

Communicating neurons: a connectionist spiking neuron implementation of stochastic diffusion search

S.J. Nasuto

School of Systems Engineering, The University of Reading, Whiteknights, Reading, UK.

e-mail: s.j.nasuto@reading.ac.uk

J.M. Bishop *

Department of Computing, Goldsmiths College, New Cross, London, UK.

e-mail: m.bishop@gold.ac.uk

K. De Meyer

Division of Engineering, Kings College, London

Keywords: Spiking Neural Networks, Selective Attention; Stochastic Diffusion Search, Swarm Intelligence, Global Search

Abstract

An information-processing paradigm in the brain is proposed, instantiated in an artificial neural network using biologically motivated temporal encoding. The network will locate within the external world stimulus, the target memory, defined by a specific pattern of micro-features.

The proposed network is robust and efficient. Akin in operation to the Swarm Intelligence paradigm, Stochastic Diffusion Search, it will find the best-fit to the memory with linear time complexity. Information multiplexing enables neurons to process knowledge as ‘tokens’ rather than ‘types’. The network illustrates possible emergence of cognitive processing from low level interactions such as memory retrieval based on partial matching.

1 Introduction

One of the important roles of metaphor in science is to facilitate understanding of complex phenomena. Metaphors should describe phenomena in an intuitively understandable way that captures their essential features. We argue that a description of single neurons as computational devices does not capture the information processing complexity of

real neurons and argue that describing them in terms of communication could provide a better alternative metaphor. These claims are supported by recent discoveries showing complex neuronal behaviour and by fundamental limitations of established connectionist cognitive models. We suggest that real neurons operate on richer information than provided by a single real number and therefore their operation cannot be adequately described in standard Euclidean setting. Recent findings in neurobiology suggest that, instead of modelling the neuron as a logical or numerical function, it could be described as a communication device.

The prevailing view in neuroscience is that neurons are simple computational devices, summing up their inputs and calculating a non-linear output function. Information is encoded in the mean firing rate of neurons which exhibit narrow specialisation - they are devoted to processing a particular type of input information. Further, richly interconnected networks of such neurons learn via adjusting inter-connection weights. In the literature there exist numerous examples of learning rules and architectures, more or less inspired by varying degrees of biological plausibility. Almost from the very beginning of connectionism, researchers were fascinated by computational capabilities of such devices [42, 56]. The revival of connectionism in the mid-eighties featured increased interest

*corresponding author

in analysing the properties of such networks [55], as well as in applying them to numerous practical problems [26]. At the same time the same devices were proposed as models of cognition capable of explaining both higher level mental processes [58] and low level information processing in the brain [22]. However, these promises were based on the assumption that the computational model captures all the important characteristics of real biological neurons with respect to information processing. We will indicate in this article that very recent advances in neuroscience appear to invalidate this assumption. Neurons are much more complex than was originally thought and thus networks of oversimplified model neurons are orders of magnitude below the complexity of real neuronal systems. From this it follows that current neural network ‘technological solutions’ capture only superficial properties of biological networks and further, that such networks may be incapable of providing a satisfactory explanation of our mental abilities. We propose to complement the description of a single neuron as a computational device by an alternative, more ‘natural’ metaphor :- we hypothesise that a neuron can be better and more naturally described in terms of communication rather than purely computation. We hope that shifting the paradigm will result in escaping from the local minimum caused by treating neurons and their networks merely as computational devices. This should allow us to build better models of the brain’s functionality and to build devices that reflect more accurately its characteristics. We will present a simple connectionist model, NEural STOchastic diffusion search netwORk (NESTOR), fitting well in this new paradigm and will show that its properties make it interesting from both the technological and brain modelling perspectives. Selman et al. posed some challenge problems for Artificial Intelligence [61]. In particular Rodney Brooks suggested revising the conventional McCulloch Pitts neuron model and instead investigating the potential implications (with respect to our understanding of biological learning) of new neuron models based on recent biological data. Further, Selman claimed that the supremacy of standard heuristic, domain specific search methods of Artificial Intelligence need to be revised and suggested that recent investigation of fast general purpose search procedures has opened a promising alternative avenue. Furthermore, in

the same paper Horvitz posed the development of richer models of attention as an important problem, as all cognitive tasks “... require costly resources” and “controlling the allocation of computational resources can be a critical issue in maximising the value of a situated system’s behaviour.” We claim that the new network presented herein addresses all three challenges posed in the above review paper [61], as it is isomorphic in operation to Stochastic Diffusion Search (SDS), a fast, generic probabilistic search procedure which automatically allocates information processing resources to search tasks.

2 Computational metaphor

The emergence of connectionism is based on the belief that neurons can be treated as simple computational devices [42]. Further, the assumption that information is encoded as mean firing rate of neurons was a base assumption of all the sciences related to brain modelling. The initial boolean McCulloch-Pitts model neuron was quickly extended to allow for analogue computations. The most commonly used framework for connectionist information representation and processing is a subspace of a Euclidean space. Learning in this framework is equivalent to extracting an appropriate mapping from the sets of existing data. Most learning algorithms perform computations which adjust neuron interconnection weights according to some rule, adjustment in a given time step being a function of a training example. Weight updates are successively aggregated until the network reaches an equilibrium in which no adjustments are made (or alternatively stopping before the equilibrium, if designed to avoid overfitting). In any case knowledge about the whole training set is stored in final weights. This means that the network does not possess any explicit internal representation of the (potentially complex) relationships between training examples other than that which implicitly exists as a distribution of weight values. We do not consider representations of arity zero predicates, (e.g. those present in NETtalk [59]), as sufficient for representation of complex relationships. These limitations result in poor internal knowledge representation making it difficult to interpret and analyse the network in terms of causal relationships. In particular it is difficult to imagine how such a system could

develop symbolic representation and logical inference (cf. the symbolic/connectionist divide). Such deficiencies in the representation of complex knowledge by neural networks have long been recognised [19, 7, 54]. The way in which data are processed by a single model neuron is partially responsible for these difficulties. The algebraic operations that it performs on input vectors are perfectly admissible in Euclidean space but do not necessarily make sense in terms of the data represented by these vectors. Weighted sums of quantities, averages etc., may be undefined for objects and relations of the real world, which are nevertheless represented and learned by structures and mechanisms relying heavily on such operations. This is connected with a more fundamental problem missed by the connectionist community - the world (and relationships between objects in it) is fundamentally non-linear.

Classical neural networks are capable of discovering non-linear, continuous mappings between objects or events but nevertheless they are restricted by operating on representations embedded in linear, continuous structures (Euclidean space is by definition a finite-dimensional linear vector space equipped with standard metric). Of course it is possible in principle that knowledge from some domain can be represented in terms of Euclidean space. Nevertheless it seems that only in extremely simple or artificial problems the appropriate space will be of small dimensionality. In real life problems spaces of very high dimensionality are more likely to be expected. Moreover, even if embedded in a Euclidean space, the actual set representing a particular domain need not be a linear subspace, or be a connected subset of it. Yet these are among the topological properties required for the correct operation of classical neural nets. There are no general methods of coping with such situations in connectionism.

Methods that appear to be of some use in such cases seem to be freezing some weights (or restricting their range) or using a 'mixtures of experts' or 'gated networks' [31]. However, there is no principled way describing how to perform the former. Mixture of experts models appear to be a better solution, as single experts could in principle explore different regions of a high dimensional space thus their proper co-operation could result in satisfactory behaviour. However, such architectures need to be individually tailored to particular prob-

lems. Undoubtedly there is some degree of modularity in the brain, however, it is not clear that the brain's operation is based solely on a rigid modularity principle. In fact we will argue in the next section that biological evidence seems to suggest that this view is at least incomplete and needs revision. We feel that many of the difficulties outlined above follow from the underlying interpretation of neuron functioning in computational terms, which results in entirely numerical manipulations of knowledge by neural networks. This seems a too restrictive scheme. Even in computational neuroscience, existing models of neurons describe them as geometric points even though neglecting the geometric properties of neurons (treating dendrites and axons as merely passive transmission cables) makes such models very abstract and may strip them of some information processing properties. In most technical applications of neural networks the abstraction is even higher - axonic and dendritic arborisations are completely neglected - hence they cannot in principle model the complex information processing taking place in these arbors [62]. We think that brain functioning is best described in terms of non-linear dynamics but this means that processing of information is equivalent to some form of temporal evolution of activity. The latter however may depend crucially on geometric properties of neurons as these properties obviously influence neuron activities and thus whole networks. Friston [21] stressed this point on a systemic level when he pointed out to the importance of appropriate connections between and within regions - but this is exactly the geometric (or topological) property which affects the dynamics of the whole system. Qualitatively the same reasoning is valid for single neurons. Undoubtedly, model neurons which do not take into account geometrical effects perform some processing, but it is not clear what this processing has to do with the dynamics of real neurons. It follows that networks of such neurons perform their operations in some abstract time not related to the real time of biological networks (We are not even sure if time is an appropriate notion in this context, in case of feedforward nets 'algorithmic steps' would be probably more appropriate).

This concerns not only classical feedforward nets which are closest to classical algorithmic processing but also many other networks with more interesting dynamical behaviour (e.g. Hopfield or other

attractor networks). Of course one can resort to compartmental models but then it is apparent that the description of single neurons becomes so complex that we have to use numerical methods to determine their behaviour. If we want to perform any form of analytical investigation then we are bound to simpler models.

Relationships between real life objects or events are often far more complex for Euclidean spaces and smooth mappings between them to be the most appropriate representations. In reality it is usually the case that objects are comparable only to some objects in the world, but not to all. In other words one cannot equip them with a 'natural' ordering relation. Representing objects in a Euclidean space imposes a serious restriction, because vectors can be compared to each other by means of metrics; data can be in this case ordered and compared in spite of any real life constraints. Moreover, variables are often intrinsically discrete or qualitative in nature and in this case again Euclidean space does not seem to be a particularly good choice. Networks implement parametrised mappings and they operate in a way implicitly based on the Euclidean space representation assumption - they extract information contained in distances and use it for updates of weight vectors. In other words, distances contained in data are translated into distances of consecutive weight vectors. This would be fine if the external world could be described in terms of Euclidean space however it would be a problem if we need to choose a new definition of distance each time a new piece of information arrives.

Potentially new information can give a new context to previously learnt information, with the result that concepts which previously seemed to be not related now become close. Perhaps this means that our world model should be dynamic - changing each time we change the definition of a distance? However, weight space remains constant - with Euclidean distance and fixed dimensionality. Thus the overall performance of classical networks relies heavily on their underlying model of the external world. In other words, it is not the networks that are 'smart', it is the choice of the world model that matters. Networks need to obtain 'appropriate' data in order to 'learn', but this accounts to choosing a static model of the world and in such a situation networks indeed can perform well. Our feeling is that, to a limited ex-

tent, a similar situation appears in very low level sensory processing in the brain, where only the statistical consistency of the external world matters. However, as soon as the top down information starts to interact with the bottom up processing the semantic meaning of objects becomes significant and this can often violate the assumption of static world representations. It follows that classical neural networks are well equipped only for tasks in which they process numerical data whose relationships can be well reflected by Euclidean distance. In other words classical connectionism can be reasonably well applied to the same category of problems which could be dealt with by various regression methods from statistics. Moreover, as in fact classical neural nets offer the same explanatory power as regression, they can be therefore regarded as its non-linear counterparts. It is however doubtful whether non-linear regression constitutes a satisfactory (or the most general) model of fundamental information processing in natural neural systems. Another problem follows from the rigidity of neurons' actions in current connectionist models. The homogeneity of neurons and their responses is the rule rather than the exception. All neurons perform the same action regardless of individual conditions or context. In reality, as we argue in the next section, neurons may condition their response on the particular context, set by their immediate surroundings, past behaviour and current input etc. Thus, although in principle identical, they may behave as different individuals because their behaviour can be a function of both morphology and context. Hence, in a sense, the way conventional neural networks operate resembles symbolic systems - both have built in rigid behaviour and operate in an a priori determined way. Taking different 'histories' into account would allow for the context sensitive behaviour of neurons - in effect for existence of heterogeneous neuron populations. Standard nets are surprisingly close to classical symbolic systems although they operate in different domains: the latter operating on discrete, and the former on continuous spaces. The difference between the two paradigms in fact lies in the nature of representations they act upon, and not so much in the mode of operation. Symbolic systems manipulate whole symbols at once, whereas neural nets usually employ sub-symbolic representations in their calculations. However, both execute pro-

grams, which in case of neural networks simply prescribe how to update the interconnection weights in the network. Furthermore, in practice neural networks have very well defined input and output neurons, which together with their training set, can be considered as a closed system relaxing to its steady state. In modular networks each of the ‘expert’ nets operates in a similar fashion, with well defined inputs and outputs and designed and restricted intercommunication between modules. Although many researchers have postulated a modular structure for the brain [20], with distinct functional areas being black boxes, others [44, 18] have realised that the brain operates rather like an open system which, due to the ever changing conditions, exhibits extensive connectivity between areas and no fixed input and output. The above taxonomy resembles a similar distinction between algorithmic and interactive systems in computer science, the latter possessing many interesting properties [70].

3 Biological evidence

Advances in neuroscience provide us with evidence that neurons are much more complex than previously thought [34]. In particular, it has been hypothesised that neurons can select input depending on its spatial location on the dendritic tree or on its temporal structure [34, 6, 23]. Some neurobiologists suggest that synapses can remember the history of their activation or, alternatively, that whole neurons discriminate spatial and/or temporal patterns of activity [23]. Various authors have postulated spike encoding of information in the brain [65, 60, 35]. The speed of information processing in some cortical areas, the small number of spikes emitted by many neurons in response to cognitive tasks [53, 57, 66], together with very variable behaviour of neurons in vivo [63], suggest that neurons would not be able to reliably estimate mean firing rate in the time available. Some results suggest that firing events of single neurons are reproducible with very high reliability and interspike intervals encode much more information than firing rates [8]. Others found that neurons in isolation can produce, under artificial stimulation, very regular firing with high reproducibility rate suggesting that the apparent irregularity of firing in vivo may follow from interneuronal interactions or may be

stimulus dependent [39]. The use of interspike interval coding enables richer and more structured information to be transmitted and processed by neurons. The same mean firing rate corresponds to a combinatorial number of interspike interval arrangements in a spike train. What would previously be interpreted as a single number can carry much more information in temporal coding. Moreover, temporal coding enables the system to encode unambiguously more information than is possible with a simple mean firing rate. Different parts of a spike train can encode qualitatively different information. All these possibilities have been excluded in the classical view of neural information processing. Even though a McCulloch-Pitts neuron is sufficient for production of spike trains, spike trains by themselves do not solve the binding problem (i.e. do not explain the mechanism responsible for integration of object features which are processed in spatially and temporally distributed manner). However, nothing would be gained, except possibly processing speed, if the mean firing rate encoding would be merely replaced by temporal encoding as the underlying framework of knowledge representation and processing still mixes qualitatively different information by simple algebraic operations. The irregular pattern of neuron activity in vivo [63] is inconsistent with temporal integration of excitatory post synaptic potentials (EPSPs) assumed in classical model neurons. It also introduces huge amounts of noise, thus making any task to be performed by neurons, were they unable to differentially select their input, extremely difficult. On the other hand, perhaps there is a reason for this irregular neuronal behaviour. If neurons are coincidence detectors rather than temporal integrators [34, 65] then the randomness of neuron firing is an asset rather than liability. One of the most difficult and as yet unresolved problems of computational neuroscience is that of binding distinct features of the same object into a coherent percept. However, in [49], Nelson postulates that it is the traditional view ‘transmission first, processing later’, that introduces the binding problem. On this view processing cannot be separated from transmission and, when entangled with transmission performed by neural assemblies spanning multiple neuronal areas, it makes the binding problem non-existent [50].

4 Communication metaphor

The brain's computational capabilities have to be understood in a metaphorical sense only. All matter, from the simplest particles to the most complex living organisms undergoes physical processes which, in most sciences, are not given any special interpretation. However, when it comes to nervous systems the situation changes abruptly. In neuroscience, and what follows in connectionism, it is assumed that neurons and their systems possess special computational capabilities, which are not attributed to other, even the most complex, biological substances (e.g. DNA). This is a very anthropomorphic viewpoint because, by definition, computation is an intentional notion and it assumes existence of some 'demon' able to interpret it. Thus we claim that the very assumption of computational capabilities of real neurons leads to homuncular theories of mind. In our opinion to say that neurons perform computations is equivalent to saying that e.g., a spring extended by a moderate force computes, according to Hook's law, how much it should deform. We need to stress that our stance does not imply that one should abandon using computational tools for modelling and analysing the brain. However, one should be aware of their limitations. On the other hand, although also metaphorical, treating neurons as communicating with each other captures their complex (and to us fundamental), capability of modifying behaviour depending on the context. Our claim is that communication as biological information processing could describe more compactly complex neuronal operations and provide us with intuitive understanding of the meaning of these operations (albeit we do not impose that this meaning would be accessible to single neurons). Although interpreting neurons as simple numerical or logical functions greatly simplifies their description, it introduces problems at the higher levels of neural organisation. Moreover, as we argued earlier recent neurobiological evidence supports the claim that the idea of neurons being simple computational devices has to be reconsidered. We argue that communication better describes neuron functionality than computation. In contrast to computation, communication is not a merely anthropomorphic projection on reality. Even relatively simple organisms, e.g. bacteria, communicate with each other or with the environment. This ability

is essential for their survival and it seems indispensable for more complex interactions and social behaviour of higher species. The role of communication in human development and in social interactions cannot be overestimated [14]. It seems therefore that communication is a common process used by living systems on all levels of their organisation. In our opinion the most fundamental qualitative properties of neurons postulated recently are their capability to select different parts of converging signals and the capability of choosing which signals to consider in the first place. Thus neurons can be said to communicate to each other simple events and to select information which they process or transmit further. The selection procedure could be based on some criteria dependent on the previous signals' properties such as where from and at what moment the information arrived. This would account for neurons' spatio-temporal filtering capacity. Also it would explain the amount of noise observed in the brain and apparent contrast between reliability of neural firing in vitro and their random behaviour in vivo. What is meaningful information for one neuron can be just noise for another. Moreover, such noise would not deter functionality of neurons that are capable of responding to selected information. One could object to our proposal using the parsimony principle - why introduce an extra level of complexity if it has been shown that networks of simple neurons can perform many of the tasks attributed to biological networks? However, we argue that such a position addresses a purely abstract problem, which may have nothing to do with brain modelling. What is possible to compute with artificial neurons is, in principle, a mathematical problem; how the same functionality is achieved in the brain is another matter. The information processing capacity of dendritic trees is a scientific fact not merely a conjecture. Instead of computational parsimony we propose an 'economical' one: the brain facilitates the survival of its owner and for that purpose uses all available resources to process information. Polychronisation, emergence of time locked but not synchronised patterns of activity or synfire chains, groups of mutually synchronised neurons which are out of synch with each other are other suggestions proposed in the literature that propose richer temporal processing capacity of neurons and their networks [30, 1]. In the next section we will outline a connectionist

architecture, in which neurons’ operation is based on communication more than on computation (the fact that communication is implemented using a form of computation should not detract us from recognition of the importance of the communication as a level at which the operation of neurons can be meaningfully described and interpreted versus computation as implementing such processes on an inherently Turing like architecture).

5 NESTOR: a connectionist Spiking Neuron Stochastic Diffusion Network

The connectionist network architecture of NESTOR consists of three one-dimensional layers: the input retina consisting of a layer of receptor neurons; a layer of matching neurons and a layer of memory neurons. In operation the function of NESTOR is to label and locate the best fit of a target pattern (hereafter called a memory and defined by the activity of memory neurons), from a particular external stimulus on the retina (defined by the pattern of activity on receptor neurons). All neurons output spike trains in which information is encoded in Inter-Spike Intervals (ISI’s). Such temporal coding allows for a much richer repertoire of information encoding than a simple mean firing rate. In particular, different parts of the spike train may encode different types of information effectively leading to multiplexing. Moreover, only parts of the spike trains may encode directly information about the stimulus (or memory); the other parts providing a modulatory input internal to the network and setting the context in which the sensory (or memory) information is processed. It is such a hypothetical encoding scheme that is utilised in NESTOR.

Matching neurons are fully interconnected to all other matching neurons. They are also fully connected to both receptor neurons and memory neurons. A schematic diagram of network connectivity is given in Figure 1. The restriction to one-dimensional layers is for clarity of exposition only and extension to two-dimensional layers is straightforward.

Matching neurons perform a set of operations on their input essentially akin to (nonlinear) filter-




Figure 1: *The architecture of NESTOR*

ing. In contrast to the operation of classical neuron models, matching neurons select the spike trains to process on the basis of information contained in the first part of the spike trains and on their own state. Subsequently, their output and state will be dependent on the information contained in the second part of the accepted spike trains. However, the output will always consist of (parts of) the accepted spike trains; no further processing of the inputs is needed. In this sense their operations amounts to filtering which information will be propagated throughout the network rather than calculating some numeric transformation of their overall input. Nonlinearity of filtering stems from its contextual nature; which parts of matching neuron’s input are being propagated depends on its own state and on the properties of other input components.

More specifically, matching neurons both maintain and output hypotheses - ‘where’ values - defining possible locations of the target pattern on the retina. Also for clarity of exposition we use the labels ‘active’ and ‘inactive’ to identify the internal ‘state’ of matching neurons, however these labels are neither used or known by other matching neurons. An active matching neuron will output its current hypothesis as a spike train; an inactive matching neuron will adopt a new hypothesis from the first spike train it receives and output this. This spike train can either be from another matching neuron or from memory and retina neurons.

Matching neurons evaluate their hypothesis - ‘where’ value - by comparing randomly selected micro-feature(s) - ‘what’ information - from the memory and the corresponding ‘what’ information from the retina. If the micro-features are the same

(the corresponding ‘what’ values are the same), this provides evidence that the hypothesis is good and so the matching neuron becomes active.

An active matching neuron initiates a spike train and retains its current hypothesis (its ‘where’ value). Thus a successful matching neuron will not switch to evaluate other locations in the search space (retina), but will continue to randomly evaluate its current hypothesis - ‘where’ value - by testing for the presence of other (randomly selected) micro-features from the target.

On the contrary, a neuron failing to discover the same micro-feature will remain inactive and will adopt a new hypothesis encoded in a spike train arriving from other matching neurons. As these spike trains are accepted on the basis of their arrival time, an inactive matching neuron can either exploit a potentially correct position signalled by an active matching neuron or explore a completely new position encoded in the spike train generated by another inactive neuron.

However, if the first spike train arrives from memory or retina, a matching neuron will simply output a spike train encoding the position defined by the input trains without further processing or changes to its [notional] state.

It is apparent that, over consecutive cycles of operation, a matching neuron evaluating the spike trains encoding the retinal region with the highest overlap of microfeatures with the target memory is more likely to retain its memory at the next cycle of activity, than it is when processing spike trains encoding other locations. This is because, by definition, the chance of matching a micro-feature from the retina and the memory is highest here. Effectively matching neurons operate by filtering the incoming spike trains on the basis of their past activity.

The location on the retina of the best fit to the memory located by the network is encoded in the dominant spike train output of matching neurons, i.e. it is the position corresponding to the mode of the distribution of matching neuron ‘where’ signals.

As the process of memory location by the network is statistical in nature, the assembly of neurons encoding the best fit solution can fluctuate dynamically both in number and identity. Even though single neurons can change their ‘where’ information with relatively high probability and their activity can be considered random when considered

in isolation, they nevertheless collectively produce stable behaviour in a quasi-deterministic manner. This form of dynamic information decoding has to be contrasted with the conventional neural networks operation where the readout from the network is possible due to the deterministic and fixed functionality of individual neurons.

As the time jitter of matching neurons efferent activity is small compared to the length of the information encoding, an assembly of neurons finding the best-fit may produce time locked or near oscillatory behaviour. *Hence in this model oscillatory behaviour may be a result of, rather than a cause of, the binding of features belonging to the same object.*

5.1 Retina and memory

The function of each receptor neuron is to communicate its position on the retina and to encode the micro-feature that stimulated it. The first ISI sent by a receptor neuron encodes information about its position on the retina (the ‘where’ information, Δt_{ret}^w) and the second ISI encodes the micro-feature that stimulated this receptor (the ‘what’ information, Δt_{ret}^m). For modelling purposes we assume the following relationship between a receptor’s retinal position and its encoding via the ISI as:

$$\Delta t_{j+1}^w \propto \Delta t_j^w + \Delta t^w \quad (1)$$

where $j, j+1$ are two adjacent retinal positions and Δt^w is the minimal increment of the ISI (see Figure 2). The above encoding is akin to the principle of preservation of receptor’s topographic organisation although its particular form is arbitrary and assumed here to simplify this exposition - any relation enabling determination of the length of the first ISIs of neighbouring receptor neurons would be equally suitable.

In operation memory neurons are functionally identical to receptor neurons. Thus at the onset of activity, their first ISI, Δt_{mem}^w encodes relative ‘where’ information and the second ISI, Δt_{mem}^m encodes memory ‘what’ information.

Hence the only difference between receptor cells and memory neurons is that receptors respond to and encode external stimuli whereas memory neurons encode and propagate internal representations of the external world. However, in operation matching neurons need to distinguish the ‘source’




Figure 2: *Encoding of information in the spike trains of receptor cells. The relative position of receptor cells in the spike train is encoded in the first Inter-Spike Intervals. The difference t , between the first Inter-Spike Interval's lengths of adjacent receptor cells is constant for any receptor cell j . The second ISI's encode micro-features that activated corresponding receptor cells.*

of afferent information (matching, memory or receptor neuron). This requires either 'labelled' input lines or some form of source-encoding within the spike train. However, both methods provide information at the expense of an increase in network complexity, therefore to aid the clarity of this short exposition, the process underlying this particular mechanism will be left open, with source information explicitly available to matching neurons as required.

5.2 Matching layer neurons

A matching neuron only stores the 'where' information (defined by the ISI encoding the value of Δt_{neur}^w), of an accepted spike train. The delay between a matching neuron firing and the arrival of the resulting spike train to efferent neurons follows from the transmission delays introduced by efferent axons.

We assume that the sum of the maximal length of the spike train, T_l and maximal time delay T_d is smaller than the length of the time interval between consecutive updates of matching neurons, T :

$$T_l + T_d \leq T \quad (2)$$

Further, it is assumed that the probability distribution of the spike train arrival times to a given matching neuron is the same for all afferent axons.

The method by which a matching neuron processes incoming information depends on the information itself and on the past activity imprinted in neuron's memory (Δt_{neur}^w) and state (either active or inactive). If a matching neuron is in an active state, it will select for processing the first spike trains from the retina and memory that fulfil the condition:

$$\Delta t_{ret}^w + \Delta t_{mem}^w = \Delta t_{neur}^w \quad (3)$$

This condition ensures that each matching neuron selects a micro-feature from memory and compares it with appropriate information from the retina, (see Figure 3). Subsequently, the matching neuron will compare the second ISIs from the corresponding spike trains, Δt_{mem}^m and Δt_{ret}^m , defining the 'what' information from the memory and the retina.

If the comparison is successful (i.e. the micro-feature propagated by the receptor neuron is identical to that propagated by the memory neuron) the matching neuron fires a spike train corresponding to the position defined by Δt_{neur}^w . In the next cycle of activity it will explore other microfeatures at that position as it will remain in an active state and its internal memory will not change. Otherwise, the matching neuron will become inactive. In the inactive state the matching neuron will accept information from either the first arriving spike train from another matching neuron or spike trains memory and the retina (albeit with no constraints on ISI's). In the latter case the matching neuron will fire a spike train corresponding to the position defined jointly by first inter spike intervals arriving from the memory and the retina. In the former case it will modify its memory, Δt_{neur}^w , to store the first ISI from another neuron and will change its state to active. This mechanism allows a matching neuron to check the position pointed to by another successful matching neuron or alternatively a completely random position in the search space derived from the spike train generated by another inactive neuron. From the fact that arrival times

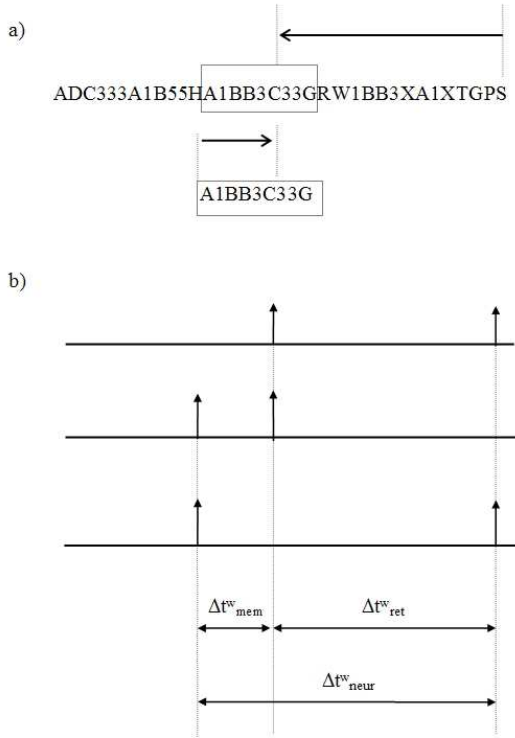


Figure 3: (a) Relative positions of the micro-feature, C , selected from the target memory, $[A1BB3C33G]$, and on the retina, $[ADC33 .. A1BB3C33G .. XTGPS]$; (b) their encoding in the length of Inter-Spike Intervals by receptor cell and memory neuron respectively, together with the resulting matching neuron target position encoding.

of spike trains are independent and identically distributed random variables it follows that the choice of the spike trains based on their arrival times to a given neuron is unbiased. It means that a matching neuron can choose information arriving from other matching neurons with equal probability.

Operation of NESTOR can be better understood if one realises that it is closely related to a distributed swarm intelligence system called Stochastic Diffusion Search described in the next section.

6 Stochastic Diffusion Search

SDS is an efficient probabilistic swarm intelligence global search and optimisation technique that has been applied to diverse problems such as site selection for wireless networks [68], mobile robot self-localisation [10], object recognition [12] and text search [11]. Additionally, a hybrid SDS and n-tuple RAM [2] technique has been used to track facial features in video sequences [12, 24].

Previous analysis of SDS has investigated its global convergence [45], linear time complexity [46] and resource allocation [48] under a variety of search conditions.

SDS is based on distributed computation, in which the operations of simple computational units, or agents are inherently probabilistic. Agents collectively construct the solution by performing independent searches followed by diffusion of information through the population. Positive feedback promotes better solutions by allocating to them more agents for their exploration. Limited resources induce strong competition from which the largest population of agents corresponding to the best-fit solution rapidly emerges.

SDS uses a population of agents. In many search problems the solution can be thought of as composed of many subparts and SDS explicitly utilises such decomposition to increase the search efficiency of individual agents. In SDS each agent poses a hypothesis about the possible solution and evaluates it partially. Successful agents repeatedly test their hypothesis while recruiting unsuccessful agents by direct communication. This creates a positive feedback mechanism ensuring rapid convergence of agents onto promising solutions in the space of all solutions. Regions of the solution space labelled by the presence of agent clusters can be

interpreted as good candidate solutions. A global solution is thus constructed from the interaction of many simple, locally operating agents forming the largest cluster. Such a cluster is dynamic in nature, yet stable, analogous to, “a forest whose contours do not change but whose individual trees do”, [3, 13]. The search mechanism can be illustrated with the following analogy.

6.1 The restaurant game

A group of delegates attends a long conference in an unfamiliar town. Each night they have to find somewhere to dine. There is a large choice of restaurants, each of which offers a large variety of meals. The problem the group faces is to find the best restaurant, that is the restaurant where the maximum number of delegates would enjoy dining (this simplistic model constructed to illustrate the SDS assumes that all diners have the same preferences). Even a parallel exhaustive search through the restaurant and meal combinations would take too long to accomplish. To solve the problem delegates decide to employ a Stochastic Diffusion Search.

Each delegate acts as an agent maintaining a hypothesis identifying the best restaurant in town. Each night each delegate partially evaluates his hypothesis by dining there and randomly selecting one of the meals from the menu. The next morning at breakfast every delegate who did not enjoy his meal the previous night, asks one randomly selected colleague to share his dinner impressions. If the experience was good, the unsatisfied diner also adopts this restaurant as his choice. Otherwise he simply selects another restaurant at random from those listed in ‘Yellow Pages’.

Using this strategy it is found that very rapidly significant number of delegates congregate around the ‘best’ restaurant in town.

Abstracting from the above algorithmic process:

By iterating through test and diffusion phases agents stochastically explore the whole solution space. However, since tests succeed more often on good candidate solutions than in regions with irrelevant information, an individual agent will spend more time examining good regions, at the same time recruiting other agents, which in turn, via a positive feedback cycle, recruit even more agents. Good candidate solutions are thus iden-

Initialisation phase

Whereby all agents (delegates) generate an initial hypothesis (select a restaurant).

loop

Test phase

Each agent evaluates evidence for its hypothesis (meal degustation).

Agents are classified as active (happy diners) or inactive (disgruntled diners).

Diffusion phase

Inactive agents adopt a new hypothesis by either communication with another agent or, if the selected agent is also inactive, there is no information flow between the agents; instead the selecting agent must adopt a new hypothesis (restaurant) at random.

endloop

tified by concentrations of a substantial population of agents.

Central to the power of SDS is its ability to escape local minima. This is achieved by the probabilistic outcome of the partial hypothesis evaluation in combination with reallocation of resources (agents) via stochastic recruitment mechanisms. Partial hypothesis evaluation allows an agent to quickly form its opinion on the quality of the investigated solution without exhaustive testing (e.g. it can find the best restaurant in town without having to try all the meals available in each).

7 Comparison of SDS and NESTOR

Under specific assumptions about the probability distribution of delays and the use of multiplexing of information encoded in the spike trains, NESTOR mimics operation of the generic SDS. The correspondence between these architectures can be established by equating matching neurons of the latter with agents. However, the important difference between the two is that in SDS the activity state of a given agent is accessible to other agents, whereas in NESTOR information about the internal state of a matching neuron is only used lo-

cally. Matching neurons do not have direct access to the internal states of other neurons. Moreover, the two algorithms differ in the way they encode information. All the information necessary for operation of matching neurons is (in principle) encoded explicitly, whereas in SDS agents can only act on the information distributed by other agents; their activity constituting an implicit information. Moreover, the information encoding in SDS is in an abstract, dimensionless form, typical for vast majority of artificial neural networks and even for more faithful neural models utilising rate coding. On the contrary, all the information processed by matching neurons in NESTOR has a physical dimension of time interval between consecutive spikes. This characteristic introduces time delays into NESTOR and thus bears on the dynamics of this network. Another difference between the systems follows from different modes of operation. SDS operates in synchronous mode and all agents go through their cycles of activity in parallel. In contrast, matching neurons in NESTOR operate asynchronously, which prevents a one to one correspondence between the largest cluster of active agents in the in SDS and the strongest invariant spike train pattern encoding the solution found by NESTOR. Both above mentioned features affect the time count in both systems - several hundreds of time steps in NESTOR correspond to a single iteration of SDS.

Figure 4 shows an example of operation of both SDS, panel (a), and NESTOR, panel (b), illustrating these differences. Thus, traces denote the number of agents in the largest active cluster in case of SDS, and the number of active neurons firing the spike train encoding the solution in NESTOR. The differences discussed above influence the quantitative differences in the dynamics of the two systems (e.g. much larger variance in case of NESTOR). Nevertheless, in spite of these differences, the qualitative behaviour of SDS and NESTOR is similar (in both the time course and quasi-steady state behaviour).

8 Discussion and conclusions

The neural network outlined in this paper performs SDS and solves the best-fit matching problem. This functionality emerges from its ability to

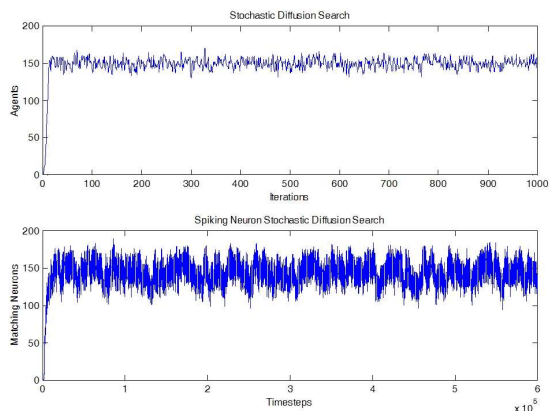


Figure 4: (a) Time evolution of the largest cluster of active agents pointing to the same position in the search space in SDS; (b) time evolution of the number of active matching neurons emitting the spike train encoding the solution found by NESTOR. The difference of the information encoding mode between the two algorithms affects different time scales as well as quantitative measures of response such as variance (see text for details). However, qualitative behaviour of both systems is similar in both time course and quasi-steady state.

self-organise in response to incoming stimuli. The network effectively uses a tagged dynamic assembly encoding for the target, and finding this within the search space results in the onset of time locked activity of spike trains within the neural assembly.

Micro-features defining the target are tagged both by their retinal encoding and their relative position on the retina. A label-tag encodes ‘what’ the micro-feature is and a location-tag defines ‘where’ it is. Thus the NESTOR network 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 vectors in Euclidean space). It has been suggested by Van de Velde [69] that ‘type’ representation schemes are a fundamental cause of many of the problems encountered when modelling symbolic processes by associative networks.

NESTOR constitutes a very simplified model qualitatively demonstrating how information may be processed in a neural architecture utilising temporal encoding and multiplexing. As such, some important issues such as encoding of memories need future elaboration. Instead, this paper argues the possibility that knowledge representation in the networks using temporal encoding of information may be fundamentally different from that used in classical nets. In our model the activity of single neurons does not suggest a semantical interpretation, as neurons respond to all features. 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. The neurons do not constitute the internal representation in themselves, as the assembly is dynamically fluctuating, but their pattern of activity does. Thus, although an assembly supporting a particular representation in different time instants can differ considerably in the number and identity of its constituent neurons, the representation 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 described by Treisman [64]. 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 ‘locks’ to another. Moreover, the network will retain one of the most fundamental properties of SDS - automatic allocation of resources, whose characteristics have been extensively investigated in [48].

References

- [1] Abeles, M. Synfire chains. In M. A. Arbib (Ed.), *The handbook of Brain theory and neural networks*. (pp. 1143-1146). Cambridge, MA: MIT Press, 2002.
- [2] I. Aleksander and T.J. Stonham, *Guide to pattern recognition using random access memories*, *Computers & Digital Techniques* 2:1 (1979), 29-40.
- [3] W.B. Arthur, Inductive reasoning and bounded rationality, (The El Farol Problem), *Amer. Econ.Rev. Papers and Proceedings* 84 (1994), 406.
- [4] G. Ascoli, Progress and perspectives in computational neuroanatomy, *The Anatomical Record (New Anat.)*, 257, pp. 195-207, 1999.
- [5] T. Back, *Evolutionary Algorithms in Theory and Practice*, Oxford University Press, 1996.
- [6] [H. Barlow, Intraneuronal information processing, directional selectivity and memory for spatio-temporal sequences. *Network: Computation in Neural Systems* 7 (1996) 251-259.
- [7] [J. Barnden & J. Pollack (eds.), *High-Level Connectionist Models*, Ablex: Norwood, NJ, (1990).
- [8] [M.J. Berry, et al., The structure and precision of retinal spike trains. *Proc. Natl. Acad. Sci. USA* 94 (1997) 5411-5416.
- [9] J.M. Bishop, *Anarchic Techniques for Pattern Classification*, Chapter 5. PhD Thesis, University of Reading, 1989.
- [10] P.D. Beattie & J.M. Bishop, *Self-Localisation in the ‘SENARIO’ Autonomous Wheelchair*,

- Journal of Intelligent and Robotic Systems, 22 1998, 255-267.
- [11] J.M. Bishop, Stochastic Searching Networks, Proc. 1st IEE Int. Conf. ANNs, London, UK, (1989), 329-331.
- [12] J.M. Bishop & P. Torr, The Stochastic Search Network, in Linggard, R., Myers, D.J., Nightingale, C. (eds.) Neural Networks for Images, Speech and Natural Language, Chapman & Hall, New York, 1992, pp. 370-387.
- [13] J.M. Bishop, S.J. Nasuto, & K. De Meyer, Knowledge Representation in Connectionist Systems. ICANN 2002, Madrid, Spain, (2002).
- [14] [R. Brown, Social psychology. Free Press, New York (1965)
- [15] P. Cariani, As if time really mattered: Temporal strategies for neural coding of sensory information, in Origins: Brain and Self-Organization, K. Pribram (ed.), Lawrence Erlbaum Assoc, 1994.
- [16] D. Corne, M. Dorigo, & F. Glover, New Ideas in Optimisation, McGraw-Hill, 1999.
- [17] E. Bonabeau, M. Dorigo, & G. Theraulaz, Swarm Intelligence: from Natural to Artificial Systems, Oxford University Press, 1999.
- [18] [M. Farah, Neuropsychological inference with an interactive brain: A critique of the locality assumption. Behavioural and Brain Sciences (1993)
- [19] [J. Fodor & Z.W. Pylyshyn, Z.W.: Connectionism and Cognitive Architecture: A Critical Analysis. In: Boden, M.A. (ed.): The Philosophy of Artificial Intelligence, Oxford University Press (1990)
- [20] [J.A. Fodor, The Modularity of Mind. MIT Press (1983).
- [21] [K.J. Friston, Transients, Metastability, and Neuronal Dynamics. Neuroimage 5 (1997) 164-171.
- [22] [K. Fukushima, Neocognitron: A hierarchical neural network capable of visual pattern recognition. Neural Networks 1 (1988), 119-130.
- [23] [R. Granger et al.: Non-Hebbian properties of long-term potentiation enable high-capacity encoding of temporal sequences. Proc. Natl. Acad. Sci. USA Oct (1991) 10104-10108.
- [24] H.J. Grech-Cini, & G.T. McKee, Locating the Mouth Region in Images of Human Faces. In P.S.Schenker (Ed.), Proceedings of SPIE - The International Society for Optical Engineering, Sensor Fusion VI 2059, Massachusetts, 1993.
- [25] [D. Goldberg, Genetic Algorithms in search, optimization and machine learning. Addison Wesley, Reading MA, 1989.
- [26] [S. Haykin, Neural Networks: A Comprehensive Foundation. Macmillan, New York (1994).
- [27] R. Hirsh, Modulatory integration: A concept capable of explaining cognitive learning and purposive behavior in physiological terms, Psychobiology, 18 (1), pp: 3-15, (1990).
- [28] J.H. Holland, Adaptation in natural and artificial systems. The University of Michigan Press Ann Arbor MIT, 1975.
- [29] M. Iosifescu, Finite Markov processes and their applications. Wiley, Chichester, 1980.
- [30] E. M. Izhikevich, Polychronization: Computation With Spikes, Neural Computation (2006) 18:245-282
- [31] [M.I. Jordan & R.A. Jacobs, Hierarchical mixtures of experts and the EM algorithm. MIT Comp. Cog. Sci. Tech. Report 9301 (1993).
- [32] [E.N. Kamas, L.M. Reder, & M.S. Ayers, Partial matching in the Moses illusion: Response bias not sensitivity, Memory & Cognition, 24:6, pp. 687-699, 1996.
- [33] J. Kennedy & R.C. Eberhart, Swarm Intelligence, Morgan Kaufman, San Francisco, 2001.
- [34] [Koch, C.: Computation and the single neuron. Nature 385 (1997) 207-210
- [35] [Koenig, P., et al.: Integrator or coincidence detector? The role of the cortical neuron revisited. Trends Neurosci. 19(4) (1996) 130-137.

- [36] K. De Meyer, Explorations in Stochastic Diffusion Search: soft- and hardware implementations of biologically inspired Spiking Neuron Stochastic Diffusion Networks, University of Reading, UK KDM/JMB/2000-1, 2000.
- [37] K. De Meyer, J.M. Bishop, & S.J. Nasuto, Small World Network behaviour of Stochastic Diffusion Search. ICANN 2002, Madrid, Spain, 2002.
- [38] K. De Meyer, J.M. Bishop, & S.J. Nasuto, Attention through Self-Synchronisation in the Spiking Neuron Stochastic Diffusion Network. *Consc. and Cogn.* 9:2 (2000), 81-81.
- [39] [Z.F. Mainen & T.J. Sejnowski, Reliability of spike timing in neocortical neurons. *Science* 168 (1995) 1503-1506.
- [40] M. Moglich, U. Maschwitz & B. Holldobler, Tandem Calling: A New Kind of Signal in Ant Communication. *Science*, 186:4168 (1974), pp. 1046-1047.
- [41] T. Morey, K. De Meyer, S.J. Nasuto & J.M. Bishop, Implementation of the Spiking Neuron Stochastic Diffusion Network on Parallel Hardware (abs). *Consciousness and Cognition*, 9:2 97-97, (2000), Academic Press.
- [42] W.S. McCulloch & W.Pitts, .: A logical calculus immanent in nervous activity. *Bulletin of Mathematical Biophysics* 5, (1943), 115-133.
- [43] C. F. Moukarzel, Spreading and Shortest Paths in Systems with Sparse Long- Range Connections. *Physical Review E* 60:6, (1999), 62636266.
- [44] [D. Mumford, Neural Architectures for Pattern-theoretic Problems. In: Koch, Ch., Davies, J.L. (eds.): *Large Scale Neuronal Theories of the Brain*. The MIT Press, London, England (1994).
- [45] S.J. Nasuto & J.M. Bishop, Convergence Analysis of Stochastic Diffusion Search. *Journal of Parallel Algorithms and Applications* 14:2, pp: 89-107, 1999.
- [46] S.J. Nasuto, J.M. Bishop & S. Lauria, Time Complexity of Stochastic Diffusion Search. *Neural Computation '98*, Vienna, Austria, 1998.
- [47] S.J. Nasuto, K. Dautenhahn & J.M. Bishop, Communication as an Emergent Methaphor for Neuronal Operation. *Lecture Notes in Artificial Intelligence* 1562 (1999), 365-380.
- [48] S.J. Nasuto, Analysis of Resource Allocation of Stochastic Diffusion Search, PhD Thesis, University of Reading, UK, 1999.
- [49] [J.I. Nelson, Visual Scene Perception: Neurophysiology. In: Arbib, M.A. (ed.): *The Handbook of Brain Theory and Neural Networks*. MIT Press: Cambridge MA (1995).
- [50] [J.I. Nelson, Binding in the Visual System. In: Arbib, M.A. (Ed.): *The Handbook of Brain Theory and Neural Networks*, MIT Press, Cambridge MA (1995).
- [51] A. Neumaier, Complete search in continuous global optimization and constraint satisfaction. In Isereles, A., (ed.) *Acta Numerica* 2004, Cambridge University Press.
- [52] M.E.J. Newman & D.J. Watts, Scaling and Percolation in the Small-World Network Model. *Physical Review E* 60:6 (1999), 7332-7342.
- [53] [D.I. Perret et al., Visual neurons responsive to faces in the monkey temporal cortex. *Experimental Brain Research* 47 (1982) 329-342.
- [54] [S. Pinker & A. Prince, On Language and Connectionism: Analysis of a Parallel Distributed Processing Model of Language Acquisition. In: Pinker, S., Mahler, J. (eds.): *Connections and Symbols*, MIT Press, Cambridge MA, (1988).
- [55] T. Poggio & F. Girosi, Networks for approximation and learning. *Proceedings of the IEEE* 78 (1990), 1481-1497.
- [56] F. Rosenblatt, *Principles of Neurodynamics*. Spartan Books, Washington DC (1962).
- [57] [E.T. Rolls & M.J. Tovee, Processing speed in the cerebral cortex and the neurophysiology of visual backward masking. *Proc. Roy. Soc. B* 257 (1994) 9-15.
- [58] [D.E. Rumelhart & J.L. McClelland, J.L., (eds.), *Parallel Distributed Processing. Explorations in the Microstructure of Cognition*, MIT Press, Cambridge MA (1986).

- [59] [T.J. Sejnowski & C.R. Rosenberg, C.R., Parallel networks that learn to pronounce English text. *Complex Systems* 1 (1987), 145-168.
- [60] [T.J. Sejnowski, Time for a new neural code ?, *Nature* 376 (1995) 21-22.
- [61] [B. Selman et al., Challenge Problems for Artificial Intelligence. Proceedings of AAAI-96, National Conference on Artificial Intelligence, AAAI Press, 1996.
- [62] [G.M. Shepherd The Synaptic Organisation of the Brain. Oxford University Press, London Toronto, 2nd Ed, (2004).
- [63] [W.R. Softky & C. Koch, The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSP. *J. of Neurosci.* 13 (1993) 334-350.
- [64] A. Treisman, Features and Objects: The fourteenth Bartlett memorial lecture, *The Quarterly Journal of Experimental Psychology*, 40A:2, pp. 201-237, 1988.
- [65] [A.M. Thomson, More Than Just Frequency Detectors ?. *Science* 275 Jan (1997) 179-180.
- [66] [S.J. Thorpe & M. Imbert, Biological constraints on connectionist modelling. In: Pfeifer, R., et al. (eds.): *Connectionism in Perspective*. Elsevier (1989).
- [67] D.J. Watts & S.H. Strogatz, Collective Dynamics of Small-World Networks. *Nature* 393 (1998), 440-442.
- [68] R.M. Whitaker and S. Hurley, An agent based approach to site selection for wireless networks. *Proc ACM Symposium on Applied Computing (Madrid)*, (2002), 574-577.
- [69] F. Van der Velde, On the use of computation in modelling behaviour, *Network: Computation in Neural Systems*, 8, pp. 1-32, 1997.
- [70] [P. Wegner, Why Interaction is More Powerful than Algorithms. *CACM* May (1997).
- [71] R. Yuste, Dendritic shock absorbers, *Nature*, 387, pp: 851-853, 1997.