

A Model of Predictive Coding Based on Spike Timing

Dana H. Ballard

Department of Computer Science
University of Rochester,
Rochester, NY 14627-0226, USA

Rajesh P. N. Rao

The Salk Institute
La Jolla,
California, USA

Zuohua Zhang

Department of Computer Science
University of Rochester,
Rochester, NY 14627-0226, USA

Abstract

Several decades of research have made many advances towards the goal of interpreting the neural spike train but a comprehensive understanding remains elusive. This paper pursues this goal in the context of a new class of models termed *predictive models*. Predictive models characterize the cortex as a memory whose parameters can be used to predict its input. This allows the input to be economically coded as a residual difference between itself and the prediction. Such models have recently had considerable success in modeling features of visual cortex. This paper shows that the predictive coding model can be extended to a lower level of detail that includes individual spikes as primitives. This is a significant improvement in perspicuity compared to the firing rate variables used by most current models. The specific model we describe exploits the use of coincidence of spike arrival times and the fact that neural representations can be distributed over large numbers of cells.

1 Introduction

The key problem for systems neuroscience has been succinctly described by [Nicholls *et al.*, 1992]: “The brain uses stereotyped electrical signals to process all the information it receives and analyzes. The signals are symbols that do not resemble in any way the external world they represent, and it is therefore an essential task to decode their significance.”

Despite the fact that we now have detailed models of spike generation within single neurons and provisional models of neural firing, there is still no comprehensive understanding of how individual spikes are used in brain computation. To attempt to break the neural code, a logical place to start is the cortex, which uses a six-layered sheet of cells to represent memories. Since the ground-breaking work of Hubel and Wiesel, research has focused on the visual cortex, particularly striate cortex. Starting with Barlow, there has been a concerted effort to interpret the response of cells in information-theoretic terms. However just recently progress has been made with predictive coding models [Olshausen and Field, 1997; Field, 1987; Rao and Ballard, 1996b; Rao and Ballard, 1998c]. In such models the synapses of cells are driven by the statistics of natural scenes. Adjustment of synaptic strength is based on the ability of a set of neurons to predict its input. Such models have been able to predict the distribution of cell properties such as receptive field size with impressive accuracy [Olshausen and Field, 1997]. However the main attractiveness of predictive coding is as a comprehensive minimal encoding strategy that extends across modalities. The critical feature of our predictive coding formulation [Rao and Ballard, 1996b] is that the structure of receptive fields can be predicted by assuming that the cortical memory costs metabolically in terms of spikes and synapses. The formulation also has other important advantages that 1) it can be used to model cortical hierarchies and 2) it can be used to solve image segmentation problems. An alternate formulation of predictive models uses a “sparseness” term [Olshausen and Field, 1997]. This cost is justified by its ability to fit experimental data. Its use to the brain is less clear, but the independence of the consequent neural representations facilitates their use in goal-directed processing.

All current predictive models have a significant drawback in that they rely on the classical interpretation of neural signaling. In the classical model of a neuron, the rate of spikes produced by a neuron is a function of the input magnitude [Rieke *et al.*, 1997]. This coding is well established in the peripheral nervous system and brainstem. It is also widely accepted as a model for cerebral neurons, such as cortical neurons in that many analyses assume rate coding as a starting point e.g. [Abbott and Dyan, 1999; Panzeri *et al.*, 1999]. The classical receptive field of a cortical neuron is defined by plotting its output spike rate in terms of stimulus parameters. The rate code model is implicitly used by experimentalists in characterizing the neural output by the post-stimulus histogram. In addition, it is widely used in computational learning models that use model neurons as primitive units. The rate code model fits experimental data well, but it has at least two crucial computational problems. First, cortical firing rates are typically low, from 10 to 100 Hertz. Thus decoding the spike rate at recipient neurons may be too slow to orchestrate complete motor responses to visual input that can occur within 200-400 milliseconds. Second, there is the “binding problem.” When a subset of neurons has to be grouped or distinguished in some way, this is very difficult to signal with a rate code.

Several groups [Prut *et al.*, 1998; von der Malsburg, 1995; Castelo-Branco *et al.*, 1998] have noted that computing with single spikes could potentially solve both the timing and binding problems. Individual spikes are generated on a timescale of milliseconds, and coincidences

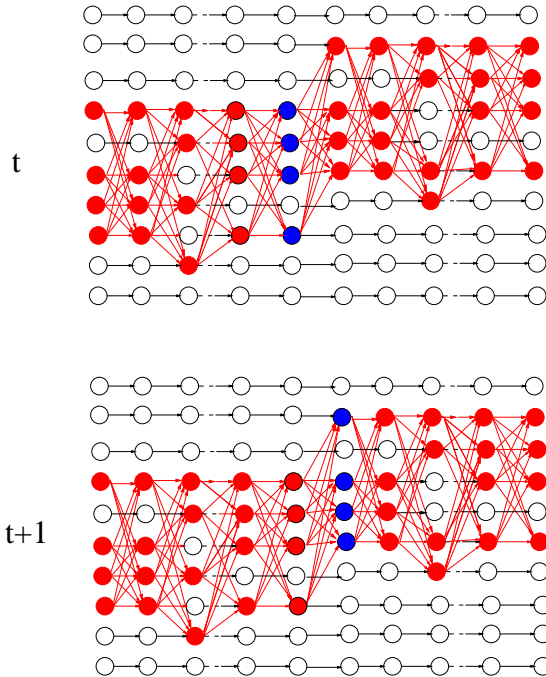


Figure 1: **A. Paths Architecture.** Signals are sent along 100-neuron wide *paths* (red) **Volleys** The signals are timed to be coincident in millisecond-synchronous *volleys* (blue).

between spikes can be used to indicate binding. Thus there is tremendous incentive for a general model that could use single spikes in computation. It would allow contact with experimental data on a more fundamental level. The model could stand in for the animal in that any measurement made on the animal system can be duplicated using the model.

We describe a model that shows how predictive coding can be extended to single spikes. As shown in Figure 1, single spikes are used in precisely timed groups termed volleys. The model shows how neurons could be used in a timesharing mode, each carrying large numbers of different volleys. In this model, the fact that the phase between a neuron's neighboring spikes is unimportant gives the appearance that timing is unimportant, but groups of neurons contain information coded in terms of precisely-timed volleys of spikes. Although such chains have been considered in terms of how they might emerge from random connections[Abeles, 1991], we contend that learning algorithms could create much larger numbers of such chains that would be observed by chance. By using a distributed encoding strategy, neurons can use many cooperating chains together in a *path*(See Figure 5). We contend that by considering the interactions between different such paths, that paths never interfere and thus can form the basis of secure parallel communication in the brain. Paths that do not interfere, even when they overlap substantially, can provide the brain with vast reserves of parallel computation. This is in sharp contrast with the limitations of the rate coding method.

Although research has been dominated for a long time by rate code models, recent research shows a huge increase in experimental work that implies that precise timing is important. At the moment this work is not organized into a coherent theory, but consists of many separate

findings. Thus there would be a great advantage if an overall theoretical framework could unify these results.

In the retina, the efficacy of synchronous codes has been noted as a way of retinal encoding [Meister, 1996]. Retinal ganglion cells have overlapping receptive fields, so that they have the possibility of communicating information via synchronous subsets of firing patterns. Meister reviews the evidence for this and analyzes the efficacy of the synchronous data as an encoding scheme. The retinal data establishes that biology has discovered the encoding scheme that we propose for cortex. Closer to the cortex, Reid, Alonso and Usrey have observed millisecond timing in cortico-thalamic connections [Reid and Alonso, 1995; Alonso *et al.*, 1996; Usrey *et al.*, 1998]. Furthermore this timing facilitates the firing of the recipient cells. Singer and colleagues have observed the synchronization of activity for visual inputs [Castelo-Branco *et al.*, 1998]. This activity seems to be part of a network as it is observed in the local field potential. In addition Shadlen and Newsome [Shadlen and Newsome, 1998] has observed reliable spike train sequences in area MT for repeated motion stimuli. Gallant *et al.* [Gallant *et al.*, 1998] has also observed repeated spike train activity for repeated stimuli. Additional evidence that volleys might be generated comes from experiments by [Sillito *et al.*, 1994]. In their anesthetized cats, neuron outputs in the LGN were correlated in the presence of the cortex, but uncorrelated (even though their firing rates were undisturbed), when the cortex was removed. In a tour de force experiment, Castelo-Branco *et al.* [Castelo-Branco *et al.*, 1998] observed correlations between the retina, LGN, and cortical areas 17 and 18 in the anesthetized cat. Their data show that the LGN exhibits correlations in the 60-120Hz range and that cortex exhibits correlations in two ranges: the 60-120Hz range seen in the LGN and also a 30-60Hz range.

Additional evidence for coincident firing comes from the Abeles group [Prut *et al.*, 1998]. They show convincingly that precise spike timing can be observed at the level of a millisecond in many motor cortex cells. Furthermore this timing is related to behavioral state, occurring in a subset of a complicated task. This kind of behavior would be expected in our model as the traffic through a particular neuron can be task dependent. The claim made by the model is that volleys are ubiquitous. When they are not found in a multicell recording it is probably because the sampling was too sparse to recover path elements. Thompson, Bichot and Schall [Thompson *et al.*, 1997], in more conventional post-stimulus histograms, has also shown that single neural recordings exhibit different task dependencies at different times. That is the spikes from a single neuron appear to reflect different features of the task at different moments (as opposed to being fixed feature detectors).

2 Predictive Coding

Since our central example of the paths model uses predictive coding, we begin with a short introduction to the relevant predictive coding concepts.

A fundamental feature of the neocortex is the reciprocity of connections between connected areas: if area 1 projects to area 2, then area 2 almost invariably projects to area 1 [Felleman and Van Essen, 1991]. The importance of feedback connections in determining neural response properties has been convincingly demonstrated in several neurophysiological experiments showing that neurons in lower areas such as striate cortex are heavily influenced by feedback from higher visual areas [Sandell and Schiller, 1982; Murphy and Sillito, 1987; Mignard and Malpeli, 1991]. An important functional role for feedback emerges if we posit

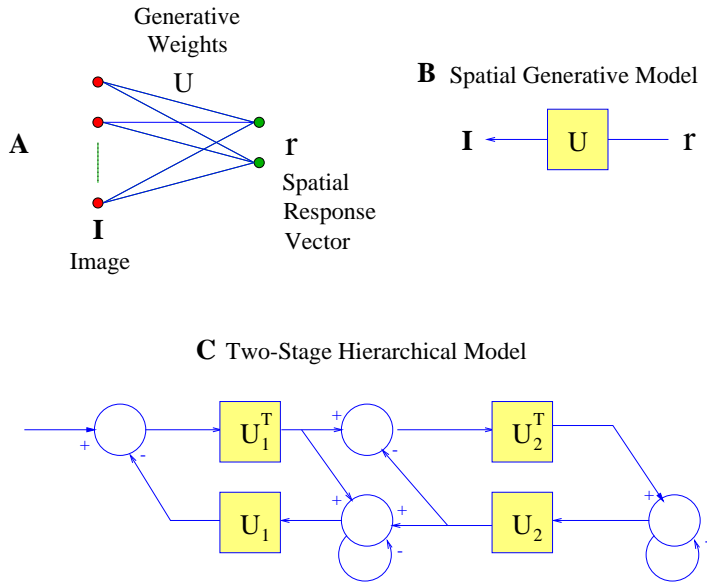


Figure 2: Predictive Coding Architecture (A) A smaller number of neurons can be used to predict input (B) The feedback circuit can be summarized in terms of a block diagram. (C) Multiple levels use the Kalman filter to propagate estimates.

that the feedback signals carry top-down predictions of lower-level inputs from higher, more abstract areas. The feedforward connections then need only convey the residuals [MacKay, 1956; Maybeck, 1979; Mumford, 1994] (or differences) between the current input and its prediction from the higher area. We have developed a hierarchical model of cortical visual representation that is based on the Minimum Description Length (MDL) principle [Rao and Ballard, 1996b; Rao and Ballard, 1998c]. The MDL principle applied to the cortex results in an error metric that trades off the cost of predicting the input against the cost of the circuitry doing the prediction. This can be specified as:

$$E = \|\mathbf{I}(\mathbf{x}) - \mathbf{U}\mathbf{r}\|^2 + \alpha\|\mathbf{r}\|^2 + \gamma\|\mathbf{U}\|^2 \quad (1)$$

where the operator $\|\cdot\|^2$ denotes the sum of squares of the elements of the vector or matrix argument. The first term represents the cost of predicting an image \mathbf{I} using neuronal firing rates \mathbf{r} operating through synapses \mathbf{U} . The second and third terms are the cost of the spikes and firing rates respectively. The firing rate dynamics are governed by

$$\dot{\mathbf{r}} = -\eta \frac{\partial E}{\partial \mathbf{r}} = 2\eta[\mathbf{U}^T(\mathbf{I} - \mathbf{U}\mathbf{r}) - \alpha\mathbf{r}]$$

The important point is the feedforward term $\mathbf{U}^T(\mathbf{I} - \mathbf{U}\mathbf{r})$ is the residual error rather than the raw signal \mathbf{I} . the implementation of these dynamics uses Kalman Filter circuits shown in Figure 2.

Most neurons in the visual cortex encode information in a highly selective fashion, typically responding only to a limited range of input stimuli such as moving bars of a particular orientation, velocity, and direction of movement [Hubel and Wiesel, 1962; Hubel and Wiesel, 1965]. These

common properties of cortical cells can be seen as small transformations which can be modeled by extending the MDL error metric to include additional terms, as summarized in Table 1 [Rao and Ballard, 1997; Rao and Ballard, 1996a; Rao and Ballard, 1998a].

Feature	Metric
Shift Invariance	$\ I - U\mathbf{r} - J\mathbf{x}\ ^2$
Spatio temporal RFs	$\Sigma \ I_t - U_t\mathbf{r}\ ^2$
Stereo	$\ I_L - I_R - \frac{dI_R}{dx}\mathbf{x}\ ^2$

Table 1: Elementary features can be created by augmenting the MDL metric.

To elaborate on one example, invariance to small shifts can be obtained by adding a first-order term. The way this can be done is to model the internal prediction as a Taylor series. Given a previously encountered reference image \mathbf{I} , a new transformed image $\mathbf{I}(\mathbf{x})$ can be predicted using a the following Taylor series expansion [Horn and Schunck, 1981; Black and Jepson, 1996]:

$$\mathbf{I}(\mathbf{x}) = \mathbf{I} + \frac{\partial \mathbf{I}}{\partial \mathbf{x}} \mathbf{x} + \text{higher order terms} \quad (2)$$

The matrix $\frac{\partial \mathbf{I}}{\partial \mathbf{x}}$ in equation 2 is an $n \times m$ matrix of partial derivatives known as the Jacobian J . As shown in [Rao and Ballard, 1998b], it is possible to approximate the Jacobian as $J = DI$ so that we can thus express equation 2 in the following equivalent form:

$$\mathbf{I}(\mathbf{x}) = U\mathbf{r} + DI\mathbf{x} + \text{higher order terms} \quad (3)$$

The goal is to estimate the object coefficients \mathbf{r} and the transformation vector \mathbf{x} for a given image and, on a longer time scale, learn appropriate basis vectors in U and D directly from the input image stream. Using equation 3, optimal estimates of the various parameters can be obtained by minimizing the following squared-error optimization function:

$$E = \|\mathbf{I}(\mathbf{x}) - U\mathbf{r} - J\mathbf{x}\|^2 + \alpha\|\mathbf{r}\|^2 + \beta\|\mathbf{x}\|^2 + \gamma\|U\|^2 + \lambda\|D\|^2 \quad (4)$$

Figure 3 shows the results of using this approximation technique with a small training set of images.

The computation of the properties observed in cortical cells may require more than one level. In a hierarchical neural network, the outputs of spatially adjacent modules at each level are also conjoined and fed as input to the next higher level. As a result, the receptive fields of model neurons become progressively larger as one ascends the hierarchical network in a manner similar to that observed in the occipitotemporal pathway [Van Essen, 1985]. At the same time, feedback connections allow top-down influences to modulate the output of lower-level modules. Such an arrangement allows a hierarchy of abstract internal representations to be learned while simultaneously allowing higher-level context to guide dynamic recognition. This setting is important for the illustration of our spike timing model, since a feature of the model is that information can be kept track of across several hierachical levels.

To illustrate the hierarchical properties of model, we trained a two-level network on a sample of natural images. Interactions between top-down and bottom-up signals as captured by the residuals in the model, were able to model complex receptive field (RF) properties analogous to

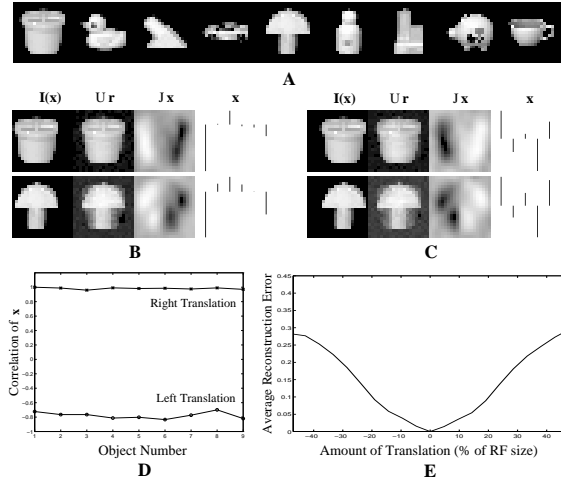


Figure 3: (A) The images (21×21 pixels) used for training the object identity network, with $\mathbf{x} = \mathbf{0}$ (B) & (C) The response of the trained networks to two different objects translated leftward and rightward respectively. In each case, the translated image is factored into the original image ($U\mathbf{r}$) and a shift ($J\mathbf{x}$) by the learning algorithm. The transformation vector \mathbf{x} (upward bar = positive value, downward bar = negative value) is approximately the same for the given translation even when different objects were translated. (D) The transformation estimates are object-invariant. The plot at the top is the correlation between the right translation vector \mathbf{x}_1 for object 1 and the right translation vectors for all other objects in the training set. The plot at the bottom is the correlation between \mathbf{x}_1 and *left* translation vectors for all training objects. (E) The basis vectors which were learned from two pixel translations of natural image patches, were tested for larger transformations. Despite being trained on only two pixel translations, the basis vectors can nevertheless represent larger translations up to ± 30 - 35% of RF size, with moderate reductions in reconstruction accuracy.(From [Rao and Ballard, 1998a])

those observed in the cortex. In the classical phenomenon of *endstopping* neurons that respond vigorously to a bar have diminished responses as the bar extends beyond the classical RF [Hubel and Wiesel, 1962; Hubel and Wiesel, 1965]. Endstopping has previously been characterized as a feedforward effect caused primarily by lateral inhibition from neighboring cells [Hubel and Wiesel, 1965; Dobbins *et al.*, 1987], but results from the model suggest that end-inhibition may be explained more generally as occurring due to predictive feedback from a higher area such as V2. As shown in Figure 4A, the typical response of a model neuron drops off sharply as a test bar extends beyond the classical RF. This occurs primarily due to diminishing level 1 residuals in the Kalman predictor network caused by better predictions from the higher level as stimulus length is incrementally increased up to the size of the larger RF of level 2 model neurons. Such an explanation is supported by the RF profiles of level 2 neurons, some of which appeared to be tuned towards long line segments. Support for the posited role of feedback in endstopping was obtained by disabling the feedback connections from level 2 to 1, which eliminated endstopping in model neurons (Figure 4B).

These simulation results suggest that many apparently complex, non-linear responses of cortical cells may be more succinctly explained by predictive feedback and cooperation within the larger context of cortical networks to which a cell is directly or indirectly connected. Though not explicitly trained on bars, the model exhibited strong endstopped responses as a consequence of the statistical structure of natural images as captured by the MDL-based objective function and its Kalman filter implementation. These model response properties were largely mediated

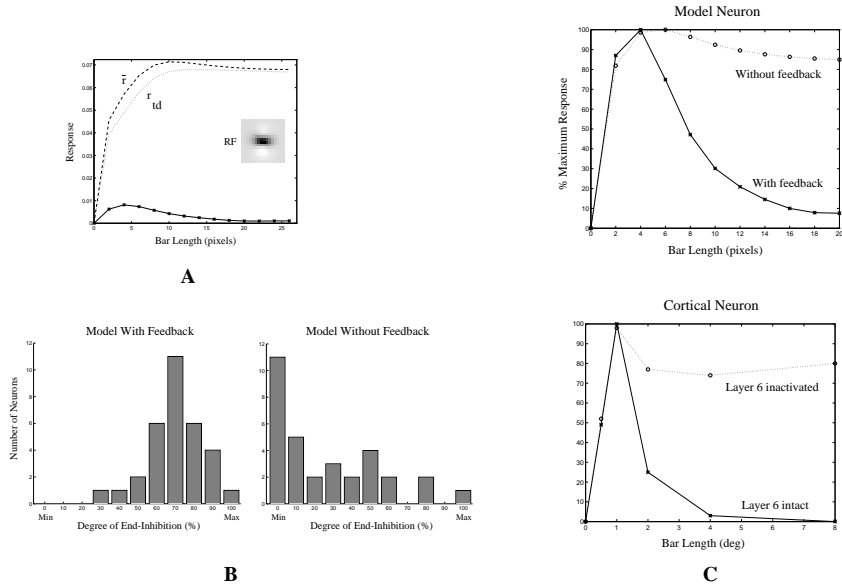


Figure 4: Simulation results. (A) A cell’s receptive field that is edge-tuned can be modulated by feedback. The uppermost trace shows the feed-forward response as a function of the length of a stimulus bar, which is modulated by the feedback response (middle trace), resulting in the phenomenon of endstopping. (B) When feedback is removed from model cells, endstopping disappears. (C) Comparisons between a model cell and a representative recording from a cell in cat visual cortex show striking agreement.

by top-down feedback. Thus, a general prediction of the model is that inactivation of a higher cortical area such as V2 or MST should cause substantial disinhibition or reduction in response suppression by stimuli from beyond the classical receptive field in lower areas such as V1 or MT, respectively.

3 Individual Spike Models

The central ideal of this paper is that spike timing can be harnessed to implement general computations. Thus all the examples in section 2 can be done with our timing model. To develop this model, we first define its general feedforward properties and then add the feedback details needed to implement predictive coding.

The central idea behind coincident firing is that neurons are vastly more efficaciously fired if their arriving spikes are precisely timed[Abeles, 1991]. Furthermore this signal can be made reliable in that each individual neuron can be fired with great certainty. Given these conditions, it is easy to see that a neuron A can be made to fire a neuron B if there is a chain of intervening neurons each of which receive synchronous input from the neurons earlier in the stage as shown in Figure 5. The synfire chain provides fast communication as the neurons in the chain are fired with a single arriving volley of spikes. As the spike velocities are on the order of 5-8cm/milliseconds, thus communication along axons (neglecting delays at the somas) takes on the order of a millisecond.

To propagate such a signal reliably, there must be several neurons at each synchronous

stage. This is to ensure that the firing chain is not broken. To formalize this let the number of neurons in each stage be N and those used to propagate the signal be w . Then the neurons so involved at each stage constitute a *path* where w is its width. A key idea is that neurons used to construct paths are shared. It turns out that this can be done in a way that different signals do not interfere. Figure 5C shows a path of width three denoted by the color red. As shown by Figure 5A, this representation scheme is very difficult to interpret at the single cell level. The traffic through a single neuron represents the total of the active paths within a given time window. Furthermore that traffic is non-stationary in that different time windows may yield very different traffic patterns. How can the message carried by a neuron be interpreted given that its collection of spikes represents many different signals? The interpretation is actually easy but requires that one have access to all the neurons at a stage on a path. If that is possible then the interpretation of the path is just the conjunction of all the spike trains. This is shown in Figure 5B. the content of the message is difficult to appreciate in general but from single cell recordings early in the cortical visual pathway one can infer properties of this encoding scheme.

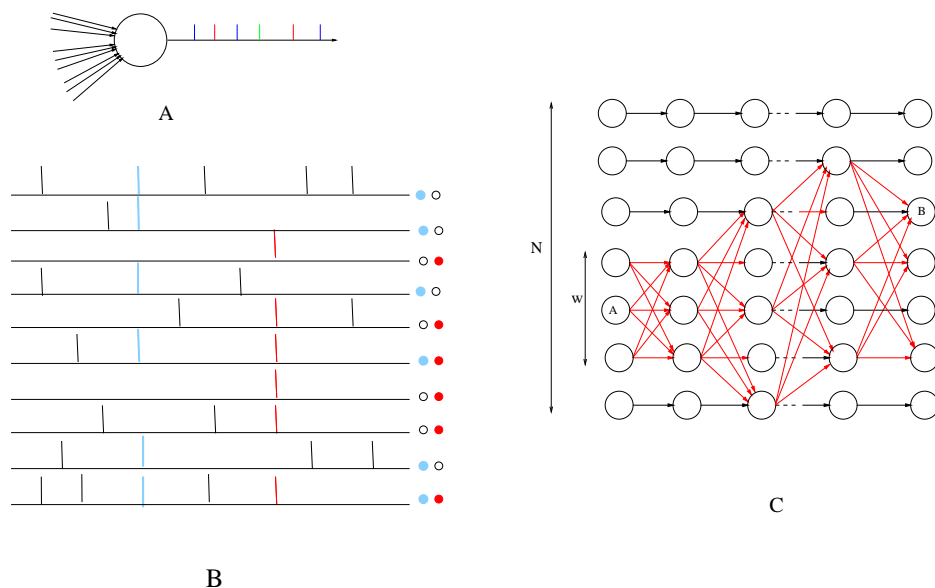



Figure 5: **Paths.****A. The idea of shared coding along paths.** The input to a neuron is in terms of synchronous volleys, indicated as colors on the diagram. These are superimposed in the spike train. The result is that the spike train contains the disjunction of all its inputs. In the general case, the different spikes can represent very different processes. **B. The basic idea of distributed coding.** The spike train from a single neuron is very difficult to interpret as its meaning is conveyed by the other neurons in the stage. However if the neurons at a stage in the path are known, then the meaning of the path can be recovered by intersecting all the individual spike trains. Here two processes denoted by red and blue can be interpreted by intersecting the spikes on their respective paths, as highlighted by the code on the extreme right. **C.** The input to a neuron is assumed synchronous. Its output is propagated reliably because it is part of a group of neurons on a path. For the red path shown, interference would come when a rival path or set of paths connected to the neurons used by those on a path, or alternatively through random connections. For very wide paths of 100 neurons, interference is extremely unlikely. We contend that through learning, lengthy, wide paths can be created that are secure against both kinds of interference.



Receiver	Sender w		
	10	30	50
10	0	3,932	16,726
20	0	0	279
100	0	0	0 (100)

Table 2: **Left. Sender and Receiver Stages Right. Expected number of nodes in a path under random conditions.** Given the number of nodes w , a biological assumption of $N = 20,000$ neurons per stage, one can calculate the number of neurons that will have a given multiplicity of connections at the recipient stage (denoted as *Receiver* given the number of neurons in the sending stage. The number in red denotes the connectivity assumed by the opur model which assumes non-random connections. [Abeles, 1991]

3.1 Paths

The performance of a path is crucially dependent upon its width. A path of width three might be easily triggered by random neural traffic unrelated to that particular path. Once triggered the connections within the path assure that the synchronous volley is propagated. While a path width of three is not secure, much wider paths are extremely secure, because the reliability of a path increases dramatically with its width. The only problem is that, as Abeles [Abeles, 1991] has shown, with random connections, paths will not have the width that we postulate. Abeles's calculations are shown in Table 2a. Assuming $N = 20,000$ neurons per stage, paths of width 100 are impossible to occur at random. In contrast, we assume that paths with many stages of width 100 can be reliably created (Table 2b). The number 100 is somewhat arbitrary in that any large number that was still a fraction of the total synapses would work (However see Figure 6). Consequently the width of a path may be expected to vary. However for the purpose of analysis we pick a fixed width of 100 at each stage. To make paths of width 100 that have a useful length, the brain must be designed to form such paths. We contend that wide paths that would be impossible to form on by chance can be formed by a synchrony-seeking learning algorithm [Ballard, 1999]. Given a Hebb-like rule for synaptic strengthening that rewards synchrony, all the synapses should tend to become sensitive to synchronous volleys. For the current exposition we assume that such a rule works and that wide and long paths are possible.

Given a total of N neurons per stage and w neurons on a path in a stage, how many paths pass through a stage? A simple analysis shows that there are $\binom{N}{w}$ of them. For $N = 10,000$ and $w = 100$ this is on the order of

$$10^{400}$$

This calculation uses a very conservative value for N . The total number of cortical cells is on the order of 10^{11} so a stage could use many more neurons. Nonetheless even with $N = 10^4$ the number is enormous. Thus the capacity as measured in the number of possible paths is for all practical purposes unlimited. This is main advantage of communicating via paths, which are subsets of the neurons in a stage, since the number of possible subsets is so huge.

The crucial and perhaps surprising property of paths is that many such active paths can co-exist without interfering with each other. Paths potentially will interfere if the neurons defining one path completely overlap at any one stage with the set of neurons in the others. Let us term the coherent signal down an active path as a *volley*, and term the paths that are carrying volleys at some interval of time as *active*. While there are a huge number of paths, they should not all be active at a time. The reason for this is that for large numbers of active paths, the probability that an inactive path will be accidentally triggered approaches certainty. As active paths are presumably used for relevant computation, the security of active paths is of essential importance. There is already evidence that the cortical memory tends to pick codings that are independent [Olshausen and Field, 1997], but here we are interested in estimating the properties of the path that would make it secure.

A set of active paths could potentially excite a particular inactive path if:

1. The path lies in the union of the other active paths, and
2. Volleys from the other active paths are timed to coincide.

The probability of both these conditions being met is difficult to compute exactly, but a ready estimate is that if the number of active processes is k and that they require width w , then the probability of path interference is bounded by

$$\left(1 - \left(1 - \frac{w}{N}\right)^k\right)^w$$

This is plotted as a contour plot in Figure 6. The figure shows that as long as the number of active processes combined with width is low enough, then the probability of interference is zero, as shown in the lower left triangle. But as either or both increase the probability of interference becomes one as shown by the upper right triangle. To derive this estimate, consider the probability that one neuron in a path will survive interference by k processes. This is $\left(1 - \frac{w}{N}\right)^k$. Thus the probability that it will be interfered with is one minus this quantity. Consequently the probability that all the neurons on a path will be interfered with is this raised to the w power. In this calculation we have neglected the additional helpful constraint of needing all the spikes to be coincident, but this is a small factor of on the order of 10^{-2} .

3.2 Synchrony and the Post-stimulus Histogram

A challenge for any proposal to that use synchronous firing is that a very large amount of experimental data, mostly in the form of post-stimulus histograms (PSHs), seems to imply that the neural code is asynchronous. Although it might seem that the synchronous coding is incompatible with the observations of classical tuning of receptive fields observed in single-cell recording experiments, the experimental observations can be reconciled if the idea of volleys is combined

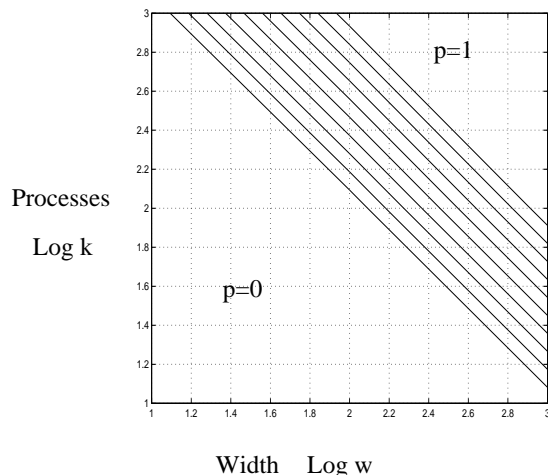


Figure 6: **The simple model** Path interference predicts enormously reliable behavior that can become catastrophically degraded as either loading (the number of coincident processes) or the path width increases. Below the contour lines the probability of interference is 0 and above the contour lines the probability is 1. The increments for the contour lines are 0.1.

with distributed parameter representations. Neurons in the cortex seem to sample parameter spaces [Ballard, 1986]. This representation is expensive for high dimensional spaces. Consider a d dimensional space with grain n . The potential number of neurons needed is n^d . This becomes prohibitive expensive for dimensions greater than about five when n is about 100.

Enormous efficiencies can be realized if all the measurement cells are not used at once. As shown by Hinton [Hinton, 1981], if overlapping receptive fields are used, then the number of cells can be reduced. The basic strategy is shown in Figure 7. The strategy depends on the fact that individual signals in the space of possible signals are sparse. In that case increased accuracy can be gained by notionally intersecting the coarse fields of cells that are firing.




Figure 7: **The basic idea of distributed coding.** An accurate parameter value can be signaled by the conjunction of neurons, each tuned to a coarse value of the parameter. A) Using small overlapping receptive fields to signal location means that a single receptive field carries the location information. B) Using receptive fields whose width is large with respect to the sampling interval means that multiple neurons carry location information.

Formally the number of neurons needed for a given accuracy are reduced by a factor of

$$\left(\frac{1}{W}\right)^{d-1}$$

where W is the width of the larger receptive field measured in units of n . For the most efficient case the number of cells that need to represent a stimulus is on the order of d , but in practise more can be used with increasing reliability.

If we assume a redundant encoding, then for any stimulus, several different sets of neurons can potentially signal the stimulus value. The stimulus can be recovered by a weighted average of the parameters that the neurons are tuned for, but in reality this need never be done and the distributed code can be used directly.

Suppose that these different and overlapping sets are connected in paths. When different paths send volleys, the neurons that share paths exhibit the combined traffic. To put it another way, any given volley will use a subset of the neurons whose receptive fields overlap the stimulus, and most important, these subsets typically intersect. An individual neuron may be a member of different subsets at different times. The traffic through an individual neuron is higher the more the neuron's receptive field is central to the stimulus parameter, and thus neurons may appear to have a "firing rate" encoding. In the terminology of paths, the number of active paths that pass through a single neuron is a function of that neuron's tuning. This effect of multiple paths on a neuron's pattern of activation is shown in Figure 8. In Figure 8B two overlapping groups of neurons signal a parameter by sending separate volleys. The neuron that is a member of both paths exhibits the combined traffic. Owing to the fact that the different volleys have start times that are not precisely controlled, the combined traffic may appear to be random.

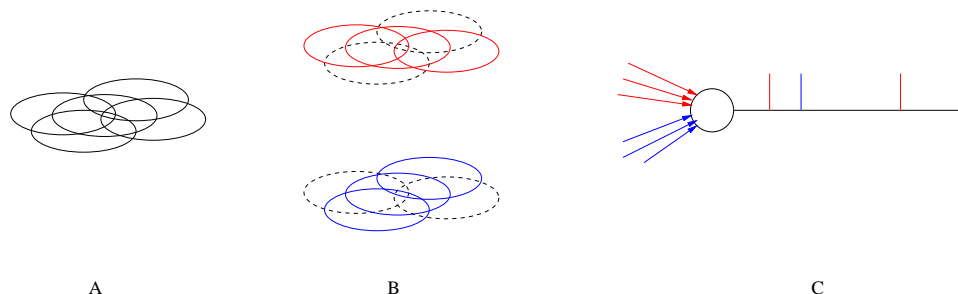


Figure 8: **How the disjunctive coding relates to classical receptive field measurements.** (A) For a given parameter such as orientation, many different neurons are tuned to a given stimulus. (B) Different groups of these neurons should signal in volleys, being connected with paths. Here the red and blue subsets of the overlapping receptive fields are assumed to fire synchronously. However the neurons that are best tuned to the parameter will typically be part of many paths. (C) The best tuned neuron is part of many subsets. Thus the spike train of the best tuned neuron reflects the disjunction of these subsets, in this example both the red and blue signals, and exhibits a "firing rate" encoding.

We can interpret the output of the paths through a neuron in certain constrained situations in the cortical sensory periphery. Consider the formation of simple cell receptive fields. The goal of the cortex is to provide a memory. The way this is assumed to be done is to encode behaviorally-relevant correlations. In the sensory periphery, such as the LGN, an image patch is encoded by ON-center and OFF-center cells. This ensemble of cells codes represents higher order correlations than are captured by individual cells, but only implicitly. What needs to be done is to have a way of discovering higher order correlations.

If it were practical to interpret the output of a cell as a firing rate, which in turn transmitted an analog value, then the mathematics of finding this correlation would be straightforward. However we posit that the higher order statistics are being communicated in a different way.

A mechanism at the LGN constrains the output to be in synchronous volleys. For each volley the outputs of the individual cells are sampled and encoded as a spike or silence. The multicell spike data from a single volley can be thought of as a binary image \mathbf{B}_t . This packet represents the higher order statistics in that the probability density function of a recipient neuron can be written as

$$\langle \mathbf{B}'_t \rangle$$

where \mathbf{B}'_t is the subset of \mathbf{B}_t that triggers the neuron. Interestingly, this is the function measured by the reverse correlation technique. See e.g. [Reid *et al.*, 1997].

It should be obvious that the sum of the volleys selected according to the receptive field interpreted as the mean of a probability density function(pdf) converges to an estimate of the mean of a probability density function(pdf). Figure 9 shows the original pdf mean and the pdf approximation mean obtained by averaging 100 volleys. The receptive field of a cell may be interpreted as a filter, passing volleys that could have been generated by its pdf. In addition, if the traffic of a neuron were repeated volleys, and the phase between volleys was not precisely controlled, the firing rate of any particular neuron could be interpreted as a rate code. Thus the synchronous firing strategy could appear to be a rate code when measured by conventional techniques.




Figure 9: **Probability Density Function Approximation.** The time average of the activated synapses is the correlation function. (A) Simulated receptive field of a neuron (B) The same receptive field measured by averaging 100 of its volleys.

4 Feedback

A crucial problem for paths is that of handling feedback. This is a challenging problem because the feedback will necessarily be delayed by several means, and cannot arrive coincident with the spike that caused it. The solution we adopt is to delay the next input by an equal amount. Thus the feedback spike arrives at the same time as the subsequent spike for the relevant path.

Of course the feedback is also encoded as part of the relevant path. What this means is that you can effectively think of the entire path including its feedback branches as a separate *sampled data system*. In standard sampled data systems, time is discrete and the system variables are analog quantities. Here is a special case where the variable are also discrete being individual spikes. Thus in a negative feedback circuit, a volley may set up a feedback chain that cancels a subsequent volley. Since the propagation time is small most of the delays must be due to the neuron's soma. Figure 10A shows this case. The important point here is that now we are

adding the supposition that the volleys, where they use feedback for modification, are sent in synchronous packets, that is sets of volleys, each separated by a delay Δ . The length of the set need only be necessary for the current computation.




Figure 10: **Handling Feedback.** Owing to processing times, there will be delays at neuron's somas. Thus in a feedback circuit, the feedback will arrive delayed with respect to the input. The solution is to delay the subsequent input volley by an equal amount. The feedback delay Δ is by design to make the sample data system updating work and should not be confused with smaller delays Δ needed to adjust synchrony. In a hierarchical system, such as the cortex, neurons in subsequent maps are connected with delay circuits as shown.

Leaving the neural context aside for a moment, in an analogue feedback circuit without delays, one can write the relationships at each node in the feedback network in terms of an update equation that reflects the current state. That is where r is the state vector, $\dot{r}(t) = F(r(t))$. In the neural model these equations must have two important specializations. In the first place they must account for the aforementioned delays needed to handle the synchronous signals. Secondly, the model must recognize that the state vector is binary, representing the presence or absence of a spike. Thus

$$r_t = F(r_{t-\Delta}), r \in \mathbf{0}, \mathbf{1}$$

The binary model means that the feedback circuits must be combined with the distributed encoding strategies for parameters to build circuits that can handle non-binary quantities. Also, the feedback system just described introduces a new potential source of error. If for any reason a cell in a path is *inhibited* at the wrong moment, the volley will be stopped. For that reason it would be prudent to use a threshold for the cell that is less than that produced by the entire volley. However for wide paths this does not effect the reliability calculations qualitatively.

By incorporating this strategy, circuits can be built that manipulate analog quantities according to algorithmic rules. In other words, this strategy provides a general communication protocol upon which behaviorally useful operations can be coded. Circuits can exploit the security of overlapping paths to run in parallel. They each need to be sampled at the rate Δ , but they can share the same networks as long as they are out of phase. Figure 10 sketches this for a hierarchical network. Thus in principle, vast numbers of paths can exist independently in the brain, even when they have feedback loops, without interfering with each other. A key assumption is that the delay Δ used in the feedback circuit, is long with respect to individual traffic. Thus in the interim specified by Δ , many other signals can be sent that represent other variables.

Let us make these concepts more concrete with a specific example. Consider the problem of learning specific models of image data. We suppose that the task is to have a predictive

model that can reconstruct incoming data. This has already been developed in the context of a rate model[Rao and Ballard, 1998c]. Here we show how the model can use individual spikes along paths as the signaling method. The additional assumption we need is that the volleys generated by the incoming data from the thalamus are repeated that is they are in the form $\{\mathbf{B}_1, \mathbf{B}_2, \dots, \mathbf{B}_T, \mathbf{B}_1, \mathbf{B}_2, \dots, \mathbf{B}_T, \mathbf{B}_1, \mathbf{B}_2, \dots, \mathbf{B}_T, \dots\}$.

In the context of the synchronous firing model, given a set of volleys $\{\mathbf{B}_1, \mathbf{B}_2, \dots, \mathbf{B}_T\}$, the goal of the model is to produce *answering feedback volleys* $\{\mathbf{r}_{1+\Delta}, \mathbf{r}_{2+\Delta}, \dots, \mathbf{r}_{T+\Delta}\}$ such that

$$\mathbf{B}_t = U\mathbf{r}_t$$

Here U is a matrix of synapses connecting the units \mathbf{r} to the units that produce the $\{\mathbf{B}_t\}$. The value of \mathbf{r}_t may be 1 or 0 reflecting spike or no spike respectively.

Following [Rao and Ballard, 1998c] one can use the Minimum Description Length principle to trade off reconstruction error cost with the reconstruction units. This leads to minimizing an error of the form

$$E = \sum_{t=0}^T ((\mathbf{B}_t - U\mathbf{r}_t)^T (\mathbf{B}_t - U\mathbf{r}_t) + \alpha \mathbf{r}_t^T \mathbf{r}_t + \beta \|U\|^2)$$

in the above expression, the first term measures reconstruction error and the next two measure the cost of the machinery doing the reconstruction. The term $\mathbf{r}^T \mathbf{r}$ measures the cost of producing spikes and the term $\|U\|^2$ measures the cost of the synaptic memory. The appropriate synaptic weights may be found by training on an image data set. Each image is converted to a set of \mathbf{B}_t which are used in training. The algorithm computes a stable set of responses \mathbf{r}_t by using a form of gradient descent on E . At each iteration, \mathbf{r}_t changes the state of its units if doing so will lower error. Once the responses converge, the adjustments are made to the synapses according to

$$\dot{U} = -\eta \frac{\partial E}{\partial U}$$

When the network has been trained, a volley \mathbf{B}_t leads to an initial \mathbf{r} via the feedforward pathway $U^T \mathbf{B}_t$. Next the feedback signal $U\mathbf{r}$ is used to try and cancel the $\mathbf{B}_{t+\Delta}$. The error $U^T (\mathbf{B}_{t+\Delta} - U\mathbf{r})$ is used to modify r and the process continues.

To illustrate these principles, a network of 32×32 input units and 100 coding units was trained on three images of toy animals. The volleys were artificially created by using the binary encoding of the image, that is

$$\mathbf{B}_i^k(x, y) = \begin{cases} 1 & \text{if the } i+1 \text{ th bit of } I^k(x, y) \text{ is } 1 \\ 0 & \text{otherwise} \end{cases}$$

To train the synapses, the three sets of binary images were used in a repeated sequence. The learning rule is based on gradient descent. The state of a neuron (0 or 1) will be changed if doing so will improve the objective function. The synapses were shared between all three images so that they have to work for each volley from each image in the set. Note that the input is assumed to be appropriately interlaced. that is, for the k -th image, the input volleys appear in repeated sequence. Thus for image k ,

$$\dots, \mathbf{B}_0^k, \mathbf{B}_1^k, \mathbf{B}_2^k, \dots, \mathbf{B}_7^k, \mathbf{B}_0^k, \mathbf{B}_1^k, \mathbf{B}_2^k, \dots, \mathbf{B}_7^k, \mathbf{B}_0^k, \dots$$

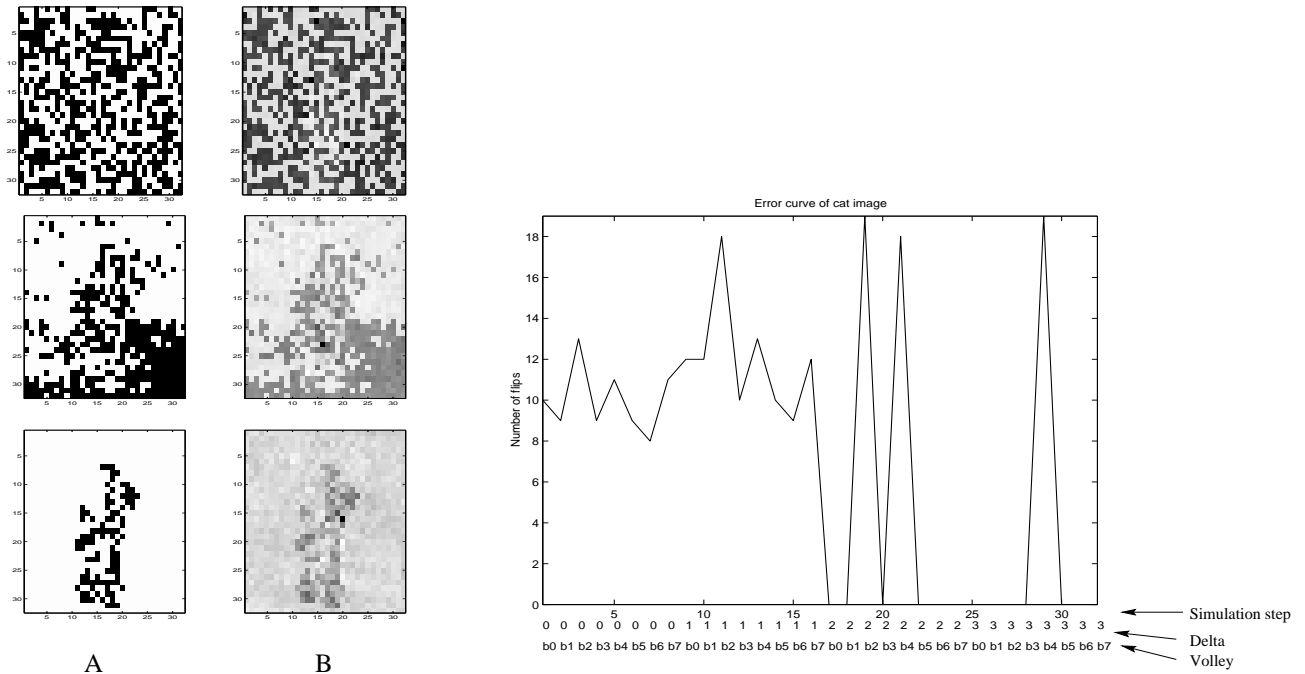


Figure 11: **Learned Reconstructions** For this example the volleys are on the order of 500 neurons. (Left After training on two input images, the network learns to cancel their input volleys. A) Binary images B_0 , B_2 , and B_4 . B) The corresponding cancelling volleys. Right The ensemble average of state changes summed over all the representational neurons. When a volley arrives at a neuron, it will send a spike or not. “Flips” counts the number of times the state changed from the previous time period. Note that most of the volleys (B_0 , B_3 , B_5 , B_6 , B_7) have converged after one iteration and all the volleys have converged after three iterations.

This is done for a time that simulates the time that the k -th image would be on the retina.

The results show clearly that after training, the network can generate feedback volleys that cancel the incoming binary images. Figure 11 shows a subset of the binary images and the return signals Ur that cancel them. The estimates Ur_t are analog values but can be thresholded to produce near exact copies of the incoming spike trains. Note that the algorithm would work for any binary encoding strategy as long as it is repeated over time for a given input.

While Figure 11 shows that predictive coding can reconstruct images, but it also contains additional structure. Since [Castelo-Branco *et al.*, 1998] observed coincident firing with multicell electrode recordings, our model should be reconcilable with their data. To compare the two we plot a cross-correlation diagram between pairs of model neurons that have exhibited coincidence. This is shown in Figure 12. The left hand side figure shows the cross-correlation using a window of 80 milliseconds (Our simulation uses one spike per 8 milliseconds). On the right hand side, the progress of the correlogram through time is shown. Here we have assumed that the overall process is ergodic, that is that the spatial averaging used in the model is equivalent to the temporal averaging used in the experimental data.

These simulations show two features that are not in the Singer laboratory data. First the signal is constant instead of phasic, and second, the correlation peaks are much more sharply tuned.

The reason that our signal is constant is that the model has only two levels. In the cortex,

the responses plotted here would be potentially recoded by higher levels. When this coding was accomplished, the signal should be much diminished, and be phasic like the experimental data.

The second difference is more interesting, as reconciling it predicts a feature of the cortex. The correlation function in the experimental data has the form of a Gabor function can be described as

$$g(t) = A \cos 2\pi\omega t \times e^{-\frac{t^2}{2\sigma^2}}$$

Since multiplication in the time domain corresponds to convolution in the frequency domain, $G(u)$ can be written as

$$G(u) = A'(e^{i2\pi\omega t} + e^{-i2\pi\omega t}) \otimes e^{-2\sigma^2\pi u^2}$$

The experimental data implies that σ is approximately 40 milliseconds, or that the frequency modulation is about ± 4 Hz. At 33Hz this implies that the temporal jitter is on the order of ± 3 milliseconds. So if our simulation data were allowed to have a small random phase drift with a standard deviation of this amount, the simulation data would be identical to the experimental data.

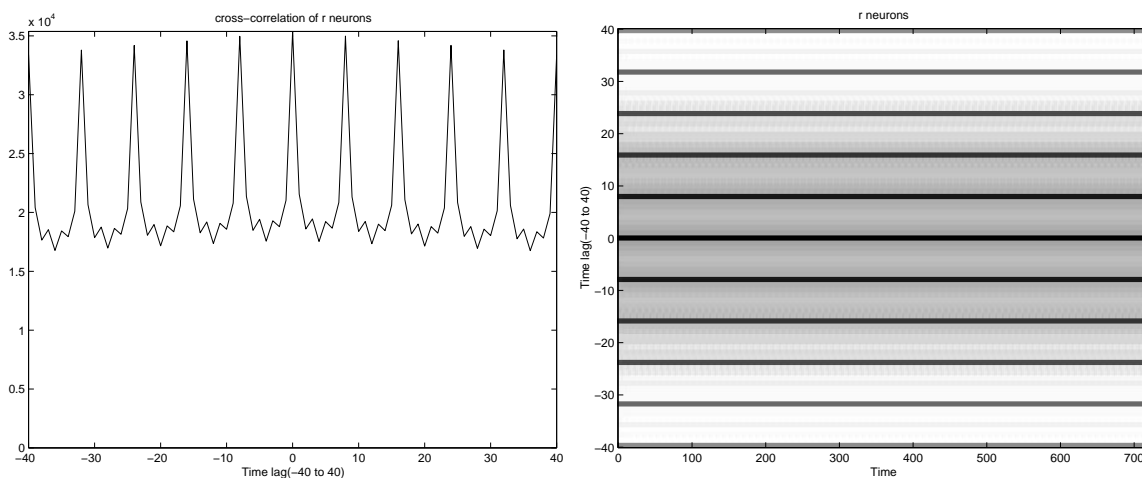


Figure 12: **Cross-correlogram for model units** **Left** Crosscorrelogram shows that spikes are extensively synchronized **Right** This correlation persists for as long as the input is on the model retina.

5 Summary

The simplest way to understand the timing model is to temporally ignore the timing issues associated with feedback. In that case the properties of the model are apparent by examining a single stage of neurons on a path. By communicating in subsets, a particular computational process can be made sufficiently unique such that one the order of one hundred other processes can share the neurons in the stage without interfering with each other or the first process. Recurrent or feedback circuits add the additional complication that some process sampling delay has to be instituted to assure that all the computations from any given process can be

kept separate from the rest of the computations. Evidence suggests that the cortex does institute delays for processes but does not have to use an exactly the same delay for all active processes. The non-interference property at each stage, allows different processes to overtake each other in time without consequence.

Our extension of predictive coding models to single spikes is controversial in that it places demands on both the functioning of the Cortico-LGN circuit and the timing of spikes. The cortico-LGN circuit has to be structured to generate repeated volleys that are coding the LGN input. Recent research suggests that the LGN may generate volleys. There are also suggestions that they can be repeated, but this evidence comes from developmental studies and the time constants are too large by a factor of 10^2 . However there is no evidence against the idea of repeated volleys on the timescales need by the model. The spike timing constraint would have once been considered controversial, but a raft of recent theoretical and experimental studies suggest that precise timing is possible. A key assumption is that neighboring spikes are relatively independent. assuming that they are part of a rate code from the outset leads to very different conclusions [Panzeri *et al.*, 1999].

Despite the above evidence of a cortical system compatible with the coincident firing models, one could still argue that the biophysical properties of the neurons themselves do not allow for the necessary precise timing. However evidence from slice recordings shows that neurons fire with millisecond reliability when the stimulated with realistic simulated input [Mainen and Sejnowski, 1995].

The tightest constraint on our model comes from the requirement that the cortico-thalamic feedback create repeated copies of the binary volleys in the feedforward pathway. The evidence for this comes from developmental experiments on ferrets. These experiments show that cortico-thalamic connections produce synchronous firing at approximately 10 sec intervals. Severing these connections eliminates the synchrony. The 10 sec interval is over 100 times longer than the 10-100 milliseconds needed by the proposed model, but the fact that the posited mechanism exists at such a slow time scale, far below the natural timescales of individual neurons, implicates its importance as a computational mechanism. There seems to be no reason why this mechanism could not simply be speeded up post-development as required by the model.

A final point on timing constraints, is made by the cortical time constants observed with feedback signals. In cortical area V1, the time to experience feedback can be as long as 100-200 milliseconds [Zipser *et al.*, 1997]. This is much longer than the theoretical time to propagate the signal through the neurons themselves. The suggestion is that the delays are planned and serve a computational purpose.

The extension of the predictive coding model to single spikes is significant for many reasons.

1. It allows simulations of the model to be directly compared to the statistics of experimental data, rather than rate codes which are a summary form that blurs temporal distinctions. Herein we showed that the simulations that the responses converged in 1 to 3 iterations, and that the Gaussian envelope in the Singer data could be indicative of temporal jitter in the sampling.
2. In other studies we have shown that rate coded version of our predictive model can be used for segmentation. The extension of this formulation to single spikes is also straightforward and would allow an understanding of how single spikes are used in vision tasks, such as image segmentation.

3. It provides an information-theoretic description of cortical hierarchies. The cortex can be understood as a biological memory that calibrates itself by trying to predict its input. In the course of doing so it trades off the goodness of the prediction against the cost of the circuitry doing the prediction.
4. Our model predicts that there are an exponential number of parallel paths in the cortex. This vast amount of available computation is more than is needed to account for any individual ongoing behavior, but is extraordinarily significant, as it opens up ways of having multiple processes active, only some of which are coupled to behavior. The additional computing can be used for mental simulation and memory management, functions usually not considered in neural models, but vital to overall brain function.

References

- [Abbott and Dyan, 1999] L. F. Abbott and Peter Dyan. The effect of correlated variability on the accuracy of a population code. *Neural Computation*, 11:91–101, 1999.
- [Abeles, 1991] Moshe Abeles. *Corticonics*. Cambridge University Press, 1991.
- [Alonso *et al.*, 1996] Jose-Manuel Alonso, W. Martin Usrey, and R. Clay Reid. Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature*, 383:815–819, 1996.
- [Ballard, 1986] Dana H. Ballard. Parameter networks. *Behavioral and Brain Sciences*, 32:333–377, 1986.
- [Ballard, 1999] Dana H. Ballard. Learning synchrony. *Tech Report in Preparation*, 1999.
- [Black and Jepson, 1996] M.J. Black and A.D. Jepson. Eigentracking: Robust matching and tracking of articulated objects using a view-based representation. In *Proc. of ECCV*, pages 329–342, 1996.
- [Castelo-Branco *et al.*, 1998] Miguel Castelo-Branco, Sergio Neuenschwander, and Wolf Singer. Synchronization of visual responses between the cortex, lateral geniculate nucleus, and retina in the anesthetized cat. *The Journal of Neuroscience*, 18:6395–6410, 1998.
- [Dobbins *et al.*, 1987] A. Dobbins, S.W. Zucker, and M.S. Cynader. Endstopped neurons in the visual cortex as a substrate for calculating curvature. *Nature*, 329:438–441, 1987.
- [Felleman and Van Essen, 1991] D.J. Felleman and D.C. Van Essen. Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1:1–47, 1991.
- [Field, 1987] D. J. Field. Relations between the statistics of natural scenes and the response properties of cortical cells. *J. Opt. Society of America*, 4:2379–2394, 1987.
- [Gallant *et al.*, 1998] J. L. Gallant, C. E. Conner, and D. C. Van Essen. Neural activity in areas v1, v2 and v4 during free viewing of natural scenes compared to controlled viewing. *Neuroreport*, 9:2153, 1998.
- [Hinton, 1981] G. Hinton. Distributed coding. *Int. Joint Conf. on AI*, 1981.

- [Horn and Schunck, 1981] B.K.P. Horn and B.G. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
- [Hubel and Wiesel, 1962] D.H. Hubel and T.N. Wiesel. Receptive fields, binocular interaction, and functional architecture in the cat’s visual cortex. *Journal of Physiology (London)*, 160:106–154, 1962.
- [Hubel and Wiesel, 1965] D.H. Hubel and T.N. Wiesel. Receptive fields and functional architecture in two non-striate visual areas (18 and 19) of the cat. *Journal of Neurophysiology*, 28:229–289, 1965.
- [MacKay, 1956] D.M. MacKay. The epistemological problem for automata. In *Automata Studies*, pages 235–251. Princeton, NJ: Princeton University Press, 1956.
- [Mainen and Sejnowski, 1995] Z. F Mainen and T. J. Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268:1503, 1995.
- [Maybeck, 1979] P.S. Maybeck. *Stochastic Models, Estimation, and Control (Vols. I and II)*. New York: Academic Press, 1979.
- [Meister, 1996] Markus Meister. Multineuronal codes in retinal signaling. *Proceedings of the National Academy of Sciences*, 93:609–614, 1996.
- [Mignard and Malpeli, 1991] M. Mignard and J.G. Malpeli. Paths of information flow through visual cortex. *Science*, 251:1249–1251, 1991.
- [Mumford, 1994] D. Mumford. Neuronal architectures for pattern-theoretic problems. In C. Koch and J.L. Davis, editors, *Large-Scale Neuronal Theories of the Brain*, pages 125–152. Cambridge, MA: MIT Press, 1994.
- [Murphy and Sillito, 1987] P.C. Murphy and A.M. Sillito. Corticofugal feedback influences the generation of length tuning in the visual pathway. *Nature*, 329:727–729, 1987.
- [Nicholls *et al.*, 1992] John G. Nicholls, A. Robert Martin, and Bruce G. Wallace. *From Neuron to Brain*. Sinauer, 1992.
- [Olshausen and Field, 1997] B. A. Olshausen and D. J. Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision Research*, 37:3311–3325, 1997.
- [Panzeri *et al.*, 1999] Stefano Panzeri, Simon R. Schultz, Alessandro Treves, and Edmund Rolls. Correlations and the encoding of information in the nervous system. *Proceedings of the Royal Society B*, 1999.
- [Prut *et al.*, 1998] Yifat Prut, Eilon Vaadia, Hagai Bergman and Iris Haalman, Hamutal Slovin, and Moshe Abeles. Spatiotemporal structure of cortical activity: Properties and behavioral relevance. *Journal of Neurophysiology*, 79:2857–2874, 1998.
- [Rao and Ballard, 1996a] R. P. N. Rao and D. H. Ballard. A class of stochastic models for invariant recognition, motion, and stereo. Technical Report 96.1, National Resource Laboratory for the Study of Brain and Behavior, Department of Computer Science, University of Rochester, June 1996.

- [Rao and Ballard, 1996b] Rajesh P.R. N. Rao and Dana H. Ballard. Dynamic model of visual processing predicts neural response properties of visual cortex. *Neural Computation*, 9:721–763, 1996.
- [Rao and Ballard, 1997] R. P. N. Rao and D. H. Ballard. Efficient encoding of natural time varying images produces oriented space-time receptive fields. Technical Report 97.4, National Resource Laboratory for the study of Brain and Behavior, Computer Sci. Dept., University of Rochester, August 1997.
- [Rao and Ballard, 1998a] R. P. N. Rao and D. H. Ballard. Development of localized oriented receptive fields by learning a translation-invariant code for natural images. *Network: Computation in Neural Systems*, 9(2):219–234, 1998.
- [Rao and Ballard, 1998b] R. P. N. Rao and D. H. Ballard. What/where networks and local receptive fields for transformation estimation. *Network*, 9:219–234, 1998.
- [Rao and Ballard, 1998c] Rajesh P. N. Rao and Dana H. Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive field effects. *Nature Neuroscience*, 2:79–87, 1998.
- [Reid and Alonso, 1995] R. Clay Reid and Jose-Manuel Alonso. Specificity of monosynaptic connections from thalamus to visual cortex. *Nature*, 378:281–284, 1995.
- [Reid *et al.*, 1997] R. C. Reid, J. D. Victor, and R. M. Shapley. The use of m-sequences in the analysis of visual neurons: Linear receptive field properties. *Visual Neuroscience*, 14:1015, 1997.
- [Rieke *et al.*, 1997] Fred Rieke, David Warland, Rob de Ruyter van Steveninck, and William Bialek. *Spikes*. MIT Press, 1997.
- [Sandell and Schiller, 1982] J.H. Sandell and P.H. Schiller. Effect of cooling area 18 on striate cortex cells in the Squirrel monkey. *Journal of Neurophysiology*, 48(1):38–48, 1982.
- [Shadlen and Newsome, 1998] Michael N. Shadlen and William T. Newsome. The variable discharge of cortical neurons: Implications for connectivity, computation and information coding. *The Journal of Neuroscience*, 18:3870, 1998.
- [Sillito *et al.*, 1994] Adam M. Sillito, Helen E. Jones, George L. Gerstein, and David C. West. Feature-linked synchronization of thalamic relay cell firing induced by feedback from visual cortex. *Nature*, 369:479–482, 1994.
- [Thompson *et al.*, 1997] K. G. Thompson, N. P. Bichot, and J. D. Schall. Dissociation of visual discrimination from saccade programming in macaque frontal eye field. *Journal of neurophysiology*, 77:1046, 1997.
- [Usrey *et al.*, 1998] W. Martin Usrey, John B. Reppas, and R. Clay Reid. Paired-spike interactions and synaptic efficacy of retinal inputs to the thalamus. *Nature*, 395:384–387, 1998.
- [Van Essen, 1985] D.C. Van Essen. Functional organization of primate visual cortex. In A. Peters and E.G. Jones, editors, *Cerebral Cortex*, volume 3, pages 259–329. Plenum, 1985.

[von der Malsburg, 1995] Christoph von der Malsburg. Binding in models of perception and brain function. *Current opinion in neurobiology*, 5(4):520, 1995.

[Zipser *et al.*, 1997] Karl Zipser, Victor A. F. Lamme, and Peter N. Schiller. Contextual modulation in primary visual cortex. *Journal of Neuroscience*, 16:7376, 1997.