The Advantages of Fat, Fashion-Conscious Supermodels...

Dr John Szymanski, February 2005-02-04
Department of Electronics, University of York, Heslington, York YO10 5DD

Introduction

The stereotypical image of (human!) models as the vacuous, underfed clotheshorses of the fashion industry is well entrenched in popular culture - everyone knows that the models are (unhealthily) thinner than they should be, yet, decade after decade, they stay thin. Can it be argued that the same is true of our mathematical and algorithmic models - have we developed a preference for overly "lean" models and, if we want to investigate complex adaptive systems, is it time to start fleshing them out a bit?

Thin models

In common with established practice for scientific design or analysis, when investigating bio-inspired computational techniques, or developing models of complex adaptive systems, the researcher cannot help but impose some elements of their specific viewpoint, opinions or prior expectations on their work. Although a particular algorithm may be "bio-inspired" to some degree, each individual person, of course, tends to find their inspiration in different ways. As researchers, traditionally we desire complexity in the outputs but, since we are usually seeking to increase our "understanding", we tend to avoid complexity within the models - in order to deduce causes from effects, we include only those selected causative processes which we both understand and believe significant.

Hence, in practice, each researcher tends to make an assortment of prior decisions (implicit or explicit, intentional or accidental) regarding the relative importance of different model elements - often to the extent of ignoring certain aspects completely in order to reduce the number of parameters, or to keep the model complexity to a level where they believe that its behaviour can be readily analysed and understood. This approach has served science well for a long time, but is inherently bound to the principle of superposition - are we bogged down in the idea that we can add characteristics to a system model one at a time, and hence to some degree analyse and "understand" the impact of each change?

It is a given that no mathematical model of any "real" system can ever be complete. And it is also undeniable that, excepting certain "thought experiments", there is little point in creating a model or algorithm that is so complicated that it cannot be implemented. This raises the question of just what is "enough" when creating models of complex systems. The argument here is that when working with non-classical computation, although we really have little choice but to work with models which incorporate some degree of simplification, we should attempt to be as inclusive as possible. In essence, whether or not we accept the "edge of chaos" viewpoint that some of the most interesting areas for investigation lie at the phase transitions between order and disorder, it is surely important that the results should display phase transitions that are in some sense "adequate", "stable" and "robust", rather than just some reduced projection or shadow of a larger picture - in short, when dealing with emergence, we cannot ever be sure of the potential importance of any of those model elements which we choose to omit.

The problem is that very interesting results *can* be produced by overly simple systems or rules - any school student can get some experience of emergence by playing with cellular automata via the game of Life, or by exploring the patterns of computed fractal clouds, trees and landscapes. But how can we evaluate the relative *quality* of the results of such overly simple systems - is one model (rule set, search strategy, cost function, etc.) in some sense better than another? Can any conclusions be transferred or generalised to other cases? There is a real danger that we will have research projects investigating literally thousands of slightly different approaches - one group concentrating on the importance of agents being able to learn from their environment and modify their behaviour, with another performing schema analyses of recombination and mutation in evolutionary algorithms, attempting to characterise when these processes might help/hinder performance, and yet another instead pursuing the significance of co-adaptation. All such projects will have the potential to deliver novel and interesting results - but which is best? Who should get the funding?

Fat, fashion-conscious, supermodels

A "thin" model may well be able to produce interesting results that exhibit emergent complexity, but the overall behaviour may change unpredictably as new, apparently trivial, features are introduced into the model. Questions arise regarding issues such as transferability, repeatability and robustness.

If we truly wish to study complex adaptive systems, perhaps we need to abandon the idea of simple, thin models and accept the fact that we need to embrace "fat, fashionable supermodels" - models that are well overdetermined in terms of their internal structure (fat), with organisms/agents that adapt their behaviour and characteristics with time and external stimuli (fashion-conscious), and that reflect not only the characteristics of the agents, but also its relationship with the environment in which the agents exist (a "supermodel"). Drawing together existing ideas of autonomous agents, co-adaptation, maturity, etc. a "fatter", more inclusive approach would need to involve multiple populations of organisms which interact with each other and with their environment, with multiple rule sets that are both dynamic and more complicated than those commonly used at present.

- The environment should:
  - represent a changing multidimensional landscape (multiple ecosystems)
  - reflect variable resource types and levels (both with time and location)
  - influence the organisms (health, happiness...)
  - be influenced by the organisms
  - be dynamic (seasonal/cyclic, random effects, windfalls, catastrophes...)
- The agents/organisms should:
  - involve multiple populations (predator, prey, symbiosis, parasite...)
  - be adaptive (with time, to environment, to other agents, in multiple ways...)
  - have multiple adaptive motivations (survival, propagation, security, desire...)
  - have multiple survival techniques (range, search strategy, retreat, fight...)
  - have multiple propagation strategies (spread, procreate - quantity, frequency, etc.)
  - have the potential for a wide range of interactions with others (switchable rule sets...)
  - have memory (age, health, happiness, maturity, boredom...)
  - be able to evolve (genetics, learning...)
  - to be able to communicate behaviour and characteristics to other agents
  - influence the environment
  - be influenced by the environment

Further, all of the parameters defining the various interactions and relationships above should be able to adapt freely, and allow switching between alternative local stimulus-response rule sets. The agents should have the potential flexibility to learn, to vary widely in behaviour, to co-operate or fight, and to switch behaviours when certain thresholds (panic, need, boredom?) are reached.

Conclusion

In essence, the point here is that we need to work towards a truly general model of adaptive agents that can be used for widely different applications - the place for the application-specific information should be in the way in which we set up the environment and its dynamics, rather than within the local strategic and stimulus-response rules that govern the behaviour of the agents.

We really need to take the idea of bio-inspired methods to its logical limit. We must note that it is the environment that determines the various forms of life that thrive in different regions, not the basic rules of genetics, search strategies or interaction processes - whether microscopic or macroscopic, in desert or in tundra, in deep water or in air, organisms survive, but take on very different forms. The conclusion is that for complex systems research we need to develop as complete and inclusive (fat) model of a fully adaptive (fashion-conscious) organism as possible, and of the ways in which it interacts with its environment (the supermodel) - and then let evolution decide which of the various rules and parameters - are most appropriate to the different (and dynamic) environments that we have specified *a priori* for different applications.

So don't try to date a supermodel - implement one instead!