Attention through Self-Synchronisation in the Spiking Neuron Stochastic Diffusion Network

K. De Meyer * J. M. Bishop
S. J. Nasuto
Department of Cybernetics
University of Reading, Whiteknights,
READING, Berkshire
United Kingdom
http://www.rdg.ac.uk/~krm/sdp.htm

ASSC4 – June 2000

Abstract

The paper discusses ensemble behaviour in the Spiking Neuron Stochastic Diffusion Network, SNSDN, a novel network exploring biologically plausible information processing based on higher order temporal coding. SNSDN was proposed as an alternative solution to the binding problem [1]. SNSDN operation resembles Stochastic Diffusion Search, SDS, a non-deterministic search algorithm able to rapidly locate the best instantiation of a target pattern within a noisy search space ([3], [5]). In SNSDN, relevant information is encoded in the length of interspike intervals. Although every neuron operates in its own time, 'attention' to a pattern in the search space results in self-synchronised activity of a large population of neurons. When multiple patterns are present in the search space, 'switching of attention' results in a change of the synchronous activity. The qualitative effect of attention on the synchronicity of spiking behaviour in both time and frequency domain will be discussed.

1 Stochastic Diffusion Search

Stochastic Diffusion Search ([3], [4], [5], [6], [7], [8]) is a parallel, non-deterministic pattern matching algorithm. It is capable of rapidly locating a specified pattern – or its best instantiation – in a noisy search space. Its operation is most easily explained by analogy.

1.1 Ant Search Analogy

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 or the minimum expenditure of energy.

*demeyer@fan.net

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 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 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 unsuccesful 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.

1.2 The Algorithm

In SDS a group of independent **Matching Agents** process information from the **Search Space** in order to find the best fit to a specified target pattern (the **Model**). The Model and the Search Space consist of sets of **micro-features**. For simple, one-dimensional string matching, Model and Search Space are strings, the individual characters of the strings being the micro-features. The Matching Agents try to locate the position of the Model string in the Search Space, by comparing a character of the Model with the corresponding character in the Search Space, defined by the Matching Agent's location (the 'where'-value) plus the offset of the character within the Model.

A Stochastic Diffusion Search involves several stages:

1. **Initialise Matching Agents:** Agents are assigned random (*in the absence of a-priori information*) locations in the search space, and their state is initialised to *inactive*.

2. **Test:** This procedure evaluates a randomly selected micro-feature of the Model at a given location in the Search Space. If the test succeeds, then the state is set to *active*; else, the state is set to *inactive*.

3. **Diffuse:** Every *inactive* Matching Agent chooses a Matching Agent at random. If the other Agent is *active*, then its location in the Search Space is copied. If the other Agent is *inactive*, a new, random location in the search space is chosen.

4. **Termination:** If the termination criteria are fulfilled (if the population of Agents has converged onto a location in the Search Space), terminate. Else go back to Step 2.

2 Text Search with the Spiking Neuron Stochastic Diffusion Network

Recent neuro-biological experiments indicate that the long-standing assumption of encoding of information in the mean firing rate of neurons cannot be the only mechanism of information encoding in the brain. Various new encoding schemes, based on the timing of spikes are being explored ([9], [10], [11], [12]).

The Spiking Neuron Stochastic Diffusion Network belongs to this latter category. It performs SDS-like pattern matching, using *communicating* neurons rather than *computing* ([2]). Its important features will be outlined here:

Temporal Encoding: Information is encoded in the time between spikes, the *Inter-Spike Interval* (ISI).

Bivariate Information: Memory Neurons (the Model) and **Retina Neurons** (the Search Space) communicate 2 kinds of information to the **Matching Neurons**: the *location* of a character in the Model or SS ('where'-value) and the *ASCII-value* of the character ('what'-value).

Matching Neurons perform SDS-like test-diffuse cycles. They communicate locations of potential solutions to other Matching Neurons in Inter-Spike Intervals.

No Forced Synchronisation: Matching Neurons operate in their own time. No global synchronisation method exists.

Randomisation: At various stages in the SDS-algorithm, micro-features or new locations in the SS have to be chosen at random. In SNSDN, this is accomplished by choosing first-incoming spikes. A good randomisation is guaranteed by *random refractory periods* at certain moments in the operation of the different kinds of neurons, and by *random axon transmission delays*.

3 Results

'Attention' for a target pattern in the Search Space results in a self-synchronised population of Matching Neurons. 'Shifting of attention' from one pattern to a second results in a change of the synchronous activity. Results are presented for an experiment with 100 Matching Neurons, a Model of 6 characters and a Search Space of 90 characters. 2 patterns with 5/6 overlap with the Model are present in the Search Space. This means each pattern has a 16.6% chance of failing a test-phase. Matching Neuron Activity and Spike Count are reported every interval of 10 timesteps; the algorithm ran for 50000 timesteps.

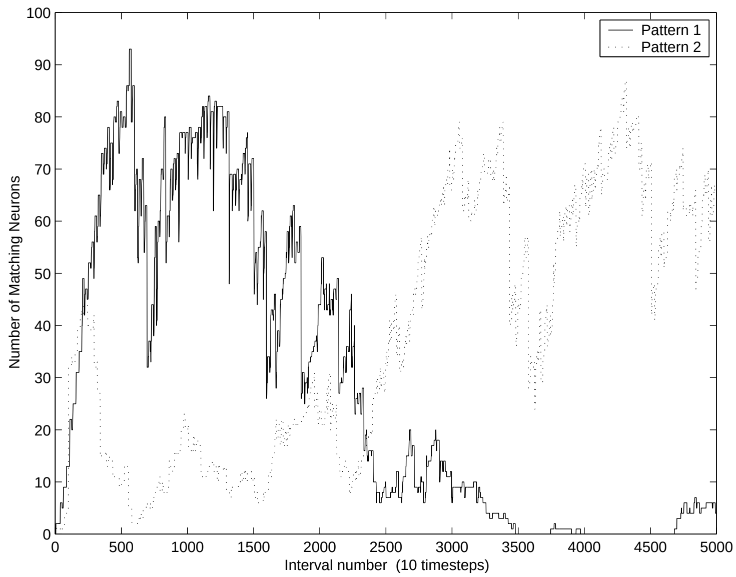

Figure 1: Attention of Matching Neurons looking for the Model in the Search Space. ‘Switching of Attention’ from one pattern to the second can be seen halfway the experiment.

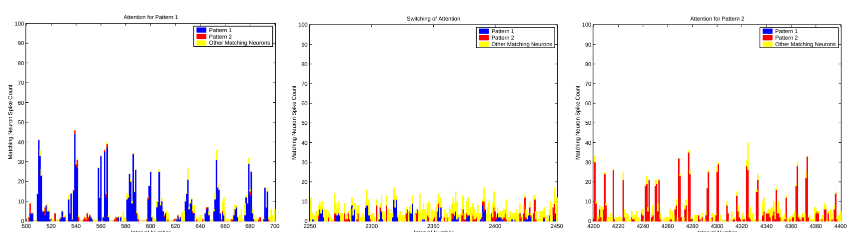

Figure 2: Spikes of Matching Neurons during several stages of attention, each for 200 intervals (2000 timesteps)

.

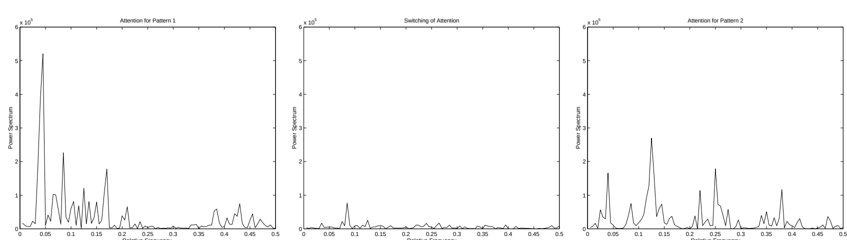

Figure 3: Power Spectra of Matching Neuron Spikes during the respective stages of attention

.

References

[1] Bishop, J.M., Nasuto, S.J., (1998), "Neural Stochastic Diffusion Search Network - a theoretical solution to the binding proble", Proc. ASSC2, Bremem.

[2] Nasuto, S.J., Dautenhahn, K., Bishop, J.M., 1999, "Communication as an Emergent Methaphor for Neuronal Operation", Lecture Notes in Artificial Intelligence, 1562:365-380, Springer.

[3] Bishop, J.M., (1989),"Stochastic Searching Network", Proc. 1st IEE Conf. ANNs, London.

[4] J.M. Bishop, P. Torr: 'The Stochastic Search Network'. In R. Linggard, D.J. Myers, C. Nightingale (eds.), Neural Networks for Images, Speech and Natural Language. New York, Chapman & Hall, 1992.

[5] Nasuto, S.J., (1999),"Resource Allocation Analysis of the Stochastic Diffusion Search", PhD Thesis, University of Reading.

[6] S.J. Nasuto, J.M. Bishop, S. Lauria: 'Time Complexity of Stochastic Diffusion Search', Neural Computation '98, Vienna, Austria (1998).

[7] S.J. Nasuto, J.M. Bishop: 'Convergence Analysis of Stochastic Diffusion Search', Journal of Parallel Algorithms and Applications 14, pp 89-107, 1999.

[8] P.D. Beattie, J.M. Bishop: 'Self-Localisation in the 'Senario' Autonomous Wheelchair', Journal of Intelligent and Robotic Systems 22, pp 255-267, Kluwer Academic Publishers (1998).

[9] Rieke, F., Warland, D., de Ruyter van Steveninck, R.R., Bialek, W., 1997. "Spikes - Exploring the Neural Code". MIT Press, Cambridge MA.

[10] Maass., W., Bishop, C.M., 1999. "Pulsed Neural Networks". MIT Press, Cambridge, MA.

[11] deCharms, R.C., Zador, A., "Neural Representation and the Cortical Code", Annu. Rev. Neurosci. 2000, 23:613-647.

[12] Reich, D.S., Mechler, F., Purpura, K.P., Victor, J.D., "Interspike Intervals, Receptive Fields and Information Encoding in Primary Visual Cortex", Journal of Neuroscience, 2000, 20(5):1964-1974.