Anatomy of a Martian Iguana

Paul Fitzpatrick
Embodied Intelligence (6.836)
MIT ID 984985527
paulfitz@ai.mit.edu

Abstract

Adaptive behavior is best understood through study of the complete action-perception cycle within which it occurs [7]. For naturally occurring behaviors, this requires detailed modeling of an animal's nervous system, body, and environment. Beer in [2] presents a simplified model agent which captures many biologically relevant qualities, while itself having no immediate analogue in the animal kingdom. He then evolves a number of visually-guided behaviors using the model, and suggests that a careful analysis of their operation could provide empirical data concerning the role representation plays in cognitive behavior. This paper provides an example of such analysis, discusses its consequences, and gives a critique of Beer's methodology.

1 Introduction

Adaptive behavior emerges from rich interactions between the nervous system, body, and environment of an animal. Computational neuroethology is a methodology for studying the neural basis of adaptive behavior within an explicitly embodied, situated framework [5]. It makes use of joint models of relevant aspects of the neural mechanism itself, the biomechanics of the body, and the external environment. Models used in computational neuroethology vary in degree of biological plausibility. The more realistic can make quantitative, experimentally verifiable predictions [15]. More qualitative work can inform further experimentation or put theoretical models from the biology literature to the test of mechanistic implementation [7]. Beer proposes an even simpler model that could be used to empirically investigate ideas that are currently argued primarily at the theoretical level [2]. His model is sufficiently complex to be able to exhibit very simple instances of what he calls "minimally cognitive behaviors", such as basic object perception and discrimination, while still being amenable to rigorous analysis.

The most sophisticated behavior examined by Beer involves a visual discrimination task where the agent tries to catch circular objects, but "runs away" from diamonds. Beer creates an environment and an agent body that are just sufficient to allow this behavior. He then evolves a dynamical neural network as a control system for the agent, such that the network will actually cause the desired behavior to be exhibited. His paper concludes with a series of questions: are there "circle" and "diamond" detectors in the agent; does the agent make use of representation in any meaningful sense; or is it more appropriate to view the role of the agent's control system as facilitating the expression of the desired behavior by the whole dynamical system within which it is embedded. His purpose in the paper is to show that these kinds of questions can be investigated within his framework, and not to actually answer them. This paper takes up that challenge. It presents an analysis of an agent evolved using the same neural network architecture, body, and environment that Beer describes, and then attempts to answer his questions. The paper then draws together the arguments as to why we can reasonably expect questions asked of an invented model to have some true relevance outside of that model, and not simply be a computationally expensive form of navel gazing.

2 The dynamical system

This section briefly summarizes the nature of the dynamical system within which Beer evolved his minimally cognitive behaviors, as described in [2]. This consists of a joint model of a "nervous system", a body, and an environment, which together form an autonomous dynamical system. The details of the model vary somewhat from behavior to behavior. Only those details relevant to the circle/diamond discrimination behavior are described here.

The environment

The simulated agent is situated on the "floor" of a rectangular environment of dimensions 400×275 units (see Figure 1). Space is empty except for the agent itself and a single object, which "falls" down from an arbitrary location at the top of the environment. The object is either a circle of diameter 30 units or a diamond with edges of length 30.

The body

The agent has a circular body of diameter 30. It is constrained to move horizontally along the bottom of the

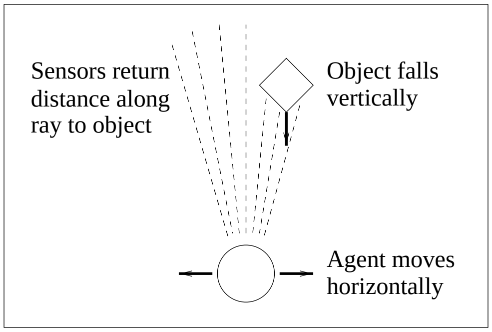

Figure 1: The agent and its environment (not to scale). Objects fall vertically down towards the agent, which is constrained to move in the horizontal direction only. A simulation run ends when the object approaches the bottom of the environment.

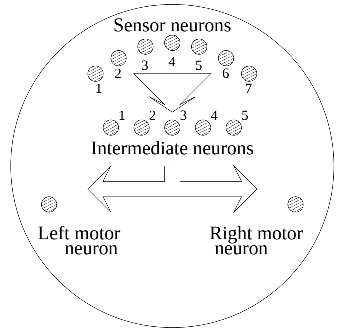

Figure 2: The "nervous system" of the agent. There are seven sensor neurons feeding forward to five fully interconnected intermediate neurons, which feed forward to a left and right motor neuron.

environment, and cannot move vertically. The agent is also constrained to face directly upwards. It has a set of sensors arranged as in Figure 1. These see along a "ray" projecting out from the body of the agent. The leftmost and rightmost ray point 30° from the vertical. There is no noise in the sensors, actuators, or any part of the dynamical system.

The nervous system

The nervous system of the agent is a continuous-time recurrent neural network, whose neurons obey:

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

$$\sigma(x) = \frac{1}{1 + e^{-x}}$$

where y is the activation level of the neuron, w_{ji} is the connection weight from the j^{th} to the i^{th} neuron, g is a gain, θ is a bias, I, when non-zero, is an external input from a sensor, τ is the time constant of the neuron, and N is the number of neurons. There is a single neuron for each sensor, which all feed forward to a set of five fully interconnected intermediate neurons, which in turn feed forward to a pair of motor neurons. Motion of the agent left and right is controlled by the difference in activation level of the two motor neurons. Bilateral symmetry is imposed on the parameters of the neural architecture. In other words, connection weights, biases, and gains associated with the neurons on the left of Figure 2 are equal to those in the symmetric location on the right.

3 The behavior

The dynamical system is now completely specified except for the actual parameters of the neural network controller. These parameters were chosen by artificial evolution such that the agent would move towards falling circular objects, and away from diamond-shaped objects. The evolutionary algorithm will only be sketched here. It relied on a real-valued encoding of the network parameters, with mutation but no cross-over. Mutation was performed by addition of a vector chosen from a uniform distribution over the unit hypersphere of appropriate dimensionality, scaled by a Gaussian centered on zero and with a specified variance. I used an algorithm due to Sibuya [14] for this. For further information on the evolutionary algorithm, or on the components of the dynamical system, see the original paper, which is refreshingly clear on the details.

Using a population size of 100, a fully competent individual emerged after several thousand generations. Figures 3 and 4 show the horizontal position of the agent over time as it responds to either of the two object shapes, circular or diamond. Clearly it achieves the desired behavior. Performance generalizes well over different initial horizontal displacements of the object. Notice that the agent approaches the object initially regardless of its shape, and then seems to decide whether to catch it or run away. The strategy is qualitatively similar to the one presented by Beer, although it differs in detail. The next section analyzes how the behavior operates, and to which aspects of the object's shape the agent is sensitive.

4 The analysis

Establishing the neural mechanism of the behavior directly is difficult. Analyzing a continuous time recurrent neural network with more than two neurons is hard (see [1] for a rigorous treatment of the one and two neuron case). This particular agent has five neurons with mutually recurrent connections. The approach I took was to perform experiments on the agent at the behavioral level, and then use the constraints this revealed to guide an analysis of the network.

Given that the agent needs to perform exactly one of two possible actions, catch or flee, it is worth asking

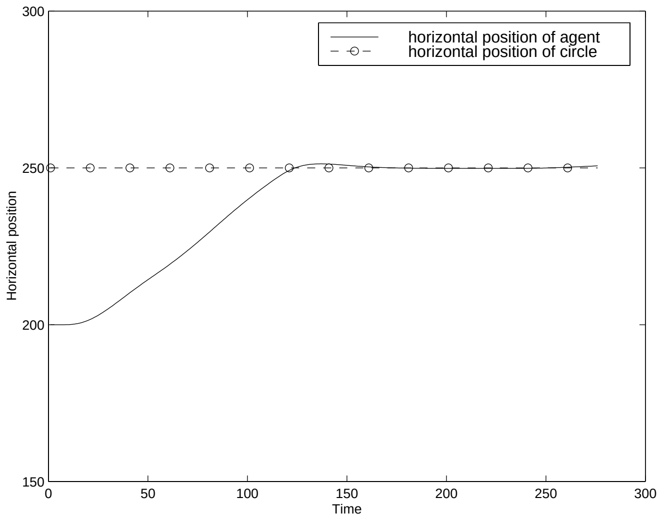

Figure 3: The response of the agent to a circular object. The graph shows the horizontal position of the agent and the object over time. Time and distance units are both arbitrary. The agent moves to match its position with the object.

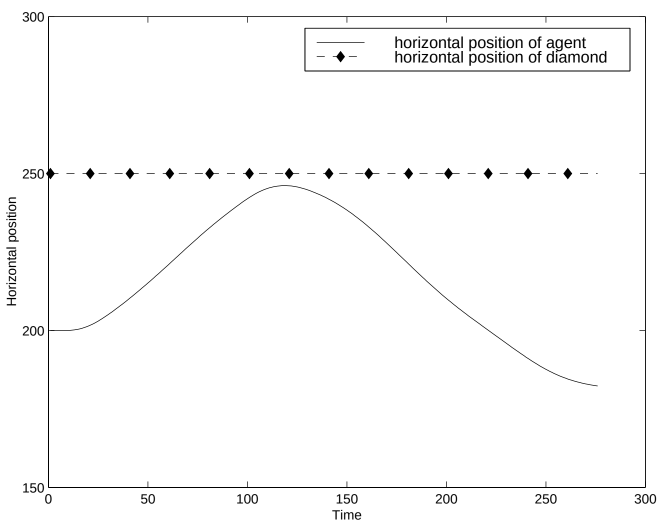

Figure 4: The response of the agent to a diamond-shaped object. The agent closes in on the object, then moves away from it again.

when it makes the commitment to one or the other. One way to test this is to initiate an experiment with the falling object having one shape, and then watch how the agent responds if the shape of the object changes in mid-experiment. Figure 5 shows that the agent will respond to the new shape only if it is introduced early enough. This holds true even when the agent has the object in full view. This suggests that the agent is receptive to the shape of the object for a limited period only, and from then on is receptive only to the cues necessary to either move closer or further away.

Another important concern is to determine the cue or

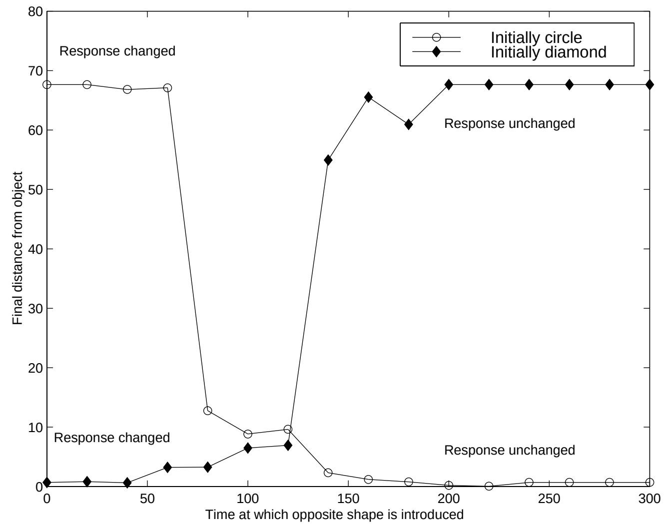

Figure 5: The effect of switching the shape of the object in mid-experiment. Low values imply that the agent finished close to the object, suggesting that it treated it like a circle. High values suggest the opposite, that the agent treated the object like a diamond.

cues that the agent uses to distinguish circles for diamonds. Beer speculates that the relevant cue may be "smooth" versus "pointy". The truth, at least for every agent I evolved, was somewhat more mundane. Recall that circular objects were specified to have a diameter of 30, and diamonds to have sides of length 30. Viewed from below, the horizontal extent of a diamond will therefore be $30\sqrt{2} = 42.4$, while the horizontal extent of a circle will be just 30. Figure 6 gives strong evidence that horizontal extent, or something very strongly correlated with it, is the cue used by the agent to distinguish between the two shapes. If a circle of diameter 42.4 is presented to the agent, it will treat it as a diamond, and if a diamond is shrunk to have a horizontal extent of 30, it will be treated as if it were a circle.

Now the question is: how does the agent's behavior come to be determined by the horizontal extent of the object? Perhaps some of the agent's sensors play a larger role in this than others. One way to find out is to "blind" each sensor in turn and see how it affects the behavior of the agent. This gives some information, but is difficult to interpret because blinding sensors affects aspects of the behavior we are not immediately concerned with, such as the agent moving towards the object to start with, and the feedback loop needed for efficient catching. The ambiguity this causes can be removed by making stronger use of the complete control we have over the simulation. My solution was to let one sensor "see" the object as the opposite shape to the rest, and then determine how strong an impact it had on the overall behavior of the agent. Now the activity of the sensor is different only when the shape of the object makes a difference, which

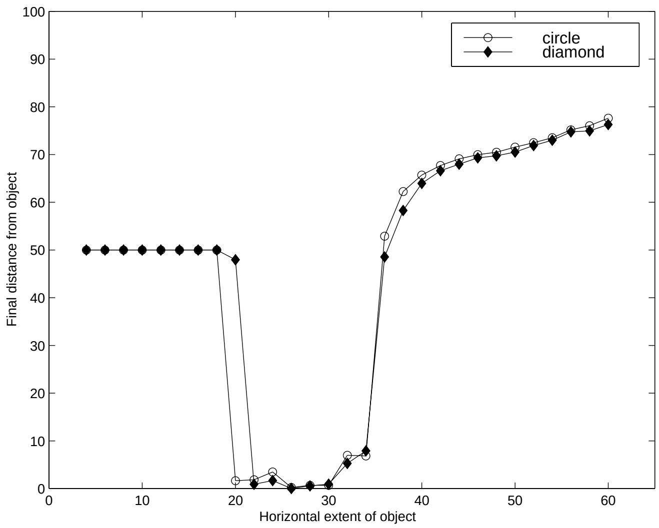

Figure 6: Response of the agent to objects with varying widths. This is the diameter of a circle, or the corner-to-corner distance of a diamond. Notice that the agent treats diamonds and circles with the same horizontal extent almost identically. Objects with widths of around 20–30 are treated as attractive, and all other widths are treated as repulsive.

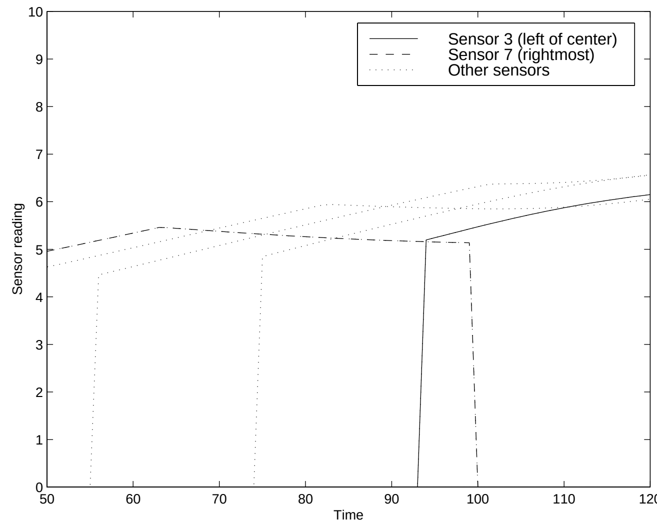

Figure 7: Sensor activity for the agent interacting with a diamond object during the "sensitive" time period identified by earlier experiments. Note that the overlap in the activity of the two sensors implicated in discrimination.

is exactly what we want. The results of this experiment were much clearer, showing that only two of the seven sensors played a significant role in discrimination. When the object started to the right of the agent, these were the right-most sensor and the sensor immediately to the left of center (sensor neuron 3 and 7 in Figure 2). When the object fell to the left of the agent, the symmetrically opposite pair of sensors were involved instead (1 and 5). This must be so, given that the network is bilaterally symmetric, and that the falling object is the only factor that breaks the symmetry of its environment[1].

Figure 7 shows the activity of the sensors implicated in discrimination within the critical time period identified in Figure 5 and refined by further experimentation. Notice that the regions where the sensors are active overlap. In contrast, for a circle there is a distinct gap between their active regions. For objects with widths varying between that of a circle and that of a diamond, the gap length correlates strongly with the behavior of the agent. More convincingly, the agent can be made to display behavior ranging from full attraction to full repulsion by artificially varying the gap length. Also, if sensor data is made binary (either high or zero, with no fine-grain distance information) the behavior of the agent is unaffected, other than becoming slightly less efficient at catching. This suggests that the agent truly is measuring a qualitative feature such as this gap length, rather than some more complex quantitative measure. That the gap length is in fact the primary feature of interest was confirmed by further experimentation and controls. Analysis of the trigonometry shows that the gap length is a good indicator of object width, although the perceived width depends on how far away the object is when the agent reaches it. The slack around the widths of circles (30) and diamonds (42.4) shown in Figure 5 means that a certain amount of error in the measurement is tolerable.

Figures 8 and 9 show the internal state of the agent as it responds to either a circle or a diamond. This state consists of the activation levels of the intermediate neurons. Both sensor and motor neurons also contain state, but this plays a role only over short time periods because they have no recurrent connections. Notice that striking differences do not arise in the internal state until quite late. These differences can be shown to be caused by the fact that the agent is close to an object, or distant from it, rather than causing the agent to be close or distant. An attempt to interpret earlier differences in a causal manner quickly runs into trouble. For example, consider the slightly different early activation levels of neuron 2. Tracing what causes this difference is extremely hard. After much effort, the best I could say is that yes, it correlates with the horizontal extent of the object. But unfortunately, so did every activation level. This is not surprising given that all the sensor neurons feed to every intermediate neuron, and every intermediate neuron feeds to every other intermediate neuron.

So much for "cause", how about "effect"? It is much easier to make definite statements about this. Earlier experiments determined that there is a specific time period during which the agent "decides" whether to catch

[1] This means it is impossible for an agent with this neural architecture to deal with objects falling directly from above, or falling outside its sensor range, because there is nothing to break symmetry and allow it to move left in preference to right or vice versa.

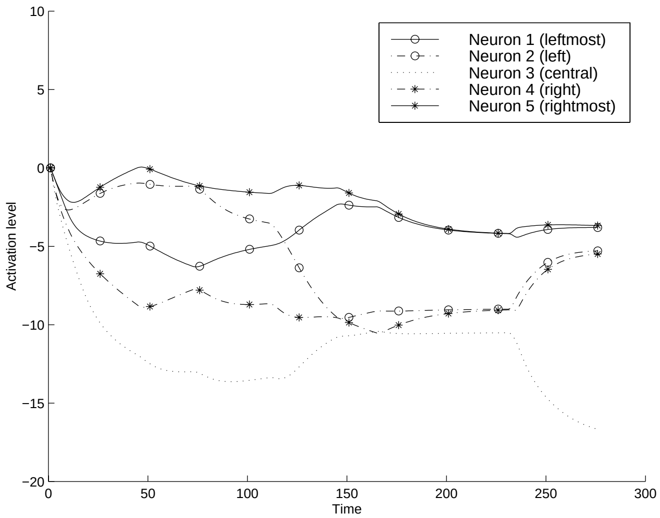

Figure 8: Activation levels of intermediate neurons over time when a circular object is presented.

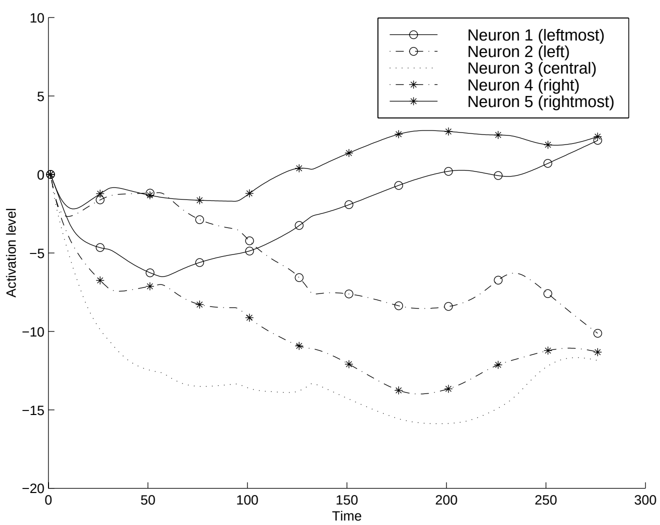

Figure 9: Activation levels of intermediate neurons over time when a diamond object is presented. Compared with Figure 8 there are no dramatic differences during the sensitive period, although there are many subtle ones.

or flee. This suggests that the dynamical system initially progresses similarly regardless of the shape of the object presented, and then diverges irreversibly into two different modes depending on some feature that correlates with the object shape. The feature could be directly sensed, or could be embedded in the accumulation of small differences in neural activation levels. The key point is that at some stage the dynamical system is likely to diverge based on a relatively small difference. So one way to work out when the dynamical system is "making a decision" is simply to test how sensitive it is to small changes in neural levels at various times. Figure 10 shows the result of doing just that. For three of the five intermediate neurons, these "kicks" have little effect at any time. The remaining two neurons are quite robust to kicks too — except in a particular range, within which a kick is sufficient to make the dynamical system act as if the object had the opposite shape. This sensitive area compares well with Figure 5. When the object starts to the agent's right, the two neurons implicated are intermediate neurons 2 and 5 in Figure 2.

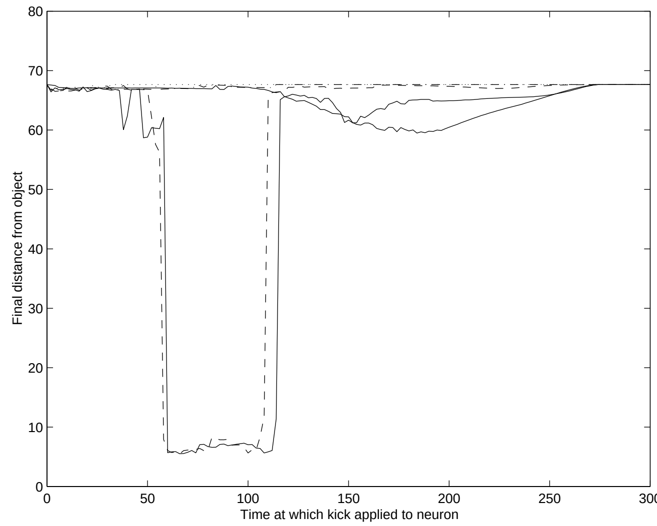

Figure 10: Effect of "kicking" intermediate neurons at different times. The object the agent is interacting with happens to be a diamond, so transitions from high final distance to low final distance show "sensitive" regions. The two neurons that show sensitivity are intermediate neurons 2 and 5.

So, although every piece of state in the agent correlates with the shape of the object, the activation levels of two specific neurons within a specific time period seem particularly crucial for determining the behavior of the agent. I will return to this point later, and discuss what it implies in terms of the use of representation by the agent. The next section briefly touches on some explorations of agents evolved to perform discrimination without the object width cue.

5 Variations on a theme

I evolved a number of agents to generate the "circle-loving" behavior as specified by Beer, and all relied on measuring the horizontal extent of the object in some way. I do not believe that Beer intended his particular choice of object dimensions to allow the agents to use this relatively simple "hack", but his specifications are too clear and complete for the problem to simply be a misinterpretation on my part. Previous work Beer was involved in showed that dynamical neural networks could indeed be used to implement an agent that distinguished between shapes by integrate sensor readings over time [17]. Perhaps if circles and diamonds with

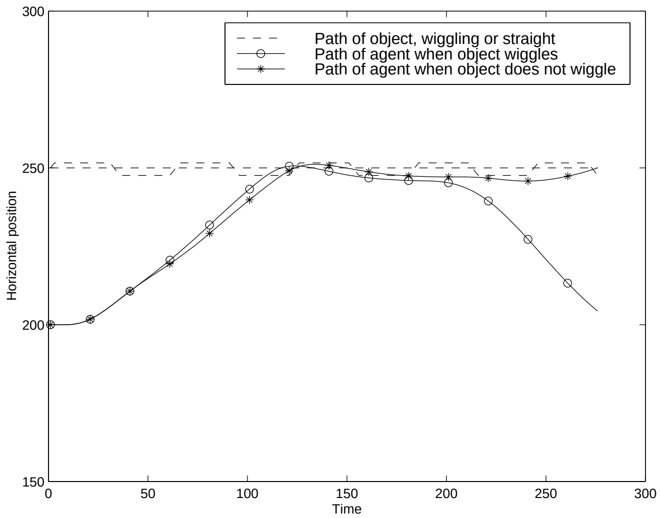

Figure 11: Path of agent responding to circles that fall in a straight line, and circles that "wiggle" as they fall.

the same width were used, so that the horizontal-extent hack was not available, the agent truly would be forced to distinguish between "pointy" and "smooth" as Beer envisaged. My attempts to evolve such an agent met with limited success. The best agent behaved qualitatively correctly, in that it eventually turning towards the object if it was a circle and turning away if it was a diamond. Its strategy was to approach the object in such a way as to arrive at a stereotyped distance from it at a well-defined time, and then to drift closer or further away depending on small differences in the sensor readings at that point. This was not particularly robust. A better approach to creating such an agent would be to evolve with test-cases that include circles and diamonds of varying sizes. However, adding another parameter to generalize over made the problem too hard for my computational resources and methodology. To perform an accurate discrete simulation of a continuous time recurrent neural network requires fine-grain time steps, and is quite computationally expensive.

I was interested in whether this type of agent could be evolved to discriminate between objects that differ not in shape but in temporal behavior. In particular, I tried to evolve an agent to catch circles that fell passively as before, but to run away from circles that "wriggled" a little as they fell. Figure 11 shows a successful agent doing just that. The strategy it adopts is to close in on the object initially, and then once the object has fallen far enough a wriggle becomes easy to detect. It is interesting that the agent relies on a simple form of active vision in both this case and in the circle-loving behavior from the previous section. This suggests that active vision is strongly advantageous for agents with limited sensing capability. Beer found similar "foveating" behavior in his agents.

6 Use of representation

I now return to Beer's question about the use of representation in the agent. Is there something in the neural network of the agent that we can point to and call a representation? Philosophically, there are difficulties even asking the question. A representation is only a representation *for* or *to* someone; the concept is meaningless without an observer (see [8] or [9, page 101]). It must be clear who or what this representation-user is. Otherwise we will have introduced a homunculus into our analysis, which will need explaining in turn.

Before tackling the actual circle-loving agent presented in Section 4, imagine that one of its neurons went "high" in the presence of a diamond and "low" in the presence of a circle[2]. As observers, we may choose to say that the activation level of the neuron represents the shape of the object. Irrespective of whether this is justified or not, it does not tell us whether the *agent* itself uses representation. We have isolated something that may be a representation; now we must identify something within the agent that can be called a representation-user. This is not easy to do in a satisfying way.

Turning to the actual agent presented in this paper, the situation is even less clear. First, there is no convenient single feature related to the shape of the object that we as observers can pick out. There is also no obvious candidate for a distributed representation (I tried various forms of dimensionality reduction in the search for one, without success), except in the weak sense that the entire state of the agent might serve as one.

Beer asks if, rather than talking about representation, it might be more appropriate to view the neural circuitry of his agents as "merely instantiating dynamics that, when coupled to their bodies and environments, give rise to effective performance of the tasks for which they were selected". The language of dynamical systems certainly seems to match reality better in this case. Dynamical systems theory also promises to provide a new set of tools for guiding our understanding of systems like this agent and its environment. See, for example, [12] for an analysis of an evolved agent in terms of state space attractors. For the circle-loving agent, even though formal analysis was difficult, useful structure could be picked out by perturbing the dynamical system and noting changes in its time evolution. That analysis showed that slightly perturbing a particular part of the network (neurons 2 and 5), during a particular time period (50 to 100 steps), strongly affected later behavior of the agent; the behavior bifurcates into either catching or running away. In a sense, this point in space and time shares some characteristics with what we think of as representation. The state of the network elements at this point correlate with some external reality, so we can interpret or assign meaning-

[2] In fact, this was the case for one agent I evolved.

ful labels to the state; the state also has a clear impact on determining the agent's behavior. There are many significant differences however. The "representation" is temporally bounded; look later or earlier at the same network elements and their states have very different interpretations, if indeed they can be interpreted meaningfully at all without reference to the rest of the network. Experimentation shows that the "representation" can occur at different times, depending on the distance to the object, and even in different parts of the network, depending on whether the object is to the left of the agent or to the right (neurons 1 and 4 are used instead of 2 and 5 if the object is to the left of the agent rather than the right). So it seems more sensible to speak in the language of dynamical theory and call a bifurcation a bifurcation, rather than trying to stretch the notion of representation to the point of ripping it.

7 Discussion

We can conclude from the previous section that representation in the conventional sense is not necessary for an agent to perform the circle-loving behavior. Can any general conclusions be drawn from this? This section examines how realistic it is to claim that experimentation with Beer's model can provide useful empirical data about the general nature of cognitive behavior.

The short answer is that at our current level of understanding, existence proofs of agents performing visually-guided tasks without the use of representation do indeed represent a step forward (consider, for example, progress in behavior-based robotics [13]). This answer is sufficient to justify the work in this paper.

But the issue of representation was just one example chosen by Beer to motivate his model. He intends the model to be generally useful for studying the "cognitive implications of adaptive behavior ideas". Is this hope justified? Clearly this can only be answered by further research — but what I can do is show that if he is wrong, he is at least in good company. When Beer compares his approach to related work, he focuses on current research on evolving or building agents. I intend to instead discuss Beer's work in relation to literature explicitly dealing with the idea of inventing simple artificial creatures to explore issues in cognitive behavior.

In this and other work, Beer's interest lies in developing general principles for a dynamical theory of adaptive behavior [3]. To this end, he choses to work with a very simple idealized model agent. It can be argued that since "general principles" by definition cannot be specific to any particular system that exhibits them, we are free to trade off biological plausibility for simplicity in the models we use to elucidate those principles. A similar point was made in the context of cognitive organization by Dennett in a much-cited commentary:

> ... one does not want to get bogged down with technical problems in modeling the cognitive eccentricities of turtles if the point of the exercise is to uncover *very* general, *very* abstract principles... So why not then *make up* a whole cognitive creature, a Martian three-wheeled iguana, say, and an environmental niche for it to cope with? [11, original italics]

In a sense this is exactly what Beer has done, and exactly the reason why he has done it. Some very early work by Toda also resonates with Beer's notion of minimally cognitive behavior. Toda sketched a design for a fungus-eating[3] robot uranium miner which lived in a simplified environment and performed a specific cognitive task (collecting uranium) involving visual orientation and discrimination [16]. His stated goal was to "begin with an environment, and attempt to design a subject with the minimal optimal qualities to function effectively in this environment." It is worth noting that a considerable amount of Toda's design relates to choosing an appropriate distribution of receptors over the robot's retina to suit its particular task. This illustrates an advantage to working with a complete model of an agent and its environment: there are now few places to sweep hard problems such as perception "under the rug", as arguably happened unwittingly in later AI work [6].

A criticism that could be leveled at Beer is that there is no reason *a priori* to expect adaptive behavior to have particularly strong general principles. The contrary could in fact be argued, given how strongly specific behaviors tend to be tied to a particular environment and agent morphology. I should stress here that Beer is cautious in his stated objectives: "the goal is to explore the space of possible dynamical organizations of agents that engage in minimally cognitive behavior". Such an exploration may reveal constraints inherent in the nature of adaptive behavior, if such exist, and if not it at least gives a glimpse of "life as it could be" [6] to contrast with the extant biological world. Again to quote Dennett, this time commenting on the philosophical underpinnings of modeling work in AI:

> ... getting the cat skinned at all can be a major accomplishment... This sort of research strategy permits highly abstract constraints and difficulties to be explored (how could *anything* learn a natural language? how could *anything* achieve a general capacity for pattern recognition in an unstereotypic environment?) ... [10, original italics]

The work of Braitenberg offers perhaps the most convincing justification for Beer's model. Braitenberg, in a celebrated series of thought-experiments [4], designed vehicles that exhibited seemingly complex behavior from

[3]Note that this research took place in the sixties.

relatively simple components. He noted that it is far easier to invent such systems than it is for an observer to analyze their mechanism after the fact. He called this observation the "law of uphill analysis and downhill invention". Applying this lesson in the current context, it suggests that it may be far simpler to construct adaptive behaviors than it is to analyze pre-existing ones. This is close to what Beer does. He designs an environment, a body, and a nervous system such that the resulting dynamical system gives rise to a desired behavior. If there are in fact general principles of adaptive behavior, Braitenberg's evidence suggests that we might discover them far more rapidly by invention rather than by the horrendously difficult task of reverse-engineering nature.

There was a little sleight of hand in that last paragraph because Beer does *not* in fact design the entire dynamical system by hand. In particular, only the architecture of the nervous system is specified, with the actual parameters being chosen by artificial evolution. Understanding the resulting agent therefore requires some "uphill analysis" after all. This is in fact an interesting extension to the situation Braitenberg considered. What happens if there is a part of a device you are trying to invent that you have no idea how to build? Artificial evolution allows us to put our ignorance in a "black box" and still make progress. If we can successfully analyze the solution that evolution came up with, that is ideal. But even if we can only partially analyze the solution, then we may at least learn something about the space of viable solutions — for example, that representation-free strategies are possible in a given context. And even if we can't analyze the solution at all, we have an existence proof that the particular model we evolved parameters for is capable of exhibiting the desired behavior.

8 Conclusions

This paper has shown that a simple "circle-loving" visual discrimination task can be performed without making use of anything identifiable as a "representation" in the conventional sense. This answers a question posed by Beer in [2]. It seems likely that the corresponding agent that Beer evolved, but did not analyze, discriminated between objects based on their width rather than on a measure of smoothness as Beer may have intended. This cannot be stated definitively because there is no guarantee that evolution led to the same strategy in both cases. It did prove possible to evolve an agent to discriminate between objects which no longer differed in width, but the best strategy found was not particularly robust.

In general, Beer's model holds promise as one way to investigate the abstract properties of adaptive behavior. By working with an agent whose body and environment we are free to invent as we choose, and in a simulated environment, we have enormous experimental control compared with what is possible in biology (consider, for example, the experiment in Section 4 where an object was effectively one shape for one sensor and a different shape for all the rest). This means that if we evolve part of our design for an agent — either because we don't yet know how to construct that part, or because we are interested in exploring the space of possible designs — we have very strong leverage with which to analyze the result.

One cause for concern might be that the model does not incorporate noise, which can significantly alter the dynamics of a system [12] and may well play a crucial role in adaptive behavior generally.

References

[1] R. D. Beer. On the dynamics of small continuous-time recurrent neural networks. In *Adaptive Behavior*, 1995.

[2] R. D. Beer. Toward the evolution of dynamical neural networks for minimally cognitive behavior. In P. Maes et al., editors, *Proceedings of the 4th International Conference on Simulation of Adaptive Behavior*, 421–429. MIT Press, 1996.

[3] R. D. Beer. The dynamics of adaptive behavior: A research program. In *Robotics and Autonomous Systems*, 20:257–289, 1997.

[4] V. Braitenberg. *Vehicles: Experiments in synthetic psychology.* MIT Press, 1984.

[5] H. D. Chiel and R. D. Beer. The brain has a body: Adaptive behavior emerges from interactions of nervous system. In *Trends in Neurosciences*, 20:553–557, 1997.

[6] D. Cliff. A.I. and A-Life: Never mind the blocksworld. In A. Cohn, editor, *ECAI94: Proceedings of the 11th European Conference on Artificial Intelligence*, 803–808. John Wiley and Sons Ltd., 1994.

[7] D. Cliff. Neuroethology, Computational. In M. A. Arbib, editor, *The Handbook of Brain Theory and Neural Networks*, 626–630. MIT Press, 1995.

[8] D. Cliff and J. Noble. Knowledge-based vision and simple visual machines. In *Philosophical Transactions of the Royal Society of London*, Series B, 352:1165–1175, 1997.

[9] D. C. Dennett. *Brainstorms: Philosophical essays on mind and psychology.* MIT Press, 1978.

[10] D. C. Dennett. Current issues in the philosophy of mind. *American Philosophical Quarterly*, 15(4):249–261, 1978.

[11] D. C. Dennett. Why not the whole iguana? *Behavioral and brain sciences*, 1:103–104, 1978.

[12] P. Husbands, I. Harvey, and D. Cliff. Circle in the round: State space attractors for evolved sighted robots. In *Robotics and Autonomous Systems*, 15:83–106, 1995.

[13] M. J. Mataric. Behavior-based control: Examples from navigation, learning, and group behavior. In H. Hexmoor, I. Horswill, and D. Kortenkamp, editors, *Journal of Experimental and Theoretical Artificial Intelligence*, 9(2–3):323–336, 1997.

[14] M. Sibuya. A method for generating uniformly distributed points on N-dimensional spheres. In *Annals of the Institute of Statistical Mathematics*, 14:81–85, 1964.

[15] B. A. Skierczynski et al. A model of the hydrostatic skeleton of the leech. In *Journal of Theoretical Biology*, 181(4):329–342, 1996.

[16] M. Toda. The design of a fungus-eater: A model of human behavior in an unsophisticated environment. *Behavioral science*, 7:164–183, 1962.

[17] B. Yamauchi and R. D. Beer. Sequential behavior and learning in evolved dynamical neural networks. In *Adaptive Behavior*, 2(3):219–246, 1994.