Multi-objectivity as a Tool for Constructing Hierarchical Complexity

Jason Teo, Minh Ha Nguyen, and Hussein A. Abbass

Artificial Life and Adaptive Robotics (A.L.A.R.) Lab,
School of Computer Science, University of New South Wales,
Australian Defence Force Academy Campus, Canberra, Australia.
{j.teo,m.nguyen,h.abbass}@adfa.edu.au

Abstract. This paper presents a novel perspective to the use of multi-objective optimization and in particular evolutionary *multi-objective optimization* (EMO) as a measure of complexity. We show that the partial order feature that is being inherited in the Pareto concept exhibits characteristics which are suitable for studying and measuring the complexities of embodied organisms. We also show that multi-objectivity provides a suitable methodology for investigating complexity in artificially evolved creatures. Moreover, we present a first attempt at quantifying the morphological complexity of quadruped and hexapod robots as well as their locomotion behaviors.

1 Introduction

The study of complex systems has attracted much interest over the last decade and a half. However, the definition of what makes a system complex is still the subject of much debate among researchers [7,19]. There are numerous methods available in the literature for measuring complexity. However, it has been argued that complexity measures are typically too difficult to compute to be of use for any practical purpose or intent [16]. What we are proposing in this paper is a simple and highly accessible methodology for characterizing the complexity of artificially evolved creatures using a multi-objective methodology. This work poses evolutionary multi-objective optimization (EMO) [5] as a convenient platform which researchers can utilize practically in attempting to define, measure or simply characterize the complexity of everyday problems in a useful and purposeful manner.

2 Embodied Cognition and Organisms

The view of intelligence in traditional AI and cognitive science has been that of an agent undertaking some form of information processing within an abstracted representation of the world. This form of understanding intelligence was found to be flawed in that the agent's cognitive abilities were derived purely from a processing unit that manipulates symbols and representations far abstracted from

E. Cantú-Paz et al. (Eds.): GECCO 2003, LNCS 2723, pp. 483–494, 2003.

the agent's real environment [3]. Conversely, the embodied cognitive view considers intelligence as a phenomenon that emerges independently from the parallel and dynamical interactions between an embodied organism and its environment [14]. Such artificial creatures possess two important qualities: embodiment and situatedness.

A subfield of research into embodied cognition involves the use of artificial evolution for automatically generating the morphology and mind of embodied creatures [18]. The term mind as used in this context of research is synonymous with brain and controller - it merely reflects the processing unit that acts to transform the sensory inputs into the motor outputs of the artificial creature. The automatic synthesis of such embodied and situated creatures through artificial evolution has become a key area of research not only in the cognitive sciences but also in robotics [15], artificial life [14], and evolutionary computation [2,10].

Consequently, there has been much research interest in evolving both physically-simulated virtual organisms [2,10,14] and real physical robots [15,8,12]. The main objective of these studies is to evolve increasingly complex behaviors and/or morphologies either through evolutionary or lifetime learning. Needless to say, the term "*complex*" is generally used very loosely since there is currently no general method for comparing between the complexities of these evolved artificial creatures' behaviors and morphologies. As such, without a quantitative measure for behavioral or morphological complexity, an objective evaluation between these artificial evolutionary systems becomes very hard and typically ends up being some sort of subjective argument.

There are generally two widely-accepted views of measuring complexity. The first is an information-theoretic approach based on Shannon's entropy [17] and is commonly referred to as statistical complexity. The entropy $H(X)$ of a random variable X, where the outcomes x_i occur with probability p_i, is given by

$$H(X) = -C \sum_{i}^{N} p_i \log p_i \tag{1}$$

where C is the constant related to the base chosen to express the logarithm. Entropy is a measure of disorder present in a system and thus gives us an indication of how much we do not know about a particular system's structure. Shannon's entropy measures the amount of information content present within a given message or more generally any system of interest. Thus a more complex system would be expected to give a much higher information content than a less complex system. In other words, a more complex system would require more bits to describe compared to a less complex system. In this context, a sequence of random numbers will lead to the highest entropy and consequently to the lowest information content. In this sense, complexity is somehow a measure of order or disorder.

A computation-theoretic approach to measuring complexity is based on Kolmogorov's application of universal Turing machines [11] and is commonly known as Kolmogorov complexity. It is concerned with finding the shortest possible computer program or any abstract automaton that is capable of reproducing a given string. The Kolmogorov complexity $K(s)$ of a string s is given by

$$K(s) = \min\{|p| : s = C_T(p)\} \tag{2}$$

where $|p|$ represents the length of program p and $C_T(p)$ represents the result of running program p on Turing machine T. A more complex string would thus require a longer program while a simpler string would require a much shorter program. In essence, the complexity of a particular system is measured by the amount of computation required to recreate the system in question.

3 Complexity in the Eyes of the Beholder

None of the previous measures are sufficient to measure the complexity of embodied systems. As such, we need first to provide a critical view of these measures and why they stand shorthanded in terms of embodied systems.

Take for example a simple behavior such as walking. Let us assume that we are interested in measuring the complexity of walking in different environments and the walking itself is undertaken by an artificial neural network. From Shannon's perspective, the complexity can be measured using the entropy of the data structure holding the neural network. Obviously a drawback for this view is its ignorance of the context and the concepts of embodiment and situatedness. The complexity of walking on a flat landscape is entirely different from walking on a rough landscape. Two neural networks may be represented using the same number of bits but exhibit entirely different behaviors.

Now, let us take another example which will show the limitations of Kolmogorov complexity. Assume we have a sequence of random numbers. Obviously the shortest program which is able to reproduce this sequence is the sequence itself. In other words, a known drawback for Kolmogorov complexity is that it has the highest level of complexity when the system is random. In addition, let us re-visit the neural network example. Assume that the robot is not using a fixed neural network but some form of evolvable hardware (which may be an evolutionary neural network). If the fitness landscape for the problem at hand is monotonically increasing, a hill climber will simply be the shortest program which guarantees to re-produce the behavior. However, if the landscape is rugged, reproducing the behavior is only achievable if we know the seed; otherwise, the problem will require complete enumeration to recreate the behavior.

In this paper, we propose a generic definition for complexity using the multi-objective paradigm. However, before we proceed with our definition, we first remind the reader of the concept of partial order.

Definition 1: Partial and Lexicographic Order Assume the two sets A and B. Assume the l-subsets over A and B such that $A = \{a_1 < \ldots < a_l\}$ and $B = \{b_1 < \ldots < b_l\}$.
A partial order is defined as $A \leq_j B$ if $a_j \leq b_j, \ \forall j \in \{1, \ldots, l\}$
A lexicographic order is defined as $A <_j B$ if $\exists a_k < b_k$ and $a_j = b_j, \ j < k$, $\forall j, k \in \{1, \ldots, l\}$

In other words, a lexicographic order is a total order. In multi-objective optimization, the concept of Pareto optimality is normally used. A solution x belongs to the Pareto set if there is not a solution y in the feasible solution set such that y dominates x (*ie.* x has to be at least as good as y when measured on all objectives and better than y on at least one objective). The Pareto concept thus forms partial orders in the objective space.

Let us recall the embodied cognition problem. The problem is to study the relationship between the behavior, controller, environment, learning algorithm, and morphology. A typical question that one may ask is what is the optimal behavior for a given morphology, controller, learning algorithm and environment. We can formally represent the problem of embodied cognition as the five sets B, C, E, L, and M for the five spaces of behavior, controller, environment, learning algorithm, and morphology respectively. Here, we need to differentiate between the robot behavior B and the desired behavior $\hat{B}$. The former can be seen as the actual value of the fitness function and the latter can be seen as the real maximum of the fitness function. For example, if the desired behavior (task) is to maximize the locomotion distance, then the global maximum of this function is the desired behavior, whereas the distance achieved by the robot (what the robot is actually doing) is the actual behavior. In traditional robotics, the problem can be seen as Given the desired behavior $\hat{B}$, find L which optimizes C subject to $E \bigcup M$. In psychology, the problem can be formulated as Given C, E, L and M, study the characteristics of the set B. In co-evolving morphology and mind, the problem is Given the desired behavior $\hat{B}$ and L, optimize C and M subject to E. A general observation is that the learning algorithm is usually fixed during the experiments.

In asking a question such as "Is a human more complex than a Monkey?", a natural question that follows would be "in what sense?". Complexity is not a unique concept. It is usually defined or measured within some context. For example, a human can be seen as more complex than a Monkey if we are looking at the complexity of intelligence, whereas a Monkey can be seen as more complex than the human if we are looking at the number of different gaits the monkey has for locomotion. Therefore, what is important from an artificial life perspective is to establish the complexity hierarchy on different scales. Consequently, we introduce the following definition for complexity.

Definition 2: Complexity is a strict partial order relation.

According to this definition, we can establish an order of complexity between the system's components/species. We can then compare the complexities of two species $S_1 = (B_1, C_1, E_1, L_1, M_1)$ and $S_2 = (B_2, C_2, E_2, L_2, M_2)$ as:

$$S_1 \text{ is at least as complex as } S_2 \textbf{ with respect to concept } \Psi \textbf{ iff}$$

$$S_2^{\Psi} = (B_2, C_2, E_2, L_2, M_2) \leq_j S_1^{\Psi} = (B_1, C_1, E_1, L_1, M_1), \forall j \in \{1, \dots, l\}, \textbf{Given}$$
$$B_i = \{B_{i1} < \dots < B_{il}\}, C_i = \{C_{i1} < \dots < C_{il}\}, E_i = \{E_{i1} < \dots < E_{il}\},$$
$$L_i = \{L_{i1} < \dots < L_{il}\}, M_i = \{M_{i1} < \dots < M_{il}\},\ i \in \{1, 2\} \tag{3}$$

where Ψ partitions the sets into l non-overlapping subsets.

We can even establish a complete order of complexity by using the lexicographic order as:

$$S_1 \text{ is more complex than} S_2 \textbf{ with respect to concept } \Psi \textbf{ iff}$$

$$S_2^{\Psi} = (B_2, C_2, E_2, L_2, M_2) <_j S_1^{\Psi} = (B_1, C_1, E_1, L_1, M_1), \forall j \in \{1, \ldots, l\}, \textbf{Given}$$
$$B_i = \{B_{i1} < \ldots < B_{il}\}, C_i = \{C_{i1} < \ldots < C_{il}\}, E_i = \{E_{i1} < \ldots < E_{il}\},$$
$$L_i = \{L_{i1} < \ldots < L_{il}\}, M_i = \{M_{i1} < \ldots < M_{il}\},\ i \in \{1, 2\} \tag{4}$$

The lexicographic order is not as flexible as partial order since the former requires a monotonic increase in complexity. The latter however, allows individuals to have similar levels of complexity; therefore, it is more suitable for defining hierarchies of complexity. Some of the characteristics in our definition of complexity here include

Irreflexive: The complexity definition satisfies irreflexivity; that is, x cannot be more complex than itself.

Asymmetric: The complexity definition satisfies asymmetry; that is, if x is more complex than y, then y cannot be more complex than x.

Transitive: The complexity definition satisfies transitivity; that is, if x is more complex than y and y is more complex than z, then x is more complex than z.

The concept of Pareto optimality is similar to the concept of partial order except that Pareto optimality is more strict in the sense that it does not satisfy reflexivity; that is, a solution cannot dominate itself; therefore it cannot exist as a Pareto optimal if there is a copy of it in the solution set. Usually, when we have copies of one solution, we take one of them; therefore this problem does not arise. As a result, we can assume here that Pareto optimality imposes a complexity hierarchy on the solution set.

The previous definition will simply order the sets based on their complexities according to some concept Ψ. However, they do not provide an exact quantitative measure for complexity. In the simple case, given the five sets B, C, E, L, and M; assume the function f, which maps each element in each set to some value called the fitness, and assuming that C, E and L do not change, a simple measure of morphological change of complexity can be

$$\frac{\partial f(b)}{\partial m}, b \in B, m \in M \tag{5}$$

In other words, assuming that the environment, controller, and the learning algorithm are fixed, the change in morphological complexity can be measured in the eyes of the change in the fitness of the robot (actual behavior). The fitness will be defined later in the paper. Therefore, we introduce the following definition

Definition 3: Change of Complexity Value for the morphology is the rate of change in behavioral fitness when the morphology changes, given that both the environment, learning algorithm and controller are fixed.

The previous definition can be generalized to cover the controller and environment quite easily by simply replacing "morphology" by either "environment", "learning algorithm", or "controller". Based on this definition, if we can come up with a good measure for behavioral complexity, we can use this measure to quantify the change in complexity for morphology, controller, learning algorithm, or environment. In the same manner, if we have a complexity measure for the controller, we can use it to quantify the change of complexity in the other four parameters. Therefore, we propose the notion of defining the complexity of one object as viewed from the perspective of another object. This is not unlike Emmeche's idea of complexity as put in the eyes of the beholder [6]. However, we formalize and solidify this idea by putting it into practical and quantitative usage through the multi-objective approach. We will demonstrate that results from an EMO run of two conflicting objectives results in a Pareto-front that allows a comparison of the different aspects of an artificial creature's complexity.

In the literature, there are a number of related topics which can help here. For example, the VC-dimension can be used as a complexity measure for the controller. A feed-forward neural network using a threshold activation function has a VC dimension of $O(WlogW)$ while a similar network with a sigmoid activation has a VC dimension of $O(W^2)$, where W is the number of free parameters in the network [9]. It is apparent from here that one can control the complexity of a network by minimizing the number of free parameters which can be done either by the minimization of the number of synapses or the number of hidden units. It is important to separate between the learning algorithm and the model itself. For example, two identical neural networks with fixed architectures may perform differently if one of them is trained using back-propagation while the other is trained using an evolutionary algorithm. In this case, the separation between the model and the algorithm helps us to isolate their individual effects and gain an understanding of their individual roles.

In this paper, we are essentially posing two questions, what is the change of (1) behavioral complexity and (2) morphological complexity of the artificial creature in the eyes of its controller. In other words, how complex is the behavior and morphology in terms of evolving a successful controller?

3.1 Assumptions

Two assumptions need to be made. First, the Pareto set obtained from evolution is considered to be the actual Pareto set. This means that for the creature on the Pareto set, the maximum amount of locomotion is achieved with the minimum number of hidden units in the ANN. We do note however that the evolved Pareto set in the experiments may not have converged to the optimal set. Nevertheless, it is not the objective of this paper to provide a method which guarantees convergence of EMO but rather to introduce and demonstrate the application of measuring complexity in the eyes of the beholder. It is important to mention that although this assumption may not hold, the results can still be valid. This will be the case when creatures are not on the actual Pareto-front but the distances between them on the intermediate Pareto-front are similar to that of creatures on the actual Pareto-front.

The second assumption is there are no redundancies present in the ANN architectures of the evolved Pareto set. This simply means that all the input and output units as well as the synaptic connections between layers of the network are actually involved in and required for achieving the observed locomotion competency. We have investigated the amount of redundancy present in evolved ANN controllers and found that the self-adaptive Pareto EMO approach produces networks with practically zero redundancy.

4 Methods

4.1 The Virtual Robots and Simulation Environment

The Vortex physics simulation toolkit [4] was utilized to accurately simulate the physical properties, such as forces, torques, inertia, friction, restitution and damping, of and interactions between the robot and its environment. Two artificial creatures (Figure 1) were used in this study.

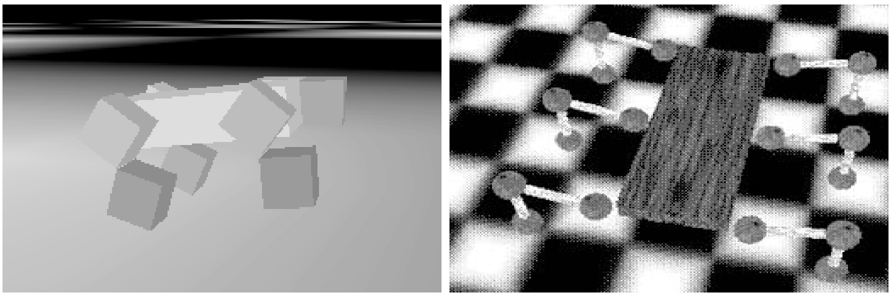

Fig. 1. The four-legged (quadruped) and six-legged (hexapod) creatures.

The first artificial creature is a quadruped with 4 short legs. Each leg consists of an upper limb connected to a lower limb via a hinge (1 degree-of-freedom (DOF)) joint and is in turn connected to the torso via another hinge joint. Each of the hinge joints is actuated by a motor that generates a torque producing rotation of the connected body parts about that hinge joint.

The second artificial creature is a hexapod with 6 long legs, which are connected to the torso by insect hip joints. Each insect hip joint consists of two hinges, making it a 2 DOF joint: one to control the back-and-forth swinging and another for the lifting of the leg. Each leg has an upper limb connected to a lower limb by a hinge (1 DOF) joint. The hinges are actuated by motors in the same fashion as in the first artificial creature.

The Pareto-frontier of our evolutionary runs are obtained from optimizing two conflicting objectives: (1) minimizing the number of hidden units used in the ANN that act as the creature's controller and (2) maximizing horizontal locomotion distance of the artificial creature. What we obtain at the end of the runs are Pareto sets of ANNs that trade-off between number of hidden units and locomotion distance. The locomotion distances achieved by the different Pareto solutions will provide a common ground where locomotion competency can be used to compare different behaviors and morphologies. It will provide a set of ANNs with the smallest hidden layer capable of achieving a variety of locomotion competencies. The structural definition of the evolved ANNs can now be used as a measure of complexity for the different creature behaviors and morphologies.

The ANN architecture used in this study is a fully-connected feed-forward network with recurrent connections on the hidden units as well as direct input-output connections. Recurrent connections were included to allow the creature's controller to learn time-dependent dynamics of the system. Direct input-output connections were also included in the controller's architecture to allow for direct sensor-motor mappings to evolve that do not require hidden layer transformations. Bias is incorporated in the calculation of the activation of the hidden as well as output layers.

The Self-adaptive Pareto-frontier Differential Evolution algorithm (SPDE) [1] was used to drive the evolutionary optimization process. SPDE is an elitist approach to EMO where both crossover and mutation rates are self-adapted. Our chromosome is a class that contains one matrix Ω and one vector ρ. The matrix Ω is of dimension $(I + H) \times (H + O)$. Each element $\omega_{ij} \in \Omega$, is the weight connecting unit i with unit j, where $i = 0, \ldots, (I-1)$ is the input unit i, $i = I, \ldots, (I+H-1)$ is the hidden unit $(i-I)$, $j = 0, \ldots, (H-1)$ is the hidden unit j, and $j = H, \ldots, (H+O-1)$ is the output unit $(j-H)$.

The vector ρ is of dimension H, where $\rho_h \in \rho$ is a binary value used to indicate if hidden unit h exists in the network or not; that is, it works as a switch to turn a hidden unit on or off. Thus, the architecture of the ANN is variable in the hidden layer: any number of hidden units from 0 to H is permitted. The sum, $\sum_{h=0}^{H} \rho_h$, represents the actual number of hidden units in a network, where H is the maximum number of hidden units. The last two elements in the chromosome are the crossover rate δ and mutation rate η. This representation allows simultaneous training of the weights in the network and selecting a subset of hidden units as well as allowing for the self-adaptation of crossover and mutation rates during optimization.

4.2 Experimental Setup

Two series of experiments were conducted. Behavioral complexity was investigated in the first series of experiments and morphological complexity was investigated in the second. For both series of experiments, each evolutionary run was allowed to evolve over 1000 generations with a randomly initialized population size of 30. The maximum number of hidden units was fixed at 15 based on preliminary experimentation. The number of hidden units used and maximum locomotion achieved for each genotype evaluated as well as the Pareto set of solutions obtained in every generation were recorded. The Pareto solutions obtained at the completion of the evolutionary process were compared to obtain a characterization of the behavioral and morphological complexity.

To investigate behavioral complexity in the eyes of the controller, the morphology was fixed by using only the quadruped creature but the desired behavior was varied by having two different fitness functions. The first fitness function measured only the maximum horizontal locomotion achieved but the second fitness function measured both maximum horizontal locomotion and static stability achieved. By static stability, we mean that the creature achieves a statically stable locomotion gait with at least three of its supporting legs touching the ground during each step of its movement. The two problems we have are:
($P1$)

$$f_1 = d \tag{6}$$

$$f_2 = \sum_{h=0}^{H} \rho_h \tag{7}$$

($P2$)

$$f_1 = d/20 + s/500 \tag{8}$$

$$f_2 = \sum_{h=0}^{H} \rho_h \tag{9}$$

where $P1$ and $P2$ are the two sets of objectives used. d refers to the locomotion distance achieved and s is the number of times the creature is statically stable as controlled by the ANN at the end of the evaluation period of 500 timesteps. $P1$ is using the locomotion distance as the first objective while $P2$ is using a linear combination of the locomotion distance and static stability. Minimizing the number of hidden units is the second objective in both problems.

To investigate morphological complexity, another set of 10 independent runs was carried out but this time using the hexapod creature. This is to enable a comparison with the quadruped creature which has a significantly different morphology in terms of its basic design. The $P1$ set of objectives was used to keep the behavior fixed. The results obtained in this second series of experiments were then compared against the results obtained from the first series of experiments where the quadruped creature was used with the $P1$ set of objective functions.

5 Results and Discussion

5.1 Morphological Complexity

We first present the results for the quadruped and hexapod evolved under $P1$. Figure 2 compares the Pareto optimal solutions obtained for the two different morphologies over 10 runs. Here we are fixing E and L; therefore, we can either measure the change of morphological complexity in the eyes of the behavior or the controller; that is, $\frac{\delta f(B)}{\delta M}$ or $\frac{\delta f(C)}{\delta M}$ respectively. If we fix the actual behavior B as the locomotion competency of achieving a movement of $13 < d < 15$,

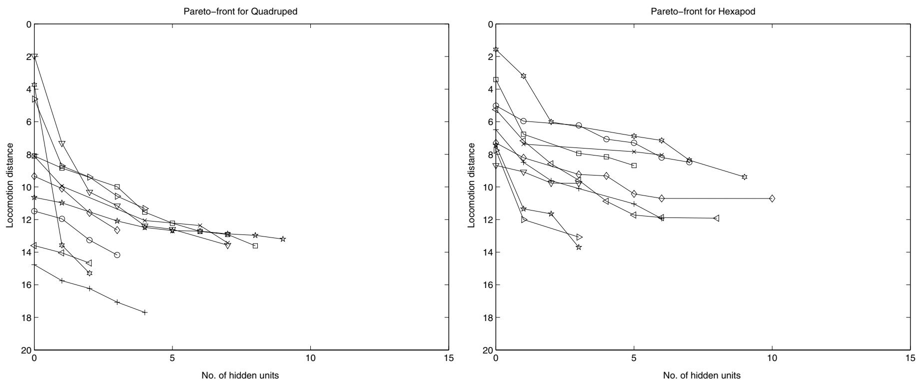

Fig. 2. Pareto-frontier of controllers obtained from 10 runs using the quadruped and hexapod with the $P1$ set of objectives.

then the change in the controller $\delta f(C)$ is measured according to the number of hidden units used in the ANN. At this point of comparison, we find that the quadruped is able to achieve the desired behavior with 0 hidden units whereas the hexapod required 3 hidden units. In terms of the ANN architecture, the quadruped achieved the required level of locomotion competency without using the hidden layer at all, that it relied solely on direct input-output connections as in a perceptron. This phenomenon has been previously observed to occur in wheeled robots as well [13]. Therefore, this is an indication that from the controller's point of view, given the change in morphology δM from the quadruped to the hexapod, there was an increase in complexity for the controller δC from 0 hidden units to 3 hidden units. Hence, the hexapod morphology can be seen as being placed at a higher level of the complexity hierarchy than the quadruped morphology in the eyes of the controller.

If we would like to measure the complexity of the morphology using the behavioral scale, we can notice from the graph that the maximum distance achieved by the quadruped creature is around 17.8 compared to around 13.8 for the hexapod creature. In this case, the quadruped can be seen as being able to achieve a more complex behavior than the hexapod.

5.2 Behavioral Complexity

A comparison of the results obtained using the two different sets of fitness functions $P1$ and $P2$ is presented in Table 1. Here we are fixing M, L and E and looking for the change in behavioral complexity. The morphology M is fixed by using the quadruped creature only. For $P1$, we can see that the Pareto-frontier offers a number of different behaviors. For example, a network with no hidden units can achieve up to 14.7 units of distance while the creature driven by a network with 5 hidden units can achieve 17.7 units of distance within the 500

Table 1. Comparison of global Pareto optimal controllers evolved for the quadruped using the $P1$ and $P2$ objective functions.

Type of Behavior	Pareto Controller	No. of Hidden Units	Locomotion Distance	Static Stability
P1	1	0	14.7	19
	2	1	15.8	24
	3	2	16.2	30
	4	3	17.1	26
	5	4	17.7	14
P2	1	0	5.2	304
	2	1	3.3	408
	3	2	3.6	420
	4	3	3.7	419

timesteps. This is an indication that to achieve a higher speed gait entails a more complex behavior than a lower speed gait.

We can also see the effect of static stability, which requires a walking behavior. By comparing a running behavior using a dynamic gait in $P1$ with no hidden units against a walking behavior using a static gait in $P2$ with no hidden units, we can see that using the same number of hidden units, the creature achieves both a dynamic as well as a quasi-static gait. If more static stability is required, this will necessitate an increase in controller complexity.

At this point of comparison, we find that the behavior achieved with the $P1$ fitness functions consistently produced a higher locomotion distance than the behavior achieved with the $P2$ fitness functions. This meant that it was much harder for the $P2$ behavior to achieve the same level of locomotion competency in terms of distance moved as the $P1$ behavior due to the added sub-objective of having to achieve static stability during locomotion. Thus, the complexity of achieving the $P2$ behavior can be seen as being at a higher level of the complexity hierarchy than the $P1$ fitness function in the eyes of the controller.

6 Conclusion and Future Work

We have shown how EMO can be applied for studying the behavioral and morphological complexities of artificially evolved embodied creatures. The morphological complexity of a quadruped creature was found to be lower than the morphological complexity of a hexapod creature as seen from the perspective of an evolving locomotion controller. At the same time, the quadruped was found to be more complex than the hexapod in terms of behavioral complexity. For future work, we intend to provide an empirical proof of measuring not only behavioral complexity but also environmental complexity by evolving controllers for artificial creatures in varied environments. We also plan to apply these measures for characterizing the complexities of artificial creatures evolved through co-evolution of both morphology and mind.

References

1. Hussein A. Abbass. The self-adaptive pareto differential evolution algorithm. In *Proceedings of the 2002 Congress on Evolutionary Computation (CEC2002)*, volume 1, pages 831–836. IEEE Press, Piscataway, NJ, 2002.
2. Josh C. Bongard. Evolving modular genetic regulatory networks. In *Proceedings of the 2002 Congress on Evolutionary Computation (CEC2002)*, pages 1872–1877. IEEE Press, Piscataway, NJ, 2002.
3. Rodney A. Brooks. Intelligence without reason. In L. Steels and R. Brooks (Eds), *The Artificial Life Route to Artificial Intelligence: Building Embodied, Situated Agents*, pages 25–81. Lawrence Erlbaum Assoc. Publishers, Hillsdale, NJ, 1995.
4. Critical Mass Labs. Vortex [online]. http://www.cm-labs.com [cited – 25/1/2002].
5. Kalyanmoy Deb. *Multi-objective Optimization using Evolutionary Algorithms.* John Wiley & Sons, Chicester, UK, 2001.
6. Claus Emmeche. *Garden in the Machine.* Princeton University Press, Princeton, NJ, 1994.
7. David P. Feldman and James P. Crutchfield. Measures of statistical complexity: Why? *Physics Letters A*, 238:244–252, 1998.
8. Dario Floreano and Joseba Urzelai. Evolutionary robotics: The next generation. In T. Gomi, editor, *Proceedings of Evolutionary Robotics III*, pages 231–266. AAI Books, Ontario, 2000.
9. Simon Haykin. *Neural networks – a comprehensive foundation.* Prentice Hall, USA, 2 edition, 1999.
10. Gregory S. Hornby and Jordan B. Pollack. Body-brain coevolution using L-systems as a generative encoding. In L. Spector et al. (Eds), *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 868–875. Morgan Kaufmann, San Francisco, 2001.
11. Andrei N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1:1–7, 1965.
12. Hod Lipson and Jordan B. Pollack. Automatic design and manufacture of robotic lifeforms. *Nature*, 406:974–978, 2000.
13. Henrik H. Lund and John Hallam. Evolving sufficient robot controllers. In *Proceedings of the 4th IEEE International Conference on Evolutionary Computation*, pages 495–499. IEEE Press, Piscataway, NJ, 1997.
14. Rolf Pfeifer and Christian Scheier. *Understanding Intelligence.* MIT Press, Cambridge, MA, 1999.
15. Jordan B. Pollack, Hod Lipson, Sevan G. Ficici, Pablo Funes, and Gregory S. Hornby. Evolutionary techniques in physical robotics. In Peter J. Bentley and David W. Corne (Eds), *Creative Evolutionary Systems*, chapter 21, pages 511–523. Morgan Kaufmann Publishers, San Francisco, 2002.
16. Cosma R. Shalizi. *Causal Architecture, Complexity and Self-Organization in Time Series and Cellular Automata.* Unpublished PhD thesis, University of Wisconsin at Madison, Wisconsin, 2001.
17. Claude E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423, 1948.
18. Karl Sims. Evolving 3D morphology and behavior by competition. In R. Brooks and P. Maes (Eds), *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pages 28–39. MIT Press, Cambridge, MA, 1994.
19. Russell K. Standish. On complexity and emergence [online]. *Complexity International*, 9, 2001.