

Sparse Coding in the Neocortex

Daniel J. Graham and David J. Field

Department of Psychology, Cornell University, Ithaca, NY 14817

Outline

6.03.1 INTRODUCTION : OPTIMALITY IN BIOLOGICAL SYSTEMS

6.03.2 DEFINING SPARSE CODING

6.03.3 PHYSIOLOGICAL EVIDENCE FOR SPARSE CODING

6.03.4 TWO VIEWS OF SPARSE CODING

6.03.4.1 *Maximizing Information Transfer*

6.03.4.2 *Metabolic Constraints*

6.03.5 CONCLUSIONS

12/2/05

Sparse Coding in the Neocortex

Daniel J. Graham

Department of Psychology, Cornell University, Ithaca, NY 14853 USA

djg45@cornell.edu (607)435-4799

and

David J. Field

Department of Psychology, Cornell University, Ithaca, NY 14853 USA

djf3@cornell.edu (607)592-9890

Keywords

efficient codes; evolution; invariance; kurtosis; metabolic costs; natural scenes; nonlinearities; optimality; selectivity; sparse coding; sparseness; visual cortex; visual system.

Glossary

Kurtosis : The fourth statistical moment of a distribution; measures the degree to which a distribution is peaked and heavy-tailed; a Gaussian has a kurtosis of 0.

Natural scene : An image of the natural world.

Sparse code : A method of representing information that shows a low activity ratio; a code for which most coding units are inactive most of the time.

Abstract

Sparse coding has been proposed as a guiding principle in neural representations of sensory input, particularly in the visual system. Because sparse codes are defined as representations with low activity ratios—i.e., at any given time a small proportion of neurons are active—they are sometimes proposed as a means to help conserve metabolic costs. Although we accept that such metabolic costs play a role in evolutionary selection, we also argue that sparse coding offers a useful solution to the problem of representing natural data because such a scheme allows the system to take advantage of the sparse structure of the sensory environment. Given the highly regular and sparse structure of natural data (e.g., natural scenes), a variety of recent studies suggest that the sparse coding in the cortex is an adaptation that has evolved to offer an efficient strategy for storing and building associations with natural inputs. However, highly selective sparse codes require a large number of neurons and can also pose problems for learning. We argue that larger brains have evolved strategies

that allow neurons to be both highly selective and invariant across irrelevant dimensions. We argue that this invariance is what allows such highly sparse codes to be useful.

6.03.1 INTRODUCTION : OPTIMALITY IN BIOLOGICAL SYSTEMS

For any biological system, an account of why the system is structured as it is requires consideration of a number of interacting forces including the uses (or goals) of the system, the environment in which the system must function, and the constraints that history and biology put on the design. It might seem reasonable to assume that the constraints of evolution and development play a large role in determining the design of any neural system. However, a range of recent studies have argued that neural systems have found highly efficient solutions for representing environmental information. These studies have explored topics from retinal coding (Sterling, 2004) to the computations provided on the semicircular canals with respect to head rotations (Squires, 2004) to the optimal cortical layout that minimizes neural wiring (Van Essen, 1997). In all these studies, there remain intriguing questions regarding how close these solutions are to optimal, and how evolution and development lead to this optimality.

In this chapter, however, we focus on a general aspect of sensory representation called sparse coding. We argue that there is widespread evidence of such coding in neural systems across a variety of species. We look briefly at the question of what is meant by sparse coding and then ask why sensory systems would profit from performing such coding. We argue that the natural environment is inherently sparse and codes that take advantage of this structure can be both metabolically efficient and useful for learning. However, the constraints involved in producing a highly sparse code can be severe: if representing each object and pose requires a different set of neurons, the system would need a very large number of neurons indeed. We believe that many vertebrates have developed a strategy of combining high sparseness with invariance, a strategy that overcomes the

combinatorial explosion of a highly sparse code. In this chapter, we address the extent to which sparseness is an optimal coding solution for natural data, and the additional processing strategies that may have shaped cortical evolution in vertebrates.

6.03.2 DEFINING SPARSE CODING

Sparse coding generally refers to a representation where a small number of neurons are active, with the majority of the neurons inactive or showing low activity (e.g., Field, 1987; 1994; Rolls and Tovee, 1995). In his influential 'single neuron doctrine,' Barlow (1972) suggested sparseness as one of the principles important to sensory representation. However, sparse coding in its extreme form results in a representation sometimes called a 'grandmother cell' code. In such a code, each object in the world (e.g., your grandmother) is represented by a single cell. One might argue that a large brain with tens of billions of neurons can certainly handle a few hundred thousand object-level neurons. And there are a range of studies showing there exist neurons that can be highly selective to faces and other objects (e.g., Kendrick and Baldwin, 1987; Quiroga, *et al.* 2005). However, those promoting the usefulness of sparse representations are not proposing that the ultimate goal is to have one neuron for every object—and certainly not for a particular view of every object. We believe that sparseness helps learning and prediction even at early stages of sensory processing, like those found in V1. But too much specificity or sparseness can actually make learning harder. We will explore this question later in the chapter.

Two lines of evidence support the notion that sparse representations are common in neural systems: the first comes from physiology, the second from computational and theoretical research. In each case, the evidence requires a definition of sparseness. There have been several definitions of sparseness and a number of ideas regarding what sparse codes actually represent.

Sparseness can be defined over a population of neurons at a given point in time (population sparseness) or it can be measured for a single neuron over some window of time—the latter is called temporal or lifetime sparseness, (Willmore and Tolhurst, 2001) and it is sometimes referred to as non-parametric selectivity (Lehky, Sejnowski and Desimone, 2004). For a given distribution of responses, we obtain a histogram of activity. One might think that the simplest definition of the sparseness of this distribution is to simply measure the proportion of active neurons, or how often a neuron is active. However, the histogram of activities is usually defined over a window of time and is therefore not binary. In response to any population of stimuli (e.g., natural scenes), one typically obtains a distribution of activities (a distribution of spike probability), so the measures of sparseness refer to the relative shape of the distribution.

Two definitions of sparseness are widespread. The first—the Treves-Rolls measure (Equation 1)—is more appropriate for measuring the sparseness over time for real spiking neurons (Rolls and Tovee, 1995). The second definition uses kurtosis (the fourth statistical moment of a distribution) as its metric (Equation 2) and it is more useful for modeled neurons and computational studies, where a comparison of different codes and transforms is necessary (Field, 1994).

$$S = \frac{\left(\frac{1}{n} \sum_i r_i \right)^2}{\frac{1}{n} \sum_i r_i^2}$$

Equation 1: Treves-Rolls sparseness measure.

$$k = \frac{1}{n} \sum_{i=1}^n \frac{(r_i - \bar{r})^4}{\sigma^4} - 3$$

Equation 2: Kurtosis sparseness measure.

These measures are applied to histograms of responses over a population of neurons, or for a given neuron over time. With the Treves-Rolls measure, the sparseness approaches zero as the neuron is either off or highly on (for a given time window). With the kurtosis measure, large values occur when a response distribution deviates maximally from the Gaussian state by being sharply peaked and heavy-tailed. Highly kurtotic behavior produces a relatively high probability of either a small or large response, and a relatively low probability of a mid-level response. Both measures can be sensitive to outliers.

6.03.3 PHYSIOLOGICAL EVIDENCE FOR SPARSE CODING

Much of the discussion in recent years regarding sparse coding has come from the computational and theoretical literature but there is considerable physiological evidence for sparse representations in most biological systems. For a neuron with a full response and refractory time of 5 msec, the maximal firing rate would be 200 Hz. A system that is providing a maximal information rate would fire at approximately half the time, i.e. at 100 Hz. Although neurons may reach such rates, neurons do not maintain these rates for more than brief periods. We do not know of any neural systems that maintain such high firing rates for extended periods.

Such a high firing rate would require considerable energy resources, so much so that Attwell and Laughlin (2001) argue that the limited biochemical energy available for producing action potentials must limit the average firing rates of neurons to less than 1 Hz. Further, Lennie (2003) estimates that the limited resources imply that at any given time, only 1/50th of any population of cortical neurons can afford to show high firing rates. Therefore, for biochemical reasons alone, we should expect a considerable degree of both lifetime and population sparseness. Olshausen and Field (2005) further argue that even in areas that have been well studied—areas like V1—these low average firing rates

imply that a significant number of neurons will have such low firing rates as to be missed entirely by the typical search strategies.

As noted by Olshausen and Field (2004), there are a number of studies suggesting that many neural systems utilize highly sparse codes. DeWeese and colleagues (2003), recording from auditory neurons in the rat, have demonstrated that neurons in A1 can reliably produce a single spike in response to a sound. Evidence from olfactory systems in insects (Perez-Orive *et al.*, 2002; Theunissen *et al.*, 2003), somatosensory neurons in rat (Brecht and Sakmann, 2002), and recordings from rat hippocampus (Thompson and Best, 1989) all demonstrate highly sparse responses. Prefrontal cortex shows similar sparseness in behaving rhesus monkeys (Abeles, Vaadia and Bergman, 1990). As the authors of the latter study say, most areas of association cortex are "not carrying out any computations for the majority of the time."

Motor neuron representations are often described as a population code, where it is proposed that the accuracy of a movement is guided by the degree of activity of a relatively large population of neurons (see Georgopoulos, 1986). Here too, we find evidence of sparse responses. Some motor neurons in layer 6 of rabbit motor cortex will produce just one spike during some movements (Beloozerova, Sirota and Swadlow, 2003). And stimulation of a single neuron in the rat is sufficient to deflect a whisker (Brecht *et al.*, 2004).

With respect to sparse coding, the most widely studied sensory system is the visual system. Much of this work, which we discuss below, has been motivated by information theoretic issues. The area also contains a wealth of experimental data. In inferotemporal cortex (IT), a wide range of studies supports the notion that neurons are selective to high-level object dimensions, and to features such as faces and hands. Such neurons are believed to show some degree of invariance over position, size, pose, brightness, etc. Nevertheless, unless we are to assume that such objects fall within the neuron's receptive field at least half of their waking life, we should expect these neurons to show a high

degree of sparseness. Indeed, Baddeley *et al.* (1997) found that cells in IT show sparse responses to natural stimuli, to a similar degree as do cells in V1.

Much of the most interesting work tying together the statistics of natural signals and physiology comes from work on the responses of V1 neurons. Vinje and Gallant (2000, 2002) found that V1 neuron responses in macaque become more and more sparse as the size of a natural stimulus is increased beyond the classical receptive field. Stimulation in the classical receptive field also produced sparseness, which could reflect the rather arbitrary nature of the classical/non-classical delineation. Moreover, stimulation in the nonclassical receptive field showed the following results:

1. Increased sparseness for individual neurons during repeated presentations (lifetime sparseness).
2. Increased sparseness across the population of neurons tested (population sparseness).
3. Decreased correlation in neighboring neurons, thereby whitening the response.

The results presented here provide examples of the sparse behavior of neurons in primate visual cortex under naturalistic conditions. Although we cannot argue that a sparse coding strategy is ubiquitous in the cortex, these results do support the implication that sparse coding is widespread in the nervous systems of the mammals tested. We also find evidence of sparse coding across a variety of non-mammalian species. In addition to the creatures mentioned above, the selectivity of sensory neurons has been supported by studies in amphibians (e.g., Ewert, 1980), turtles (Ammermüller *et al.*, 1995) and insects (e.g., Perez-Orive *et al.*, 2002; Strausfeld and Lee, 1991; Srinivasan *et al.*, 1993; Lehrer, 1994).

Any form of selectivity implies that neurons will show a degree of sparseness in the natural environment, since selectivity by definition means neurons respond only to a portion of the possible environmental stimuli. Answering the question of why sensory systems show highly selective responses will require innovative ways of thinking about sparseness. In the next section we consider

an information theoretic approach to sparse coding and we compare this to an approach that argues for sparse coding as a result of the metabolic constraints on neural firing. We argue that both approaches will likely be needed to explain all aspects of sparse coding in neural systems.

6.03.4 TWO VIEWS OF SPARSE CODING

6.03.4.1. Maximizing Information Transfer

Consider a neuron, or population of neurons, that has a mean firing rate and some distribution around that mean. With a limit on the range of possible firing rates, the maximum information transfer occurs when all states of the channel are used with equal frequency—a flat distribution of firing rates. If the bound on a distribution is instead the variance of the responses (rather than the range of responses), the distribution with the greatest information rate (greatest entropy) is a Gaussian. The visual system appears to follow neither of these models. A sparse code means that neural firing rates will show a highly peaked, non-Gaussian distribution; i.e., one that does not produce maximum information transfer. Moreover, as sparseness increases, the information rate drops. Why might this be a good idea?

In general, we argue that the information rate of the system should match the information rate of the input. Natural stimuli are not random. As far as the visual system is concerned, natural scenes are not arrays of random points of light. Rather, they are constrained by strong correlations between neighboring regions, and image discontinuities are usually defined by edges. This predictability implies that a high information rate is simply unnecessary. By taking advantage of the redundant properties of images, sensory codes can get away with sending less information and using fewer spikes.

For example, consider a collection of 5 x 5 pixel images that each contain one block letter of the alphabet. If we looked at the histogram of any given pixel, we might discover that the pixel was on

roughly half the time. However, if we were to represent these letters with templates that respond uniquely to each letter, each template would respond just 1/26th of the time. This letter code is more sparse—and more efficient—relative to a pixel code. Although no information is lost, the letter code would produce the lowest information rate.

Moreover, a representation that was letter-based would provide a more efficient means of learning about the associations between letters. If the associations were between individual pixels, a relatively complex set of statistical relationships would be required to describe the co-occurrences of letters (e.g., between the Q and U). Sparseness can assist in learning since each unit is providing a relatively complete representation of the local structure (Field, 1994).

Of course, the natural world does not consist of letters. Natural scenes are highly structured and can be modeled to a first approximation as a sparse collection of local features (e.g., edges). Early work by one of us (Field, 1987; 1994) showed that if the receptive fields of V1 neurons are modeled as a collection of linear templates, then the responses of those neurons to natural scenes are highly sparse. Furthermore, when the parameters of the modeled neurons were altered from those of V1 neurons, the sparseness dropped, suggesting that the parameters were near to optimal given the constraint of having a linear array of model neurons.

A stronger test of this hypothesis was developed by Olshausen and Field (1996), who trained a neural network to find the most sparse representation for a population of image patches drawn from natural scenes. The network was trained to develop a set of filters that would maximize sparseness and losslessness in its representation of natural scenes. Given these two criteria, the network settled on a set of filters with considerable similarity to simple cell receptive fields in V1.

Insert Figure 1 about here

This basic sparse coding algorithm is quite similar to techniques that search for independent solutions by minimizing response entropy (that is, by searching for non-Gaussian response histograms). The family of these techniques has been called independent components analysis or ICA (Bell and Sejnowski, 1997). The name ICA is a bit unfortunate since the solutions are almost never independent, given natural input. But the approach has been applied to a wide range of problems and it has been employed to account for a variety of properties in early visual neurons including spatio-temporal tuning (e.g., van Hateren and Ruderman, 1998) and spatio-chromatic properties (e.g., Hoyer and Hyvarinen, 2000), and generalization of the approach has been used to model some nonlinear properties of neurons (Hyvarinen and Hoyer, 2000). In early vision, Graham, Chandler and Field (2005) show evidence that sparseness is likely a factor contributing to the utility of center-surround receptive field organization, along with decorrelation and response gain. Furthermore, sparse codes of natural sounds have been shown to produce temporal response profiles with properties similar to those of early auditory neurons (Lewicki, 2002).

This entire line of work suggests that sparseness in neural firing is primarily a result of efficiently matching the neuron's response properties with the sparse statistical structure of the environment. There remain a number of questions as to how a biological system might achieve this efficient representation. Although the particular method used by Olshausen and Field (1996) may not be biologically plausible, it is argued that there exist learning algorithms that could train a system to achieve a sparse code. The advantage of learning such a code is that the properties of individual neurons need not be predetermined or coded in the genome of the organism. Rather, what is required is that the system evolve only a general sparse coding strategy. Given appropriate input, a learning algorithm would serve to produce the proper neural response properties and it would help to tile the neurons in a way that allows the correct spacing as a function of the different parameters

of selectivity (e.g., position, scale, orientation). Although there is little evidence of such a learning algorithm in insects and animals with relatively simple brains, the evolution of a sparse learning algorithm may be widespread in larger brains. However, the argument is not that the system must necessarily learn from the natural world. Rather, one intriguing possibility is that the system learns from the patterned structure of the spontaneous activity in the developing organism (e.g., Wong, 1999).

It is not yet clear whether the patterns of spontaneous activity is sufficient for generating the receptive field properties found in the newborn. However, such an approach would provide a relatively simple means for producing a large number neurons with efficient tuning. The learning algorithm also has the advantage that it can accommodate a large variation in cell number. If any evolutionary mutation or developmental change results in a larger (or smaller) brain, the learning algorithm should be capable of adapting to the system and adjusting the tiling appropriately.

Could all neurons in higher areas from V1 to IT simply be exhibiting sparse coding as a means of efficiently representing the natural world? This may be possible. A number of investigators are exploring ways to extend these ideas of efficient coding to higher levels. It is clear that these more complex representations require a form of nonlinear coding that makes the tests of efficiency considerably more difficult, though attempts to model nonlinear behavior of neurons with such efficient coding approaches have met with some success (e.g., Wiskott and Sejnowski, 2002; Schwartz and Simoncelli, 1997).

However, there is a penalty that applies to learning if the system is too sparse. An extremely sparse code (one in which neurons are highly selective for specific objects in specific poses, lighting etc.) would have neurons that fired quite rarely. In order to effectively learn about the world, any system must keep track of the relative probability of co-occurrences. No matter how a neural system keeps track of these co-occurrences, if they occur too rarely it would be impossible to determine

whether any feature is statistically related to any other feature. We cannot learn about how “faces” behave in particular situations, if we have a neuron for every unique face. It is important the system be invariant at some level so that we can collapse across instances of the category. Most presentations of objects or events will occur just once or not at all during development if the object is defined too precisely.

From the perspective of learning, then, a code that is too sparse becomes intractable even if there were enough neurons to accommodate it. Consequently, a system must also build invariance into the code in order to develop and function efficiently. With high-level objects like faces, this learning constraint would require that the face-selective neurons be invariant to dimensions along which the face varies in different settings (lighting, pose, size, etc.). Thus, both invariance and selectivity are necessary for achieving an efficient, sparse representation of sparse natural input. Indeed, this invariance is a known property of visual neurons. As neurons are found to be more selective, we find greater degrees of invariance. The complex cells in V1, show selectivity to scale (spatial frequency), and orientation, and they show small amounts of invariance to position. Higher level neurons in IT and medial temporal cortex may show much higher selectivity to faces, hands, etc. However, they also show much greater invariance to lighting, pose, and size (Rolls, 2000). In one study of human medial temporal cortex (Quiroga *et al.*, 2005), neurons were found that were selective to particular actresses (Jennifer Aniston) while invariant to pose, lighting and position in the image.

We therefore argue that although the evolution of large brains may allow a larger number of highly selective neurons, the constraints of learning require that the selectivity go hand in hand with a greater degree of invariance. Although there have been a number of proposals regarding how invariance is achieved in the mammalian systems (Olshausen *et al.*, 1993), no firm answer has emerged. We know that some of the simpler visual systems, such as those of *Drosophila*, do not

show such invariance (Dill, Wolf, and Heisenberg, 1993). However, it remains unclear how such invariance has evolved, and what we might expect from systems that show only partial invariance.

6.03.4.2 Metabolic Constraints

We conclude our discussion by returning to the issue of metabolic constraints. Could we argue that primary evolutionary pressure driving towards sparse coding is one related to the metabolic costs of neural firing? As noted earlier, both Attwell and Laughlin (2001) and Lennie (2003) argue that there are not enough resources to achieve anything but a low-activity system. Moreover, when we find sparse activity in frontal cortex (Abeles *et al.*, 1990), it is more difficult to argue that the sparse activity must arise because it is mapping the sparse structure of the world. Even at early levels, if sparseness were metabolically desirable, there are a number of ways of achieving sparseness without matching the structure of the world. Any one of a wide variety of positively accelerating nonlinearities would do. Simply giving the neurons a very high threshold would achieve a sparse code, but the system would lose information. We argue that the form of sparse coding found in sensory systems is useful because such codes maintain the information in the environment, but do so more efficiently. We argue that the evolutionary pressure to move the system towards a sparse code comes from the representational power of sparse codes.

However, we do accept that metabolic constraints are quite important. It has been demonstrated that at the earliest levels of the visual system, ganglion cells (Berry, Warland and Meister, 1997) and lateral geniculate nucleus cells (Reinagel and Reid, 2000) show sparse (non-Gaussian) responses to temporal noise. A linear code, no matter how efficiently it was designed would not show such sparse activity, so we must assume that the sparseness is at least in part due to the nonlinearities in the system and not due to the match between the receptive fields and the sparse structure of the input.

As with the results show sparse responses in non-sensory areas, we must accept that metabolic may also be playing a significant role.

6.03.5 CONCLUSIONS

We are therefore left with a bit of a puzzle. We know that higher levels of the visual system show considerable sparseness: neurons fire at rates far below their maximal rate. However, we cannot conclude that sparseness is only a result of an efficient mapping of the sparse structure of the world. Metabolic efficiency must also be considered, independent of the statistical structure of the world. In addressing the question of why a system is sparse, we must accept that widespread sparseness in cortex is due to several factors. Many of us believe that the metabolic constraints are secondary, however, and that artificial visual systems will someday incorporate much of the coding we find in neural systems. The constraints of evolution, metabolism, anatomy and development all play a role in determining why the nervous system is designed the way it is. But one should not presume that these constraints force the system towards some non-optimal solution. At present, the evidence suggests that the nervous system has evolved a highly efficient learning algorithm for discovering and representing the structure in the world. And the sparse responses of neurons are an integral component of that efficient representation.

Acknowledgements

This work was supported by the following grants and contracts: DJG was supported by NIH EY015393-01 (Kirschstein-NRSA) and NSF DGE-9870631 (IGERT Program in Nonlinear Systems fellowship) and DJF was supported by NGA contract HM 1582-05-C-0007.

References

- Abeles, M., Vaadia, E. and Bergman, H. 1990. Firing patterns of single units in the prefrontal cortex and neural network models. *Netw.* 1, 13-25.
- Ammermüller, J., Müller, J. F. and Kolb, H. 1995. The organization of the turtle inner retina. II. Analysis of color-coded and directionally selective cells. *J. Comp. Neurol.* 358, 35-62.
- Attwell, D. and Laughlin, S. B. 2001. An energy budget for signaling in the grey matter of the brain. *J Cereb. Blood Flow Metab.* 21, 1133-1145.
- Baddeley, R., Abbott, L. F., Booth, M. C., Sengpiel, F., Freeman, T., Wakeman, E. A. and Rolls, E. T. 1997. Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proc. Roy. Soc. Lond. B* 264, 1775-1783.
- Barlow, H. B. 1972. Single units and sensation: A neuron doctrine for perceptual psychology? *Perception* 1, 371-394.
- Bell, A. J. and Sejnowski, T. J. 1997. The 'independent components' of natural scenes are edge filters. *Vision Res.* 37, 3327-3338.
- Beloozerova, I. N., Sirota, M. G. and Swadlow, H. A. 2003. Activity of different classes of neurons of the motor cortex during locomotion. *J. Neurosci.* 23, 1087-1097.
- Berry, M. J., Warland, D. K. and Meister, M. 1997. The structure and precision of retinal spike trains. *Proc. Nat. Acad. Sci. USA* 94, 5411-5416.
- Brecht, M. and Sakmann, B. 2002. Dynamic representation of whisker deflection by synaptic potentials in spiny stellate and pyramidal cells in the barrels and septa of layer 4 rat somatosensory cortex. *J. Physiol.* 543, 49-70.
- Brecht, M., Schneider, M., Sakmann, B. and Margrie, T. W. 2004. Whisker movements evoked by stimulation of single pyramidal cells in rat motor cortex. *Nature* 427, 704-710.
- DeWeese, M., Wehr, M. and Zador, A. 2003. Binary spiking in auditory cortex. *J. Neurosci.* 23, 7940-7949.
- Dill, M., Wolf, R. and Heisenberg, M. 1993. Visual pattern recognition in *Drosophila* involves retinotopic matching. *Nature* 365, 751-753.
- Ewert, J. P. 1980. *Neuroethology*. New York: Springer-Verlag.
- Field, D. J. 1987. Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am. A* 4, 2379-2394.
- Field, D. J. 1994. What is the goal of sensory coding? *Neural Comput.* 6, 559-601.
- Georgopoulos, A. P., Schwartz, A. B. and Kettner, R. E. 1986. Neuronal population coding of movement direction. *Science* 233, 1416-1419.
- Graham, D. J., Chandler, D. M. and Field, D. J. 2005. Can the theory of "whitening" explain the center-surround properties of retinal ganglion cell receptive fields? *Vision Res.* under review.
- Hoyer P. O. and Hyvärinen, A. 2000. Independent component analysis applied to feature extraction from colour and stereo images. *Netw.* 11, 191-210.
- Hyvärinen, A. and Hoyer, P. O. 2000. Emergence of phase and shift invariant features by decomposition of natural images into independent feature subspaces. *Neural Comput.* 12, 1705-1720.
- Kendrick, K. M. and Baldwin, B. A. 1987. Cells in temporal cortex of conscious sheep can respond preferentially to the sight of faces. *Science* 236, 448-450.
- Lehrer, M. and Srinivasan, M. V. 1992. Freely flying bees discriminate between stationary and moving objects: performance and possible mechanisms. *J. Comp. Physiol. A* 171, 457-467.
- Lehky, S. R., Sejnowski, T. J. and Desimone, R. 2005. Selectivity and sparseness in the responses of striate complex cells. *Vision Res.* 45, 57-73.
- Lennie, P. 2003. The cost of cortical computation. *Curr. Biol.* 13, 493-497.

- Lewicki, M. S. 2002. Efficient coding of natural sounds. *Nature Neurosci.* 5, 356-363.
- Olshausen, B. A., Anderson, C. H. and Van Essen, D. C. 1993. A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *J. Neurosci.* 13, 4700-4719.
- Olshausen, B. A. and Field, D. J. 1996. Emergence of simple cell receptive field properties by learning a sparse code for natural images. *Nature* 381, 607-609.
- Olshausen, B. A. and Field, D. J. 2004. Sparse coding of sensory inputs. *Curr. Opin. Neurobiol* 14, 481-487.
- Olshausen, B. A. and Field, D. J. 2005. How close are we to understanding V1? *Neural Comput.* 17, 1665-1699.
- Perez-Orive, J., Mazor, O., Turner, G. C., Cassenaer, S., Wilson, R. I. and Laurent, G. 2002. Oscillations and sparsening of odor representations in the mushroom body. *Science* 297, 359-365.
- Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C. and Fried, I. 2005. Invariant visual representation by single neurons in the human brain. *Nature* 435, 1102-1107.
- Reinagel, P. and Reid, R. C. 2000. Temporal coding of visual in the thalamus. *J. Neurosci.* 20, 5392-5400.
- Rolls, E. T. and Tovee, M. J. 1995. Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *J. Neurophysiol.* 73, 713-726.
- Rolls, E. T. (2000). Functions of the primate temporal lobe cortical visual areas in invariant visual object and face recognition. *Neuron* 27, 205-218.
- Schwartz, O. and Simoncelli, E. P. 2001. Natural signal statistics and sensory gain control. *Nature Neurosci.* 4, 819-825.
- Squires, T. M. 2004. Optimizing the vertebrate vestibular semicircular canal: Could we balance any better? *Phys. Rev. Lett.* 93, 198106.
- Sterling, P. 2004. How retinal circuits optimize the transfer of visual information. In: *The Visual Neurosciences*, Vol. II (eds. L. Chalupa and J. S. Werner), p. 235-259. MIT Press: Cambridge.
- Strausfeld, N. J. and Lee, J-K. 1991. Neuronal basis for parallel visual processing in the fly. *Visual Neurosci.* 7, 13-33.
- Theunissen, F. E. 2003. From synchrony to sparseness. *Trends Neurosci.* 26, 61-64.
- Thompson, L. T. and Best, P. J. 1989. Place cells and silent cells in the hippocampus of freely-behaving rats. *J. Neurosci.* 9, 2382-2390.
- Van Essen, D. C. 1997. A tension-based theory of morphogenesis and compact wiring in the central nervous system. *Nature* 385, 313-318.
- van Hateren, J. H. and Ruderman, D. 1998. Independent component analysis of natural image sequences yields spatio-temporal filters similar to simple cells in primary visual cortex. *Proc. Roy. Soc. Lond. B* 265, 2315-2320.
- van Hateren, J. H. and van der Schaaf, A. 1998. Independent component filters of natural images compared with simple cells in primary visual cortex. *Proc. Roy. Soc. Lond. B* 265, 359-366.
- Vinje, W. E. and Gallant, J. L. 2000. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287, 1273-1276.
- Vinje, W. E. and Gallant, J. L. 2002. Natural stimulation of the nonclassical receptive field increases information transmission efficiency in V1. *J. Neurosci.* 22, 2904-2915.
- Willmore, B. and Tolhurst, D. J. 2001. Characterizing the sparseness of neural codes. *Netw.* 12, 255-270.
- Wiskott, L. and Sejnowski, T. J. 2002. Slow feature analysis: unsupervised learning of invariances. *Neural Comput.* 14, 715-770.
- Wong, R. O. L. 1999. Retinal waves and visual system development. *Ann. Rev. Neurosci.* 22, 29-47.




Figure 1. Results of a neural network that searches for a sparse code using filters to describe 12x12 pixel image patches drawn from a collection of natural scenes (Olshausen and Field, 1996). The collection of filters shown represent the 256 templates that the network found for describing each patch. When any given natural scene patch is multiplied by this family of filters, one finds that most of the responses are near zero and a small set of filters produce large responses (those matched to the image structure). The templates have been shown to provide a good first-order account of the responses of cortical V1 simple cells suggesting that such simple cells are optimized to provide a sparse code for natural scenes. Adapted from Olshausen and Field, 1996.

Websites

Field Lab: <http://redwood.psych.cornell.edu/>

Tolhurst Lab: <http://www.physiol.cam.ac.uk/staff/tolhurst/>

Van Hateren Lab (natural scene stimuli collection): <http://hlab.phys.rug.nl/>

Olshausen Lab: <http://redwood.ucdavis.edu/bruno/>