Toward the Evolution of Dynamical Neural Networks for Minimally Cognitive Behavior

Randall D. Beer

Santa Fe Institute
1399 Hyde Park Rd.
Santa Fe, NM 87501
beer@santafe.edu

Dept. of Computer Engineering and Science
and Dept. of Biology
Case Western Reserve University
Cleveland, OH 44106
beer@alpha.ces.cwru.edu

Abstract

Current debates regarding the possible cognitive implications of ideas from dynamical systems theory and adaptive behavior research would benefit greatly from a careful study of simple model agents that exhibit minimally cognitive behavior. This paper describes one such agent, and presents the results of preliminary experiments on the evolution of dynamical neural networks for visually-guided orientation, object discrimination and accurate pointing with a simple manipulator to objects appearing in its field of view.

1 Introduction

Many of the key ideas emphasized in adaptive behavior research are beginning to have a significant impact on cognitive science. For example, adaptive behavior research in general, and the dynamical perspective on adaptive behavior that is often taken in such research in particular, have begun to significantly influence the growing debates concerning the nature and necessity of notions of representation and computation in explaining cognitive behavior (Brooks, 1991; Clark & Toribio, 1994; Beer, 1995b; Port & van Gelder, 1995). Likewise, the important roles played by an agent's body and its environment in the generation of its behavior, long emphasized in adaptive behavior research, parallel a renewed concern for embodiment and situatedness in cognitive science (Suchman, 1987; Lakoff, 1987; Damasio, 1994; Haugeland, 1995; Hutchins, 1995). The sorts of decentralized and distributed control mechanisms that are utilized in adaptive behavior research are also strongly reminiscent of the picture of human brain activity that is emerging from neuroanatomical and brain imaging studies in cognitive neuroscience (Posner & Raichle, 1994; Gazzaniga, 1995), as well as from more detailed studies of invertebrate nervous systems (Altman & Kien, 1989). Finally, dynamical and adaptive behavior ideas are beginning to significantly impact work in developmental psychology (Rutkowska, 1994; Thelan & Smith, 1994).

Despite this widespread impact on cognitive science, most of the empirical work in adaptive behavior research to date has focused on relatively simple sensorimotor behavior, such as obstacle avoidance and wall following. While there are clear advantages to this strategy, and much is still not understood about the design and analysis of even these simple behaviors, it is unfortunate that so much of the discussion concerning the cognitive implications of adaptive behavior ideas is being carried out in the absence of concrete models. For example, current debate on the role of representation in cognition is mostly occurring at a philosophical level, with intuition and analogy rather than the careful study of concrete models leading the way. But one of the major advantages of the adaptive behavior approach is that, by grounding an agent's behavior in an environment and making far fewer a priori assumptions about the necessary design of its internal control mechanisms, this approach provides a much broader intellectual playing field on which to explore these issues through the development and analysis of concrete models.

My own recent work in this area has focused on the use of evolutionary algorithms to evolve dynamical neural networks for controlling the behavior of model agents and then analyzing the dynamics of the resulting agent-environment systems. Here, an agent's nervous system is viewed not as an information processing system, but rather as a dynamical system which, in conjunction with the dynamics of the agent's body and its environment, is capable of producing effective behavior in that environment. To date, this approach has been successfully applied to chemotaxis, legged locomotion, sequential behavior and learning (Beer & Gallagher, 1992;

Yamauchi & Beer, 1994). The principal motivation behind this work has been the development and analysis of simpler idealized models of adaptive behavior for the purpose of elucidating the essential principles of a dynamical theory of adaptive behavior. To this end, dynamical analyses have been performed on many of these evolved circuits, and a preliminary theoretical framework for adaptive behavior has been proposed (Beer, 1995a; Beer, 1995b).

The goal of the work described in this paper is to begin to explore the applicability of this approach to the design and analysis of more cognitive behavior. Section 2 describes a visually-guided agent whose capabilities are both rich enough to begin to explore cognitive behavior yet simple enough to be tractable to evolution and analysis. The particular neural network model and evolutionary algorithm used here are described in Section 3. Sections 4-6 describe the results of preliminary experiments in the evolution of dynamical neural networks for visually-guided orientation, object discrimination and accurate pointing, respectively. Finally, related work is briefly discussed in Section 7 and Section 8 summarizes the results of this paper and suggests some directions for future research.

2 A Visually-Guided Agent

What sorts of agents and behaviors should we attempt to study? On the one hand, the capabilities and behavior of the model agents that we study must be rich enough to be cognitively interesting, so that they raise the sorts of issues that we would like to explore. For example, if we wish to explore the nature and necessity of the notion of representation in cognitive behavior, then we must examine tasks that are sufficiently "representation-hungry" (Clark & Toribio, 1994). On the other hand, these model agents must be simple enough to be computationally and analytically tractable, so that we have some hope of evolving and analyzing them using techniques that are at most an incremental step beyond what is currently known to be feasible.

Generally speaking, visually-guided behavior provides an excellent arena in which to explore the cognitive implications of dynamical and adaptive behavior ideas. Visually-guided behavior includes such phenomena as visual orientation, object perception and discrimination, visual attention, perception of self-motion, object-oriented action, and visually-guided motion and manipulation. However, despite this complexity and richness, significant progress on understanding the processes and neural architectures underlying visually-guided behavior is beginning to be made in cognitive neuroscience (Posner & Raichle, 1994; Gazzaniga, 1995). Furthermore, a relatively simple model agent can be designed that supports simplified versions of all of these phenomena.

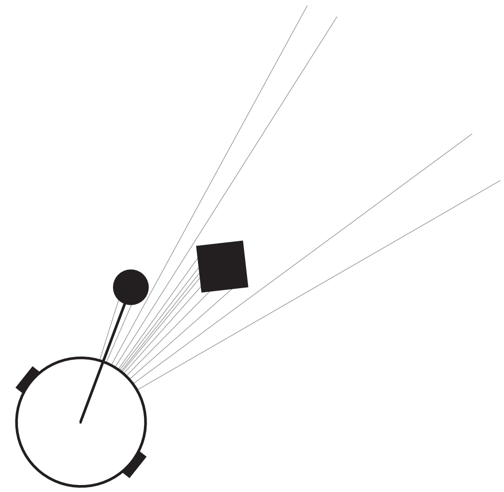

Figure 1: Basic design of a visually-guided agent.

The model agent is illustrated in Figure 1. This two-dimensional agent possesses an "eye" consisting of a foveated array of distance sensors, two "motors" that produce 2D movement of the entire body, and a simple transparent 2 degree-of-freedom "arm" and opaque 1 degree-of-freedom "hand" for manipulating objects. Note that the intent here is not to model in any depth the particular visually-guided behavior of any real animal or robot. Rather, the goal is to explore the space of possible dynamical organizations of agents that engage in minimally cognitive behavior. Thus, there is no particular need to strive for physical realism in these experiments. For example, the agent's "vision" is certainly not intended as a serious model of the actual physics of light or photoreception. However, it does raise similar issues in the perception of objects using a spatially-structured array of distal sensors. Likewise, while the agent's hand does not realistically model the limbs of any animal or robot, it does raise some of the same issues in the visual control of manipulation.

Given its sensory and motor capabilities, what sorts of cognitively interesting behavior might this agent engage in? This agent could perceive the two-dimensional structure of objects and organize its behavior in accordance with their shapes (e.g., orienting to a novel object and discriminating one object from another). It could navigate around obstacles in two dimensions, deciding which gaps its body can fit through and which it cannot. It could also exhibit simple forms of object persistence (e.g., continuing to pursue a goal object that is momentarily occluded by another object). At a more sophisticated level, selectively interacting with one object from among a set of objects raises interesting focus-of-attention issues. In addition, the fact that its hand is opaque raises interesting issues in the discrimination of self from nonself. This agent could also actively manipulate objects in its environment (e.g., building simple structures out of the objects in its environment). Finally, one could imagine groups of these agents engaging in simple cooperative tasks, such as tossing an object back and forth. Thus, the behavior of this agent is potentially of some cognitive interest.

The remainder of this paper describes three sets of experiments aimed at an initial exploration of some of the simplest capabilities of the agent sketched above. Specifically, these experiments are designed to establish the basic soundness and feasibility of the proposed agent, to explore how difficult it is to evolve dynamical neural network controllers for these tasks, and to determine the neural architectures and evolutionary algorithm configurations that are best suited to them.

3 Methods

In all of the following experiments, the agent has a circular body with a diameter of 30 (in an environment of size 400 × 275), with an eye consisting of either 5 or 7 rays of maximum length 220 uniformly distributed over a visual angle of $\pi/6$. An intersection between a ray and an object causes an input to be injected into the corresponding sensory neuron. The magnitude of the injected input is inversely proportional to the distance to the object. When rays are at their maximum length, no input is injected, while a maximum input of 10 is injected for rays of zero length.

The spatial resolution of the agent's eye is determined by a number of factors. Resolution obviously depends on the number of rays and the visual angle over which they are distributed. Resolution also clearly depends on how far away an object is, since a more distant object will intersect fewer rays. Finally, the spatial resolution of the eye is very dependent on the values of the bias and gain parameters of the ray sensory neurons. If the biases are too high or too low, then objects will give either a saturated response or no response, respectively. If the gains are too low, each ray will show very little difference in response regardless of its length. If the gain is too high, each ray will essentially give a binary response at a very narrow range of distances (which will make the ray biases very difficult to evolve). While these issues will be discussed in greater detail in the pointing experiments described in Section 6, they are important to keep in mind for all of the experiments described in this paper.

The agent's behavior is controlled by a continuous-time recurrent neural network with the following state equation:

$$\tau_i \dot{y}_i = -y_i + \sum_{j=1}^{N} w_{ji}\sigma\left(g_j(y_j + \theta_j)\right) + I_i \quad i = 1,\ldots,N$$

where y is the state of each neuron, τ is its time constant, w_{ji} is the strength of the connection from the j^{th} to the i^{th} neuron, g is a gain, θ is a bias term, $\sigma(x) = 1/(1+e^{-x})$ is the standard logistic activation function, and I represents an external input (e.g., from a sensor). States were initialized to 0 and circuits were integrated using the forward Euler method with an integration step size of 0.1.

The evolutionary algorithm used in the experiments described in this paper is similar to a very simple evolutionary strategy (Bäck & Schwefel, 1993). A population of individuals is maintained, with each individual encoded as a vector of real numbers (representing the connection weights w_{ji}, the biases θ_i and the time constants τ_i, as well as the gains g_i of the ray sensory neurons, with all other gains fixed to 1). Initially, a random population of vectors is generated by initializing each component of every individual to random values uniformly distributed over the range ±1. Individuals are selected for reproduction using fitness proportional selection with linear fitness scaling with a fitness scaling multiple of 2 (Goldberg, 1989). A selected parent is mutated by adding to it a random displacement vector whose direction is uniformly

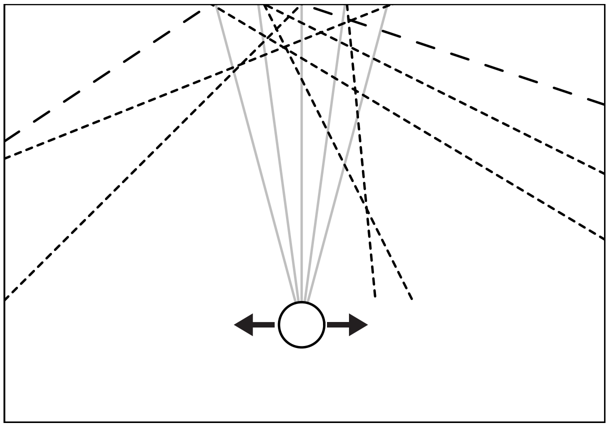

Figure 2: Experimental setup for orientation experiments. The agent moves horizontally. Its rays are shown in gray. Dotted and dashed lines denote the paths of circular objects used to evaluate the agent's performance during evolution, as described in the text.

distributed on the M-dimensional hypersphere (Knuth, 1981, p. 130) and whose magnitude is a Gaussian random variable with 0 mean and variance σ^2. For each slot in the new population, the child is chosen if its performance is greater than or equal to that of the parent, otherwise the parent is copied. All random numbers were generated using the routine `ran1` described in (Press et al, 1994, p. 280), which has a period greater than 10^8 and uses a shuffling algorithm to remove low-order serial correlations.

In the experiments described in this paper, search vector components were mapped to circuit parameters using linear maps from ±1 to a given range of circuit parameter values. Unless otherwise stated, these circuit parameter ranges were as follows: circuit biases $\in$ [-5,5], time constants $\in$ [1,2], and connection weights $\in$ [-5,5]. Because ray sensor gain and bias ranges varied across experiments, they will be reported separately below. Gains were clipped to be greater than 0 and time constants were clipped to be greater than 1.

4 Orientation Experiments

One of the most basic capabilities required by any visually-guided agent is the ability to orient to a visual stimulus. In the first set of experiments to be described, agents were evolved that could use their vision to adjust their horizontal position so as to catch falling objects (Figure 2).

These agents had 5 rays. Their horizontal velocity was proportional to the sum of the opposing horizontal forces produced by each motor (with a constant of proportionality of 0.2). Circular objects with a diameter of 26 were dropped from the top of the environment with an initial horizontal offset from the center of the agent in the range ±70, a horizontal velocity in the range ±6, and a vertical velocity in the range [0.5,5].

The performance measure to be maximized was:

$$200 \quad - \quad \sum_{i=1}^{NumTrials} d_i \Big/ NumTrials$$

where d_i is the horizontal distance between the centers of the object and the agent when their leading edges meet on the i^{th} trial, and *NumTrials* is the total number of trials.

In the first set of orientation experiments, bilaterally symmetric feedforward networks with 5 ray sensory neurons and 2 motor neurons were evolved (for a total of 8 parameters). All time constants were fixed to 1. Ray sensor biases were in the range [-10,-5] and ray sensor gains were in the range [1,5]. Populations of 25 individuals were evolved for 50 generations with a mutation variance σ^2 of 10. Six evaluation trials were used (shown as dotted lines in Figure 2). Note that when objects reach one of the walls, their horizontal motion stops but their vertical motion continues.

This turned out to be a fairly simple task, and agents with a mean fitness of 99.31% (N = 5) quickly evolved (throughout this paper, fitness will be reported as a percentage of the maximum attainable performance). The movement of a typical circuit on one trial is shown in Figure 3a. Note how the agent quickly orients to the object and then tracks it as it falls. Interestingly, these circuits generalized poorly to 100 random trials, with mean fitness dropping to 79.60%. An examination of the qualitative behavior of these agents revealed that they were failing to respond quickly enough to objects with large horizontal and small vertical velocities.

In an attempt to improve this deficiency, two additional evaluation trials were used (shown as dashed lines in Figure 2) and the 5 experiments were repeated. The resulting agents had a mean fitness of 87.29% on the 8 evaluation trials, and a mean fitness of 90.25% on 100 random trials. Closer examination revealed that these new agents were indeed more sensitive to objects with large horizontal and small vertical velocities, primarily because the mean bias of their ray sensors (-0.85 ± 0.18 s.e.) was significantly larger

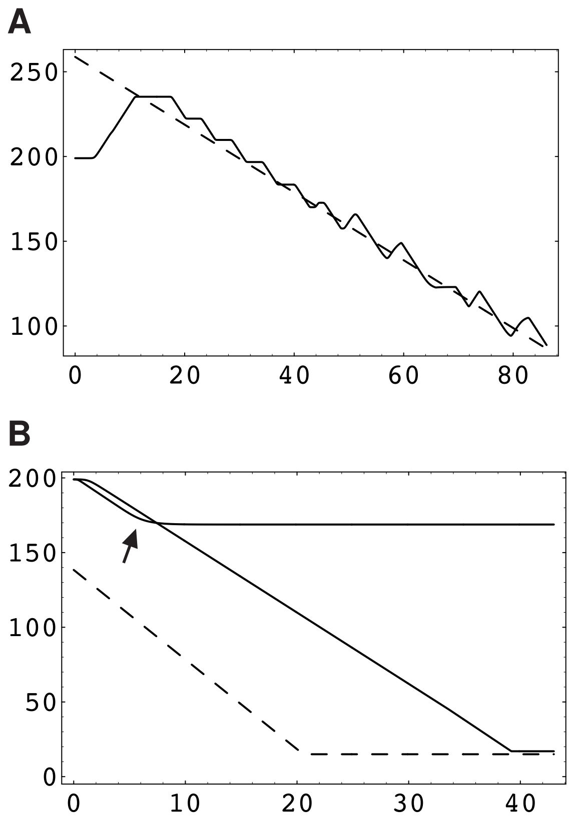

Figure 3: Plots of the horizontal positions over time of evolved orientation agents (solid line) attempting to catch a circular object (dashed line). (a) Path of a typical reactive agent on a successful trial. (b) A comparison of the paths taken by typical reactive and dynamical agents for an object with large horizontal and small vertical velocity. While the reactive agent ceases to move as soon as the object leaves its visual field (arrow), the dynamical agent continues to pursue the object, eventually catching it against the left wall.

than the mean bias of the ray sensors in the first set of agents (-2.06 ± 0.33 s.e.) ($p < 0.05$ using Welch's approximate *t*-test). However, these agents still missed some of these objects because they would quickly pass out of the agent's field of view and these agents would not pursue objects that they could no longer see. An example of this behavior is shown in Figure 3b.

Thus, these agents are faced with a simple example of an object persistence problem. Since both the agent and the objects are constrained to remain within the "walls" of this environment, these agents should continue to pursue objects that have momentarily passed out of their field of view. However, because these agents are controlled by static feedforward networks, they are purely reactive; they cannot organize their behavior according to sensory stimuli that are no longer present. Note that adding any number of interneurons is not going to solve this problem. Rather, what is needed is for the circuits controlling these agents to have internal dynamics.

In order to address this object persistence problem, a final set of orientation experiments was run using a dynamical elaboration of the feedforward circuit used earlier, now with evolvable time constants and bilaterally symmetric self and recurrent connections between the motor neurons (for a total of 12 parameters). All other aspects of these experiments were identical to the second set of experiments. In this case, the resulting agents had a mean fitness of 99.09% on the 8 evaluation trials and a mean fitness of 96.60% on 100 random trials. An examination of the qualitative behavior of these dynamical agents revealed that they would indeed continue to pursue objects that had momentarily disappeared from their field of view (Figure 3b). Similar results were obtained with circuits in which the ray sensory neurons and the motor neurons were fully interconnected. Thus, even in this relatively simple task and circuit, internal dynamics can offer significant advantages to an agent.

5 Discrimination Experiments

In order to selectively interact with different objects, a visually-guided agent must be capable of visually discriminating between them. In a second set of experiments, agents were evolved which could discriminate between circles and diamonds and between circles and horizontal lines, catching circles as in the orientation experiments while avoiding the other objects (Figure 4).

These agents had 7 rays. The experimental setup was similar to that used in the orientation experiments, with agents moving horizontally as objects fall from above. In this case, objects fell straight down with an initial horizontal offset in the range ±50 and a vertical velocity of either 3 or 4. Circular objects had a diameter of 30, diamonds had sides of length 30 and lines had a length of 30.

The performance measure to be maximized was:

$$\sum_{i=1}^{NumTrials} p_i \Big/ NumTrials$$

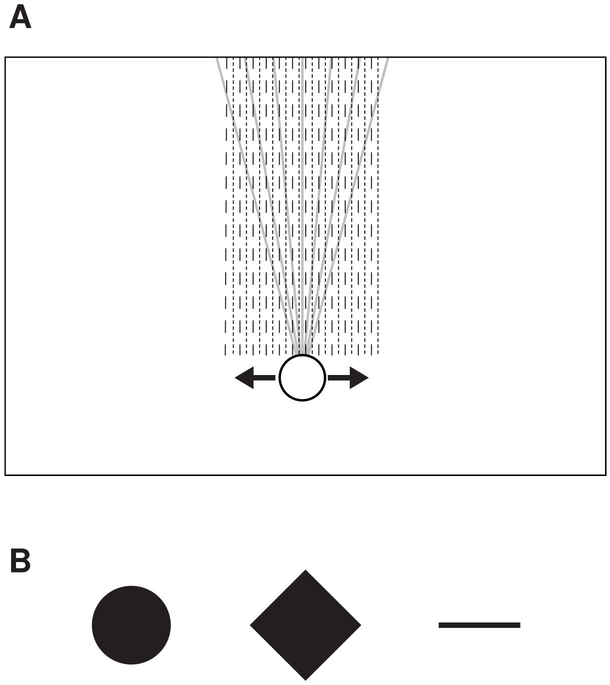

Figure 4: Experimental setup for discrimination experiments. (a) The agent moves horizontally. Its rays are shown in gray. Broken lines denote the paths of circular (dotted) and other (dashed) objects used to evaluate the agent's performance during evolution. (b) Evaluation objects used in the discrimination experiments.

where $p_i = 1 - d_i$ for a circular object and $p_i = d_i$ for the other objects, d_i is the horizontal distance between the centers of the object and the agent when their leading edges meet on the i^{th} trial (clipped to *MaxDistance* and normalized to run between 0 and 1), *NumTrials* is the total number of trials, and *MaxDistance* is 1.5 times the sum of the radii of the object and the agent. The reason that d_i was clipped to *MaxDistance* is to prevent the avoidance of, for example, diamonds by large distances from dominating the fitness at the expense of accuracy in catching circles. A total of 24 evaluation trials were used during evolution, uniformly distributed over the range of horizontal offsets and alternating between circular objects and either diamonds or lines (Figure 4). The large number of trials was necessary in order to ensure good generalization.

The circuit architecture was bilaterally symmetric, with 7 ray sensory neurons projecting to 5 fully interconnected interneurons which in turn projected to the two motor neurons controlling horizontal motion (for a total of 47 parameters). Ray sensor gains were in the range [1,5]. Many different ray sensor bias ranges were used, from [-10,0] to [-4,-2], with the narrower ranges generally giving better results. All other ranges were identical to those used in the orientation experiments.

Agents that could visually discriminate between objects were much more difficult to evolve than agents that could simply orient to an object. In these experiments, populations of from 300 to 400 individuals were evolved for from 100 to 200 generations with a mutation variance σ^2 of 15. The behavior of the best discriminator of circles and diamonds and the best discriminator of circles and lines will be described here.

The best circle/diamond discriminator had a mean fitness of 99.83% on the 24 evaluation trials and a mean fitness of 98.96% on 100 random trials. Qualitatively, all objects were correctly classified on the 24 evaluation trials and only one circle was incorrectly classified during the 100 random trials.

The behavior of this agent is shown in Figure 5, both while catching a circle (Figure 5a) and while avoiding a diamond (Figure 5b) dropped from the same horizontal position. Note that, in both cases, the agent initially foveates the object in the first 20 time steps, actively scans it for approximately the next 20 time steps, and then either centers it in the case of the circle or avoids it in the case of the diamond. This foveate-and-decide strategy was fairly typical of the evolved circle/diamond discriminators, although it was not universal. Even though the rays are uniformly distributed across the visual field, foveating the object still has the advantage of bringing the maximum number of rays to bear on the object and of placing the object in a standard position with respect to the agent.

It is important to emphasize that this agent is not merely centering and then statically pattern-matching an object. Rather, its strategy seems to dynamic, with active scanning apparently playing an essential role. The importance of an agent having control of its own gaze direction has been a major theme in active vision research (Ballard, 1991; Churchland et al, 1994). Given that this task necessarily involves relative motion of the agent and the object (because the agent must express its decision by moving either toward or away from an object as it falls), it is perhaps not too surprising that dynamics plays an essential role in the operation of this controller.

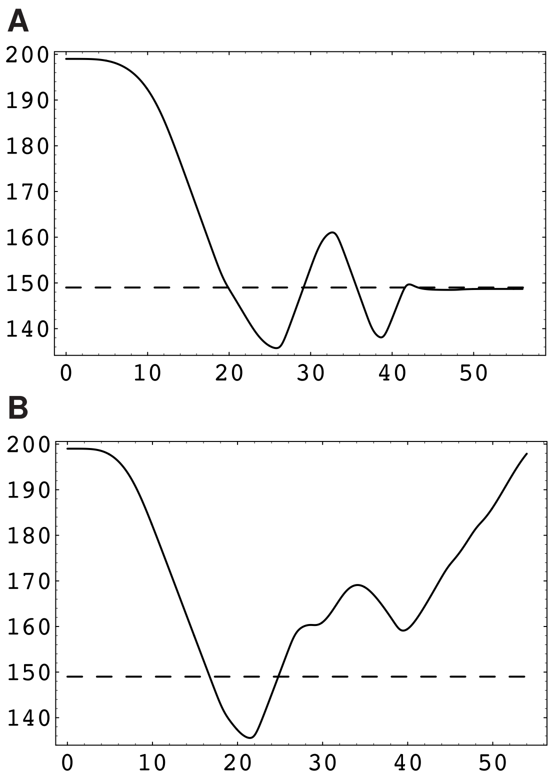

Figure 5: Plots of the horizontal positions over time of an evolved agent that can visually discriminate between circles and diamonds. The path of the agent is shown as a solid line, while that of the object is shown as a dashed line. (a) Path of the agent catching a circle. (b) Path of the agent avoiding a diamond.

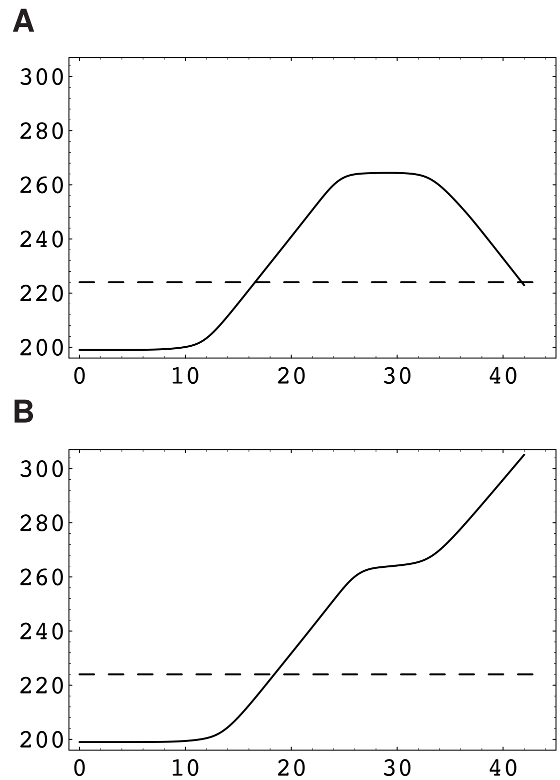

Figure 6: Plots of the horizontal positions over time of an evolved agent that can visually discriminate between circles and lines. The path of the agent is shown as a solid line, while that of the object is shown as a dashed line. (a) Path of the agent catching a circle. (b) Path of the agent avoiding a line.

The best circle/line discriminator had a mean fitness of 99.26% on the 24 evaluation trials and a mean fitness of 97.85% on 100 random trials. Qualitatively, all objects were correctly classified during both the 24 evaluation trials and the 100 random trials.

The behavior of this agent is shown in Figure 6, both while catching a circle (Figure 6a) and while avoiding a line (Figure 6b). The strategy here is rather different than for the circle/diamond discriminator described above. In this case, the agent initially *antifoveates* both the circle and the line, moving so that the object lies near the opposite periphery of its field of view. The agent then pauses and, as the object nears, the agent either centers it in the case of the circle or continues to move away in the case of a line. This antifoveate-and-decide strategy was fairly typical of the other circle/line discriminators that evolved, although it was not universal.

6 Pointing Experiments

Before a visually-guided agent can manipulate objects, it must be capable of coordinating the movement of its manipulator with objects appearing in its visual field. In a final set of experiments, stationary agents with a 1 degree-of-freedom transparent manipulator were evolved to point to the centers of circular objects appearing at arm's length in their visual field (Figure 7).

These agents had 5 rays and were unable to move. Instead, they had a transparent manipulator of length 45 with 1 angular degree of freedom and an angular range of $\pm\pi/4$. The angular velocity of the manipulator was proportional to the sum of two opposing torques (with a constant of proportionality of 0.1). Static circular objects of diameter 26 appeared at arm's length.

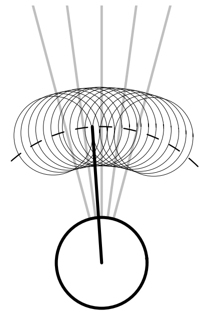

Figure 7: Experimental setup for pointing experiments. The agent is fixed in place and has an arm that swings along the dashed arc. Its rays are shown in gray. Overlapping circles denote the locations of the objects used to evaluate the agent's performance during evolution.

The performance measure to be maximized was the average angular accuracy of pointing, defined as:

$$1 - \frac{\sum_{i=1}^{NumTrials} \left| \theta_i^{object} - \theta_i^{arm} \right|}{MaxEvaluationAngle * NumTrials}$$

where θ_i^{object} is the angle of the object relative to the center of the body and θ_i^{arm} is the angle of the arm at the end of the i^{th} trial, *MaxEvaluationAngle* is the angular width over which objects can appear (π/4 in these experiments), and *NumTrials* is the total number of trials. A total of 15 evaluation trials were used during evolution, each of duration 20 time steps, with the arm centered at the beginning of each trial.

To date, these pointing experiments have employed a bilaterally symmetric feedforward network with 5 ray sensory neurons projecting to 3 interneurons which in turn project to two motor neurons controlling the arm. In addition, the motor neurons received weighted inputs from two arm angle sensors (giving a total of 18 parameters). Arm angle sensors gave a linear response from 0 (arm centered) to 1 (arm at an extreme angle on the side opposite the sensor). Time constants were fixed to 0.75.

Pointing agents with reasonably good accuracy were relatively easy to obtain by evolving population sizes as small as 50 for 25 generations with a mutation variance σ^2 of 5. The pointing behavior of a typical agent evolved with ray sensor bias $\in$ [-10,0] and ray sensor gain $\in$ [1,15] is shown in Figure 8b. This agent had a mean accuracy of 98.13% on the 15 evaluation trials and

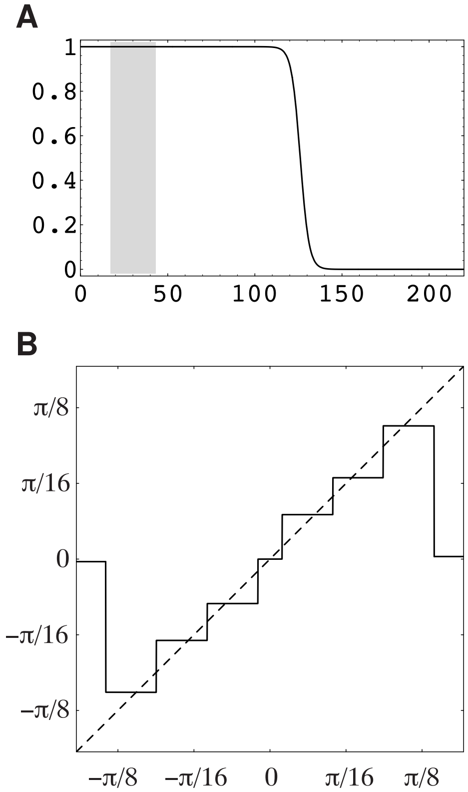

Figure 8: An evolved pointing agent. (a) A plot of the evolved ray sensor response (output vs. distance). The location of the target objects is shown in gray. (b) A plot of the final arm angle as a function of the angle of the target object. The response of a perfect pointer is shown as a dashed line.

a mean accuracy of 98.06% on 100 random trials. As Figure 8b clearly shows, the pointing resolution of this agent is limited to a number of discrete angles. The reason for this is that this agent evolved a ray sensor response that is essentially binary in nature: each ray sensor gives a response near 1 if that ray intersects the object and a response near 0 if it does not (Figure 8a). This sort of response function was typical of experiments using wider ranges of ray sensor biases.

In order to achieve higher accuracy, we must give more careful consideration to the ray sensor biases and gains. Best results were obtained for a ray sensor response such as that shown in Figure 9a. Here, the ray sensor biases and gains have been chosen so that each ray sensor goes from an output of almost 0 near the center of the object to an output of almost 1 near the front, giving a maximum dynamic range. Best results were obtained by simply fixing the ray sensor biases and gains to these "near-optimal" values.

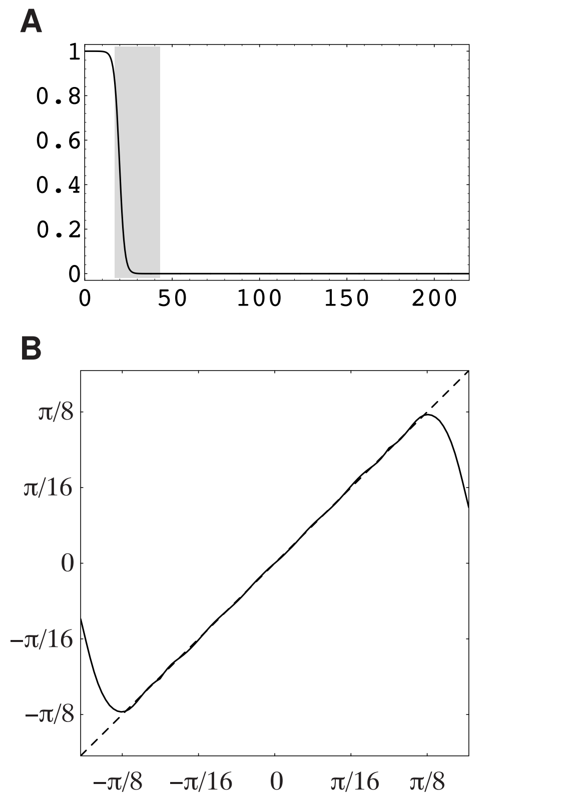

Figure 9: Another evolved pointing agent. (a) A plot of the response of an "near-optimal" ray sensor having a bias of -9.1 and a gain of 15. (b) A plot of the final arm angle as a function of the angle of a target object.

The pointing behavior of one of the best agents evolved under these conditions is shown in Figure 9b (using a population size of 100 evolved for 200 generations with a σ^2 of 5). This agent had a mean accuracy of 99.87% on the 15 evaluation trials and a mean accuracy of 99.88% on 100 random trials. Note that this accuracy corresponds to an angular error of 0.054 degrees, which is very good considering the small number of rays and interneurons used.

7 Related Work

Perhaps the work most closely related to that described in the present paper is ongoing research at the University of Sussex on visually-guided behavior (Cliff, Harvey & Husbands, 1993; Harvey, Husbands & Cliff, 1994). This work also employs a class of dynamical neural network models and evolutionary algorithms to develop agents that can center themselves in an enclosed circular arena, visually orient to an object, track a moving object, and visually discriminate between two objects. However, the emphasis of this work has been on the evolution of neural controllers for actual visually-guided robots rather than the theoretical issues being explored here. Accordingly, they are concerned with the actual transmission and detection of light rays and they are working with a physical robot. In addition, their work to date has not considered visually-guided manipulation.

Broadly speaking, the COG project (Brooks & Stein, 1994), Brooks' attempt to apply his behavior-based robotics methodology to the construction of a humanoid robot head and torso, shares the goal of exploring the cognitive implications of adaptive behavior ideas. However, the COG project has not employed dynamical neural networks and evolutionary algorithms, and it does not especially emphasize dynamical ideas. On the other hand, the goals of the COG project are much more ambitious than those of the present work, and it is also using an actual physical robot.

8 Conclusions

This paper has sketched a visually-guided agent and demonstrated the evolution of dynamical neural networks for simplified versions of visually-guided orientation, object discrimination, and accurate pointing in this agent. The first two of these experiments also illustrated the importance of internal dynamics for more complex behavior. While I do not wish to claim that the behaviors described here are cognitive, these results do represent a first step toward the evolution of more cognitively sophisticated agents. Future work will attempt to extend these capabilities in a variety of directions: selective orientation in the presence of multiple objects, discriminating among a larger set of objects, catching moving objects with an opaque hand, 2 degree-of-freedom movement and 3 degree-of-freedom manipulation, etc. Ultimately, these individual capabilities will need to be integrated into the complete agent sketched in Section 2.

However, it is worth noting that even the modest capabilities reported here already begin to raise some cognitively interesting questions. Consider, for example, the circle/diamond discrimination agents, which foveate and actively scan objects before catching or avoiding them. How do these agents achieve such a high accuracy with so few interneurons? Can we identify "circle" and "diamond" (or "smooth" and "pointy") detectors in these circuits? Will the notion of distributed representation that has been developed for static feedforward networks apply to agents controlled by recurrent dynamic circuits which actively control their perception? Or is it most appropriate to view these circuits as merely instantiating dynamics that, when coupled to their bodies and environments, give rise to effective performance of the tasks for which they were selected? Rather than debating competing intuitions in the abstract, experiments such as those described here provide concrete models within which such questions can be precisely framed and answered. Accordingly, as this research progresses, detailed studies of the operation of the circuits that evolve will be a major focus of attention.

Acknowledgments

I would like to thank Melanie Mitchell and Hillel Chiel for their detailed comments on an earlier draft of this paper. This paper has also benefited from discussions with Andy Clark. This work was supported in part by the Santa Fe Institute, and in part by grant N00014-90-J-1545 from the Office of Naval Research.

References

Altman, J.S. & Kien, J. (1989). New models for motor control. *Neural Computation* **1**:173-183.

Bäck, T. & Schwefel, H. (1993). An overview of evolutionary algorithms for parameter optimization. *Evolutionary Computation* **1**(1): 1-23.

Ballard, D. (1991). Animate vision. *Artificial Intelligence* **48**:57-86

Beer, R.D. (1995a). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence* **72**:173-215.

Beer, R.D. (1995b). Computational and dynamical languages for autonomous agents. In (Port & van Gelder, 1995).

Beer, R.D. & Gallagher, J.G. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior* **1**(1):91-122.

Brooks, R. (1991). Intelligence without representation. *Artificial Intelligence* **47**:139-159.

Brooks, R. & Stein, L.A. (1994). Building brains for bodies. *Autonomous Robots* **1**:7-25.

Churchland, P.S., Ramachandran, V.S. & Sejnowski, T.J. (1994). A critique of pure vision. In C. Koch & J. Davies (Eds.), *Large-Scale Neuronal Theories of the Brain*. MIT Press.

Clark, A. & Toribio, J. (1994). Doing without representing? *Synthese* **101**:401-431.

Cliff, D., Harvey, I. & Husbands, P. (1993). Explorations in evolutionary robotics. *Adaptive Behavior* **2**(1):73-110.

Damasio, A. (1994). *Descartes' Error*. Putnam.

Gazzaniga, M.S., Ed. (1995). *The Cognitive Neurosciences*. MIT Press.

Goldberg, D.E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley.

Harvey, I., Husbands, P. & Cliff, D. (1994). Seeing the light: Artificial evolution, real vision. In Proceedings SAB94, pp. 392-401.

Haugeland, J. (1995). Mind embodied and embedded. In Yu-Houng & Jih-Ching Ho (Eds.), *Mind and Cognition* (pp. 3-38). Taipei, Taiwan: Academia Sinica.

Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.

Knuth, D. (1981). *Seminumerical Algorithms,* 2nd Edition (The Art of Computer Programming, Vol. 2). Addison-Wesley.

Lakoff, G. (1987). *Women, Fire and Dangerous Things*. University of Chicago Press.

Port, R. & van Gelder, T. (1995). *Mind as Motion: Explorations in the Dynamics of Cognition*. MIT Press.

Posner, M.I. & Raichle, M.E. (1994). *Images of Mind*. W.H. Freeman.

Press, W.H., Teukolsky, S.A., Vetterling, W.T. & Flannery, B.P. (1994). *Numerical Recipes in C,* 2nd Edition. Cambridge University Press.

Rutkowska, J.C. (1994). Scaling up sensorimotor systems: Constraints from human infancy. *Adaptive Behavior* **2**(4):349-373.

Suchman, L. (1987). *Plans and Situated Actions*. Cambridge University Press.

Thelan, E. & Smith, L. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action.*

Yamauchi, B. & Beer, R.D. (1994). Sequential behavior and learning in evolved dynamical neural networks. *Adaptive Behavior* **2**(3):219-246.