

Towards Machine Learning Control of Chemical Computers

Adam Budd,¹ Christopher Stone,¹ Jonathan Masere,² Andrew Adamatzky,¹

Ben De Lacy Costello² and Larry Bull¹

Learning Classifier Systems Group Technical Report – UWELCSG06-004

¹Faculty of Computing, Engineering & Mathematics, University of the West of England

²Faculty of Applied Sciences, University of the West of England

Coldharbour Lane, Frenchay

Bristol BS16 1QY, U.K.

Abstract

The behaviour of pulses of Belousov-Zhabotinski (BZ) reaction-diffusion waves can be controlled automatically through machine learning. By extension, a form of chemical network computing, i.e., a massively parallel non-linear computer, can be realised by such an approach. In this initial study a light-sensitive sub-excitable BZ reaction in which a checkerboard image comprising of varying light intensity cells is projected onto the surface of a thin silica gel impregnated with tris(bipyridyl) ruthenium (II) catalyst and indicator is used to make the network. As a catalyst BZ solution is swept past the gel, pulses of wave fragments are injected into the checkerboard grid resulting in rich spatio-temporal behaviour. This behaviour is shown experimentally to be repeatable under the same light projections. A machine learning approach, a learning classifier system, is then shown able to direct the fragments to an arbitrary position through dynamic control of the light intensity within each cell in both simulated and real chemical systems.

1 Introduction

There is growing interest in research into the development of hybrid wetware-silicon devices focused on exploiting their potential for ‘non-linear computing’. The aim is to harness the as yet only partially understood intricate dynamics of non-linear media to perform complex ‘computations’ (potentially) more effectively than with traditional architectures and to further the understanding of how such systems function. The area provides the prospect of radically new forms of machines and is enabled by improving capabilities in wetware-silicon interfacing. We are developing an approach by which networks of non-linear media – reaction-diffusion systems - can be produced to achieve a user-defined computation in a way that allows control of the media used. Machine Learning algorithms are used to design the appropriate network structures by searching a defined behavioural space to create a computing resource capable of satisfying a given objective(s). In this paper we examine the underlying dynamics of the chosen Belousov-Zhabotinsky (BZ) [Zhakin & Zhabotinsky, 1970] reaction-diffusion system in which the networks are created via light and present initial results from the general control/programming methodology.

Excitable and oscillating chemical systems have been used to solve a number of computational tasks such as implementing logical circuits [Steinbock et al., 1996; Toth et al. 1994], image processing [Kuhnert et al., 1989], shortest path problems [Steinbock et al., 1995] and memory [Motoike et al., 2001]. In addition chemical diodes [Agladze et al., 1996], coincidence detectors [Gorecki et al., 2003] and transformers where a periodic input signal of waves may be modulated by the barrier into a complex output signal depending on the gap width and frequency of the input [Sielewiesiuk & Gorecki, 2002] have all been demonstrated experimentally. However, to some degree the lack of compartmentalisation in these simple chemical systems limits the domain of solvable tasks thus making it difficult to realise general-purpose computing. This proposed methodology of utilising networks of coupled oscillating chemical reactions may provide a solution. The fact that these coupled oscillators can be controlled via the application of external fields such as light provides a possible method for undertaking a number of complex computations provided an effective methodology for realising large scale networks can be found.

A number of experimental and theoretical constructs utilising networks of chemical reactions to implement computation have been described. These chemical systems act as simple models for networks of coupled oscillators such as neurons, circadian pacemakers and other biological systems [Kawato & Suzuki, 1980]. Over 30 years ago the construction of logic gates in a bistable chemical system was described by Rossler [Rossler 1974]. Ross and co-workers [Hjelmfelt et al., 1991; 1992] produced a theoretical construct suggesting the use of “chemical” reactor systems coupled by mass flow for implementing logic gates neural networks and finite-state machines. In further work Hjelmfelt et al. [Hjelmfelt & Ross, 1993; Hjelmfelt et al., 1993] simulated a pattern recognition device constructed from large networks of mass-coupled chemical reactors containing a bistable iodate-arsenous acid reaction. They encoded arbitrary patterns of low and high iodide concentrations in the network of 36 coupled reactors. When the network is initialized with a pattern similar to the encoded one then errors in the initial pattern are corrected bringing about the regeneration of the stored pattern. However, if the pattern is not similar then the network evolves to a homogenous state signalling non-recognition.

In related experimental work Laplante et al. [1995] used a network of eight bistable mass coupled chemical reactors (via 16 tubes) to implement pattern recognition operations. They demonstrated experimentally that stored patterns of high and low iodine concentrations could be recalled (stable output state) if similar patterns were used as input data to the programmed network. This highlights how a programmable parallel processor could be constructed from coupled chemical reactors. This described chemical system has many properties similar to parallel neural networks. In other work Lebender and Schneider [1994] described methods of constructing logical gates using a series of flow rate coupled continuous flow stirred tank reactors (CSTR) containing a bistable nonlinear chemical reaction. The minimal bromate reaction involves the oxidation of cerium(III) (Ce^{3+}) ions by bromate in the presence of bromide and sulphuric acid. In the reaction the Ce^{4+} concentration state is considered as “0” “false” (“1” “true”) if a given steady state is within 10% of the minimal (maximal) value. The reactors were flow rate coupled according to rules given by a feedforward neural network run using a PC. The experiment is started by feeding in two “true” states to the input reactors and then switching the flow rates to generate “true”-“false”, “false”-“true” and “false”-“false”. In this three coupled reactor system the AND (output “true” if inputs are both high Ce^{4+} , “true”),

OR (output “true” if one of the inputs is “true”), NAND (output “true” if one of the inputs is “false”) and NOR gates (output “true” if both of the inputs are “false”) could be realised. However to construct XOR and XNOR gates two additional reactors (a hidden layer) were required. These composite gates are solved by interlinking AND and OR gates and their negations. In their work coupling was implemented by computer but they suggested that true chemical computing of some Boolean functions may be achieved by using the outflows of reactors as the inflows to other reactors i.e. serial mass coupling.

As yet no large scale experimental network implementations have been undertaken mainly due to the complexity of analysing and controlling so many reactors. That said there have been many experimental studies carried out involving coupled oscillating and bistable systems [Stuchl & Marek, 1982; Crowley & Field, 1986; Bar-Eli & Reuveni, 1985; Bar-Eli, 1985; Crowley & Epstein, 1989; Holz & Schneider, 1993]. The reactions are coupled together either physically by diffusion or an electrical connection or chemically, by having two oscillators that share a common chemical species. The effects observed include multistability, synchronisation, in-phase and out of phase entrainment, amplitude or “oscillator death”, the cessation of oscillation in two coupled oscillating systems, or the converse, “rythmogenesis”, in which coupling two systems at steady state causes them to start oscillating [Dolnik & Epstein, 1996].

In this paper we adapt a system described by Wang et al. [1999] and explore the computational potential based on the movement and control of wave fragments. In the system they describe the application of Gaussian noise (where the mean light level is fixed at the subexcitable threshold of the reaction) in the form of light projected onto a thin layer of the light sensitive analogue of the BZ reaction was observed to induce wave formation and subsequently “avalanche behaviour” whereby a proliferation of open ended excitation wave fragments were formed. Interestingly calcium waves induced in networks of cultured glial cells [Jung et al. 1998] display similar features to the ones identified in this chemical system which the authors postulated may provide a possible mechanism for long-range signalling and memory in neuronal tissues.

Machine Learning techniques, such as Evolutionary Algorithms (EAs)[e.g., Michalewicz & Fogel, 1999] and Reinforcement Learning (RL) [Sutton & Barto, 1998], are being increasingly used in the design of

complex systems. Example applications include data mining, time series analysis, scheduling, process control, robotics and electronic circuit design. Such techniques can be used for the design of computational resources in a way that offers substantial promise for application in non-linear media computing since the algorithms are almost independent of the medium in which the computation occurs. This is important in order to achieve effective non-linear media computing since they do not need to directly manipulate the material to facilitate learning and the task itself can be defined in a fairly unsupervised manner. In contrast, most traditional learning algorithms use techniques that require detailed knowledge of and control over the computing substrate involved. In this paper we control the BZ network via a reinforcement learning approach which uses evolutionary computing to create generalizations over the state-action space – Holland’s Learning Classifier System [Holland, 1986], in particular a form known as XCS [Wilson, 1995].

The paper is arranged as follows: The next section describes the sub-excitable BZ system which forms the basis of this chemical computing research and results from investigations into its basic properties. The next section describes a computational model of the system and the chosen machine learning approach. Initial results from using the machine learner to control the simulated and real chemical system are then presented.

2 Experimental System

2.1 Materials and Equipment

Sodium bromate, sodium bromide, malonic acid, sulphuric acid, tris(bipyridyl) ruthenium (II) chloride, 27 % sodium silicate solution stabilized in 4.9 M sodium hydroxide were purchased from Aldrich and used as received unless stated otherwise.

An InFocus Model Projector was used to illuminate the computer-controlled image. Images were captured using a Panasonic NV-GS11 digital video camera. The microscope slide was immersed in the continuously fed reaction solution contained in a custom-designed Petri dish, designed by Radleys, with a water jacket

thermostatted at 22 °C. A Watson Marlow 205U multi-channel peristaltic pump was used to pump the reaction solution into the reactor and remove the effluent.

2.2 Experimental Procedures

2.2.1 Making Gels

A stock solution of the sodium silicate solution was prepared by mixing 222 mL of the purchased sodium silicate solution with 57 mL of 2 M sulphuric acid and 187 mL of deionised water, similar to the procedure used by Wang et al. [1999]. Pre-cured solutions for making gels were prepared by mixing 5 mL of the acidified silicate solution with a solution consisting of 1.3 mL of 1.0 M sulphuric acid and 1.2 mL of 0.025 M tris(bipyridyl) ruthenium (II) chloride. Using capillary action, portions of this solution were transferred onto microscope slides with 100 µm shims and Plexiglas covers. The transferred solutions were left for 3 hours to permit complete gellation after which the covers and shims were removed and the gels washed in deionised water to remove residual tris(bipyridyl) ruthenium (II) chloride and the sodium chloride byproduct. The gels were 26 mm by 26 mm, with a wet thickness of approximately 100 µm. The gels were stored under water and rinsed right before use.

2.2.2 Catalyst-Free Reaction Mixture

The bulk of the catalyst-free reaction mixture was freshly prepared in 300 mL batches, which involved the *in situ* synthesis of stoichiometric bromomalonic acid from malonic acid and bromine generated from the partial reduction of sodium bromate. The catalyst-free reaction solution consisted of the 0.36 M sodium bromate, 0.0825 M malonic acid, 0.18 M sulphuric acid and 0.165 M bromomalonic acid. To minimize the deterioration during the experiment, this solution was kept in an ice bath. This solution was continuously fed into the thermostatted reactor, with a reactor residence time of 30 minutes.

2.2.3 Experimental setup

The spatially distributed excitable field on the surface of the gel was made possible by the projection of a 10-by-10 cell checkerboard grid pattern generated using a computer. After Wang et al. [1999], the checkerboard image comprised the heterogeneous network of Gaussian distributed light levels, with a mean at the sub-excitable threshold, from a low of 0.394 mW cm^{-2} to a high of 9.97 mW cm^{-2} intensity cells in 256 bins, representing excitable and non-excitable domains respectively. A digital video camera was used to capture the chemical wave fragments. A diagrammatic representation of the experimental setup is shown in Figure 1.




Figure 1: A block diagram of the experimental setup where the computer, projector, mirror, microscope slide with the catalyst-laden gel, thermostatted Petri dish, peristaltic pump, thermostatted water bath, reservoir of catalyst-free reaction solution, digital camcorder and, the effluent flow are designated by A, B, C, D, E, F, G, H, I and J, respectively. The catalyst-free reaction solution reservoir was kept in an ice bath during the experiment.

2.2.4 Data Capturing and Image Processing

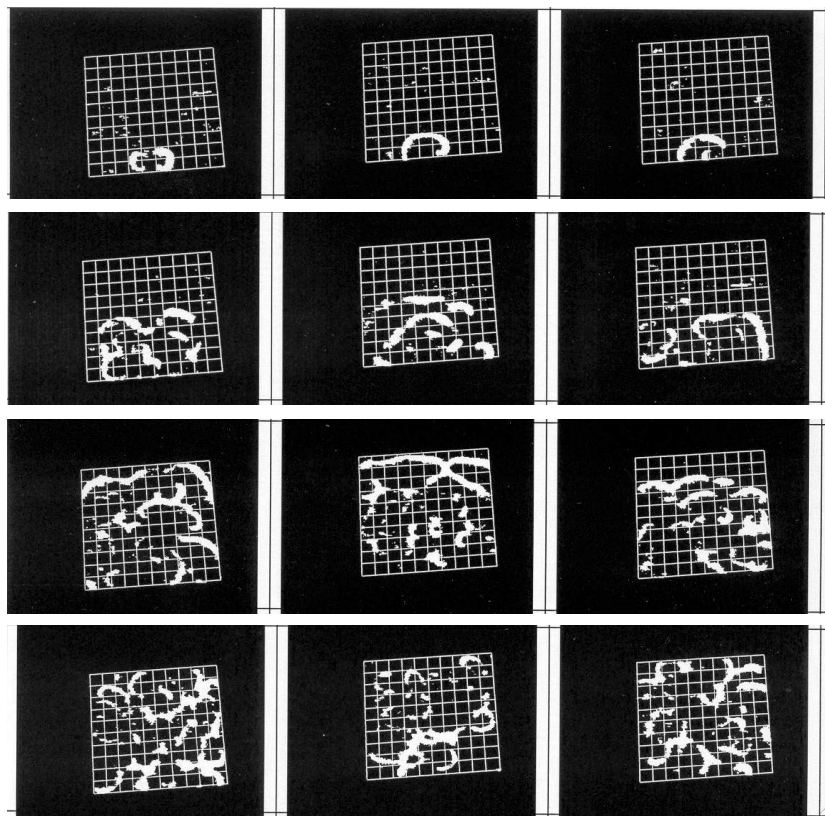
A checkerboard grid pattern was projected onto the catalyst-laden gel through a 455nm narrow bandpass interference filter and 100/100mm focal length lens pair and mirror assembly. The size of the projected grid was approximately 20mm square. Every 10 seconds, the checkerboard pattern was replaced with a uniform grey level of 9.97 mW cm^{-2} for 400 ms during which time an image of the BZ waves on the gel was captured. The purpose of removing the grid pattern during this period was to allow activity on the gel to be more visible to the camera and assist in subsequent image processing of chemical activity.

Captured images were processed to identify chemical wave activity. This was done by differencing successive images on a pixel by pixel basis to create a black and white thresholded image. Each pixel in the black and white image was set to white, corresponding to chemical activity; if the intensity of the red or blue channels differed in successive images by more than 1.95%. Pixels at locations not meeting this criterion were set to black. The thresholded images were automatically despeckled and manually edited to remove artefacts of the experiment, such as glare from the projector and bubbles from the oxidative decarboxylation of malonic acid and bromomalonic acid. The images were cropped to the grid location and the grid superimposed on the thresholded images to aid analysis of the results.

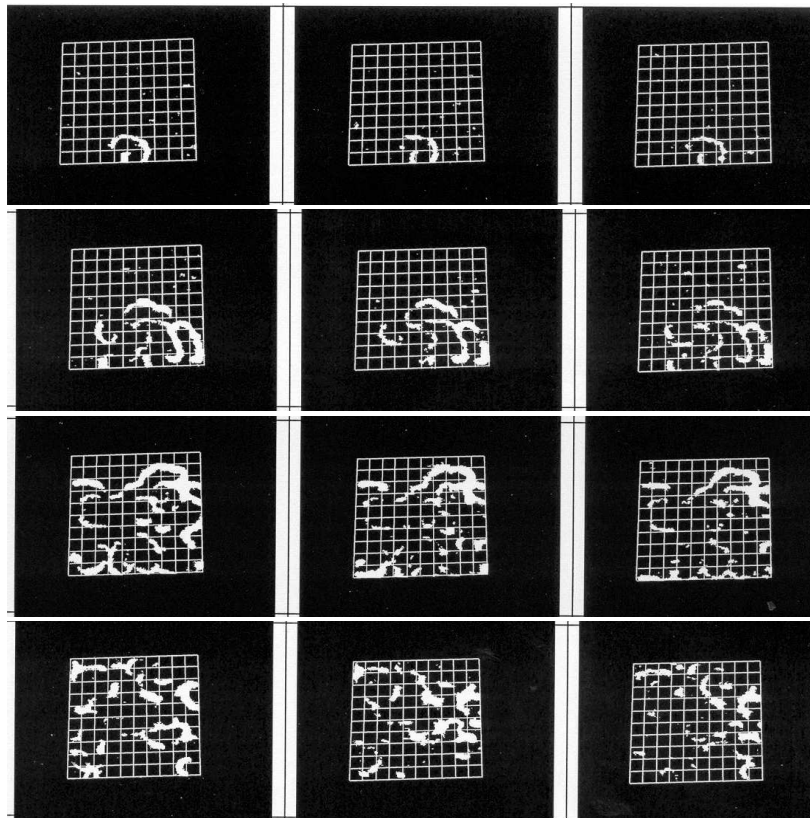
2.3 Experimental System Behaviour

Computationally, the described system has 100 inputs and potentially as many outputs. In this study we have examined the use of single pulses of excitation into the bottom of the grid. Each pulse expands and grows as a fragmented front while traversing the network of excitable domains of varying light intensity. Multiple disproportionate collisions result in daughter fragments, mutual fragment extinctions, etc. (Figure 2). That is, oscillatory/clustering behaviour can be observed – wave fragments and their collisions equate to information transfer [Adamatzky, 2001] – the fragment dynamics are mediated by local light levels, collisions, boundaries, trajectories of previous fragments, diffusion processes in adjacent cells and beyond, etc.

The spatio-temporal evolution of the system is affected by the series of light levels projected into each cell. To be useful as a computational formalism, the behaviour of the system under a given light “program”, i.e., a series of levels in each given cell, must be sufficiently repeatable. That is, it must be possible to exert control at some appropriate level over the evolution and dynamics of the excitable wave fragments. Figure 2 shows example results from three random light programs and three runs using the same light programs where we aimed to control the evolution of the fragmentation pattern from a single input, thereby enabling us to assess the repeatability/stability of the chemical system. As can be seen, in the former case (Figure 2a) there is no obvious correspondence in the spatio-temporal behaviour whereas the behaviour under the same program is remarkably similar (Figure 2b).



(a)



(b)

Figure 2: Example spatio-temporal evolution of the chemical system under different light programs (a) and the same program (b).

3 Computational System

3.1 Model

The features of this system were simulated using a two-variable Oregonator model modified to account for the photochemistry [Field & Noyes, 1973; Krug, 1990]:

$$\frac{\partial u}{\partial t} = \frac{1}{\varepsilon} \left(u - u^2 - (fv + \Phi) \frac{u - q}{u + q} \right) + D_u \nabla^2 u$$
$$\frac{\partial v}{\partial t} = u - v .$$

The variables u and v represent the instantaneous local concentrations of the bromous acid autocatalyst and the oxidized form of the catalyst, HBrO_2 and tris (bipyridyl) Ru (III), respectively, scaled to dimensionless quantities. The ratio of the time scales of the two variables, u and v , is represented by ε , which depends on the rate constants and reagent concentration; f is a stoichiometric coefficient. The rate of the photo-induced bromide production is designated by Φ , which also denotes the excitability of the system in which low light intensities facilitate excitation while high intensities result in the production of bromide that inhibits the process, experimentally verified by Kádár et al. [1997]. The scaling parameter, q , depends on reaction rates only. The system was integrated using the Euler method with a five-node Laplacian operator, time step $\Delta t=0.001$ and grid point spacing $\Delta x=0.15$. The diffusion coefficient, D_u , of species u was unity, while that of species v was set to zero as the catalyst was immobilized in gel.

3.2 Machine Learning Algorithm

XCS is a relatively recent development of Holland's Learning Classifier System formalism and has been shown able to tackle many complex tasks effectively (see [Bull, 2004] for examples). It consists of a

limited size population [P] of classifiers (rules). Each classifier is in the form of “IF condition THEN action” (*condition*→*action*) and has a number of associated parameters. Conditions traditionally consist of a trinary representation, {0,1,#}, where the wildcard symbol facilitates generalization, and actions are binary strings.

On each time step a match set [M] is created. A system prediction is then formed for each action in [M] according to a fitness-weighted average of the predictions of rules in each action set [A]. The system action is then selected either deterministically or stochastically based on the fitness-weighted predictions (usually 0.5 probability per trial). If [M] is empty a covering heuristic is used which creates a random condition to match the given input and then assigns it to a rule for each possible action.

Fitness reinforcement in XCS consists of updating three parameters, ϵ , p and F for each appropriate rule; the fitness is updated according to the relative accuracy of the rule within the set in five steps:

- i) Each rule’s error is updated: $\epsilon_j = \epsilon_j + \beta(|P - p_j| - \epsilon_j)$ where $0 \leq \beta \leq 1$ is a learning rate constant.
- ii) Rule predictions are then updated: $p_j = p_j + \beta(P - p_j)$
- iii) Each rule’s accuracy κ_j is determined: $\kappa_j = \alpha(\epsilon_0/\epsilon)^v$ or $\kappa=1$ where $\epsilon < \epsilon_0$
 v , α and ϵ_0 are constants controlling the shape of the accuracy function.
- iv) A relative accuracy κ_j' is determined for each rule by dividing its accuracy by the total of the accuracies in the action set.
- v) The relative accuracy is then used to adjust the classifier’s fitness F_j using the moyenne adaptive modifee (MAM) procedure: If the fitness has been adjusted $1/\beta$ times, $F_j = F_j + \beta(\kappa_j' - F_j)$. Otherwise F_j is set to the average of the values of κ' seen so far.

In short, in XCS fitness is an inverse function of the error in reward prediction, with errors below ϵ_0 not reducing fitness. The maximum $P(a_i)$ of the system’s prediction array is discounted by a factor γ and used

to update rules from the previous time step and an external reward may be received from the environment. Thus XCS exploits a form of Q-learning [Watkins, 1989] in its reinforcement procedure.

A Genetic Algorithm (GA) [Holland, 1975] acts in action sets [A], i.e., niches. Two rules are selected based on fitness from within the chosen [A]. Two-point crossover is applied at rate χ and point mutations at rate μ . Rule replacement is global and based on the estimated size of each action set a rule participates in with the aim of balancing resources across niches. The GA is triggered within a given action set based on the average time since the members of the niche last participated in a GA (after [Booker, 1989]).

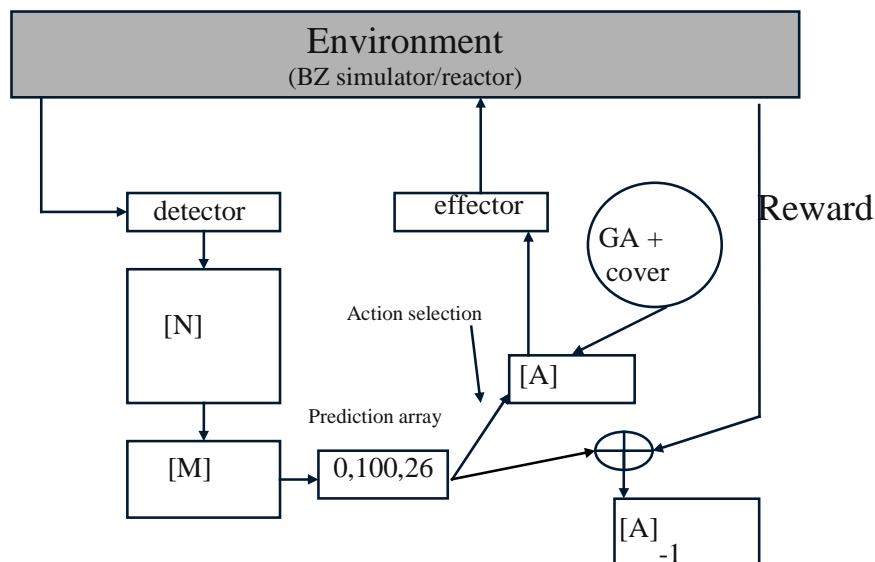


Figure 3: Schematic of the XCS Learning Classifier System.

The intention in XCS is to form a complete and accurate mapping of the problem space through efficient generalizations. In reinforcement learning terms, XCS learns a value function over the complete state/action space. In this way, XCS represents a means of using temporal difference learning on complex problems where the number of possible state-action combinations is very large (other approaches have been suggested, such as neural networks – see [Sutton & Barto, 1998] for an overview). The reader is referred to

[Butz & Wilson, 2002] for an algorithmic description of XCS and [Bull & Kovacs, 2005] for an overview of current formal understanding of its operations.

3.3 XCS Control: Simulator

The aforementioned model of the BZ system has been interfaced to an implementation of XCS in a way which approximates the envisaged hardware-wetware scenario. A 3x3 grid is initialized with a pulse of excitation in the bottom middle cell as in the wetware experiments described in section 2. Two light levels have been used thus far: one is sufficiently high to inhibit the reaction; and, the other low enough to enable it. The modeled chemical system is then simulated for 10 seconds of real time. A 9-bit binary description of the 3x3 grid is then passed to the XCS. Each bit corresponds to a cell and it is set to true if the average level of activity within the given cell is greater than a pre-determined threshold. The XCS returns a 9-bit action string, each bit of which indicates whether the high ($\Phi=0.197932817$) or low ($\Phi=0.007791641$) intensity light should be projected onto the given cell. Another 10 seconds of real time are then simulated, etc. until either a maximum number of iterations has passed or the emergent spatial-temporal dynamics of the system match a required configuration. In this initial work, a fragment is required to exist in the middle left-hand cell of the grid *only*. At such a time, a numerical reward of 1000 is given the system and the system reset for another learning trial. To be able to obtain this behavior reliably, it has been found beneficial to use an intermediate reward of 500 in the presence of a fragment in the target cell, regardless of the activity on the rest of the grid (see [Dorigo & Colombetti, 1998] for related discussions).

The XCS parameters used for this were (largely based on [Wilson, 1995]): $N=30,000$, $\beta=0.2$, $\mu=0.04$, $\chi=0.8$, $\gamma=0.71$, $\theta_{del}=20$, $\delta=0.1$, $\epsilon_0=10$, $\alpha=0.1$, $v=5.0$, $\theta_{mna}=512$, $\theta_{GA}=25$, $\rho_I = \epsilon_I = F_I = 10.0$, $p_{\#}=0.33$. Other parameters for the BZ model were $\epsilon=0.022$, $f=1.4$, $q=0.002$.

Figure 4 shows a typical light program and associated wave fragment behaviour sequence, here taking 9 steps to solve the problem, which appears optimal with the given parameters of the simulator and allowed time between XCS control steps. Figure 5 shows the reward received by the learner per trial, averaged over 3 runs. As can be seen, this approaches the maximum of 1000 in the time considered; the XCS controller is reliable in its ability to develop a fragment controller in the given scenario.

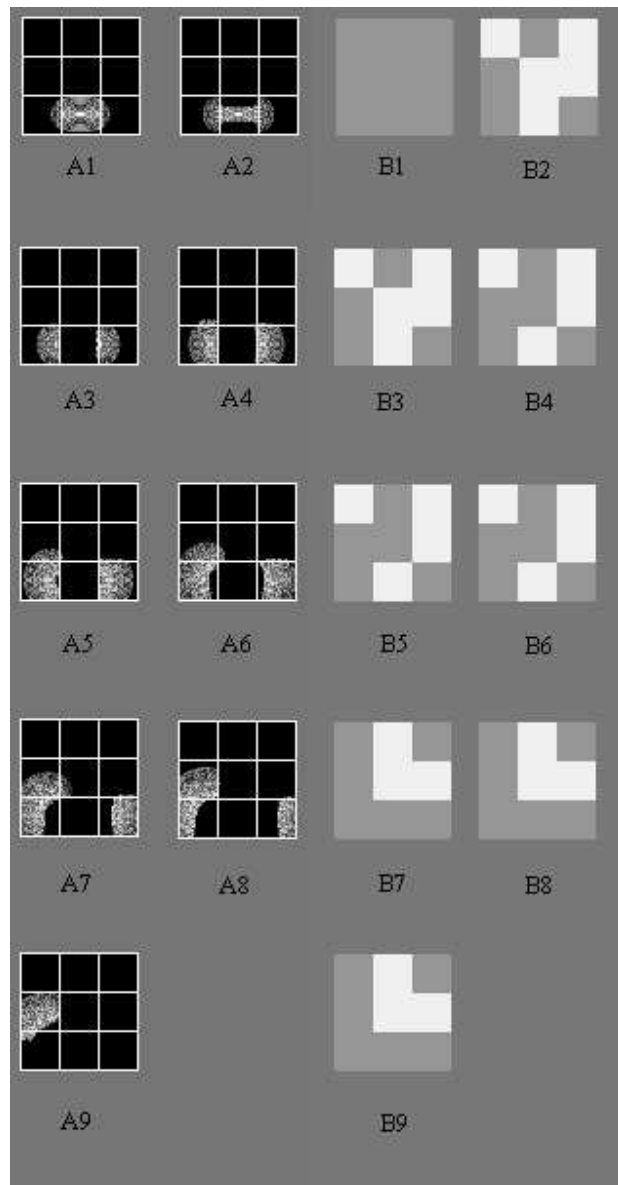


Figure 4: Example control of the simulated chemical system (A) under the learned light programs (B).

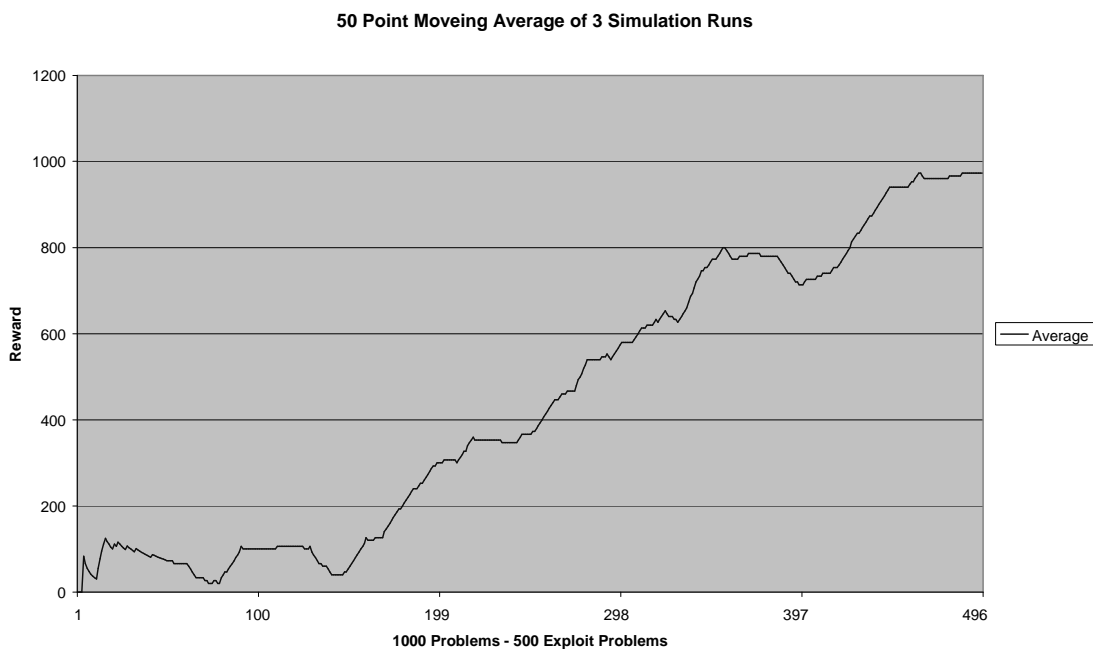


Figure 5: Showing the average reward received by the XCS controllers with increasing number of learning cycles or problems.

3.3 XCS Control: Chemistry

Given the success of the simulation experiments, the machine learning system described above was connected to the chemical system described in section 2. The scenario for the chemical experiments was the same as for the simulations, although it must be noted that there is a slightly longer delay before the XCS is able to control the initial light levels once the pulse is added to the grid due to the image processing required. Figure 6 shows an example result using the same parameters as before. This solution was discovered on trial 4 and then refined over the subsequent three trials to that shown. It can be seen that XCS has learnt to control the fragment in the real chemical system as it did in the simulation.




Figure 5: Example control of the real chemical system (A) under the learned light programs (B). Skewing of cell images is due to camera set-up.

4 Conclusions

Excitable and oscillating chemical systems have previously been used to solve a number of computational tasks. However we suggest that the lack of compartmentalisation in the majority of the previous systems limits the domain of solvable tasks thus making it difficult to realise general-purpose computing. We propose that utilising networks of coupled oscillating chemical reactions will open chemical computing to wider domains. In this paper we have presented initial results from a methodology by which to achieve the

complex task of designing such systems - through the use of machine learning techniques. We have shown using both simulated and real systems that it is possible control the behaviour of a light-sensitive BZ reaction using XCS.

Current work is extending the scenario presented to include larger grids containing many more concurrent fragments of excitation.

5 Acknowledgements

This work was supported by EPSRC grant no. GR/T11029/01. Thanks to Annette Taylor and Rita Toth for useful discussions during the course of this work.

6 References

Adamatzky, A. (2001) . *Computation in Nonlinear Media and Automata Collectives*. IoP Publishing.

Agladze K, Aliev RR, Yamaguhi T & Yoshikawa K. (1996) Chemical Diode. *J Phys Chem*. 100: 13895-13897.

Bar-Eli, K. & Reuveni, S. (1985) Stable stationary-states of coupled chemical oscillators. Experimental evidence, *J Phys Chem*. 89: 1329-1330.

Bar-Eli, K. (1985) On the stability of coupled oscillators. *Physica D* 14: 242-252.

Booker, L. (1989) Triggered Rule Discovery in Classifier Systems. In Schaffer (ed.) *Proceedings of the International Conference on Genetic Algorithms*. Morgan Kaufmann, pp265-274.

Bull, L. (2004)(ed) *Applications of Learning Classifier Systems*. Springer.

Bull, L. & Kovacs, T. (2005)(eds) *Foundations of Learning Classifier Systems*. Springer.

Butz, M. & Wilson, S.W. (2002) An Algorithmic Description of XCS. *Soft Computing* 6(3-4): 144-153.

Crowley, M.F. & Field, R.J. (1986) Electrically coupled Belousov-Zhabotinskii oscillators 1. Experiments and simulations, *J Phys Chem*. 90: 1907-1915.

Crowley, M.F. & Epstein, I.R. (1989) Experimental and theoretical studies of a coupled chemical oscillator: phase death, multistability and in-phase and out of phase entrainment. *J Phys Chem*. 93: 2496-2502.

- Dolnik, M. & Epstein, I.R. (1996) Coupled chaotic oscillators. *Physical Review E* 54: 3361-3368.
- Dorigo, M. & Colombetti, M. (1998) *Robot Shaping*. MIT Press.
- Field, R. J. & Noyes, R. M. (1974) Oscillations in chemical systems. IV. Limit cycle behavior in a model of a real chemical reaction. *J. Chem. Phys.* 60, 1877-1884.
- Gorecki, J., Yoshikawa, K. & Igarashi, Y. (2003) On chemical reactors that can count. *J Phys Chem. A* 107: 1664-1669.
- Hjelmfelt, A. & Ross, J. (1993) Mass-coupled chemical systems with computational properties. *J Phys Chem.* 97: 7988-7992.
- Hjelmfelt, A., Schneider, F.W. & Ross, J. (1993) Pattern-recognition in coupled chemical kinetic systems. *Science* 260: 335-337.
- Hjelmfelt, A., Weinberger, E.D. & Ross, J. (1991) Chemical implementation of neural networks and Turing machines. *PNAS* 88: 10983-10987.
- Hjelmfelt, A., Weinberger, E.D. & Ross, J. (1992) Chemical implementation of finite-state machines. *PNAS* 89: 383-387.
- Holland, J.H. (1975) *Adaptation in Natural and Artificial Systems*. University of Michigan Press.
- Holland, J.H. (1986). Escaping brittleness: the possibilities of general-purpose learning algorithms applied to parallel rule-based systems. In Michalski, Carbonell, & Mitchell (eds) *Machine learning, an artificial intelligence approach*. Morgan Kaufmann.

Holz, R. & Schneider, F.W. (1993) Control of dynamic states with time-delay between 2 mutually flow-rate coupled reactors *J Phys Chem.* 97: 12239.

Jung, P., Cornell-Bell, A., Madden, K.S. & Moss, F.J. (1998) Noise-induced spiral waves in astrocyte syncytia show evidence of self-organized criticality. *Neurophysiol.* **79**: 1098-1101.

Kawato, M. & Suzuki, R. (1980) Two coupled neural oscillators as a model of the circadian pacemaker. *J Theor Biol*, 86: 547-575.

Krug, H.-J., Pohlmann, L. & Kuhnert, L. (1990) Analysis of the modified complete Oregonator accounting for oxygen sensitivity and photosensitivity of Belousov-Zhabotinsky systems, *J. Phys. Chem.* 94, 4862.-4866.

Kuhnert, L., Agladze, K.I. & Krinsky, V.I. (1989) Image processing using light sensitive chemical waves, *Nature* 337: 244-247.

Laplante, J.P., Pemberton, M., Hjelmfelt, A. & Ross, J. (1995) Experiments on pattern recognition by chemical kinetics. *J Physical chemistry* 99: 10063-10065.

Lebender, D. & Schneider, F.W. (1994) Logical gates using a nonlinear chemical reaction, *J. Phys. Chem.* 98: 7533-7537.

Michalewicz, Z. & Fogel, D. (1999) *How to Solve it*. Springer Verlag.

Motoike, I.N., Yoshikawa, K., Iguchi, Y. & Nakata, S. (2001) Real time memory on an excitable field. *Physical Review E* 63 (036220): 1-4.

- Rossler, O.E. (1974) In M. Conrad, W. Guttinger & M. Dal Cin (eds) Lecture notes in Biomathematics. Springer, pp 399-418 and 546-582.
- Sielewiesiuk, J. & Gorecki, J. (2002) Passive barrier as a transformer of chemical frequency. *Journal of physical chemistry A* 106: 4068-4076.
- Steinbock, O., Kettunen, P. & Showalter, K. (1996) Chemical wave logic gates. *J Phys Chem*, 100: 18970-18975.
- Steinbock, O., Toth, A. & Showalter, K. (1995) Navigating Complex Labyrinths: Optimal Paths from chemical waves. *Science* 267: 868-871.
- Stuchl, I. & Marek, M. (1982) Dissipative structures in coupled cells: experiments. *J. Phys Chem.* 77: 2956-63.
- Sutton, R. & Barto, A. (1998) *Reinforcement Learning*. MIT Press.
- Toth, A., Gaspar, V. & Showlater, K. (1994) Signal transmission in Chemical Systems: Propagation of chemical Waves through capillary tubes. *J. Phys Chem.* 98: 522-531.
- Wang, J., Kadar, S., Jung, P. & Showalter, K. (1999) Noise driven Avalanche behaviour in subexcitable media. *Physical Review Letters* 82: 855-858.
- Watkins, C.J. (1989) Learning from Delayed Rewards. Ph.D. Thesis, Cambridge University.
- Wilson, S.W. (1995) Classifier Fitness Based on Accuracy. *Evolutionary Computing* 3: 149-175.