Spatiotemporal Neural Processes of Cognition: a review and proposal

Francis Jeanson
Carleton University

January 9th 2010

Introduction

Recent developments in cognitive science have emphasized the necessity to see cognition as a phenomenon which emerges from the ongoing interaction between the environment, body and brain (Brooks, 1991; Varela et al. 1991; Hutchins, 1995; Clancey, 1997; Clark, 1997). In part due to the unsatisfactory theoretical and experimental accounts given by computational approaches, this revival is also driven by the extraordinary advances made in the domain of related psychological and biological disciplines such as psycholinguistics, cognitive psychology, neuroscience, evolutionary and developmental sciences, etc. Indeed, by exploring the environmental and bodily processes that play a constitutive role in cognition our theories may sufficiently mature so to make visible the fundamental principles behind intelligent behaviour. In particular, it is increasingly apparent that the high complexity of the environmental/body/brain transaction will require a sufficiently sophisticated framework to explain both simple stimuli induced reactions or thoughts as well as complex continuous mindful behaviours. It should be useful then to adopt a program that sets focus on some of the fundamental solutions found by nature to accomplish this: neural coding.

The present paper aims to address two central questions. What are the known constitutive processes of the brain? How can we test their applicability to real world problems? From the pioneering histological observations made by Ramon-y-Cajal in the late 1800's to the advent of sophisticated brain imaging techniques we have gained a tremendous body of knowledge regarding the role played by the nervous system in cognition. These breakthroughs have stimulated a number of theoretical perspectives on the functional properties of brain structures. In particular early perspectives have postulated that neural integration of spike train frequencies is the central encoding mechanism of the brain (Hebb 1949; Barlow 1972). More recently alternative coding mechanisms have been postulated to take place as a result of intricate spatiotemporal properties of neural cell firings (Von der Malsburg 1981). Although increasing experimental evidence suggests that spatiotemporal coding takes place throughout the brain, our understanding of their precise nature is still poor. Furthermore, little direct behavioural evidence that such spatiotemporal codes are important for intelligent behaviour exists. In the following I will begin by reviewing the dominant rate coding paradigm as interpreted by the neurosciences. I will expose its strengths and weaknesses in order to motivate the idea that other brain mechanisms may be at work. The alternative spatiotemporal coding paradigm will then be introduced. This will be complemented by a brief exploration of the theoretical issues this alternative coding scheme may address as well as empirical evidence of its presence in the mammalian brain based on both biological and simulated results. A proposal for its investigation will then be developed. In particular, I will argue in favour of the evolutionary robotics approach as a framework suitable for the identification of spatiotemporal mechanisms and their involvement in embodied cognition.

Rate Coding and Beyond

Since early recordings of neurone activity the dominant view of information coding in the brain sees rates of neural firing as the sole neural mechanism underlying perception, thought and action. Typically experimental recordings have focussed on the activity of cells over several hundred millisecond time scales. This has made possible the correlating of statistical firing rates with psychological performances. This view has been strongly supported by work on motor neurone activity where firing rates were found to have a clear correspondence with muscle flexion (Robinson 1975; Shidara et al. 1993). Shidara et al. for instance have shown that the high frequency of firing of Purkinje cells in the cerebellum allows the cells to behave as integrators of spike numbers where firing rate modulations control eye positions and accelerations (Shidara et al. 1993). Two distinct perspectives take rate coding as the central means for information coding in the nervous system: *place coding* and *distributed coding*.

First, the place coding view sees information as being encoded by a single cell's mean spike rate over a particular psychological time during an experiment (Barlow 1972). This notion is characterized by the interpretation of cell function based on both their objective meaning: the spatial location of the firing cell, and subjective meaning: the rate of cell firing (Kruger and Becker 1991; Fujii et al. 1996). Here firing rates are averaged over the duration of psychological experiments which often last several hundred milliseconds. Furthermore, experiments are adjusted to elicit maximum firing rates of the recorded cells. From this view the functional role of a cell is inferred from an optimum stimulus: a stimulus that elicits a maximum average firing rate. Place coding thus motivates the hypothesis that subordinate cells transmit their signals to grandmother cells. These grandmother cells are thus understood to be feature specific and that they must play a central role in pattern detection or action initiation. Face recognition cells for instance have been hypothesized to reside in the inferotemporal (IT) lobe (Rolls et al. 1990). Also in IT an icon alphabet is claimed to be implemented where single cells respond selectively to perceived characters (Tanaka 1992). Although experimental recordings show reliable firing rates given specific stimuli, place coding remains restrictive regarding its view of the mode of operation of brain function. As noted by Fujii et al., averaging spike rates over entire experiments can make it impossible to detect intricate firing patterns (Fujii et al. 1996). For instance, temporal fluctuations in firing rates are commonly observed during psychological time scales. Yet statistical variations disappear when averages are computed over the time of an experiment. Furthermore, such studies typically average over many experiments to obtain significance and therefor can't explain how the brain solves problems from a single presentation of the stimulus. Kruger and Becker have provided strong evidence against the notion of the functional significance of mean firing rates over psychological times (Kruger and Becker 1991). While replicating the experimental findings of Hubel and Wiesel, which originally provided support to the place coding view with orientation-selective cells in the primary visual cortex (V1), Kruger and Becker found that spike averages over a 300ms time slot of a psychological experiment was insufficient to predict the activation of the recorded cells given a single stimulus (Hubel and Wiesel 1968). However, smaller time slots in the range of 20ms to 80ms could accurately do so with a performance of up to 80%, thus reinforcing the view that firing rate averages over long periods and many experiments does not capture the informational principle of the brain for cognition. More problematic to the place coding view is the assumption that grandmother cells exist for each representation of every object and each of those object's features. Such a hierarchical view thus implies that dedicated cells should be formed or tuned for each novel stimuli. According to von der Malsburg such a mechanism would lead to a combinatorial explosion where cardinal cells would have to exist or be created and respond preferentially to a specific stimulus by integrating information coming from all modalities (von der Malsburg 1981). The implausibility of such a mechanism and lack of evidence that such spatial specificity in the brain exists is unsupportive of the place coding view.

The second alternative to rate coding takes into account the combined effect of a large number of cells. The original concepts surrounding the distributed coding view emerged from Donald Hebb's cell assembly hypothesis (Hebb 1949). According to this hypothesis the fundamental brain code is based on the group activity of neural cells; the assembly serves to produce a mean firing rate and is identified by it; and assemblies are formed by the Hebbian learning rule: the simultaneous increase in firing rate of a set of cells (von der Malsburg 1981). This view sees the group activity of an ensemble as corresponding to a "symbol of the mind" and has served as the foundational principle to the connectionist approach to artificial intelligence (Rumelhart et al. 1986; Fujii et al. 1996). In contrast to the place coding view, distributed coding does not appeal to the single cell = single function theory. Here a cell may play more than one role by participating in the coding of multiple patterns. This suggests that a wide variety of inputs can be encoded in various subsets of a single assembly thereby escaping the issue of combinatorial explosion. However, cell assemblies are still understood to code stimuli using first order firing rate statistics. Because of this, the ability for such networks to detect variation between two inputs depends essentially on the change of firing rate of a cell assembly. This coding mechanism can thus lead to a "superposition catastrophe" and fails to account for the problem of binding stimuli (von der Malsburg 1981). Superposition for this view is problematic because it is impossible to distinguish two assemblies activated in the same area by two simultaneously presented stimuli if only the change in average firing rate is detectable from an assembly's activity. In other words, it is impossible to determine which cell belongs to which assembly if individual cells fire during the same period of time at arbitrary rates. Although grandmother cell hierarchies have been proposed to resolve this controversy, they do so at the cost of making stark claims regarding the architectural nature of the brain. Indeed, selective response to assembly activity by superordinate cells would require a hierarchical organization of knowledge within cell assemblies which goes against the Hebbian cell assembly hypothesis and brings us back to a place coding view (Fujii et al. 1996). Related to the superposition catastrophe, the binding problem also introduces difficulties to distributed processing views (von der Malsburg 1981). The binding problem can be briefly introduced by the following question: How is it that the shape and colour of objects are bound in a reliable way and not intermixed when more than one object is perceived? The visual receptive field can selectively respond in a distributed manner to input stimuli such that the colour green is detected and the colour red is detected simultaneously from parallel processes. It can also be the case that an apple is detected by a 'shape-detecting' assembly and a pickle in another. But distributed coding cannot account for a mechanism that can integrate from the simultaneous presentation to the receptive field the bond between red and apple and between green and pickle based on average firing rates alone. Because both red and green are active in the colour area and pickle and apple in the object area super-imposing cell assemblies are co-active and indistinguishable given that firing rates increase simultaneously. Finding how the brain solves this problem is important however because it is generally accepted that the binding of parallel informational fragments must take place to form an integrated whole for higher cognitive processing.

Underlying the problem of co-active populations lies the anatomical fact that neural assemblies are formed within static connections between neurones. However, selective binding of active patterns cannot be achieved by fixed connections because changes in stimuli elicit variation in the neural activity (Gerstner et al. 1997). Instead, Gerstner et al. suggest that the physically connected population should be seen as the super-set of possible connections that may potentially need to interact. Thus according to this view some dynamical mechanisms of neural coordination must exist which not only allow the formation of co-active populations within the subset of connected parts, but also elicit selective responses of either a target population or change within the active population based on the current firing pattern. Amongst the first to take note of such a requirement was von der Marlsburg who introduced the idea that subsets of neural populations within a connected network may contribute functionally distinct roles from other subsets through the synchronized firing of neurones[1] (von der Malsburg 1981). More generally, the notion that a second order statistics in neural populations codes may reveal alternate mechanisms of signal integration and cell recruitment has been gaining ground. These supplementary coding forms belong to the the family of spatiotemporal codes.

---

[1] This was central to von der Malsburg's Correlation Hypothesis (von der Malsburg 1981).

Spatiotemporal Coding

A prima facie the nervous system is essentially constituted by a dense population of neural cells interconnected via a sophisticated network of dendritic and axonal fibres which transmit electrical and chemical signals to each other. Of course these are tied to specialized sensory receptive cells, chemical releasing neurones and muscle contracting cells in order to allow an agent to perform intelligent actions. Because of this our understanding of a neural code is generally related to the notion of transforming signals or a state of activity into another signal or alternate state of activity within this neural population. There are numerous dimensions of complexity that a biological neural network can theoretically bare which have functional significance in this transformational process. For instance, the sheer number of cells as the sole fluctuating parameter in a network can be manipulated in such a way that a critical population size will give rise to a sudden "explosion" of spiking activity. Also, the complexity of the network topology can greatly constrain potential communication channels yet favour the development of specialized function as seen during development of the child's nervous system. Luciana et al. for instance observed a significant improvement of prefrontal working memory in children between the ages of 5 and 7 years (Luciana and Nelson 1998). Even more complexity can be found in the intricacy of individual cell qualities, such as synaptic facilitations, and the sophisticated network of heterogeneous cells that are formed from them. But fundamentally, change in network activity can arise from either change in the activity level of each cell or change in the set of active cells. Changes in cell activity suggests that integration via change in firing rate or signal amplitude is at work. Changes in the set of active cells, however, suggests that the mechanisms responsible for recruiting new cells or inhibiting active cells should be carefully considered. These mechanisms can either take the form of integration or coincidence detection (Abeles 1982).

Long considered the prime mechanism of neural coding, integration is a direct outcome of cell sensitivity to incoming firing rates. The frequency pulses that reach a neurone can either be integrated in the form of a firing count, a firing rate derivative, or a change in conductance amplitude proportional to the incoming firing rate depending on the kind of synapse being targeted[2] (Gerstner et al. 1997). Looking beyond single cell average firing rates over psyhcological times and cell assembly average firing rates, Kruger

---

[2] Two broad classes of synapses exist: those which facilitate transmission and those which depress it.

and Becker suggested that spatiotemporal integration could take place via modulation of firing rates over short time periods (Kruger and Becker 1991). As discussed earlier, their findings from recordings in V1 suggests that averaging firing rates over short term windows of about 20ms to 80ms allows the detection of variation in rate activity during an experiment and reliably infer stimulus type from a single set of cell activations. Change in cell activity through temporal modulation of the firing rate would then seem to take place as an effective means of integrating incoming signals. The lack of further experimental evidence for this coding mechanisms makes it unclear if rate modulation could account for the problem of superposition of stimuli or feature binding. Furthermore it does not refute the potential existence of complementary modes of coding.

As Fujii et al. note, when looking for the effect of stimuli on firing rates only, an experimenter may overlook the fine temporal structures that may arise from the correlated activity of cells (Fujii et al. 1996). The potential temporal and spatial patterns of excitation or inhibition of a neurone through synaptic or direct somatic connections is enormous. With over 5000 synaptic connections on average per cell, cortical neurones are subjected to a bombardment of signals often originating from more than one pre-synaptic cell. Activity of the individual synapses are thus likely to affect the overall response of the post-synaptic cell regardless of firing rate. The cooperative timing and location of these incoming signals may thus have a significant impact on the activation or inhibition of the target cell. Generally, this alternate spatiotemporal form can be termed *coincidence detection*. Here, a neurone's detection of a threshold number of quasi-simultaneous incoming signals will cause it to fire (Abeles 1982; Braitenberg 1988; Fujii et al. 1996). Hence, the phase relation of input spike trains are encoded by the coincidence detecting cell and determine whether an output spike will be emitted in turn. The are a number of interesting physiological simplifications that can arise from such a mechanism. For instance, a cell may be recruited (caused to fire) by many weak synaptic connections. Whereas rate coding requires that the incoming signal be transmitted through sufficiently strong synaptic connections, coincidence detectors require no minimal synaptic weights but can instead rely on the a large amount of coincident stimuli. Furthermore, the frequency of incoming signals has no direct effect on the target cell's firing pattern. Although it would appear that increased firing frequency may highten the chance of coincidence to take place, it seems likely that more finely tuned spike times can arise from appropriate circuitry. More importantly, the organization of networks of coincidence detecting cells can allow for rich dynamics while the physiological structure remains fixed. As Fujii et al. illustrate, a simple network (figure **??**) with two coincidence detection neurones A and B that are both excitable via channels p and q can exhibit varying dynamics depending on the temporal structure of the spike trains that are emitted from p and q (Fujii et al. 1996). For instance, it is possible that the superposition of spike trains from p and q when arriving at cell A and B are coincident and make both cells fire. However, if the phase of the spike trains p and q that reach A is different from the phase that reaches B it is possible to have A fire but not B or vice versa. Change in phase would occur from varying onset of the incoming spike train from p or q. This, for instance, can be accomplished by a difference in propagation delay between channels and cells. Hence, depending on the pattern of the spike trains p and q it is possible to dynamically reconfigure the state of firing of both cells A and B. This however can be accomplished in virtue of the spike firing patterns alone. No network reconfiguration or additional connection is required. However, if A and B were firing rate integrating neurones it would be impossible to have A fire but not B without adding or removing connections. In this scenario spike rates from p and q that are below threshold would trigger neither cell yet both would fire past a certain excitatory threshold. Hence in contrast to firing rate neural networks, coincidence detection networks allow for a dynamic recruitment or inhibition of cells.

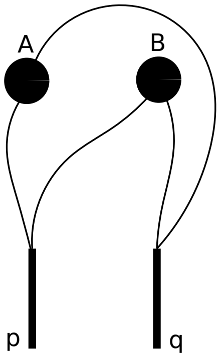

Figure 1: **Simple two cell assembly**

In virtue of this active recruitment it is possible to see how coincidence detection networks may allow neural assemblies to temporarily form by the recruitment of neurones into an active group while other connected cells remain 'dormant'. Whereas dormant cells would exhibit sub-threshold activity, recruited cells would play a role by emitting spikes within its assembly and/or to other regions. Importantly, the functional specificity of these networks could be transient implying that neurones may not only be recruited by more than one assembly but also be actively engaged in multiple assemblies during a same time period. This suggests that activity emanating from two distinct active assemblies may bind their informational content as long as they belong to a connected superset of cells. Two contending theories offer a spatiotemporal account to the superposition catastrophe and related binding problem based on coincidence detection. As mentioned earlier the superposition catastrophe arises from the difficulty to distinguish two or more stimulated assemblies within the same receptive field.

The first spatiotemporal approach, and most widely accepted, suggests that oscillatory synchronous firing of cells can help distinguish groups of active cells within a population (Singer 1994). If cells in assembly M fire synchronously as do cells in assembly N but with different rates from one another, then the identification of the corresponding stimuli within the same receptive field is possible. Indeed, synchronous activity within numerous areas of the brain have been identified experimentally (Gray and Singer 1987; Eckhorn et al. 1988; Gray et al. 1990; Engel et al. 1991; Murthy and Fetz 1992). Furthermore, the long lasting debates regarding the origin of these oscillations where conciliated by realistic neural simulations. Pauluis et al. for instance, demonstrated that synchrony could emerge from networks of cortical cells alone (Pauluis et al. 1999). Similarly, synchronous group activity is claimed to resolve the binding problem (Kruger 1991; Gochin et al. 1994). Here colours and shapes are thought to be bound by inter-areal synchrony of the assemblies in visual areas via their connected pathways.

Alternatively, a second approach maintains that large scale synchronization is an unlikely mechanism due to the fast transient nature of recorded synchrony (Gray et al. 1992; Kreiter and Singer 1992; Fujii et al. 1996). Gray et al. showed that synchrony between two cortical areas 7mm apart lasted no more than 100-200ms at a time interspersed with asynchronous states. Instead, this view known as the Dynamical Cell Assembly Hypothesis (Fujii et al. 1996) maintains that assemblies of coincident detecting cells can form sufficiently complex circuits based on four propositions: (1) the network structure and spatiotemporal nature of the input spikes spontaneously and dynamically organize a cell assembly; (2) dynamic cell assemblies can be identified by their cross-correlation; (3) dynamic cell assemblies can serve as detectors of spatiotemporal patterns; (4) multiple cell assemblies can be instantiated within the same neuronal pool. It is important to note that this hypothesis accepts oscillatory synchrony as a special case of cell assembly dynamics because both are founded on the principle of coincidence detection. However, the superposition catastrophe and binding problem are indirectly offered a solution via proposition 3 and 4 without recall to synchrony. Proposition 4 states that assemblies within the same receptive field can be instantiated, whereas proposition 3 suggests that firing patterns from incoming pathways can be detected as preferred patterns through a particular active configuration of an assembly within the target population. This view thus offers an alternate perspective which states that features need not be bound by clearly identifiable first order properties such as synchronous firing rates, cell assembly hierarchies, or grandmother cells but instead that binding and identification may be implicit to the fine temporal structures of a cell assembly.

Hence a better understanding of the functional role played by these two spatiotemporal schemes still needs to be developed. Indeed, since the discovery that oscillatory activity can emerge within cortical networks independently of subcortical structures there has been increasing agreement that synchronous oscillatory activity must play a fundamental role in cognitive function (Pauluis et al. 1999; Buszaki 2006). However, recent developments tend to reaffirm the fundamental notions behind dynamic cell assembly. In particular, Izhikevich's notion of polychronous coding stems from a body of experimental simulations which suggests that complex firing patterns based on propagation delays between cells can lead to multiple temporal codes within the same population (Izhikevich 2006). Taken together these efforts support the prospect of developing a further understanding of the kinds of fine temporal structures that can arise in the brain. Historically, this effort has been motivated by experimental findings in neural cell recordings but the improved modelling techniques and computing power in more recent years has lead to a flourishing new research field engaged in the realistic modelling and simulation of neural networks.

From In-Vivo to In-Silico Experimentation

Early experimental work on the neural basis of cognition has focussed mainly on the firing rate averages of spike trains and neural assemblies. However, new analytical techniques have emerged which provide either indirect or direct evidence of the presence of a spatiotemporal code in the brain. Although these techniques have been originally developed to study anatomical recordings, they are becoming increasingly applied to simulated results which are now not only obtainable for remarkably accurate cell models but for group dynamics as well.

In 1968 Calvin and Stevens measured rate-normalized coefficients of lumbosacral motorneurons of cats. Computing the coefficient variation C_v revealed low variation between 0.05 and 0.1 indicating that firing rates in motorneurons are reliable and not random (Calvin and Stevens 1968). However, when measuring coefficients in V1 and MT of awake monkeys, Softky and Koch recorded high C_v variations between 0.5 and 1.0 suggesting that in these areas firing is similar to a random Poisson process (Softky and Koch 1993). According to Softky and Koch the irregularities of spike trains in cortical regions supports the idea that firing rates are less likely to be responsible for information coding in these areas. Instead, they suggest that there must exist a high degree of nonlinear and fast dendritic signals (the active dendrite hypothesis) or that a strong synchronization between synaptic events takes place (the effective coincidence detector hypothesis) (Softky 1994; Softky 1995). Both of which suggest that a spatiotemporal component is at play. Although alternative explanations for these irregularities based on firing rate neurones have been proposed, Softky showed through single cell simulation how a coincidence detection mechanism could produce such irregularities (Shadlen and Newsome 1994; Usher et al. 1994; Softky 1995).

As discussed earlier, the presence in the brain of fine temporal structures within milliseconds of activity would provide direct evidence for the existence of a spatiotemporal mode of information coding. Traditionally, experimenters who seek to detect average firing rates record cell activity and compute their rates using peri-stimulus time histograms (PSTH). In 1992, Ahissar et al. applied this technique when recording the activity of two cells in the auditory cortex (Ahissar et al. 1992). When presented with sounds moving from either left to right or right to left both cells began firing at high frequencies regardless of the direction of the sound. However, after computing the cross-correlogram of the same two spike trains, they found that the cells would fire together when the sound came from one direction, but avoid firing together when the sound came from the other direction. This direction selectivity is thus made evident by examining the precise temporal onsets of spike trains which would otherwise be ignored from looking at firing rates alone.

Furthermore, direct evidence for the special case of coherent oscillations was made in the late 1980's. By making local field potential (LFP) recordings in visual cortices a number of researchers have shown through computed auto-correlograms that synchronization of oscillating firing neurones can occur across distant sites (Gray and Singer 1987; Eckhorn et al. 1988; Gray et al. 1989; Engel et al. 1991). Similarly, coherent oscillations have been attributed to sensorimotor integration in attentional tasks in awake monkeys (Murthy and Fetz 1992). These findings, however, could not confirm whether synchronous activity could emerge independently from subcortical nuclei. In particular, a long standing debate as to whether specific thalamic circuits could be responsible for coherent oscillations in the cortices could not be resolved without a method that would examine cortical circuits dissociated from other areas. Pauluis et al. managed to demonstrate through simulations based on realistic properties of cortical neurones that synchronized oscillations can in fact emerge from cortical circuitry alone (Pauluis et al. 1999). In addition their work emphasized the importance of inhibition and connection delays for these oscillations to occur.

Increasingly the role played by computer simulation in neural modelling has offered powerful means to verify functional properties of cell and networks of cells based on anatomical features. In parallel, artificial neural network paradigms have also developed over the years with the purpose of understanding the sub-symbolic functional principles of cognition. In its early beginnings, the connectionist approach of the 1980's made significant breakthroughs regarding the possibility of constructing neural models for learning, speech perception, spatial recognition, memory, etc. (Rumelhart et al. 1986). What I aim to motivate here though is that an adequate approach to neural modelling which can reveal the relevance of various coding mechanisms to intelligent behaviour must be developed. Strict connectionism, however, favours abstract signal integrating neurone models organized in complex network architectures that follow Hebb's cell assembly hypothesis. This aversion to alternate coding schemes and the specific model abstractions made prevent this approach from accounting for the fine temporal nature of firing patterns within a complex behavioural context.

As an alternative, a special class of recurrent neural networks named continuous time recurrent neural networks (CTRNN) were developed so to enable networks to exhibit a rich set of dynamics (Beer 1990). Based on traditional connectionist integrating neurones CTRNNs possess an activation decay that requires cell activity to be integrated over time. Although, seemingly minimal in change, recurrent connections of this sort can give rise to highly sophisticated dynamics. Beer has in fact shown that small circuits can exhibit stable, cyclic and even chaotic attractors (Beer 1995). The amazing richness in dynamical diversity that even simple CTRNNs can generate has stimulated a significant enthusiasm from dynamical systems approaches to cognition. This network paradigm has been successfully employed for the design of motor control, agent behavioural adaptation, and dynamic memory modelling to name a few (Cliff et al. 1993; Beer 1996; di Paolo 2000). However, despite recurrent connections and the use of an activation decay, it is not clear how CTRNNs could take advantage of sophisticated temporal dynamics based on coincidence detection. First, activation decays seem to be insufficient for the generation of sophisticated coincidence timing patterns. Even if decays were set independently for each artificial neurone, this would only alter the speed at which a cell's signal is reduced. At most, the timing at which a cell would start or stop emitting a passing threshold signal could possibly be detected. Second, the activation decay affects all output signals emitted by the neurone. This suggests that all receptive neurons will receive simultaneous stimulation from the source neurons thereby reducing the landscape of possible spatiotemporal patterns proportionally to the average number of connections of each neurone. Hence, although CTRNNs offer appealing temporal properties for dynamics at the scale of an entire network, the lack of a temporal mechanism at the level of individual connections suggests a large number of spatiotemporal patterns cannot be encoded by such networks.

The failure of traditional artificial neural network paradigms to account for spatiotemporal coding urges that a more careful look should be taken to biological processes of the nervous system so to identify the fundamental criteria behind these coding forms. As suggested by Fujii et al. and from the simulation results of Pauluis et al., signal delays between individual cells could indeed play a central role in the tuning of fine temporal structures (Fujii et al. 1996; Pauluis et al. 1999). By taking the neurophysiological properties of real networks into consideration, the conduction delays between cells should play a fundamental role in spatiotemporal coding based on coincidence detection. Beyond connection delays, however, it remains unclear whether further temporal components of cell may functionally impact their ability to encode information in sufficiently diverse ways. As seen with CTRNNs, the activation decay of cells may potentially play a secondary role in the temporal adjustment of firing patterns. In realistic Hodgkin-Huxley based cell models, membrane time constants, compartment conduction times, spike rise times, and refractory periods could all potentially play an important role in this matter. However if the population of neurons remains mostly homogeneous in terms of individual cell properties, it is fair to claim that connection delays remain the only functionally differentiating factor that could affect fine temporal structures. Exploring connectivity delays should thus constitute an important first step for the experimental exploration of spatiotemporal coding for cognition.

Spatiotemporal Codes for Intelligent Action

As presented in the earlier sections spatiotemporal neural activity appears to play an important part in many brain areas. In parallel, the theoretical implication of this sort of activity in the solution to the superposition catastrophe and binding problem has triggered the development of numerous theories regarding its role in binding visual features, spatial orientation of sound, sensorimotor coordination, memory, and more. Although experimental work has been successful at identifying the particular cellular loci involved in some of these cognitive phenomena, their functional specifications still remain largely unanswered. That is, even if objective experimental observations of neural responses can detect cells and their assembly with respect to a set of stimuli, it remains extremely difficult, if not impossible, to construct a cohesive view of the subjective meaning of these localized encodings as they appear to the rest of the network. In addition, the emergent and fast transient qualities of spatiotemporal codes observed in experiments and simulations of various neural regions responsible for divers cognitive capabilities suggests that such patterns of activity elicit complex self-organizing mechanisms. By appealing to an embodied approach for the investigation of the informational content of networks, it should become feasible to identify the functional and informational qualities of minimal-to-complex circuits with respect to perception, cognition and action (Brooks 1991; Varela et al. 1991; Pfeifer and Scheier 1999). In particular, this strategy should reveal the degree to which spatiotemporal codes play a role in sense-making for intelligent action.

To accomplish this, a proper level of network implementation will need to be devised. From the discussion in previous sections, it follows that an effective coincidence detection neural network will at a minimum need to meet the following criteria:

Cells must fire only under a certain threshold of activation. Without a firing threshold, activation values are susceptible of recruiting target neurons continuously. This would prevent the formation of any distinct cell assembly which goes against the original hypothesis. To accomplish this, nonlinear activation functions such as a logistic function in a connectionist based model could be used. More biologically realistic spiking neural network models that take into account cell membrane capacitance as a threshold value would especially suit this criteria. Finally, it may even be possible to obtain complex coding schemes with very simple binary activation artificial neurons. Here for instance an active signal would be represented by the numeric value 1 and an inactive signal by 0. But this should not suggests that a single input will be sufficient to activate a target cell. On the contrary, an adequate integrative mechanism of input signals will be required regardless of the type of activation.

Connections between cells must propagate with delays. This is arguably the central criteria for sophisticated temporal codes to occur. Although I've argued in the previous section that other physiological factors may affect the temporality of signals, it would seem that connection delays are the most important candidates for the coincidence of spikes to occur in temporally interesting ways.

By starting from these minimal criteria, investigating the types of architectures that can effectively exploit spatiotemporal patterns should help set the stage for answering a number of pending questions, including: How large should an assembly be to encode functional patterns? If activity of a cell is affected by both synaptic weights and the number of incoming neurones are there any organizational principles that determine their relative importance based on function? Do 'families' of patterns exist for particular spatiotemporal purposes? How much information can be superimposed onto the same population of cells? etc...

Because the approach that I propose here suggests that embodied agent modelling should provide a means to answer these questions, it also suggests that variably complex networks may need to be generated to satisfy the embodied task. Because the total number of parameters which serve to define the set of parameters of a simple neural network can be large, finding a proper combination of settings can quickly become an arduous task. To resolve this issue, evolutionary robotics suggests that encoding network parameters into an artificial genotype can facilitate the selection of fit network configurations (Nolfi and Floreano 2000).

Evolutionary and Developmental Perspectives

Discovering network characteristics that exhibit finely tuned temporal structures in embodied tasks may be realizable from an evolutionary robotic standpoint. This is accomplished by evolving a population of genotypes that encode complete robot controllers through a selection process based on phenotypic and/or behavioural criteria. Computationally, this can be realized using a genetic algorithm (Holland 1975). By doing so, finding effective connectivity patterns, delay values, synaptic weights, etc. becomes possible without the necessity of traditional learning algorithms such as backpropagation (Rumelhart et al. 1986), reinforcement learning (Barto et al.

1995), classifier systems (Booker et al. 1989), self-organizing maps (Kohonen 1982), etc. As Nolfi and Floreano indicate, an evolutionary approach differs in two ways from other learning approaches. First, it demands less supervision than other algorithms by indirectly specifying adequate parametrization via general task fitness criteria. Second, it places less constraints on the self-organizing process because it allows the experimenter to encode any model parameters within the genotype (Nolfi and Floreano 2000 p. 4).

Significantly, this self-organizing characteristic is particularly important to spatiotemporal networks. If connection delays are to be encoded in the genome this could cause the genetic search space to be much too large for adequate solutions to be found within a reasonable amount of generations. Worse, the lack of compactness of the genome could affect the result of the evolutionary process (Nolfi and Floreano 2000 p. 225). Indeed, the addition of a single gene results in the exponential increase of the genetic search space. To avoid this problem one could encode individual cell properties and synaptic weights in the genotype while keeping delays out of the genotype. Instead, delays could be generated via some heuristic method based on network topology. Although some interesting temporal patterns may emerge from this technique, no fine tuning of the delays with respect to behavioural fitness can take place. Alternatively it may be of interest to explore the potential of indirect genotype-to-phenotype transcription. Nolfi et al. for instance, encoded within a genotype the instructions for axonal growth, synaptic weights, and branching of artificial neurons for robot control (Nolfi et al. 1994). By employing similar kinds of developmental processes it becomes possible to maintain genotype compactness, promote the expressive power of the genome, and improve evolvability (Nolfi and Floreano 2000 p. 225-226). Hence, by encoding growth rules such as reaction diffusion properties of chemical axons or more simple topological distributions the tuning of connection delays becomes possible while evolutionary efficient. This could result in precisely tuned and effective circuits for spatiotemporal coding in cognitive and behavioural tasks. Furthermore, this may open a new field of exploration concerned with the evolutionary and developmental traits of biological nervous systems from an embodied cognition standpoint.

Now that an overall methodology has been outlined, a central challenge still remains to be addressed. What are the potential behavioural or cognitive tasks that could elicit neural coding that is spatiotemporal in nature?

Finding Spatiotemporal Tasks

An evolutionary robotics approach to neural controller design requires minimal supervision, yet a selection criteria that can distinguish fit individuals from unfit ones must be still defined. But because the project here requires the identification of behaviours that are both efficient and prompt spatiotemporal coding, the constraints on task selection are twofolds. The evolved agent's neural controller must not only be capable of finding solutions to a task but do so while exhibiting fine temporal patterns of neural activity. Three different sources of theoretical and empirical work can be drawn upon for the identification of suitable tasks.

Neuroscientific recordings of neural response to stimuli have a traditional appeal. Empirical results on behavioural tasks that exhibit fine temporal patterns and synchronized oscillatory activity can indicate a general problem domain. Murthy and Fetz for instance, demonstrated how coherent gamma oscillations appeared in the somatosensory and somatomotor areas during a sensorimotor task where monkeys would have to search through tactile exploration for raisins placed in a container outside their visual field (Murthy and Fetz 1992). But when raisins where present in the visual field no oscillations where recorded. This suggests that this complex search procedure placed particular demand on the brain in which temporal codes were elicited. Although the complex cognitive apparatus of the ape which allows it to search and grasp hidden grapes is certainly not directly approachable from an embodied modelling standpoint, it may be possible to devise a functionally equivalent but simplified task for which neural controllers could be evolved.

The extensive body of work achieved by evolutionary robotics itself is an obvious source of embodied cognitive tasks. Indeed, reactive behaviours such as obstacle avoidance, foraging, etc. to more complex problem solving such as perceptual aliasing, sensory ambiguity, sequential organization of sub-behaviours, etc. can constitute a wealth of tasks that are directly amenable to embodied investigation. These tasks, however, have already proven themselves as solvable by non spatiotemporal coding means of the sort I wish to investigate. It may be possible though that by evolving networks in which coincidence detection is permitted to occur more efficient solutions to some of these problems will be found. In particular, it may be the case that spatiotemporal codes could significantly contribute to problems requiring complex internal dynamics and modular architectures. Fortunately, sample work in this area using traditional coding schemes has already been completed which could inspire novel tasks (Nolfi 1997).

Finally, work in cognitive neuroscience may not only help identify suitable embodied tasks but it also addresses those kinds of tasks that may provide a more integrative view of the role of spatiotemporal coding in complex embodied cognition. For instance, the proposal by LeDoux of the existence of a dual-pathway in the mammalian nervous system for rapid behavioural switching has stimulated work in agent simulations to demonstrate the functional plausibility of such networks for fast switching between predator evasion and food foraging (LeDoux 1986; den Dulk et al. 2003; Heerebout and Phaf 2009). In fact, by altering the orginal architecture proposed by den Dulk et al., Heerbout and Phaf obtained individuals that exploited oscillatory activation resulting in increased fitness. Originally, they intended to implement a memory component via recurrent connections to the hidden layer of a simple feedforward network. However, by enforcing recurrent signals to propagate with single step time delays they obtained robust stimuli induced oscillations within the network. This activity allowed agents to significantly increase the speed of their response to predator or food to the point of nearly doubling their fitness score (Heerebout and Phaf 2009). This work emphasizes the potential significance of delayed recurrent connections that are both excitatory and inhibitory for oscillatory activity in simple neural networks. Expanding on work of this nature by introducing new architectural or environmental complexities could reveal the limitations and requirements of complex spatiotemporal activity.

Overall, existing empirical work supports the notion that spatiotemporal codes are located in diverse areas of the nervous system. In addition theoretical work suggests that this coding form should apply to a number of tasks. Hence, categories of cognitive and behavioural processes should be identifiable given an appropriate experimental approach.

Conclusion

The prospect of discovering the cognitive implications of spatiotemporal codes in embodied tasks should help reveal the nature of their informational value and the degree to which they can participate in complex cognitive processes for intelligent action. This paper aimed to address the empirical and theoretical context in which this coding scheme can be vindicated. By first introducing the role of rate coding paradigms a better understanding of their limitations and scope was brought forth. This established an appropriate context for the introduction of spatiotemporal coding as a complementary form of neural processing. In particular, their theoretical plausibility was defended, their manifestation in both dynamical cell assembly and synchronized versions was introduced, and their potential ability to address the superposition catastrophe and binding problem was developed. Important examples of their empirical validity was subsequently given. Increasingly, however, it was shown that simulations of both accurate biological and artificial network models can provide important insights regarding the parametric conditions that favour this coding scheme. An argument in favour of gaining a better understanding of their informational content via embodied modelling was supplemented. In particular an evolutionary approach is favoured whereby the self-organization of spatiotemporal networks via natural selection could not only give rise to yet unforeseen applications of coincidence detection temporal codes, but also validate their use within complex environmental tasks. Finally, the issue of identifying adequate tasks was brought forth. This difficulty, however, is offset by the existence of a large body of empirical work exploring spatiotemporal codes on the one side, and embodied cognitive tasks on the other. By building upon these efforts I hope that insights will be made regarding the role that self-organizing neural mechanisms bare on the cognitive processes of intelligent behaviour.

References

Abeles, M. (1982). Role of the cortical neuron: integrator or coincidence detector? Israel Journal of Medical Science, 18, 83-92.

Ahissar, M., Ahissar, E., Bergman, H., Vaadia, E. (1992). Encoding of sound-source location and movement: Activity of single neurons and interactions between adjacent neurons in the monkey auditory cortex. Journal of Neuroscience, 67, 203-215.

Barlow, H. B. (1972). Single units and sensatioperceptual psychology? Perception, 1, 371-394.

Barto, A. G., Bradtke S. J., Singh, S. P. (1995). Learning to act using real-time dynamic programming. Artificial Intelligence, 6:105-122.

Beer, R.D. (1990). Intelligence as Adaptive Behavior: An Experiment in Computational. Neuroethology. Academic Press.

Beer, R.D. (1995). On the dynamics of small continuous-time recurrent neural networks. Adaptive Behavior 3(4):469-509.

Beer, R. D. (1996). Toward the evolution of dynamical neural networks for minimal cognitive behavior. From Animals to Animats 4: Proceedings of the 4th international conference on Simulation of Adaptive Behavior. Cambridge, MA: MIT Press, p.421-429.

Booker, L., Godberg, D. E., Holland, J. H., (1989). Classifier systems and genetic algorithms. Artificial Intelligence, 40:235-282.

Breitenberg, V. (1988). Doscourse on current issues in multi-neuron studies. In A. M. H. J. Aertsen, W. von Seelen, G., Shaw & U. M. Leinhos (Eds.), Organization of neural networks. Weinheim: VCH.

Brooks, R. A. (1991). Intelligence without representation. Artificial Intelligence, 47, 139-159.

Buszaki, G. (2006). Rhythms of the Brain. Oxford University Press.

Calvin, W., Stevens, C. (1968). Synaptic noise and other sources of randomness in motoneuron interspike intervals. Journal of Neurophysiology, 31, 574-587.

Clancey, W. J. (1997). Situated Cognition: on Human Knowledge and Com- puter Representations. Cambridge University Press.

Cliff, D., I. Harvey and P. Husbands (1993). Explorations in evolutionary robotics. Adaptive Behavior 2(1): 73-110.

Di Paolo, E. A. (2000). Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions. In Meyer, J.-A., Berthoz, A., Floreano, D., Roitblat, H., and Wilson, S. (Eds.), From Animals to Animats 6: Proceedings of the Sixth International Conference on the Simulation of Adaptive Behavior Paris, France. Cambridge MA: MIT Press.

Den Dulk, P., Heerebout, B. T., Phaf, R. H. (2003). A computational study into the evolution of dual-route dynamics for affective processing. Journal of Cognitive Neuroscience, 15, 194-208.

Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., Reiboek, H. (1988). Coherent oscillations: A mechanism for feature linking in the visual cortex?. Biol. Cybern. 60: 121-130.

Engel, A.K, Kreiter, A.K., Konig, P., Singer, W. (1991). Synchronization of oscillatory neuronal responses between striate and extrastriate visual cortical areas of the cat. Proc. Natl. Acad. Sci. USA 88: 6048-6052.

Fujii, H., Ito, H., Aihara, K., Ichinose, N., Tsukada, M. (1996). Dynamical Cell Assembly Hypothesis - Theoretical Possibility and Spatio-temporal Coding in the Cortex. Neural Networks, 9, 8:1303-1350.

Gerstner, W., Kreiter, A., Markram, H., Herz, A. V. M. (1997). Neural codes: Firing rates and beyond. Proc. Natl. Acad. Sci. USA, 94, 12740-12741.

Gray, C. M., Singer, W. (1987) Stimulus-specific neuronal oscillations in the cat visual cortex: A cortical functional unit. Society of Neuroscience Abstracts. 13, 403.3.

Gray, C. M., Engel, A. K., Konig, P., Singer, W. (1990). Stimulus-dependent neuronal oscillations in cat visual cortex: receptive field properties and feature dependence. European Journal of Neuroscience, 2, 607-619.

Gray, C. M., Engel, A. K., Konig, P., Singer, W. (1992). Synchronization of oscillatory neuronal responses in cat striate complex: Temporal properties. Visual Neurosciences, 8, 337-347.

Gochin, P. M., Colombo, M., Dorfman, G. A., Gerstein, G. L., Gross, C. G. (1994). Neural ensemble coding in inferior temporal cortex. Journal of Neurophysiology, 71, 2325-2337.

Hebb, D.O. (1949). Organization of behavior. New York: Wiley.

Heerebout B. T., Phaf, R. H. (2009). Emergent Oscillations in Evolutionary Simulations: Oscillating Networks Increase Switching Efficacy. Journal of Cognitive Neuroscience. MIT Press, p. 1-17 (doi: 10.1162/jocn.2009.21205)

Holland, J. H. (1975). Adaptation in Natural and Artificial Systems. Ann Arbor, MI: University of Michigan Press.

Hubel, D. H. Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. Journal of Physiology, 195, 215-243.

Hutchins, E. (1995). Cognition in the wild. Cambridge, MA, USA: MIT Press.

Izhikevich, E. M. (2006). Polychronization: Computation with spikes, Neural Computation. 18, 245-282.

Kreiter, A. K., Singer, W. (1992). Oscillatory neuronal responses in the visual cortex of the awake macaque monkey, European Journal of Neuroscience, 4, 369-375.

Kruger, A. K., Becker, J. D. (1991). Recognising the visual stimulus from neuronal discharges. Trends in Neuroscience, 14, 282-285.

Kruger, J. (1991). Spike train correlations on slow scales in monkey visual cortex. In J. Kruger (Eds.), Neural cooperativity. Berlin: Springer-Verlag.

Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. Biological Cybernetics, 43:59-69.

LeDoux, J. E. (1986). Sensory systems and emotion: A model of affective processing. Integrative Psychiatry, 4, 237-248.

Luciana, M., Nelson, C. A. (1998). The functional emergence of prefrontally-guided working memory systems in four- to eight-year-old children. Neuropsychologia 36: 273-93.

Murthy, N. V., Fetz, E. E., (1992). Coherent 25- to 35-Hz oscillations in the sensorimotor cortex of awake behaving monkeys. Proc. Natl. Acad. Sci. Neurobiology. 89: 5670-5674.

Nolfi, S., Floreano, D. (2000). Evolutionary Robotics: the biology, intelligence, and technology of self-organizing machines. MIT Press.

Nolfi, S., Miglino, O., Parisi, D. (1994). Phenotypic Plasticity in Evolving Neural Networks, in: D. P. Gaussier and J-D. Nicoud (eds.) Proceedings of the Intl. Conf. From Perception to Action, Los Alamitos, CA: IEEE Press.

Nolfi, S., Parisi, D., (1997). Learning to adapt to changing environments in evolving neural networks. Adaptive Behavior, 5: 99-105.

Pauluis, Q., Baker, N. S., Olivier, E. (1999). Emergent oscillations in a realistic network: The role of inhibition and the effect of the spatiotemporal distribution of the input. Journal of Com- putational Neuroscience, 6: 25-48.

Pfeifer, R., Scheier, C. (1999). Understanding Intelligence. Cambridge, MA: MIT Press.

Robinson, D. A. (1975). Oculomotor control signals. In P. Bach-y-Rita, & G. Lennerstrand (Eds.), Basic mechanisms of ocular motility and their clinical implications. Oxford: Pergamon Press, 23, 337-374.

Rolls, E. T., Treves, A. (1990). The relative advantage of sparse versus distributed encoding for associatie neuronal networks in the brain. Network, 1, 407-421.

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

Singer, W. (1994). Putative functions of temporal correlations in neocortical processing. In C. Koch, & J.L. Davis (Eds.), Large scale neuronal theories of the brain. Cambridge, MA:MIT Press.

Shalden, M. N., Newsome, W. T. (1994). Noise neural codes and cortical organization. Current Opinion in Neurobiology, 4, 569-579.

Shidara, M. Kawano, K. Gomi, H., Kawato, M. (1993). Inverse-dynamics model eye movement control by Purkinje cells in the cerebellum. Nature, 365, 50-52.

Softky, M. N., (1994). Sub-millisecond coincidence detection in active dendritic trees. Neuroscience, 13, 13-41.

Softky, W. R. (1995). Simple codes versus efficient codes. Current Opinion in Neurobiology, 5, 239-247.

Softky, W. R., Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random ESPSs. Journal of Neuroscience, 13, 334-350.

Tanaka, K. (1992). Inferotemporal cortex and higher visual functions. Current Opinion in Neurobiology, 2, 502-505.

Usher, M., Stemmler, M., Koch, C., Olami, Z. (1994). Network amplification of local fluctuations causes high spike rate variability, fractal firing patterns and oscillatory local field potentials. Neural Computation, 6, 795-836.

Varela, F.J., Thompson, E., Rosch, E. (1991). The embodied mind: Cognitive science and human experience. MIT Press, Cambridge, MA, USA.

Von der Marlsburg, C. (1981). The correlation theory of brain function. (Internal Report 81-2). Goettingen: Max-PLanck-Institute for Biophysical Chemistry.