

A Model Chemical Memory in an Evolved Animat

Kyran Dale

Centre for Computational Neuroscience and Robotics (CCNR)

Informatics

University of Sussex Brighton BN1 9QH, UK

kgd20@sussex.ac.uk

Abstract

This paper describes work carried out to investigate whether a classic simulated reaction-diffusion (RD) system could be used to control a 'minimally cognitive' animat during the course of a simple memory test. This test required the animat to remember an arbitrary signal and adapt its behaviour as a result. A further requirement was that the effects of the first signal be reversed by a second signal, returning the animat to its default behaviour. In this way the two signals combined to form a behavioural-switch, regulated by a memory-trace preserved in the homogeneous chemical substrate. The reaction-diffusion system chosen was that first described by Gray and Scott (Gray-Scott) and the minimally cognitive behavior of a class introduced by Beer et. al, involving the fixation and avoidance of a falling circular object by a whiskered animat. The parameters of this RD-controller were evolved using an evolutionary, or genetic, algorithm (GA).

Introduction

The study of memory in simple cognitive models has focussed primarily on the artificial neuron (AN) as a main component. In this paradigm networks of these ANs, connected by excitatory or inhibitory signalling links, mediate behaviour. The AN building blocks are usually heterogeneous, having variables such as activation time-constants or signalling thresholds which are adjusted on an individual basis. This paper deals with a different type of controller, which we call a reaction-diffusion controller (RDC), consisting of a one-dimensional array of cellular automata (CA) implementing a classic chemical model of reaction-diffusion. The constituent cells of this CA are homogeneous, sharing global defining variables.

Within Evolutionary Robotics the prominent model dynamical system is the continuous time recurrent neural network (CTRNN) (8; 9; 4). Many examples testify to the rich dynamics of which CTRNNs are capable (7; 9), such as generating the patterns to regulate legged robot gaits, and controlling such simple cognitive tasks as navigation and shape-discrimination. Many classic reaction-diffusion (RD) systems also display rich dynamics, manifesting the full range of classic qualities such as Hopf bifurcation, stable and unstable limit-cycles, chaotic boundaries etc.. The main

motivation for the work described in this paper was to see whether the tried and tested technique of evolving neural-network controllers for simple robotic behavior could be adapted to harnessing some of the rich dynamics displayed by these RD systems. In this sense the interest was both methodological, to show that evolutionary algorithms could be used successfully with a different class of non-linear system, but also focused on exploring the ability of RD controllers. For example the ability to sustain spatio-temporal patterns suggests a role in controlling gaited movement but can systems be tuned to particular requirements? Given the difference between the essentially 'spaceless' CTRNNs and the necessarily spatial RD systems there is also the intriguing possibility that they might be able to complement one another. By placing artificial neurons in an excitable medium with which they can interact, CTRNNs might be able to exploit the spatio-temporal properties of the medium. Given the dynamical potential of these RD systems there has been very little work dedicated to exploring it (1; 2).

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 Fig. 1. Output from whisker-like proximity sensors was fed to the cells in the RD-ring via weighted links, perturbing the concentration 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 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.

Reaction-diffusion Models

Perhaps the best known example of a reaction-diffusion model is that proposed by Alan Turing (6) 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, patterning of nautilus shells and many other natural patterns.

Within the class of model reaction-diffusion 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 paper (10) and extensively analyzed by Pearson in his 1993 paper (5). A variant of the autocatalytic Selkov model of glycolysis (5) the Gray-Scott model corresponds to the following reactions:



Both reactions are reversible so p is an inert product. A feed term for u introduces a non-equilibrium constraint with the feed process removing both u and v . This results in the following reaction-diffusion equations, expressed in dimensionless units:

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

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

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 Method section below for specific details). A trivial steady state of $u = 1, v = 0$ exists for all values of F and k . Gray-Scott proves a very robust simulation, showing no qualitative difference when implemented by forward Euler integration over a broad range of spatial and temporal scales (5).

When suitably perturbed Gray-Scott exhibits a large variety of spatio-temporal patterns that have to be seen to be appreciated. Pearson's paper is replete with beautiful images but the simulation is best appreciated in real-time with a two-dimensional simulation and a suitable colour-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 two-dimensional phase-diagram described shows regions associated with specific spatio-temporal patterns, ranging from spot replication and stripes in a continuous transition to traveling waves and spatio-temporal chaos.

Visually-Guided Agents

The choice of an evolved animat model, for example to demonstrate the potential of a novel reaction-diffusion 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. ((8) p.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 $\mathbf{s}+$ a whiskered animat, capable of moving along the floor of a two-dimensional arena (in the xz plane), is required to orientate toward and track a circular object falling from the arena's ceiling with a large range of vertical and horizontal speeds. In the absence of $\mathbf{s}+$ or if $\mathbf{s}+$ is followed by the second signal $\mathbf{s}-$ the animat is required to avoid the falling object. (see Figs. 2 and 3 for details).

Beer evolved continuous time recurrent neural networks (CTRNNs) to control his animats, a control-system the author has some experience of (4). His subsequent analysis (9) 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 a reaction-diffusion medium.

Evolving Controllers

The Gray-Scott model, in keeping with most reaction-diffusion systems, is highly non-linear, at least unintuitive and often counter-intuitive¹. 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 which escapes us. In cases such as this, where we require a controller capable of exploiting even a relatively simple dynamical system, it would seem that the need is pressing to leverage the increasing computer power at our disposal and

¹The speed of modern processors makes it possible to interact in real-time with 2D implementations of these reaction-diffusion systems. Having implemented and played with just such a model of, among others, Gray-Scott, we can attest to its counter-intuitiveness.

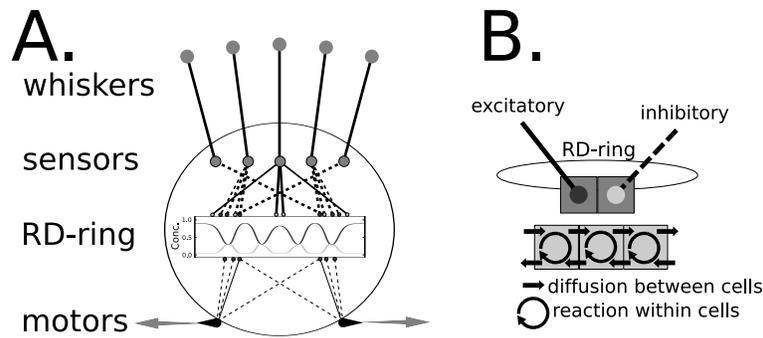


Figure 1: (A) The animat model. Output from the proximity sensors is fed, via weighted links, to the reaction-diffusion ring (*RD-Ring*) where it perturbs the cellular concentration of chemicals u and v . Solid links increase the chemical concentration in the cell while dashed links decrease it. The 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 reaction-diffusion 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 while inhibitory links (dashed) decrease it. In this way the whisker sensors affect the chemistry of the *RD-Ring* which in turn affects the motors.

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 reaction-diffusion 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, poor performing controllers, to high, where the controllers are (much) better performing. This image leaves out important details, particularly the concept of neutral-networks², but the key detail is captured. From random parameters and allowing for a suitable encoding scheme, it should be possible to automatically produce good controllers by applying evolutionary pressure. The work described in this paper and elsewhere (8; 9) is testament to that fact.

Method

To a large extent details from Beer's earlier simulations (8) were preserved and the required behaviors essentially the same. The arena was 400 units long by 275 units high (Fig. 2) 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 base.

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

Fig. 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 reaction-diffusion ring (*RD-ring*) consisting of 128 cells subject to intra-cellular reaction and inter-cellular diffusion between near-neighbours (see the chemical reactions 1, 2 and 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 the rate equations 3 and 4). Input via links to the cells perturbed the specified chemical's concentration by a simple multiple of time-step (0.1), sensor activation $\in [0, 1]$ and link weight $\in [-1, 1]$. The cellular concentration of u and v was bounded within the range $\in [0, 1]$.

The animat's motors received input from cells in the *RD-ring*. Input from a individual link was a product of 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.

Diffusion rates d_u and d_v were fixed at the standard values (5) of 2×10^{-5} and 10^{-5} respectively and the length 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 reaction-diffusion system (see subsection *Reaction-diffusion Models* above).

The GA consisted of a population of thirty 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 or desirable to have the genotype describe a fully connected controller (1408 links in all) so the number of links was pre-set. The starting number for the orientation experiment was 8 sensor→RD-ring, 4 RD-ring→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 operations. The numbers on the genotype were in the range $\in [0, 1]$, being mapped onto their respective controller parameters. Mutation consisted of the addition of a normally distributed random value with average 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 four trials (see Fig. 3) where the two values d_s and d_e specify the absolute horizontal distance between animat and shape at the trial's start and end, marked by the shape reaching the arena floor, respectively:

$$f(d_s, d_e) = \begin{cases} 1 - \frac{d_s}{d_e} & \text{if } d_e < d_s \\ \max(\frac{d_s - d_e}{50}, -1) & \text{if } d_e \geq d_s \end{cases}$$

The value of this fitness function is highest if the animat fixates the object centrally and lowest if the animat avoids the object, to a maximum at distance 50. Where the trial required the animat to avoid the falling object (T1, T3, T4 in Fig. 3) 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 towards the circle and the total fitness calculated thus:

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

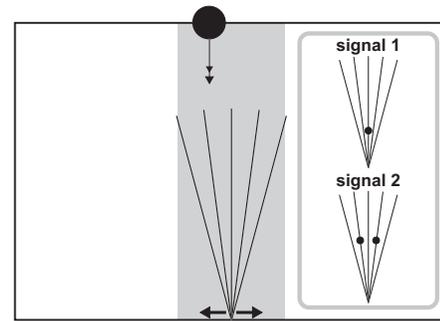


Figure 2: The memory experiment (to scale). An animat with five whiskers spread over a 30° span is placed at the centre of the arena's floor. During a trial a circle was placed at the top of the arena within the grey drop zone on a straight downward trajectory of between 3 and 4 units per second. Prior 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.

This function was designed to encourage a switching of behavior over T2 while avoiding evolutionary local minima³ 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.

Training Protocol

A number of trials were conducted to assess the ability of individual controllers and allot their respective genotypes a fitness score. Fig. 2.a shows the trial set-up for the orientation experiment. The grey 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

Memory Previous work (3) has shown that motor feedback to the RD-ring can be used by the animat to stabilize behaviour. Purely diffuse controllers were able to use motor feedback in this way to maintain a memory trace. To prevent this the animats were not allowed to form links from motors to 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

³A simple explanation...

⁴In keeping with Beer's original model (8) the simulation was noiseless, meaning that the symmetrical animat controller was incapable of breaking symmetry without stimulus from the whiskers.

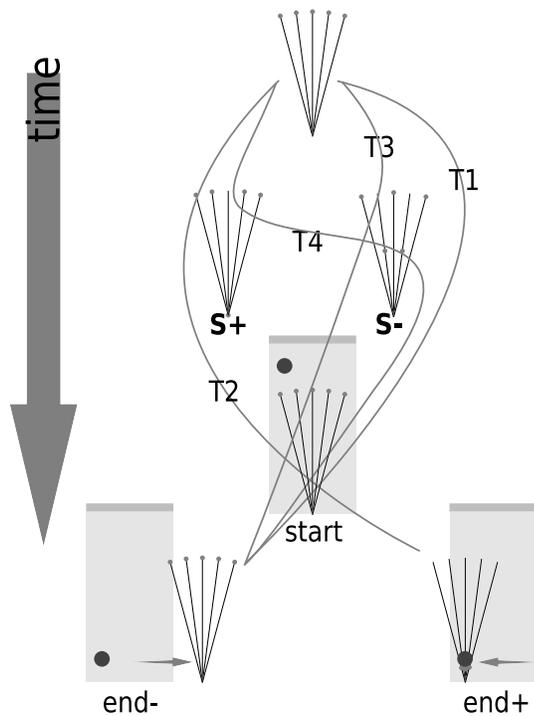


Figure 3: The four trials (**T1-4**) providing the components of the memory task's fitness function. **end+** and **end-** show the desired end positions of the animat relative to the falling object. The duration of all phases is the range $\in [400, 600]$ and the stimuli last for 10 time units. (**T1**) 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 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.

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 .

Fig. 3 shows the four trials T1-4 that comprised 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 the circle.

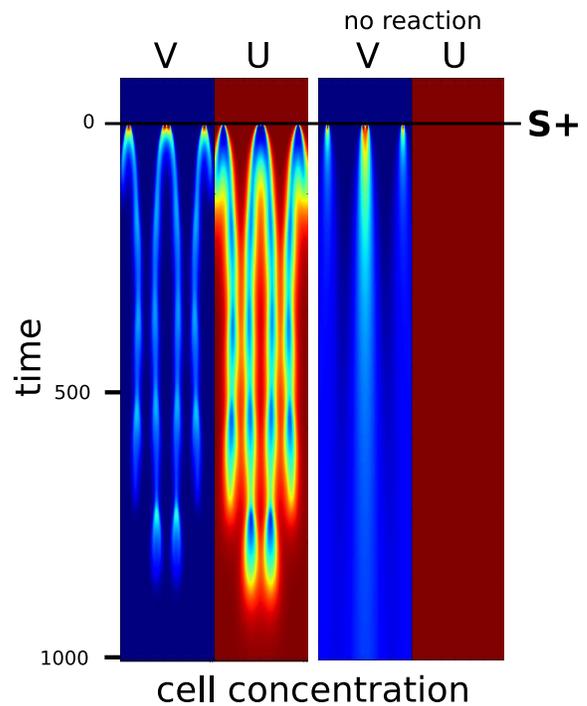


Figure 5: The colour-mapped (red high, blue low) changes in concentration over time of chemicals u and v in response to the signal **s+**. In the trace on the right the reaction component of the RD system has been disabled. When allowed to react the two chemicals maintain a strong, autocatalytic memory trace for approximately 750 time units, long enough to reliably remember the signal and score highly on the task. In the absence of reaction between u and v initially high concentrations of v diffuse away while u , unable to autocatalyse, 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 paralysed.

Results

In this section we focus on the best performing animat of a single evolved population. Populations with near-optimal solutions were readily evolved and the choice of this one is arbitrary but informed by pedagogic considerations. Similar mechanisms to those described were found in the large majority of those animats analyzed.

An important consideration in this kind of evolutionary modelling is that one be hopeful that the models produced will submit to analysis. As mentioned above, evolving a fully connected RD-controller would be infeasible, given the processing power available, but another objection might be that by encouraging evolution to distribute the animat's cognition within an overly complex structure one ends up with a model which is too complex to understand. One of the benefits of allowing network connectivity to be an evolutionary variable is that we can encourage evolution to search

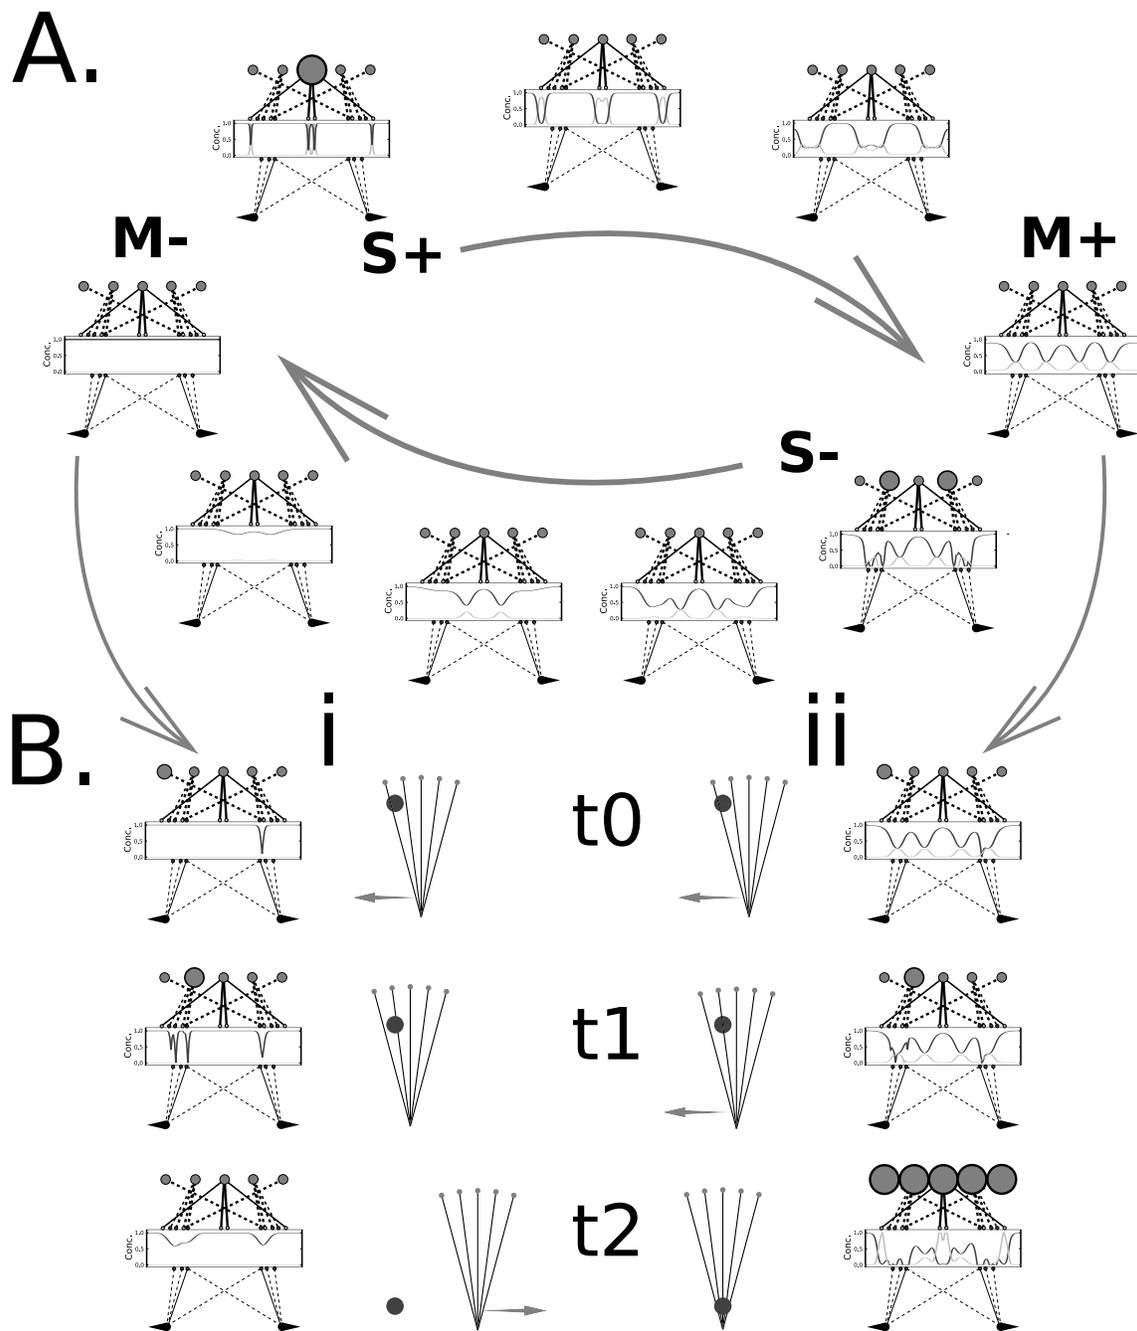


Figure 4: (A.) A chemical memory-switch. Moving clockwise from the default settled state **M-**, the animat receives the stimulus **S+** causing an auto-catalytic cycle which leads to the semi-stable state **M+**. While in state **M+** the resetting signal **S-** disrupts the RD-ring's structure, causing it to return to **M-**. (B.) shows 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.

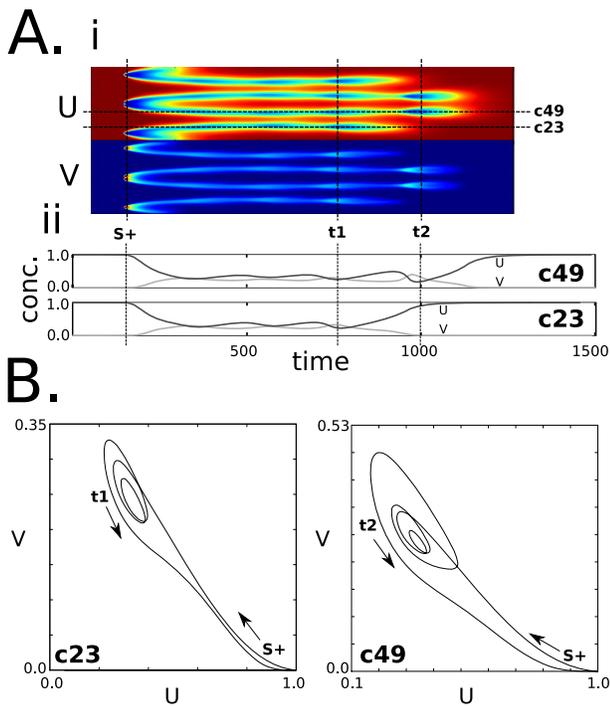


Figure 6: The change in U/V concentration over time of the two RD-ring cells **c23** and **c49** in response to the animat receiving stimulus $s+$. Two other significant times are marked, **t1** and **t2**, the return of **c23** and **c49** to default. (Ai) shows the colourmapped (red high, blue low) concentrations of chemicals U and V, with the position of cells **c23** and **c49** indicated, for U, by dashed horizontal lines. (Aii) the change in concentration over time for chemicals U (dark grey) and V (light grey) 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.

for simpler solutions. With this in mind, the animat population of these results, having achieved a close to optimal fitness, was further evolved with a single change made to the GA. The probability of adding a link to the controller during mutation was set to zero, while the probability of deleting a link remained the same. In this way evolution is 'locked' from exploring bigger networks while being able to randomly wander through or test smaller ones. This technique has been found to rapidly reduce the size of networks while maintaining their fitness and has the added advantage that one does not have to introduce an arbitrary component into the fitness function to encourage simplicity.

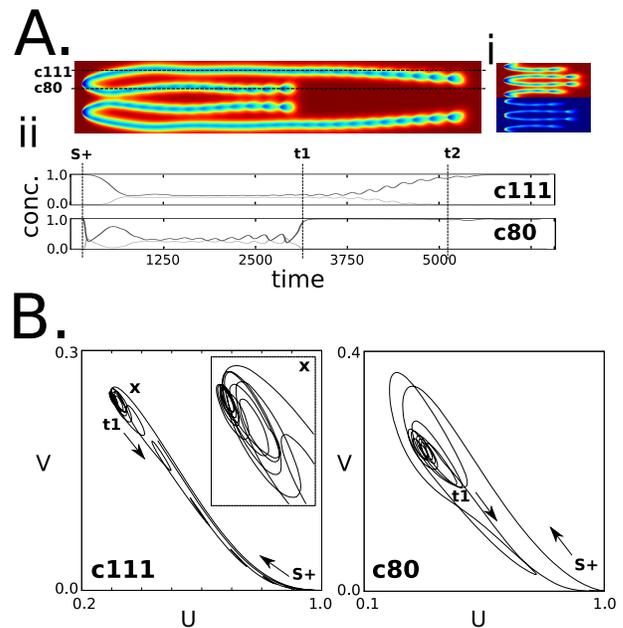


Figure 7: The reaction of two selected cells, **c80** and **c111**, to stimulus $s+$ after a further period of evolution, wherein the animat was required to increase the length of its memory. The top plot of (Ai) shows, to scale, the previous memory trace following $s+$. (B) the attractor cycles are now more densely packed and greater in number.

A Chemical Switch

The evolved model, with 22 symmetrically arranged links, is summarized in Fig. 4. Fig. 4A shows the switching cycle which 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 application of $s-$, a half maximum stimulus of the first proximal whiskers, at **M+** disrupts these waves, bringing the chemical system back to **M-**.

Fig. 4B shows the response of the animat to the same falling object trajectory whilst 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, imbalance the left and right motors, drawing the animat towards the falling object. This behaviour at $t0$ is mirrored in **M+** (Bii). At time $t1$ the behaviours start to diverge, in response to stimulation of the whisker second from left. In Bi 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 towards its centre 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 Fig. 5. 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.

The Stability of the Chemical Memory

As shown in Figs. 5 and 6 the memory trace produced by signal $s+$ is unstable. The memory is long enough to enable the animat to pass the test requirements, 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 Fig. 6 we focus on the change in concentration of two cells over the course of the memory trace. These cells are roughly at the centre of two of the four symmetrical peaks and troughs of v and u respectively (see Fig.4 $M+$). Fig. 6B shows that at these points in the chemical ring the interaction of u and v 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.

Extending the Memory's Duration

In order to see whether the duration of the animat's memory-trace was coincidental to the requirements of the task, the animat population was further evolved (see subsection *Evolutionary Controllers* 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. Fig. 7 shows the successful result of an evolutionary run which required the animat to maintain a memory five times longer than that of the original task, whose trace is shown to scale in Ai. At this point the simulations became impracticably long but there was no indication in this or other of a hard constraint on the possible duration of memory. Note that the interaction of u and v now 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.

Discussion

Although Cajal's neuron doctrine is predominant in cognitive studies and, by definition, neuroscience computational and otherwise, it does beg 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 behaviour and exhibit the classical learning models. We would not suggest the models described in this paper hold any answers to the larger questions of animal cognition but the ability of these simple chemical systems to mediate simple cognitive tasks required memory is intriguing. Extending these model systems, used extensively and successfully in Biology to explain such phenomena as cardiac rhythmia, animal patternation, morphogenetical development etc. to the cognitive realm could prove fruitful.

Acknowledgments

Particular thanks must go to Phil Husbands for green-lighting the direction taken in this work, the paper's referees and to the EPSRC (grant GR/T11043/01).

References

- [1] Breyer, J., Ackerman, J., McCaskill, J.: Evolving Reaction-diffusion Ecosystems with Self-assembling Structures in Thin Films. *Alife* 4: 25–40. (1998)
- [2] Adamatzky, A., De Lacy Costello, B., Asai, T.: *Reaction-diffusion Computers*. Elsevier, (2006).
- [3] Dale, K.: Evolving Reaction-diffusion Controllers for Minimally Cognitive Animats. *Proceedings of Simulation of Adaptive Behaviour 2006* 498–509.
- [4] Dale, K., Collett, T.S.: Using Artificial Evolution and Selection to Model Insect Navigation. *Current Biology* 11: 51–62. (2001)
- [5] Pearson, John E.: Complex Patterns in a Simple System. eprint arXiv:patt-sol/9304003 (04/1993)
- [6] Turing, A.M.: The Chemical Basis of Morphogenesis. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, Volume 237, Issue 641, pp. 37–72 (1952)
- [7] Beer, R.D.: On the Dynamics of Small Continuous-time Recurrent Neural-networks. *Adaptive Behaviour* 3:469–510
- [8] Beer, R.D.: Toward the Evolution of Dynamical Neural Networks for Minimally Cognitive Behavior In P. Maes, M. Mataric, J. Meyer, J. Pollack and S. Wilson (Eds.), *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (pp. 421–429). MIT Press.
- [9] Beer, R.D.: The Dynamics of Active Categorical Perception in an Evolved Model Agent. *Adaptive Behavior* 11(4):209–243
- [10] Gray, P., Scott, S.K.: Autocatalytic Reactions in the Isothermal Continuous Stirred Tank Reactor Oscillations and Instabilities in the System $A + 2B \rightarrow 3B$, $B \rightarrow C$. *Chem. Eng. Sci.*, **39**, 1087–1097 (1984)