarala M. Wimalaratne, Frank T. Bergman, Ralph Gauges, Peter Ghazal, Hideya Kawaji, Lu Li, Yukiko Matsuoka, Alice Villéger, Sarah E. Boyd, Laurence Calzone, Melanie Courtot, Ugur Dogrusoz, Tom C. Freeman, Akira Funahashi, Samik Ghosh, Akiya Jouraku, Sohyoung Kim, Fedor Kolpakov, Augustin Luna, Sven Sahle, Esther Schmidt, Steven Watterson, Guanming Wu, Igor Goryanin, Douglas B. Kell, Chris Sander, Herbert Sauro, Jacky L. Snoep, Kurt Kohn, and Hiroaki Kitano. The systems biology graphical notation. *Nature Biotechnology*, 27(8):735–741, 2009.

[Linderman *et al.* 2010] Jennifer J. Linderman, Thomas Riggs, Manjusha Pande, Mark Miller, Simeone Marino, and Denise E. Kirschner. Characterizing the dynamics of CD4+ T cell priming within a lymph node. *Journal of Immunology*, 184(6):2873–2885, 2010.

[Lindquist *et al.* 2004] Randall L. Lindquist, Guy Shakhar, Diana Dudziak, Hedda Warde-mann, Thomas Eisenreich, Michael L. Dustin, and Michel C. Nussenzweig. Visualizing dendritic cell networks in vivo. *Nature immunology*, 5(12):1243–1250, 2004.

[Lu *et al.* 2006] Linrong Lu, Miriam B. F. Werneck, and Harvey Cantor. The immunoregulatory effects of Qa-1. *Immunological Reviews*, 212:51–59, 2006.

[Luke *et al.* 2004] Sean Luke, Claudio Cioffi-Revilla, Liviu Panait, and Keith Sullivan. MASON: A new multi-agent simulation toolkit. In *Proceedings of the 2004 Swarmfest Workshop*, 2004.

[Luke 2009] Sean Luke. *Essentials of Metaheuristics*. Lulu, 2009. Available for free at <http://cs.gmu.edu/~sean/book/metaheuristics/>.

[Madakamutil *et al.* 2003] Loui T. Madakamutil, Igor Maricic, Eli Sercarz, , and Vipin Kumar. Regulatory T cells control autoimmunity in vivo by inducing apoptotic depletion of activated pathogenic lymphocytes. *The Journal of Immunology*, 170:2985–2992, 2003.

[Madakamutil *et al.* 2008] Loui T. Madakamutil, Igor Maricic, Eli Sercarz, and Vipin Kumar. Immunodominance in the TCR repertoire of a TCR peptide-specific CD4+ Treg population that controls experimental autoimmune encephalomyelitis. *The Journal of Immunology*, 180(1):4577–4585, 2008.

[Manabe & Bishai 2000] Yukari C. Manabe and William R. Bishai. Latent Mycobacterium tuberculosis- persistence, patience, and winning by waiting. *Nature Medicine*, 6(12):1327–1329, 2000.

[Marino *et al.* 2008] Simeone Marino, Ian B. Hogue, Christian J. Ray, and Denise E. Kirschner. A methodology for performing global uncertainty analysis and sensitivity analysis in systems biology. *Journal of Theoretical Biology*, 254(1):178–196, 2008.

[Marino *et al.* 2010] Simeone Marino, Amy Myers, JoAnne L. Flynn, and Denise E. Kirschner. TNF and IL-10 are major factors in modulation of the phagocytic cell environment in lung and lymph node in tuberculosis: A next-generation two-compartmental model. *Journal of Theoretical Biology*, 265(4):586–598, 2010.

[Matzinger 1994] Polly Matzinger. Tolerance, danger, and the extended family. *Annual Review of Immunology*, 12:991–1045, 1994.

[Matzinger 2002] Polly Matzinger. The danger model: A renewed sense of self. *Science*, 296(5566):301–305, April 2002.

[Matzinger 2007] Polly Matzinger. Friendly and dangerous signals: is the tissue in control? *Nature Immunology*, 8(1):11–13, 2007.

[McKay *et al.* 1979] M. D. McKay, R. J. Beckman, and W. J. Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2):239–245, 1979.

[Medzhitov & Janeway, Jr 2000] Ruslan Medzhitov and Charles A. Janeway, Jr. How does the immune system distinguish self from nonself? *Seminars in Immunology*, 12:185–188, 2000.

[Meier-Schellersheim & Mack 1999] M. Meier-Schellersheim and G. Mack. SIMMUNE, a tool for simulating and analyzing immune system behaviour (arxiv:cs/9903017v1). Available at <http://arxiv.org/abs/cs/9903017v1> (accessed on 12th June 2011), March 1999.

[Menezes *et al.* 2007] Juscilene S. Menezes, Peter van den Elzen, Jordan Thornes, Donald Huffman, Nathalie M. Droin, Emanual Maverakis, and Eli E. Sercarz. A public t cell clonotype within a heterogeneous autoreactive repertoire is dominant in driving eae. *The Journal of Clinical Investigation*, 117(8):2176–2185, August 2007.

[Milanesi *et al.* 2009] Luchiano Milanesi, Paolo Romano, Gastone Castellani, Daniel Remondini, and Pietro Liò. Trends in modeling biomedical complex systems. *BMC Bioinformatics*, 10(Suppl 12):I1, 2009.

[Owen *et al.* 2010] Jennifer Owen, Susan Stepney, Jonathan Timmis, and Alan F. T. Winfield. Exploiting loose horizontal coupling in evolutionary swarm robotics. In Marco Dorigo, Mauro Birattari, Gianni A. Di Caro, René Doursat, Andries Petrus Engelbrecht, Dario Floreano, Luca Maria Gambardella, Roderich Groß, Erol Sahin, Hiroki Sayama, and Thomas Stützle, editors, *Proceedings of the 7th International Conference on Swarm Intelligence ANTS 2010. Volume 6234 of Lecture Notes in Computer Science*, pages 432–439, 2010.

[Parijs & Abbas 1998] L. van Parijs and A. K. Abbas. Homeostasis and self-tolerance in the immune system: turning lymphocytes off. *Science*, 280(5361):243–248, 1998.

[Parnes 2004] Ohad Parnes. From interception to incorporation: degeneracy and promiscuous recognition as precursors of a paradigm shift in immunology. *Molecular Immunology*, 40(14–15):985–991, 2004.

[Polack *et al.* 2008] Fiona A. C. Polack, Tim Hoverd, Adam T. Sampson, Susan Stepney, and Jon Timmis. Complex systems models: Engineering simulations. In *ALife XI, Winchester, UK, September 2008*, pages 482–489. MIT Press, 2008.

[Polack *et al.* 2010] Fiona A. C. Polack, Paul S. Andrews, Teodor Ghetiu, Mark Read, Susan Stepney, Jon Timmis, and Adam T. Sampson. Reflections on the simulation of complex systems for science. In *ICECCS 2010: 15th IEEE International Conference on Engineering of Complex Computer Systems, Oxford, UK, March 2010*, pages 276–285. IEEE Press, 2010.

[Polack 2010] Fiona A. C. Polack. Arguing validation of simulations in science. In Susan Stepney, Peter H. Welsh, Paul S. Andrews, and Adam T. Sampson, editors, *Proceedings of the 2010 Workshop on Complex Systems Modelling and Simulation*, pages 51–74. Luniver Press, 2010.

[Prlic *et al.* 2007] Martin Prlic, Matthew A. Williams, and Michael J. Bevan. Requirements for CD8 T-cell priming, memory generation and maintenance. *Current Opinion in Immunology*, 19(3):315–319, 2007.

[Raivich & Banati 2004] Gennadij Raivich and Richard Banati. Brain microglia and blood-derived macrophages: molecular profiles and functional roles in multiple sclerosis and animal models of autoimmune demyelinating disease. *Brain Research Reviews*, 48(3):261–281, 2004.

[Ray *et al.* 2009] J. Christian J. Ray, JoAnne L. Flynn, and Denise E. Kirschner. Synergy between individual TNF-dependent functions determines granuloma performance for controlling mycobacterium tuberculosis infection. *Journal of Theoretical Biology*, 182(6):3706–3717, 2009.

[Read *et al.* 2009a] Mark Read, Paul S. Andrews, Jon Timmis, and Vipin Kumar. A domain model of experimental autoimmune encephalomyelitis. In Susan Stepney, Peter Welch, Paul S. Andrews, and Jon Timmis, editors, *Proceedings of the 2009 Workshop on Complex Systems Modelling and Simulation, York, UK, August 2009*, pages 9–44. Luniver Press, 2009.

[Read *et al.* 2009b] Mark Read, Paul S. Andrews, Jon Timmis, and Vipin Kumar. Using UML to model EAE and its regulatory network. In Paul S. Andrews, Jon Timmis, Nick D. L. Owens, Uwe Aickelin, Emma Hart, Andrew Hone, and Andy M. Tyrrell, editors, *ICARIS*, volume 5666 of *Lecture Notes in Computer Science*, pages 4–6. Springer, 2009.

[Read *et al.* 2011] Mark Read, Paul S. Andrews, Jon Timmis, and Vipin Kumar. Towards qualifying the implications of epistemic uncertainty on simulation based experimentation through calibration, uncertainty and sensitivity analysis. Accepted to Mathematical and Computer Modelling of Dynamical Systems, 2011.

[Rodgers & Cook 2005] John R. Rodgers and Richard G. Cook. MHC class Ib molecules bridge innate and acquired immunity. *Nature Reviews Immunology*, 5(6):459–471, 2005.

[Rumbaugh *et al.* 2005] James Rumbaugh, Ivar Jacobson, and Grady Booch. *The unified modelling language reference manual*. Addison-Wesley, 2 edition, 2005.

[Sadot *et al.* 2008] Avital Sadot, Jasmin Fisher, Dan Barak, Yishai Admanit, Michael J. Stern, E. Jan Albert Hubbard, and David Harel. Toward verified biological models. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 5(2):223–234, 2008.

[Saltelli *et al.* 2000] A. Saltelli, K. Chan, and E. M. Scott, editors. *Sensitivity Analysis*. Wiley series in probability and statistics. Wiley, 2000.

[Sampson *et al.* 2009] Adam T. Sampson, John Markus Bjørndalen, and Paul S. Andrews. Birds on the wall: Distributing a process-oriented simulation. In *2009 IEEE Congress on Evolutionary Computation (CEC 2009)*, pages 225–231. IEEE Press, 2009.

[Segovia-Juarez *et al.* 2004] Jose L. Segovia-Juarez, Suman Ganguli, and Denise Kirschner. Identifying control mechanisms of granuloma formation during M. tuberculosis infection using an agent-based model. *Journal of Theoretical Biology*, 231(3):357–376, 2004.

[Shevach *et al.* 2001] Ethan M. Shevach, Rebecca S. McHugh, Ciriaco A. Piccirillo, and Angela M. Thornton. Control of T-cell activation by CD4⁺ CD25⁺ suppressor T cells. *Immunological Reviews*, 182:58–67, 2001.

[Silverstein & Rose 2000] Arthur M. Silverstein and Noel R. Rose. There is only one immune system! The view from immunopathology. *Seminars in Immunology*, 12:173–178, 2000.

[Smith & Kumar 2008] Trevor R.F. Smith and Vipin Kumar. Revival of CD8⁺ Treg mediated suppression. *Trends in Immunology*, 29(7):337–342, 2008.

[Smith *et al.* 2009] Trevor R. Smith, Xiaolei Tang, Igor Maricic, Zacarias Garcia, Shaohsuan Fanchiang, and Vipin Kumar. Dendritic cells use endocytic pathway for cross-priming class Ib MHC-restricted CD8 $\alpha\alpha^+$ TCR $\alpha\beta^+$ T cells with regulatory properties. *Journal of Immunology*, 182(11):6959–6968, 2009.

[Smith *et al.* 2010] Trevor R. Smith, Igor Maricic, Francesco Ria, Susan Schneider, and Vipin Kumar. CD8 α^+ dendritic cells prime TCR-peptide-reactive regulatory CD4⁺FOXP3[−] T cells. *European Journal of Immunology*, 40(7):1906–1915, 2010.

[Sorathiya *et al.* 2009] Anil Sorathiya, Pietro Liò, and Luca Sguanci. Mathematical model of HIV superinfection and comparative drug therapy. In Paul S. Andrews, Jon Timmis, Nick D. L. Owens, Uwe Aickelin, Emma Hart, Andrew Hone, and Andy M. Tyrrell, editors, *Proceedings of the 8th International Conference on Artificial Immune Systems (ICARIS), Lecture Notes in Computer Science volume 5666*, pages 41–53. Springer, 2009.

[Sousa 2006] Caetano R. e Sousa. Dendritic cells in a mature age. *Nature Reviews Immunology*, 6(6):476–83, 2006.

[Stepney *et al.* 2008] Susan Stepney, Fiona Polack, and Peter Welch, editors. *Proceedings of the 2008 Workshop on Complex Systems Modelling and Simulation, York, UK, September 2008*. Luniver Press, 2008.

[Storn & Price 1995] Rainer Storn and Kenneth Price. Differential evolution - a simple and efficient adaptive scheme for global optimization over continuous spaces. Technical Report TR-95-012, Berkeley, CA, 1995.

[Swerdlin *et al.* 2008] Na’aham Swerdlin, Irun R. Cohen, and David Harel. The lymph node B cell immune response: dynamic analysis in silico. *Proceedings of the IEEE*, 96(8):1421 – 1443, 2008.

[Tambuyzer *et al.* 2009] Bart R. Tambuyzer, Peter Ponsaerts, and Etienne J. Nouwen. Mi-croglia: Gatekeepers of central nervous system immunology. *Journal of Leukocyte Biology*, 85(3):352–370, 2009.

[Tang *et al.* 2005] Xiaolei Tang, Trevor R. F. Smith, and Vipin Kumar. Specific control of immunity by regulatory CD8 T cells. *Cellular and Molecular Immunology*, 1(2):11–19, 2005.

[Tang *et al.* 2006] Xiaolei Tang, Igor Maricic, Nikunj Purohit, Berge Bakamjian, Lisa M Read-Loisel, Tara Beeston, Peter Jensen, and Vipin Kumar. Regulation of immunity by a novel population of Qa-1-restricted CD8 $\alpha\alpha^+$ TCR $\alpha\beta^+$ T cells. *The Journal of Immunology*, 177:7645–7655, 2006.

[Tang *et al.* 2007] Xiaolei Tang, Igor Maricic, and Vipin Kumar. Anti-TCR antibody treatment activates a novel population of nonintestinal CD8 $\alpha\alpha^+$ TCR $\alpha\beta^+$ regulatory t cells and prevents experimental autoimmune encephalomyelitis. *Journal of Immunology*, 178(10):6043–6050, 2007.

[Vargha & Delaney 2000] András Vargha and Harold D. Delaney. A critique and improvement of the cl common language effect size statistics of McGraw and Wong. *Journal of Educational and Behavioral Statistics*, 25(2):101–132, 2000.

[Warrender *et al.* 2006] Christina Warrender, Stephanie Forrest, and Frederick Koster. Modeling intercellular interactions in early Mycobacterium infection. *Bulletin of Mathematical Biology*, 68(8):2233–2261, 2006.

[Wath & Liò 2008] Richard C. van der Wath and Pietro Liò. A stochastic single cell based model of BrdU measured hematopoietic stem cell dynamics. In Monika Heiner and Adelinde Uhrmacher, editors, *Proceedings of Computational Methods in Systems Biology, Lecture Notes in Computer Science volume 5307*, pages 387–401, 2008.

[Welch & Barnes 2005] Peter H. Welch and Frederick R. M. Barnes. Communicating mobile processes: Introducing occam-pi. In A. E. Abdallah, C. B. Jones, and J. W. Sanders, editors, *In 25 years of CSP, Lecture Notes in Computer Science volume 3525*, pages 175–200, 2005.

[Williams *et al.* 2011] Richard Williams, Mark Read, Jon Timmis, Paul S. Andrews, and Vipin Kumar. In silico investigation into CD8Treg mediated recovery in murine Experimental Autoimmune Encephalomyelitis. In Pietro Liò, Giuseppe Nicosia, and Thomas Stibor, editors, *Proceedings of the 10th International Conference on Artificial Immune Systems (ICARIS), Lecture Notes in Computer Science volume 6825*, pages 51–54, 2011.

[Zamvil & Steinman 1990] S. S. Zamvil and L. Steinman. The T lymphocyte in experimental allergic encephalomyelitis. *Annual Review of Immunology*, 8(1):579–621, 1990.

[Zepp *et al.* 2011] Jarod Zepp, Ling Wu, and Xiaoxia Li. IL-17 receptor signaling and t helper 17-mediated autoimmune demyelinating disease. *Trends in Immunology*, 32(5):232–239, 2011.

BIBLIOGRAPHY

[Zheng *et al.* 2008] Huan Zheng, Bo Jin, Sarah E. Henrickson, Alan S. Perelson, Ulrich H. von Andrian, and Arup K. Chakraborty. How antigen quantity and quality determine T-cell decisions in lymphoid tissue. *Molecular and Cellular Biology*, 28(12):4040–4051, 2008.

[Zinkernagel 2000] Rolf M. Zinkernagel. Localization dose and time of antigens determine immune reactivity. *Seminars in Immunology*, 12:163–171, 2000.