

Complexity from a Simple Bag of Tricks

Martin C. Martin

Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

<http://www.ai.mit.edu>



The Problem: This work explores the idea that complex, adaptive behavior can arise through the judicious use of a few simple mechanisms. Evolution in a rich environment, a childhood, learning, and an open ended fitness function are each powerful techniques. Is the combination of them greater than the sum of the parts? This work is aimed at finding out.

Motivation: While robotics has made great advances, even the most advanced robots lack the flexibility and adaptability of living things. Previous work in evolving the bodies and brains of robots has demonstrated exciting results [4], but has so far only attempted simple behaviors such as walking or jumping.

Living things accomplish much more. And while their bodies and brains are much more complicated than current robots, their set of genes is much less complicated than the final result. Nature appears to achieve much of this using a few simple principles. An environment with many niches and an open-ended notion of “success” facilitates the exploration of a wide array of solutions. Learning has long been recognized as an important way for beings to gain competency in the environment. A development stage can greatly reduce the representation needed in a genome by allowing the commonality between elements to be expressed simply.

Previous Work: Karl Simms performed early captivating work in this area [5], in which simple locomotion, point following, and competing to cover a block were evolved. More recently, a group at Brandeis has evolved similar creatures that are then automatically built on a shape deposition machine [4].

There has been some work using development in the evolution of bodies and brains [3], as well as the use of learning to adapt a neural network in an evolved creature [1]. Others have evolved neural networks with properties similar to those of real neurons and a development stage inspired by recent advances in the genetics of development [2].

Approach: The evolved bodies will be constrained to use parts typical of machines, such as rigid cylinders, metallic plates and electric motors. Existing rigid body simulators are well suited to this task. The world will be much richer than existing work, containing areas of water, land and air, as well as varied terrain in each area. Later other variations may be added, such as day/night cycles and tides.

Previous work has largely focused on neural networks as the representation for brains, but an alternate representation could lead to behaviors of a much greater complexity. Reinforcement learning will be used to allow the creatures to adapt to their environment, with the details of the learning framework under genetic control. This will allow the complexity of the robot to mirror the complexity of the world, rather than forcing that complexity to be present in the genome. It will also allow the brain to better adapt to changes in the body due to mutation. This will allow more mutations to be explored, ideally allow evolution to be more efficient. Finally, learning changes the topology of the search space by making competent creatures easier to express. For example, it's often easier to define a good state (e.g. healthy and not hungry) than to list all good actions (e.g. all actions which increase health and/or reduce hunger).

A development stage similarly allows a compact description of a complex organism, in part by expressing common elements only once. If the developmental process can interact with the environment, many of the advantages of learning can apply. In combination, these principles – rich environment, learning and development – may each help the other to provide far more than the sum of the parts.

Impact: A success in this work could provide a new way to design machines of a greater complexity than is possible at the moment. Successful creatures could be reverse engineered to determine how they work. This could

lead to insights into the proper method of combining development and learning, possibly providing new paradigms for traditional, hand designed machines. The work could also point the way to giving important properties of living systems to machines, and shed light on the nature of those qualities.

References:

- [1] David Ackley and Michael Littman. Interactions between learning and evolution. *Artificial Life II*, pages 487–509, 1992.
- [2] J. C. Astor and C. Adami. A developmental model for the evolution of artificial neural networks. *Artificial Life*, 6(3):189–218, 2000.
- [3] Josh C. Bongard and Rolf Pfeifer. Repeated structure and dissociation of genotypic and phenotypic complexity in artificial ontogeny. *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 829–836, 2001.
- [4] H. Lipson and J.B. Pollack. *Nature*, pages 974–978, 2000.
- [5] Karl Simms. Evolving 3d morphology and behavior by competition. *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pages 28–39, 1994.