# Evolution and Analysis of Dynamical Neural Networks for Agents Integrating Vision, Locomotion, and Short-Term Memory

**John C. Gallagher**
Dept. of Computer Science
SUNY Inst. Of Tech. at Utica/Rome
Utica, NY 13504-3050
Email: johng@fang.cs.sunyit.edu
Phone: (315) 792-7538
Fax: (315) 792-7399

**Randall D. Beer**
Dept. of Electrical Engineering and Computer Science and Dept. of Biology
Case Western Reserve University
Cleveland, OH 44106
Email: beer@alpha.ces.cwru.edu
Phone: (216) 368-2816
Fax: (216) 368-2801

## Abstract

The use of evolutionary approaches to create dynamical "nervous systems" for autonomous agents is becoming increasingly widespread. In previous work, we have successfully applied this approach to chemotaxis, walking, learning, and such minimally cognitive behavior as visually-guided orientation, object discrimination and pointing. In this paper, we extend this approach to the integration of visually-guided orientation and walking and to an object orientation task that requires short-term memory. In addition, we examine the neural dynamics underlying the operation of some of these evolved agents.

## 1 INTRODUCTION

The potential of evolutionary approaches for autonomous agents is widely recognized (Beer & Gallagher, 1992; Brooks, 1992; Cliff et al., 1993; Nolfi et al., 1994). While much of this work has emphasized relatively low-level motor behavior, there is a growing interest in applying evolutionary approaches to more sophisticated kinds of behavior, such as predator/prey interactions, visually-guided behavior, and simple forms of language, and memory (Werner & Dyer, 1991; Harvey et al., 1994; Cliff & Miller, 1996; Parisi, 1997; Nakahara & Doya, 1997).

Our own early work on the evolution of dynamical nervous systems for autonomous agents initially focused on chemotaxis and legged locomotion (Beer & Gallagher, 1992). However, more recently we have turned to tasks involving learning (Yamauchi & Beer, 1994) and what we have called minimally cognitive behavior, such as visually-guided orientation, object discrimination, and pointing (Beer, 1996). By "minimally cognitive behavior", we mean the simplest behavior that raises cognitively interesting questions. This work has been motivated both by our desire to evolve neural controllers for autonomous robots (Gallagher et al., 1996) and by our interest in analyzing the operation of evolved neural circuits in order to uncover general principles in the dynamics of adaptive behavior (Beer, 1997).

In this paper, we extend this work in two ways. First, we explore the integration of a nontrivial perceptual task with a nontrivial motor task by evolving and analyzing dynamical neural networks for visually-guided walking. Second, we evolve dynamical neural networks for a sensorimotor task that requires short-term memory.

## 2 METHODS

In both of the experiments described in this paper, the agent moves along a horizontal line through a 500 × 500 world while square objects of size 10 fall at a constant velocity of 1 from above (Figure 1). An array of 15 proximity sensors of maximum length 450 with an angular spacing of $\pi/80$ allow the agent to perceive the falling objects. If the object intersects any given proximity sensor, the output of that sensor is inversely proportional to the separation between the object and the agent, with values ranging from 0 (no intersection) to 10 (no separation).

The agent's behavior is controlled by a continuous-time recurrent neural network (CTRNN; Beer, 1995) with the following state equation:

$$\tau_i \frac{dy_i}{dt} = -y_i + \sum_{j=1}^{N} w_{ji}\sigma\left(y_j + \theta_j\right) + s_i I_i(t) \quad i = 1,\ldots,N$$

where $y$ is the state of each neuron, $\tau$ is its time constant, $w_{ji}$ is the strength of the connection from the $j^{\text{th}}$ to the $i^{\text{th}}$ neuron, $\theta$ is a bias term, $\sigma(x) = 1/(1+e^{-x})$ is the standard logistic activation function, and $I$ represents a weighted sensory input with strength $s$. States were initialized to uniform random numbers in the range ±0.1 and

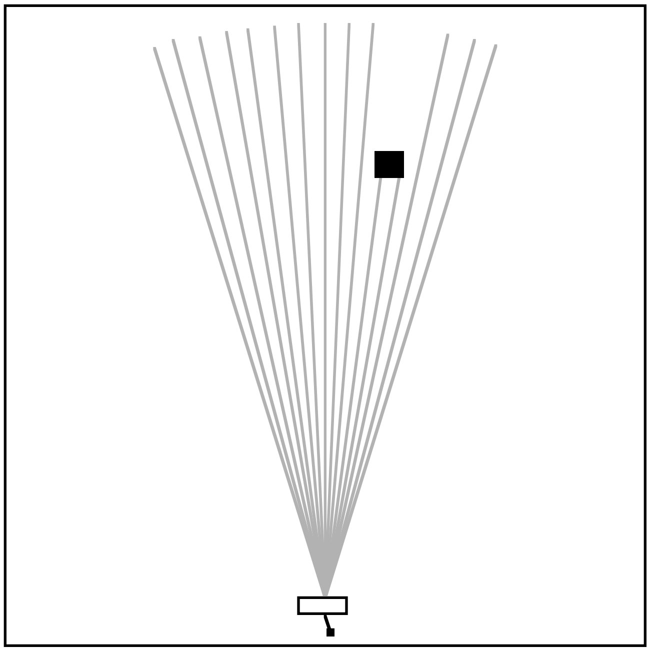

**Figure 1:** Experimental setup for visually-guided walking experiments.

circuits were integrated using the forward Euler method with an integration step size of 0.1.

We used the public domain genetic algorithm package GAucsd[1] (version 1.4) to evolve CTRNNs. Each circuit parameter was encoded in four bits, with time constants $\in$ [-0.5,10], and biases, connection weights and sensor weights in the range ±16. Note that GAucsd employs a technique known as dynamic parameter encoding (DPE) that zooms the mapping between a fixed set of bits and a real parameter based on statistics gathered during a run (Schraudolph & Belew, 1992). GAucsd parameters were set as follows: Total Trials = 220,000; Population Size = 100; Crossover Rate = 0.0; Mutation Rate = 0.01; Generation Gap = 1; Scaling Window = -1; Structures Saved = 9; Max Gens w/o Eval = 2; Options = Aclue; Maximum Bias = 0.99; Max Convergence = 376; Conv Threshold = 0.91; DPE Time Constant = 50; Sigma Scaling = 2.

## 3 VISUALLY-GUIDED WALKING EXPERIMENTS

An important step in the evolution of more sophisticated agents is the integration of multiple behaviors. In previous work, we have successfully evolved CTRNNs for legged locomotion (Beer & Gallagher, 1992) and visually-guided orientation to novel objects using pure force effectors (Beer, 1996). In the first set of experiments to be described here, we sought to evolve agents that could integrate visually-guided orientation with legged locomotion to catch objects appearing in their field of view. Unlike in our previous work on walking, which focused on optimal forward locomotion, this agent must be able to walk in either direction with movements directed toward the goal object. Unlike in our previous work on orientation, this agent must control the movements of a rhythmic leg in order to catch an object rather than simply apply a force in the correct direction.

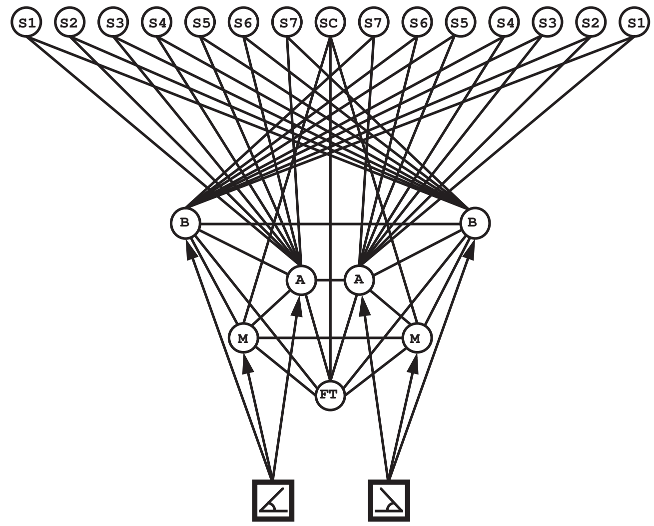


**Figure 2:** Circuit architecture for visually-guided walking experiments.

This agent has a single leg that allows it to move horizontally while objects fall from above (Figure 1). The torques produced by left and right swing effectors are summed. When the foot is up, this net torque causes the leg to swing relative to the body. When the foot is down, this net torque produces a force on the body that causes it to translate horizontally according to Newtonian dynamics. A supporting leg may stretch outside its normal range of operation ($\pm\pi/6$), but provides no translational force if it does. Whenever the foot is lifted or the leg stretches outside its maximum range (±0.99), the agent's velocity immediately drops to zero. Two leg angle sensors provide feedback of the leg position. The leg angle was randomized at the start of each trial.

The CTRNN architecture used in these experiments is shown in Figure 2. It is bilaterally symmetric and contains 22 neurons total. Fifteen sensory neurons receive inputs from the agent's proximity sensors. These connect to 4 interneurons, which in turn connect to 3 motor neurons controlling the leg. The A interneurons connect to the proximity sensors on the same side of the body, while the B interneurons connect to the proximity sensors on the opposite side of the body. Only the central proximity sensor SC connects directly to the leg motor neurons. The $M_L$ and $M_R$ motor neurons control the torque with which the leg swings to the left or right, respectively, while the FT motor neuron

[1] URL: ftp://ftp.aic.nrl.navy.mil/galist/src/GAucsd14.sh.Z

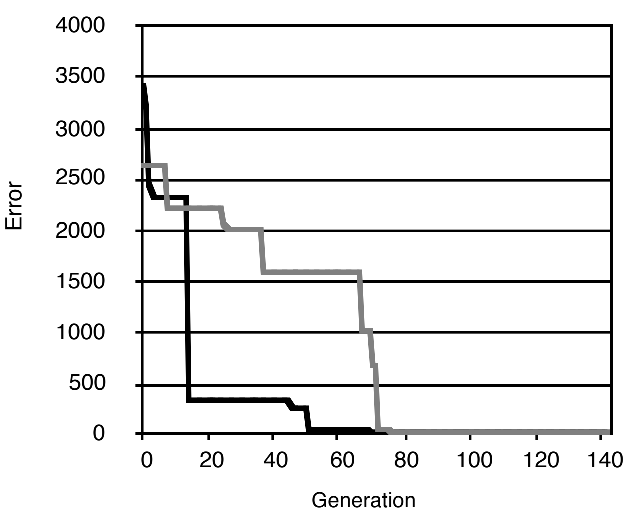


**Figure 3:** Plots of error of the best individual in each generation for two typical experiments.

controls the state of the foot. The motor neurons also receive weighted inputs from two leg angle sensors. Except for the angle sensor inputs, all connections shown are bidirectional, giving a total of 94 parameters encoded in 376 bits.

During evolution, fitness was evaluated by dropping objects with fixed velocities from different horizontal positions.

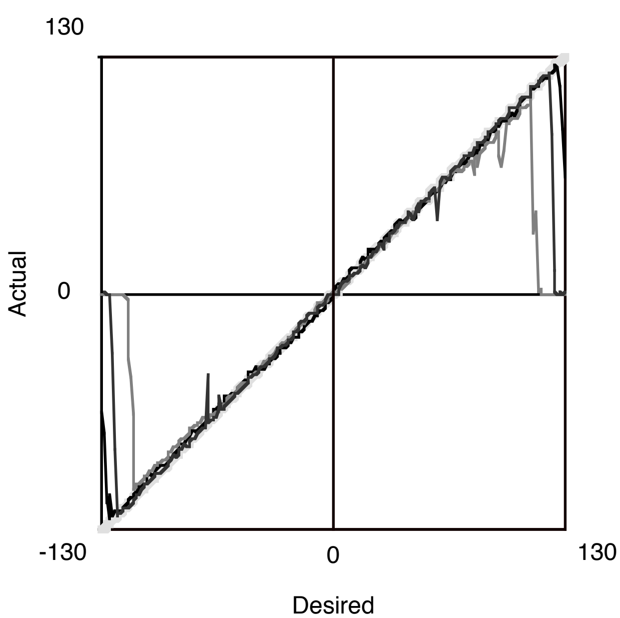


**Figure 4:** Plots of actual vs. desired final position for the top three walking agents. Note how closely the performance of these agents approximates a diagonal line over the field of view.

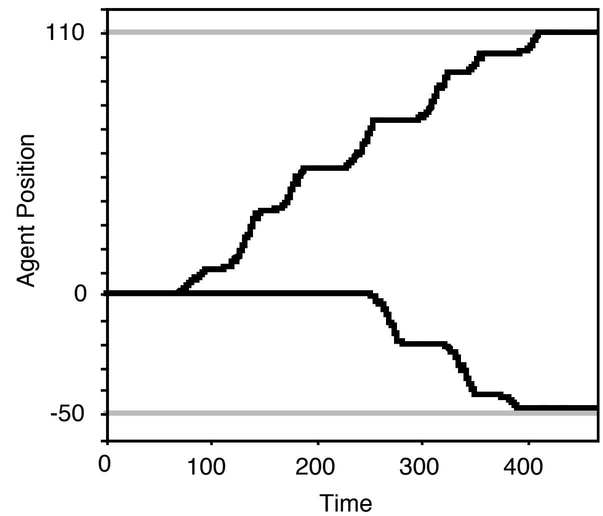


**Figure 5:** Plots of agent (black lines) and object (gray lines) position vs. time for two different initial horizontal object positions.

The error measure to be minimized was:

$$Average\ Error = \sum_{i=1}^{NumTrials} d_i^2 \Big/ NumTrials$$

where $d_i$ is the horizontal distance between the object and the agent when the object reaches the agent on the $i^{\text{th}}$ trial, and *NumTrials* is the total number of trials. Thus, a perfect agent would have an error of zero because the horizontal positions of the objects and the agent would always coincide when the object reached the agent. We used six trials with objects dropped at horizontal distances of 10, 30, 50, 70, 90 and 110 to the left of the agent (due to the bilateral symmetry of the body and CTRNN architecture, only trials on one side of the agent are necessary).

Five experiments were run. Each experiment was terminated after 143 generations because each run required approximately five days on an unloaded Sun Ultra 1. Two typical runs are shown in Figure 3. All runs exhibited a very similar progression. Initially, agents evolve that can only catch objects falling near the agent's initial position (i.e., for which very little movement is necessary). However, as the evolution progresses, the agents begin to catch objects increasingly more peripheral to the their initial position. In addition, as the evolution progresses, the accuracy with which any given object is caught improves. By the end of the search, the best agents can reliably catch objects falling from all six initial positions.

The top three agents achieved final errors of 1.40, 1.65 and 12.73, respectively on the six evaluation trials, and errors of 9.19, 33.42 and 6.87, respectively on 2000 random trials. All three best agents generalized quite well to initial object positions other than the six they were evolved on, catching

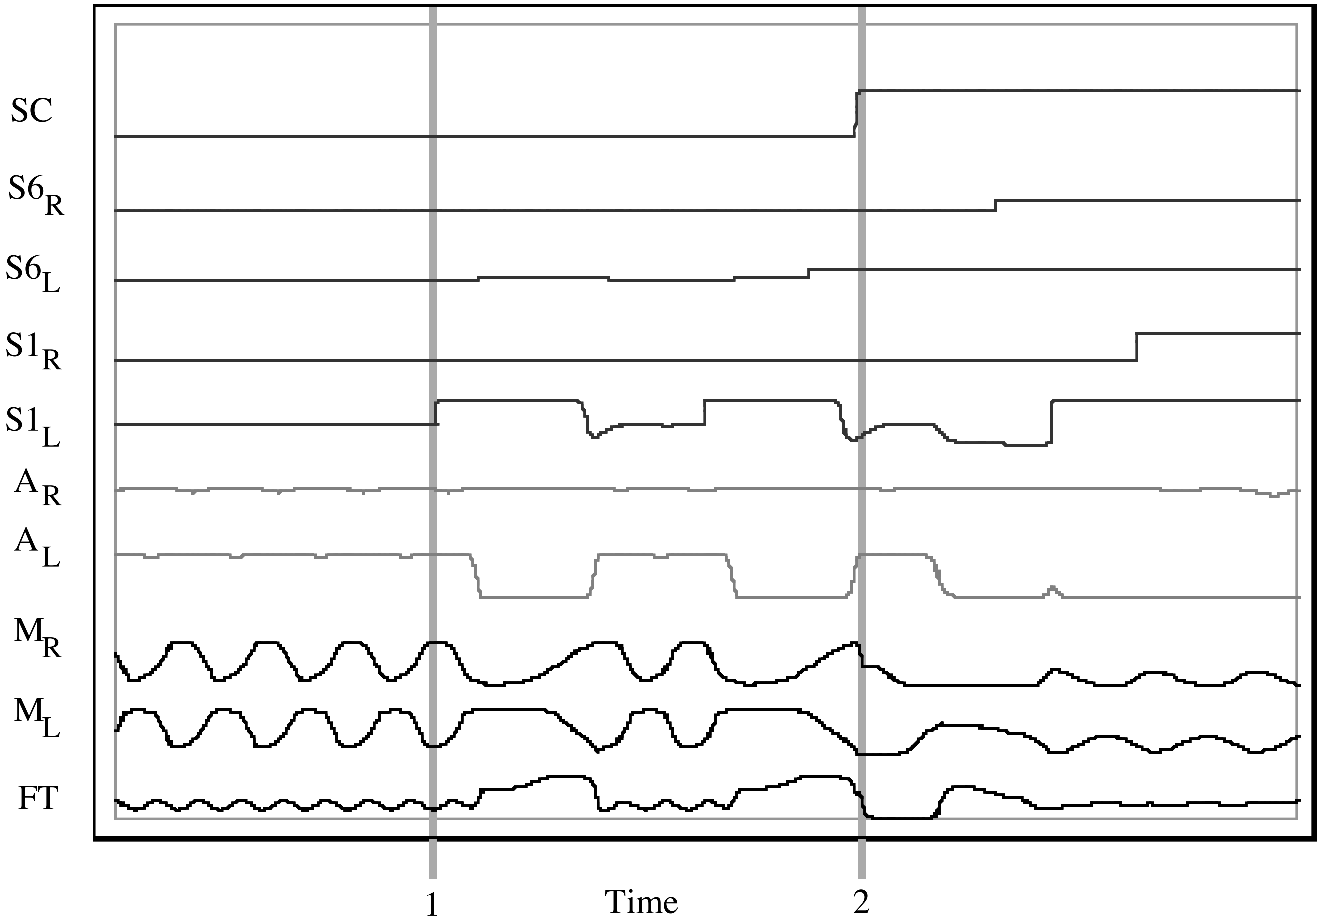


**Figure 6:** Neural activity of the key neurons in the best evolved visually-guided walker while it catches an object (corresponding to the lower agent position trace in Figure 5). Before 1, the agent waits in place swinging its leg. At 1, the object intersects the left-most proximity sensor $S1_L$ and the agent begins walking to the left. At 2, the object intersects the central proximity sensor SC and the agent once again swings its leg in place. See text for a more detailed explanation of the neural interactions underlying the switches between these different phases of behavior.

objects with good accuracy anywhere within the range over which they evolved (Figure 4).

The behavior of the best agent on two trials is shown in Figure 5. The best agent uses the following strategy. When no object is visible, the agent remains in place with its foot raised and its leg swinging back and forth. When an object appears, the agent begins to alternate stance phases (where the foot is down and the swinging leg moves the agent toward the object) and swing phases (when the foot is up and swinging back to prepare for the next stance) to walk toward the object. Note that, for some initial conditions, the agent waits longer to respond to the object than for others. Also note that the pauses between steps vary in duration. As the agent nears the object, the durations of the stance phases decrease until the foot once again is held up while the leg swings back and forth. The agent then waits in place with the foot raised and the leg swinging back and forth until the object reaches it.

Four of the five best agents evolved in these experiments utilize a very similar strategy. However the worst of them exhibits slightly different behavior. Instead of swinging its leg back and forth while waiting for an object, this fifth agent paces back and forth a small distance until it observes a falling object.

## 4 ANALYSIS OF A VISUALLY-GUIDED WALKER

How do the evolved visually-guided walkers work? Analyses of evolved dynamical nervous systems for model agents can give us both insights into general principles of evolved biological neural circuits and insights into the design and evolution of more sophisticated artificial agents (Beer, 1997). Although space does not permit a rigorous mathematical analysis (for examples of such analyses, see (Gallagher & Beer, 1993; Beer, 1995; Chiel et al., in press)), we will give a qualitative explanation for the operation of the best evolved visually-guided walker (Figure 6). There are three questions that we would like to answer: How is the basic leg oscillation generated? How do the proximity sensors control the stance phase of this cycle so as to position the agent under a falling object? How is walking terminated once the agent is accurately positioned?

By selectively eliminating neurons from the circuit, it can be shown that the basic leg oscillation is primarily generated by the $M_L$ and $M_R$ motor neurons and the leg angle sensors. This subcircuit operates like other evolved sensory-driven walking circuits that we have previously analyzed (Gallagher & Beer, 1993; Gallagher, 1998), with feedback from the leg angle sensor switching the $M_L/M_R$ circuit through a multistable regime separating two different fixed point attractors. Note that this oscillation is symmetric. When the agent is actually walking, this circuit exhibits an asymmetry between stance and swing phase duration due to the different inertias of the body and the leg (Beer et al., in press).

How is the $M_L/M_R$ subcircuit switched between swinging the leg in place and walking when an object appears? In this circuit, only SC and the left/right pairs of the S1 and S6 proximity sensors are sensitive to objects. When any of the S1 pairs or S6 pairs intersects an object, its output goes to nearly 1 (Figure 6, marker 1). This strongly inhibits the A interneuron on the same side of the body, turning it off. Normally, the foot motor neuron exhibits small oscillations during swinging behavior which are below the threshold for putting the foot down. A lack of inhibition from A allows the foot motor neuron FT to turn on, putting the foot down and generating a stance phase that moves the agent. Interestingly, because the A interneuron is phase-locked to the oscillations in $M_L/M_R$, S1 can only shut it off (and hence force the foot down) when the leg is swinging toward the same side, generating a step in the correct direction. With the foot motor neuron on, excitation to interneuron A causes A to turn back on, re-establishing the swinging behavior until the object once again intersects one of the active proximity sensors and the stepping cycle repeats.

Note that the fact that only SC and the left/right pairs of S1 and S6 proximity sensors are sensitive to objects explains the earlier observation that the agent begins to respond to objects at differing distances for different initial positions (Figure 5). If an object falls between S1 and S6, it won't be detected by the agent until it falls far enough to intersect S1. This also explains the fact that the agent sometimes pauses between steps (Figure 5). Because the other proximity sensors are not sensitive to objects, a step that places the object between S1 and S6 will produce a pause until the S1 proximity sensor is activated as the object continues to fall. These episodes of swinging behavior can easily be seen between the two steps shown in Figure 6 (between markers 1 and 2).

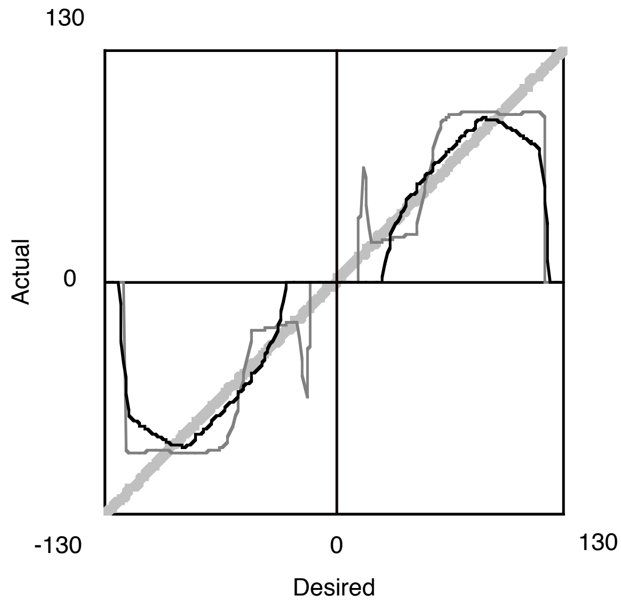


**Figure 7:** Plots of actual vs. desired final position for two short-term memory agents. Note that one approximates a diagonal line for a wide range (thin black trace) and the other appears to have memorized a set of positions (thin gray trace)

How is walking terminated once the agent is accurately positioned beneath an object? When the center proximity sensor SC jumps to 1 due to an object intersection, it strongly inhibits the $M_L$ and $M_R$ motor neurons (Figure 6, marker 2). In addition, when $A_L$ comes back on, it further inhibits the $M_L/M_R$ oscillator. These combined effects shift the $M_L/M_R$ oscillator and its interaction with $A_L$ and, through $A_L$, $S1_L$ into a new mode in which the foot motor neuron can never get above threshold. Therefore, the foot no longer goes down and the agent simply waits in place, swinging its leg.

Qualitatively, the best agents evolved in the other four experiments work in a similar fashion. Each of them also exhibits a sensory-driven swing oscillator that is switched into a stepping mode by the interneurons. The main difference is that, in the agent that paces back and forth while waiting for an object, the activity of FT sometimes go above threshold even when no object is present. Of course, the details of which proximity sensors are sensitive to objects and how the interneurons switch the swing oscillator into walking mode vary greatly from agent to agent.

## 5 SHORT-TERM MEMORY EXPERIMENTS

Most nontrivial behavior requires the integration of experiences across time and the ability to initiate actions independent of an agent's immediate circumstances. In other words, they require some sort of internal state or memory. Of course, because any CTRNN is a dynamical system, its response to a given input will automatically be dependent on

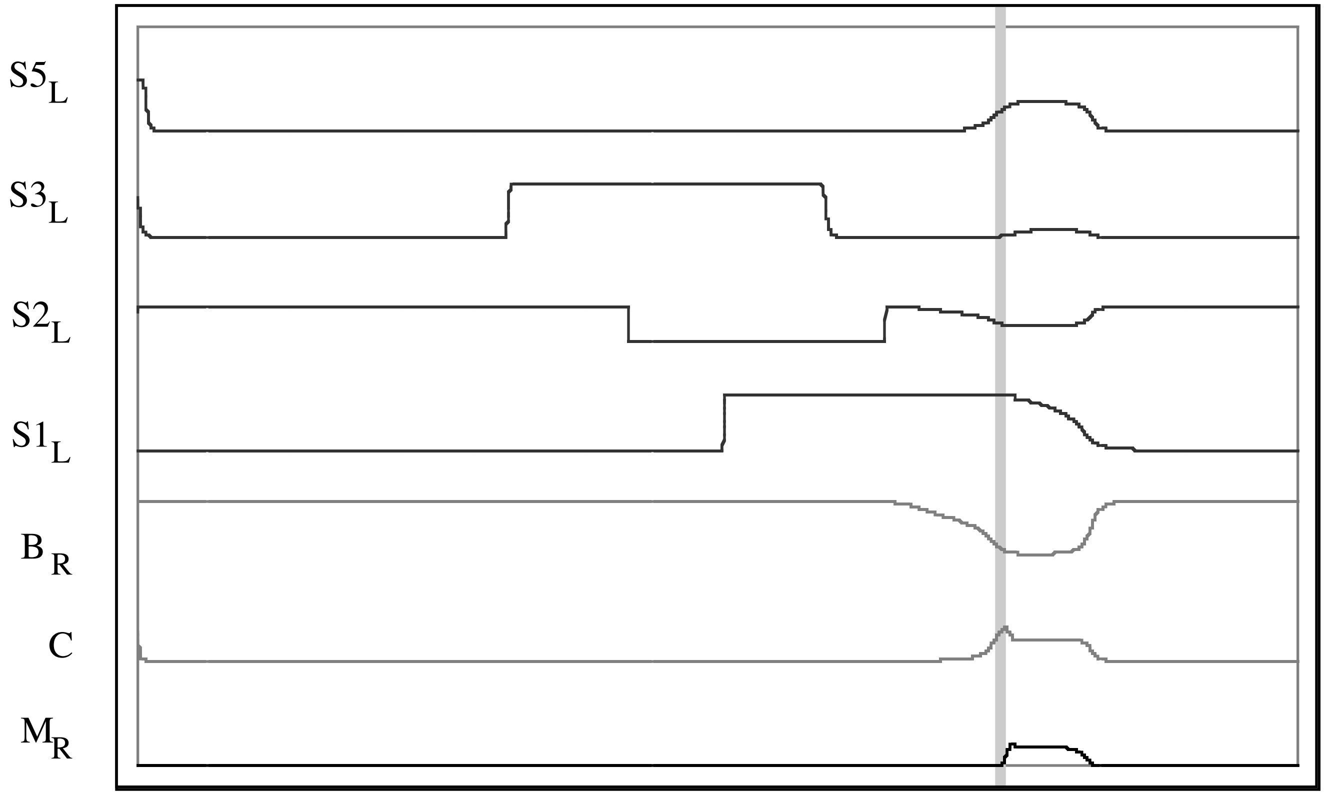


**Figure 8:** Neural activity of the key neurons in the best evolved short-term memory agent while it catches an object. All traces are labeled with the names of associated neurons as defined in Figure 2 except that FT has been renamed to reflect the fact that, for short-term memory agents, that neuron is merely an interneuron with no pre-designated function. The gray vertical bar denotes the time at which the agent begins to move and loses its sight.

its state and thus its previous history of interactions. In this sense, even the visually-guided walkers depend on a form of short-term memory for their operation. Some have argued that only tasks that require an agent to coordinate its behavior with environmental features that are not immediately present are sufficiently “representation hungry” to be of cognitive interest (Clark, 1997). As a first step in this direction, we sought to evolve agents that could catch objects after only briefly observing their initial position and then moving to the correct location while blind to the object’s subsequent motion.

The experimental setup for the short term memory experiments was nearly identical to that for the visually-guided walking experiments (Figure 1). However, Newtonian dynamics were replaced by first-order dynamics ($F = mv$) and the leg was replaced by a pair of pure velocity effectors. The circuit architecture was also identical to that used in the visually-guided walking experiments (Figure 2), except for the following changes. First, the outputs of $M_L$ and $M_R$ now control the left and right effectors which sum to produce a net velocity. Second, FT is now an interneuron with no pre-specified function. Third, the angle sensors were removed along with the leg. The proximity sensors operated as before except that once the agent exceeds a velocity threshold of 0.01 the outputs of all sensors are permanently set to zero. This requires the agent to anticipate when it has reached an object’s position even though it cannot directly observe the object.

Three experiments were run using the same error measure described earlier. Each experiment was terminated after approximately 200 generations, requiring approximately one week on an unloaded Sun Ultra 1. Two of the three experiments evolved agents that could perform the task fairly well, with final errors of 164.86 and 222.53 on the six evaluation trials and 237.58 and 138.59 respectively on 2000 random trials. While not as good as the visually-guided walkers (Figure 4), the accuracy of these two agents is still reasonably good over much of the range over which they were evolved (Figure 7), especially given the fact that they have no perception of the falling object as they move.

Each of the best agents evolved in the two successful experiments adopted slightly different strategies. Both agents wait at their initial position and watch the falling object for some period of time. They then move to their final position with a one time activation of the appropriate effector that decays to zero at the right position. One of the two agents seems to have memorized a number of stereotyped positions to which it can move and is incapable of interpolating between them (gray trace in Figure 7). This agent had the smaller error on the six evaluation trials, but the higher error on the random trials. The other agent seems to be able to catch objects reasonably accurately across most of the range, erring only on objects that are near the center or the periphery of its field of view (black trace in Figure 7). Because this agent did not overspecialize on the six evaluation trials, it scored worse on them but better on the 2000 random trials. The neural activity of the key neurons in this agent as it catches an object is shown in Figure 8.

# 6 CONCLUSIONS

In this paper, we have extended our previous work on the evolution of dynamical "nervous systems" for minimally cognitive behavior in two ways. First, we have examined a visually-guided walking task which integrates an object orientation problem with a nontrivial motor control problem. Second, we have examined a task requiring short-term memory, in which an agent's behavior in the present must be coordinated with an object it has observed in the past. These results demonstrate the potential of this approach and suggest that it may be possible to incrementally extend it to more sophisticated behaviors. Current work is focused on evolving dynamical neural networks for such capabilities as focus of attention, object manipulation, and self/nonself discrimination.

We have also examined the operation of some of the neural circuits that we have evolved. Previously, the mathematical tools of dynamical systems theory have been successfully employed in the analysis of a variety of CTRNNs (Gallagher & Beer, 1993; Beer, 1995; Chiel et al., in press; Beer et al., in press). Work is currently underway to extend these analyses to more sophisticated evolved agents, such as those capable of visually-guided object discrimination, which foveate and actively scan objects before deciding whether to catch or avoid them (Beer, 1996). One thing that is clear from these analyses is that the dynamics of evolved circuits is usually subtle and counter-intuitive. The memory agents would be especially interesting to analyze. How do these agents "remember" the location of an object that they can no longer see? Is this location represented in the activation of an individual neuron or a group of neurons? If so, how does the agent's initial perception of the object store its location and how is this stored location updated as the agent moves? If not, then how are we to explain the dynamical basis of these agents' anticipatory behavior? How might this explanation change if the objects fell diagonally rather than vertically, requiring a more sophisticated form of prediction? These are the kinds of questions that we hope to answer as this research progresses.


### Acknowledgments

This work was supported in part by grant RG0084/1997-B from the Human Frontier Science Program, grant DMI97-20309 from the National Science Foundation, and grant N00014-96-10708 from the Office of Naval Research. Additional support provided by The Research Foundation of the State University of New York for and on behalf of the SUNY Institute of Technology.


### References


Beer, R.D. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior* **3**(4):469-509.

Beer, R.D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In P. Maes, M. Mataric, J. Meyer, J. Pollack and S. Wilson (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (pp. 421-429). MIT Press.

Beer, R.D. (1997). The dynamics of adaptive behavior: A research program. *Robotics and Autonomous Systems* **20**:257-289.

Beer, R.D., Chiel, H.J. & Gallagher, J.C. (in press). Evolution and analysis of model CPGs for walking II. General principles and individual variability. To appear in *J. Computational Neuroscience*.

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

Brooks, R.A. (1992). Artificial life and real robots. In F.J. Varela and P. Bourgine (Eds), *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life* (pp. 3-10). MIT Press.

Chiel, H.J., Beer, R.D., & Gallagher, J.C. (in press). Evolution and analysis of model CPGs for walking I. Dynamical modules. To appear in *J. Computational Neuroscience.*

Clark, A. (1997). *Being There: Putting Brain, Body and World Together Again*. MIT Press.

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

Cliff, D. & Miller, G.F. (1996). Co-evolution of pursuit and evasion II: Simulation methods and results. In P. Maes, M. Mataric, J. Meyer, J. Pollack and S. Wilson (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (pp. 506-515). MIT Press.

Gallagher, J.C. & Beer, R.D. (1993). A qualitative dynamical analysis of evolved locomotion controllers. In J. Meyer, H. Roitblat and S. Wilson (Eds.), *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 71-80). MIT Press.

Gallagher, J.C., Beer, R.D., Espenschied, K.S. & Quinn, R.D. (1996). Applications of evolved locomotion controllers to a hexapod robot. *Robotics and Autonomous Systems* **19**(1):95-103.

Gallagher, J.C. (1998). *A Dynamical Systems Analysis of the Neural Basis of Behavior in an Artificial Autonomous Agent*. Ph.D. Dissertation, Department of Computer Engineering and Science, Case Western Reserve University, Cleveland, OH.

Harvey, I., Husbands, P. & Cliff, D. (1994). Seeing the light: Artificial evolution, real vision. In D. Cliff, P. Husbands, J. Meyer, and S. Wilson (Eds.), *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior* (pp. 392-401). MIT Press.

Nakahara, H. & Doya, K. (1997). Near-saddle-node bifurcation behavior as dynamics in working memory for goal-directed behavior. *Neural Computation* **10**:113-132.

Nolfi, S., Floreano, D., Miglino, O. & Mondada, F. (1994). How to evolve autonomous robots: Different approaches in evolutionary robotics. In R.A. Brooks and P. Maes (Eds.), *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems* (pp. 190-197). MIT Press.

Parisi, D. (1997). Artificial life and higher level cognition. *Brain and Cognition* **34**: 160-184.

Schraudolph, N.N. & Belew, R.K. (1992). Dynamic parameter encoding for genetic algorithms. *Machine Learning* **9**(1).

Werner, G. M. & M. G. Dyer (1991). Evolution of communication in artificial organisms. In C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen (Eds.), *Artificial Life II* (pp. 659-687). Addison-Wesley.

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