Baddeley, R., Vincent, B., & Attewell, D (2011) Information theory and perception: the role of constraints, and what do we maximise information about?. In G. Terzis & R. Arp (Eds.), Information and Living Systems: Philosophical and scientific perspectives (289-307). Cambridge, Mass: MIT Press.

10

Information Theory and Perception: The Role of Constraints, and What Do We Maximize Information About?

Roland Baddeley, Benjamin Vincent, and David Attewell

Because human and nonhuman primates are highly visual animals, an important step in understanding how their brains operate is to understand how they see: that is, how they transform the structured, colored, and dynamic pattern of light surrounding them into knowledge about the world they live in. Over the last twenty years, a number of problems associated with understanding vision have used information-theoretic techniques to help us gain insight into perception (Shannon 1948). This chapter will describe, in a biased way, one particular strand of this work: how information theory has been used to shed light on the processes and representations present at the earliest stages of vision.

There are three reasons why early vision is a particularly good system to which to apply information theoretic concepts. First, we know a lot about what it does, even if we don't always know precisely why it does it. The pioneering physiological work of Hubel and Wiesel (1962) gave us a good grasp of the basic properties of the representations of early vision, and since then, the neurons in the earliest and largest cortical area involved in vision, V1 (sometimes also known as the striate cortex or area 17), have probably been more studied than those in any other part of the brain. These cells are sensitive to certain parameters of the local pattern of light arriving at the areas of the retina they are interested in, such as color, orientation, and motion. The parameter values to which a given cell is sensitive is known as its *receptive field* (Lennie 2003a). Using increasingly sophisticated methods, such cell properties have been not only qualitatively but also quantitatively described. In addition to our good physiological understanding of early vision, the fact that displaying simple but very well-controlled visual stimuli is relatively easy means that early vision has been studied using the methods of psychophysics more than any other sense. As a result of all this intensive work, we have a lot of data, both physiological and psychophysical, with which we can compare the predictions of any information theory–inspired model of early vision.

The second reason why early vision is amenable to study using information theory is that it is possible to generate estimates of its "natural" input. Information theory calculations often require estimates of the statistics of the signal to be transmitted (Simoncelli and Olshausen 2001) and, for vision, these estimates can be made from natural images and videos. This can be done in more or less sophisticated ways, from simply approximating input in terms of some "representative natural images" taken from the web, to approaches that take into account eye movements (Rucci, Edelman, and Wray 2000; Rucci and Casile 2005), the animal's environmental niche (Warrant 1999), and possibly the effects of the observer's motion through the environment. Importantly, the logistic and theoretical problems associated with making approximations to the input probability distribution for vision are not insurmountable. In contrast, applying information theory approaches to problems in the cognitive domain, for example, may be more challenging due to the difficulty of defining precisely the input to this system.

The last reason why an information theoretical approach may be appropriate is that at least at the earliest stages of vision (e.g., the transmission of information from the eye to the cortex), there exist rather severe bottlenecks for information transmission (discussed shortly). A system that is under severe constraints about how much information can be transmitted is more likely to be, in some sense, efficient. Although this is not so for all stages of early vision, at least some parts of the system have, on the face of it, more information to transmit than would be trivially possible without some sort of efficient (in terms of information theory) recoding. Comparing representations of the world observed in early vision with various optimal information theory-based representations has, therefore, at least some chance of success.

The rest of this chapter will therefore begin with a brief description of some pertinent properties of early vision. We will then describe the basic approach that has been employed when using information theory–based ideas to gain a better understanding of early visual representations, and illustrate this approach with three basic models. These models all share a number of assumptions about the nature of early visual representations and all use approximations to the visual input based on sampling from natural images. They differ, however, in that they propose that different constraints are important in understanding early perception. Last, we will give a very brief overview of some problems with the approach, together with some potentially interesting new directions to pursue.

The Nature of the Early Visual System

The main modeling effort that we will describe is concentrated on understanding the earliest stages of visual information processing. These are possibly the most intensively studied of all brain functions, and any short summary is bound to smooth over many potentially important details. However, some background is required to understand the subsequent modeling.

Early visual processing can be thought of as having three stages. In the first stage, the light in the retina is detected (via rod and cone cells), and the resulting signal is normalized. In the second stage, the signal is recoded, and transmitted from the retina to the cortex. This information is transmitted first using cells known as *retinal ganglion cells*, and then via a subcortical area called the *lateral geniculate nucleus* (LGN), whose exact function, despite a large number of theories, is not well understood (Derrington 2000). We will therefore ignore it. In the third and final stage, the information arrives in the first cortical area (V1), and is represented by a vastly greater number of cells than previously, which make a number of properties (or "features") of the local visual input explicit rather than implicit. In most layers of V1, the visual input is represented in terms of the presence of orientated edges, with cells firing at a high rate when presented with an edge or line within its local "receptive field," and not when no such edge or line is present.

On average, we point our eyes to a new location three times a second, and the information inherent in the light coming from the new direction has to be communicated to higher levels of the cortex. The basic problem the retina faces during the first stage of early vision is that it has to deal (given time to adapt) with light intensities that vary over nine orders of magnitude. This is beyond the dynamic range of all the subsequent processing stages. So, after detecting the light, the light level at every location is transformed with a transform that is not badly approximated over the most important range by a logarithm (Valeton and van Norren 1983). This is then followed by luminance adaptation; the detected light level is normalized by subtracting the average (log) light level in the local region over the recent past from the signal. This processing means that even with very large changes in the illuminant (e.g., the sun going behind a cloud), we can still see. This normalized version of the input needs then to be transmitted to later stages for further processing.

Though the eyes are (obviously) at the front of the head, the first cortical area involved in vision (V1) is at the back of the head. This means that the information has to be transmitted over a comparatively long distance (for the cortex). Given that there are very large numbers of rods and cones, it is unfeasible to send this information as a set of raw image measurements. Instead, the image is compressed by squeezing it through a narrow bottleneck of only ~1 million retinal ganglion cells per eye, which form the optic nerves. Each of these neurons is in charge of representing only a small portion of the visual field. Each neuron represents the difference between the sum of two local luminance averages of different spatial extents. This kind of sensitivity profile is known as a *center-surround* receptive field organization, and the transformed representation it produces is transmitted, via retinal ganglion cells and the LGN, to the early visual cortex (V1). This center-surround representation is fairly fundamental in vision and is one of the features that we would like to explain using information theory arguments.

The earliest representations of the image are not orientation-sensitive. Only when the visual information reaches V1 can the representation used by the majority of cells be thought of as coding for the orientation of local edge or line segments. The neurons in most layers of V1 are also sensitive to many other aspects of the visual input (motion, color, stereo, etc.), but it is this basic local edge-detection property that most modeling has attempted to explain.

This, then, gives us the basic phenomena to be explained: (1) a local center-surround representation of the input coming from the retina to the cortex, and (2) a representation in terms of local edges in the earliest area of the cortex. How can information theory help us here?

The Basic Logic of the Approach

The three models we next present all have the same basic structure. They all propose that these early visual representations maximize the amount of information transmitted about the world, subject to some set of constraints. They all agree that it is important to have a realistic model of the input distribution (the images we normally see), rather than taking as our signal random variations of light. They also all agree that representations should be thought of in terms of linear filters: modeling receptive fields as weighted sums of the pixels of the represented images they are to process. The one way that they differ is in the proposed nature of the constraints that prevent the trivial solution of maximizing information by simply having an output representation identical to the input.

Minimizing the Number of Units or the Amount of Time to Code: Principal Components Analysis

This model is based on the following assumptions. The first is that the input to the system—the probability distribution of natural images that are usually viewed by the visual system—can be well approximated in terms of a multivariate Gaussian (normal) distribution. This is equivalent to saying that if we know the means of the image intensities and the correlation between image measurements made at different locations by the photoreceptors, averaged across a large set of representative images, then we know everything that is to be known about the distribution—there is no additional non-Gaussian structure. This assumption is not in fact true (images do have such non-Gaussian structure; it's what stops them all from looking a bit like clouds), but it does have the virtue of simplicity.

The second assumption is that the noise on each of the receptors in the retina is Gaussian and independent between receptors.

The third and last assumption, and one that requires more explanation, is that the main constraint on information transmission is the limited number of linear filters available to represent the images. This is counterintuitive, as there are in fact far more neurons in V1 than cells in the retina. There are three ways to justify this assumption.

The first justification, which we will call the basic Gaussian justification, is based on its simplicity: given these assumptions, the optimal way of representing the input is in terms of the principal components of the input measurements. This is a well-understood statistical technique that provides us with the opportunity of an analytical understanding of the optimal solution, rather than simply a numerically optimized representation. Even if certain assumptions are not met (such as not having an obvious constraint on the number of output units), knowing the basic Gaussian solution will potentially be informative.

The second justification, which we call the Hebbian learning justification, comes about because of a connection between the optimal solution (the principal components solution), and the result of a number of models of low-level visual development. The well-known Hebb rule (Hebb 1949) specifies how the synapses (weight parameters) in a neural network should be changed depending on experience. That is, it states that neurons that fire together should wire together, and is usually interpreted in terms of a weight update rule (the rate of change of a weight) based on the correlation between two units.

Provided that a neural network is linear, and has some means to make sure that all units represent unrelated features, if it is trained using a number of variants of Hebb's rule, then it will converge when the weights span the same space as the principal components. Because Hebb's rule is a very popular model of visual development, knowing the optimum for this scenario is of interest.

The last reason is more subtle, and we call it the Time Limited Performance Justification. This is based on work by Korutcheva, Parga and Nadal (1997). They analyzed a binary system in terms of maximizing information transmission when presented with Gaussian input. A remarkable property of such a binary system is that as the ratio between the number of output and input units increases, binary processing becomes equivalent to linear processing; it is equivalent to principal components analysis. There are far more cells in V1 (output units) than there are cells in the retina (input units). Neurons communicate with spikes, so, when observed over a very short time window, neurons can either spike or not—they are effectively binary. This means that even if neurons in V1 are not well approximated by linear filters, as long as there are a lot of them (and there are), and we consider their behavior only over very short time scales (so they can either fire or not), if they are maximizing information with their input, then this will be equivalent to the whole system operating as if it were performing principal components analysis. Note that this statement does not make any predictions about the receptive fields observed physiologically or about psychophysical performance when subjects have a large amount of time, but does make strong predictions about the system's performance when operating over very short timescales (less than 30 ms).

This model leads us to a hypothesis: the representation of an image in early vision is in terms of its principal components. How do we go about testing this idea? One possible method is as follows: (1) Collect a set of images representative of the kind of world that vision is used to process. (2) Randomly sample local patches (say windowed 16 × 16-pixel patches) from these images as if eye movements were being made to them. (3) Calculate the principal components of these patches. (4) Compare these with what is known about the physiology or psychophysics of early vision. This is the procedure that Baddeley and Hancock (1991) followed, and the resulting two-dimensional filters are shown in figure 10.1.

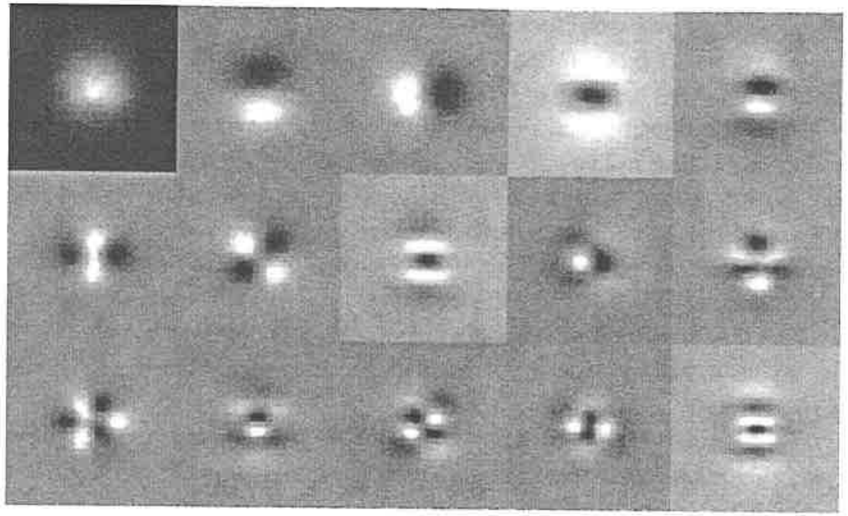

Figure 10.1
The first fifteen local principal components of a large collection of natural scenes based on work reported in Baddeley and Hancock 1991. The components are presented in terms of the amount of image variance they account for, with the component at the top left accounting for the most and components on the top row explaining more than lower rows. As can be seen, the early components (e.g., components 2–7) could be thought of as edge or line detectors of various types (responding optimally when stimulated with a line). Despite this, they do not resemble in detail receptive fields found in V1 (or retinal ganglion cell receptive fields).

How do these filters compare to what is known about physiology? Though there is some resemblance (see components 2–7), many of the components are a poor match. More importantly, although components 2–7 *qualitatively* could be described as edge detectors, *quantitatively* they are far from a good match to what is known about receptive fields. This limitation is important, because developmental models such as that of Linsker (1986) converge to the principle components of the input. The approximate "edge detector–like" properties of these models have been argued to be evidence, but the actual properties of even the components that do resemble receptive fields are quantitatively rather different (see later discussion for models that do make more accurate predictions).

This leaves the last suggestion: that the entire system, when given only very short periods of time, appears to resemble the principal components. To understand how to test this idea, a few words on what determines the nature of the components are in order. It turns out that there are three important factors that determine the nature of the components of natural images.

First, the power spectra of the input images are important. Principle components attempt to capture as much of the variance as possible, and the power spectra of the images tells us where this variance is. It is a relatively robust characteristic of natural images, from a wide range of sources, that the power at a given frequency is inversely proportional to its frequency (Field 1987): the vast majority of the power in natural images is at low frequencies. This means that the components also have to be tuned to low spatial frequencies in order to capture input variance.

The second determinant of the components is the nature of the window used when sampling local image patches. Though this determinant has nothing to do with image statistics, it can have a large effect on the derived components. Simply sampling square patches from images always means that any structure will be aligned to the sampling window. Therefore, to avoid artificially imposing structure on components, we windowed the samples with a Gaussian approximation before subsequent processing. This leaves the last characteristic: how anisotropic the image statistics of the world are. In particular, how rapidly does the correlation between image intensity measurements decay as a function of angle?

In artificial images, the correlation decays can be equal for all directions, but in naturalistic images, particularly those of wide open landscapes, the correlation decays considerably quicker in the vertical direction compared to the horizontal (Baddeley 1996). This difference results in the structure of vertical and horizontal components being different. In particular, two components can be thought of as line detectors: in figure 10.1, components 4 and 6 respond optimally when a light or dark line at the appropriate orientation is placed in the center. These two components are not simply rotations of each other, but have slightly different structure. If the components provide a good description of the whole system operating at very short intervals, a match to this difference in structure should be measurable psychophysically. In fact, this proved to be true. In some extremely labor-intensive experiments, in which subjects had to detect a single differently orientated line in the presence of a large number of other lines, Foster and Ward (1991) found that at very short presentation times (40 ms) subjects' performance was well summarized by two line-detection mechanisms operating. Figure 10.2 shows the response of these two mechanisms as a function of orientation, together with the orientation tuning of the two line-detecting principle components. As can be seen, the match is essentially perfect.

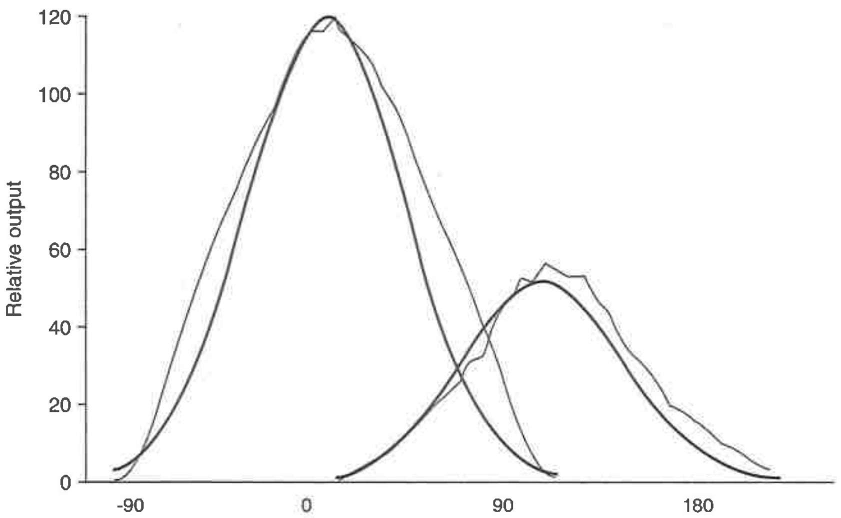

Figure 10.2
Does the orientation tuning of components 4 and 6 in any way match that estimated with short time presentation psychophysics? Shown is the response of these two components to an optimal bar rotated through 180 degrees (thin line), and for comparison the orientation tuning inferred experimentally by psychophysical methods for very short time presentations (thick line), based on the work of Foster and Ward (1991). As can be seen, though the components do not match individual receptive fields, the match of the system's orientation sensitivity is nearly perfect.

To conclude this section on principal components-based models: models that propose that information about images is maximized, and that the main constraint on communication is the number of filters to represent the image, result in a model that essentially extracts the principal components of images. The early components vaguely quantitatively resemble what is known about early receptive fields, but qualitatively they are not a good match. When people have only extremely short amounts of time to view things, then aspects of the system as a whole do appear to be consistent with information maximization. However, we very rarely only have 40 ms to view something. Given longer viewing times, human performance changes (improves), and the use of a binary code is no longer justified. Therefore, the principal components model is not, in general, a satisfactory model of human performance.

Beginning to Make the Model More Realistic: Metabolic Constraints on Firing Rate (or Sparsity)

That there are more neurons in V1 than receptors in the retina forms a specific coding challenge (Lewicki and Sejnowski 1998), and this abundance of V1 neurons is problematic for the approach described previously. Thus, to respond to this challenge, we are led to a second model in which instead of constraining the *number* of neurons representing the image, we place a constraint on their *level of activity*.

There are a number of justifications for this. By far the simplest is a metabolic argument: given that there are a vast number of neurons in V1 (the largest for any cortical area), they will have a significant cost in terms of energy consumption. As well as minimizing energy consumption, simply to minimize the amount of food we need to find and eat, highly metabolically active neural tissue also presents both resource transport and heat dissipation problems (Falk 1992; Corrard 1999). The brain is a highly metabolically active organ (Rolfe and Brown 1997), and any representation that required less energy would be favored by evolution, as long as it did not hurt performance too much.

There are a number of factors that determine the energy consumption of a piece of the brain (Attewell and Laughlin 2001); indeed, simply maintaining neurons and using synapses uses resources (discussed shortly). Here, however, we reduce energy consumption by minimizing the activity of the neurons; that is, we minimize their average firing rate. The idea is simple: rather than simply looking for the representation that captures as much information about the image as possible, we simultaneously try to minimize the amount of energy used to represent images, where the main form of energy use we concentrate on is the energy used to fire a neuron. If we again use a Gaussian approximation to the information, this optimization is very simple to implement in terms of a neural network learning rule. This was done first in Fyfe and Baddeley (1995), where we used a neural network to extract the filters with very peaked output distributions. In these simulations, we optimized a statistical quantity called *kurtosis* rather than minimizing the absolute output, but the effect is very similar. We, however, calculated only the first three optimal filters, and the nature of the full solution was, therefore, not completely clear. Our work was followed by that of two groups (Harpur and Prager 1996; Olshausen and Field 1996), again both using neural networks, and although the algorithm proposed by Harpur and Prager has computational advantages over both ours and the Olshausen and Field algorithm (one of their methods does not require a prewhitening filter), the results from the Olshausen and Field model were presented particularly nicely (and appeared in *Nature*), so the result is usually attributed to them. The basic result, the optimal receptive fields, when optimized to maximize information transmission while minimizing energy consumption, is shown in figure 10.3.

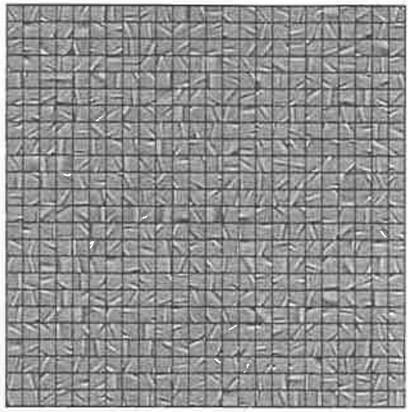

Figure 10.3
"Edge detectors" as derived by simultaneously optimizing the information transmitted about images, while minimizing the average firing rate of the cells. As can be seen, the result of such an optimization is a large number of local, oriented edge detectors that provide a good qualitative match to the receptive fields found in V1. If motion is also taken into account, as shown by van Hateren (1983), the match is quantitative as well.

This, although not obvious to someone not working in the field, was a major breakthrough. Unlike the principal components, which only very approximately resembled the receptive fields found in V1, when the fact that real-world images move was taken into account, the modeled receptive fields provided both a good qualitative and quantitative account of those found in V1 (van Hateren and van der Schaaf 1998).

Given this match, are there any other characteristics of V1 that lend weight to this interpretation? One prediction concerns the probability distribution of firing rates of cells when stimulated with natural scenes (rather than the artificial stimuli more commonly used). It can be shown that, given a constraint on the average firing rate, the neuron firing rate distribution that maximizes the output entropy (and hence potential information transmission), is an exponential distribution (Lennie 2003b). This prediction is easy to test, and in Baddeley et al. (1997) we carried out the simplest experiment: we recorded the distribution of firing rates

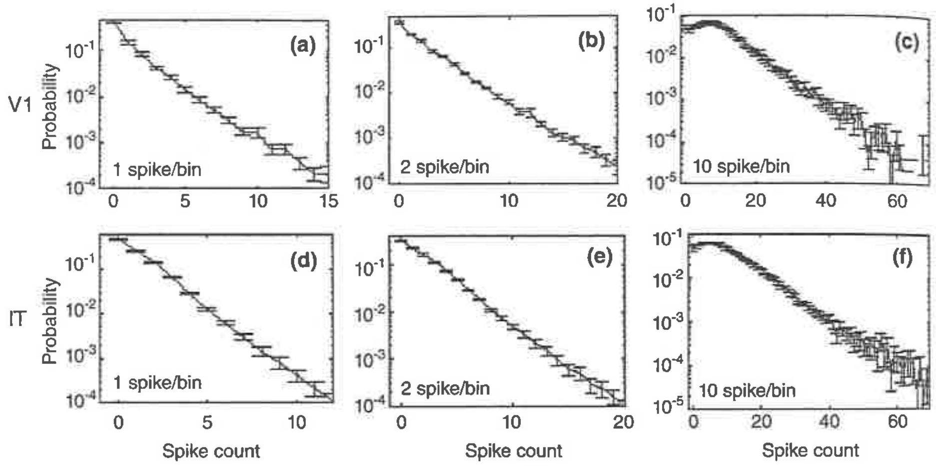

Figure 10.4
The firing rate distributions of cells in two cortical areas (V1, and IT, an area involved in object recognition), when these cells are processing nature videos. In the noiseless limit, given a constraint of the average firing rate, the cells will be transmitting the maximum amount of information when these distributions are exponential (or, because the y-axis is logarithmic, these plots form straight lines). As can be seen, over an order of magnitude range of time, the spike count distributions are indeed well approximated by an exponential distribution lending support to the idea that cells in V1 (and other areas) are maximizing their information transmission while minimizing their average firing rate. Based on Baddeley et al. 1997.

in cells in V1 while playing those videos judged to be representative of the animals' natural environment. The results are shown in figure 10.4. As can be seen, the distributions are very well fit by exponential distributions, and though there are systematic differences, when the fact that the system is not completely noise-free is taken into account, the fit is essentially perfect.

Before we move onto the final model, one point is worth mentioning. Though we have here interpreted the derived receptive fields in terms of minimizing energy consumption, this is far from the most popular interpretation. More usually, they are interpreted in terms of maximizing *sparsity*. Despite one common measure of sparsity being equivalent to minimizing energy consumption, we are not fans of this interpretation. Sparsity traditionally makes sense when you have a system that has only a very few active elements (in this case, neurons), while the majority are silent. At first, this appears an appropriate definition: neurons are in fact either far more on, or far more off, than predicted if their outputs were Gaussian-distributed. In fact, though, when averaged over even short periods, neurons are hardly ever silent, and compared to distributions with the same firing rate, their distribution is about as varied as it could possibly be, showing very little bimodality as would be predicted by a silent versus nonsilent interpretation. In a way, both interpretations state that neurons should minimize their average activity. But although minimizing metabolic activity unambiguously specifies how to quantify this, sparsity has a rather large number of definitions (which, though they give very similar results, is less satisfying), and does not predict that the optimal distribution will be exponentially distributed, as observed, but simply very long-tailed.

Minimizing Synaptic Rather than Firing Rate Energy Consumption

This, then, leaves us with a reasonably satisfying explanation of the receptive field in V1, but we appear to have missed a stage. The retinal ganglion cells do constitute an information bottleneck, and it would seem plausible that their receptive fields could be explained in a similar framework. Their center-surround receptive fields, however, do not resemble those found when minimizing average firing rate. Based on previous research (Vincent and Baddeley 2003; Vincent et al. 2005), our working hypothesis is to assume that because the information transmission bottleneck operating is so extreme, no meaningful energy saving can be made by constraining the firing rate. In contrast, as shown in figure 10.5, rather large energy consumption savings can be made by minimizing the energy used by synapses (which is quantified by the sum of the filter coefficients).

Constraints on synaptic energy consumption, which uses up a significant proportion of the entire brain's energy budget, were explored by Vincent and Baddeley (2003). When a cost function that takes into account synaptic (rather than firing rate) energy use is maximized, the optimal representations are very different, and are shown in figure 10.6. Quantitatively (and qualitatively), these representations provide a very good characterization of the observed receptive fields of retinal ganglion cells.

Furthermore, the explicitly energy-saving approach was applied to simultaneously calculate the optimal receptive fields for a simplified two-layer retinocortical visual system (Vincent et al. 2005). Although it is commonly assumed that the oriented V1 receptive fields derive appropriate input from multiple center-surround neurons, as far as we are aware, this was the first study to confirm that such an arrangement is in fact an optimal solution to something. But in order to do this, it was insufficient

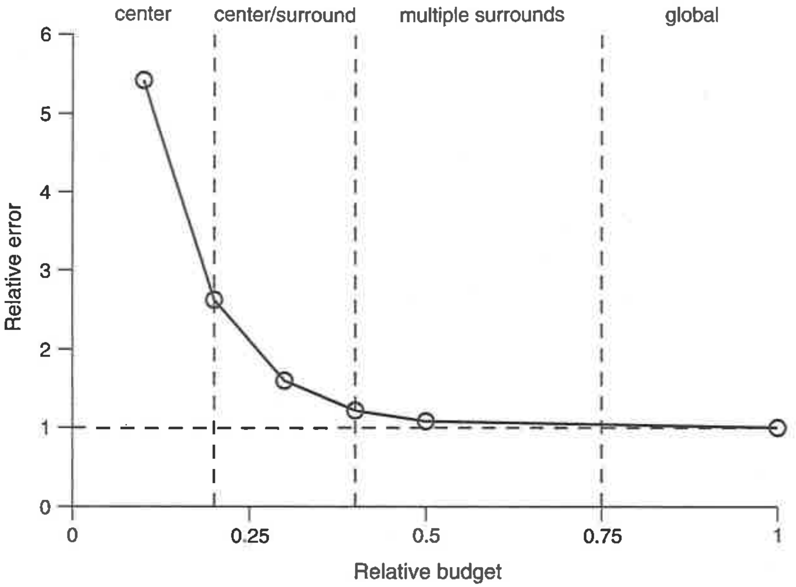

Figure 10.5
Large saving of the synaptic budget can be made at relatively little cost in terms of the amount of information transmitted. Shown are the results of six filter systems, each optimized with a different "budget" to be spent on synapses, where a budget of one corresponds to the optimal unconstrained solution. As can be seen, a 70 percent saving of metabolic cost associated with synapses can be made without severely affecting information transmission. When there is no constraint, the network simply spans the space of the principal components (and does not match observed physiology). In contrast, when only 20 to 40 percent of the energy is spent on synapses, the receptive fields in contrast form localized receptive fields. Examples of these are shown in figure 10.6. Based on Vincent and Baddeley 2003.

to simply maximize information; metabolic costs also had to be considered.

In summary, maximizing information about the natural input has provided a reasonable way of approaching the nature of the early cortical representations. If we assume that the main constraint is simply the number of units (or that the system is operating over a very short time, so neuronal outputs can be treated as binary), the optimal solution is the principal components of the input. This does not account for the form of any observed early receptive fields, but does explain short presentation time psychophysical performance. If we allow a larger number of neurons to represent an image, but minimize their average firing rate, then the optimal receptive fields provide a very good account of those measured

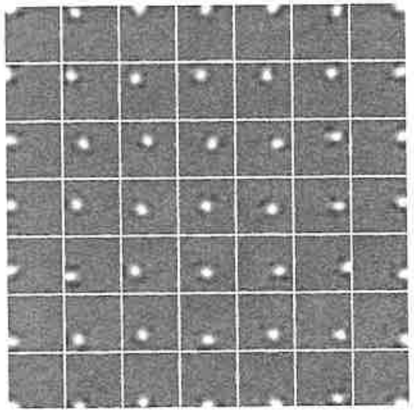

Figure 10.6
The optimal receptive fields of a system simultaneously maximizing information transmission while minimizing the amount of energy used on synaptic transmission. As can be seen, the receptive fields have a local center-surround organization, and the parameters of these receptive fields provide a good match to the properties of retinal ganglion cells. Based on Vincent and Baddeley 2003.

in V1. If, however, the constraint is placed on the metabolic cost of the synapses, then the optimal representation provides a good match to the receptive fields of the retinal ganglion cells.

Problems and Where Next?

The results presented in this discussion provide at least a potentially interesting framework for understanding early perception, but there is one very important theoretical worry that needs to be more directly addressed if we are to make progress. The worry is that we have treated the input as the pattern of light that arrives at the eye. It is easy to forget when studying vision; animals fundamentally do not care about light, but about the world that generated that light. This means that maximizing the information about the light is a sensible strategy only if all aspects of the variation of light are equally informative about behaviorally important aspects of the world. This is very unlikely to be true.

Unfortunately, approaches based on maximizing the information about the behaviorally important aspects of the world require estimates not of the statistics of natural images (which is easy), but the statistics of the behaviorally relevant aspects of the world that the animal needs to know about. It is the information about the world that should be

(a) Natural Image

(b) Signal (Reflectance) Image

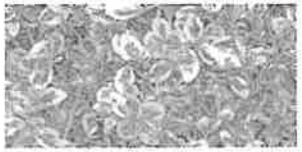

(c) Noise (Illumination) Image

(d)

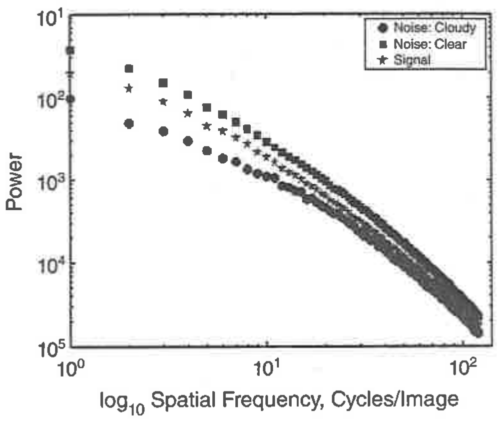

Figure 10.7
Information about the world is more important than information about light. Although at an early stage of research, we have been attempting to decompose images into variation that is relevant to properties of the world, and variation that is not informative. This figure shows part of the process, where, using parallel illumination of a scene, we can decompose it into the contribution due to reflectance, and the contribution due to illumination. The three panels to the left show: a natural image (a), which can be regarded as the product of the reflectance signal (image b), and illumination noise (image c). Panel (d) shows the power spectra of natural signal and noise images, averaged over horizontal and vertical orientations. As can be seen, the majority of variation in natural images is due to variations in illumination and does not tell us about the properties of the world. The spatial frequency spectra of illumination variation and reflectance variation are also very similar, meaning that it cannot be removed by simple linear filtering. This result is robust across a number of illumination environments. Representations that maximize the information about the indirectly observed reflectance will have to perform processing on the images (we are using Bayesian techniques) rather than simply recode them in order to maximize information.

maximized, not the information about the light generated by that world. This presents two difficulties: specifying what aspects of the world are behaviorally relevant, and quantifying them. Despite these technical difficulties and a number of false starts, we believe we are making progress on one particularly simple aspect of the world: the surface reflectance properties of objects in the world.

This chapter is not the place to describe our latest information-theoretical measurements, but figure 10.7 shows the most successful method we have developed to measure, in naturalistic settings, not only the characteristics of images, but also of the illumination and objects' surface reflectance properties that are associated with those images. This is just an example, but we believe that making models that maximize the information presented not about images, but about the behaviorally important aspects of those images, will be required if we are to make further progress in understanding how and why we have the early representations of the world that we do.

References

Attewell, David, and Simon B. Laughlin. 2001. An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow and Metabolism* 21 (10): 1133–1145.

Baddeley, Roland J., and Peter J. B. Hancock. 1991. A statistical analysis of natural images matches psychophysically derived orientation tuning curves. *Proceedings of the Royal Society B: Biological Sciences* 246 (1317): 219–223.

Baddeley, Roland J. 1996. The correlational structure of natural images and the calibration of spatial representations. *Cognitive Science* 21 (3): 351–372.

Baddeley, Roland, J. L. F. Abbott, Michael C. A. Booth, Frank Sengpiel, Tobe Freeman, Edward A. Wakeman, and Edmund T. Rolls. 1997. Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proceedings of the Royal Society B: Biological Sciences* 264 (1389): 1775–1783.

Corrard, François. 1999. Selective brain cooling. *Archives de Pediatrie* 6 (1): 87–92.

Derrington, Andrew M. 2000. Vision: Can colour contribute to motion? *Current Biology* 10 (7): R268–R270.

Falk, Dean. 1992. *Braindance: New discoveries about human origins and brain evolution.* New York: Henry Holt.

Field, David J. 1987. Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America. A, Optics and Image Science* 4 (12): 2379–2394.

Foster, David H., and Patrick A. Ward. 1991. Horizontal-vertical filters in early vision: Predict anomalous line-orientation identification frequencies. *Proceedings of the Royal Society B: Biological Sciences* 243 (1306): 83–86.

Fyfe, Colin, and Roland J. Baddeley. 1995. Finding compact and sparse- distributed representations of visual images. *Network* 6 (3): 333–344.

Harpur, George, and Richard Prager. 1996. Development of low entropy coding in a recurrent network. *Network* 7 (2): 277–284.

Hebb, D. O. 1949. *Organization of behavior: A neuropsychological theory*. New York: Wiley.

Hubel, David H., and Torsten N. Wiesel. 1962. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology* 160 (1): 106–154.

Korutcheva, Elka, Nestor Parga, and Jean-Pierre Nadal. 1997. Information processing by a noisy binary channel. *Network* 8 (4): 405–424.

Lennie, Peter. 2003a. Receptive fields. *Current Biology* 13 (6): 216–219.

Lennie, Peter. 2003b. The cost of cortical computation. *Current Biology* 13 (6): 493–497.

Lewicki, Michael S., and Terrence J. Sejnowski. 1998. Learning overcomplete representations. *Neural Computation* 12 (2): 337–365.

Linsker, Ralph. 1986. From basic network principles to neural architecture: Emergence of spatial opponent cells. *Proceedings of the National Academy of Sciences of the United States of America* 83 (19): 7508–7512.

Olshausen, Bruno A., and David J. Field. 1996. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381 (6583): 607–609.

Rolfe, David F. S., and Guy C. Brown. 1997. Cellular energy utilization and molecular origin of standard metabolic rate in mammals. *Physiological Reviews* 77 (3): 731–758.

Rucci, Michele, and Antonino Casile. 2005. Fixational instability and natural image statistics: Implications for early visual representations. *Network (Bristol, England)* 16 (2–3): 121–138.

Rucci, Michele, Gerald M. Edelman, and Jonathan Wray. 2000. Modeling LGN responses during free-viewing: A possible role of microscopic eye movements in the refinement of cortical orientation selectivity. *Journal of Neuroscience* 20 (12): 4708–4720.

Shannon, Claude E. 1948. A mathematical theory of communication. *Bell System Technical Journal* 27 (3, 4): 379–423, 623–656.

Simoncelli, Eero P., and Bruno A. Olshausen. 2001. Natural image statistics and neural representation. *Annual Review of Neuroscience* 24 (March): 1193–1216.

Valeton, J. Matthee, and Dirk van Norren. 1983. Light adaptation of primate cones: An analysis based on extracellular data. *Vision Research* 23 (12): 1539–1547.

van Hateren, J. H. 1983. Spatiotemporal contrast sensitivity of early vision. *Vision Research* 33 (2): 257–267.

van Hateren, J. Hans, and Arjen van der Schaaf. 1998. Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society B: Biological Sciences* 265 (1394): 359–366.

Vincent, Benjamin T., and Roland Baddeley. 2003. Synaptic energy efficiency in retinal processing. *Vision Research* 43 (11): 1283–1290.

Vincent, Benjamin T., Roland Baddeley, Tom Troscianko, and Iain Gilchrist. 2005. Is the early visual system optimised to be energy efficient? *Network* 16 (2–3): 175–190.

Warrant, Eric J. 1999. Seeing better at night: Life style, eye design and the optimum strategy of spatial and temporal summation. *Vision Research* 39 (9): 1611–1630.