The Evolution of Reaction-Diffusion Controllers for Minimally Cognitive Agents

Kyran Dale*,**
University of Sussex

Phil Husbands**
University of Sussex

Abstract   This article describes work carried out to investigate whether a classic reaction-diffusion (RD) system could be used to control a minimally cognitive animat. The RD system chosen was that first described by Gray and Scott, and the minimally cognitive behaviors were those used by Beer et al. involving the fixation and discrimination of diamond and circle shapes by a whiskered animat. A further task was added, which required the RD controllers to maintain and use a chemical memory. The parameters of these controllers were evolved using an evolutionary, or genetic, algorithm.

I   Introduction

Reaction-diffusion (RD) systems are at the heart of many biological processes, from growth and development mechanisms [32, 41, 43], to the generation of biochemical oscillations and cellular rhythms [26], to neural signaling [37, 39]. RD models are also very widely used in theoretical biology to better understand phenomena as diverse as evolutionary population dynamics [20], animal and plant patternation [12, 46], cardiac arrhythmia [35], and structure formation in RNA evolution [11]. However, these potentially highly intricate dynamic systems may also be capable of supporting a minimal form of cognition. This is interesting from both a theoretical perspective—it sheds further light on the behavior-generating capabilities of embodied dynamic systems [40, 49]—and from a biological one.

There are many examples of unicellular organisms, such as amoeboids and bacteria, that are capable of complex and plastic behaviors. Such organisms engage in sophisticated sensorimotor activity and interact with their surroundings to the extent that they can: identify and distinguish between various elements of their environment, make choices, adapt to changes, take part in coordinated group behavior, and make structured changes to their environment [24, 31]. They can be said to act purposefully, and there is a good case to be made for them being minimally cognitive [6, 33, 44, 47]. All of this is achieved without a nervous system. So how is behavior generated? There is strong evidence that, at least in some species, RD systems play an important role in coordinating biosignaling between spatially distributed sensors and actuators [50]. It is also interesting to note that many of the molecular pathways used by single-celled organisms have been conserved by evolution and play key roles in brain functions in higher animals [25, 27]; the biochemical mechanisms present in unicellular organisms are the origins of natural cognition.

Although progress has been made in developing a systems-level understanding of some of the mechanisms involved in microorganism behavior generation, we do not yet have sufficient knowledge of the spatiotemporal properties of the RD systems at play in these processes to model them at the level of interacting diffusing chemicals. Hence a more abstract investigation of the minimally cognitive behavior-generating potential of an embodied RD system seems highly relevant. The work described in this article is intended as an initial exploration in that direction. It acts as a probing study of the potential of RD systems in this context.

Beer and colleagues introduced a number of autonomous agent models and tasks to explore minimally cognitive behaviors involving aspects of sensorimotor coordination, active categorical perception, and memory [7, 8, 42]. The behaviors explored in this article are firmly based on Beer's, but whereas that work used evolved continuous recurrent neural networks as the "nervous system" of the agent, the research presented here employs a computational model of the well-known Gray-Scott RD system in place of a neural controller. The agent's sensors and actuators are coupled to the RD system by spatially distributed chemical sensors and probes. Chemical sensors, which measure the concentration of reagents in the chemical medium at a specific location, are connected to the agent's actuators via weighted links. Thus the chemical concentrations can control the motor outputs. The chemical probes, which are able to change the concentration of a given chemical at a specific location, are connected to the agent's sensors via weighted links. Thus sensory input can perturb the excitable medium. The weights and chemical specificity of these links are evolved. As we shall see, such a system is able to produce cognitively interesting behaviors, including the use of an evolved chemical memory.

Following a discussion of related work, the RD and agent models are presented. The section after that details the experimental method, describing orientation, discrimination, and memory tasks. Results are then presented and analyzed before the article closes with a discussion.

2   Related Work

There is a small but very interesting set of examples of prior research that has made use of RD systems in the control of mobile agents, albeit in a markedly different way from the work described in this article. Adamatzky et al. successfully demonstrated a methodology for robot navigation based on mapping the robot environment onto an excitable medium in which sites representing targets and obstacles are used as the sources of local excitation waves [2, 3]. The target is deemed to generate attractive waves, and the obstacles repulsive ones. The robot control system detects traveling and colliding wave fronts in the excitable medium and uses their positions and speeds to calculate a direction of motion toward the target that avoids obstacles. In their initial implementation repulsive waves were generated in a real Belousov-Zhabotinsky (BZ) reactor, with optical sensors detecting wave fronts, while attractive waves were generated in a cellular-automaton-based simulation of an excitable medium. The BZ reaction is too slow for a practical real-time application of the system using a chemical reactor, but a later version using a hardware-based cellular nonlinear network (CNN) implementation of excitable media was able to successfully navigate in real time [1]. CNNs, originally introduced to build large-scale analogue signal-processing systems [14, 15], are 2D arrays of locally interconnected cells. Each cell contains an identical simple nonlinear dynamic circuit whose output is related to the cell's state and possibly to those of neighboring cells. A cell's state, which is a continuous variable, depends on its history and nonlinear interactions with its neighbors. By designing the CNN state equations appropriately, such devices can be used to model RD systems in much the same way that standard computational RD simulations, such as the ones used later in this article, employ a regular spatial grid of cells coupled by sets of difference equations. In the CNN case the cells are individually implemented in hardware, allowing massively distributed parallel implementations.

Arena and colleagues used an RD CNN system to generate locomotion in a hexapod robot by exploiting traveling autowaves to provide rhythmically coordinated motor signals [5]. A later distributed CNN-based controller [4], using an architecture based on Cruse's stick-insect-inspired neural controllers [16], and similar in some respects to other locomotion controllers [10], was also successful. However,

this later controller did not make use of explicit RD-type dynamics. In a less direct application of RD systems to robot control, Trevai et al. used an RD model as part of an algorithm to calculate efficient paths for robots engaged in a cooperative exploration task [45].

All of the approaches outlined above make use of an RD system as part of a carefully designed overall behavior-generating mechanism; layers of pre- or postprocessing are required to integrate the RD system with other components of the mechanisms. In contrast, by coupling actuators and sensors directly to an RD system, the work presented in this article employs the RD system as the *entire* behavior-generating mechanism. Designing this coupling by hand in order to achieve interesting results is not feasible, so an evolutionary search algorithm has been used to explore the space of RD-controlled agents. The use of an evolutionary robotics methodology [23, 29, 34] to explore the class of systems also marks this work out from the approaches described above. The behaviors investigated in this article are also quite different from those explored previously, in that they involve explicitly cognitive aspects such as sensory discrimination and memory.

As far as we know, this is the first example of an evolved RD-based agent control system, although there has been previous work on evolved neurocontrollers that incorporated a model of diffusing neuromodulators [30, 38]. However, that work made use of a very abstract model of diffusion and cannot be regraded as an RD system proper. In the work described in the current article we are interested in exploring the behavior-generating capabilities of an unadorned classical RD system, although an interesting extension we intend to pursue in the future will involve networks of neurons interacting with the excitable medium.

3 The Agent Models

In place of the continuous-time recurrent neural network used by Beer we used a one-dimensional ring of cells within which the concentration of two coupled chemicals changed according to two differential equations describing intra-cell reactions and inter-cell diffusion (Figure 1). Output from whiskerlike proximity sensors was fed to the cells in the RD ring via weighted links, perturbing the concentrations of the two chemicals. Weighted links in turn allowed the concentration of particular chemicals in designated cells to specify motor activation, completing a sensor-motor loop. Links were made symmetrically about
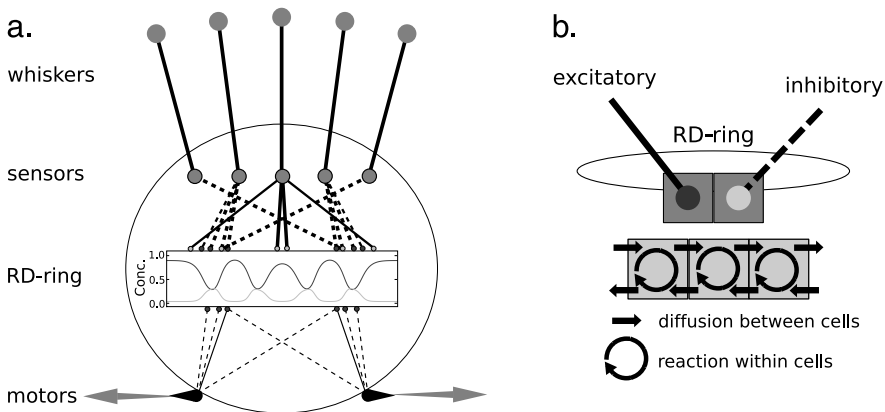

Figure 1. (a) The animat model. Output from the proximity sensors is fed, via weighted links, to the reaction-diffusion (RD) ring, where it perturbs the cellular concentration of chemicals u and v. Solid links increase the chemical concentration in the cell; dashed links decrease it. The effects of any particular link are specific to one of the two chemicals u and v, this specificity being under evolutionary control. Following a number of RD cycles, the chemical concentration levels in designated cells are in turn fed via weighted links to activate the animat's motors. Activation at a motors is summed and multiplied by a constant (10) to produce an output. The combined output of oppositional left and right motors is used to move the animat. (b) Excitatory links from the sensors increase chemical concentration in the cell specified; inhibitory links (dashed) decrease it. In this way the whisker sensors affect the chemistry of the RD ring, which in turn affects the motors.

the animat's longitudinal axis. Parameters specifying the weighted links between cells, motors, and sensors were evolved, as were the values of a dimensionless feed rate and rate constant for the RD system.

3.1 Comparison with ANN Controllers

Within evolutionary robotics and much of Alife, ANN controllers have achieved a position of dominance. There are many flavors to be found, but possibly the most common is the continuous-time recurrent neural network (CTRNN). CTRNNs have been studied fairly extensively [8, 9] and shown to possess rich dynamics that, with parameters suitably tuned over the course of evolution, can mediate minimally cognitive behaviors. As generally conceived, these CTRNNs are spaceless, with neurons influencing each other through highly specific weighted links, unconstrained by such factors as proximity or topological interference. This specificity allows for fine tuning of individual neuron dynamics, both through point-to-point rewiring and through adjustments to the linkage strengths and to intraneuron parameters such as time constants and activation thresholds.

   If we allow, however loosely, the analogy between these ANNs and their biological counterparts, a premise is that the gross functionality of the brain lies in the specific connections between neurons and not the medium in which they reside. But apart from the obvious constraints that spatialization places on biological neural connectivity and development, there is an increasing appreciation that spatial, volumetric signaling plays a part in the function of the brain and that extra-neural properties of the medium are an important element of this [13, 21]. Whereas axons and dendrites bring specificity for neuron communication, volumetric signaling—for example, gaseous diffusion of nitric oxide (NO) [37]—brings a degree of generality, allowing only certain classes of neuron to be susceptible to this global influence. Communication patterns that might be difficult to achieve through neural wiring, such as broad-scale influence of large groups of neurons, are a natural consequence of allowing active properties of the medium to change. One of the main motives for the research described in this article was to investigate the properties of a particularly well-known active medium, namely an RD system, to see whether they could be harnessed and put to some effect as a control system. The complementarity between ANNs and RD controllers suggests that they might work well together, in a system of neural networks residing in and interacting with an active medium. This direction will be explored in future work, along with a more detailed comparison with more standard neural network methods.

3.2 Reaction-Diffusion Models

Perhaps the best-known example of an RD model is that proposed by Alan Turing [46] as an attempt to explain cellular differentiation in early biological development. It is also one of the first examples of the use of a computer to solve differential equations. Turing was trying to understand how the chemicals in arrays, in this case one-dimensional, of identical cells could, by reacting within the cells and diffusing between them, form stable patterns. He was able to show that by constraining the chemical reactions within cells and the relative rate of diffusion between them one could guarantee a stable pattern. Subsequent work has shown analogous systems responsible for leopards' stripes, patternation of nautilus shells, and many other natural patterns.

   Within the class of model RD systems defined by two coupled chemicals (two rate equations), Turing was interested in those tending toward a stable configuration. But by altering the governing reactions and diffusion rates many other systems are possible, displaying a wide variety of spatio-temporal properties. One of the most intriguing is that proposed by Gray and Scott in their 1984 article [28] and extensively analyzed by Pearson in his 1993 article [36]. A variant of the autocatalytic Selkov model of glycolysis [36] the Gray-Scott model corresponds to the following reactions:

$$u + 2v \longrightarrow 3v \tag{1}$$

$$v \longrightarrow p \tag{2}$$

Both reactions are irreversible, so p is an inert product. The system is stoichiometrically conservative in that a feed term for u introduces a nonequilibrium constraint with the feed process removing u, v, and p. This results in the following RD equations, expressed in dimensionless units:

$$\frac{\partial u}{\partial t} = d_u \nabla^2 u - uv^2 + F(1 - u) \tag{3}$$

$$\frac{\partial v}{\partial t} = d_v \nabla^2 v + uv^2 - (F + k)v \tag{4}$$

where k is a dimensionless rate constant and F a dimensionless feed constant. d_u and d_v are the diffusion rates for the two chemicals (see Section 4 below for specific details). A trivial steady state $u = 1$, $v = 0$ exists for all values of F and k. The Gray-Scott simulation proves very robust, showing no qualitative difference when implemented by forward Euler integration over a broad range of spatial and temporal scales [36].

When suitably perturbed, the Gray-Scott simulation exhibits a large variety of spatiotemporal patterns that have to be seen to be appreciated. Pearson's article is replete with beautiful images, but the simulation is best appreciated in real time with a 2D diagram and a suitable color map. By fixing the diffusion rates of the chemicals and using F and k as control parameters, Pearson was able to show that within suitable limits the 2D phase diagram described shows regions associated with specific spatiotemporal patterns, ranging from spot replication and stripes in a continuous transition to traveling waves and spatiotemporal chaos.

3.2.1   Other Reaction-Diffusion Models

In order to investigate the degree to which the results were dependent on specific properties of the Gray-Scott model and to test the ability of the evolutionary method to exploit spatiotemporal dynamics in alternative systems, the experiments described below were repeated with another classic modeled chemical system, the two-variable *Oregonator*, designed to approximate the Belousov-Zhabotinsky (BZ) reaction [22]. The Oregonator is arguably a more plausible biochemical model, with patterns of target and re-entrant spiral waves analogous to those seen in such natural phenomena as cardiac rhythms or forest fires.

RD controllers based on the Oregonator were able to complete all the tasks described below save that requiring the animat to maintain a memory. It appears that, in the case of a connected ring of cells, this failure is accounted for by the difficulty the BZ reaction has in forming semistable, stationary one-dimensional waves.

3.3   Visually Guided Agents

The choice of an evolved animat model, for example to demonstrate the potential of a novel RD controller, should be informed by two key considerations. The behavior in question must be cognitively interesting, and there should be a reasonable expectation that resultant controllers can be analyzed and understood:

> The term "minimally cognitive behavior" is meant to connote the simplest behavior that raises cognitively interesting issues.
> Generally speaking, visually-guided behavior provides an excellent arena in which to explore the cognitive implications of dynamical and adaptive behavior ideas, since it raises a host of issues of immediate interest. [7: 422]

In keeping with Beer's thesis, we chose for our memory test a visual-guidance task conforming to the requirements of minimal cognition. After priming by an arbitrary signal **s**+, a whiskered animat, capable of moving in one dimension along the bottom of a vertical plane as illustrated in Figure 2, is
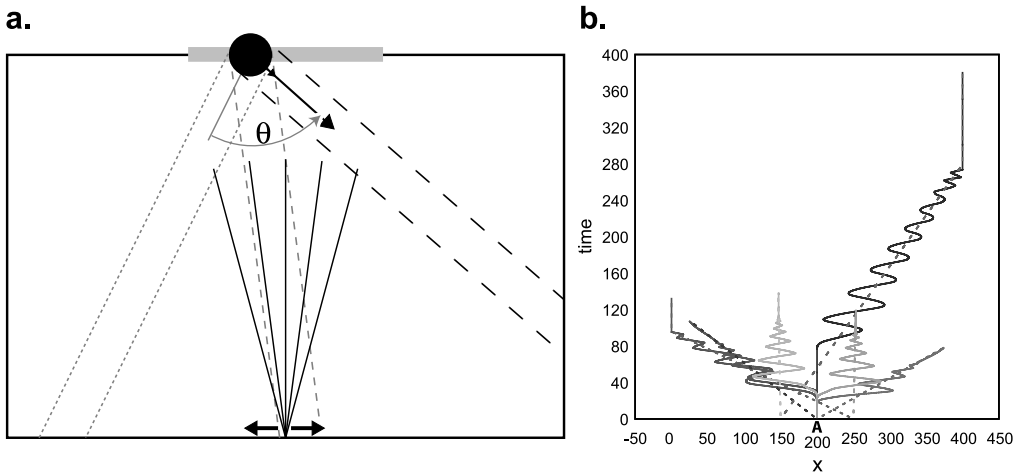
a.

b.

Figure 2. (a) The fixation experiment (to scale). An animat with five whiskers spread over a 30° span is placed at the center of the arena's floor. During a trial a circular object was placed in the starting zone (gray line) on a trajectory within the limits defined by the animat's distal whiskers (θ). This was to ensure that the animat received some stimulation from the falling object. The object's speed is indicated by the relative position of large and small arrows, lying between 0.5 and 7 units per second. The end of the trial was signaled by the object reaching the arena floor, at which point the distance between object and animat was used, along with their relative start points, to calculate a fitness score for the animat. (b) Six animat-circle trajectory pairs, with the circle's trajectory dashed.

required to orient toward and track a circular object falling from the top of the plane with a large range of vertical and horizontal speeds. In the absence of **s**+, or if **s**+ is followed by the second signal **s**−, the animat is required to avoid the falling object. (see Figures 4 and 5 later in this article for details).

Beer evolved CTRNNs to control his animats—a control system one of the authors has some experience of [19]. Beer's subsequent analysis [8] of the CTRNNs' dynamics makes them probably the best-understood of all animat controllers, evolved or otherwise. This represents a useful benchmark and an obvious model to emulate. The use of such canonical models to provide a common point of reference would seem to be an efficient way to exploit the resources available. Broadly speaking, this work preserves the details of Beer's model while replacing the CTRNN controller with a novel one using an RD medium.

3.4 Evolving Controllers

The Gray-Scott model, in keeping with most RD systems, is highly nonlinear, at least unintuitive, and often counterintuitive.[1] It is not immediately clear how one could hand-wire such a controller, but it would require an intuition about the rich dynamics of the system that escapes us. In cases such as this, where we require a controller capable of exploiting even a relatively simple dynamic system, it would seem that the need is pressing to leverage the increasing computer power at our disposal and automate the process of discovery. This approach is particularly appropriate to a robot that is intended to remain in silico. The search algorithm employed here is a genetic algorithm (GA). A simplistic, but initially useful, way of understanding how a GA works is to picture the parameter space, describing in this case the details of our RD controller, such as linkage points and weights, as a fitness landscape. Every point in this landscape describes an animat controller, and height above ground corresponds to fitness. If the landscape is reasonably well ordered, it should be possible for the GA to find its way from low ground initially, corresponding to randomly wired, poorly performing controllers, to high ground, where the controllers are (much) better performing. This image leaves out important details, particularly the concept of neutral networks,[2] but the key detail is captured. From random parameters and allowing for a suitable encoding

---

1 The speed of modern processors makes it possible to interact in real time with 2D implementations of these RD systems. Having implemented and played with just such models (among others, the Gray-Scott one), we can attest to its counterintuitiveness.

2 A complex subject, highlighting our poor intuition of movement in higher-dimensional space.

scheme, it should be possible to automatically produce good controllers by applying evolutionary pressure. The work described in this article and elsewhere [7, 8, 17, 18] is a testament to that fact.

4 Method

4.1 The Animat Model

To a large extent details from Beer's earlier simulations [7] were preserved and the required behaviors were essentially the same. The arena was 400 units long by 275 units high (Figure 4, discussed later) in all the experiments. The animat's five whisker sensors were 220 long and uniformly spaced over a $30°$ spread. Activation of the whiskers was a simple linear function with a minimal value of 0 when the whisker was unimpinged and 1 when it was intersected at the base.

Figure 1 shows a diagram of the animat. Activation from the sensors $\in [0, 1]$ was fed through weighted links $\in [-1, 1]$ to the one-dimensional RD ring, consisting of 128 cells subject to intracellular reaction and intercellular diffusion between near neighbors [see the chemical reactions (1) and (2) and the rate equations (3) and (4)]. The weighted links were specific to either chemical u or v; this specificity being under evolutionary control.

The sensors, motors, and input to the RD ring were updated using the forward Euler method with an integration step size of 0.1. During this time step each cell in the RD ring was updated twice, using discretized versions of the rate equations (3) and (4). At each time step (0.1), input via links to the cells perturbed the specified chemical's concentration. The sensor activation was in the range (0, 1), and the links' weights in the range $(-1, 1)$. The cellular concentration of u and v was bounded within the range [0, 1].

The animat's motors received input from cells in the RD ring. Input from an individual link was a product of a link weight $\in [-1, 1]$ and the concentration of the evolutionarily specified chemical in the cell. To update the animat's position, the activation of the oppositional motors was subtracted (right − left), and the result multiplied by 10. This multiplier was fairly arbitrary, taking into account the need for the animat to move fast enough to catch objects with a maximal horizontal velocity around 5. It worked well enough, but is probably too large. On reflection this value should probably have been an evolutionarily specified parameter, but given the fitness scores generated, any gains could only have been very marginal.

4.2 The Reaction-Diffusion Ring

The diffusion rates d_u and d_v were fixed at the standard values [36] of 2×10^{-5} and 10^{-5}, respectively, and the circumference of the RD ring was 0.32. Each animat genotype specified a value for the rate constant k and feed constant F (Equations 3 and 4), which were seeded at values 0.055 and 0.02, respectively, in the otherwise randomly generated initial populations. By moving through this F, k parameter space, evolution had some control over the properties of the RD system (see Section 3.2 above).

4.2.1 Simulating a Reaction-Diffusion System Using a 1D Cellular Automaton

In order to approximate a continuous chemical system, the ring was divided into 128 discrete cells, each with concentrations $\in [0.0, 1.0]$ of chemicals u and v. Using the two differential equations [(3) and (4)], Euler integration with a step size of 0.1 was applied at each time step to advance the system iteratively:

1. For each cell n, we calculate the change in chemical concentration C produced by diffusion across the boundaries of neighboring cells $n-1$ and $n+1$:

$$dC_n = D(C_{n-1} + C_{n+1} - 2C_n) \tag{5}$$

where D is a cell's density, calculated with respect to the spatial resolution of the simulation. This approximates the Laplacian operator ∇^2 in the differential equations 3 and 4.
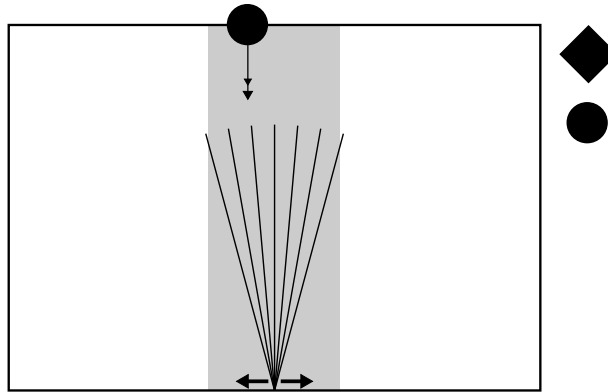
Figure 3. The discrimination experiment (to scale). An animat with seven whiskers spread over a 30° span is placed at the center of the arena floor. During a trial, an object was placed at the top of the arena within the gray drop zone on a straight downward trajectory at between 3 and 4 units per second. A pair of trials with the same starting position and speed was performed for diamond and circle objects. The animat was rewarded for its ability to fixate on the circle and avoid the diamond.

2. The remaining reaction component of the equations is calculated using the current chemical concentrations and the constants F and k.

3. For each cell reaction and diffusion, components of the differential equations are added to the current chemical concentrations, and the results thresholded $\in [0.0, 1.0]$ to give new concentrations.

4. This process is iterated over the duration of the trial.

4.3 The Genetic Algorithm

The GA consisted of a population of 30 animat genotypes, which were updated generationally according to rank-based selection. The genotypes were essentially a list of weighted, chemically specific links, describing the wiring of an animat controller. As the animat controllers were symmetrical, each link on the list corresponded to two links on the controller. At each generation these lists were converted into their respective animat controllers and assigned a fitness value according to how well the controller performed its task. It was neither practical nor desirable to have the genotype describe a fully connected controller (1,408 links in all), so the number of links was preset. The starting numbers for the orientation experiment were eight sensors → RD ring, four RD rings → motor, making 24 symmetrically arranged links in all.

At the end of each generation a new generation was formed from the old and subjected to mutation. The numbers on the genotype were in the range [0, 1], being mapped onto their respective controller parameters. Mutation consisted of the addition of a normally distributed random value with mean 0 and standard deviation 0.25. A second mutation operator was applied to each genotype with a probability of 10%, randomly deleting a link from or adding a link to the list. The link-addition operator allowed two links to share start and end points and chemical specificity.

The same fitness function $f(d_s, d_e)$ was used to evaluate all trials (see Figures 3 and 5, discussed later), where the two values d_s and d_e specify the absolute horizontal distance between animat and the shape at the trial's start and end, marked by the shape reaching the arena floor, respectively:

$$f(d_s, d_e) = \begin{cases} 1 - \frac{d_e}{d_s} & \text{if } d_e < d_s \\ \max\left(\frac{d_s - d_e}{50}, -1\right) & \text{if } d_e \geq d_s \end{cases}$$
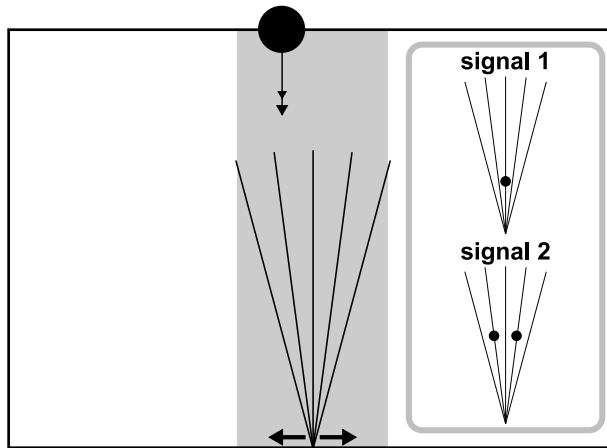
Figure 4. The memory experiment (to scale). An animat with five whiskers spread over a 30° span is placed at the center of the arena floor. During a trial, a circle was placed at the top of the arena within the gray drop zone on a straight downward trajectory of speed between 3 and 4 units per second. Prior to the circle drop the animat received either signal 1 or signal 2 or signals 1 and 2 consecutively. The signals consisted of an arbitrary pattern applied to the animat's whiskers, after which the system was allowed to settle. The animat was rewarded for its ability to reverse behavior on receiving signal 1 — for example, switching from a circle fixator to a circle avoider.

The value of this fitness function is highest if the animat fixates the object centrally, and lowest if the animat avoids the object, having a maximum at distance 50. Where the trial required the animat to avoid the falling object (T1, T3, T4 in Figure 5, discussed later), the resultant fitness was multiplied by -1.0.

To evaluate the animat's performance at the memory task we used an amalgam of the fitnesses $F1-4$ over four trials $T1-4$. The animat was required to show the opposite behavior over T2 to the other trials, dependent on a prior stimulus. The four trials were allotted a fitness for orientation toward the circle and the total fitness was calculated as follows:

$$f(F1, F2, F3, F4) = (F1 - F2) \times (F3 + F4)$$

This function was designed to encourage a switching of behavior over T2 while avoiding evolutionary local minima[3] that might result by a simple adding of $F1-4$. The function does not specify whether T2 should show avoidance or fixation behavior, only that it is opposite to that seen in the other trials. To disambiguate these two possibilities the amalgamated fitness was multiplied by -1.0 in those cases where the animat showed aversion to the object in trial T2.

4.4   Training Protocol

A number of trials were conducted to assess the ability of individual controllers and allot their respective genotypes a fitness score. Figure 4 shows the trial setup for the orientation experiment. The gray drop zone delimits the possible circle trajectories during a trial. The object trajectories were constrained so as to ensure some whisker stimulation for the animat.[4]

4.4.1   Orientation

A number of trials were conducted to assess the ability of individual controllers and allot their respective genotypes a fitness score. Figure 2(a) shows the trial setup for the orientation experiment. The gray

---

3 Where the genetic algorithm drives the population of genotypes to a suboptimal region of the fitness landscape, wherein they are trapped by evolutionary pressure.
4 In keeping with Beer's original model [7], the simulation was noiseless, meaning that the symmetrical animat controller was incapable of breaking symmetry without stimulus from the whiskers.

zone at the top indicates potential starting points for the circle during a trial, and θ the limits within which the trajectory line is set. In keeping with Beer's original model [7], the simulation was noiseless, meaning that the symmetrical animat controller was incapable of breaking symmetry without stimulus from the whiskers. The trajectories of falling objects were controlled to make sure the animat received some sensory stimulus during the course of a trial.

4.4.2 Discrimination

To test a controller for the ability to orient toward the circle, the gray starting zone (± 70 from the animat's starting point $x = 200$) was broken into four equally spaced regions, and two starting points were chosen randomly from each region. Two random trajectories within the limits set by θ were chosen for each starting point, and each of these tested at two random speeds, within the range [0.5, 7]. This makes a total of 16 trials per assessment. Two assessments were carried out, and the average was returned as a fitness score.

Figure 3 shows the setup for the diamond-circle discrimination trials. Animat performance with dropped circles and diamonds was compared for eight random points starting within the drop zone (the gray region ± 50 from the animat's starting point $x = 200$) and two random speeds $\in [3, 4]$ per point, making a total of 32 trials in all for a single assessment. Two assessments were made for each controller, and the average returned as a fitness score for the genotype.

4.4.3 Memory

Motor feedback to the RD ring can be used by the animat to stabilize behavior. Purely diffuse controllers were able to use motor feedback in this way to maintain a memory trace. To prevent this exploitation, in the memory experiment, the animats were not allowed to form links from motors to the RD ring over the course of evolution. A second consideration is the length of time between the animat receiving one of the priming stimuli **s**+ and **s**− and its response to the falling object. The limits on this random wait time $\in [400, 600]$ were set to oblige the animat to use the reaction between u and v to sustain a memory.

Figure 5 shows the four trials T1−4 that constituted a single fitness test:

T1   No signal

T2   Signal **s**+

T3   Signal **s**−

T4   Signals **s**+ followed by **s**−

The animat was required to avoid the falling object in all trials but T2, where, after receiving signal **s**+, it was required to fixate on the circle.

5 Results

5.1 Orientation

Animats capable of orienting toward and fixating on a falling circle were easily evolved within 200 generations on all runs, almost invariably with close to optimal fitness for the best members. Perhaps not so surprising, given the symmetrical nature of controller and task, it is nevertheless remarkable how many randomly generated animat controllers were able to achieve respectable scores ab initio. It is far too early to say, but it appears that the symmetrically wired RD rings have a natural tendency to resist perturbation and normalize left and right sensor input. It should be stressed that for the harder discrimination task no animat populations showed any ability ab initio.

Figure 2(b) shows the tracking performance of a typical animat for three pairs of random circle trajectories, with evenly spaced start points. This animat has been chosen as being truly representative of the best solutions found in all runs: Similar mechanisms to those described were found in the vast
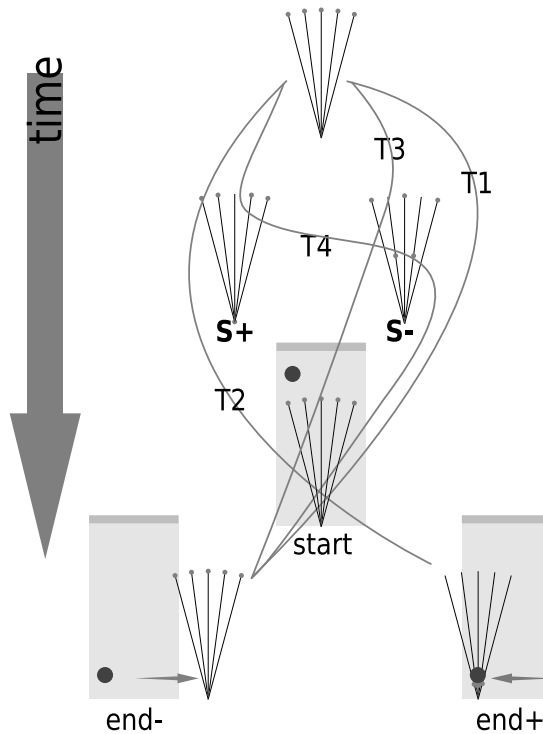
Figure 5. The four trials (T1 – 4) providing the components of the memory task's fitness function. **end+** and **end−** show the desired end positions of the animat relative to the falling object. The duration of all phases is in the range [400, 600], and the stimuli last for 10 time units. T1: In this trial, the animat receives no signal. The falling circle should elicit an aversion response. T2: In this trial, the animat receives the **s+** signal. As a consequence the animat should fixate on the falling circle. T3: In this trial, the animat receives the reset signal **s−**. This should not affect the animat's aversion to a falling circle. T4: In this trial, the animat receives two signals: the priming **s+** followed by **s−**. The second **s−** signal should reset the animat's response, causing it to avoid the falling circle.

majority of those animats analyzed. The dashed lines show circle trajectories, and the solid lines show the animat's response. Starting at $x = 200$, the animat remains stationary until its whiskers are stimulated by the falling object. It then oscillates back and forth below the falling object, continually overshooting its position. These oscillations are damped as the object gets closer to the animat and the sensor stimulus increases. The plots resemble attractor cycles converging on the intersection of the animat and the falling shape at the ground. This animat controller had sixteen links in total. After an initial 200 generations, during which near-optimal performance was achieved, the animat was evolved for a further 1000 generations with the link-addition operator turned off but the link-deletion operator still active. In this way evolution was prevented from exploring bigger networks while being able to test smaller ones. This technique was found to rapidly reduce the size of networks while maintaining their fitness and has the added advantage that it avoids the need to introduce an arbitrary component into the fitness function to encourage simplicity. A major motivation for encouraging evolution to discover simpler networks was the analytic tractability of solutions thus found.

5.2 Discrimination

Animats capable of distinguishing between circles and diamonds were readily evolved within 200 generations (see Figure 3 for details of the task). From a typical batch of 20 runs at least a third achieved fitnesses greater than 1.8 (from a maximum of 2) over 32 random trajectories. Much better results were achieved if the animats were trained initially to fixate on a circle before being introduced to the discrimination task.

Figure 6 shows the network of a typical near-optimal animat after the removal of most redundant links by an extended period of evolution with the addition mutator turned off [see Figure 6(b) for details]. A total of 26 links are used, with fairly even sampling of the seven available whisker sensors. Six links from the RD ring to the motors dictate the animat's movement. The inner two links have converged on the same cell, allowing the cell to strongly influence the animat's movement by exciting one motor while inhibiting its opposite. Interestingly, although they allowed feedback from the motors to the RD ring, these links have been pruned away as unnecessary. Again, this animat was chosen as representative of the majority of near-optimal solutions examined; most used very similar mechanisms.

5.3 Memory

Using a fitness of 3.7 (out of the maximum of 4.0) as an acceptable solution, animats using the Gray-Scott RD controllers and capable of completing the memory task proved easy to evolve. The increased trial length and number of generations limited us to 10 evolutionary runs, of which four produced an acceptable solution within 500 generations. All successful RD controllers used a semi-stable standing wave to preserve a chemical memory trace over the period between training stimuli and the presentation of falling circles. In this sense, the animat discussed below is representative of this successful group. In order to make the animat easier both to analyze and to describe, the population of animats was further evolved to reduce the size of the controllers while preserving fitness [see Figure 6(b) for an explanation of the pruning method]. While this further period of evolution was very effective in reducing the number of links used by the RD controller, it did not qualitatively alter the solution.

5.3.1 A Chemical Switch

The evolved model, with 22 symmetrically arranged links, is summarized in Figure 7. Figure 7(a) shows the switching cycle that allows the animat to flip from state **M−** to state **M+** and back again in response to signals **s+** and **s−**. Clockwise from **M− s+**, a maximum stimulus of the central whisker increases, via v-specific positive links, the concentration of chemical v at four cells in the RD ring. This establishes autocatalytic waves, as u and v react and diffuse, which roughly stabilize at **M+**. The
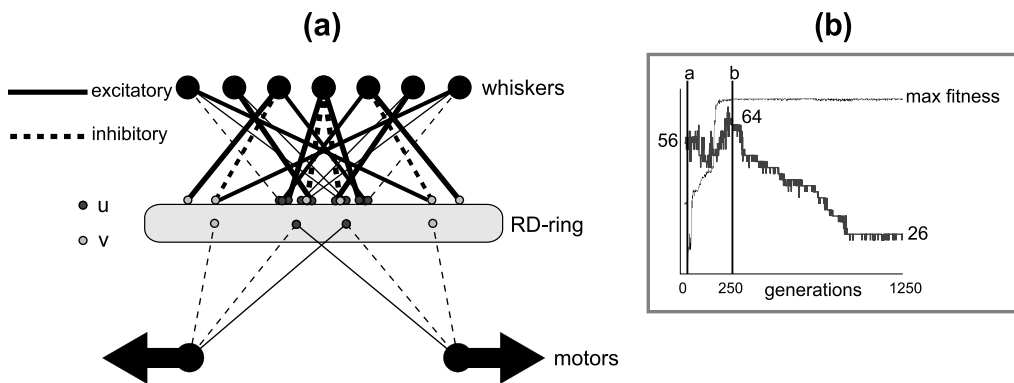

Figure 6. (a) The complete reaction-diffusion controller for an animat capable of discriminating diamond and circle objects. Colored nodes are used to indicate the chemicals affected by the whisker sensors and the chemicals affecting the motors, respectively. This network has no feedback from motors to the reaction-diffusion ring. (b) Simplifying the network of a diamond-circle discriminator. The figure shows superimposed plots for the number of weighted links in the reaction-diffusion controller (thick line) and the fitness of the best individual in the population (thin line). After a short period during which the animat is evolved to fixate on a circle, the fitness drops as the animat is introduced to the diamond-circle discrimination task (line a). At this point, the controller consists of 56 weighted links. Over 250 generations, the GA is allowed to randomly insert and delete links from the controller genotype until, at line b, the network has 64 links and the animat has achieved near-optimal fitness. At this point, the GA's insertion mutator is turned off while continuing to allow evolution to delete links randomly. Over the next 1000 generations, the animat maintains a close to optimal fitness while losing links, ending with 26 links, less than half the number it started with.
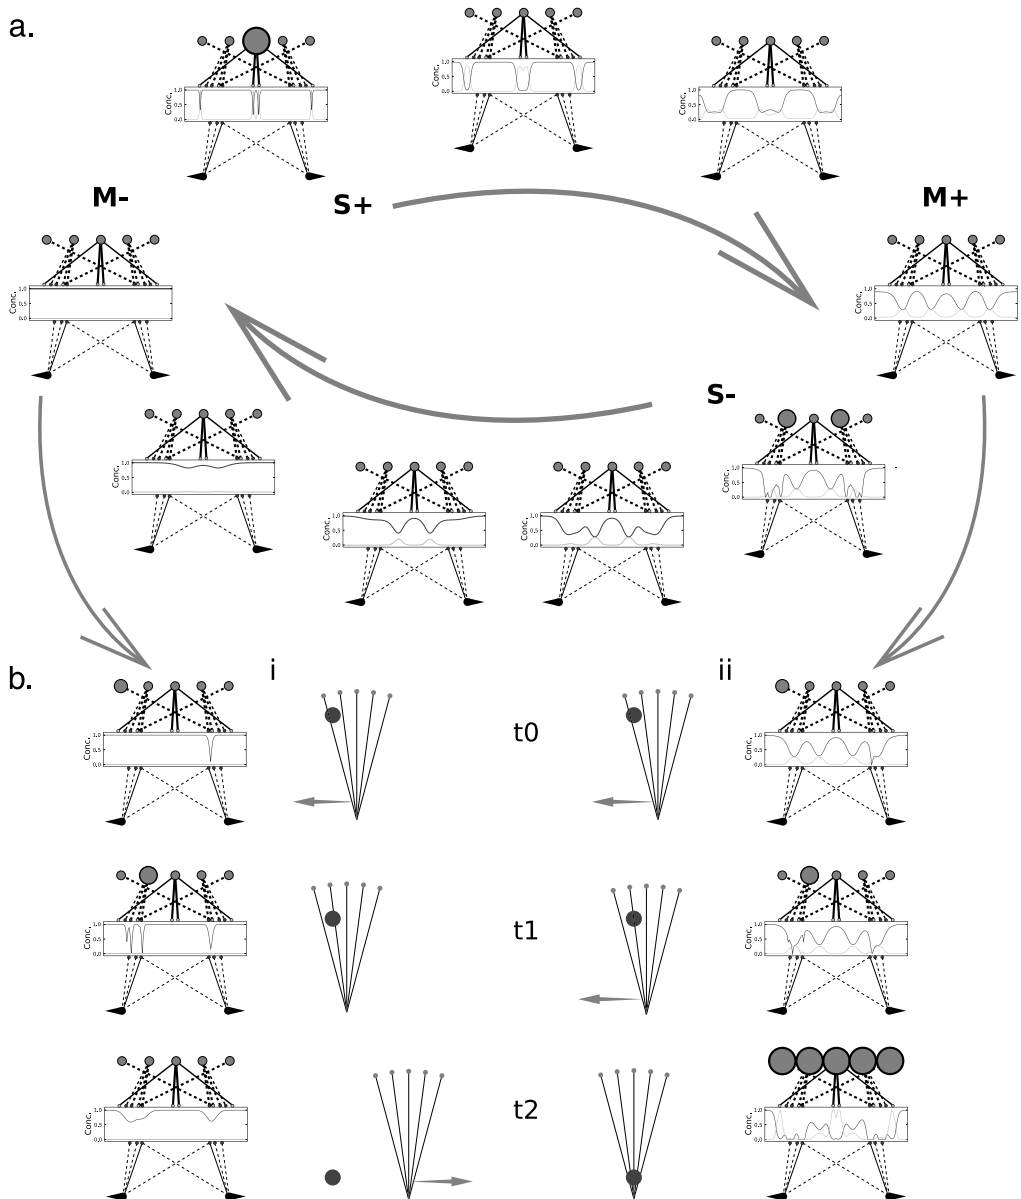
Figure 7. (a) A chemical memory switch. Moving clockwise from the default settled state **M−**, the animat receives the stimulus **s+**, causing an auto catalytic cycle, which leads to the semi-stable state **M+**. While in state **M+**, the resetting signal **s−** disrupts the RD ring's structure, causing it to return to **M−**. (b) The response of the animat, for states **M+** and **M−**, to a falling object. (i) In the default settled state **M−** the animat displays aversion to the falling circle. (ii) In the stimulated state **M+** the same falling circle elicits a fixation response.

application of **s−**, a half-maximum stimulus of the first proximal whiskers, at **M+** disrupts these waves, bringing the chemical system back to **M−**.

Figure 7(b) shows the response of the animat to the same falling object trajectory while in default state **M−** (i) and primed state **M+** (ii). In **M−**, stimulation of the leftmost whisker reduces the concentration of v via an inhibitory connection. Motor links, sensitive to changes in v's concentration, unbalance the left and right motors, drawing the animat toward the falling object. This behavior at t0 is mirrored in **M+** [Figure 7(b)ii]. At time t1 the behaviors start to diverge, in response to stimulation of

the whisker second from the left. In Figure 7(b)i this engages a strong avoidance response in the animat, causing it to move quickly away from the object. The animat primed by **s+** does the opposite, moving the object toward its center, where, thus fixated, it remains through the course of the trial.

The dependence of the memory state **M+**, induced by signal **s+**, on interaction between chemicals *u* and *v* is highlighted in Figure 8. In the absence of a reaction component, the memory is not established. Purely diffusive controllers were thus unable to evolve a solution to this task.

5.3.2   The Stability of the Chemical Memory

As shown in Figures 8 and 9, the memory trace produced by signal **s+** is unstable. The memory is long enough to enable the animat to pass the test requirements, its duration being roughly defined by the upper bound of the randomized time between the application of a stimulus and the first engagement of the falling circle with the animat's whisker. In Figure 9, we focus on the change in concentration of two cells over the course of the memory trace. These cells are roughly at the centers of two of the four symmetrical peaks and troughs of *v* and *u*, respectively (see Figure 7, **M+**). Figure 9(b) shows that at these points in the chemical ring, the interaction of *u* and *v* makes them describe orbits around an area of the phase-space. At time t1, this orbit cannot be sustained and the cell c23 returns to default. The destruction of the c23's orbit presages the destruction of the longer cycle of c49 at time t2. The dependence of c49's cycle on c23's emphasizes the global nature of this memory.

5.3.3   Extending the Memory's Duration

In order to see whether the duration of the animat's memory trace was adequate to the requirements of the task, the animat population was further evolved (see Section 3.4 above), under standard conditions, while the time between the application of stimuli and the dropping of the object was gradually increased, requiring the animat to maintain a longer memory of the stimuli. Figure 10 shows the
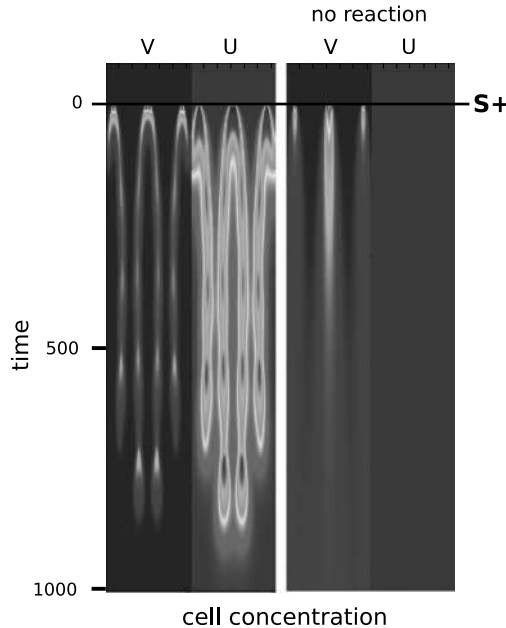

Figure 8. The changes in concentration over time of chemicals *u* and *v* in response to the signal s+. In the trace on the right, the reaction component of the RD system has been disabled. When allowed to react, the two chemicals maintain a strong, autocatalytic memory trace for approximately 750 time units, long enough to reliably remember the signal and score high on the task. In the absence of reaction between *u* and *v*, initially high concentrations of *v* diffuse away while *u*, unable to autocatalyze, shows no change in activity; the links to the animat's motors are all *u*-specific, so in the absence of reaction, the animat is paralyzed.
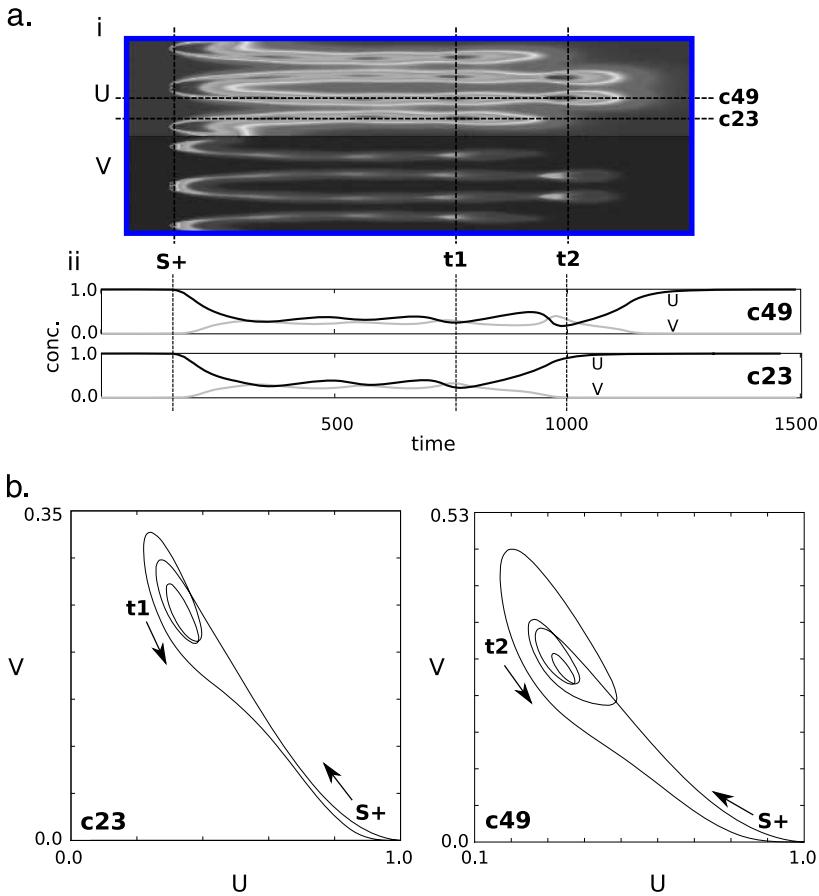
Figure 9. The change in u and v concentrations over time for the two RD ring cells c23 and c49 in response to the animat receiving stimulus s+. Two other significant times are marked, t1 and t2, the return of c23 and c49 to default. (a)i shows the concentrations of chemicals u and v, with the positions of cells c23 and c49 indicated, for u, by dashed horizontal lines. (a)ii shows the change in concentration over time for chemicals u (dark gray) and v (light gray) in cells c23 and c49. (b) Plotting the activation of the two chemicals u and v against each other reveals the unstable attractor cycles at c23 and c49 that characterize the memory trace. The orbits, though unstable, are maintained throughout the course of the trial, allowing the stimulated animat to respond differently to the falling object.

result of a successful evolutionary run that required the animat to maintain a memory five times longer than that of the original task, whose trace is shown to scale in Figure 10(a)i. At this point, the simulations became impracticably long, but there was no indication in this or in other simulations of a hard constraint on the possible duration of memory. Note that the interaction of u and v now makes them describe many more and tighter orbits in phase space, maintaining the memory trace for much longer. It should be stressed here that the necessary limits to spatial and temporal resolution in the simulated chemical RD ring probably play a part in the precise characteristics of the memory trace and that these results should be interpreted qualitatively.

Full analysis of the network will take some time but some provisional observations can be made. Figure 11 shows two trajectories made by the animat controller shown in Figure 6 in response to a diamond and circle dropped at the extreme end of its evolved range (see Figure 3) at speed 3.5 and along line s. The left plot shows the animat avoiding the diamond (dashed line) while fixating on the circle (solid line). The middle two plots show the concentration of chemical in the RD ring's cells over the course of the two trajectories, and the plot on the right is the result of subtracting the two middle plots. This plot shows that following a period of relatively similar activity, during which time the animat is fixating on both the diamond and the circle, there is a large difference in activity at time 56 around cells
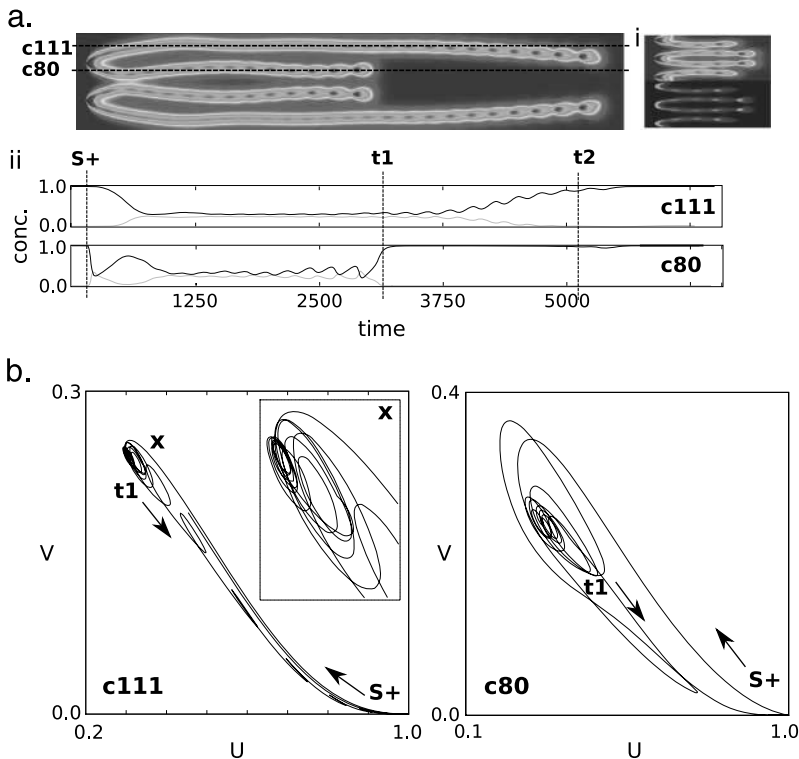
Figure 10. The reaction of two selected cells, c80 and c111, to stimulus **s+** after a further period of evolution, wherein the animat was required to increase the length of its memory. The top plot of (a)i shows, to scale, the previous memory trace following **s+**. (b) The attractor cycles are now more densely packed and greater in number.
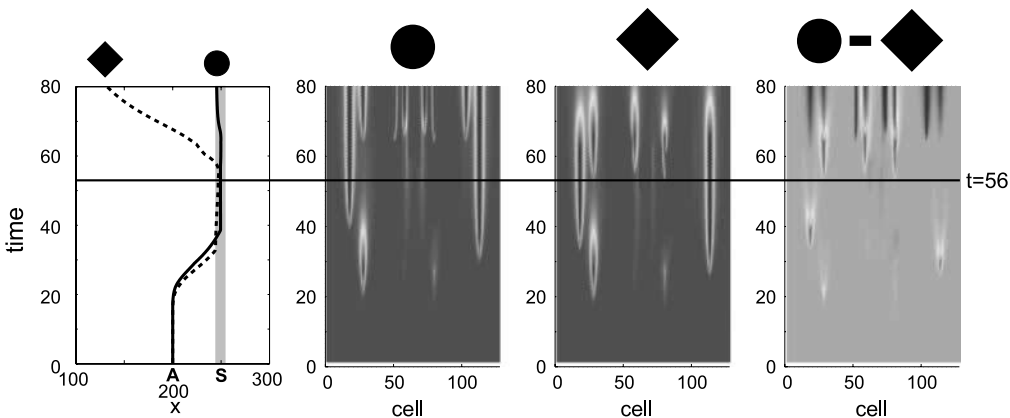

Figure 11. The leftmost plot shows the *x* position over time for two superimposed animat trajectories, both starting at *A* (*x* = 200), in response to a shape dropped from *S* (*x* = 250) at speed 3.5. The dashed line shows the animat's response to a dropped diamond; the solid line, to a dropped circle. The middle plots show activation of the reaction-diffusion cells during the circle and diamond trials. The rightmost plot shows the result of subtracting the activations of circle and diamond trials. Around time 56, a strong difference is seen in chemical concentrations in the reaction-diffusion rings, corresponding to the animat's evasion of the diamond (left plot). In this sense the right plot shows traces of the animat's active shape detection.

27 and 79, these falling under the remit of the rightmost whisker. While the circle-plot concentrations remain symmetrical along the ring, an asymmetry can be seen in the diamond plot. This corresponds to the initiation of a strong leftward avoidance tactic, taking the animat over a hundred units from the diamond at trial's end (maximum fitness is achieved for a distance of 50 units). Observing the diamond trajectory, it is clear that the animat's attempt to fixate on the diamond as it does the circle causes it to keep the diamond to the right of its central line, introducing an asymmetry as the right whisker stimulation is unbalanced by the left. This loss of balance initiates avoidance.

5.4 Other Reaction-Diffusion Systems

The orientation, discrimination, and memory tests were repeated using the two-variable Oregonator model, designed to approximate the classic Belousov-Zhabotinsky (BZ) reaction [22]. Solutions for the orientation and discrimination tasks were easily evolved, but the BZ-based controllers failed to find solutions to the memory task. This suggests that in the case of this ringed cellular automaton, properties of the Gray-Scott system, missing in the Oregonator, enabled the creation of a memory trace.

The results above (Section 5.3) show that the ability of the Gray-Scott controllers to perform the memory task relied on the formation of a semi-stable, stationary wave, which was maintained by the RD dynamics over the course of a memory trial. No such standing wave was observed in the evolved BZ-based controllers. Linear stability analysis of such standing waves is a difficult problem, but can be done analytically in the case of the Gray-Scott model [48], the localized solutions resembling those seen in the successful memory controllers. Experiments with 2D BZ systems have shown that under the right conditions they can serve as writable memory devices [48], but less work has been done with 1D models. Steady state solutions can be achieved using specially selected parameters [51], but it appears that the chance of discovering such configurations is significantly less than in the Gray-Scott case.

6 Discussion

Although Cajal's neuron doctrine is predominant in cognitive studies and, by definition, in neuroscience (computational or not), it does raise a very big question. How does a single-celled animal, of a set representing the larger biomass of the animal kingdom and those evolutionary precursors of all multi-cellular lifeforms (including ourselves), negotiate its world and engage cognitively with it? Any explanation cannot involve neurons—single cells in themselves—but must explain how a seemingly homogeneous blob of chemicals can produce robust behavior and exhibit the classical learning models. We would not suggest that the models described in this article hold any answers to the larger questions of animal cognition, but the ability of these simple chemical systems to mediate simple cognitive tasks requiring memory is intriguing. Extending these model systems—used extensively and successfully in biology to explain such phenomena as cardiac rhythmia, animal patternation, and morphogenetical development—to the cognitive realm could prove fruitful.

Acknowledgments

References
1. Adamatzky, A., Arena, P., Basile, A., Carmona-Galan, R., De Lacy Costello, B., Fortuna, L., Frasca, M., & Rodriguez-Vazquez, A. (2004). Reaction-diffusion navigation robot control: From chemical to VLSI analogic processors. *IEEE Transactions on Circuits and Systems*, *51*(5), 926–938.

2. Adamatzky, A., & De Lacy Costello, B. (2003). Reaction-diffusion path planning in a hybrid chemical and cellular-automaton processor. *Chaos, Solitons, Fractals*, *16*, 727–736.

3. Adamatzky, A., De Lacy Costello, B., Melhuish, C., & Ratcliffe, N. (2003). Experimental reaction-diffusion chemical processors for robot path planning. *Journal of Intelligent Robotic Systems*, *37*, 233–249.

4. Arena, P., Cruse, H., & Frasca, M. (2002). Cellular nonlinear network-based bio-inspired decentralized control of locomotion for hexapod robots. *Adaptive Behavior*, *10*(2), 97–111.

5. Arena, P., Fortuna, L., & Branciforte, M. (1999). Reaction-diffusion CNN algorithms to generate and control artificial locomotion. *IEEE Transactions on Circuits and Systems I*, *46*(2), 253−260.

6. Barandiaran, X., & Moreno, A. (2006). On what makes certain dynamical systems cognitive: A minimally cognitive organization program. *Adaptive Behavior*, *14*(2), 171−185.

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

8. Beer, R. D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, *1*(4), 209−243.

9. Beer, R. D. (2003). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, *1*(3), 469−510.

10. Beer, R. D., Quinn, R. D., Chiel, H. J., & Ritzmann, R. E. (1997). Biologically-inspired approaches to robotics. *Communications of the ACM*, *40*(3), 30−38.

11. Breyer, J., Ackermann, J., & McCaskill, J. (1998). Evolving reaction-diffusion ecosystems with self-assembling structures in thin films. *Artificial Life*, *4*(1), 25−40.

12. Britton, N. F. (1986). *Reaction-diffusion equations and their applications to biology*. London: Academic Press.

13. Changeux, J.-P. (1993). Chemical signalling in the brain. *Scientific American*, *269*(5), 58−62.

14. Chua, L., & Yang, L. (1988). Cellular neural networks: Applications. *IEEE Transactions on Circuits and Systems 1*, 35, 1273−1290.

15. Chua, L., & Yang, L. (1988). Cellular neural networks: Theory. *IEEE Transactions on Circuits and System 1*, 35, 1257−1272.

16. Cruse, H., Kindermann, T., Schumm, M., Dean, J., & Schmitz, J. (1998). Walknet—A biologically inspired network to control six-legged walking. *Neural Networks*, 11, 1435−1447.

17. Dale, K. (2006). Evolving reaction-diffusion controllers for minimally cognitive animats. In *Proceedings of the International Conference on Simulation of Adaptive Behavior (SAB), 2006* (pp. 498−509).

18. Dale, K. (2008). A model chemical memory in an evolved animat. In S. Bullock, J. Noble, R. Watson, & M. A. Bedau (Eds.), *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems* (pp. 142−149). Cambridge, MA: MIT Press.

19. Dale, K., & Collett, T. S. (2001). Using artificial evolution and selection to model insect navigation. *Current Biology*, 11, 1305−1316.

20. Dockery, J., Hutson, V., Mischaikow, K., & Pernarowski, M. (1998). The evolution of slow dispersal rates: A reaction diffusion model. *Journal of Mathematical Biology*, 37, 61−83.

21. Edelman, G. M., & Gally, J. A. (1992). Nitric oxide: Linking space and time in the brain. *Proceedings of the National Academy of Sciences of the USA*, 89, 11651−11652.

22. Field, R. J. (1975). Limit cycle oscillations in the reversible oregonator. *Journal of Chemical Physics*, *63*, 2289−2296.

23. Floreano, D., Husbands, P., & Nolfi, S. (2008). Chapter 61: Evolutionary robotics. In B. Siciliano & O. Khatib (Eds.), *Springer handbook of robotics* (pp. 1423−1451). Berlin: Springer.

24. Fuqua, C., Glazier, J., Brun, Y., & Alber, M. (2004). Proceedings of Biocomplexity VI: Complex behavior in unicellular organisms. *Biofilms*, *1*(4).

25. Gerhart, J., & Kirschner, M. (1997). *Cells, embryos and evolution*. Cambridge, MA: Blackwell Science.

26. Goldbeter, A. (1997). *Biochemical oscillations and cellular rhythms: The molecular bases of periodic and chaotic behaviour*. Cambridge, UK: Cambridge University Press.

27. Goldstein, L. (2001). Kinesin molecular motors: Transport pathways, receptors, and human disease. *Proceedings of the National Academy of Sciences of the USA*, *98*(13), 6999−7003.

28. Gray, P., & Scott, S. K. (1952). Autocatalytic reactions in the isothermal continuous stirred tank reactor oscillations and instabilities in the system $a + 2b \rightarrow 3b$, $b \rightarrow c$. *Chemical Engineering Science*, 39, 1087−1097.

29. Harvey, I., Di Paolo, E., Wood, R., Quinn, M., & Tuci, E. (2005). Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life*, *11*(1−2), 79−98.

30. Husbands, P., Smith, T., Jakobi, N., & O'Shea, M. (1998). Better living through chemistry: Evolving GasNets for robot control. *Connection Science*, *10*(4), 185–210.

31. Madigan, M., Martinko, J., Dunlap, P., & Clark, D. (2008). *Brock biology of microorganisms* (12th ed.). Upper Saddle River, NJ: Pearson Higher Education.

32. Maini, P., Baker, R., & Chuong, C.-M. (2006). Developmental biology: The Turing model comes of molecular age. *Science*, *314*(5804), 1397.

33. Maturana, H. R., & Varela, F. (1980). *Autopoiesis and cognition: The realization of the living*. Boston: D. Reidel.

34. Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence and technology of self-organizing machines*. Cambridge, MA: MIT Press.

35. Panfilov, A., Keldermann, R., & Nash, M. (2007). Drift and breakup of spiral waves in reaction-diffusion mechanics systems. *Proceedings of the National Academy of Sciences of the USA*, *104*(19), 7922–7926.

36. Pearson, J. E. (1993). Complex patterns in a simple system. eprint arXiv:patt-sol/9304003.

37. Philippides, A. O., Husbands, P., & O'Shea, M. (2000). Four-dimensional neuronal signaling by nitric oxide: A computational analysis. *Journal of Neuroscience*, *20*(3), 1199–1207.

38. Philippides, A. O., Husbands, P., Smith, T., & O'Shea, M. (2005). Flexible couplings: Diffusing neuromodulators and adaptive robotics. *Artificial Life*, *11*(1–2), 139–160.

39. Philippides, A. O., Ott, S. R., Husbands, P., Lovick, T., & O'Shea, M. (2005). Modeling co-operative volume signaling in a plexus of nitric oxide synthase-expressing neurons. *Journal of Neuroscience*, *25*(28), 6520–6532.

40. Port, R., & van Gelder, T. (Eds.). (1995). *Mind as motion: Explorations in the dynamics of cognition*. Cambridge, MA: MIT Press.

41. Sick, S., Reinker, S., Timmer, J., & Schlake, T. (2006). WNT and DKK determine hair follicle spacing through a reaction-diffusion mechanism. *Science*, *314*(5804), 1447–1450.

42. Slocum, A. C., Downey, D. C., & Beer, R. D. (2000). Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. In J.-A. Meyer, A. Berthoz, D. Floreano, H. Roitblat, & S. Wilson (Eds.), *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior* (pp. 430–439). Cambridge, MA: MIT Press.

43. Smith, S., & Armstrong, J. (1993). Reaction-diffusion control of heart development: Evidence for activation and inhibition in precardiac mesoderm. *Developmental Biology*, *160*(2), 535–542.

44. Stewart, J. (1996). Cognition = life: Implications for higher-level cognition. *Behavioural Processes*, *35*, 311–326.

45. Trevai, C., Fukazawa, Y., Ota, J., Yuasa, H., Ai-ai, T., & Asama, H. (2003). Cooperative exploration of mobile robots using reaction-diffusion equation on a graph. In *Proceedings of the 2003 IEEE International Conference on Robotics and Automation* (pp. 2269–2274). Piscataway, NJ: IEEE Press.

46. Turing, A. M. (1952). The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London B*, *237*, 37–72.

47. Van Duijn, M., Keijzer, F., & Franken, D. (2006). Principles of minimal cognition: Casting cognition as sensorimotor coordination. *Adaptive Behavior*, *14*(2), 157–170.

48. Vanag, K. V., & Epstein, I. R. (2007). Localized patterns in reaction-diffusion systems. *Chaos*, *17*(037110).

49. Wheeler, M. (2005). *Reconstructing the cognitive world*. Cambridge, MA: MIT Press.

50. Yamada, H., Nakagaki, T., Baker, R. E., & Maini, P. K. (2007). Dispersion relation in oscillatory reaction-diffusion systems with self-consistent flow in true slime mold. *Journal of Mathematical Biology*, *54*(6), 745–760.

51. Yang, L., Dolnik, M., Zhabotinsky, A. M., & Epstein, I. R. (2002). Pattern formation arising from interactions between Turing and wave instabilities. *Journal of Chemical Physics*, *117*(15).