

# **Temporal codes and sparse representations : a key to understanding rapid processing in the visual system**

Guyonneau R, VanRullen R and Thorpe SJ

**Where neural information processing is concerned, there is no debate about the fact that spikes are the basic currency for transmitting information between neurons. *How* the brain actually uses them to encode information remains more controversial. It is commonly assumed that neuronal firing rate is the key variable, but the speed with which images can be analysed in the visual system poses a major challenge for rate-based approaches. We will thus expose here the possibility that the brain makes use of the spatio-temporal structure of spike patterns to encode information. We then consider how selective neural responses can be generated so rapidly through spike-timing-dependent plasticity (STDP) and how these selectivities can be used for visual representation and recognition. Finally, we show how temporal codes and sparse representations may very well arise one from another and explain some of the remarkable features of processing in the visual system.**

Keywords: visual processing, temporal codes, distributed representations, STDP.

## **1. Introduction**

Ever since Adrian demonstrated in the 1920's that information about the outside world enters the nervous system in the form of series of action potentials, spiking activity has been thought of as the primary form of neural communication. His observations indicated that the nervous system uses firing rates to encode information because as intensity of the stimulus was increased, sensory neurons increased the number of action potentials they emitted (Adrian, 1928). Since then, rate-based codes have, for example, been demonstrated to play a role in the perception of motion direction (Salzman, 1990); they can also constitute the relevant variable used in observing brain phenomena (Roelfsema, 1998 for an experimental example in visual attention) and investigating theoretical hypotheses (Song et al., 2000). But, more importantly, rate codes have dominated the way in which we think about the brain (Barlow, 1972). In addition to being a simple and practical means of measuring neural activity, a reason behind this state resides in that spikes do not seem to be precise enough to carry information other than through their mere number (Perkel & Bullock, 1968; Softky & Koch, 1993; Shadlen & Newsome, 1994).

Yet, a growing number of experimental studies have questioned this postulate, at least where sensory modalities are concerned: spike times can indeed be reproducible, meaning that presenting the organism with the same stimulus elicits spikes with the same precise timing, in the order of the millisecond (Mainen & Sejnowski, 1995 for a biophysical model study; Berry et al., 1997 for experimental evidence; a review in VanRullen et al., 2004). It thus seems that reliable information can be encoded in the timing of individual spikes, possibly using the degree of synchrony across subsets of neurons (Abeles, 1991; Singer, 1999). Another possibility has to do with the time-to-first-spike (Thorpe, 1990), a theoretical proposal which has recently received experimental support: Johansson and Birznieks (2004) reported that on the basis of the timing of the first spikes generated in the median nerve within the upper arm, one can infer with reasonable certitude which stimulus was present in the environment, in that case the direction of the force applied to the fingertip. Although it still remains to be

experimentally proven that the brain actually uses a coding scheme based on the fine spatio-temporal structure of activity, spike times – and especially *first* spike times – nonetheless encode more information than firing rates (Gautrais & Thorpe, 1998; Petersen et al., 2001 for an experimental proof).

There are also cases where rate-based models simply have a hard time explaining some remarkable aspects of brain performances, namely those in which speed is of the essence. Experimental studies of neurons in various parts of the the monkey brain have demonstrated selective responses to complex visual stimuli such as faces, food and familiar 3D objects only 100 to 150 ms after stimulus onset (Perrett et al., 1982; Rolls et al., 1979; Rolls et al., 1982). Similarly, Ultra Rapid Visual Categorization (URVC) experiments show that a complex natural scene can be categorized based on the presence or absence of a target object (animal, vehicle) in around 150 ms (Thorpe et al., 1996; VanRullen and Thorpe, 2001a; VanRullen and Thorpe, 2001b) and even more rapidly in monkeys (Fabre-Thorpe et al., 1998). Such temporal constraints pose a major challenge to theories of coding if we are to account for rapid processing in the visual system. In order to reach the inferior temporal cortex (IT) where high-level visual descriptions are represented, information about the stimulus would need to cross something like 10 neuronal processing stages on the way from retinal photoreceptors. That leaves only about 10 milliseconds of processing at each stage, a small temporal window which comprises synaptic delays (~1ms), post-synaptic integration, spike generation and propagation to the next stage. These propagation times may be considerable because intracortical conduction velocities may often be limited to only 1-2 m.s<sup>-1</sup> (Bullier et al., 1988). Furthermore, firing rates in the visual system typically do not exceed 100Hz implying that, on average, few cells can fire more than 1 spike in the critical 10 millisecond window. It follows that neurons only have time to fire a single spike to achieve rapid visual processing tasks where complex categories such as animals can be detected, effectively discarding any codes based on more than 1 spike (Thorpe & Imbert, 1989).

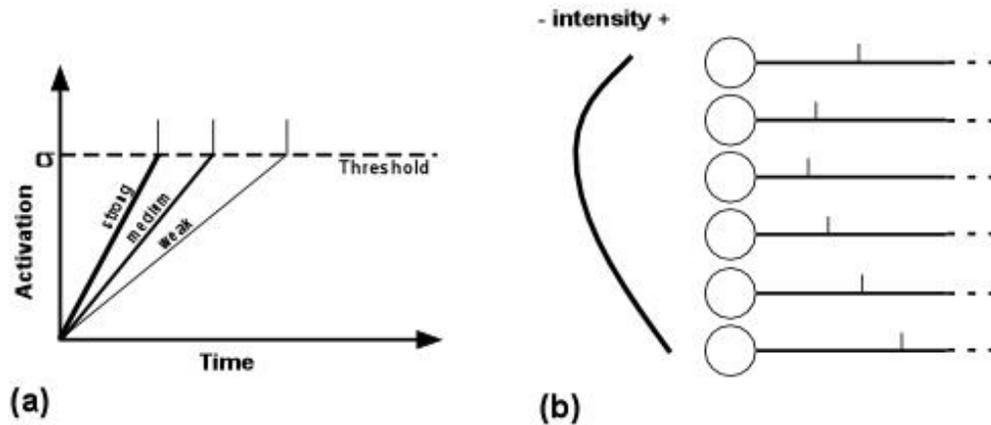
Under this constraint of one-spike-per-cell coding, how could the brain possibly encode information to achieve such speed-demanding tasks? And how are these selective responses generated? Here, we will discuss the possibility that information resides in the spatio-temporal structure of spike waves (one-spike-per-afferent asynchronous patterns) and argue that, coupled to a biological hebbian learning rule whose intrinsically suited for temporal codes (spike-timing-dependent plasticity, STDP), neurons can be made to be highly selective to particular aspects of the visual scene. In addition to the implications of STDP for feedforward visual processing, we will also raise hypotheses about its role in establishing connectivities between different areas of the visual connections, especially where feedback connections are concerned.

## 2. Spike times codes

What does a neuron tell us when emitting action potentials? Usually, we say that its firing rate reflects an analogue value such as stimulus intensity. Following the trail blazed by Barlow, we can even take a step further in reflexion and state that a single unit not just codes for the presence of a certain feature but for the *probability* of its presence – or degree of certainty (Barlow, 1972). The answer thus looks very simple: the higher the firing rate, the higher the stimulus intensity.

But at the same time, the rate-only solution tends to overlook an even simpler fact: a neuron, even according to the simplistic integrate-and-fire model, reaches its threshold fastest when the incoming stimulus matches its selectivity. This is the basic idea at the heart of asynchronous coding schemes that propose that, effectively, neurons can be thought of as

acting as analog-to-delay converters (Thorpe & Imbert, 1989; **Fig. 1a**). In other terms, the activation, or stimulus intensity, would determine the firing latency, with respect to a reference signal which can be either external - the stimulus onset (Celebrini et al., 1993; Panzeri et al., 2001; Johansson & Birznieks, 2004); or internal – possibly determined by low-field potentials (LFP) at a particular frequency (Hopfield, 1995). For example of the latter case, the location of a rat within its environment can be retrieved from the precise phase relationship of action potentials fired by “place cells” with respect to the hippocampal EEG theta rhythm (O’Keefe & Recce, 1993; Mehta et al., 2002).



**Figure 1 Latency encoding of stimulus**

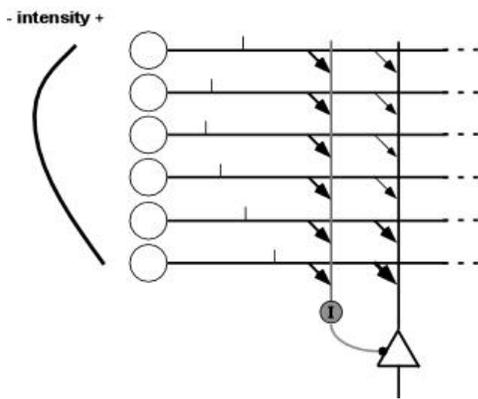
- (a) For a single neuron, the weaker the stimulus, the longer the time-to-first spike.  
 (b) When exposed to a population of neurons, the stimulus evokes a spike wave, which asynchrony encodes the information within.

Interestingly, retinal ganglion cells can be seen as analog-to-delay units (Uzzell and Chichilnisky, 2004) and display the favourable feature of emitting precisely timed spikes (Reich et al., 1997; Berry et al., 1998). Considering that the retina tiles the visual scene with an array of such cells, it follows that what is visually presented to the organism can be translated into a wave of first spikes whose times are more or less precisely repeated from one presentation of the same stimulus to the other. Said otherwise, the order of firing within a population of cells, induced by the relative spike times, can be used to encode information with a single spike per neuron (**Fig. 1b**). VanRullen and Thorpe provided theoretical evidence supporting this view by demonstrating that, when applied to retinal ganglion cells, such a coding scheme allowed identification of many stimuli when as few as 1.0% of the cells have fired (VanRullen & Thorpe, 1998).

At this point, it is worth emphasizing that this sparse level of diagnostic activation in the input layer correspond to the very first spikes, meaning that after stimulus onset the information needed to achieve identification can very rapidly reach the cortex for subsequent categorisation in the ventral pathway. If that small of information is sufficient, and given the high metabolic cost of spike generation (Levy & Baxter, 1996; Lennie, 2003), it could be useful to restrict firing in the input layer to economise energy. How might this be achieved? Here again, temporal coding can provide a simple solution in the form of a feedback, disinaptic inhibitory mechanism acting as a k-Winner-Take-All (k-WTA) operation on the input: depending on the threshold for its activation, inhibition shuts down the entire population once the first  $k$  cells have fired in the input layer. The kind of time-locked recurrent circuit proposed here exists in CA1 of the hippocampus, that produces inhibition in

response to the onset of a series of spikes (Pouille & Scanziani, 2004). Specifically, Pouille and Scanziani showed the ability of such circuits to differentially process the first action potential in a spike train from its rate, at least as far as feedback inhibition is concerned.

But now, if information is encoded in the rank order, how could an efferent neuron decode it? A simple feedforward shunting inhibition mechanism could in principle be used to produce a selective response (**Fig. 2**). Progressively desensitizing the output neuron with each incoming spike allows the first inputs to be fully effective while later ones produce less and less activation. The total amount of activation will thus depend on how well the order of firing within the input pattern matches the corresponding synaptic weights, maximal activation happening when they superimpose precisely. Hence, if the synaptic weights are set according to the stimulus to be identified (see section 3), the corresponding efferent neuron would only respond optimally in the presence of that particular stimulus. This idea has recently been tested within a biological model of the lateral geniculate nucleus; it successfully demonstrated that asynchronous excitatory spike waves and disynaptic feedforward inhibition generate orientation selectivity in the primary visual system while providing contrast invariance through automatic gain control (Delorme, 2003).



**Figure 2 Shunting inhibition for rank-order decoding**

A pyramidal neuron (triangle cell) receives excitatory input from the afferent units through synapses with variable weights. Unit I receives an excitatory input from the same population but through synapses with equal weights. This inhibitory cell generates shunting inhibition that progressively desensitizes the pyramidal cell as more and more of the inputs fire. First inputs are effective while later ones produce less and less activation.

How plausible is the suggestion that shunting inhibition could be used to implement rank-order decoding? To answer this, let us take a look at two requirements of shunting inhibition in the context of rank-order coding: (1) quasi-instantaneity, so that each excitatory input be followed by its respective inhibitory desensitisation; (2) non-selectivity, to happen whenever an input fires regardless of the information it bears. Interestingly, recent experimental work on fast spiking interneurons in the somatosensory cortex supports the existence of such circuits in the brain. Indeed, Swadlow and Gusev characterized the fast spiking inhibitory interneurons as having very small somas that allow for a very rapid reaction to the activation of their inputs. Furthermore, they do receive strong inputs and show essentially no stimulus selectivity (Swadlow & Gusev, 2002). Moreover, these cells are known to be electrically coupled: these connections will tend to make the population fire together, thus considerably reducing any selectivity its cells could possess (Galaretta & Hestrin, 2001). Finally, evidence in the visual system shows that shunting inhibition targets neurons of cat primary visual cortex as evidenced by the three-fold increase in soma conductance recorded very rapidly after presentation of a visual stimulus (Borg-Graham et al., 1998).

A temporal neural code based on latency rank-order coding thus provides a theoretical explanation of the speed with which the brain is able to categorize complex visual scenes. All along the ventral pathway, a wave of spikes propagates in a feedforward manner, initiated at the level of the retinal ganglion cells, going through areas V1, V2 and V4 to reach IT where complex shapes are selectively encoded. At each stage of this hierarchy, information is

contained within the asynchronous firing pattern, with the first spikes corresponding to the most salient features for the next stage (VanRullen & Thorpe, 2002). The question remains as to how these “features” are shaped. In the next section, we will explore the possibility that neurons in the visual system can be made to be selective to particular aspects of the visual scene using the temporal code described so far, coupled with spike-timing dependent plasticity (STDP).

### 3. Temporal learning with STDP

How can neurons be made selective? This question directly refers to the neural correlates of learning and synaptic plasticity (a review in Abbott & Nelson, 2000). And since we are concerned with the way neurons encode information in their activity, it specifically relates to the field of activity-driven synaptic modifications (also termed “Hebbian” learning in reference to Donald O. Hebb’s original postulate – Hebb, 1949).

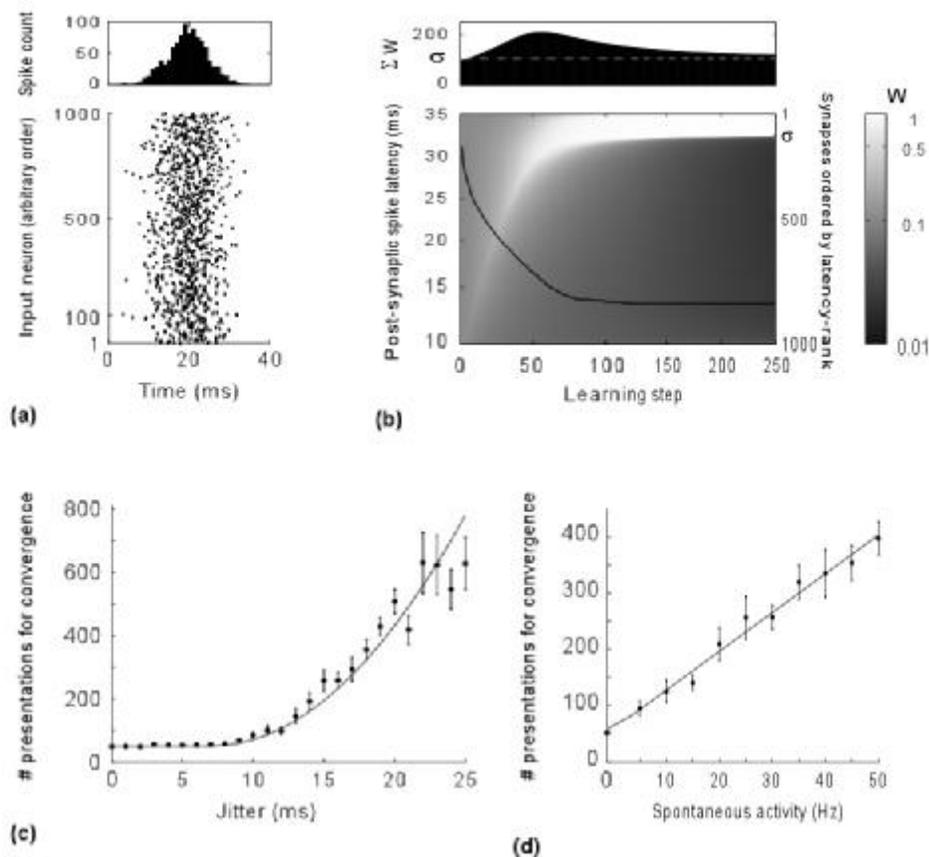
The rules for long-term potentiation (LTP) and depression (LTD) were initially described in terms of rate-based approaches. Thus, after having been stimulated at high firing rates, an excitatory synapse can show potentiation (LTP) (Bi & Poo, 2001; Bear, 2003 for reviews). Interestingly, recent experimental studies have brought these rules into the temporal domain as it has been found that LTP and LTD could be induced depending on the temporal relationship of a pre-synaptic spike relative to a post-synaptic one: when the former precedes the latter, the synapse is reinforced; it is depressed when the postsynaptic spike is emitted before a spike hits the synapse (Markram et al., 1997). STDP became even more “time-friendly” when it was later discovered that the amount of modification depended on the precise delay between the two spikes: maximal when the pre- and post-synaptic spikes are close together, the effects gradually decrease and disappear with intervals in excess of a few tens of milliseconds (Bi and Poo, 1998; Zhang et al., 1998; Feldman, 2000). This discovery not only served to revive the debate on timing in neural codes in a simple yet elegant way but more importantly has the important characteristic of allowing weight modifications to be regulated locally. This can be seen in the role of back-propagating action potentials in signalling the occurrence of a post-synaptic potential at the synaptic locus (Magee & Johnston, 1997).

By providing an experimental basis for synaptic learning rules based on spikes rather than firing rates, spike timing-dependent synaptic plasticity has become the subject of numerous theoretical investigations. While it may be involved in the formation, as well as the refinement, of cortical maps (Song & Abbott, 2001), temporal pattern recognition or coincidence detection (Kepecs et al., 2002 for a review), its most interesting feature resides in its competitive and stabilizing nature: it strengthens correlated inputs while being insensitive to firing rates or the degree of variability of a given synaptic input (Song et al., 2000; Gerstner & Kistler, 2002a).

However, these modelling studies have mostly been conducted using firing rates as the vector for the neural code whereas STDP is, intrinsically, sensitive to precise spike times. What of STDP effects with temporal codes? We will first consider that reproducibility in stimulus-locked neural responses implies that neurons could potentially be repeatedly exposed to the same input patterns (**Fig. 3a**). This might occur as the result of multiple exposures with the same stimulus at different times in life, or alternatively, a single stimulus exposure could also potentially result in a sequence of similar processing waves through the rhythmic activity of cortical oscillations (Hopfield, 1995). In either of these conditions, theoretical simulations have shown that two concomitant phenomena will take place: a reduction in the latency of the post-synaptic response along with the selective reinforcement of the synapses receiving the

earliest repeated inputs (Song et al., 2000; Gerstner & Kistler, 2002b; Guyonneau et al., 2004; **Fig. 3b**). The dynamical consequence of the asymmetrical, retrograde form of STDP (retrograde because potentiation affects what happened *before* the post-synaptic spike, thus favoring a “back-in-time” motion), this trend is simplistically explained as follows: for one given input pattern presentation, the input spikes elicit a post-synaptic response, triggering the STDP rule: synapses carrying input spikes just preceding the post-synaptic one are potentiated while later ones are weakened. The next time this input pattern is re-presented, firing threshold will be reached sooner which implies a slight decrease of the post-synaptic spike latency. Consequently the learning process, while depressing some synapses it had previously potentiated, will now reinforce different synapses carrying even earlier spikes than the preceding time. By iteration, it follows that upon repeated presentation of the same input spike pattern, the post-synaptic spike latency will tend to stabilize at a minimal value while the first synapses become fully potentiated and later ones fully depressed.

Remarkably, stabilization occurs when enough of the earliest inputs are potentiated so as to evoke a post-synaptic spike with a single spike per input. Hence the output neuron threshold ( $\sigma$ ) determines the number of excitatory synapses actually selected (**Fig. 3b, top**), which can be evaluated to 10 to 40 – out of a ten thousands – according to electrophysiological measures (Shadlen & Newsome, 1994; Nowak & Bullier, 1997).



**Figure 3** Temporal learning for an asynchronous spike wave

With an identically repeated spike wave, the neuron learns to react faster to its target. Synaptic weights converge onto the earliest firing afferents. **(a) Typical incoming activity.** Bottom, raster plot of an asynchronous spike wave (mean 20ms, 5ms standard deviation)

being presented as such from one step to the other; when it spikes, the STDP learning rule is applied and its potential reset to 0 before going to the next presentation. Prior to presentation, the pre-synaptic neurons do not fire any spikes. *Top*, the corresponding post-stimulus time histogram (PSTH) showing the induced gaussian form corresponding to the reproduced spike wave. **(b) Dynamics of repeated STDP.** *Top*, sum of all synaptic weights at each presentation (the dashed line represents the output neuron threshold). The sum of the synaptic weights stored in the afferents stabilizes at threshold value. *Bottom*, the horizontal axis corresponds to the number of presentations (i.e. the learning step). The black line refers to the left axis and shows the reduction of post-synaptic latency during the course of learning. The background image refers to the right axis where each synapse weight is mapped by a gray-level index (see the corresponding bar on the right). Synapses are ordered by spiking latency of the corresponding neuron within the reproducible input pattern. During learning, earliest synapses become fully potentiated and later ones are weakened. **(c) Effect of jitter.** Jitter is generated by a gaussian distribution. Increasing its standard deviation does not affect convergence until about 10ms. From there, it slows the system roughly quadratically. **(d) Effect of spontaneous firing rate.** Increasing background activity slows convergence approximatively linearly.

Remarkably, this trend still arises in biologically realistic conditions as neither jitter (**Fig. 3c**) nor spontaneous activity (**Fig. 3d**) can prevent it from occurring (Guyonneau et al., 2004). Input reproducibility is the main reason behind STDP being able to selectively reinforce the earliest inputs. Jitter main effect will precisely be to perturb the fine temporal structure: increasing its width increases the probability of not only piling up consecutive effects but also in inverting them, thus slowing down convergence.

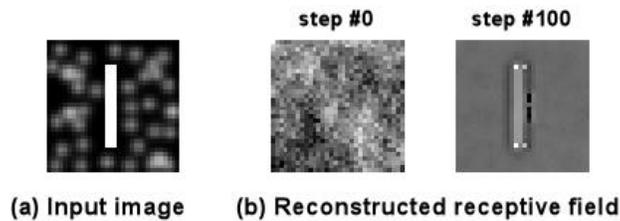
Spontaneous activity translates into additional spikes being randomly present in the LTP part of the STDP window, or in its LTD part, in proportion with the respective time-span of these modification windows (respectively 20 ms for LTP and 22ms for LTD – see Guyonneau et al., 2004 for details). It follows that in the course of learning, the noise-induced total potentiation affecting a given synapse is slightly inferior to the noise-induced total depression because, statistically, more spikes will have triggered LTD than LTP. And the more spontaneous activity, the bigger the off-balance in favour of the depression process. Potentiation being what makes selective reinforcement reach for the earliest inputs in the presence of the reproducible stimulus, it follows that, as the rate of spontaneous activity is increased, learning will increasingly take more time to converge.

Hence, STDP would naturally lead a neuron to respond rapidly to a precise pattern on the basis of the first few spikes in its afferents, even in a “noisy” system. Additionally, if we are ready to accept that selectivity to a stimulus can be expressed in the temporal domain if a neuron responds earlier to a given stimulus than to any other, then learning through STDP does make the neuron more selective (Guyonneau et al., 2004). This theoretical result thus supports the idea that information can be potentially be encoded in a single spike wave. But how would the phenomenon apply in the case of the visual system? What can be learnt in terms of visual representations in the ventral pathway when mimicking experience-dependent synaptic shaping?

#### 4. STDP in Vision: Feedforward processing

To make the situation more realistic, imagine how the array of retinal ganglion cells would respond to a flashed bar (**Fig. 4a**). Using the fact that the most strongly activated cells reach threshold first, presenting an oriented bar would evoke a wave of spikes whose first action

potentials corresponds to the *ordered* borders of the stimulus. Now suppose that a neuron in visual cortex receives initially weak inputs from all the retinal afferents within a particular receptive field area (**Fig. 4b** – step#0), and that in response to the flash, it finally goes over threshold and fires a spike. According to the trend described in section 3, repeated presentation of the same stimulus would progressively refine its set of afferents until the high weights are all on the earliest firing inputs. Followingly, we can infer that orientation selectivity in V1 could be explained by an ordered alignment of LGN inputs (Hubel & Wiesel, 1962; **Fig. 4b** – step#100) using latency rank order coding as demonstrated in Delorme (2003).



**Figure 4 Shaping of orientation selectivity**

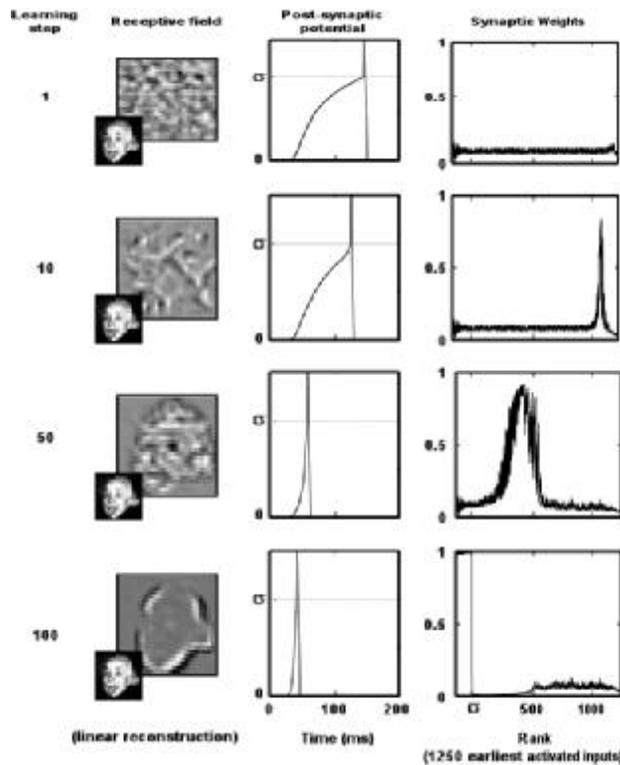
A model of the retina is repeatedly presented with an oriented bar (**a**). The evoked activity is propagated towards a neuron in visual cortex via the LGN. Initially its inputs are very weak (**b** – step #0) but the spike timing-dependent plasticity process refine them until selectively reinforcing the first ones (**b** – step #100). As these correspond to the ordered borders of the stimulus, the selectivity of the visual neuron is thus composed of an ordered alignment of LGN inputs.

The receptive field is linearly reconstructed based on the synaptic weights and the selectivity of the afferent neuron (ON- and OFF-center receptive fields).

As proposed in section 2, rapid visual processing could be achieved in the form of an asynchronous wave of spikes, initiated in the retina and regenerated at each stage of the hierarchy to be propagated throughout the ventral pathway (VanRullen & Thorpe, 2002). Hence, neurons could be exposed to reproducible spike waves at each level of the hierarchy. Interestingly, experimental observations showed that V1 cells display a degree of selectivity in their latencies, respective to stimulus onset (Celebrini et al., 1993). Then, one can extend the present architecture to take V1 cells responses and selectivities into account, generating spikes at times that depend on the local orientation, similarly to RGCs responses in the former model.

Presenting this population with the same image repeatedly exposes a neuron in a later stage of the hierarchy to a characteristic temporally organized wave of spikes at the output from V1. As expected, the post-synaptic potential progressively becomes steeper, and stabilizes with a minimal latency (**Fig. 5**, middle column) while at the same time the earliest inputs are finally selected (**Fig. 5**, right column). But, unlike the results in the earlier theoretical section (section 3), here we can get a glimpse of the representations used through

the linear reconstruction of the receptive field, based on the set of synaptic weights and the selectivity of the corresponding afferent neurons (**Fig. 5**, left column): in the course of learning, the reconstructed “optimal stimulus” for the neuron does not seem to make much sense. But as the neuron “matures”, a structured representation emerges, that is built upon the earliest afferents of the input spike wave. As a consequence, with as few as a hundred presentations, the neuron can “learn” to recognize a particular image, in this case, a well known image of Einstein.



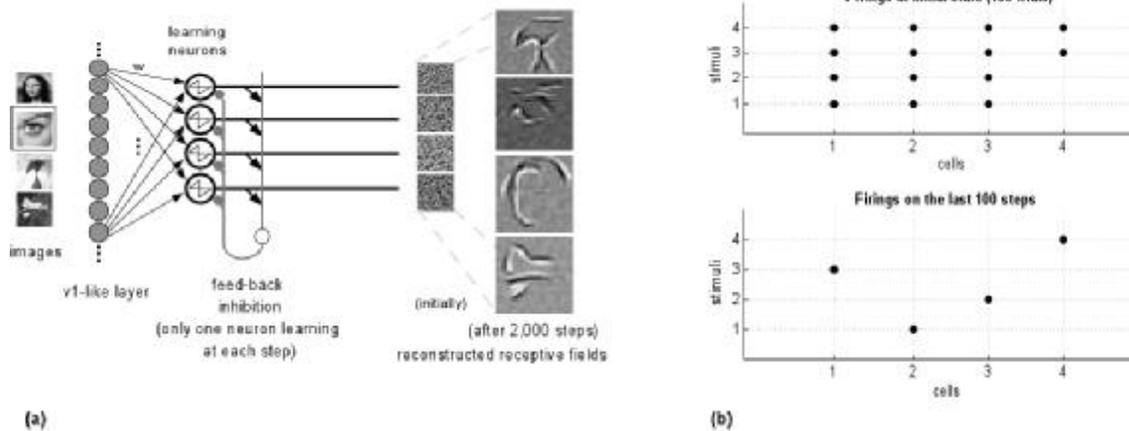
**Figure 5 Einstein: STDP learning of a V1-filtered face**

A population of V1-like cells encodes an orientation for each pixel in the image presented to the network (here, Einstein’s face); each cell acts as an analog-to-delay converter where the latency of its first spike depends on the strength of the orientation in its receptive field. Time taken to achieve recognition of the stimulus decreases (middle column) while a structured representation emerges and stabilizes (left column) that is built upon the earliest afferents of the input spike wave (right column).

The receptive field is linearly reconstructed based on the synaptic weights and the selectivity of the afferent neuron (here, orientations). It is thus redundant with the distribution of effective weights.

Of course, there will never be just one neuron to listen to the output of V1. And normally, the same input stimulus will not be presented over and over again. Imagine what would happen if we extend the preceding simulation to more input images, randomly presented to a population of as many potentially learning cells. Suppose that each time one of the cells fires, it inhibits its neighbourhood - the other cells in the population - through a feedback (FB) mechanism. As a result of this competition, only one cell learns the v1-filtered spiking activity induced by the stimulus at each step. The results are remarkable: while the neurons obviously do not know which image was presented, they nevertheless retain one stimulus and one only, by blindly separating the sources (**Figure 5a**). When considering the responses at the population level, one can also note that learning makes the responses selective by implementing a sparse code (**Figure 5b**; Olshausen 2002). Indeed, the rapid responses displayed by the neuron when presented with the stimulus used for learning (Guyonneau et al., 2004) translates within the present architecture into a situation where the neuron that fires

first inhibits its neighbors, and thus becomes the only one to fire. The local, autonomous<sup>1</sup> STDP learning rule can then lead to rapid, highly selective visual responses reminiscent of the notorious “grandmother cells” (Barlow, 1972).



**Figure 6** Population learning: Emergence of selective responses

**(a) Architecture of the network and receptive fields.** At each step, one of the 4 images on the left is presented to the network. The V1-like layer generates one spike per cell just like in Figure 4. Spike times are then jittered (5 ms gaussian delay) and 5 Hz Poisson-inspired spontaneous activity is added to the spike pattern at each presentation. The changing incoming activity is propagated towards the next layer where 4 neurons integrate it. The first to fire inhibits its neighbours and triggers the STDP learning rule. After 2,000 random presentations, each neuron has learned one stimulus and one only, in the same manner as in Figure 4. **(b) Population response.** A thick dot indicates if the neuron fired first when presented with the corresponding stimulus. *Top*, initially, each neuron is likely to respond to any of the stimuli when tested for first response on 100 trials without STDP learning (plasticity was shut off for the test). *Bottom*, at the end of training, any input image is clearly identifiable based on which neuron fires (learning steps #1901 to #2000).

Within these feed-forward visual architectures, neural selectivities at a given level are carved through STDP on the ones from the preceding stage (**Fig. 4**), according to the statistics of the external world, just like letters (ON- and OFF-center inputs for example) compose words (orientations). In that sense, the fact that STDP is a purely local hebbian mechanism allows the whole procedure to be recurrently applied, *ad libitum*, provided “forms” at the initial level are precisely known and the architecture respected. Hence, arbitrarily complex selectivities, for visual representation and recognition, can be autonomously generated in later stages, through exposition to patches of natural scenes for example.

## 5. STDP in Vision: feedback connectivity

<sup>1</sup> As far as STDP is concerned, the neuron’s spiking activity is the main determinant of the synaptic modifications. These regulations in turn affect integration of the presynaptic spikes before producing postsynaptic ones. Thus, the learning rule makes the cell autonomous, at least functionally speaking.

So far, the architectures that we have been considering are very simple ones, with essentially only feed-forward connections – the exception being the feed-back inhibition used to perform the k-WTA operation. What would happen if the STDP learning rule was applied in the case of architectures with a richer set of connections? The connections between different areas may be classified into 3 different types according to their preferential termination: FF connections typically terminate in layer 4, lateral ones in a columnar way and FB ones prefer a multilaminar pattern excluding layer 4. As a consequence, cortical areas can be organized in a hierarchy according to the pattern of reciprocal connections between them (Felleman & van Essen, 1991). However, it has already been suggested that the rank of a given area in the anatomical hierarchy does not necessarily correspond to that of the order of activation. For instance, area MT may be placed at the top of the anatomical hierarchy while yielding early responses that would qualify it for the earliest stages of visual processing; area MT, but also subsequent area MST and frontal eye field (FEF), would thus be part of the “fast brain”, where median latencies typically occur between 40 and 80 ms ; as opposed to the “slow brain” (100-150 ms responses) that groups most areas of the temporal lobe and some areas of the frontal cortex located rostral to the frontal eye field (Nowak & Bullier, 1997).

According to this hypothesis, latency can be seen as a dynamical principle of hierarchical organisation in the visual system, which would have consequences in refining laminar patterns of connection. We have shown here that STDP tunes neurons to the earliest spikes they receive: latencies basically dictate which afferents are selected. The similarities are such that one can wonder how STDP might influence connections between visual areas. For instance, the first, fastest wave of visual information, channelled by the magnocellular (M) pathways, reaches MT via V1 in a FF manner; from there, the M input goes “backwards” to V1 in time to meet the slower by 20ms, parvocellular (P) input (Nowak et al., 1995). It follows that P neurons in the superficial layers would be influenced by M inputs from MT because they arrive sufficiently early to be enhanced by STDP. Asymmetrically, since the M neurons have already responded on the first pass, any connections from MT inputs would tend to be depressed. Hence, the M pathway would indeed act as an ignition device for the P flow in V1 (Ullman, 1995) or even, if acting selectivity-wise too, as an active blackboard by modulating V1 and V2 responses (Bullier, 2001); it would also be temporally decoupled: one wave of activity in the M system would not influence the next one.

Similarly, an apparent “anomaly” of the visual system might be explained by STDP. According to the classical view, the frontal eye fields (FEF) are located high in the visual hierarchy, in a FB position compared to V4 (Felleman and van Essen, 1991). However, FEF-V4 connections display a high percentage of supragranular layers neurons retroactively marked when injecting in V4, suggesting that FEF-to-V4 are in fact FF in nature (Batardiere et al., 2002). From the anatomical point of view, FEF initially possesses neurons in the supragranular and infragranular layers both connecting to V4 layer 4. In terms of latency responses after stimulus presentation, FEF tends to fire before V4. The first wave of activity hitting V4 would thus come from FEF supragranular neurons and elicit the earliest responses which would consequently be strengthened. At the same time, the supragranular activity in FEF propagates locally in its lower layers which would in turn fire slightly later. Inputs from the infragranular layers arriving after the first spikes in FEF, these connections would be depressed. Hence, FEF would hierarchically be placed *before* V4.

## 6. Conclusion

When we aim at understanding how the brain works, we are faced with the puzzling challenge of penetrating its essential complexity (ironically, making sense of what makes

sense of the things in the world is yet another big issue). We can however single out its most significant performances, and provided we can explain them easily, unveil a piece of the mind game. In that sense, Thorpe and colleagues have put lights on remarkable visual performances that help us producing knowledge about the neural code. Given the physiological properties of nerve cells, the known anatomy of the visual system and the primate performances in rapid visual categorisation tasks, they hypothesize that fast neural information processing is based on a single spike per neuron. As exposed here, this proposition, somewhat peculiar in regard to common knowledge yet mandatory by its factual roots, underlines the key role of temporal codes and sparse representations in understanding some not-so-obvious inner workings of the brain.

Hence, neural representations, as they emerge through spike timing-dependent plasticity and simple, biologically plausible architectures, would be built on a very small subset of afferents, characterized by their correlated repetition and most of all early arrival. Said otherwise, experience at the neuronal level selects inputs characterized by their ability to transmit relevant information in the fastest way possible. Precisely, expertise, or maturation, would correlate with the development of quickness as displayed by the latency reduction exposed here and in visual categorisation during childhood (Batty & Taylor, 2002). In short, one usually says that practice makes perfect; well, repetition may very well make swift.

In conclusion, the temporal approach described here does in no way discard the evident role of firing rates. In fact, the brain may very well use two different strategies that need not be opposed but rather considered in a complementary way, as experimental evidence shows that neurons can make the distinction between these (Pouille & Scanziani, 2004). Rates-based models nonetheless get most of the neuroscientific attention, mostly because temporal coding still lacks a clean-cut experimental demonstration; at the same time, it yields fascinating proposals.

## References

Abbott, L. F., and Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nat Neurosci* 3 *Suppl*, 1178-1183.

Abeles, M. (1991). *Corticonics: Neural circuits of the cerebral cortex*, Cambridge University Press edn (Cambridge).

Adrian, E.D. (1928) *The Basis of Sensation*, Londres, Christophers.

Bair, W., and Koch, C. (1996). Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comput* 8, 1185-1202.

Barlow, H. B. (1972). Single units and sensation: a neuron doctrine for perceptual psychology? *Perception* 1, 371-394.

Batardiere, A., Barone, P., Knoblauch, K., Giroud, P., Berland, M., Dumas, A. M., and Kennedy, H. (2002). Early specification of the hierarchical organization of visual cortical areas in the macaque monkey. *Cereb Cortex* 12, 453-465.

Batty, M., and Taylor, M. J. (2002). Visual categorization during childhood: an ERP study. *Psychophysiology* 39, 482-490.

- Bear, M. F. (2003). Bidirectional synaptic plasticity: from theory to reality. *Philos Trans R Soc Lond B Biol Sci* 358, 649-655.
- Berry, M. J., Warland, D. K., and Meister, M. (1997). The structure and precision of retinal spike trains. *Proc Natl Acad Sci U S A* 94, 5411-5416.
- Bi, G. Q., and Poo, M. M. (2001). Synaptic modification by correlated activity : Hebb's postulate revisited. *Annu Rev Neurosci* 24, 139-166.
- Borg-Graham, L.J., Monier, C. and Fregnac, Y. (1998). Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature* 393, 369-373.
- Bullier, J., McCourt, M. E., and Henry, G. H. (1988). Physiological studies on the feedback connection to the striate cortex from cortical areas 18 and 19 of the cat. *Exp Brain Res* 70, 90-98.
- Bullier, J. (2001). Integrated model of visual processing. *Brain Res Brain Res Rev* 36, 96-107.
- Celebrini, S., Thorpe, S., Trotter, Y., and Imbert, M. (1993). Dynamics of orientation coding in area V1 of the awake primate. *Vis Neurosci* 10, 811-825.
- Fabre-Thorpe, M., Richard, G., and Thorpe, S. J. (1998). Rapid categorization of natural images by rhesus monkeys. *Neuroreport* 9, 303-308.
- Feldman, D. E. (2000). Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron* 27, 45-56.
- Felleman, D. J., and Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex* 1, 1-47.
- Galarreta, M. and Hestrin, S. (2001). Electrical synapses between GABA-releasing interneurons. *Nature Rev. Neurosci* 2, 425-433.
- Gautrais, J., and Thorpe, S.J. (1998). Rate coding versus temporal order coding: a theoretical approach. *Biosystems* 48, 57-65.
- Gerstner, W., and Kistler, W. M. (2002a). Mathematical formulations of Hebbian learning. *Biol Cybern* 87, 404-415.
- Gerstner, W., and Kistler, W. M. (2002b). Learning to be fast. In *Spiking neuron models*, C. U. Press, ed. (Cambridge), pp. 421-432.
- Guyonneau, R., VanRullen, R., and Thorpe, S.J. (2004). Neurons tune to the earliest spikes through STDP. *Neural Computation* (accepted).
- Hebb, D.O. (1949). *The organization of behavior: a neuropsychological theory*. (Wiley, NY).

Hopfield, J.J. (1995). Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376, 33-35.

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

Johansson, R. S., and Birznieks, I. (2004). First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nat Neurosci* 7, 170-177.

Kepecs, A., van Rossum, M. C., Song, S., and Tegner, J. (2002). Spike-timing-dependent plasticity: common themes and divergent vistas. *Biol Cybern* 87, 446-458.

Lennie, P. (2003). The cost of cortical computation. *Curr Biol* 13, 493-497.

Levy, W. B., and Baxter, R. A. (1996). Energy efficient neural codes. *Neural Comput* 8, 531-543.

Liu, R. C., Tzonev, S., Rebrik, S., and Miller, K. D. (2001). Variability and information in a neural code of the cat lateral geniculate nucleus. *J Neurophysiol* 86, 2789-2806.

Magee, J. C., and Johnston, D. (1997). A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. *Science* 275, 209-213.

Mainen, Z. F., and Sejnowski, T. J. (1995). Reliability of Spike Timing in Neocortical Neurons. *Science* 268, 1503-1506.

Markram, H., Lubke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213-215.

Mehta, M. R., Lee, A. K., and Wilson, M. A. (2002). Role of experience and oscillations in transforming a rate code into a temporal code. *Nature* 417, 741-746.

Nakamura, K. (1998). Neural processing in the subsecond time range in the temporal cortex. *Neural Comput* 10, 567-595.

Nowak, L. G., Munk, M. H., Girard, P., and Bullier, J. (1995). Visual latencies in areas V1 and V2 of the macaque monkey. *Vis Neurosci* 12, 371-384.

Nowak, L. G., and Bullier, J. (1997). The timing of information transfer in the visual system. In *Extrastriate visual cortex in primates*, K. S. Rockland, J. H. Kaas, and A. Peters, eds. (New York, Plenum Press), pp. 205-241.

O'Keefe, J., and Recce, M.L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3, 317-330.

Olshausen, B. A. (2002). Sparse codes and spikes. In *Probabilistic Models of the Brain: Perception and Neural Function*, R. P. Rao, B. A. Olshausen, and M. S. Lewicki, eds. (MIT Press), pp. 257-272.

- Panzeri, S., Petersen, R.S., Schultz, S.R., Lebedev, M., and Diamond, M.E. (2001). The role of spike timing in the coding of stimulus location in rat somatosensory cortex. *Neuron* 29, 769-777.
- Perkel, D.H., and Bullock, T. (1968). Neural coding. *Neurosciences Research Program Bulletin* Vol:221-348.
- Perrett, D. I., Rolls, E. T., and Caan, W. (1982). Visual neurones responsive to faces in the monkey temporal cortex. *Exp Brain Res* 47, 329-342.
- Petersen, R. S., Panzeri, S., and Diamond, M. E. (2001). Population coding of stimulus location in rat somatosensory cortex. *Neuron* 32, 503-514.
- Pouille, F. and Scanziani, M. (2004). Routing of spike series by dynamic circuits in the hippocampus. *Nature* 429, 717-723.
- Reich, D. S., Victor, J. D., Knight, B. W., Ozaki, T., and Kaplan, E. (1997). Response variability and timing precision of neuronal spike trains in vivo. *J Neurophysiol* 77, 2836-2841.
- Richmond, B. J., and Optican, L. M. (1990). Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. II. Information transmission. *J Neurophysiol* 64, 370-380.
- Roelfsema, P. R., Lamme, V. A., and Spekreijse, H. (1998). Object-based attention in the primary visual cortex of the macaque monkey. *Nature* 395, 376-381.
- Rolls, E. T., Sanghera, M. K., and Roper-Hall, A. (1979). The latency of activation of neurones in the lateral hypothalamus and substantia innominata during feeding in the monkey. *Brain Res* 164, 121-135.
- Rolls, E. T., Perrett, D. I., Caan, A. W., and Wilson, F. A. (1982). Neuronal responses related to visual recognition. *Brain* 105, 611-646.
- Rubin, J., Lee, D. D., and Sompolinsky, H. (2001). Equilibrium properties of temporally asymmetric Hebbian plasticity. *Phys Rev Lett* 86, 364-367.
- Salzman, C. D., Britten, K. H., and Newsome, W. T. (1990). Cortical microstimulation influences perceptual judgements of motion direction. *Nature* 346, 174-177.
- Shadlen, M. N., and Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Curr Opin Neurobiol* 4, 569-579.
- Singer, W. (1999). Time as coding space? *Curr Opin Neurobiol* 9, 189-194.
- Softky, W. R., and Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13, 334-350.
- Song, S., Miller, K. D., and Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci* 3, 919-926.

Song, S., and Abbott, L. F. (2001). Cortical development and remapping through spike timing-dependent plasticity. *Neuron* 32, 339-350.

Swadlow, H.A., and Gusev, A.G (2002). Receptive-field construction in cortical inhibitory interneurons. *Nat Neurosci* 5, 403-404.

Thorpe, S.J., and Imbert, M. (1989). Biological constraints on connectionist modelling. In *Connectionism in perspective*, R. P. e. al., ed. (Elsevier), pp. 63-93.

Thorpe, S.J. (1990). Spike arrival times: a highly efficient coding scheme for neural networks. In *Parallel processing in neural systems* (Eckmiller, R. et al., eds.), pp 91-94, Elsevier.

Thorpe, S.J., Fize, D., and Marlot, C. (1996). Speed of processing in the human visual system. *Nature* 381, 520-522.

Thorpe, S.J., Delorme, A., and Van Rullen, R. (2001). Spike-based strategies for rapid processing. *Neural Netw* 14, 715-725.

Ullman, S. (1995). Sequence seeking and counter streams: a computational model for bidirectional information flow in the visual cortex. *Cereb Cortex* 5, 1-11.

Uzzell, V.J., and Chichilnisky, E.J. (2004). Precision of Spike Trains in Primate Retinal Ganglion Cells. *J Neurophysiol* 92, 780-789.

VanRullen, R., and Thorpe, S.J. (2001a). Is it a bird? Is it a plane? Ultra-rapid visual categorisation of natural and artificial objects. *Perception* 30, 655-668.

VanRullen, R., and Thorpe, S.J. (2001b). The time course of visual processing: from early perception to decision-making. *J Cogn Neurosci* 13, 454-461.

VanRullen, R., and Thorpe, S. J. (2002). Surfing a spike wave down the ventral stream. *Vision Res* 42, 2593-2615.

Victor, J. D., and Purpura, K. P. (1996). Nature and precision of temporal coding in visual cortex: a metric-space analysis. *J Neurophysiol* 76, 1310-1326.

Zhang, L. I., Tao, H. W., Holt, C. E., Harris, W. A., and Poo, M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature* 395, 37-44.