

Communication as an emergent metaphor for neuronal operation

Slawomir J. Nasuto¹, Kerstin Dautenhahn¹, Mark Bishop²

Department of Cybernetics, University of Reading,
Reading, RG2 6AE, UK,

¹{sjn, kd}@cyber.rdg.ac.uk,

²J.M.Bishop@reading.ac.uk

Abstract. The conventional computational description of brain operations has to be understood in a metaphorical sense. In this paper arguments supporting the claim that this metaphor is too restrictive are presented. A new metaphor more accurately describing recently discovered emergent characteristics of neuron functionality is proposed and its implications are discussed. A connectionist system fitting the new paradigm is presented and its use for attention modelling briefly outlined.

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 [1,2].

The revival of the connectionism in the mid-eighties featured increased interest in analysing the properties of such networks [3], as well as in applying them to numerous practical problems [4]. At the same time the same devices were proposed as models of cognition capable of explaining both higher level mental processes [5] and low level information processing in the brain [6].

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 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 compliment 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 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.

In a recent paper [7], Selman et al. posed some challenge problems for Artificial Intelligence. In particular Rodney Brooks suggested revising the conventional McCulloch Pitts neuron model and investigation of 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 [7], as it is isomorphic in operation to Stochastic Diffusion Search, a fast, generic probabilistic search procedure which automatically allocates information processing resources to search tasks.

Computational metaphor.

The emergence of connectionism is based on the belief that neurons can be treated as simple computational devices [1]. 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 internal representation of the (potentially complex) relationships between training examples. Such information exists only as a distribution of weight values. We do not consider representations of arity zero predicates, (e.g. those present in NETalk [8]), 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 [9,10,11].

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 an 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 restriction of their range) or using a 'mixture of experts or gated networks' [12]. However, there is no a 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 problems. 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 although 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 [13].

We think that the 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 [14] 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 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 extent, 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 programs, 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 [15], with distinct functional areas being black boxes, more recently some [16, 17] have realised that the brain operates rather like an open system. And due to the ever changing conditions a system with 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 [18].

Biological evidence.

Recent advances in neuroscience provide us with evidence that neurons are much more complex than previously thought [19]. In particular it has been hypothesised that neurons can select input depending on its spatial location on dendritic tree or temporal structure [19,20,21]. 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 [21].

Various authors have postulated spike encoding of information in the brain [22,23,24]. The speed of information processing in some cortical areas, the small number of spikes emitted by many neurons in response to cognitive tasks [25,26,27], together with very random behaviour of neurons *in vivo* [28], suggest that neurons would not be able to reliably estimate mean firing rate in the time available. Recent results suggest that firing events of single neurons are reproducible with very high reliability and interspike intervals encode much more information than firing rates [29]. 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 [30].

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 constituting an 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* [28] is inconsistent with temporal integration of excitatory post synaptic potentials (EPSP's) 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 [19,22] 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 [31], 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 [32].

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 that 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 however problems at the higher levels of neural organisation. Moreover, recent neurobiological evidence supports our 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 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 [33]. 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 parsimony principle - why to 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 it 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.

Architecture of NESTOR.

Taking into account the above considerations we adopt a model neuron that inherently operates on rich information (encoded in spike trains) rather than a simple mean firing rate. Our neuron simply accepts information for processing dependent on conditions imposed by a previously accepted spike train. It compares corresponding parts of the spike trains and, depending on the result, further distributes the other parts. Thus neurons do not perform any numerical operations on the obtained information - they forward its unchanged parts to other neurons. Their power relies on the capability to select appropriate information from the incoming input depending on the context set by their history and the activity of other neurons.

Although we define a single neuron as a functional unit in our architecture we are aware that the debate on what constitutes such a unit is far from being resolved. We based this assumption on our interpretation of neurobiological evidence. However, we realise that even among neuroscientist there is no agreement as to what constitutes such elementary functional unit, (proposals range from systems of neurons or microcircuits [34], through single neurons [35] to single synapses [13]). In fact it is possible that qualitatively similar functional units might be found on different levels of brain organisation.

In the characteristics of this simple model neuron we have tried to capture what we consider to be fundamental properties of neurons. Although our model neurons are also dimensionless, nevertheless in their information processing characteristics we included what might follow for real neurons from their geometric properties (namely ability to distinguish their inputs - spatio-temporal filtering).

A network of such model neurons was proposed in [36]. The NEural STochastic diffusion search netWOrk (NESTOR) consists of an artificial retina, a layer of fully connected matching neurons and retinotopically organised memory neurons. Matching neurons are fully connected to both retina and memory neurons.

It is important to note that matching neurons obtain both ascending and descending inputs. Thus their operation is influenced by both bottom-up and top-down information. As Mumford [16] notices, systems which depend on interaction between feedforward and feedback loops are quite distinct from models based on Marr's feedforward theory of vision.

The information processed by neurons is encoded by a spike train consisting of two qualitatively different parts - a tag determined by the relative position of the receptor on the artificial retina and a feature signalled by that receptor. The neurons operate by introducing time delays and acting as spatiotemporal coincidence detectors.

Although we exclusively used a temporal coding, we do not mean to imply that firing rates do not convey any information in the brain. This choice was undertaken for simplicity of exposition and because in our simplified architecture it is not important how the information about the stimulus is encoded. What is important is the possibility of conveying more information in spike trains than it would be possible if information is only encoded in a single number (mean firing rate). As far as we are aware there are no really convincing arguments for eliminating one of the possible

encodings and in fact both codes might be used in the brain - mean firing for stimulus encoding and temporal structure of spike trains for tagging relevant information.

NESTOR uses a dynamic assembly encoding for the target. Finding it in the search space results in onset of time locked activity of the assembly. Different features of the same object are bound by their relevant position in the search space and synchronisation of activity within the assembly may follow as a result of binding. Thus binding in the network is achieved by using additional information contained in tags.

Effectively NESTOR implements the Stochastic Diffusion Search (SDS) [37] - a matching algorithm whose operation depends on co-operation and competition of agents which were realised here as model neurons. Therefore in the next section we will describe the network operation in terms of the underlying *generic* mechanism of SDS.

Stochastic Diffusion Search.

SDS consists of a number of simple agents acting independently but whose collective behaviour locates the best-fit to a predefined target within the specified search space. Figure 1 illustrates the operation of SDS on an example search space consisting of a string of digits with the target - a pattern '371' - being exactly instantiated in the search space.

It is assumed that both the target and the search space are constructed out of a known set of basic microfeatures (e.g. bitmap pixel intensities, intensity gradients, phonemes etc.). The task of the system is to solve the best fit matching problem - to locate the target or if it does not exist its best instantiation in the search space. Initially each agent samples an arbitrary position in the search space, checking if some microfeature in that position matches with corresponding microfeature of the target. If this is the case, then the agent becomes active otherwise it is inactive. Activity distinguishes agents which are more likely to point to a correct position from the rest.

Next, in a diffusion phase, each inactive agent chooses at random another agent for communication. If the chosen agent is active, then its position in the search space will be copied by the inactive agent. If, on the other hand, the chosen agent is also inactive then the choosing agent will reallocate itself to an arbitrary position in the search space.

This procedure iterates until SDS reaches an equilibrium state, where a maximal stable population of active agents will point towards common position in the search space. In the most general case convergence of SDS has to be interpreted in statistical sense [38]. The population supporting the solution will fluctuate, identities of particular agents in this population will change but nevertheless the system as a whole will exhibit a deterministic behaviour. From such competition and co-operation between weakly randomly coupled agents emerges the deterministic behaviour of SDS. This self-organisation in response to an external stimulus incoming from the search space is one of the most important properties of SDS.




Fig. 1. SDS consisting of five agents searching in the string of digits for a pattern '371'. Active agents point to corresponding features with (*solid arrows*). Inactive agents are connected to the last checked features by (*dashed lines*). Agents pointing to the correct position are encircled by (*ovals*). The first number in the agent denotes position of the potential solution and the second number - the relative position of the checked microfeature

The time complexity of SDS was analysed in [39] and shown to be sublinear in the presence of no noise when the perfect match is present. Further work has confirmed that this characteristic also holds in more general conditions. As noted in [39] this performance is achieved without using heuristic strategies, in contrast to the best deterministic one- and two-dimensional string searching algorithms or their extensions to tree matching [40], which at best achieve time linearity.

Attention modelling with NESTOR.

Conventional models of visual attention are based on concepts of separate feature maps, which are composed of neurons selective to the appropriate feature only [41].

However recent research [42] suggests that in most visual cortical areas neurons respond to almost any features, implying a multiplexing problem. Moreover, a majority of cells responding to a particular feature often reside outside of the area supposed to be responsible for extracting this feature from the scene.

Information processing by assemblies spanned by intercommunicating neurons from distant areas of the brain has already been postulated [32] as the fundamental operation mode of the brain. This view, together with findings on long range interactions resulting in receptive fields spanning multiple cortical areas [43], in fact reduces the division of the cortex into many separate areas to a mere neuroanatomical taxonomy. It also supports the hypothesis that local interactions are not the most important feature of real biological networks. The most recent findings suggest that, contrary to assumptions of some researchers [41], attention may be operating on all levels of visual system with the expectation of the whole system directly influencing

cell receptive fields and, as a result, information processing by single neurons (for an excellent exposition see [44] and references therein).

These findings are qualitatively reflected in the architecture of NESTOR. Although network architecture and neuron properties only very approximately correspond to the architecture of the visual system and properties of real neurons, nevertheless, in the light of the cited evidence, we think that it is an interesting candidate for modelling visual attention.

The formation of a dynamic assembly representing the best fit to the target corresponds to an attentional mechanism allocating available resources to the desired object.

The analysis of properties of our model suggests that both parallel and serial attention may be just different facets of one mechanism. Parallel processing is performed by individual neurons and serial attention emerges as a result of formation of an assembly and its shifts between interesting objects in the search space.

Conclusions.

Much new evidence is emerging from the neuroscience literature. It points to the neuron as a complex device, acting as a spatio-temporal filter probably processing much richer information than originally assumed. At the same time our understanding of information processing in the brain has to be revised on the systems level. Research suggests that communication should not be disentangled from computation, thus bringing into question the usefulness of 'control-theoretic' like models based on clearly defined separate functional units.

We claim that this new evidence suggests supplementing the oversimplistic McCulloch-Pitts neuron model by models taking into account such a communication metaphor. It seems more accurate and natural to describe emergent neuron operations in terms of communication - a vital process for all living organisms - exhibiting 'computations' only as a mean of implementing neuron functionality in biological hardware. In this way we will avoid several problems lurking behind computational metaphor, such as homunculus theories of mind and the binding problem.

We propose a particular model neuron and discuss a network of such neurons (NESTOR) effectively equivalent to the Stochastic Diffusion Search. NESTOR shows all the interesting properties of SDS and moreover we think that it serves as an interesting model of visual attention. The behaviour of neurons in our model is context sensitive and the architecture allows for extending to heterogeneous neural populations.

Although the model advanced in this paper is based solely on exploring the communication metaphor we argue that it shows interesting information processing capabilities - fast search for the global optimum solution to a given problem and automatic allocation of resources, maintaining in parallel exploration and exploitation of the search space.

In this article we focus on the implications of communication for information processing of single neurons, which enable us to make first steps in the analysis,

analogous to advances in analysis of purely computational models. However, we are aware that the model proposed here occupies an opposite end, with respect to the McCulloch Pitts model, of an entire spectrum of alternatives. It seems reasonable that the most realistic model neurons would enjoy properties of both the computational McCulloch Pitts and our communication based model. Nonetheless we hope that adopting a communication metaphor will result in more adequate models of the brain being developed, eventually helping us to better exploit the brain's strengths and avoid its weaknesses in building artificial systems which aim to mimic brain functionality.

Acknowledgments.

Authors would like to thank anonymous referee for critical comments which helped us to refine and improve our paper.

References.

1. McCulloch, W.S., Pitts, W.: A logical calculus immanent in nervous activity. *Bulletin of Mathematical Biophysics* 5 (1943) 115-133
2. Rosenblatt, F.: *Principles of Neurodynamics*. Spartan Books, Washington DC (1962)
3. Poggio, T., Girosi, F.: Networks for approximation and learning. *Proceedings of the IEEE* 78 (1990) 1481-1497
4. Haykin, S.: *Neural Networks: A Comprehensive Foundation*. Macmillan, New York (1994)
5. Rumelhart, D. E., McClelland, J.L. (eds.): *Parallel Distributed Processing. Explorations in the Microstructure of Cognition*, MIT Press, Cambridge MA (1986)
6. Fukushima, K.: Neocognitron: A hierarchical neural network capable of visual pattern recognition. *Neural Networks* 1 (1988) 119-130
7. Selman, B. et al.: *Challenge Problems for Artificial Intelligence*. *Proceedings of AAAI-96, National Conference on Artificial Intelligence*, AAAI Press, 1996
8. Sejnowski, T.J., Rosenberg, C.R.: Parallel networks that learn to pronounce English text. *Complex Systems* 1(1987) 145-168
9. Fodor, J., Pylyshyn, Z.W.: *Connectionism and Cognitive Architecture: A Critical Analysis*. In: Boden, M.A. (ed.): *The Philosophy of Artificial Intelligence*, Oxford University Press (1990)
10. Barnden, J., Pollack, J. (eds.): *High-Level Connectionist Models*, Ablex: Norwood, NJ, (1990)
11. Pinker, S., Prince, A.: 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)
12. Jordan, M.I., Jacobs, R.A.: Hierarchical mixtures of experts and the EM algorithm. *MIT Comp. Cog. Sci. Tech. Report* 9301 (1993)
13. Shepherd, G.M.: *The Synaptic Organisation of the Brain*. Oxford University Press, London Toronto (1974)

14. Friston, K.J.: Transients, Metastability, and Neuronal Dynamics. *Neuroimage* 5 (1997) 164-171
15. Fodor, J.A.: *The Modularity of Mind*. MIT Press (1983)
16. Mumford, D.: 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)
17. Farah, M.: Neuropsychological inference with an interactive brain: A critique of the locality assumption. *Behavioural and Brain Sciences* (1993)
18. Wegner, P.: Why Interaction is More Powerful than Algorithms. *CACM* May (1997)
19. Koch, C.: Computation and the single neuron. *Nature* 385 (1997) 207-210
20. Barlow, H.: Intraneuronal information processing, directional selectivity and memory for spatio-temporal sequences. *Network: Computation in Neural Systems* 7 (1996) 251-259
21. Granger, R., 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
22. Thomson, A.M.: More Than Just Frequency Detectors ?. *Science* 275 Jan (1997) 179-180
23. Sejnowski, T.J.: Time for a new neural code ?. *Nature* 376 (1995) 21-22
24. Koenig, P., et al.: Integrator or coincidence detector? The role of the cortical neuron revisited. *Trends Neurosci.* 19(4) (1996) 130-137
25. Perret, D.I., et al.: Visual neurons responsive to faces in the monkey temporal cortex. *Experimental Brain Research* 47 (1982) 329-342
26. Rolls, E.T., Tovee, M.J.: Processing speed in the cerebral cortex and the neurophysiology of visual backward masking. *Proc. Roy. Soc. B* 257 (1994) 9-15
27. Thorpe, S.J., Imbert, M.: Biological constraints on connectionist modelling. In: Pfeifer, R., et al. (eds.): *Connectionism in Perspective*. Elsevier (1989)
28. Softky, W.R., Koch, Ch.: The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSP. *J. of Neurosci.* 13 (1993) 334-350
29. Berry, M. J., et al.: The structure and precision of retinal spike trains. *Proc. Natl. Acad. Sci. USA* 94 (1997) 5411-5416
30. Mainen, Z.F., Sejnowski, T.J.: Reliability of spike timing in neocortical neurons. *Science* 168 (1995) 1503-1506
31. Nelson, J.I.: Visual Scene Perception: Neurophysiology. In: Arbib, M.A. (ed.): *The Handbook of Brain Theory and Neural Networks*. MIT Press: Cambridge MA (1995)
32. Nelson, J.I.: Binding in the Visual System. In: Arbib, M.A. (Ed.): *The Handbook of Brain Theory and Neural Networks*, MIT Press, Cambridge MA (1995)
33. Brown, R.: *Social psychology*. Free Press, New York (1965)
34. Douglas, R.J., Martin, K.A.C.: Exploring cortical microcircuits. In: McKenna, Davis, Zornetzer, (eds.): *Single Neuron Computation*. Academic Press (1992)
35. Barlow, H.B.: Single units and sensation: A neuron doctrine for perceptual psychology?. *Perception* 1 371-394
36. Nasuto, S.J., Bishop, J.M.: Bivariate Processing with Spiking Neuron Stochastic Diffusion Search Network. *Neural Processing Letters* (at review).
37. Bishop, J.M.: Stochastic Searching Networks. *Proc. 1st IEE Conf. Artificial Neural Networks*, pp. 329-331, London (1989).
38. Nasuto, S.J., Bishop, J.M.: Convergence Analysis of a Stochastic Diffusion Search. *Parallel Algorithms and Applications* (in press).
39. Nasuto, S.J., Bishop, J.M., Lauria, S.: Time Complexity Analysis of Stochastic Diffusion Search, *Proc. Neural Computation Conf.*, Vienna, Austria (1998)

40. van Leeuwen, J. (ed.): Handbook of Theoretical Computer Science. MIT Press: Amsterdam (1990)
41. Treisman, A.: Features and Objects: The fourteenth Bartlett Memorial Lecture. *The Quarterly Journal of Experimental Psychology* 40A(2) (1998) 201-237
42. Cowey, A.: Cortical Visual Areas and the Neurobiology of Higher Visual Processes. In: Farah, M.J., Ratcliff, G.: *The neuropsychology of high-level vision*. (eds.): LEA Publishers (1994)
43. Spillmann, L., Werner, J.S.: Long range interactions in visual perception. *Trends Neurosci.* 19(10) (1996) 428-434
44. McCrone, J.: Wild minds. *New Scientist* 13 Dec (1997) 26-30