Ontogenetic Development of Skills, Strategies and Goals for Autonomously Behaving Systems

Derek Harter and Robert Kozma
Department of Mathematical Sciences, University of Memphis
Memphis, TN 38152, USA

ABSTRACT

Biological organisms display an amazing ability during their ontogenetic development to adaptively develop solutions to the various problems of survival that their environments present to them. Dynamical and embodied models of cognition are beginning to offer new insights into how the numerous, heterogeneous elements of neural structures may self-organize during the development of the organism in order to effectively form adaptive categories and increasingly sophisticated skills, strategies and goals. The ontogenetic development of behavior in biological organisms represents a significant level of improvement over current approaches to machine learning.

In this paper we discuss the possibility of building action selection mechanisms for autonomous agents based upon new insights into how exactly biological organisms manage to self-organize patterns of behavior during their ontogenetic development. We present a simple task environment that, nevertheless, affords opportunities for the hierarchical development of increasingly complex behaviors in humans. We present some results of standard machine learning mechanisms on performing this task. And finally we discuss future plans for developing models of the ontogenetic development of behavior for autonomous agents in the task environment.

Keywords: Adaptive Intelligent Systems, Ontogenetic Development, Dynamical Systems

INTRODUCTION

The ability of biological organisms to self-organize patterns of behavior during development is an amazing feat, one unsurpassed by current models of learning for constructed artifacts. How biological organisms manage to self-organize hierarchical patterns of behavior simply through interaction with and experience in their task domains is unknown. However new insights into the processes of self-organization in dynamical systems [1, 2, 3, 4, 5, 9, 10, 14, 15, 16, 18] is beginning to offer possible clues as to how this may be accomplished.

The development of behavior in biological organisms is primarily a self-organizing phenomenon. Organisms are born with a basic repertoire of motor skills and instinctive needs. These are often tied to simple action-loops [1], which provide a basic repertoire of simple pattern completion and instinctive behaviors that can begin to satisfy the intrinsic drives of the organism. As the organism develops both physically and behaviorally, however, these instinctive behavior patterns begin to be associated with more general sensory stimuli. The organism learns to recognize patterns in the environment that are important and useful affordances for beneficial behaviors [8]. Increasingly complex patterns of behavior are organized around the solutions that are discovered at earlier stages of development.

Thelen and Smith [18, 17] view development as a shifting ontogenetic landscape of attractor basins. As physical and behavioral patterns develop the landscape is continually reformed and reshaped. Each developed behavior opens up many possibilities for new more complex patterns of behavior, while closing off possibilities for others. Even relatively simple tasks can provide opportunities for the development of increasingly complex strategies in order to improve performance. For example in the simple task we present in the next section, humans develop higher level strategies for improving their performance.

Many theories of the development of behavior in biological organisms are beginning to view it in terms of a self-organizing dynamical system [18, 10, 9]. The organization of patterns of behavior is viewed, in some sense, as the formation and evolution of attractor landscapes. Some research [16, 5, 4, 6, 7, 12] also indicates that chaotic dynamics may play an essential part in the formation of perception and behavior in biological organisms. For example, one role that chaos may play in the development of behavior is as a type of variability upon which selection may operate in the successive generation of increasingly improved performances.

TASK ENVIRONMENT

The hierarchical self-organization of behavior in biological organisms occurs naturally in almost any task that the organisms wishes to perform. Humans are capable of displaying this type of development when undertaking any unfamiliar and new task. For example, in the task domain described here, which is a simplified form of the Tetris game environment [11], humans learn the necessary skills to become proficient in the task simply by performing the task. Such learning displays a typical progression in ability. Initial learning is rapid as the human develops

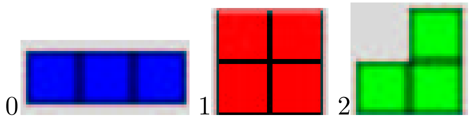

Figure 1: The shapes used in the packing task.

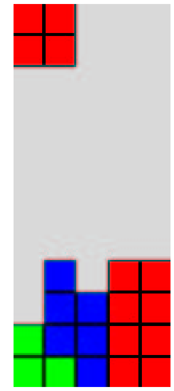

Figure 2: An example packing task trial. Shapes enter from the top and must be positioned and rotated before they are dropped. Performance is evaluated by the height and the density of the packing of the shapes.

basic perceptual and motor skills necessary for performing the task. Once such basic skills are in place, further abilities develop based upon earlier competencies. The progression follows the typical curve of diminishing returns for learning, where initial progress is quite rapid, followed by slower, more incremental improvements.

In the packing task, humans are given a series of 10 shapes of differing configurations. The human player is allowed to position and orient the shape over a playing field before dropping it. The shape falls onto the playing field once dropped. The goal of the player in this task is to maximize the density of the packing achieved. This is made difficult because of various reasons. Shapes are presented to the player at random and the player does not know in advance which shapes he will be asked to pack. Therefore optimal packings are not possible since the player has no knowledge of how the future environment may affect or change current behavior.

The packing task can also be made more difficult with the addition of further constraints. For example logical constraints may be added on the total number of moves and rotations that the player is allowed to make. In this case the player is constantly forced to make constraint satisfaction decisions on the tradeoff of immediate optimizations versus long range conservation of resources. Another interesting variation is the addition of real-time constraints on the performance of the player. In such a scenario the player given a limited amount of time in which to position and orient a shape before it is automatically dropped (typically 1 sec). Real-time constraints add interesting dynamics to the development of behavior for a task domain. Biological organisms have evolved to develop behavior in real-time demanding task environments. The natural mode of operation for such developmental mechanisms is in real-time environments. Real-time constraints may not simply be interesting constraints on tasks, but may in some sense turn out to be necessary to the developmental mechanisms of biological organisms [9, 18].

EXPERIMENTS

In the packing task, the behaving system is presented with a series of shapes, one shape at a time. In this packing task, which is a simplified form of the Tetris game [11], the system can be presented with one of 3 shapes as shown in figure 1. The goal of the task is to move and rotate a shape before allowing it to drop onto a playing field in such a way as to end up with as compact of a packing as possible. An example of a packing trial in progress can be found in figure 2.

The behaving system does not know in advance what sequence of shapes it will be given. In our version of the packing task, the system is given random sequences of 10 shapes. The performance of the system on the packing task is evaluated by examining the density of their packing and by examining the total height of the resulting packing.

The system can produce two types of behavior. It must specify where to position (or move) the shape in the playing field, and how to rotate the shape. Once the system has specified the position and rotation of the shape, it is allowed to fall down onto the playing field. The shape settles into place and the next shape is presented to the behaving system to be positioned and rotated.

Encoding

We now present an example of a standard neural network that learns to perform the simple packing task. The neural network needs to be given some sense of the current state of the environment. For the experiments performed here, two pieces of input were given to the network: the type of shape that has appeared, and a perception of the contours of the current playing field.

The encoding of the type of shape is relatively simple. In the packing task environment there are 3 different shape types. We used 2 bits to encode the type of shape. The shapes in figure 1 were given numbers (from left to right) of 0, 1 and 2 and were encoded as 00, 01 and 10 respectively.

The perception of the state of the playing field (the environment) is necessary in order to produce good behavior on where and how to position the shape before dropping it. In our reduced packing task, the playing field consisted of 5 columns. We sensed the height of each column currently in the environment, and encoded this for training and testing the networks. The lowest point in the playing field is used as a baseline and is encoded as having a height of 0. All other heights are calculated from the baseline depth. We used 2 bits to encode the height of each column, and simply ignored perception of columns that were greater than 3 units above the baseline.

For example in figure 2, the leftmost column has the lowest depth in the playing field, and would be encoded with height 0. The next column to the right has a height 2 units above the baseline. So from left to right, the height of the columns in figure 2 would be encoded as 0, 2, 1, 2, 2. The type of shape shown in figure 2 about to be dropped is shape number 1. As stated before we used 2 bits to encode the shape type, and 2 bits for each of the column heights, for a total of 12 bits of input. The situation shown in figure 2 would be encoded as:

```
Type Col1 Col2 Col3 Col4 Col5
0 1  0 0  1 0  0 1  1 0  1 0
```

For the output of the system we developed the following encoding. We encoded the position to place the shape from the left edge in 3 bits. We need to be able to specify up to 5 units of displacement, thus we needed 3 bits to encode the 5 possibilities. The shapes can be rotated in increments of 90 degrees. Shape 2 (the L shape) can be rotated into 4 different distinct orientations. Therefore we also needed 2 bits to encode all possible specifications of rotation.

Training

We trained standard backpropogation networks using the encoding described above. For training data we had a human perform 50 packing trials, and we captured and encoded the input and the output of the behavior that the human produced when performing the packing task. We also captured a similar set of data for testing. We trained and tested the networks with many different configurations of number of hidden nodes and epochs trained. We then chose the best configurations in order to evaluate the performance of the networks on the packing task as discussed in the next section. The neural network performed best with 50 hidden nodes.

Evaluation

We then used our packing task testbed in order to evaluate the performance of the networks on simulated packing trials. We gave the networks 100 random trials and measured their performance by calculating the packing density and height that they achieved. Packing height is simply a measure of the highest column of blocks in the playing field. Packing density is measured by looking at the ratio of the number of filled spaces in the packing to the total area of the packing. In figure 2 the packing has a height of 4 and a density of 17 / 20 or 0.85. The lower the height of the packing is the better the performance and similarly the denser the packing is the better the performance.

Results

A human learned the packing task and was asked to perform the task for 100 trials. Similarly the resulting neural networks were run on 100 trials of the packing task. Table 1 shows a comparison of the average performance on the 100 trials of the neural network and the human.

Table 1: Comparison of average height and density performance measures on 100 simulated packing tasks

	Height	Density
Human	7.62	0.8748
Neural Network	8.18	0.8261

Basic neural networks perform adequately on the packing task, but obviously are not quite as good at packing as humans, even for this simplified task domain. Humans, when performing this task, alter their strategies as the task progresses. Early on in a packing trial, a human is willing to leave open opportunities for particular shapes. People know intuitively that, even though they see shape types at random, they are likely to see the particular shape type needed if it is still early in the trial. However, as the trial progresses, strategies shift to those that will simply minimize the height of their packing.

This shift in strategies causes confusion for simple backpropogation networks. They see this as conflicting output patterns for the same input. Strategies that would possibly correct this deficiency for basic neural networks and other solutions will be discussed in the next section.

DISCUSSION

The basic neural networks presented here are not quite capable of human level performance in the packing task. The primary reason for this deficiency is an inability to perceive the changes in circumstances that cause a shift in the behavior of the human trainers. We have no doubt that adding on more contextual input (such as the current height of the packing, or a count of the number of shapes packed so far) would improve the performance of the basic network, though it remains to be seen if it could equal human performance. Also, other methods such as recurrent, dynamical neural networks, or genetic algorithm optimizations, should be capable of bringing standard methods of machine learning up to human level performance on this simple task.

The point is not to equal human performance in this simplified domain, but to begin to create models that can develop behavior on their own in a cognitively plausible manner, and that display some of the flexibility of biological development. Most standard methods of machine learning should be able to competently handle the packing task environment in its simplified form but inevitably will break down as we add complexity and real time constraints to the task.

KIII is a dynamical memory device, which has been used successfully to solve difficult classification problems in vague, and noisy environments [12]. The KIII model incorporates several KII sets, which can be interpreted as units generating limit cycle oscillations in an autonomous regime. High-dimensional aperiodic and chaotic behavior does not emerge until the complete KIII system is formed.

KIII has a multi-layer architecture with excitatory and inhibitory lateral, feed-forward, and feedback connections. KIII models can grasp the essence of the observed dynamic behavior in certain biological neural networks. It seems feasible to build a simplified version of KIII for the action selection task addressed in this work. We call it 3*KII model, as it consists of 3 mutually interconnected KII sets. Each KII set has a well-defined oscillation frequency. The complete 3*KII model, however, may exhibit high-dimensional, aperiodic oscillations as the result of competing, incommensurate frequencies of the KII components.

The advantage of 3*KII is that it allows a self-organized encoding of behavioral patterns into localized wings of a high-dimensional attractor. Therefore, we can obtain flexible and noise-resistant transitions among the states of the system, self-organized into a sequence of elementary actions of phase transitions. It is expected that by defining a more challenging packing task with a larger number and more complicated set of block patterns and also a larger playing field; the application of a dynamical encoding and action selection mechanism such as 3*KII would prove to be beneficial. Also the emergence of self-organized action patterns would be imminent and complex behavioral patterns could be generated and studied.

CONCLUSION

The development of behavior, even in a simplified environment such as the packing task, can shed light on the mechanisms of biological development and learning. Biological organisms are able to effectively develop increasingly complex skills and strategies simply by interacting with and solving problems in their environment. The dynamic, self-organization of behavior in biological organisms is a powerful model of learning that, if better understood, would provide great opportunities for improved artificial behaving and learning systems. Development of behavior in biological organisms can be viewed as a self-organizing dynamical system. Some research also indicates the importance of chaotic modes of organization in the development of behavior.

References

[1] Andy Clark. *Being There: Putting Brain, Body, and World Together Again.* The MIT Press, Cambridge, MA, 1997.

[2] Gerald M. Edelman. *Neural Darwinism: The Theory of Neuronal Group Selection.* Basic Books, New York, NY, 1987.

[3] Gerald M. Edelman and Giulio Tononi. *A Universe of Consciousness: How Matter Becomes Imagination.* Basic Books, New York, NY, 2000.

[4] Walter J. Freeman. Consciousness, intentionality and causality. In Núñez and Freeman [13], pages 143–172.

[5] Walter J. Freeman. *How Brains Make Up Their Minds.* Weidenfeld & Nicolson, London, 1999.

[6] Walter J. Freeman and Robert Kozma. Local-global interactions and the role of mesoscopic (intermediate-range) elements in brain dynamics. *Behavioral and Brain Sciences*, 23(3):401, 2000.

[7] Walter J. Freeman, Robert Kozma, and Paul J. Werbos. Biocomplexity: Adaptive behavior in complex stochastic dynamical systems. *BioSystems*, 2000.

[8] J. J. Gibson. *The Ecological Approach to Visual Perception.* Houghton Mifflin, 1979.

[9] Horst Hendriks-Jansen. *Catching Ourselves in the Act: Situated Activity, Interactive Emergence, Evolution and Human Thought.* The MIT Press, Cambridge, MA, 1996.

[10] J. A. Scott Kelso. *Dynamic Patterns: The Self-organization of Brain and Behavior.* The MIT Press, Cambridge, MA, 1995.

[11] David Kirsh and Paul Maglio. Reaction and reflection in tetris. In J. Hendler, editor, *Artificial Intelligence Planning Systems: Proceedings of the First Annual International Conference (AIPS92)*, San Mateo, CA, 1992. Morgan Kaufman.

[12] Robert Kozma and Walter J. Freeman. Chaotic resonance - methods and applications for robust classification of noisy and variable patterns. *International Journal of Bifurcation and Chaos*, 11(6), 2001.

[13] Rafael Núñez and Walter J. Freeman, editors. *Reclaiming Cognition: The Primacy of Action, Intention and Emotion.* Imprint Academic, Bowling Green, OH, 1999.

[14] Rolf Pfeifer and C. Scheier. *Understanding Intelligence.* The MIT Press, Cambridge, MA, 1998.

[15] Robert F. Port and Timothy van Gelder, editors. *Mind as Motion: Explorations in the Dynamics of Cognition.* The MIT Press, Cambridge, MA, 1995.

[16] Christine A. Skarda and Walter J. Freeman. How brains make chaos in order to make sense of the world. *Behavioral and Brain Sciences*, 10:161–195, 1987.

[17] Esther Thelen. Time-scale dynamics and the development of an embodied cognition. In Port and van Gelder [15], chapter 3, pages 69–100.

[18] Esther Thelen and Linda B. Smith. *A Dynamic Systems Approach to the Development of Cognition and Action.* The MIT Press, Cambridge, MA, 1994.