Communicating Neurons
- an alternative connectionism -

J.M.Bishop and S.J.Nasuto.

Department of Cybernetics, University of Reading,
Whiteknights, Reading, Berkshire, UK.
RG6 6AY

Key Words: Connectionism; feature extraction, visual search.

Abstract. *Since the advent of the McCulloch/Pitts model, the prevailing metaphor at the core of most connectionist thought is that of neuronal operation being defined in terms of computation. At the WnnW '93 conference a novel weightless network (the Hybrid Stochastic Diffusion Network,* ***HSDN****) was introduced. It has since been shown that this network will locate a target pattern on a retina with probability one in a time growing at most linearly with retinal size. Within the Hybrid Stochastic Diffusion Network, neurons communicate bi-variate information - indeed the central feature of the HSDN is neural operation based on communication rather than computation. This paper outlines the case for a shift of neuronal metaphor from the prevailing computational view to one more grounded in communication.*

Corresponding author:

Dr. J.M.Bishop,
Department of Cybernetics,
University of Reading,
Whiteknights,
Reading,
Berkshire, UK.
RG6 2AY.

Telephone: *(UK) 0118 9316388*
Fax. *(UK) 0118 9318220*
E-mail: *J.M.Bishop@reading.ac.uk*

Communicating Neurons: an alternative connectionism

J.M.Bishop and S.J.Nasuto
Dept. Cybernetics, University of Reading, UK.

Abstract. *Since the advent of the McCulloch/Pitts model, the prevailing metaphor at the core of most connectionist thought is that of neuronal operation being defined in terms of computation. At the WnnW '93 conference a novel weightless network (the Hybrid Stochastic Diffusion Network,* ***HSDN****) was introduced. It has since been shown that this network will locate a target pattern on a retina with probability one in a time growing at most linearly with retinal size. Within the Hybrid Stochastic Diffusion Network, neurons communicate bi-variate information - indeed the central feature of the HSDN is neural operation based on communication rather than computation. This paper outlines the case for a shift of neuronal metaphor from the prevailing computational view to one more grounded in communication.*

Key Words: Connectionism; visual search; feature extraction.

1. Introduction.

Over the hundred years since the publication of James' Psychology (1891), neuroscientists have attempted to define the fundamental features of the brain and its information processing capabilities in terms of mean firing rates at 'points in the brain cortex' (neurons) and computations. After Hubel and Wiesel (1962), the function of the neuron as a specialist feature detector was treated as established doctrine. From this followed the functional specialisation paradigm, mapping different areas of the brain to specific cognitive function, reincarnating an era of modern phrenology.

Connectionism mapped well onto the above assumptions. At its heart the MCP model (McCulloch-Pitts, 1943), defines the neuron as a simple computational device mapping real valued (continuous or discrete) inputs to a real valued output. Using a suitable learning algorithm, a multi-layer network of such cells (e.g. a Multi-Layer Perceptron, MLP) can learn an approximate realisation of any continuous mapping (Kolmogorov, 1942). Such connectionist networks are effectively analysed within the framework of Euclidean space:

1. The free parameters of the network map to a point in $\Re^n$.
2. Knowledge is represented by vectors (*or subsets of an appropriate Euclidean space*).
3. Learning is the process of traversing a specific spatial trajectory over $\Re^n$, with consecutive steps being determined by the particular training data and learning algorithm used.
4. A MLP and its learning rule combine to effectively perform non-linear regression on the input training data and the target output vectors.

Connectionist networks have attracted a lot of interest. They have been successfully applied to various engineering problems (Tarashenko, 1998), and as metaphors of concepts drawn from neuroscience, have also been offered as models of both high (Rumelhart & McClelland, 1986) and low level (cf. Ahmad, 1991; Van de Laar, 1997) cognition. However current connectionist models of high level cognition have also been strongly criticised (cf.

Dinsmore, 1990) and the situation at the domain of low level modelling is little better (Abbot, 1990). This is especially true in the light of the most recent advances in understanding the neuroanatomy and biophysics of single neuron processing.

Two examples from recent neuroanatomical studies serve to highlight fundamental problems with current connectionist models of brain architecture which are grounded upon centralised control of distinct functional units, via inhibition and negative feedback. Firstly the assumption that the thalamus in some way supervises the operation of visual cortical areas has not been born out by empirical investigation of the primary cortex of cats and primates. Here observations have indicated that excitatory input from subcortical structures in the thalamus (eg. LGN) constitute no more than 20% of total number of excitatory synapses made onto the layer IV spiny stellate cells of the cortex (Douglas, 1997).

Further, many connectionist models make extensive use of some form of inhibition (Sirosh, 1996). In particular, a form of lateral inhibition is often used as a normalising mechanism and to prevent signal saturation, (Douglas, 1997), and to define topological structure (Kohonen, 1995). However recent neurological studies reported in Douglas and Martin, (1991), clearly indicate that 85% of synapses in the cortex are excitatory, and moreover about 85% of the synapses made by excitatory neurons are onto excitatory neurons.

Yet further evidence against the widespread use of inhibitory mechanisms in the brain is the recent discovery that the principal mediators of synaptic inhibition, GABAA receptors, when intensely activated, excite rather than inhibit neuron firing (Staley et al., 1995). Thus, if connectionist models of cognition are to maintain a claim to biological plausibility, then the mechanisms at the heart of the paradigm need substantial revision.

In (Nasuto, Dautenhahn & Bishop, 1998b; 1999) we discuss in more detail the implications of recent neurological findings and propose a reinterpretation of the functionality of the single neuron away from its historical computational regression metaphor, to one grounded in terms of communication and search. In this article we discuss this new Communicating Neuron metaphor, in which each neuron processes more complex (bi-variate) information than that implicit in the standard MCP model, with respect to the novel hybrid n-tuple/Stochastic Diffusion Network - first presented to the n-tuple/weightless network community at WnnW '93, (Bishop & Mitchell, 1993).

In the next section we introduce the generic Stochastic Diffusion Process, **SDP**, (Bishop, 1989), on which our Communication Neuron metaphor is based, before defining a Hybrid Stochastic Diffusion Network, **HSDN**, (Bishop, 1990), as a specific instantiation of this metaphor. We subsequently contrast issues of knowledge representation inherent in traditional uni-variate networks with that found in the HSDN bi-variate network. The final section contains our discussion and concludes that a network grounded upon the Communicating Neuron metaphor, such as the HSDN, potentially overcomes many of problems, as outlined by Fodor and Pylyshyn (1988), that are traditionally associated with the established connectionist Artificial Intelligence paradigm.

2. Stochastic Diffusion Processes.

In 1976 Newell and Simon claimed that the necessary and sufficient means for 'general intelligent action' is a physical symbol system and heuristic search, however this claim has

been repeatedly questioned for defining a closed system without means to ground its representations (*cf. Law after Lakoff*). It will be shown that Stochastic Diffusion Processes grounded upon a quasi-symbol system and 'best-fit' search, overcome these problems by merging Newell and Simon's insight onto an extended, open, connectionist framework.

Although a SDP describes its target by a set of *micro-features*, the particular format of this decomposition is unimportant. Hence the process is quasi-symbolic, for although it may be possible to label component micro-features symbolically – as a nose; a letter 'A'; a cardinal '1' etc. – it is not necessary that this should be the case. For example a 3D object, X, could be described using Binford's Generalised Cones (1971) or by its 2D projection. In the former case the component micro-features would be the parameters defining the set of cones required to describe X, whereas in the latter case they would be the ordered set of pixel colour/intensity values or n-tuple combinations thereof (Aleksander & Stonham, 1979).

It has previously been demonstrated that one emergent property of an evolving Stochastic Diffusion Process is the dynamic solution to dynamic best-fit search tasks (Nasuto & Bishop, 1999). This is most easily seen by analogy.

2.1 Ant Search

Consider the following example of hypothetical ant-like creatures searching for a good nutrient source in a dynamic environment. Each ant seeks to locate some food and return it to the nest. The colony as a whole seeks to maximise the rate of return of food for the minimum expenditure of energy.

With no a-priori information on the likely location of food, each searching ant will leave the nest and perform a random walk around the local terrain. If in the course of its explorations an ant finds some food, it returns to the nest a positive ant otherwise it is labelled negative.

On its return the positive ant could broadcast the location of the new food source throughout the colony, causing all other ants to investigate it. In computing terminology, this is isomorphic to a 'greedy' solution to the search problem. However a small or transient food source would result in the colony committing all its resources for poor return.

But consider what would occur if, on its return to the nest, each positive ant simply tells the first searching ant it meets the location of its find. If the food source is good (i.e. it is temporally stable and bountiful), over a relatively small period of time the nest will automatically allocate more and more of its resources (ants) to exploiting it. Whereas if the resource is poor, any positive ants that are initially attracted to it will sooner or later not find anything and revert back to being searching ants. Conversely since unsuccessful ants which meet on their return to the nest do not exchange resource-location information, they simply re-commence their random search.

The above is a simple hypothetical example of a self-organising Stochastic Diffusion Process which will efficiently converge (allocate ants/resources) to the best point (food) in the search space.

2.2 Stochastic Diffusion Search

Functionally, ant search is equivalent to Stochastic Diffusion Search, **SDS**, (Bishop, 1989). In SDS a group of independent agents process information from the search space in order to find the best-fit to a specified target pattern. Each agent searches for a micro-feature of the target and once found competes to attract other agents to evaluate this position. In this way the SDS explores the whole search space. Due to the emergent co-operation of agents pointing to the same solution, interesting areas in the search space (those that share many micro-features with the target) are more thoroughly exploited than background areas.

Agents are divided into two charges: positive and negative. A positive agent has successfully found a micro-feature from the target in the search space; a negative agent has not. Thus the charge label identifies agents more likely to point to an instantiation of the target than to background. Negative agents utilise this activity information when deciding whether to communicate with a randomly selected agent in a subsequent phase of processing. Communication only occurs if the selected agent is positive and results in the flow of 'where' information from the positive agent to the negative. Conversely, if the selected agent is also negative, then there is no information flow between agents; instead, a new random 'where' position is adopted.

In this way positive agents attract more resources to examining promising regions of the search space. Dependent on the quality of the target instantiation (i.e. how many micro-features it has in common with the target), positive agents will spend different periods of time examining particular positions in the search space. On average, those agents pointing to the best instantiation of the target will spend longer examining the same location than agents pointing to a background or poor quality positions (i.e. those areas of the search space with few micro-features in common with the target). Hence such agents will have more possibilities to communicate their search position to others, and in this way a population of agents will converge onto the current best instantiation of the target in the search space. The continuing exploration of the search space by negative agents ensures that the process will eventually discover, and converge to, the best-possible fit of the target. This convergence to the global-best solution is demonstrated in (Nasuto & Bishop, 1999) and its time complexity discussed in Nasuto et al, (1998a).

2.3 The Stochastic Diffusion Search Algorithm

A Stochastic Diffusion Search involves several stages and can be succinctly summarised in algorithmic form as follows:

```
INITIALISE (agents);
WHILE NOT TERMINATE (agents) DO
  TEST (agents);
  DIFFUSE (agents);
END;
```

1: **INITIALISE**: This assigns agents to examine specific (*in the absence of a-priori information, random*) locations in the search space (*mappings*).

2: **TERMINATE**: This evaluates the activity of the current population of agents to establish if the chosen termination criteria has been met. One such termination criteria is the *Strong Halting Criteria*, (Nasuto & Bishop, 1999).

3: **TEST**: This procedure simply evaluates a randomly selected micro-feature of the model with the corresponding micro-feature from the search space at the position defined by the agents *mapping*. If the TEST establishes that the micro-feature is present at this position then the agent is defined as being in a positive state, otherwise it is defined as negative.

4: **DIFFUSE**: This stage re-assigns agent mappings by a stochastic diffusion process. Each negative agent randomly selects another agent from the population. If the selected agent is positive then its mapping is copied across to the negative agent. If the the selected agent also was negative, then a new mapping is randomly selected.

3. The Hybrid Stochastic Diffusion Neural Network

An example of a problem involving pattern invariance in which the Hybrid Stochastic Diffusion Network has been successfully applied is that of locating eye positions within grey scale images of human faces (Bishop & Torr, 1992). An adult human has little difficulty in locating eye positions from a facial image, even though these positions could be anywhere within the image field. This is a difficult task in machine vision since facial features do not generally comprise of hard edges and lines. The Hybrid Stochastic Diffusion Network, a combination of a Stochastic Diffusion Network and an n-tuple network, was shown to accurately locate all eye features on which it has been trained, and also achieve over 60% success in the general problem of locating eye features on which it has not been explicitly trained.

The Hybrid Stochastic Network comprises of two systems (see figure 1), both of which can be easily implemented in dedicated hardware.

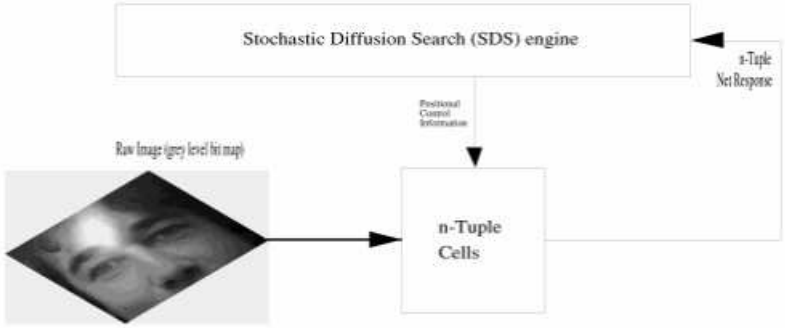

Figure 1. The Hybrid Stochastic Diffusion Network.

3.1 n-Tuple Micro-feature classification

The n-Tuple method of pattern classification, first proposed by Bledsoe & Browning (1959) and implemented in hardware as the WISARD system (Aleksander et al, 1984), has significant advantages over traditional methods of simple pattern classification (Aleksander &

Stonham, 1979). The data to be classified are stored as an array of boolean values. This array is then sampled, according to some predefined schema (usually either randomly or sequentially), with [n] such samples forming one n-Tuple. The data can be either exclusively sampled, such that each value is used only once, or oversampled by a factor of (m), where each value is used (m) times.

To learn a pattern, the (n) boolean values from each sample are used to form an n-bit address into an array (or RAM), where there is one such RAM per n-Tuple sample, and a boolean TRUE value is stored at this location. To classify an image, the data are sampled as before. However if, instead of storing a TRUE value at each defined location, a count of the TRUE values stored at these addresses is maintained, the value of the count gives the system response to an unknown pattern. After training with a set of patterns, such a network is able to classify patterns not in training set. If these new patterns are "similar" to those in the training set, then the system has generalised successfully.

3.2 Searching for the target

The simplest form of feature location would be to scan the classifier sequentially over the search space, however a video image of 256*256 pixels is a large space to traverse, requiring 65356 sequential comparisons (iterations). Even if a dedicated hardware system is used, working at 50 frames a second, it would still take 1300 seconds to scan the entire image. However, by using a Stochastic Diffusion Process to control placement of n-tuple RAMs over the image (as used in the HSDN), a typical analysis only requires around 30 iterations (Bishop & Torr, 1992).

Figure 1 shows a schematic block diagram of the Hybrid Stochastic Network. The input video image is first digitised to a resolution of 256*256 pixels at an accuracy of 8 bits. A selection of points from this grey level data is processed by an array of Minchinton cells (Bishop, Mitchell & Minchinton, 1991), the selection of points being controlled by the Stochastic Diffusion Search engine mapping cells. Each n-tuple RAM thus receives binary data from a given mapping into the search space defined by the mapping cells. The response of the RAM to this data comprises the TEST (or evaluation) phase of the Stochastic Diffusion Search. The DIFFUSION phase then reallocates mappings contingent upon their success as described in section 2.3.

3.3 The equivalence of the HSDN and SDS

The equivalence of the Hybrid Stochastic Diffusion Network and a simple Stochastic Diffusion Search follows from the equivalence of the neuron model to the agent of the SDS. The only difference between the two is that evaluation of a micro-feature is performed by an n-tuple RAM neuron in place of a simple equivalence test.

Due to the equivalence between SDS and the HSDN, both its properties of global convergence and time complexity are preserved. Hence the network will find the best possible solution with probability one (Nasuto & Bishop, 1999) in a time, dependent on the size of the network, either independent of, or growing linearly with, the search space size (Nasuto et al. 1998a).

4. Knowledge Representation in the Hybrid Stochastic Diffusion Network.

Functionally, the information processing performed by a HSDN is most easily illustrated in the visual domain. Here it seeks to label and locate the 'best-fit' of an external world stimulus (*the target*), defined by the specific pattern of micro-features learnt by the n-tuple network (*the memory*), detected at the retina (*the search space*). This requires the micro-features defining the target to be identified (tagged) both by their n-tuple value and their relative position on the retina. Each n-tuple labels 'what' each micro-feature is and the corresponding locate-tag defines 'where' on the retina it is. Thus the HSDN processes knowledge as 'Tokens' and not as 'Types'. This contrasts information processing in most associative networks, where knowledge is represented as simple types (defined by a vector in Euclidean space). It has been suggested (Van de Velde, 1997) that 'Type' representation schemes are a fundamental cause of many of the problems encountered when modelling symbolic processes by associative networks.

However there are some connectionist models of higher cognitive functions that do process knowledge as tokens (e.g. McClelland's 1981, Interactive Activation & Competition network). In such models the tokens are materially defined by nodes, each of which represents a particular feature of the environment. But these networks are limited to representation of arity zero predicates and Dinsmore (1990), following Fodor & Pylyshyn (1988), suggests this is too strong a restriction for representing the complexity of the real world.

5. Discussion and conclusions.

The neural network outlined in this paper demonstrates that it is possible to construct useful connectionist models without standard inhibitory mechanisms. It has been shown that the Hybrid Stochastic Diffusion Network performs Stochastic Diffusion Search, to solve the best-fit matching problem. This functionality emerges from its ability to self-organise in response to incoming stimuli. The network effectively uses a tagged dynamic assembly encoding for the target.

Knowledge representation within the HSDN is fundamentally different from that used in classical nets. In our model the activity of a single neuron does not possess a semantical interpretation. However, an assembly of such neurons, locked to a particular location in the search space, does acquire a semantical interpretation, as it supports a tagged-tokenised internal representation of the object it attends to. Thus the neurons do not constitute an internal representation in themselves, as the assembly is dynamically fluctuating, however their collective pattern of activity does. Hence, although an assembly supporting a particular representation in different time instants can differ considerably in the number and identity of its constituent neurons, the collective representation so formed is continuous and stable over time.

Further, the allocation of neurons to the best-fit of the target is analogous to an attention mechanism switching computational resources between target objects. This suggests a possible solution to the classical parallel/serial divide problem of attention theory (Treisman, 1988). In our model both types of attention coexist. Single neurons process information from the search space in parallel and serial attention emerges when an assembly of neurons that have been focused on one area of the retina migrates to another.

References.

Abbot, L., (1990), *Learning in Neural Network Memories*, Network: Comp. Neural Sys., **1**, pp. 105-122.

Ahmad, S., (1991), *VISIT: An efficient computational model of human visual attention*, PhD Thesis, University of Illinois.

Aleksander, I. & Stonham, J., (1979), *A Guide to Pattern Recognition Using Random Access Memories*, Computers & Digital Techniques **2**, pp. 29-40.

Aleksander, I., Thomas, W.V. & Bowden, P.A., (1984), *WISARD: a Radical Step Forward in Image Recognition*, Sensor Review, pp: 120-124.

Binford, T.O., (1971), *Visual Perception by Computer*, Proc. IEEE Systems & Control, Miami.

Bishop, J.M., (1989), *Stochastic Searching Networks*, Proc. 1st IEE Int. Conf. On Artificial Neural Networks, IEE Press: London.

Bishop, J.M. & Torr, P.H., (1992), *The Stochastic Search Network*, in Lingard, D.J. & Nightingale, C., (eds), Neural Networks for Vision, Speech and Natural Language, Chapman Hall.

Bishop, J.M., Minchinton, P.R., Mitchell, R.J., (1981), *Real Time Invariant Grey Level Image Processing using Digital Neural Networks*, Proc. IMechE Conf., EuroDirect '91 - Computers in the Engineering Industry, pp. 187-189, Birmingham, UK.

Bishop, J.M. & Mitchell, R.J., (1993), *Optinisation of Tuple size in a Hybrid Stochastic Search*, Proc. WnnW '93, pp. 139-145, York 1993.

Bledsoe, W.W. & Browning, I., (1959), Pattern Recognition and Reading by Machine, Proc. Eastern Joint Computer Conference

Dinsmore, J., (1990), *Symbolism and Connectionism: a difficult marriage*, Technical Report TR 90-8, Southern University of Carbondale, USA.

Douglas, R., (1997), Private communication.

Douglas, R, & Martin, K., (1991), *Opening the grey box*, TINS, **14:7**.

Fodor, J.A. & Pylyshyn, Z.W., (1988), *Connectionism and cognitive architecture: A critical analysis*, Cognition **28**, pp. 3-72.

Hubel, D.H. & Wiesel, T.N., (1962), *Receptive fields, binocular interaction and functional architecture in the cat's visual cortex*, Journal of Physiology, **160**, pp. 106-154.

James, W., (1891), *Psychology (briefer course)*, Holt: New York.

Kohonen, T., (1988), *Self Organising and Associative Memory*, Springer Verlag: New York..

Kolmogorov, A.N., (1942), *Interpolation and Extrapolation of Stationary Random Sequences*, Trans. RAND corporation 1962, Santa Monica.

McClelland, J.L., (1981), *Retrieving General Knowledge from Stored Knowledge of Specifics*, Proc. 3rd Conf. Cognitive Science Society, Berkeley.

McCulloch, W.S. & Pitts, W., (1943), *A logical calculus immanent in nervous activity*, Bulletin of Mathematical Biophysics, **5**, pp. 115-133.

Nasuto, S.J., & Bishop, J.M., (1999), *Convergence of the Stochastic Diffusion Search*, Parallel Algorithms, (in press).

Nasuto, S.J., Bishop, J.M. & Lauria, S., (1998a), *Time Complexity Analysis of the Stochastic Diffusion Search*, Neural Computation '98, Vienna.

Nasuto, S.J., Dautenhahn, K. & Bishop, J.M., (1998), *Communication as an emergent metaphor for neuronal operation*, in: Nehaniv, C., (ed.), Plenary Working Papers in Computation for Metaphors, Analogy and Agents, pp. 65-70, Aizu, Japan.

Nasuto, S.J., Dautenhahn, K. & Bishop, J.M., (1999)., *Communication as an emergent metaphor for neuronal operation*, Lecture Notes in Computer Science **1562**, Springer-Verlag, (in press).

Newell, A. & Simon, H.A., (1976), *Computer Science as Empirical Enquiry: Symbols and Search*, Communications of the ACM **19**.

Rumelhart, D.E. & McClelland, J.L., (1986), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, MIT Press: Cambridge MA.

Sirosh, J., et al. (eds.), (1996), *Lateral Interactions in the Cortex: Structure and Function*, Hypertext book: www.cs.utexas.edu/users/nn/wcb-pubs/htmlbook96: ISBN 0-9647060-0-8.

Staley, K., et al., (1995), *Ionic Mechanisms of Neuronal Excitation by Inhibitory GABAA Receptors*, Science, **269**, pp. 977-981.

Tarassenko, L., (1998), *A Guide to Neural Computing Applications*, Arnold: London.

Treisman, A., (1988), *Features and Objects: The fourteenth Bartlett memorial lecture*, The Quarterly Journal of Experimental Psychology, **40A(2)**, pp. 201-237.

Van de Laar, P., et al., (1997), *Task-dependent Learning of Attention*, Neural Networks, 10:1, pp. 1-21.

Van de Velde, F., (1997), *On the use of computation in modelling behaviour*, Network: Comput. Neural Syst., **8**, pp. R1-R32.