

1 Introduction

Recently there has been a huge increase in experiments that imply that neurons communicate with synchronous volleys. Evidence from slice recordings shows that neurons fire with millisecond reliability when stimulated with realistic simulated input [MS95]. In the retina, the efficacy of synchronous codes has been noted as a way of retinal encoding [Mei96]. Meister shows how retinal ganglion synchronous codes trade off communication rate with increased signal fidelity. Reid, Alonso and Usrey have observed precise timing in cortico-thalamic connections [RA95, AmUR96, URR98]. Additional evidence that volleys might be generated comes from experiments by [SJGW94]. In 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, Casteo-Branco et al [CBNS98] observed correlations between the retina LGN and cortical areas 17 and 18 in the anesthetized cat.

To interpret this data in a general way requires a theoretical framework. Recently progress has been made with predictive coding models [OF97, Fie87, RB96, RB98]. The brain has to have a way of assuring the usefulness of its representations; one check as to usefulness is whether or not the representations can predict the current input. In such models the 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. The critical feature of the predictive coding formulation is that the structure of receptive fields can be predicted by assuming that the cortical memory costs metabolically in terms of spikes and synapses. This allows the general interpretation that the cortex trades off the accuracy of its computations with the amount of circuitry used to carry them out.

Predictive models have had a significant drawback in that they relied on the classical rate code interpretation of neural signaling. Our research extended predictive coding so that it could use volleys of synchronous spikes instead of a rate code [BRZ99]. The studies reported herein extend this work by showing that, in a detailed LGN model, the features of the extended model account for features observed in experimental data, but provide a different interpretation: We think that temporal features of receptive fields reflect the convergence of predictive coding rather than signaling motion features.

2 Predictive Coding Using Timed Volleys of Spikes

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 [RB98]. In a 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* $\{\hat{\mathbf{B}}_{1+\Delta}, \hat{\mathbf{B}}_{2+\Delta}, \dots, \hat{\mathbf{B}}_{T+\Delta}\}$ such that $\mathbf{B}_1 = \hat{\mathbf{B}}_{1+\Delta}$ and so on. We do this by using an internal model consisting of units \mathbf{r}_t with the goal that

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

Here U is a matrix of synapses connecting the units \mathbf{r}_t 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 [RB98] one can use the Minimum Description Length principle to trade off reconstruction error cost with the cost of the reconstruction units. This leads to minimizing an error of the form

$$E = \sum_{t=1}^T \{ \|\mathbf{B}_t - U\mathbf{r}_t\|^2 + \alpha \|\mathbf{r}_t\|^2 \} + \beta \|U\|^2$$

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 $\|\mathbf{r}_t\|^2$ 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. For purposes of simulation, each eight-bit gray level image is

converted to a set of eight volleys. The i -th volley consists of the i -th bits in the gray level code. These \mathbf{B}_t are then used repeatedly in training. The MDL-based learning algorithm computes a stable set of responses \mathbf{r}_t by using a form of gradient descent on E .

To train the synapses, sets of binary images were used in a repeated sequence. The synapses were shared between all the 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 a repeated sequence of what we term packets. Thus for image k ,

$$\dots, \mathbf{B}_1^k, \mathbf{B}_2^k, \mathbf{B}_3^k, \dots, \mathbf{B}_8^k, \mathbf{B}_1^k, \mathbf{B}_2^k, \mathbf{B}_3^k, \dots, \mathbf{B}_8^k, \mathbf{B}_1^k, \dots$$

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

Results reported in [BRZ99] showed clearly that after training, the network can generate feedback volleys that cancel the incoming binary images.

3 Simulation Results

To illustrate these principles using natural images, a network of 12×12 ON-center and 12×12 OFF-center input units and 36 coding units was trained on approximately 1000 natural images. The volleys were artificially created by using the binary encoding of the image, that is for an ON-center cell $\mathbf{B}_i^k(x, y) =$ the i -th bit of $[I^k(x, y)]^+$.

The sampling strategy from [RB98] was used. Patches from natural images will be filtered using the high pass filter described therein. Next that result was converted to volleys by thresholding the resultant image. Image samples greater than zero were each assigned to an ON-center LGN model cell and image samples less than zero were each assigned to an OFF-center cell. This input was used to train the network by using the algorithm described above.

Figure 1 shows the receptive fields from a sample of four model LGN cells from the trained network determined by reverse correlation(see figure text). The results show striking agreement with actual LGN recordings done by [RA95, AmUR96, URR98]. In addition we can compare the correlations that are observed after training in the network with actual correlations observed in the LGN. These comparisons are shown in Table 1.

Overlap of RF centres	Reid, Usrey exptl results		Our results	
	Per cent positively correlated	Strength of positive correlations	Per cent positively correlated	Strength of positive correlations
No overlap	0% (0/73)	-	5.09% (1584/31104)	2.33%
Partial overlap	17% (28/167)	1.9%	10.09% (813/8056)	5.8%
Total overlap	48% (20/42)	10%	38.61% (837/2168)	44.97%
and same sign and similar size	79%(10/13)	23%	82.71%(837/1012)	96.35%

Table 1: Comparison of experimental results from Reid/Usrey to results generated by our model




Figure 1: **Reverse Correlation Results** Once the network has been trained, it can be tested using random images as input. For each LGN model cell, the set of images that produced its spikes are saved as a function of different time delays. These are then summed and averaged to provide an indication of the cell's spatio-temporal receptive field. Above, four cells are shown in vertical rows. For each cell the temporal response is shown every 20 milliseconds.

4 Conclusion

The comparisons in Table 1 show excellent agreement between the simulation and experiment. The “strength of positive correlations” may appear to be different but the methods of determining strength were not commensurate. The reader can verify that the two columns are in appropriate ratio. In addition our overlap test is slightly less stringent than that of Reid et al so a small number of “no overlap” cells show positive correlations. The major consequence of the good fit is that the temporal dynamics observed experimentally may reflect the convergence of a predictive coding circuit.

Our synchronous predictive model requires that cortico-thalamic feedback create repeated volleys in the feedforward pathway. The evidence for this comes from developmental experiments on ferrets[WK99]. These experiments show that cortico-thalamic connections produce synchronous firing and that severing these connections eliminates the synchrony between left and right eye signals. In addition we know that in cortical area V1, the time to experience feedback can be as long as 100-200 milliseconds[ZLS97]. 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, as also suggested by our model.

References

- [AmUR96] 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.

- [BRZ99] Dana H. Ballard, Rajesh P. N. Rao, and Zuohua Zhang. A single-spike model of predictive coding. *Conference on Computational Neuroscience*, 1999.
- [CBNS98] 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.
- [Fie87] 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.
- [Mei96] Markus Meister. Multineuronal codes in retinal signaling. *Proceedings of the National Academy of Sciences*, 93:609–614, 1996.
- [MS95] Z. F. Mainen and T. J. Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268:1503, 1995.
- [OF97] 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.
- [RA95] R. Clay Reid and Jose-Manuel Alonso. Specificity of monosynaptic connections from thalamus to visual cortex. *Nature*, 378:281–284, 1995.
- [RB96] 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.
- [RB98] 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.
- [SJGW94] 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.
- [URR98] 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.
- [WK99] M. Weliky and L.C. Katz. Correlational structure of spontaneous neuronal activity in the developing lateral geniculate nucleus in vivo. *Science*, 265:599, 1999.
- [ZLS97] Karl Zipser, Victor A. F. Lamme, and Peter N. Schiller. Contextual modulation in primary visual cortex. *Journal of Neuroscience*, 16:7376, 1997.