Submitted to the Sixth International Conference on Simulation of Adaptive Behavior, Sept. 11-15, 2000, Paris, France

Further Experiments in the Evolution of Minimally Cognitive Behavior: From Perceiving Affordances to Selective Attention

Andrew C. Slocum[1], Douglas C. Downey[1] and Randall D. Beer[1,2]
[1]Dept. of Electrical Engineering and Computer Science
[2]Dept. of Biology
Case Western Reserve University
Cleveland, OH 44106
{acs2,dcd5}@po.cwru.edu, beer@eecs.cwru.edu

Abstract

In this paper, we extend previous work on the evolution of continuous-time recurrent neural networks for *minimally cognitive behavior* (the simplest behavior that raises issues of genuine cognitive interest). Previously, we evolved dynamical "nervous systems" for orientation, reaching, and discrimination. Here we evolve agents that can judge the passability of openings relative to their own body size, discriminate between visible parts of themselves and other objects in their environment, predict and remember the future location of objects in order to catch them blind, and switch their attention between multiple distal objects.

1. Introduction

Notions of situatedness, embodiment and dynamics are playing an increasingly influential role in both the foundations and the practice of cognitive science (Clark, 1997; Beer, 2000). This offers an excellent opportunity for work in adaptive behavior to contribute to cognitive science, and many such efforts are underway (Brooks & Stein, 1994; Almássy et al., 1998; Pfeiffer & Scheier, 1999). Evolutionary approaches are particularly fruitful because they allow an exploration of possible cognitive architectures relatively unencumbered by *a priori* assumptions. Such approaches are being successfully applied to increasingly more sophisticated behavior (Harvey et al., 1994; Cliff & Miller, 1996; Nakahara & Doya, 1997; Parisi, 1997; Di Paolo, 1997).

In our own work, we use genetic algorithms to evolve dynamical "nervous systems" for model agents and then analyze the dynamics of the resulting systems (Beer, 1997). We believe that such simpler idealized models can serve as "frictionless planes" in which basic theoretical principles of the dynamics of agent-environment systems can be worked out. Because we are ultimately interested in cognitive questions, we have begun to focus this work on *minimally cognitive behavior*, the simplest possible agent-environment systems that raise issues of genuine cognitive interest (Beer, 1996). If we hope to evolve and ultimately analyze in detail model agents exhibiting genuinely cognitive behavior, it is essential to focus on the simplest possible agent-environment systems that exhibit the cognitive behavior of interest. It is also important that the cognitive issues arise in a natural way in the context of the agent's behavior, rather than being posed abstractly. In previous work, we introduced a simple visually-guided agent that could potentially be used to address a wide range of basic cognitive phenomena and demonstrated the evolution of dynamical "nervous systems" for orientation and reaching to objects and discrimination between objects (Beer, 1996).

In this paper, we extend this work to a wider range of more complicated tasks. First, we explore agents that must visually decide which openings their bodies can and cannot fit through. Second, we evolve agents that must distinguish between visible parts of themselves and objects in their environment. Third, we examine a task that requires an agent to predict and remember the future location of a target object in order to catch it blind. Finally, we explore a task in which an agent must switch its attention between multiple distal objects.

2. Methods

In all of the experiments described in this paper, an array of proximity sensors allowed an agent to perceive distal objects that fall toward it from above. If an object intersected a proximity sensor, the output of that sensor was inversely proportional to the separation between the object and the agent, with values ranging from 0 (no intersection) to 10 (no separation). The agent moved according to first-order dynamics, with motor neurons directly specifying the velocity of movement.

The agent's behavior was controlled by a continuous-time recurrent neural network (CTRNN) 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.

A real-valued genetic algorithm (Mitchell, 1996) was used to evolve CTRNN parameters. A population of individuals was maintained, with each individual encoded as a length M vector of real numbers. Initially, a random population of

vectors was generated by initializing each component of every individual to random values uniformly distributed over the range ±1 (they could move outside this range during evolution). Individuals were selected for reproduction using a linear rank-based method. A specified elitist fraction of top individuals in the old population were simply copied to the new one. The remaining children were generated by either mutation or crossover with an adjustable crossover probability. A selected parent was mutated by adding to it a random displacement vector whose direction was uniformly distributed on the *M*-dimensional hypersphere and whose magnitude was a Gaussian random variable with 0 mean and variance σ^2. The expression derived in the Appendix was used as a guideline for setting the mutation variance. A neuron's time constant, bias, gain and input weights were treated as a module during crossover.

Unless otherwise indicated, search parameters in the range ±1 were mapped linearly into CTRNN parameters with the following ranges: connection weights ∈ [-5,5], biases ∈ [-10,0], gains ∈ [1,5], and time constants ∈ [1,2]. All proximity sensors shared the same time constant. All CTRNNs were bilaterally symmetric. While this often made trials involving nearly-centered objects difficult, it reflects the symmetry of the agent and the tasks, and it halves the number of parameters that must be evolved.

3. Perceiving Affordances

Any situated, embodied agent must be sensitive to the relationship of its own body to its surroundings and it must be able to perceive the actions that this environment affords in somatic terms (Gibson, 1979). For example, in order for an agent to perceive whether or not an aperture is passable, it must judge the aperture's width relative to its own body (Warren & Wang, 1987). In our first set of experiments, we evolved agents that could accurately distinguish between passageways and obstacles in a falling wall, passing through openings wide enough to accommodate their bodies while avoiding openings that were too narrow.

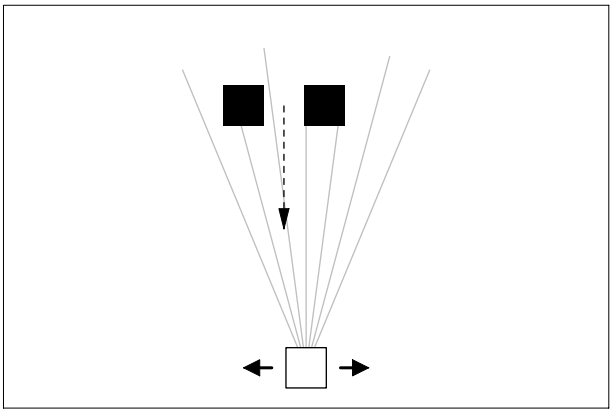

Figure 1: Experimental setup for the passability experiments. The agent moves horizontally while a wall with an adjustable aperture falls from above. The rays of the agent's proximity sensors are shown in gray.

Square agents of size 20 had 7 proximity sensors of maximum length 160 uniformly distributed over a visual angle of π/4 (Figure 1). Their horizontal velocity was proportional to the sum of opposing forces produced by a bilateral pair of effectors (with a constant of proportionality of 8). Walls consisting of two squares of width 20 separated by an aperture whose width was in the range [16,24] dropped from above with a vertical velocity of 4 and a horizontal offset of ±50 relative to the agent.

The circuit architecture was bilaterally symmetric, with 7 sensory neurons projecting to 6 fully interconnected interneurons that in turn projected to two motor neurons controlling horizontal motion (for a total of 71 parameters). Populations of 100 individuals were evolved for 2000 generations with a mutation variance σ^2 of 0.3, a crossover probability of 0.5 and an elitist fraction of 5%.

The performance measure to be maximized was:

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

where $p_i = \begin{cases} 2|d_i| & \text{if agent collides with wall} \\ 100 & \text{otherwise} \end{cases}$

for an opening too narrow for the agent to pass through and

$$p_i = \begin{cases} \max\left(0, 80 - 4|d_i|\right) & \text{if agent collides with wall} \\ 100 & \text{otherwise} \end{cases}$$

for an aperture wide enough for the agent to pass through, and d_i is the final horizontal separation between the center of the agent and the center of the aperture at the end of the i^{th} trial. This fitness measure assigns near-zero fitness to incorrect actions and linearly penalizes near-misses. Since making the correct decision without hitting the wall results in a significantly higher score, this performance measure

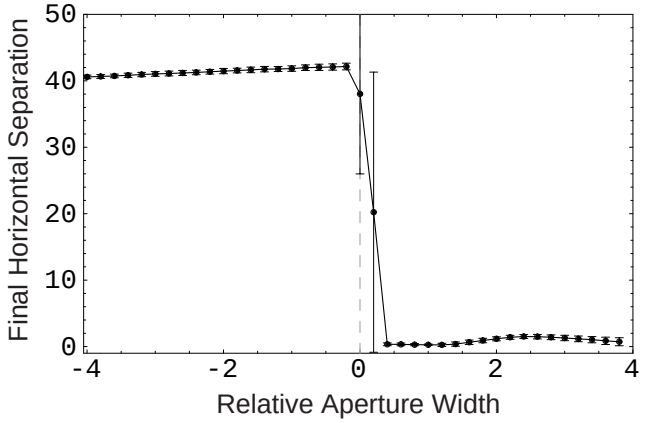

Figure 2: Categorization of apertures into passable and impassable by the best passability agent. The final horizontal separation between the agent and the center of the aperture (mean ± s.d., N = 101 trials) is plotted against the aperture width relative to the agent's size.

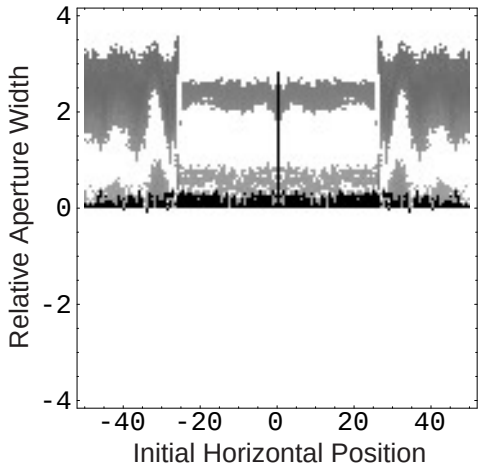

Figure 3: Generalization performance of the best passability agent. The performance as a function of initial horizontal position and aperture width relative to the agent is shown as a density plot. The highest performance is shaded white and the lowest performance is shaded black.

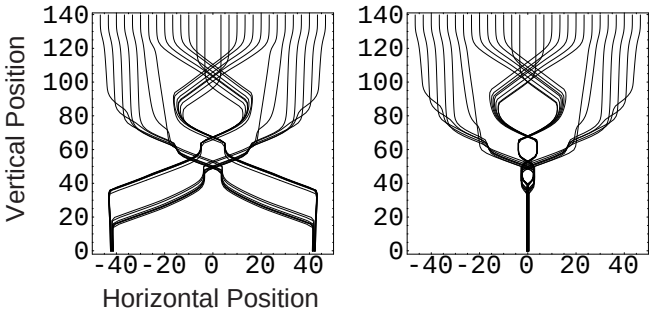

Figure 4: Behavior of the best passability agent. The wall's horizontal and vertical position over time relative to the agent is plotted for an aperture 1 unit smaller than the agent (left) and 1 unit larger than the agent (right). Trials begin at top and time increases from top to bottom.

also rewards accuracy. The search began with 4 easy and 2 difficult test cases. Every time that the best agent's average fitness exceeded 90%, two additional trials were added until a total of 30 trials was reached.

Out of a total of 30 runs, 6 produced agents that achieved an average performance greater than 90% on all 30 trials. The best agent had a mean fitness of 99.2% on the 30 evaluation trials and 96% on 1000 random trials. By plotting the mean final horizontal separation between the agent and the center of the aperture as a function of aperture width (Figure 2), we see that the agent makes a very sharp categorical distinction between passable and impassable apertures at its own body width (dashed gray line). Since the evaluation trials are characterized by only 2 parameters (initial horizontal offset and aperture width), we can directly visualize the agent's generalization performance as a 2-dimensional density plot (Figure 3). Note that this agent can accurately discriminate aperture size differences smaller than 0.5 (2.5% of its body size), which is quite good given the small number of proximity sensors that it possesses. The gray areas in this plot indicate trial parameter regimes where the agent brushes the wall when passing through an aperture. Not surprisingly, the majority of these gray regions fall between the 30 evaluation trials used during evolution. Interestingly, this agent prefers to err on the side of avoiding apertures that are just barely large enough to fit through rather than trying to fit through apertures that are too small.

The behavior of the best agent is shown in Figure 4 for aperture widths just below (left) and just above (right) the agent's width. Note that the two plots exhibit a great deal of qualitative similarity. In both cases, the trajectories of motion are grouped into two distinct bundles depending on the initial horizontal offset of the wall. The distinguishing feature seems to be whether or not the wall intersects the outermost ray shortly after a trial begins. More central trials result in scans that repeatedly cross the midline, while the more peripheral trials result in a slower centering movement that crosses the midline only once. The two plots qualitatively differ only in the final orientation movement, with the agent eventually moving to avoid an impassable aperture (left), but centering a passable aperture (right).

The strategies of the other top agents varied widely. Some agents distinguished between central and peripheral trials, while others did not. Some agents initially foveated the aperture, while others initially foveated one side of the wall instead. Some scanned the wall multiple times before deciding, while others scanned only once. However, all of the top agents were quite decisive. If they received a less than perfect score, it was generally because they made the wrong decision or they were slightly sloppy in centering a barely passable opening, not because they equivocated between centering and avoidance and ended up colliding with the middle of one of the wall blocks.

4. Self/NonSelf Discrimination

Any agent possessing a spatially extended body has the potential to perceive this body with its own distal sensors. This possibility raises the problem of distinguishing self from nonself depending on which of the objects in an agent's field of view are under its direct control (Neisser, 1993). In a second set of experiments, we evolved agents that could catch moving objects with an opaque hand. These experiments extend previous work on evolving agents that could point to stationary objects with a transparent manipulator (Beer, 1996).

Agents of size 20 had 7 proximity sensors of maximum length 160 uniformly distributed over a visual angle of $\pi/4$ (Figure 5). They were unable to move, but had an opaque hand of size 5 centered on the end of a transparent arm of length 25 with one angular degree of freedom having an angular range of $\pm\pi/2$. The angular velocity of the arm was proportional to the sum of two opposing torques produced by a bilateral pair of effectors (with a constant of proportionality of 0.15). Circular objects of diameter 20 dropped from above with a vertical velocity of 4 and an initial hori-

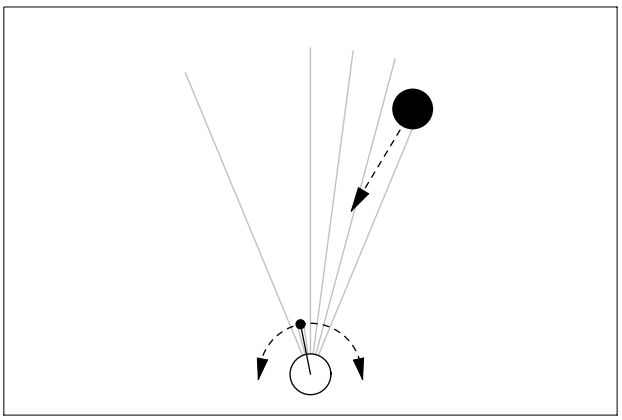

Figure 5: Experimental setup for self/nonself discrimination experiments. The agent is stationary, but can swing an arm with an opaque hand along an arc while objects fall from above.

zontal offset in the range ±24 relative to the agent. Given the initial horizontal offset, an object's horizontal velocity was drawn from a range than guaranteed its path would intersect the arc of the hand.

The circuit architecture was bilaterally symmetric, with the same architecture as in the previous experiments with the addition of a bilateral pair of arm angle sensors (for a total of 74 parameters). Each sensor was sensitive to arm displacement in one direction only, with an output of 0.1 when the arm was centered and an output of 0.9 at the edges of vision ($\pm\pi/8$). Populations of 100 individuals were evolved for 1000 generations using a mutation variance σ^2 of 0.1, a crossover probability of 0.5 and an elitist fraction of 5%.

The performance measure to be maximized was:

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

where

$$p_i = 1 - \frac{\max\left(\frac{\pi}{4}, |\theta_i|\right)}{\pi/4}$$

and θ_i is the angular error at the end of the i^{th} trial. The search began with 5 initial trials. Every time that the best agent's average fitness exceeded 90%, 5 additional trials were added until a total of 30 trials was reached.

Out of a total of 16 runs, 7 produced agents that achieved an average performance greater than 90% on all 30 trials. The best agent had a mean fitness of 97.6% on the 30 evaluation trials and 95% on 1000 random trials. As shown in Figure 6, its accuracy is very good.

The behavior of the best agent catching objects at the midline is shown in Figure 7. This agent oscillates its hand back and forth at one of 4 different locations before moving to the midline as the object nears. Note that the hand can slip from one location to another. Interestingly, the hand generally oscillates on the same side of the agent as the object appears initially, repeatedly occluding the falling object.

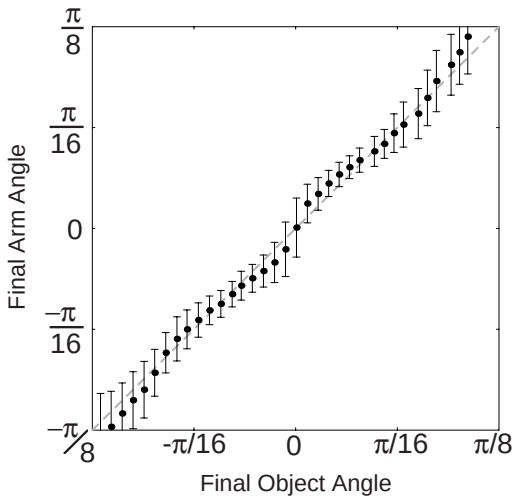

Figure 6: Mean catching accuracy of the best self/nonself discrimination agent. The final angular position of the hand is plotted against the final angular position of the object (mean ± s.d., N = 150 trials). Note that the average behavior closely approximates the ideal (dashed gray line).

Also, note that there seems to be a small but systematic error in the final hand position. The hand tends to end a bit too far to the left for objects that originate from the left and vice versa. In addition, although not shown in this figure, it turns out that the most difficult objects for this agent to catch are those that completely cross its field of view (i.e., they originate at the extreme left but intersect the arc of the hand at the extreme right, or vice versa). The arm angle sensors play a crucial role in this agent's behavior. If these sensors are lesioned, the arm initially follows the same trajectories shown in Figure 7, but then swings out of the field of view and never returns. In contrast, making the hand transparent has only a small effect on the agent's behavior (the two arm angle trajectory bundles on each side of the

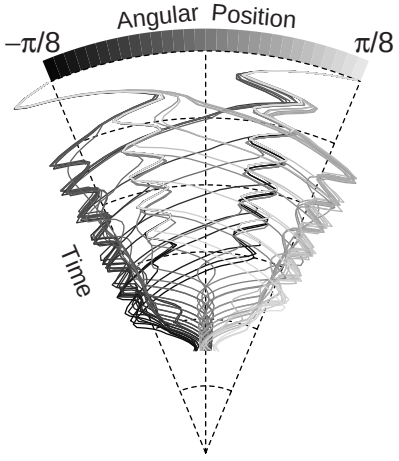

Figure 7: Arm angle trajectories over time of the best self/nonself discrimination agent catching objects at the midline from initial hand positions at either the left or right edge of the visual field. The trajectories are shaded according to the initial angular position of the object as indicated at the top of the plot.

agent merge into one halfway between them) and the performance is almost identical. This suggests that the agent uses its arm angle sensors to discount the presence of its own hand in its field of view. Although the detailed arm trajectories of the other top agents varied a great deal, they all moved the hand back and forth within the field of view, and they all exhibited a similar pattern of sensitivity to arm angle sensor lesions and insensitivity to hand transparency.

5. Short-Term Memory

A minimally cognitive agent must be able to transcend its immediate environment by allowing past experiences to influence its future actions. Indeed, it has been argued that only agents that can coordinate their behavior with environmental features that are not immediately present are sufficiently "representation hungry" to be of cognitive interest (Clark, 1997). In previous work, we evolved agents that could catch objects falling vertically after only briefly observing their position and then moving to the correct location while blind to the object's subsequent motion (Gallagher & Beer, 1999). Here we probe the strategies that evolve for objects falling vertically, and extend this work to objects exhibiting horizontal motion as well.

In our first set of short-term memory experiments, agents of diameter 30 had 9 proximity sensors of maximum length 205 uniformly distributed over a visual angle of $\pi/6$ (Figure 8). Their horizontal velocity was proportional to the sum of opposing forces produced by a bilateral pair of effectors (with a constant of proportionality of 5). Circular objects of diameter 26 dropped from above with a vertical velocity of 2 and an initial horizontal offset of ± 50. As soon as an agent began to move, the input to all proximity sensors was permanently set to 0, so that the agent's subsequent behavior could only depend on observations of the object collected before movement began.

The CTRNN architecture was bilaterally symmetric, with 9 sensory neurons, 4 fully interconnected interneurons and 2 fully interconnected motor neurons (for a total of 56 parameters). The sensory neurons projected to both the interneurons and the motor neurons and the interneurons and motor neurons were fully interconnected. All proximity sensors shared a single gain and bias, interneuron and motor neuron biases were in the range [-5,5], and motor neurons had gains fixed to 1. Gains were clipped to be greater than 0 and time constants were clipped to be greater than 1. Populations of 100 individuals were evolved for 500 generations with a mutation variance σ^2 of 0.4, a crossover probability of 0 and an elitist fraction of 2%.

The performance measure to be maximized was:

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

where d_i is the final horizontal separation between the center of the agent and the center of the object at the end of the i^{th} trial. Twelve evaluation trials were used, evenly spaced over the range [0,55].

Out of a total of 5 runs, all produced agents that achieved an average performance greater than 97% on all 12 trials. The best agent had a mean fitness of 99% on the 12 evaluation trials and 99.3% on 1000 random trials. The accuracy of this agent in catching objects as a function of their horizontal position is shown in Figure 9. Despite its blindness, the agent's accuracy is nearly perfect except for small errors around the midline and at the periphery of the visual field.

The behavior of this agent is shown in Figure 10. This agent waits in place until the object intersects one of its outermost rays. After a short delay, the agent begins to move in the direction of the intersection with an average velocity designed to bring it to the object's horizontal position at the time the object reaches the agent. The agent's motor response has two components: an initial transient

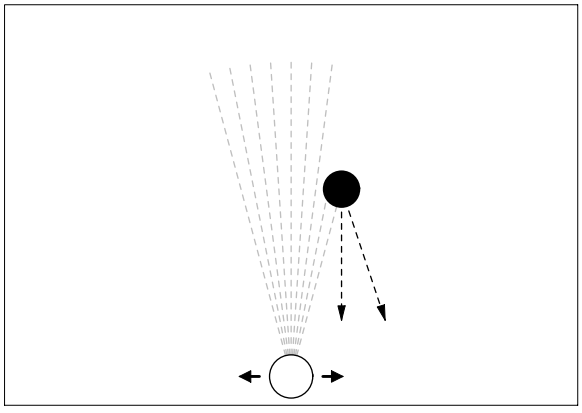

Figure 8: Experimental setup for short-term memory experiments. The agent can move horizontally while objects fall either vertically or diagonally from above. The rays are dashed because, as soon as the agent begins to move, it goes blind.

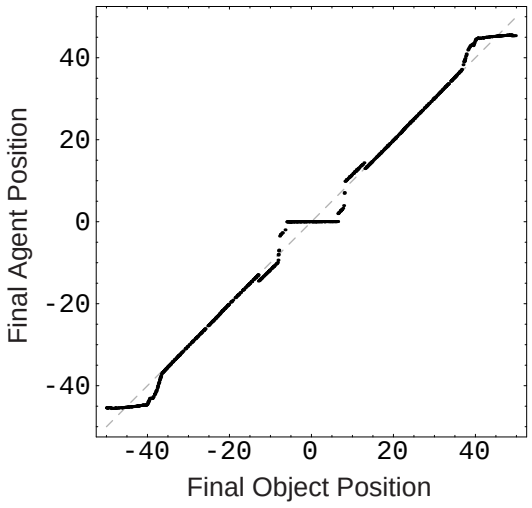

Figure 9: Accuracy of the best short-term memory agent for vertically falling objects. The final horizontal position of the agent is plotted against the final horizontal position of the object. Note that the average performance closely approximates the ideal (dashed gray line) except at the midline and periphery.

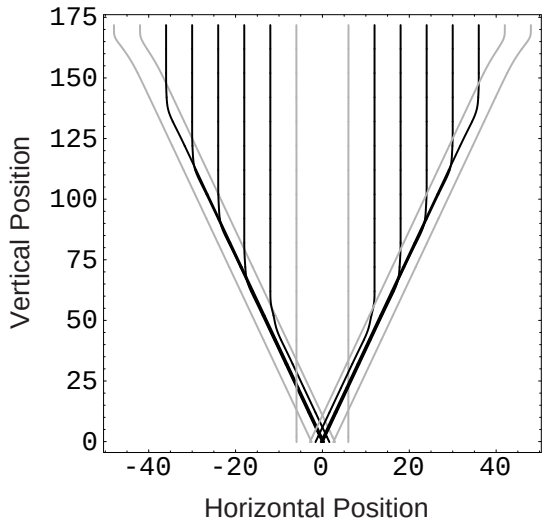

Figure 10: Behavior of the best short-term memory agent for vertically falling objects. Trajectories of motion relative to the agent of objects falling vertically from several different initial horizontal offsets are shown. Gray trajectories are those for which the agent's strategy begins to break down.

phase and a final constant phase. The steady-state horizontal velocity achieved in the constant phase was 0.53, which is very close to the required horizontal velocity predicted by the vertical velocity of the object and the angle of the outermost ray: $2/(\tan 5\pi/12) = 0.54$. The transient phase appears to be designed to correct for the fact that a falling circle will intersect the outermost ray not at the circle's leading edge, but to one side. Thus, the object's actual horizontal position is somewhat more central than the point of intersection with the ray would indicate. Since the agent can no longer see the object once it begins to move, both the transient and steady-state velocity produced by the agent is an internally-

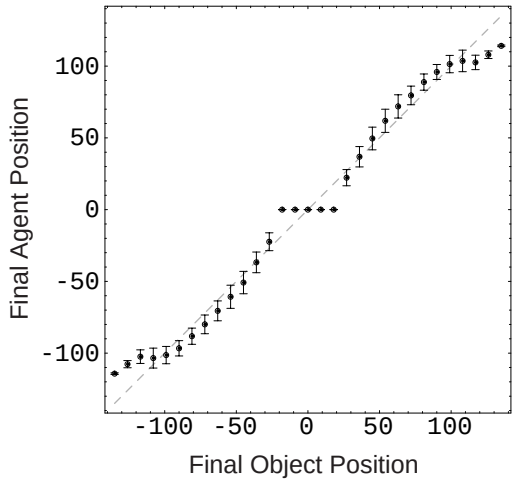

Figure 11: Mean accuracy of the best short-term memory agent for diagonally falling objects. The final horizontal position of the agent is plotted against the final horizontal position of the object (mean ± s.d., $N = 200$ trials). Note that the average performance closely approximates the ideal (dashed gray line).

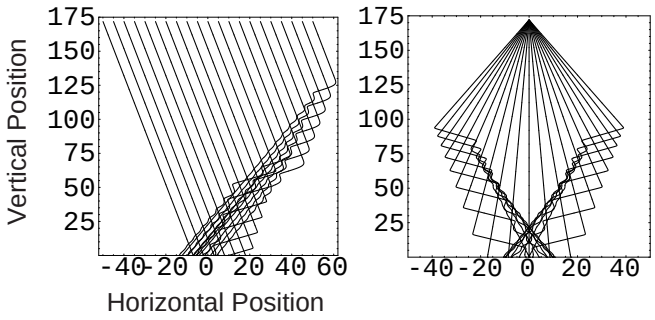

Figure 12: Behavior of the best short-term memory agent for diagonally falling objects. (Left) Trajectories of motion relative to the agent of objects falling diagonally with a horizontal velocity of 0.5 from several different initial horizontal positions are shown. (Right) Trajectories of motion relative to the agent of objects falling diagonally from the midline with several different horizontal velocities are shown.

generated, temporally-extended response to the initial perturbation of the object intersecting an outermost ray.

From Figures 9 and 10, it is clear that this strategy begins to break down near the midline and near the periphery. An object falling near the midline will not intersect the outermost ray until it has nearly reached the agent, allowing insufficient time for the agent to respond. Fortunately, objects falling near the midline require little movement to catch. A sufficiently peripheral object will immediately intersect the outermost ray, forcing the agent to treat all such cases identically. Three of the top five agents use a similar strategy, whereas the other two use different strategies.

In our second set of short-term memory experiments, we evolved agents that could catch diagonally-moving objects despite going blind during movement. Note that the strategy described above for vertical objects will not work here because objects with different horizontal velocities can intersect a peripheral ray at the same time and vertical distance. The experimental setup for our second set of experiments was identical to that for the first except for the following differences. Objects now fell diagonally with a horizontal velocity in the range ± 1. The CTRNN architecture now has 6 interneurons (for a total of 82 parameters). Experiments were now run for 5000 generations with a mutation variance σ^2 of 0.5. The performance measure was the same as for vertical objects, except that a 20 point penalty was assessed if the agent began moving within the first 5 time steps of the trial. Twenty-eight evaluation trials were used, with the initial horizontal position in the range [0,50] and the horizontal velocity in the range ± 1.

Out of a total of 5 runs, 4 produced agents that achieved an average performance greater than 94% on all 28 trials. The best agent had a mean fitness of 96.2% on the 28 evaluation trials and 95.7% on 1000 random trials. The accuracy of this agent is shown in Figure 11. Despite the horizontal motion of the objects and despite the fact that the agent goes blind when it moves, this agent is quite accurate except for small errors at the midline and periphery. Its be-

havior is shown in Figure 12. It begins to move about the time that the object loses contact with the outermost ray, with an oscillating velocity that slowly decays to a constant value. Generally speaking, the initial peak of the oscillation varies systematically with the object's horizontal velocity, with larger object velocities producing larger peak agent velocities. This would cause the agent to move a greater horizontal distance for objects with larger horizontal velocities. Thus, it appears that the diagonal short-term memory agents use a variation of the strategy employed by the vertical short-term memory agents, but with a more complex transient structure to account for the horizontal velocities of the objects. All of the top agents used a similar strategy.

6. Selective Attention

A complex environment often contains many more objects than an agent can simultaneously interact with. This requires a minimally cognitive agent to be able to focus its attention on one object while ignoring others. Indeed, attentional mechanisms are fundamental components of many other cognitive systems (Posner, 1995). In a final set of experiments, we evolved agents that could catch two objects moving at different horizontal and vertical velocities.

Agents of diameter 30 had 9 proximity sensors of maximum length 205 uniformly distributed over a visual angle of $\pi/6$ (Figure 13). Their horizontal velocity was proportional to the sum of opposing forces produced by a bilateral pair of effectors (with a constant of proportionality of 5). Two circular objects of diameter 26 dropped from above. One of the objects had a vertical velocity in the range [3,4] and the other had a vertical velocity in the range [1,2]. The horizontal velocities of the objects were in the range ±2. The velocities and initial positions of the two objects were constrained such that $|x_1 - x_2| / |t_1 - t_2| \leq 5\alpha$, where x and t represent the final horizontal positions and times of impact of the two objects, 5 is the maximum horizontal velocity of the agent and $\alpha = 0.7$ was chosen to ensure that the agent had a reasonable chance of reaching the second object after catching the first one.

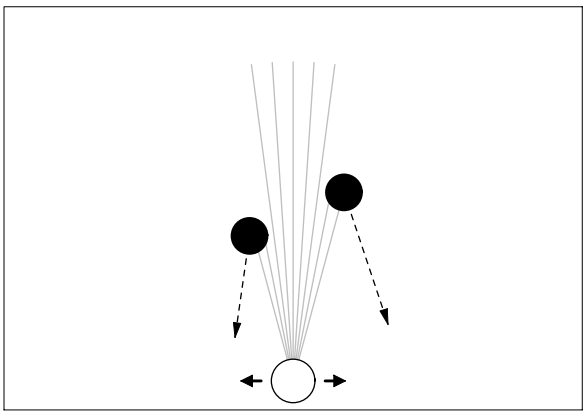

Figure 13: Experimental setup for selective attention experiments. The agent can move horizontally while 2 objects fall from above. As described in the text, the initial positions and velocities of the two objects are constrained so that the agent has a reasonable chance of catching both.

The CTRNN architecture was bilaterally symmetric, with 9 sensory neurons, 10 fully interconnected interneurons, and 2 fully interconnected motor neurons (for a total of 146 parameters). The sensory neurons projected to both the interneurons and the motor neurons and the interneurons and motor neurons were fully interconnected. All proximity sensors shared a single gain and bias, interneuron and motor neuron biases were in the range [-5,5], and motor neurons had gains fixed to 1. Gains were clipped to be greater than 0 and time constants were clipped to be greater than 1. Populations of 100 individuals were evolved for 9000 generations, with a mutation variance σ^2 of 1, a crossover probability of 0 and an elitist fraction of 2%.

The performance measure to be maximized was:

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

where $p_i = |d_{i,1}| + |d_{i,2}|$ and $d_{i,1}$ and $d_{i,2}$ are the final horizontal separations between the center of the agent and the center of the first and second objects on the i^{th} trial.

The total number of trials grew from 1 to 35 in the course of an evolution. A new trial was added every time the best performance exceeded a threshold T or 600 generations passed without the addition of a new trial. Here $T = 198 - n/14 - gen/2500$, where n is the current number of evaluation trials and gen is the current generation number. These trials were chosen to include as many difficult cases as possible. Once the number of trials reached 35, new trials were generated randomly until one was found for which the current best agent scored lower than 170. This new trial then replaced the agent's current highest scoring trial of the 35. This process continued until the search was terminated.

This is the most difficult problem that we have attempted to date. First and foremost, the agent must somehow avoid being distracted by one object while it is orienting to another. If it simply orients to the average position of the two objects, it will miss both in general. Second, which object to attend to first is not always obvious from the outset because an object that is initially farther away can still reach the agent first if it is falling faster. We will call this the *passing objects* (PO) problem. Finally, in the course of catching one object, the other object can pass entirely out of the agent's field of view, raising an *object permanence* (OP) problem. Furthermore, this can occur before a faster but more distant object passes a closer but slower one, forcing the agent to choose which to pursue based only on a prediction of which object will reach it first. Thus, a successful agent must be able to partially decouple its behavior from its immediate circumstances while still remaining sensitive to them. It must also be capable of making and remembering predictions about the future configuration of objects based on observations of their past motion.

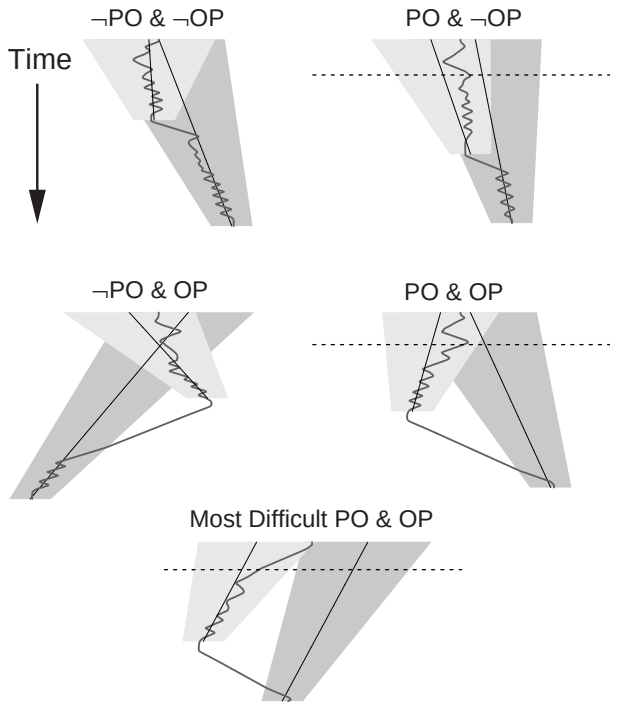

Figure 14: Behavior of the best selective attention agent. Example motion plots are shown for different combinations of passing objects (PO) and object permanence (OP). In addition, an especially difficult PO&OP case, in which one object disappears from view before the more distant but faster falling object passes it, is shown at bottom. Each plot shows the horizontal positions of the two objects (straight black lines) and the agent (gray line) over time. The shaded regions correspond to positions and times in which the faster-falling circle (light gray) and the slower-falling circle (darker gray) can be seen by the agent. The dashed lines indicate the time at which the first object overtakes the second in passing objects cases.

Out of a total of 7 runs, 3 produced agents that scored higher than 90% on 1000 random trials. The best agent had a mean performance of 94.2% on the random trials, with the following breakdown: 97.5% on trials involving neither passing objects nor object permanence (227/1000 trials), 97.28% on trials involving only passing objects (152/1000), 93.64% on trials involving only object permanence (329/1000), and 90.72% on trials involving both passing objects and object permanence (292/1000). Clearly, trials involving object permanence were difficult, and those involving both object permanence and passing objects were even more difficult. The most difficult cases of all were those in which one object disappeared from the field of view before the more distant but faster-falling object passed the closer but more slowly-falling object (Figure 14, bottom). There were no such cases in the first 1000 random trials, and only 8 such cases in 10,000 random trials. On these 8 cases, the best agent had a mean performance of 79.8%. The other top agents exhibited a similar pattern of difficulty.

The general behavior of the best agent in each of the four cases, as well as the especially difficult PO&OP case, is illustrated in Figure 14. In all cases, the agent tries to keep both objects in view as long as possible by making large sweeps back and forth. As objects fall, the agent eventually tightens its scan on the object it will catch first. Note that the agent is very decisive in all cases except the PO&¬OP case. Here, the objects are fairly close together and on converging paths, so they probably form one contiguous object in the agent's field of view. This seems to initially confuse the agent, since its tight scan only slowly drifts toward the closer object. The other two top agents produced similar behavior, but they varied in the specific pattern of scanning they used and how fast they moved. However, detailed analyses of the strategies employed remain to be done.

7. Conclusion

In this paper, we have extended our previous work on the evolution of minimally cognitive behavior to a significantly wider range of tasks, including the perception of body-scaled affordances, self/nonself discrimination, short-term memory and selective attention. Our results demonstrate that CTRNNs can be evolved for a wide range of cognitively interesting behavior using a relatively simple evolutionary algorithm. As we have attempted more difficult tasks, shaping by incrementally adding new test cases as evolution progresses has become an essential part of our methodology. It is also interesting to note that active scanning is frequently observed in the successful agents that we have evolved (the short-term memory agents are an exception because scanning is impossible in that case). While we believe that active perception is likely to be a common feature of distal sensing in situated, embodied agents, it probably arises in our experiments due to the relatively coarse spatial resolution of our "visual" sensors.

As we move toward the evolution of increasingly more cognitive behavior, the most interesting challenge we face is understanding how the evolved CTRNNs work. While we have made substantial progress on the analysis of evolved neural circuits for sensorimotor control (Beer et al., 1999), the analysis of evolved CTRNNs for more sophisticated behavior still poses a significant challenge. As internal state mediates between perception and action in increasingly more sophisticated ways, the agent's behavior can become increasingly decoupled from its immediate circumstances while still remaining sensitive to them. Given that there has been some skepticism regarding how well intuitions grounded in the dynamics of situated action will carry over to more cognitive behavior (Clark, 1997), it will be very interesting to see how evolution shapes this internal dynamics in order to accomplish minimally cognitive tasks. For this reason, the analysis of evolved CTRNNs for minimally cognitive behavior is a major focus of ongoing work.

Acknowledgments

We would like to thank Alan Calvitti and Hillel Chiel for their feedback on an earlier draft of this paper. This work

was supported in part by grant RG0084/1997-B from the Human Frontier Science Program.

References

Almássy, N., Edelman, G.M. and Sporns, O. (1998). Behavioral constraints in the development of neuronal properties: A cortical model embedded in a real-world device. *Cerebral Cortex* **8**:346-361.

Beer, R.D. (2000). Dynamical approaches to cognitive science. To appear in *Trends in Cognitive Sciences*.

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

Beer, R.D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In P. Maes et al. (Eds.), *From Animals to Animats 4: Proc. 4th Int. Conf. on Simulation of Adaptive Behavior* (pp. 421-429). MIT Press.

Beer, R.D., Chiel, H.J. and Gallagher, J.C. (1999). Evolution and analysis of model CPGs for walking: II. General principles and individual variability. *J. Computational Neuroscience* **7**:119-147.

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

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

Cliff, D. & Miller, G.F. (1996). Co-evolution of pursuit and evasion II: Simulation methods and results. In P. Maes et al. (Eds.), *From Animals to Animats 4: Proc. 3rd Int. Conf. on Simulation of Adaptive Behavior* (pp. 506-515). MIT Press.

Di Paolo, E.A. (1997). An investigation into the evolution of communication. *Adaptive Behavior* **6**:285-324.

Gallagher, J.C. and Beer, R.D. (1999). Evolution and analysis of dynamical neural networks for agents integrating vision, locomotion, and short-term memory. In W. Banzhaf et al. (Eds.), *Proc. Genetic and Evolutionary Computation Conf.* (pp. 1273-1280).

Gibson, J.J. (1979). *The Ecological Approach to Visual Perception*. Houghton Mifflin.

Harvey, I., Husbands, P. & Cliff, D. (1994). Seeing the light: Artificial evolution, real vision. In D. Cliff et al. (Eds.), *From Animals to Animats 3: Proc. 3rd Int. Conf. on Simulation of Adaptive Behavior* (pp. 392-401). MIT Press.

Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. 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.

Neisser, U. (1993). The self perceived. In U. Neisser (Ed.), *The Perceived Self: Ecological and Interpersonal Sources of Self-Knowledge* (pp. 3-21). Cambridge University Press.

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

Pfeifer, R. and Scheier, C. (1999). *Understanding Intelligence*. MIT Press.

Posner, M.I. (1995). Attention in cognitive neuroscience: An overview. In M.S. Gazzaniga (Ed.), *The Cognitive Neurosciences* (pp. 615-624). MIT Press.

Warren, W.H. Jr. and Wang, S. (1984). Visual guidance of walking through apertures: Body-scaled information for affordances. *J. Experimental Psychology: Human Perception and Performance* **13**:371-383.

Appendix

Choosing the mutation variance σ^2 appropriate to a given problem representation and fitness function is nontrivial. In lieu of a detailed understanding of the statistical structure of the search space, it would at least be useful to know the average change to each parameter produced by a mutation. In this appendix, we derive an expression for this expected magnitude of change for our evolutionary algorithm.

Let **p** be the random perturbation vector added to an individual parameter vector **x** of size N by mutation. Then we seek $E\{|p_i|\}$, the expected magnitude of any component p_i of **p**. Due to the spherical symmetry of mutation, this expectation is identical for all components, so we will focus on the first component p_1 without loss of generality. As described in Methods above, the magnitude m of **p** is drawn from a normal distribution with 0 mean and variance σ^2, while the direction of **p** is given by a unit vector **u** uniformly distributed over the unit hypersphere S^{N-1}. Because these two distributions are independent:

$$E\{|p_1|\} = E\{|m|\}\, E\{|u_1|\}$$

If $m \sim N(0, \sigma^2)$, then

$$E\{|m|\} = \int_{-\infty}^{\infty} |x| \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-x^2}{2\sigma^2}} dx = \sqrt{\frac{2\sigma^2}{\pi}}$$

In order to motivate our derivation of $E\{|u_1|\}$, we will first consider the special cases $N = 2$ and $N = 3$ before moving to general N. For $N = 2$, we want the average magnitude of the projection of **u** onto x_1 (Figure A1, left). At any angle ϕ from x_1, there are two unit vectors whose projection onto x_1 is cos ϕ. If we sum up the magnitudes of all such projec-

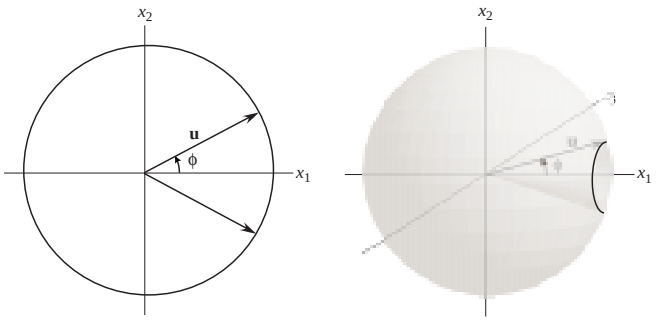

Figure A1: For $N = 2$ (left), there are two vectors at an angle ϕ from x_1 whose projection onto x_1 is cos ϕ. For $N = 3$ (right), there is a circle of vectors at an angle ϕ from x_1 whose projection onto x_1 is cos ϕ. The radius of this circle is sin ϕ.

tions for the entire unit circle and divide by the total "number" of vectors (given by the circumference of the unit circle), we obtain the desired average:

$$E\{|u_1|\} = \frac{2\int_0^{\pi/2} 2\cos\phi\, d\phi}{2\pi} = \frac{2}{\pi}$$

Note that, in order to account for the contribution of the vectors lying on the left semicircle (whose projections are –cos ϕ), we have merely doubled the result for the right semicircle.

For $N = 3$, we find a *circle* of vectors at angle ϕ from x_1 whose projection onto x_1 is cos ϕ (Figure A1, right). The radius of this circle of vectors is sin ϕ and the number of such vectors is just the circumference of this circle or 2π sin ϕ. Thus, the desired average can be found by summing up the magnitudes of the projections of all such circles of vectors for the entire unit sphere and dividing by the total "number" of vectors (given by the surface area of the unit sphere):

$$E\{|u_1|\} = \frac{2\int_0^{\pi/2} 2\pi\sin\phi\cos\phi\, d\phi}{4\pi} = \frac{1}{2}$$

For general N, we will find an (N-2)–sphere of vectors at angle ϕ from x_1, all of whose projections onto x_1 are cos ϕ. The radius of this (N-2)–sphere will be sin ϕ and the number of such vectors will just be the "surface area" of this (N-2)–sphere. The desired average can be found by summing up the magnitudes of the projections of all such (N-2)–spheres of vectors for the entire unit hypersphere S^{N-1} and dividing by the total "number" of vectors (given by the "surface area" of S^{N-1}):

$$E\{|u_1|\} = \frac{2\int_0^{\pi/2} A(N-1)\sin^{N-2}\phi\cos\phi\, d\phi}{A(N)}$$

$$= \frac{2\,A(N-1)}{A(N)}\frac{1}{N-1}$$

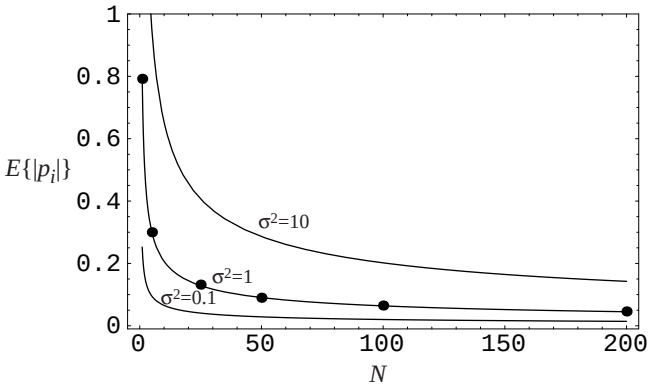

Figure A2: Variation of $E\{|p_i|\}$ with N for several different values of σ^2. Data from our evolutionary algorithm is plotted for the case $\sigma^2 = 1$ (mean of 5000 trials).

where $A(N)$ is the "surface area" of the unit hypersphere S^{N-1}:

$$A(N) = \frac{N\pi^{N/2}}{\Gamma\left(\frac{N}{2}+1\right)}$$

and Γ is Euler's gamma function. Then we have that

$$E\{|u_1|\} = \frac{\Gamma\left(\frac{N}{2}\right)}{\sqrt{\pi}\,\Gamma\left(\frac{N+1}{2}\right)}$$

and thus $E\{|p_1|\} = E\{|m|\}\,E\{|u_1|\} = \dfrac{\sqrt{2\sigma^2}\,\Gamma\left(\frac{N}{2}\right)}{\pi\,\Gamma\left(\frac{N+1}{2}\right)}$

The variation of $E\{|p_i|\}$ with N for representative values of σ^2 is shown in Figure A2, along with data from our evolutionary algorithm for the case $\sigma^2 = 1$, showing that this expression provides an accurate fit.

If $\mathcal{L}_i$ is the linear map from a search parameter p_i in the range ±1 to a CTRNN parameter c_i, then the expected magnitude of change $E\{|\Delta c_i|\}$ in c_i would be given by $\mathcal{L}(E\{|p_i|\})$. For example, for a connection weight in the range ±5, $E\{|\Delta c_i|\} = 5\,E\{|p_i|\}$. Putting this all together, if a given expected magnitude of change $E\{|\Delta c_i|\}$ is desired in a CTRNN parameter c_i, then the mutation variance σ^2 should be set as:

$$\sigma^2 = \left(\mathcal{L}_i^{-1}\left(E\{|\Delta c_i|\}\right)\frac{\pi\,\Gamma\left(\frac{N+1}{2}\right)}{\sqrt{2}\,\Gamma\left(\frac{N}{2}\right)}\right)^2$$

Of course, in general the desired expected change and the linear map will be different for different classes of CTRNN parameters. In this case, some compromise between the different required σ^2 must be selected.