Hebbian Learning by a Simple Gene Circuit

Chrisantha Fernando
Systems Biology Centre
University of Birmingham
Birmingham, Edgbaston, B15 2TT, UK
C.T.Fernando@cs.bham.ac.uk

Jon Rowe
School of Computer Science
University of Birmingham
Birmingham, Edgbaston, B15 2TT, UK
J.E.Rowe@cs.bham.ac.uk

ABSTRACT

We present a simple design for a gene circuit capable of Hebbian learning. Differential equation models of the circuit are tested on four tasks; sensitization, associative learning, self-organzation of gene receptive fields by lateral inhibition, and supervised training of a gene perceptron to make a simple linear classification. We suggest that such motifs may exist in complex eukaryotic cells which have been found to contain significant RNA networks. However, the behavioural evidence for associative learning in simple eukaryotes e.g. paramecia, is contradictory. Natural adaptive uses for a Hebbian learning circuit could include robustness of gene expression to developmental noise, and classical conditioning in complex environments. Engineering applications include design of unicellular biosensors that learn the principle components of signals in a complex environment, or the production of cells capable of expression of a gene under the conditions for which that gene is trained to be expressed.

1. INTRODUCTION

An unexpectedly large amount of RNA is non-coding, i.e. is not translated into proteins. For example as much as ten times more of the genomic sequence in chromosome 21 and 22 is transcribed than accounted for by predicted and characterized exons [14]. Understanding the function of such non-coding RNA networks is a major scientific problem [4]. In some cases non-coding RNA is conserved across species, so it is likely to be functional [6]. Small interference RNA and micro RNAs have been shown to be important in eukaryotic gene regulation [7]. The ratio of non-coding to protein-coding DNA rises as a function of developmental complexity, suggesting a causal relationship [18]. It has been proposed that complex RNA networks could be undertaking computations in eukaryotic cells [21]. To answer how such RNA networks could do their computations the neural network metaphor has been applied to in vitro transcriptional circuits, for example, Kim et al describe a DNA transcriptional switch [16] and the resulting network that is formally equivalent to the Hopfield neural network model [12] [1]

Of what kinds of learning are single-celled eukaryotes capable? Even bacteria are capable of implicit memory [13] such as habituation and sensitization, e.g. perfect adaptation in chemotaxis [20]. However, associative learning, e.g. classical, instrumental conditioning and reinforcement learning, are more challenging. Early experiments claimed to have demonstrated classical conditioning in paramecia [10] [9]. Paramecia were unaffected by a vibratory stimulus but rapidly swam away when exposed to an electric shock. If individual Paramecia were conditioned by exposing them to an electric shock paired with a vibratory stimulus, they displayed rapid swimming movements when later given the vibratory stimulus alone. Other experiments by Bramstedt, Grabowski and Mirsky and Katz [15] have been less conclusive; they used an UCS (unconditioned stimulus) of heat and a CS (conditioned stimulus) of light. Finally the finding that paramecia could learn to escape from capillary tubes better after experience in previous tubes has been shown not to be due to associative learning [11]. Experiments were also conducted with bacteria to investigate behavior during complex chemotaxis tasks [1]. It has been proposed that such early single cell learning and perceptual mechanisms may have been the predecessors of multi-cellular neural and hormonal signaling [17] [8]. We are not aware of any other experimental efforts to demonstrate associative learning in single-celled eukaryotes. In this paper we hypothesize a simple motif in a transcription network that could approximate Hebbian learning. We use models to demonstrate how such a motif could be used in unsupervised and supervised learning tasks.

2. METHODS

Hebbian learning in its simplest form follows the rule

$$\tau_w \frac{\mathbf{dw}}{dt} = v\mathbf{u} \tag{1}$$

where $\mathbf{w}$ is a vector of synaptic strengths or weights, $\mathbf{u}$ is a vector of input (pre-synaptic) activities and v is the output (post-synaptic) activity, and τ_w is a small constant training rate. Although modifications of this rule exist, e.g. the

[1]This is a leaky-integrator network similar to the CTRNNs used by Randall Beer [2].

covariance rule and the BCM rule [5], we implement an approximation of this simplest rule as a transcriptional module shown in figure 1

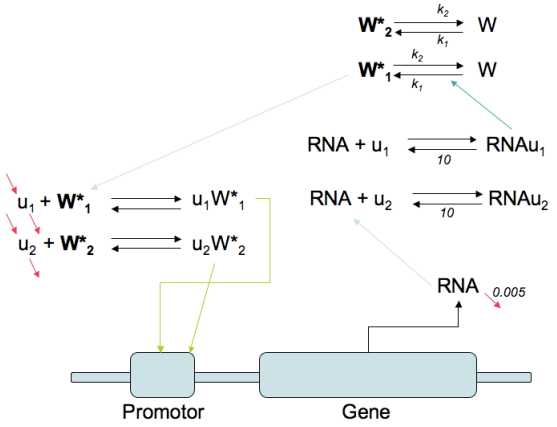

Figure 1: A transcription network capable of approximating Hebbian learning. Two TFs, u_1 and u_2 can bind to the promotor only if they are bound to their specific W molecule. W molecules represent the weight, see text for details.

Transcription factor (TF) concentrations, u_1 and u_2 are analogous to the presynaptic activities. $[RNA]$ is analogous to the post-synaptic activity. The $u_i W_i^*$ complex acts to stabilize the RNA polymerase on the promotor, and hence to produce more RNA. The Hill equation is used to describe gene activation by this complex. The RNA produced binds to the transcription factors forming specific $RNAu_i$ complexes, again in a rapid equilibrium. These complexes slowly catalyze the conversion of W precursor into activated specific W_i^*. W_i^*also decays very slowly, so that its concentration stores information about the history of $[RNA] \times [u_i]$ received over a long time period. Cellerator for mathematica was used to simulate the system [19]. The cellerator representation is shown for 2 different transcription factors, u_1 and u_2, see figure 2.

```
{u1i → u1, 0.01},
{u2i → u2, 0.01},
{u1 → ∅, 0.01},
{u2 → ∅, 0.01},

{RNA + u1 ⇌ RNAu1, Kf, Kb*10},
{RNA + u2 ⇌ RNAu2, Kf, Kb*10},
{RNAu1 + w1 → RNAu1 + w1s, 0.0001},
{RNAu2 + w2 → RNAu2 + w2s, 0.0001},
{w2s → w2, Kb/10000000},
{w1s → w1, Kb/10000000},
{u1 + w1s ⇌ uw1s, Kf, Kb},
{u2 + w2s ⇌ uw2s, Kf, Kb},
{uw1s ↦ RNA, hill[vmax → 0.1, nhill → 1, khalf → 1]},
{uw2s ↦ RNA, hill[vmax → 0.1, nhill → 1, khalf → 1]},
{RNA → ∅, 0.005}
```

Figure 2: Cellerator representation of chemical reactions. The forward (and backward) rates of the reaction is shown on the right of the reaction. Kf and Kb are set to 1. Note the slow production of W^* from W and the even slower reverse reaction from W^* back to W. The null symbol in the last reaction indicates decay of RNA to nothing. The two penultimate reactions represent the $u_i W_i^*$ complex activating RNA production according to the Hill equation.

3. RESULTS

The above circuit is tested in three unsupervised and one supervised task.

3.1 Unsupervised Hebbian Learning

First is demonstrated the capability of W_i^* molecules to mediate long term potentiation in a Hebbian manner, see figure 3. We conducted an experiment with two (transcription factor) TF inputs acting on a non-coding gene. During a training period the TFs were applied as bolus injections, the masses of the two TFs in each injection were chosen from a 2D Gaussian distribution with equal variance and unequal means, resulting in the elliptical distribution of points in the space {[TF1],[TF2]}. TF1 was applied in boluses of mean 2 (c.u) concentration units and standard deviation 0.1 c.u whereas TF2 was applied with mean 1 c.u and s.d. 0.1 c.u. After 10 stimulations by simultaneous injections of TF1 and TF2, separate test injections of TF1 and TF2 were carried out at unit concentration, and the resulting RNA spike was measured. Even though the test injections of TF1 and TF2 were identical in concentration, more RNA was produced upon injection of equal quantities of TF1 than TF2. This is because of Hebbian learning in which $[W_1^*]$ increased to 0.0009 at the end of the training phase, whereas $[W_2^*]$ only increased to 0.0005 at the end of the training phase. $[W_i^*]$ acts as a long term memory of the history of TF activation.

The basic Hebb rule is susceptible to saturation of synaptic weights [5]. So is the above implementation in which all the W_i precursor molecules could potentially be used up to produce the maximum concentration of W_i^* molecules, after which no further information about stimulus history could be encoded. To prevent saturation, the extent to which the $RNAu_i$ complex catalyses the production of W_i^* must be adjusted to suit the intensity of TF input. A modification of the network is to use a common W precursor, rather than separate precursors for each W_i^* molecule. This modification implements competition which is analogous to synaptic normalization rules that model the limited set of LTP (long-term potentiation) producing resources available at the synapse.

Secondly, the above network is used to demonstrate how associative learning could straightforwardly take place at one gene in a manner analogous to associative learning in a neuron [22]. Let several TFs potentially act at the promotor of this gene, for now we consider just two. Let TF1 represent the presence of a food molecule in the environment (the unconditioned stimulus), so let it bind strongly and innately to the promotor. However let the other TF signals bind very weakly; they represent other external chemicals that could potentially be associated with the food molecule (potential conditioned stimuli), and even precede the food molecule temporally in certain environments. Assume for example that TF2 slightly precedes TF1 in an environment.

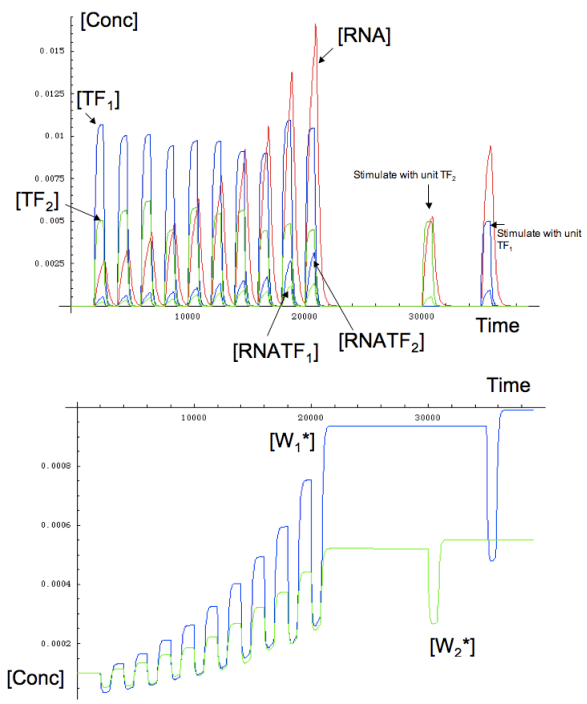

Figure 3: Long term potentiation at a promotor is demonstrated by the greater sensitivity of the promotor to TF1 than TF2 after a history of stimulation with a greater amount of TF2 than TF1. Initial concentration of W^* molecules is 0.0001, and initial concentration of inactive W molecules is 1.

Figure 4 shows that the above network can mediate associative learning.

Thirdly we demonstrate an analogy to the formation of ocular dominance stripes, orientation domains, and other self-organized maps in the cortex in which individual neurons develop specific sensitivity for a subset of input features based on competitive learning [5]. An equivalent problem can be envisaged to exist if an intra-cellular RNA network is to form representations of complex environments outside the cell, that may differ significantly between lifetimes. It may be adaptive that the same gene be activated to different extents by different transcription factors depending on the developmental history of the cell. We consider a very simple case here. Imagine now two non-coding genes, both receiving inputs from TF1 and TF2 wired up in the same way as in figure 1, except that the RNAs show 'lateral inhibition', i.e. the RNA from one gene destroys the RNA of the other, and vice versa. How does the 'receptive field' of the two genes depend on the correlation between stimulation by TF1 and TF2? Figure 5 shows a simulation of the above network stimulated by TF inputs of equal mean and variance $[1, 0.1]$ and in a regular in-phase spike train, with and without lateral inhibition.

3.2 Supervised Hebbian Learning

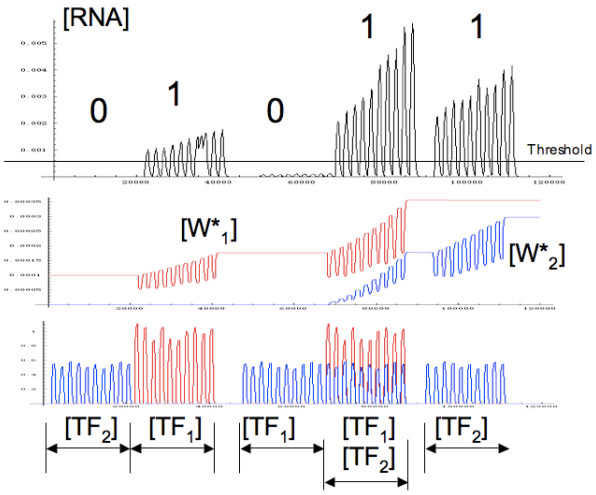

Figure 4: There are 5 phases to the experiment. Initial $W_1^* = 0.0001$, Initial $W_1^* = 10^{-9}$ Phase 1: CS stimulus with TF_2, Phase 2: UCS stimulus with TF_1, Phase 3: CS stimulus with TF_2, Phase 4: Paired CS and UCS, Phase 5: CS stimulus with TF_2. $[RNA]$ represents the output, which is 1 if above the threshold shown. The UCS has been associated with the CS. Note $[W_2^*]$ only increases significantly in phase 4 when UCS and CS are paired.

Supervised learning tasks involve either learning to classify inputs into two categories or function approximation, in which the output is trained to be a particular function of the input. This can be achieved using Hebbian learning in which a molecule is applied "externally" by a teacher in response to the TF inputs in order to pattern the weights so that a proper classification results. A single gene with many TFs subject to Hebbian learning has computational abilities similar to a perceptron, if the TFs act independently, although more complex transfer functions are possible [3]. Training is accomplished by injecting an inhibitor molecule that produces LTD (long term depression) rather than the standard LTP (long term potentiation) when the desired output state is low RNA (e.g. class II rather than class I). When this molecule is present, weights decrease in proportion to the product of $[u_i]$ and $[RNA]$. Figure 6 shows a modification of the network in figure 1 that allows training to take place. Molecule I is the inhibitor, which converts W_i^* molecules back to W_i molecules. It therefore produces LTD for class II stimuli. Figure 7 shows the training and testing of a gene perceptron on a 2D problem in which a simple linear discrimination is learned.

4. DISCUSSION

The stability of Hebbian learning in all the above cases depends on the ability of the cell to maintain W_i^* concentrations over long periods of time without decay. More complex Hebb type rules [5] could conceivably be implemented to maximize robustness. Such modifications could involve additional homeostatic mechanisms that implement the chemical equivalent of "synaptic normalization". We have demonstrated that by a very simple positive feedback

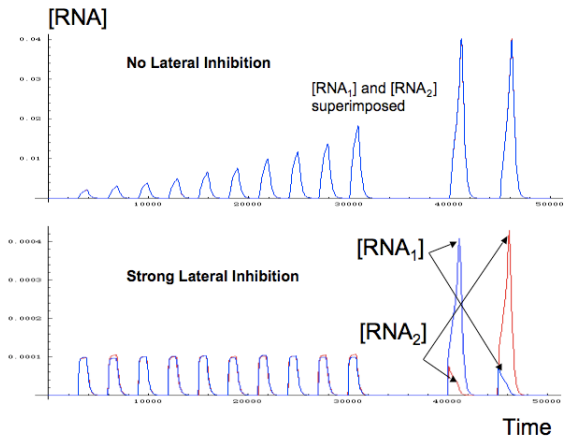

Figure 5: Lateral inhibition results in different genes being preferentially activated by different TF inputs, whereas without lateral inhibition there is equal representation by the gene of both inputs. A bolus of TF_1 was input at time 40000, and a bolus of TF_2 was input at time 45000, to test the response. Lateral inhibition is implemented by the reactions $RNA_1 + RNA_2 \rightarrow RNA_1$, and $RNA_2 + RNA_1 \rightarrow RNA_2$, occurring at a very high rate 1000.

process from RNA to the stability of binding of its transcription factors, many of the features of Hebbian learning that have been investigated in theoretical neuroscience may apply in Systems Biology. When Donald O. Hebb proposed his principle in the book entitled "The Organization of Behaviour" in 1949, there was no evidence for it, and no means to verify it. We suggest that researchers working on RNA networks should explicitly search for Hebbian mechanisms acting in single celled and multi-cellular Eukaryotes. Also, applications immediately suggest themselves in which cells can be engineered to express a gene under conditions modifiable by training, for example a drug could be programmed to be released in a manner which is specific to the individual patient.

4.1 Acknowledgments

Thanks to the European 6th Framework Grant on the Evolution of Cell Signaling Networks (ESIGNET) for funding. Thanks to Dov Stekel for useful discussions.

5. REFERENCES

[1] J. Adler and W. Tso. 'decision'-making in bacteria: chemotactic response of escherichia coli to conflicting stimuli. *Science*, 184:1292–1294, 1974.

[2] R. Beer. Parameter space structure of continuous-time recurrent neural networks. *Neural Computation*, 18:3009–3051, 2006.

[3] L. Bintu, N. Buchler, H. Garcia, T. H. Ulrich Gerland, J. Kondev, and R. Phillips. Transcriptional regulation by the numbers: models. *Current Opinion in Genetics and Development*, 15:116–124, 2005.

[4] J.-M. Claverie. Fewer genes, more noncoding rna. *Science*, 309:1529–1530, September 2005.

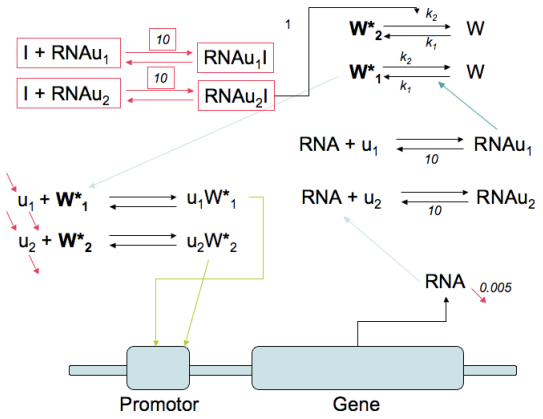

Figure 6: The inhibitor I is injected during presentation of stimuli destined for class II. It binds to the $RNAu_i$ complex. The resulting molecule $RNAu_iI$ inactivates W_i^* back into W_i, so reducing the weight in proportion to the product of $[RNA]$ and $[u_i]$.

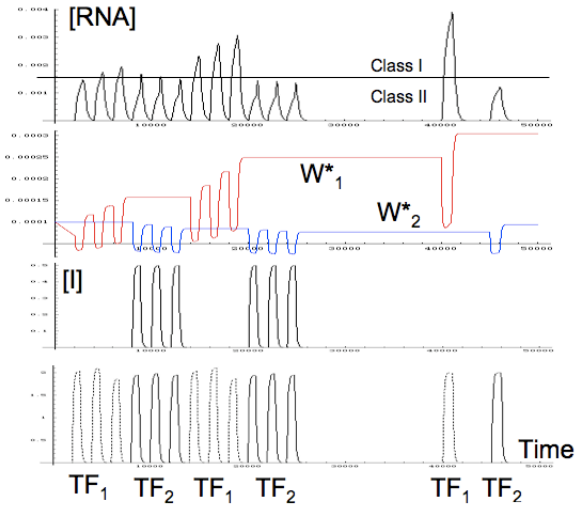

Figure 7:

[5] P. Dayan and L. Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. The MIT Press, Cambridge, Massachusetts, 2001.

[6] C. Dennis. The brave new world of rna. *Nature*, 418:122–124, July 2002.

[7] A. Fire. Potent and specific genetic interference by double-stranded rna in caenorhabditis elegans. *Nature*, 391:806–811, 1998.

[8] J. Haldane. La signalisation animale. *Annee Biol*, 58:89–98, 1954.

[9] T. Hanzel and W. Rucker. Trial and error learning in paramecium: a replication. *Behavioural Biology*, 7(6):873–80, 1972.

[10] T. Hennessey, W. Rucker, and C. McDiarmid. Classical conditioning in paramecia. *Animal Learning and Behavior*, 7:419–23, 1979.

[11] D. Hinkle and D. Wood. Is tube-escape learning by

protozoa associative learning? *Behav Neurosci*, 108(1):94–99, 1994.

[12] J. Hopfield. Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences*, 81:3088–3092, 1984.

[13] E. Kandel, J. Schwartz, and T. Jessell. *Principles of Neural Science*. Elsevier, 1986.

[14] P. Kapranov. Large-scale transcriptional activity in chromosomes 21 and 22. *Science*, 296:916–919, 2002.

[15] M. Katz and W. Deterline. Apparent learning in paramecium. *Journal of Comparative Physiology and Psychology*, 51(2):243–247, 1958.

[16] J. Kim, J. Hopfield, and E. Winfree. Neural network computation by in vitro transcriptional circuits. In *Advances in Neural Information Processing Systems: NIPS proceedings*, pages 681–688, March 2004.

[17] D. Koshland. Bacterial chemotaxis in relation to neurobiology,. *Annual Review of Neurosciences*, 3:43–75, 1980.

[18] J. Mattick. Rna regulation: a new genetics? *Nature Reviews Genetics*, 5:316–323, 2004.

[19] B. Shapiro, A. Levchenko, E. Meyerowitz, B. Wold, and E. Mjolsness. Cellerator: extending a computer algebra system to include biochemical arrows for signal transduction simulations. *Bioinformatics*, 22(5):677–678, March 2003.

[20] T. Shimizu, S. Aksenov, and D. Bray. A spatially extended stochastic model of the bacterial chemotaxis signalling pathway. *Journal of Molecular Biology*, 329:291–309, 2003.

[21] E. Tannenbaum. An rna-centered view of eukaryotic cells. *Biosystems*, pages 316–323, to appear 2005.

[22] T. Trappenberg. *Fundementals of Computational Neuroscience*. Oxford University Press, Oxford, UK, 2002.