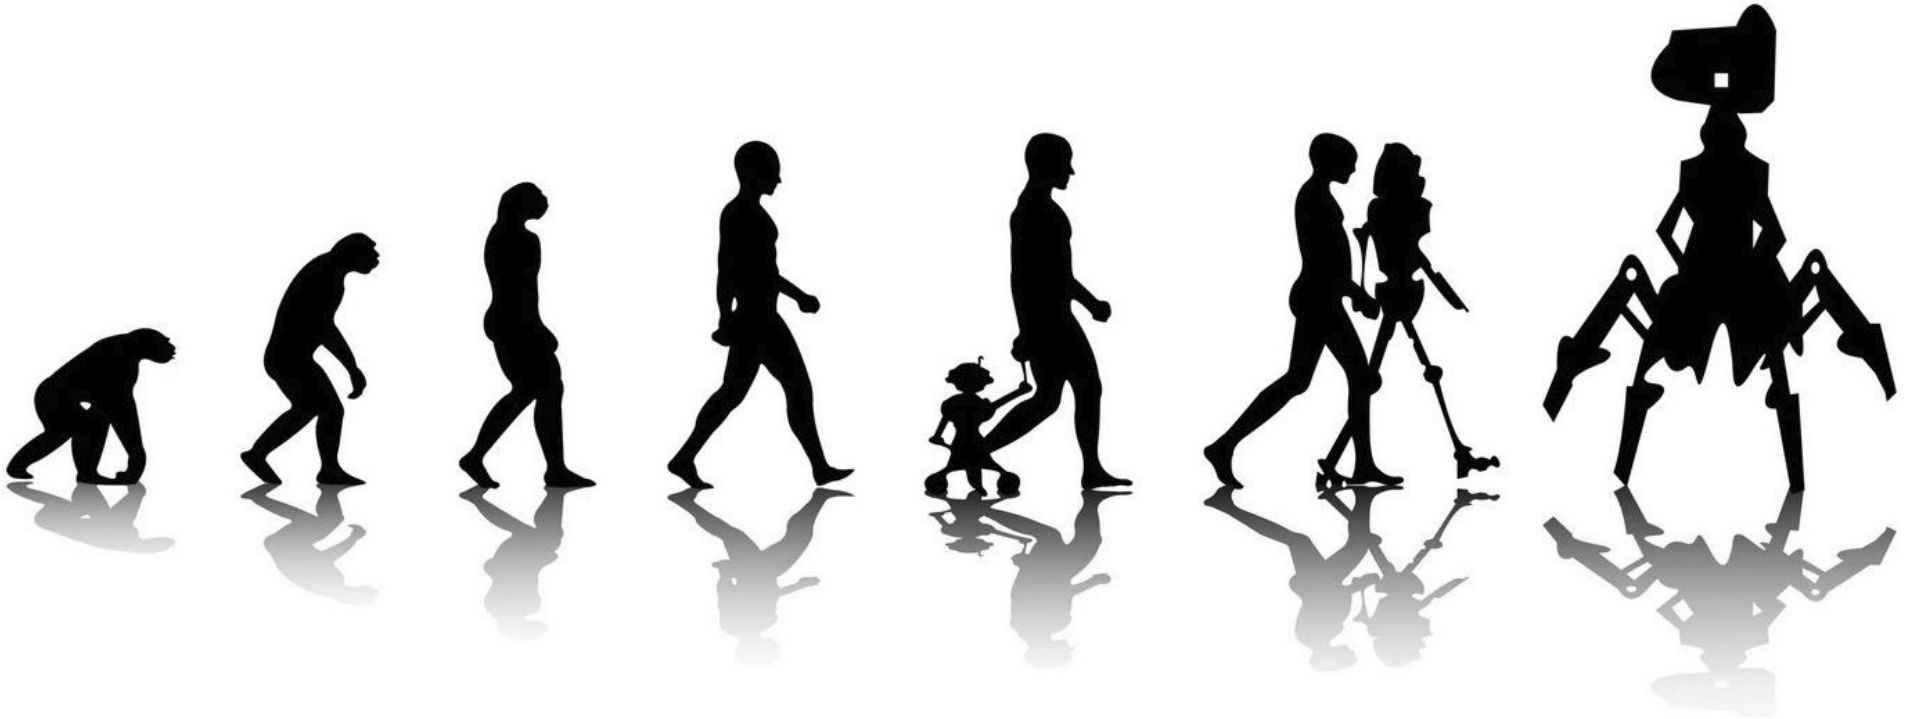

Modern Robotics: Evolutionary Robotics

COSC 4560 / COSC 5560

Professor Cheney

2/9/18

Minimally Cognitive Systems

We said last class that adaptive behavior is built around sensory-motor actions

Just how simple of a sensory motor system do we need to get interesting “cognitive” behavior?

Is it simple enough that we can still understand and interpret the underlying brain/behavior?? (like we did for the Braitenberg vehicles)

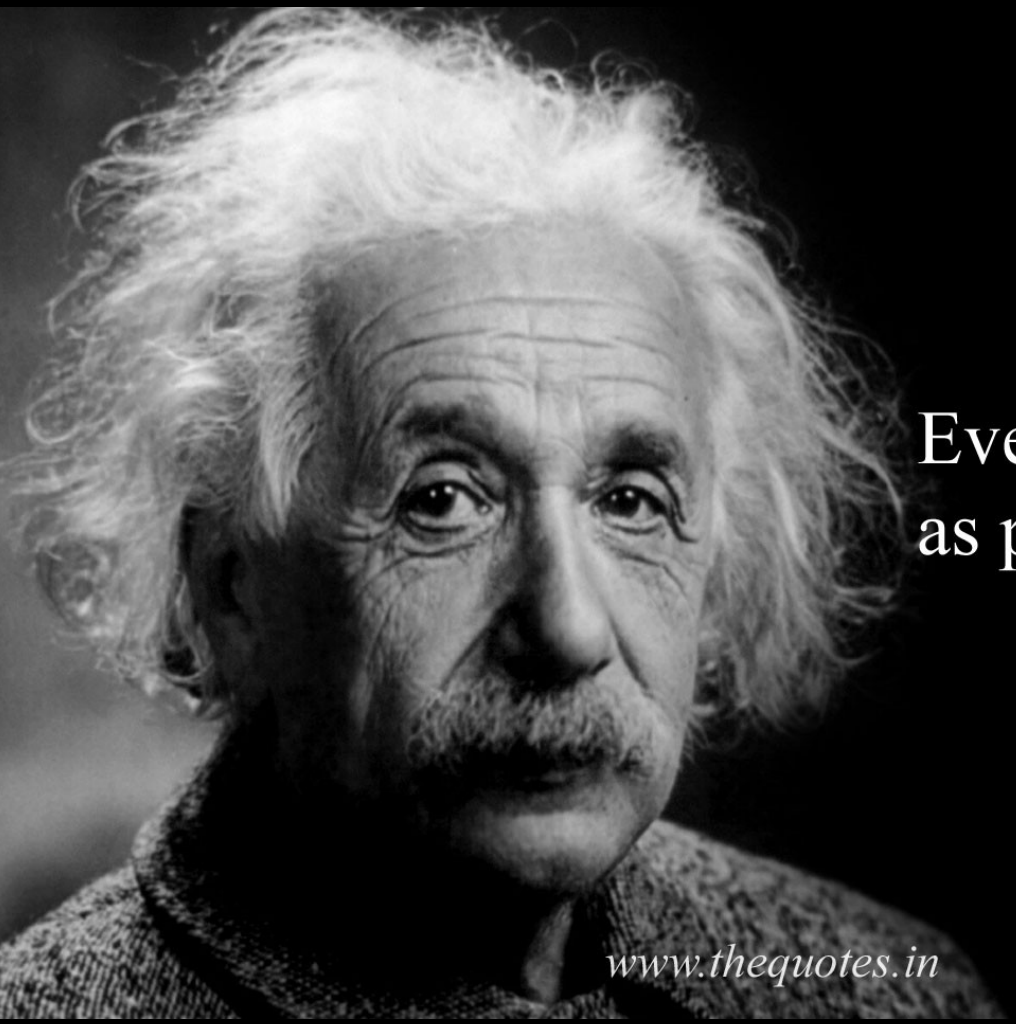

Everything should be made as simple as possible, but not simpler.

Albert Einstein


ALDEBARAN ROBOTICS
11
NAO
PR2

What is the simplest setup
(e.g. task/environment/agent)
to evolve a cognitive agent?

The Dynamics of Brain–Body–Environment Systems: A Status Report

Randall D. Beer

Cognitive Science Program, Department of Computer Science and Department of Informatics, Indiana University, Bloomington, IN, USA

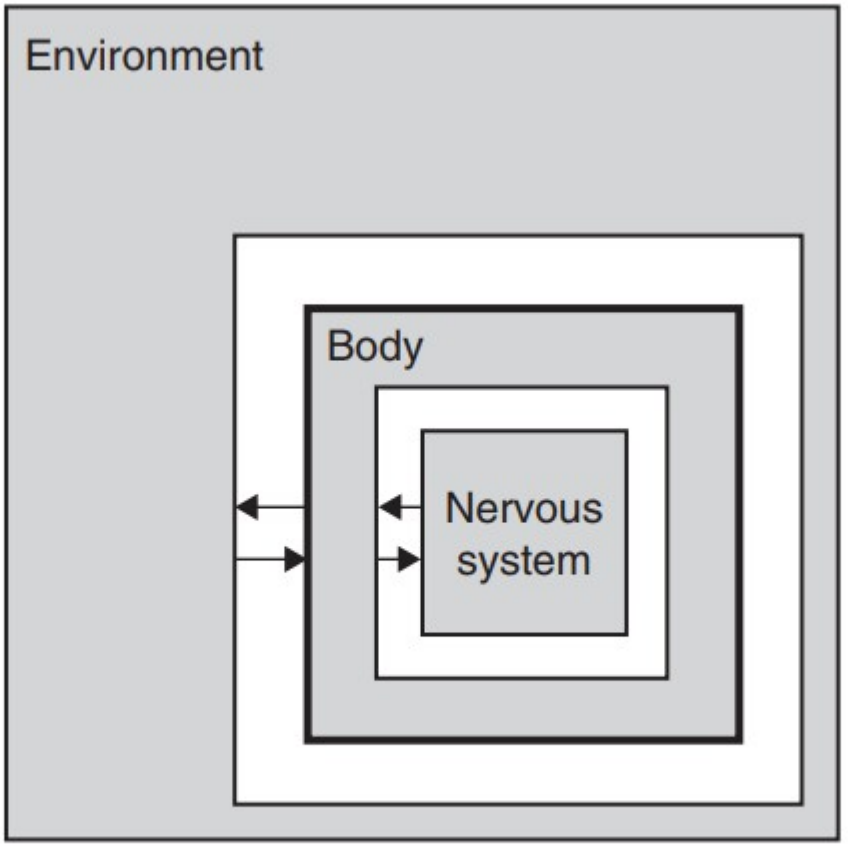

FIGURE 6.1 An agent and its environment as coupled dynamical systems. The agent in turn is composed of coupled nervous system and body dynamical systems.

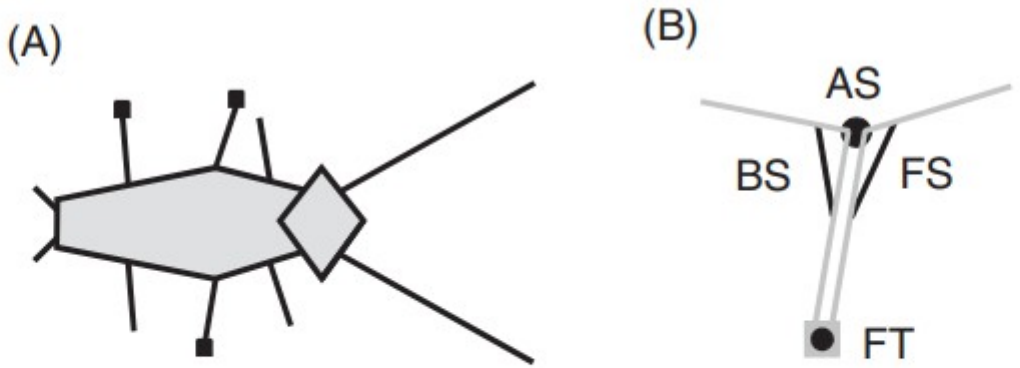

FIGURE 6.3 The walking scenario. (A) The body model. (B) Leg detail. Each leg possesses a binary foot effector (FT) and an antagonistic pair of effectors for swinging the leg: backward swing (BS) and forward swing (FS). In some experiments, an angle sensor (AS) was also utilized.

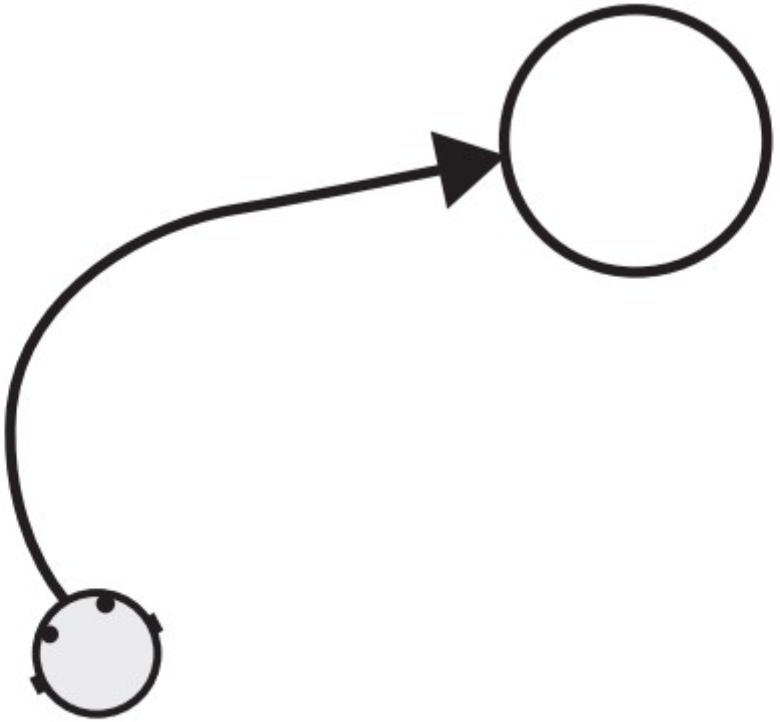

FIGURE 6.2 The chemotaxis scenario. The agent has a bilateral pairs of chemosensors (black disks) and motors (black rectangles). Its task is to navigate to the chemical source whose intensity falls off as the inverse square of distance.

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

CTRNN

Continuous Time
Recurrent Neural Network

Continuous Time Neural Networks

synaptic weights

bias

input

$$\frac{dy}{dt} = \frac{1}{\tau_i}\left(-y_i + \sum_{j=1}^{N} w_{ji}\sigma(g(y_j + \theta_j))\right) + I_i$$

time constant

gain

transfer function

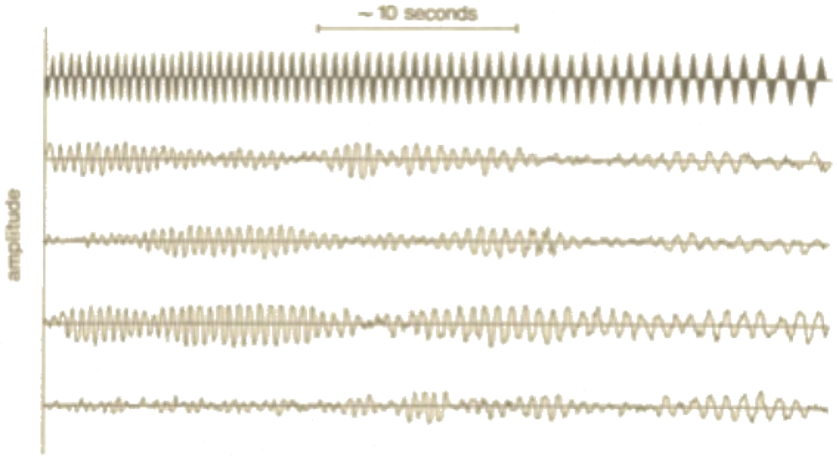

Higher time constant = more stable

Lower time constant = faster reaction

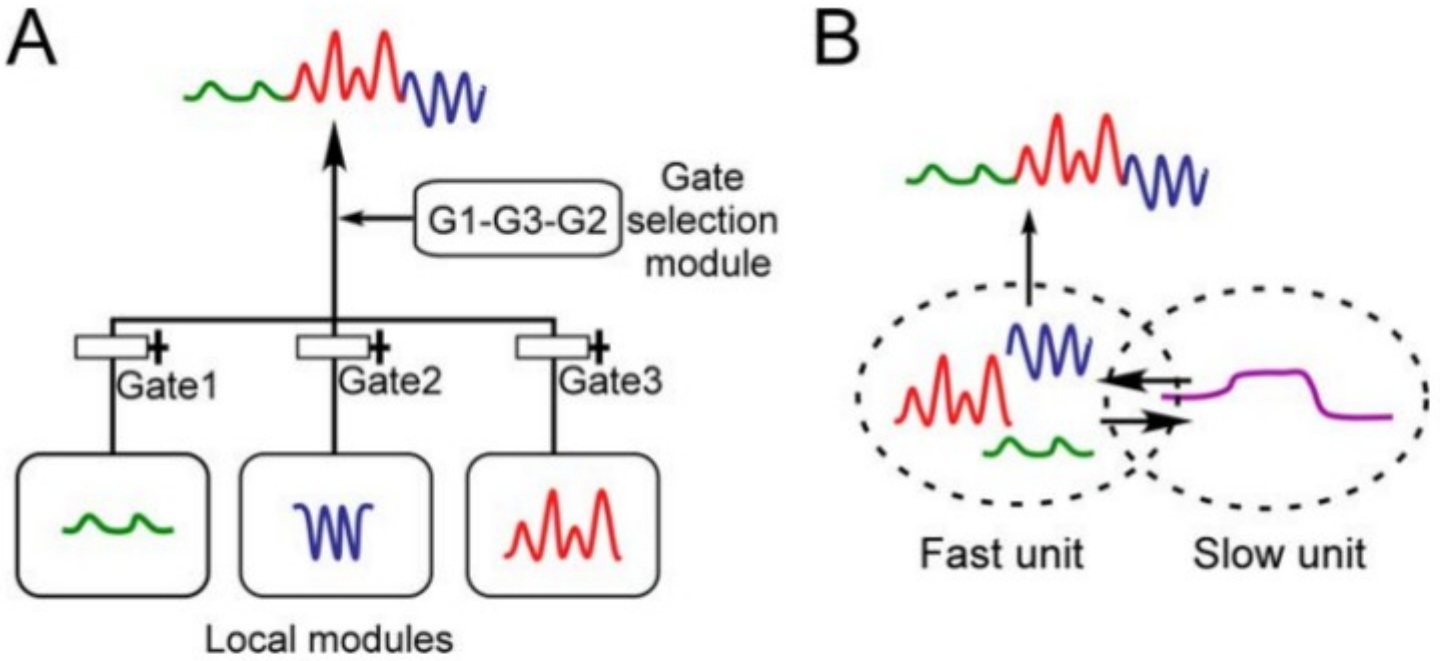

Figure 1. Schematic drawings of (A) local representation model and (B) multiple timescale model. (A) Curves colored red, blue, and green represent sensori-motor sequences corresponding to motor primitives. Output of the system consists of behavior sequences made up of combinations of these primitives. In the local representation model, functional hierarchy is realized through the use of explicit hierarchical structure, with local modules representing motor primitives in the lower level, and a higher module representing the order of motor primitives switched via additional mechanisms such as gate-selection. (B) In the multiple timescale model, primitives are represented by fast context units whose activity changes quickly, whereas sequences of primitives are represented by slow context units whose activity changes slowly.
doi:10.1371/journal.pcbi.1000220.g001

Perceiving Affordances

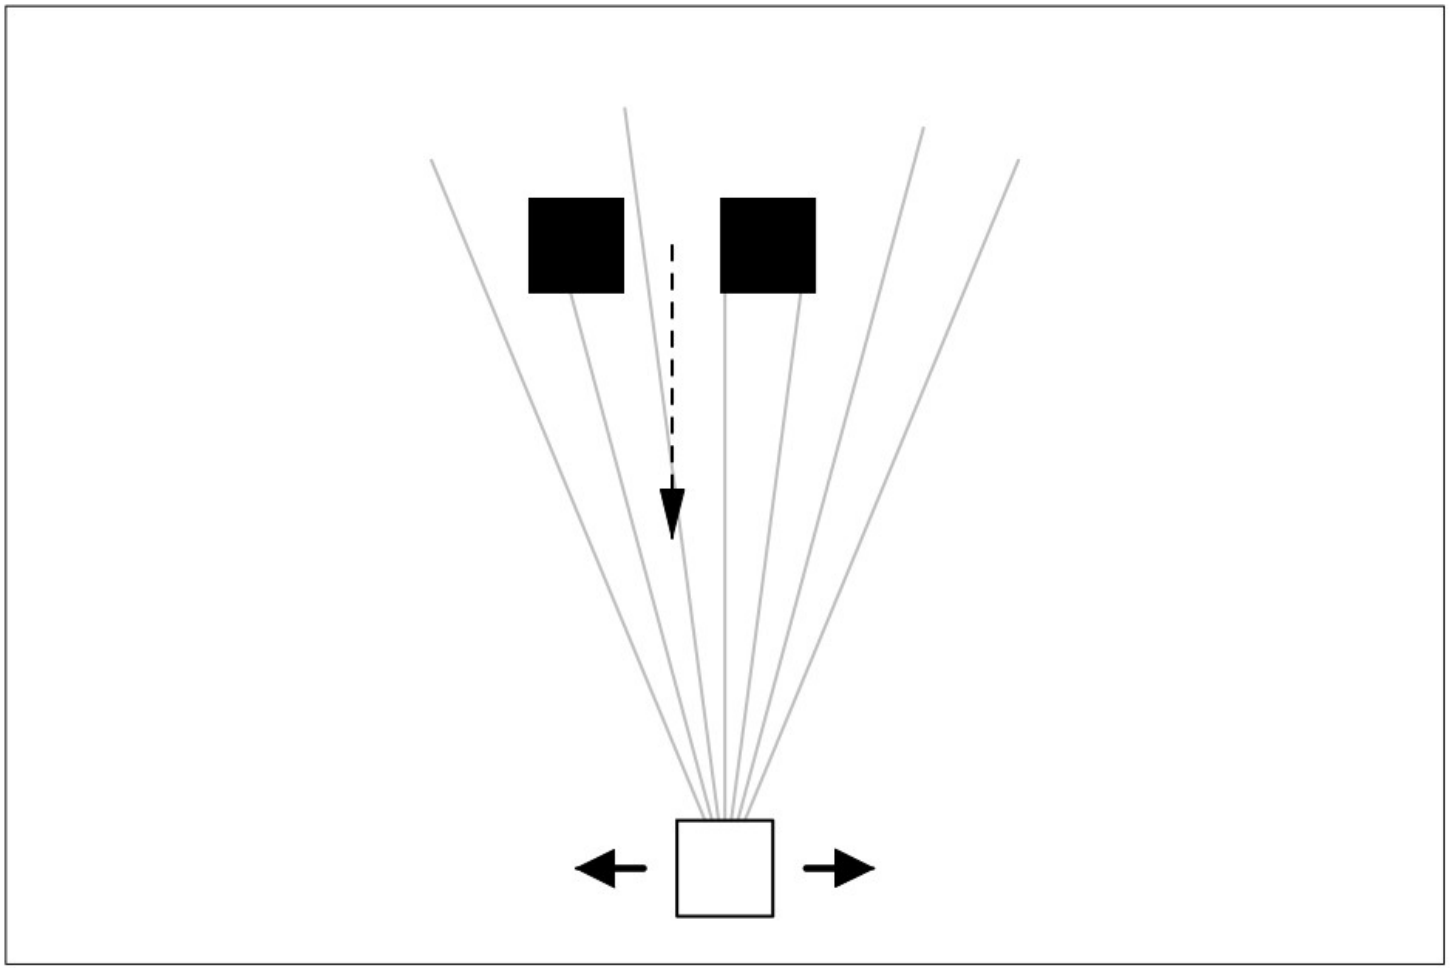

Figure 1: Experimental setup for the passability experiments. The agent moves horizontally while a wall with an adjustable aperture falls from above. The rays of the agent's proximity sensors are shown in gray.

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.

Square agents of size 20 had 7 proximity sensors of maximum length 160 uniformly distributed over a visual angle of $\pi/4$ (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

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

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.

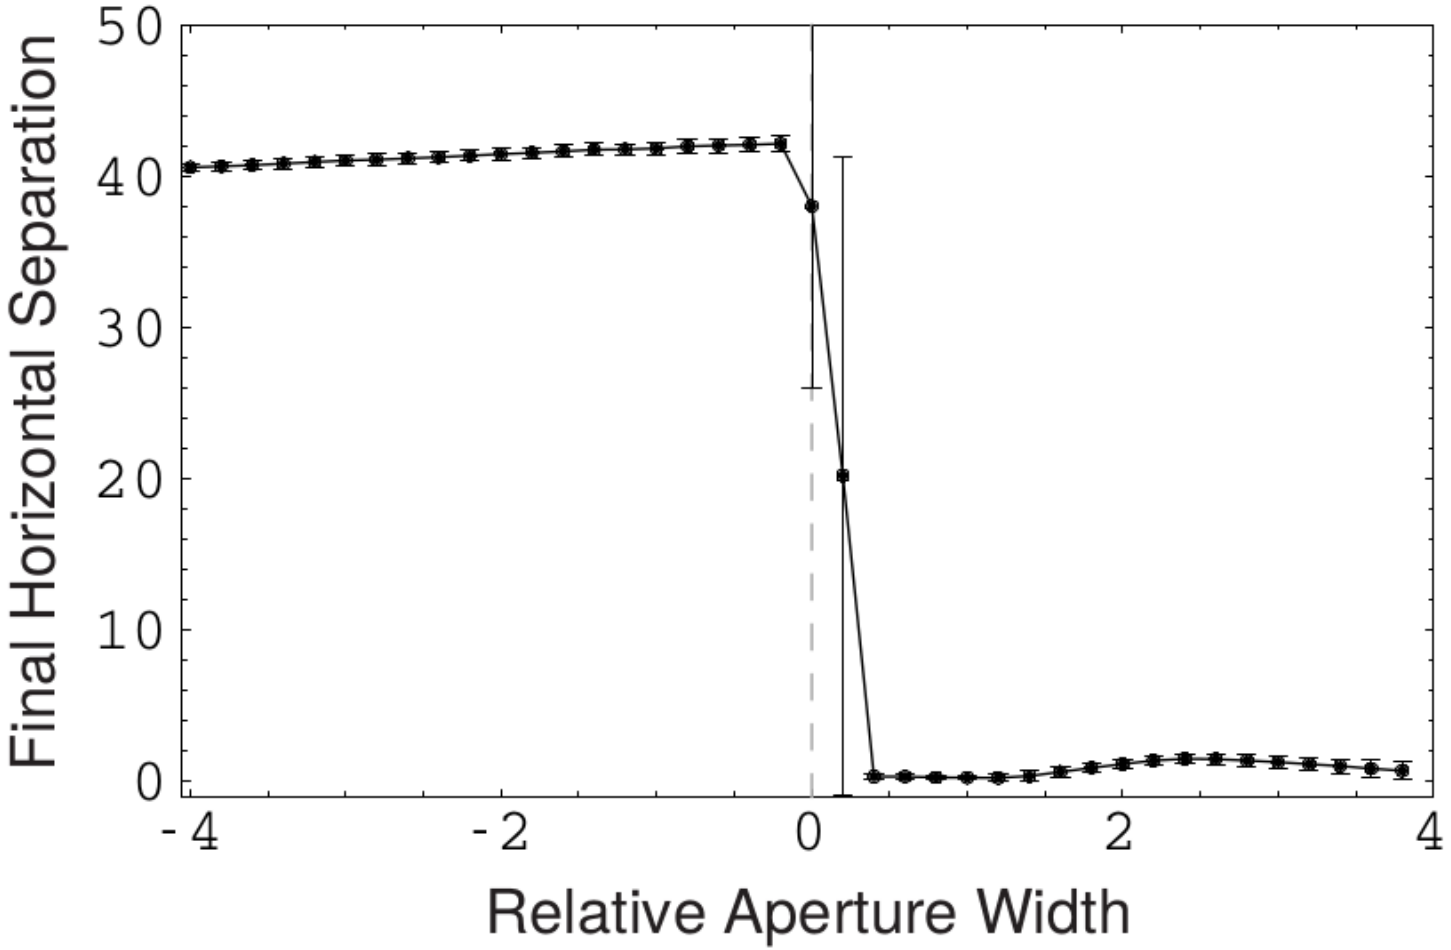

Figure 2: Categorization of apertures into passable and impassable by the best passability agent. The final horizontal separation between the agent and the center of the aperture (mean $\pm$ s.d., $N = 101$ trials) is plotted against the aperture width relative to the agent's size.

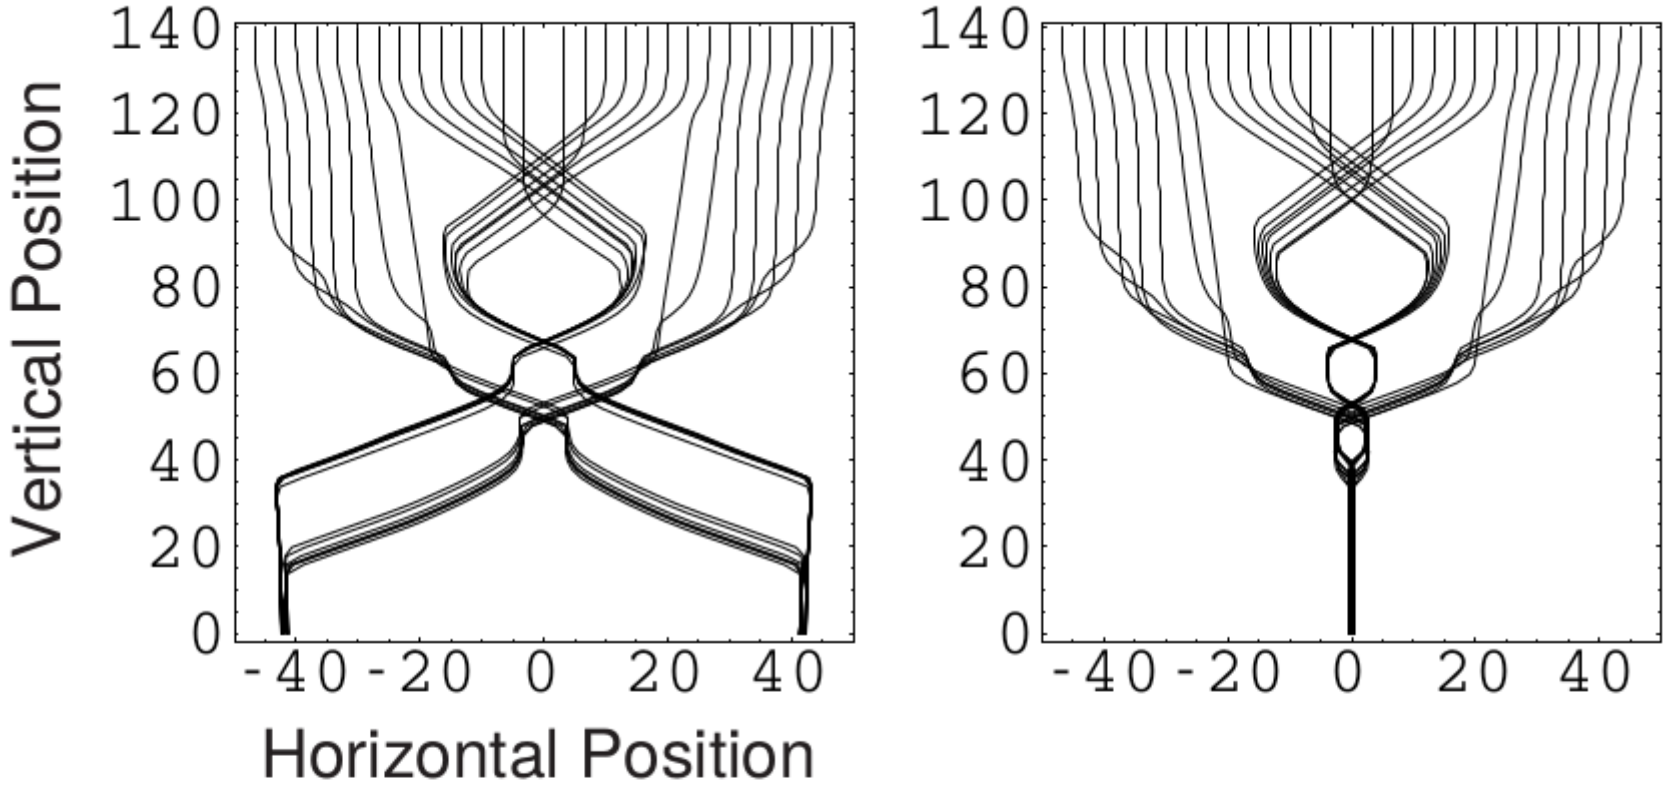

Figure 4: Behavior of the best passability agent. The wall's horizontal and vertical position over time relative to the agent is plotted for an aperture 1 unit smaller than the agent (left) and 1 unit larger than the agent (right). Trials begin at top and time increases from top to bottom.

Self/NonSelf Discrimination

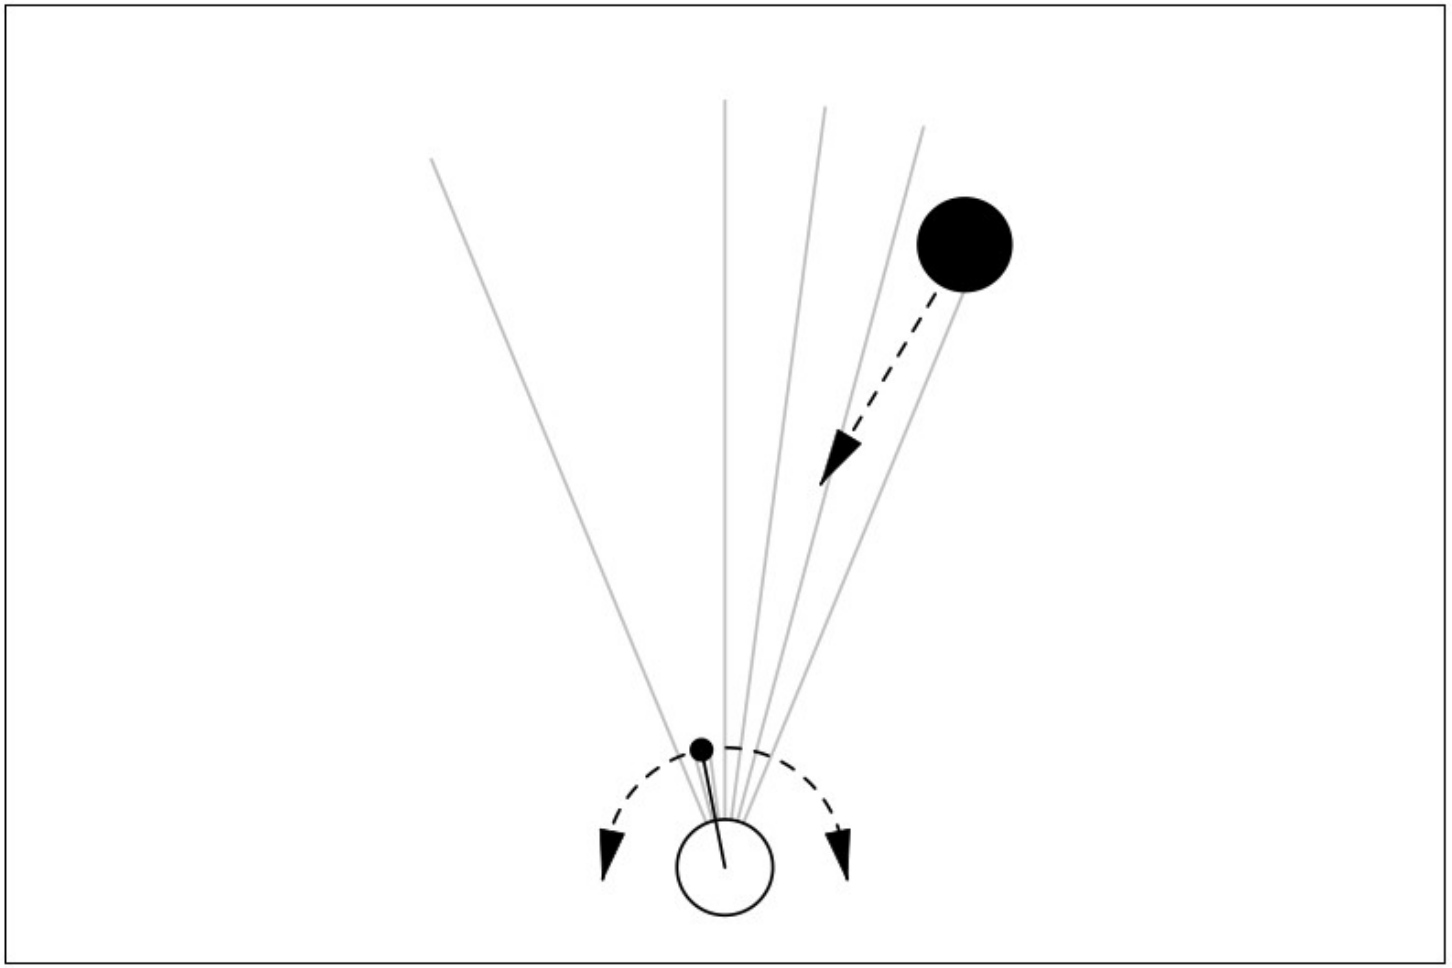

Figure 5: Experimental setup for self/nonself discrimination experiments. The agent is stationary, but can swing an arm with an opaque hand along an arc while objects fall from above.

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.

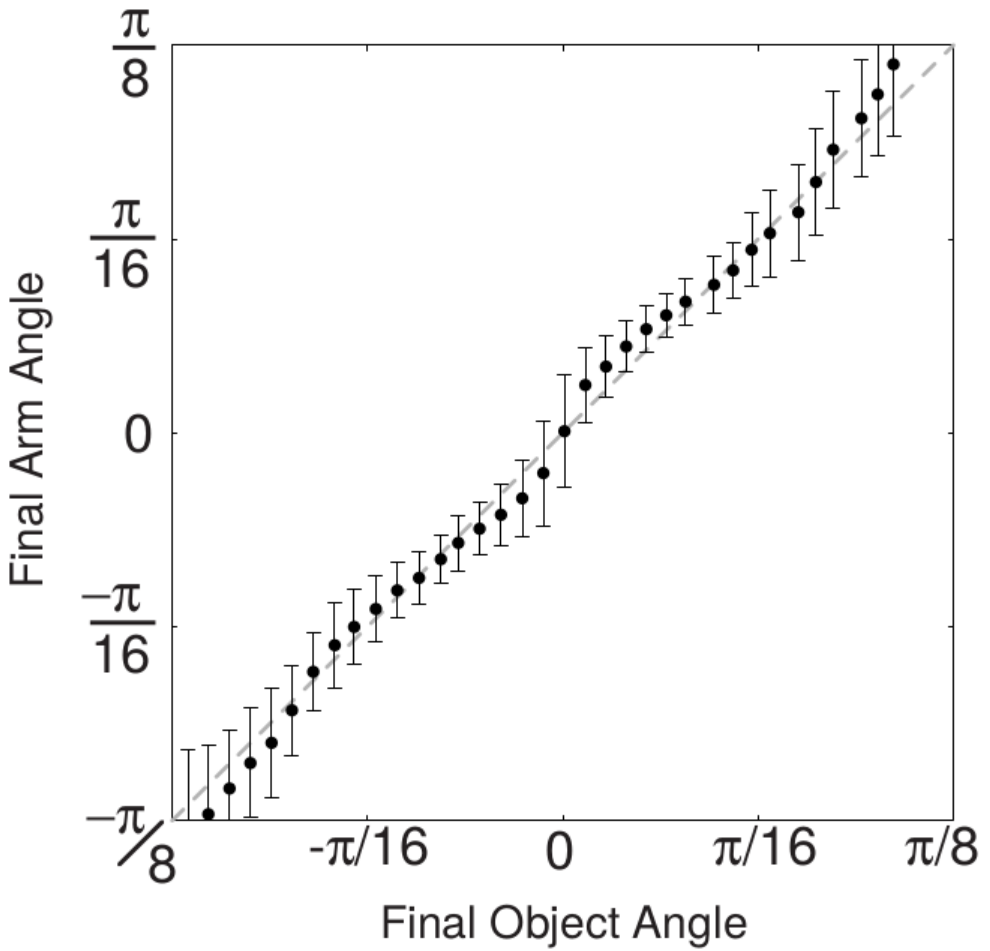

Figure 6: Mean catching accuracy of the best self/nonself discrimination agent. The final angular position of the hand is plotted against the final angular position of the object (mean ± s.d., N = 150 trials). Note that the average behavior closely approximates the ideal (dashed gray line).

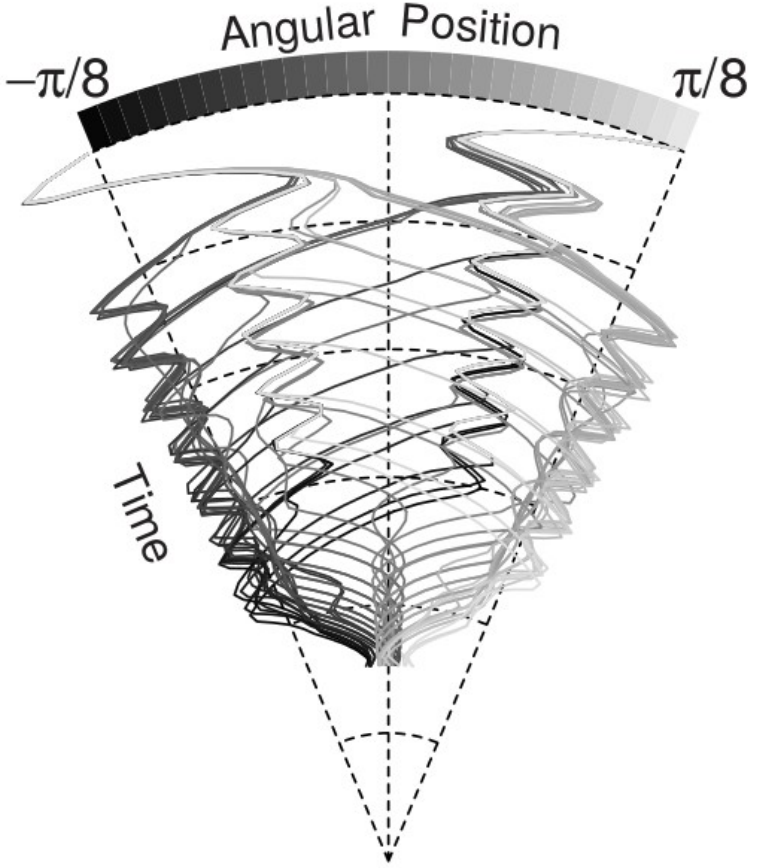

Figure 7: Arm angle trajectories over time of the best self/nonself discrimination agent catching objects at the midline from initial hand positions at either the left or right edge of the visual field. The trajectories are shaded according to the initial angular position of the object as indicated at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

Short-Term Memory

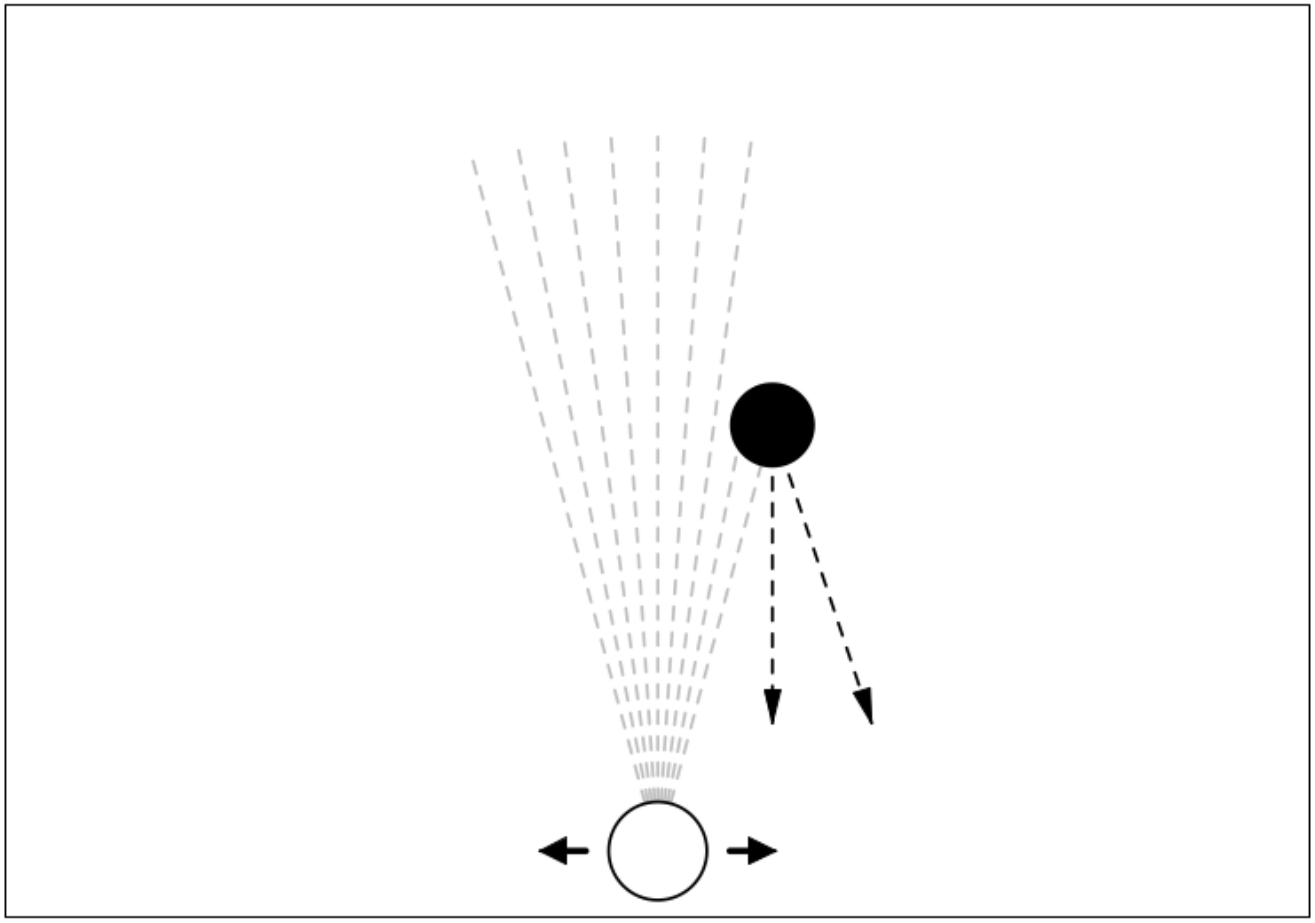

Figure 8: Experimental setup for short-term memory experiments. The agent can move horizontally while objects fall either vertically or diagonally from above. The rays are dashed because, as soon as the agent begins to move, it goes blind.

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.

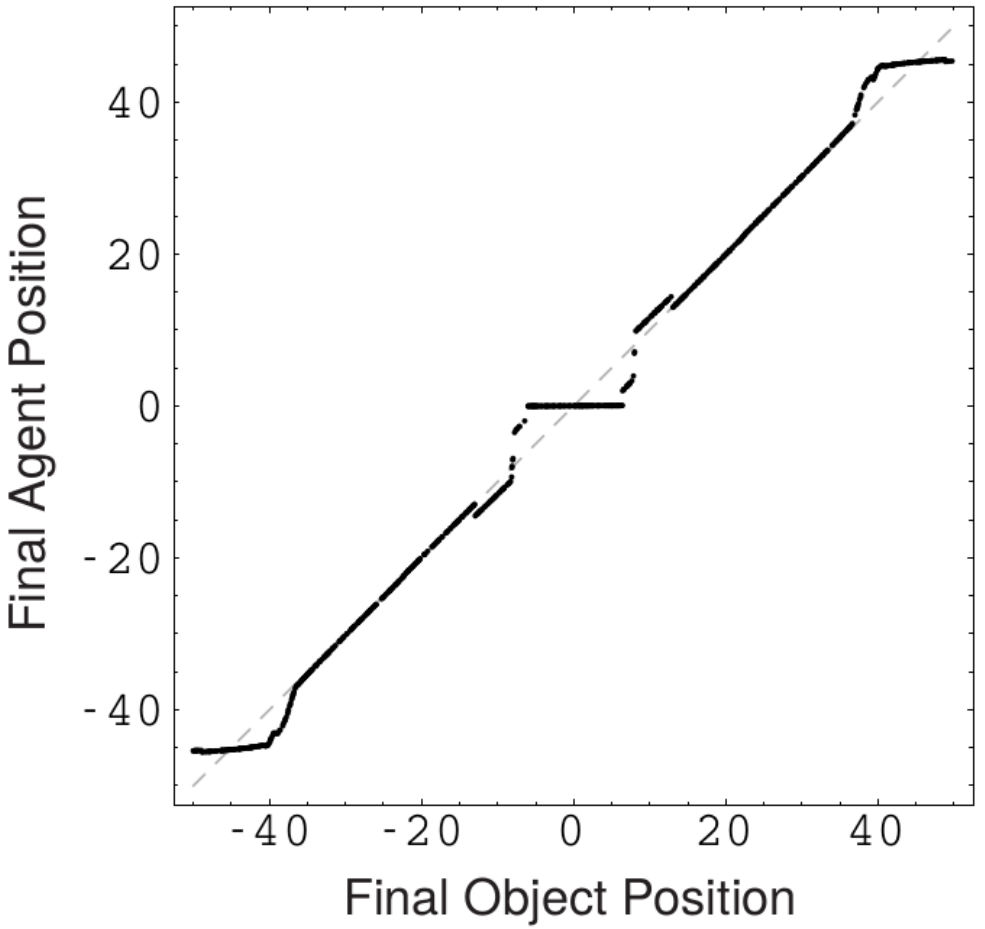

Figure 9: Accuracy of the best short-term memory agent for vertically falling objects. The final horizontal position of the agent is plotted against the final horizontal position of the object. Note that the average performance closely approximates the ideal (dashed gray line) except at the midline and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

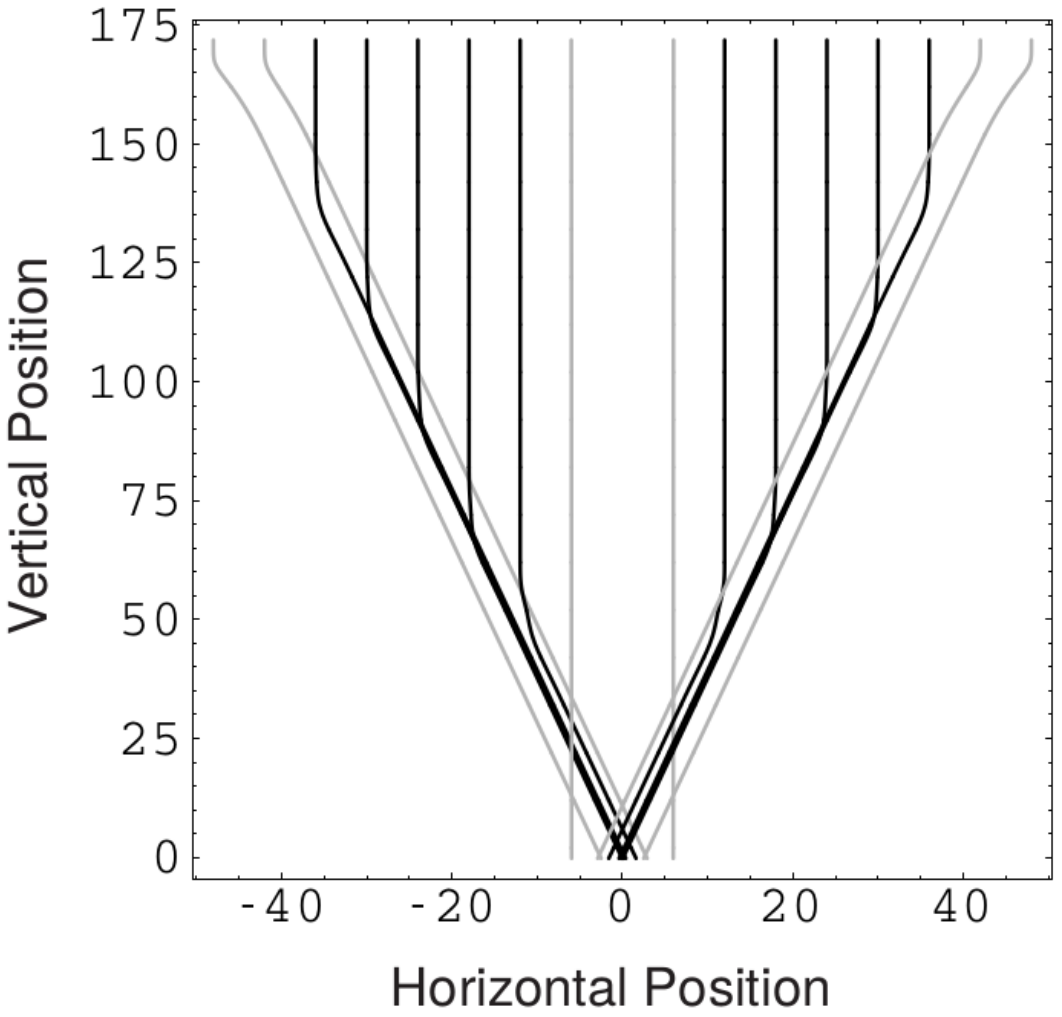

Figure 10: Behavior of the best short-term memory agent for vertically falling objects. Trajectories of motion relative to the agent of objects falling vertically from several different initial horizontal offsets are shown. Gray trajectories are those for which the agent's strategy begins to break down.

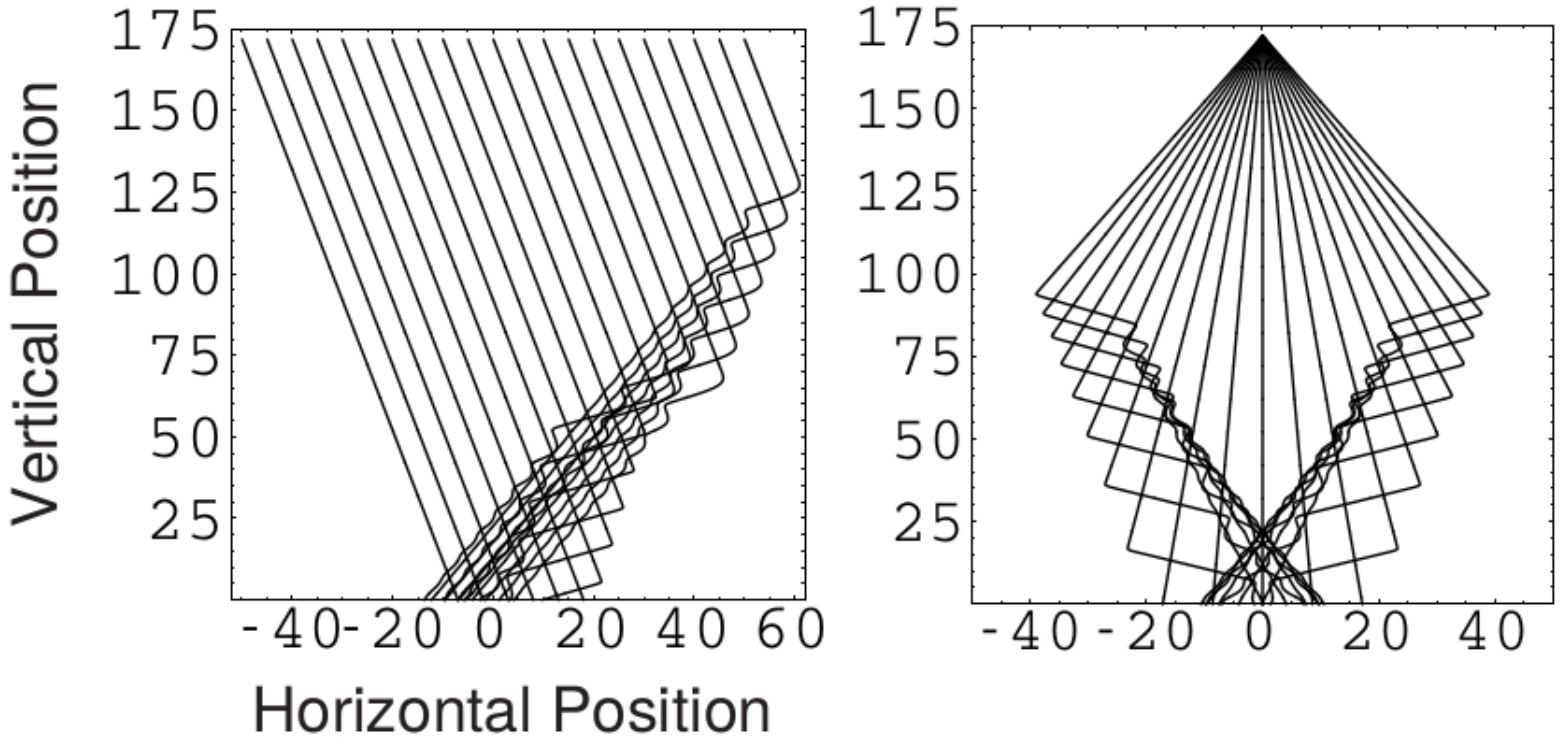

Figure 12: Behavior of the best short-term memory agent for diagonally falling objects. (Left) Trajectories of motion relative to the agent of objects falling diagonally with a horizontal velocity of 0.5 from several different initial horizontal positions are shown. (Right) Trajectories of motion relative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

Selective Attention

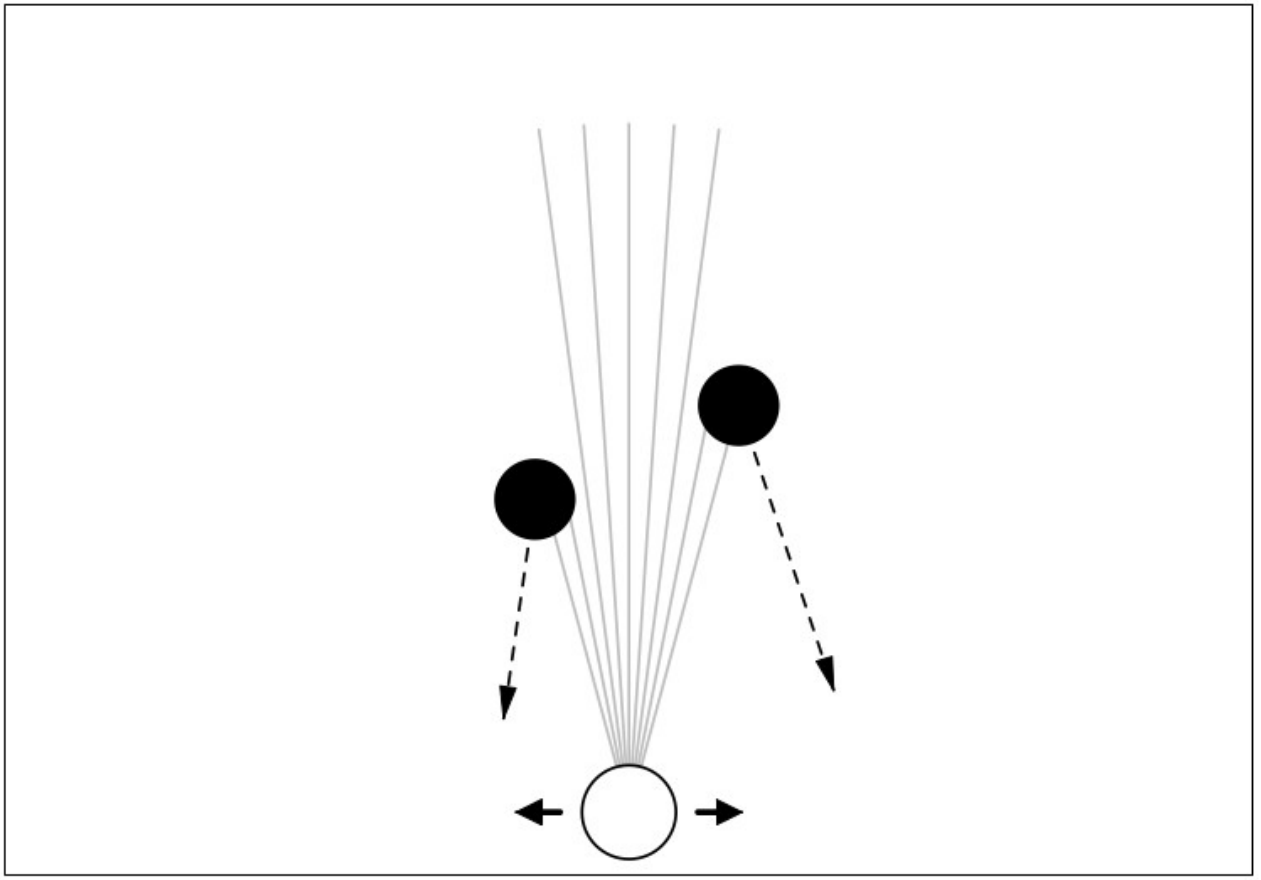

Figure 13: Experimental setup for selective attention experiments. The agent can move horizontally while 2 objects fall from above. As described in the text, the initial positions and velocities of the two objects are constrained so that the agent has a reasonable chance of catching both.

The 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.

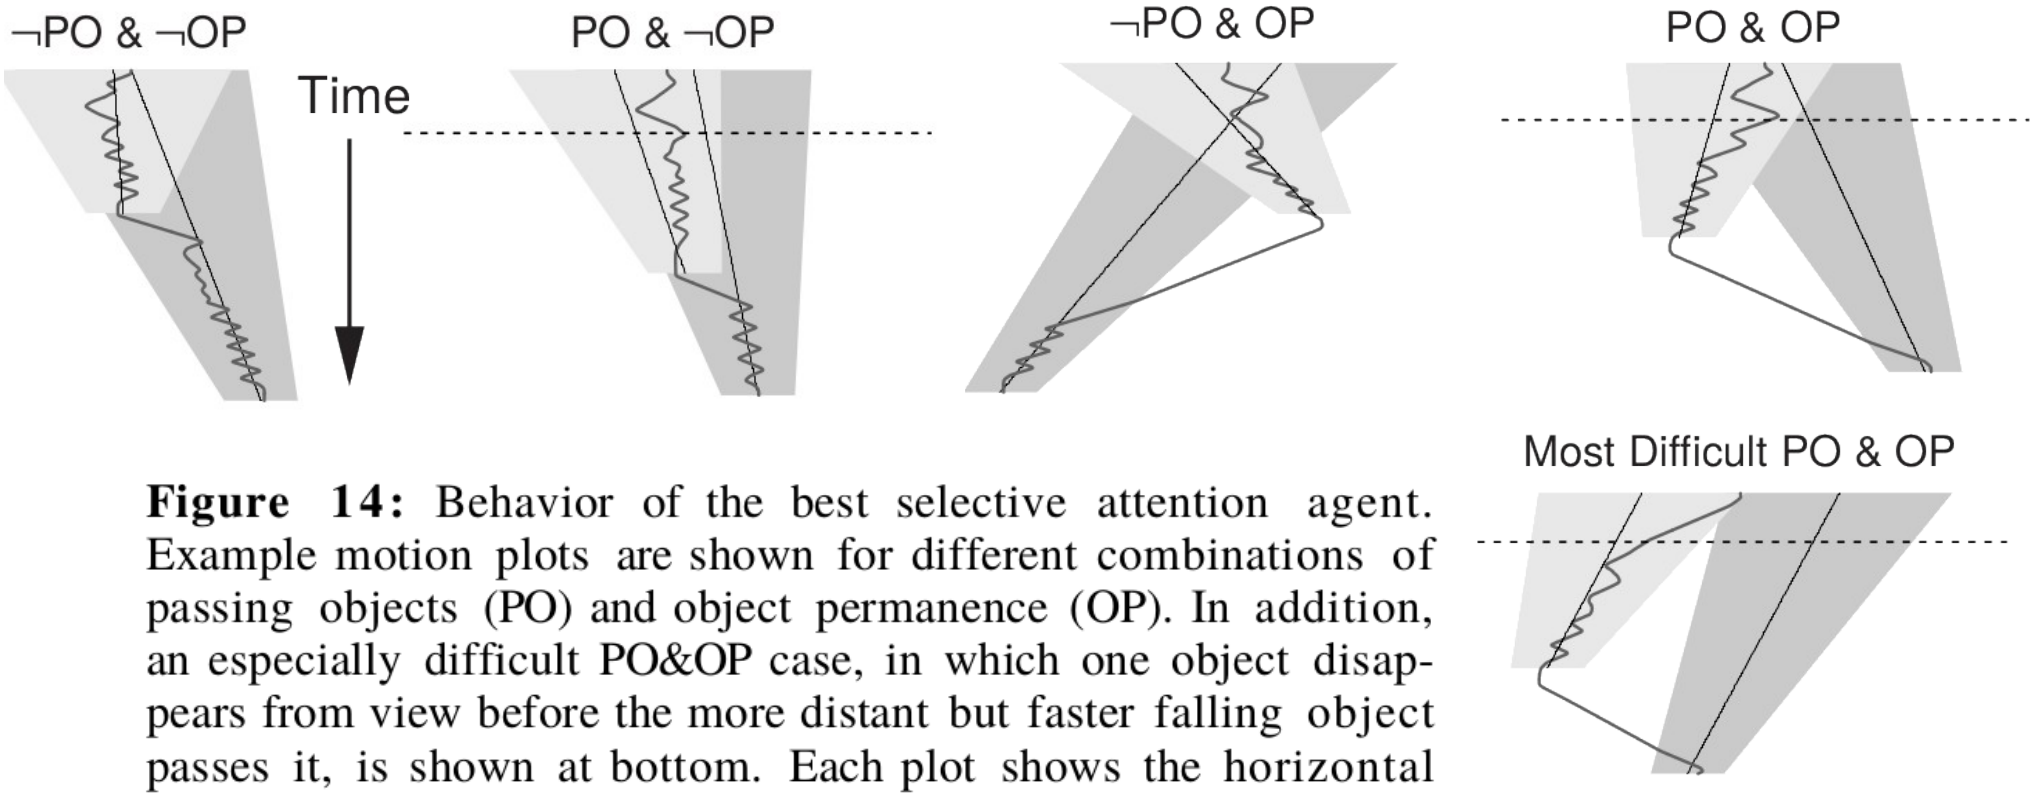

Figure 14: Behavior of the best selective attention agent. Example motion plots are shown for different combinations of passing objects (PO) and object permanence (OP). In addition, an especially difficult PO&OP case, in which one object disappears from view before the more distant but faster falling 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.