ESANN'1993 proceedings - European Symposium on Artificial Neural Networks
Brussels (Belgium), 7-8-9 April 1993, D-Facto public., ISBN 2-9600049-0-6, pp. 65-70

Population Coding in a Theoretical Biologically Plausible Network

Gregory R. Mulhauser

Department of Philosophy, University of Edinburgh
Edinburgh, Scotland EH8 9JX

Abstract. I propose a simple network architecture and unsupervised learning algorithm which avoids some weaknesses common to other networks and which I claim may produce low cost population coding. The network models the synchronising influence in biological systems of ephaptic interactions and gap junctions between physically proximal neurons through virtual lateral connections within a single layer. The spike frequency adaptation behaviour of some real neurons is implemented by means of an idealised fatigue factor for each node in the network, and the connection weights are updated through a simple correlation rule which does not require backpropagation.

1. Introduction

As their name suggests, artificial neural networks are inspired by the networks of nerve cells in brains but are designed for use on electronic computers. Unfortunately, in the quest to simplify artificial neural models and reduce the computational overhead required for their execution, a number of poorly understood characteristics of biological neural networks have been abstracted away or ignored altogether. The model I propose showcases novel approaches for implementing certain of these neuronal characteristics which may return performance improvements when integrated into existing artificial neural networks. The model makes minimal computational demands and is unique in combining lateral connections within a layer to represent gap junctions and ephaptic interactions, a fatigue factor for each node, and a Hebbian (1) learning rule.

While I am not yet able to report results of empirical testing of the network architecture and learning algorithm I describe, structural similarities to other networks suggest that we might expect the net to produce population coding which resembles something like a cross between Kohonen's modified competitive learning strategy (2) and principal-components algorithms.

2. Preliminary Architectural Details

The network I describe consists of an input layer completely connected to a second layer that includes lateral connections. It is within this layer that we should expect population coding to emerge. Connection weights are in the range [-1...1] and node outputs are in the range [0...1]. Node outputs, except for the fatigue modification which I describe in section 4, may be calculated from the linear sum of weighted inputs according to a sigmoid curve or other standard function. Note that no derivatives are taken of the output function. Thus it may be desirable to choose an output function which allows more information effectively to be represented in a node's output value, much as information may be represented in the output frequency of a real neu-

ESANN'1993 proceedings - European Symposium on Artificial Neural Networks
Brussels (Belgium), 7-8-9 April 1993, D-Facto public., ISBN 2-9600049-0-6, pp. 65-70

ron. This might even provide a few of the benefits of pulse coded networks. [3, 4, 5, 6] Biases are not specifically addressed but may be implemented in standard ways if desired.

3. Architecture

In biological neural networks, local neuronal activity is partially synchronised by means of gap junctions and ephaptic interactions. The former are actual physical connections between adjacent cells made by large macromolecules which extend through both cell membranes and contain water-filled pores. [7, 8] Ephaptic interactions do not require a physical connection between nerve cells, but they have a similar effect: the electrical currents set up by the flow of ions across the membrane of one neuron may induce electrical currents in nearby cells. [8, 9]

While the strategy used by Kohonen to enable physically adjacent nodes in a competitive learning network to code similar input patterns might be seen as one possible abstraction of this feature, I propose a simpler abstraction in which the input layer is completely connected to a single layer, within which each node is directly connected to its three nearest neighbours. The architecture may be viewed geometrically as a regular pattern of hexagons, as in Figure 1. For simplicity in implementing the connections as a data structure, the pattern can simply be stretched as shown in Figure 2.

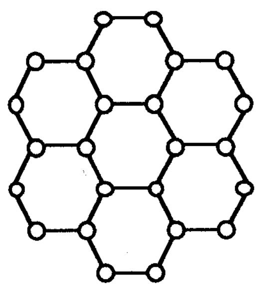

Fig. 1. Horizontal Connections

For some applications, it may be desirable to maintain the pattern of three lateral connections per node by warping the structure into a third dimension and connecting nodes shown here with only two connections to the corresponding nodes on the opposite side. Nodes might even be arranged in patterns of both hexagons and pentagons and the whole structure transformed into a three dimensional football shape.

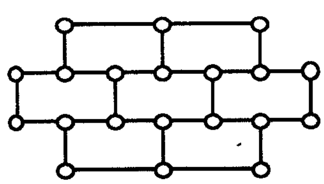

Fig. 2. Simplified Horizontal Connections

In implementing this connection pattern, we may treat lateral synapses as "virtual connections" and view all connections as being *between* layers. This is accomplished

ESANN'1993 proceedings - European Symposium on Artificial Neural Networks
Brussels (Belgium), 7-8-9 April 1993, D-Facto public., ISBN 2-9600049-0-6, pp. 65-70

by introducing a third layer with a number of nodes identical to the layer in which we are implementing "virtual connections". We may then dispense with horizontal connections between nodes and calculate their effects in the subsequent layer. Each node is connected to "itself" in the next layer by a connection whose strength is permanently set to one. This node also receives inputs from each of its neighbours in the preceding layer, as in Figure 3. We may thus calculate first the response of each node to the firing patterns of the input layer and then calculate separately the influence on it of the activity of the nodes in its horizontally local neighbourhood. There may be different advantages to setting each of the virtual lateral connections to a uniform initial strength such as 1/3 or simply to randomising them across the net.

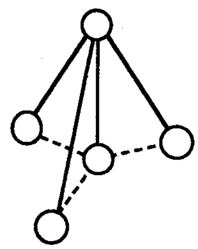

Fig. 3. Implementing Horizontal Connections with an Extra Layer

4. Learning

Because of its role in the function for updating connection strengths, I describe first the fatigue factor associated with each neuron. Fatigue in real neurons, known as spike frequency adaptation, is the tendency of some kinds of neurons to decrease their firing frequency during sustained depolarisations. This behaviour may prevent the same neurons from becoming active in representations of too many distinct input patterns. For the present model, I suggest a fatigue value ϕ_x in the range $[0\ldots1]$ (where larger numbers indicate more fatigue) which is applied to calculate a "real output" O_x for a node x by a straightforward modification of an output out_x calculated with a sigmoid or other function from the weighted sum of inputs to the node:

$$O_x = \text{out}_x(1 - \phi_x) \tag{1}$$

When calculating a new fatigue value ϕ'_x for a node x, we should like to take into account its recent firing history as represented by the previous fatigue value as well as the magnitude of its present output. If its present output is low relative to its fatigue value, the fatigue should fade quickly. Likewise, if its present output is high, fatigue should achieve significance rapidly. This can be modelled as in equations (2) and (3),

$$\phi'_x = \phi_x + \varepsilon(O_x - \phi_x)^2 \qquad \text{when } O_x > \phi_x \tag{2}$$

$$\phi'_x = (\phi_x)^2 \qquad \text{when } O_x \leq \phi_x \tag{3}$$

where ε is a parameter set in advance to adjust how significant an influence fatigue is allowed to become. Fatigue values should be set to 0 before the network is trained.

ESANN'1993 proceedings - European Symposium on Artificial Neural Networks
Brussels (Belgium), 7-8-9 April 1993, D-Facto public., ISBN 2-9600049-0-6, pp. 65-70

The learning rule I describe is inspired by Hebb's theory that synaptic coupling increases when the activity of converging network elements is coincident. We would like to increase the weight of a connection between two nodes when the nodes produce similar output. A high correlation paired with a low present connection strength should correspond to a large increase in weight, while a low correlation paired with a high connection strength should correspond to a large decrease in connection strength. We seek a relationship something like Figure 4, where the x-axis represents the present connection weight, the y-axis the correlation between the nodes' outputs, and the z-axis the amount by which the connection should be updated.

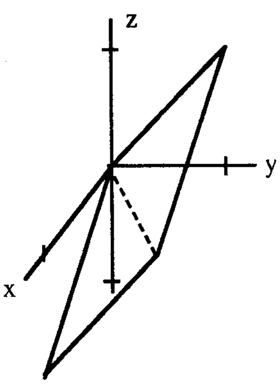

Fig. 4. Weight Updating

The relationship might be nonlinearised as in Figure 5 to increase the responsiveness of the network, although such a nonlinearisation would need to be applied carefully to avoid increasing the sensitivity of the network to a level where it becomes unable to converge on a stable pattern of connection strengths.

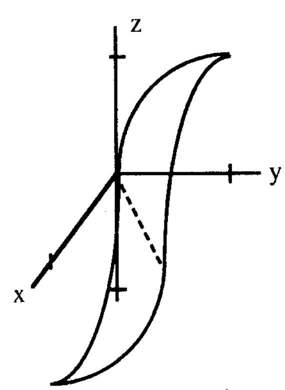

Fig. 5. Nonlinearised Weight Updating

I suggest that equation (4) establishes the relationship indicated in Figure 4 between a $w_{x,y}$ scaled to $[0\ldots1]$ (x-axis), representing the present connection weight between nodes x and y, $1 - |O_x - O_y|$ (y-axis), representing the correlation between the

ESANN'1993 proceedings - European Symposium on Artificial Neural Networks
Brussels (Belgium), 7-8-9 April 1993, D-Facto public., ISBN 2-9600049-0-6, pp. 65-70

outputs O of nodes x and y, and $\delta(w_{x,y})$ (z-axis), the amount by which to update the connection strength.

$$\delta(w_{x,y}) = \sigma\left[1 - |O_x - O_y| - \frac{w_{x,y} + 1}{2}\right] \tag{4}$$

Here $w_{x,y}$, which ranges over [-1...1], is scaled to the range [0...1] by the last term inside the brackets.

The σ term relates to the implementation of the fatigue factor. We are seeking a strategy which will avoid large modifications to a connection weight when both nodes in question are highly fatigued. This might occur, for instance, in situations where a similar input pattern has been presented over and over again and the nodes coding it have begun to decrease their outputs in response to fatigue. Equation (5) indicates one method of implementing the strategy:

$$\sigma = [1 - \min(\phi_x, \phi_y)]\,[\max(O_x, O_y)] \tag{5}$$

The second bracketed term in (5) also ties the degree to which a weight is altered to the strength of the nodes' outputs. Thus, high correlation must be coupled with high output magnitude to invoke a maximal change in connection strength.

5. Drawbacks and Directions for Future Research

While the techniques described here are inspired by theories of how real living neurons function, they are not without their limitations, in part because these theories may not accurately reflect real neural behaviour. The onset and decay of fatigue in real neurons is poorly understood, and the effects of ephaptic interactions and gap junctions are difficult to quantify. Dendritic growth characteristics are actively researched, but it is impossible to say yet just how correct is Hebb's theory of synaptic development.

The network implementation makes minimal demands on computational resources and may lend itself to the construction of larger but faster nets, yet the specific strategies at work in the net might prove problematic. For instance, the network may prove too sensitive to the choice of an ε value in equation (2). Equations (2) and (3) may allow fatigue to become significant too rapidly or to decay too slowly. Also, depending on the types of data presented to the input layer, it may be necessary to place a bound on the value of the virtual horizontal connection weights to prevent interactions within this layer from becoming too significant. Finally, the virtual horizontal connections may produce unusual relationships between the number of nodes in the input layer and the desirable total number of connections in the network.

This architecture and learning algorithm invite further research into the performance of the network I have defined as well as the performance of other networks modified with some of the techniques I have described. Virtual horizontal connections as a simple means of encouraging population coding should be tested and compared with other established methods, and the usefulness of fatigue factors to promote uniform distributions of activity may be examined in any net where population coding is the main aim. The learning algorithm itself may be improved and tested as a substitute for less biologically plausible algorithms currently enjoying widespread use.

Another promising avenue for research with this type of network is the exploration of schemes whereby *new* horizontal connections could be "grown" between highly cor-

ESANN'1993 proceedings - European Symposium on Artificial Neural Networks
Brussels (Belgium), 7-8-9 April 1993, D-Facto public., ISBN 2-9600049-0-6, pp. 65-70

related neurons with a speed related to the distance between the neurons in question. This dynamic approach might allow synchronisation and lateral reinforcement between spatially separated but functionally related network nodes. Finally, a feedback mechanism might be implemented to help the net converge on stable population codings. An exact copy of the connection weights between the input and second layer might be used to extract an output pattern from the layer in which we are seeking the population coding. This pattern might then be averaged with the original input pattern and fed through the net. With this method, responses consistent with this "reconstruction" standard of representation could be reinforced.

References

1. D.O. Hebb: The organization of behavior. New York: Wiley 1949
2. T. Kohonen: Self-organization and associative memory. New York: Springer-Verlag 1984
3. Mats Gustafsson, Lars Asplund, Olle Gällmo, and Ernst Nordström: Pulse coded neural networks for hardware implementation. Presented 9 September 1992 at the First Swedish National Conference on Connectionism in Skövde, Sweden
4. Bart Kosko: Neural Networks and fuzzy systems: a dynamical approach to machine intelligence. Englewood Cliffs, New Jersey: Prentice-Hall International 1992
5. G. Duchateau and Anders Lansner: A bayesian artificial neural network with spiking units. Report TRITA-NA-P9101, Royal Institute of Technology (Sweden) 1991
6. Mark A. Gluck, D.B. Parker, and E.S. Reifsnider: Learning with temporal derivatives in pulse-coded neuronal systems. Neural Information Processing Systems 1, pp. 195-203 (1989)
7. David A. McCormick: Membrane properties and neurotransmitter actions. In: Gordon M. Sheperd (ed.): The synaptic organization of the brain, 3 ed. Oxford: Oxford University Press 1990, pp. 32-66
8. S.W. Kuffler, J.G. Nicolls, and A.R. Martin: From neuron to brain. Sunderland, Massachusetts: Sinauer Associates 1984
9. B. Hille: Ionic channels of excitable membranes. Sunderland, Massachusetts: Sinauer Associates 1984

Acknowledgement

I am grateful to HM Government's Marshall Aid Commemoration Commission for financial support of my research during the time this paper was conceived.