To appear: Current Opinion in Neurobiology, Vol 13, April 2003

Vision and the Statistics of the Visual Environment

Eero P Simoncelli

Summary
It is widely believed that visual systems are optimized for the visual properties of the environment inhabited by the host organism. A specific instance of this principle known as the Efficient Coding Hypothesis holds that the purpose of early visual processing is to produce an efficient representation of the incoming visual signal. The theory provides a quantitative link between the statistical properties of the world and the structure of the visual system. As such, specific instances of this theory have been tested experimentally, and have be used to motivate and constrain models for early visual processing.

Address
4 Washington Place, Room 809, New York, NY 10003-6604, USA; phone: (212) 998-3938; fax: (212) 995-4011; email: eero.simoncelli@nyu.edu

Introduction

One of the primary roles of theory in the sciences is to provide fundamental principles that explain why the natural world is constructed as it is. In biology, the main example of such a principle is the theory of evolution by natural selection. In the context of vision, this sort of philosophy was championed by David Marr, who argued that it is essential to consider the visual system at an abstract computational level in order to understand its design [1]. Two specific instances of this philosophy provide a quantitative link between the statistical properties of the visual environment and the structure of biological visual systems: The so-called *Efficient Coding Hypothesis*, and the formulation of early vision problems in terms of *Bayesian estimation or decision theory.* As such, specific instances of both theories may be tested experimentally, and may be used to motivate and constrain models for vision.

The Efficient Coding Hypothesis has been a central topic in a variety of recent workshops and meetings [2], as well as several review articles [3, 4] and a special journal issue [5]. Recent results are interesting, albeit controversial. Because of this flurry of recent interest, and because Bayesian theories of vision have been reviewed in a number of other places [e.g., 6, 7, 8, 9, 10, 11], I willl focus almost entirely on efficient coding in this article, with emphasis on articles published in 2001 or later.

Efficient Coding

The theory of information plays a natural role in models of neural systems, by providing abstract but unique quantitative definitions for information [12]. Barlow [13] recognized the importance of information theory in this context, and hypothesized that the efficient coding of visual information could serve as a fundamental constraint on neural processing. That is, a group of neurons should encode information as compactly as possible, in order to most effectively utilize the available computing resources. Mathematically, this is expressed as a desire to maximize the information that neural responses provide about the visual environment.

The simplest form of this hypothesis (in particular, ignoring the noise in neural responses) decouples naturally into two separate statements: one regarding the statistics of individual neural responses and a second regarding the joint statistics of the response of a population [14, 3, 15]:

- The responses of an individual neuron to the natural environment should fully utilize its output capacity, within the limits of any constraints on the response (e.g., maximum firing rate).
- The responses of different neurons to the natural environment should be statistically independent of each other. In other words, the in-

formation carried by each neuron should not be redundant with that carried by the others. This is also consistent with a notion that the visual system strives to decompose a scene into statistically independent constituents (e.g., individual objects).

More detailed discussions of these ideas, as well as the role of noise, may be found in many other references [16, 13, 17, 18, 19, 20, 21, 22, 3, 15].

Criticisms

A variety of criticisms have been voiced regarding the Efficient Coding Hypothesis. A number of these represent misconceptions about the theory, some are aimed at particular variants of the theory, some are about practical experimental issues, whereas others are more fundamental. Below is brief discussion of some of these (see the recent review article by Horace Barlow for additional discussion [4]).

The purpose of vision. It has often been argued that efficient coding of visual information is irrelevant because the purpose of vision is not to encode or reconstruct the visual world. There is some truth to this criticism, in that the hypothesis does not take into account how the information that has been extracted is to be used. This may be viewed as either an advantage (because one does not need to assume any specific visual task or goal, and does not even need to specify what is being represented) or a limitation (because tasks and goals are clearly relevant for visual processing). More complete theories, such as that given by Bayesian estimation and decision, can take into account both the statistical structure of the environment and the visual task or goal.

Relevance of information theory. A second criticism of the Efficient Coding Hypothesis is that “information theory is irrelevant because the brain is not concerned with bits.” Bits are just a standard choice of unit for information, but the abstract definition of information is well-motivated, unique, and is most certainly relevant to the brain.

Experimentally observed dependency. Another criticism of the Hypothesis is that “some experimental data from multi-neuron recordings show correlation, synchronization, or other forms of statistical dependency between neurons”. Most such experiments do not use naturalistic stimuli, and thus dependencies in the neural responses are not directly relevant to the hypothesis. In addition, recent studies suggest that responses to natural stimuli in primary visual cortex are relatively independent[23, 24, 25]. Finally, even if one were to observe dependencies in neural responses under natural stimulus conditions, the hypothesis states only that the system strives for independence: The constraints of neural processing may prevent actual achievement. Perhaps a more realistic expectation, then, is that successive stages of processing (e.g., along an ascending sensory pathway) should reduce statistical dependence [26].

Over-representation in cortex. A further criticism is based on a comparison of the number of retinal ganglion cells to the number of neurons in primary visual cortex. Critics argue that “the number of neurons devoted to processing sensory information seems to expand as one goes deeper into the system, suggesting that the brain increases redundancy.” This argument usually assumes, however, that the coding capacity of all neurons (and in particular those in retina and cortex) is the same. The distribution of information amongst more neurons does not necessitate more redundancy if the form of neural coding employed by those neurons is allowed to differ. For example (as Barlow points out in [4]) cortical neurons tend to have lower firing rates, and may well use a different form of code than retinal neurons. In addition, cortical neurons have more complex temporal dynamic properties (e.g., adaptation) that may serve to encode information over longer timescales. Although the redundancy of retinal and V1 neurons has not been experimentally compared, a related comparison in the auditory system was able to demonstrate a reduction of redundancy [26].

Experimental impracticality. Many authors have pointed out that "Estimation of information-theoretic quantities requires enormous amounts of data, and is thus impractical for experimental verification." This is a significant problem, especially since commonly used estimators of information are also known to be heavily biased. Nevertheless, cases of successful experimental measurement give reason for optimism (see below).

Definition of input and output. The Hypothesis depends critically on the probability distribution of natural images and the definition of neural response, both of which are underconstrained. Thus the theory is not as assumption-free as one is led to believe. In my opinion, this is the most fundamental problem with the hypothesis. The input distribution is typically not defined explicitly , but is assumed to be well represented by a collection of calibrated "naturalistic" images. One must also specify which neurons are meant to satisfy the hypothesis (e.g., neurons within a particular cell class, or within a specific visual area, or across multiple visual areas), and how their responses are to be measured (e.g., mean firing rates vs. individual spike times). Again, cases of successful experimental measurement give reason for optimism.

Importance of noise. A final criticism of the Efficient Coding Hypothesis is that "commonly used versions of the theory that ignore noise and other physical constraints are too simplistic." This is a valid criticism, but in many cases may not constitute a fatal flaw. In particular, even simplistic forms of the theory seem to make interesting predictions, and many authors have developed more sophisticated versions of the theory that do include physical constraints such as noise (see below).

Testing the Hypothesis

Although the efficient coding hypothesis is roughly fifty years old, it has only recently been explored quantitatively. This recent progress is due to three fundamental improvements: (1) we have a much better understanding of early sensory processing; (2) mathematical and engineering tools have been developed to describe and manipulate more complex statistical models; and (3) advances in computing and imaging technologies allow us to gather and manipulate vast quantities of image data, both for statistical modeling purposes, and for use as experimental stimuli.

There are two basic methodologies for testing and refining the efficient coding hypotheses. The direct approach is to examine the statistical properties of neural responses under natural stimulation conditions [e.g. 27, 21, 28, 29, 23]. An alternative approach is to use the statistical properties of natural images to constrain or derive a model for early sensory processing [e.g. 30, 31, 32, 33, 34, 19, 35, 36, 37, 38]. Below, I'll review some recent examples of each of these.

Experimental Tests

In recent years, there have been a number of interesting experimental articles examining neural responses to naturalistic images or image sequences (see [39] for a review). These authors measure efficiency in a variety of different ways, and while some of these seem confirmatory of the efficient coding, others seem inconsistent. But on the whole, these recent results are advancing our understanding of the issues.

For example, Baddeley et. al. [29] had shown that firing rate distributions of cat V1 neurons and Monkey IT neurons were exponential under naturalistic conditions, which implies optimal information transmission for a fixed average rate. A subsequent study of monkey IT neurons found that only a minority were well described by an exponential firing distribution [40]. De Polavieja later argued that the discrepancy was due to the fact that the exponential solution is correct only in the noise-free case, and showed that by taking noise into consideration, one could account for the data [41].

Similarly, previous results suggested that retinal ganglion cells exhibit strong correlations in firing, and that these patterns could provide useful information [42] (although these experiments did not directly address efficient coding, as they were not based on natural image stimuli). A recent article by Nirenberg et. al. [24] argues that retinal ganglion cells act as independent encoders, based on multi-cell recordings with natural stimuli, and a novel (but controversial) choice of redundancy measure. Reich et. al. [43] find that responses of V1 neurons are nearly independent under non-natural (approximately white noise) stimulation. This seems inconsistent with efficient coding, since one typically expects efficiency to degrade for non-natural stimuli (e.g., [21]). Wiener et. al. [44] find that the information encoded by the spike counts of V1 neurons over moderate length intervals is nearly the same for various artificial stimuli as it is for naturalistic stimuli, but Vinje and Gallant [25] recently reported that the presence of natural stimuli in the nonclassical receptive field increases several measures of informational efficiency. Finally, although the study concerns the auditory system, it is worth mentioning an interesting article by Chechick et.al. [26], that provides one of the first direct tests of the reduction of redundancy as one ascends a sensory pathway.

A number of studies also suggest that the visual system exhibits improved performance under naturalistic input conditions. For example, Lewen et. al. [45] demonstrated that the H1 neuron in the blowfly responds over a substantially broader range of velocities for outdoor scenes as compared with indoor scenes. Kern et. al. [46] show that blowfly neurons can encode turning directions independent of environmental texture and spatial structure, but only when the environment is “natural”. Although the connection has not been carefully established, the efficient coding hypothesis would seem to suggest that optimal characterization of a neural system might be best performed using naturalistic stimuli. Experiments to demonstrate this have been performed [e.g., 47], and recently developed analysis techniques may provide the necessary tools for characterization [48, 49].

Optimal Models

The second method of testing the hypothesis is to derive a model for efficient coding of the environment and then compare this with physiological data. Many of the original attempts to accomplish this were constrained to linear filtering and second-order statistical modeling [e.g., 30, 31, 32, 33, 34, 19]. Recently, a number of authors have established relationships between higher-order statistical properties and linear filtering [e.g., 35, 36]. Others have used various forms of nonlinear processing such as divisive gain control [38, 50, 51, 52]. In addition, some authors are beginning to explore cascades of redundancy-reduction stages. For example, Hoyer and Hyvarinen [53] have developed a model that forms contour-like receptive fields on the basis of the statistical properties of V1 complex cells layer under natural stimulus conditions.

Extensions

There are a variety of auxilliary constraints or extensions that authors have recently used to augment the efficiency hypothesis. Several authors have tried to incorporate metabolic costs as a constraint [54, 55]. Recently, Balasubramanian and Berry [56] demonstrated that retinal ganglion cells in tiger salamander are optimized to transmit visual information at minimal metabolic cost, assuming the symbols of the neural code are represented by spike bursts of a given length. Although the result is interesting, the data were recorded and under artificial visual stimulation conditions, and thus do not bear directly on the issue of environmental statistics. Similarly, Gottschalk shows that a hyperbolic ratio provides an optimal neural response nonlinearity in that it maximizes information transmission subject to a simple model of neural costs that includes both synaptic strength and spike rate [57]. Again, the result is implicitly based on white noise input stimuli, and thus does not bear directly on the Efficient Coding Hypothesis.

It has been proposed that efficient coding might also apply to adaptive processes [58, 59, 60]. Bialek and colleagues [61, 62] have demonstrated that adaptation of the fly's H1 neuron to the variance of a white noise stimulus appears to optimize information trasmission.

Finally, a number of authors have augmented efficient coding with other constraints or principles. Several authors have derived models for V1 complex cells by maximizing the independence or temporal coherence of cells constructed using nonlinear combinations of linear subunits (e.g., sums of squared linear filter responses) [63, 64, 65]. Balboa and Grzywacz argue that lateral inhibition in the retina retina is inconsistent with pure information maximization, and provide an alternative hypothesis that the system is designed to detect and represent image features in the presence of photon absorption noise [66]. And Twer and Mcleod showed that color representation in monkey ganglion cells is best understood in the context of optimizing the average discriminability of color signals in the natural environment [67].

Conclusion

The recent resurgence of interest in efficient coding has produced a number of interesting experimental and theoretical results. But rather than merely solidifying or disproving the basic hypothesis, this body of work demonstrates the complexity and subtlety of establishing firm quantitative connection between the properties of natural stimuli and neural response. Most of the difficulty lies in the definition of the input (what is a "natural" image?) and the output (which neurons should be considered, and how does one define "neural response"?), as well as the complexities of incorporating realistic constraints (e.g., noise, metabolic costs) and realistic computational goals. As we wrestle with these issues, it becomes clear that the value of efficient coding has already far exceeded its role as a hypothesis to be confirmed or disproven. The effort directed at studying this hypothesis has significantly advanced the field by increasing the general level of quantitative investigation, by leading us to carefully study and model the visual environment from a statistical standpoint, by forcing us to consider the influence of environmental statistics on neural response, and by encouraging us to examine neural response to natural stimuli. And although it seems improbable that Efficient Coding will suffice as the sole principle for understanding sensory system design, it is clear that it will continue to play an important role.

References

[1] David Marr. **Vision: A Computational Investigation into the Human Representation and Processing of Visual Information**. W. H. Freeman and Company, San Fransisco, 1982.

[2] Recent workshops on Natural Scene statistics and their relevance to biology:
http://redwood.ucdavis.edu/bruno/NIPS96/abstracts.html
http://www.klab.caltech.edu/~pam/cns96/summaries/natscenes96.html
http://goethe.klab.caltech.edu/~pam/nssmeeting
http://www.cnl.salk.edu/~rao/workshop.html
http://www.klab.caltech.edu/~pam/nss2000.html
http://gandalf.psych.umn.edu/~kersten/kersten-lab/CompNeuro2002/index.htm
http://www.klab.caltech.edu/~pam/NSS2002.htm
http://www.ski.org/Carandini/symposium.

[3] E P Simoncelli and B Olshausen. **Natural image statistics and neural representation**. *Annual Review of Neuroscience*, 24:1193–1216, May 2001.

[4]•• H B Barlow. **Redundancy reduction revisited**. *Network: Computation in Neural Systems*, 12:241–253, 2001.
A retrospective look at the efficient coding hypothesis by one of its original proponents. The author argues that the original statement of the hypothesis overemphasized issues of compression, but was basically correct in stating the importance of redundancy.

[5] P Reinagel and S Laughlin. **Editorial: Natural stimulus statistics**. *Network: Computation in Neural Systems*, 12(3), Aug 2001.

[6] D C Knill and W Richards, editors. **Perception as Bayesian Inference**. Cambridge University Press, 1996.

[7] R P N Rao and D H Ballard. **Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects**. *Nature Neuroscience*, 2(1):79–87, 1999.

[8] D Kersten and P Schrater. **Pattern inference theory: A probabilistic approach to vision**. In D Heyer and R Mausfeld, editors, *Perception and the Physical World*, chapter 7, pages 191–228. Wiley, West Sussex, UK, 2002.

[9] P Schrater and D Kersten. **Vision, psychophysics and Bayes**. In R P N Rao, B A Olshausen, and M S Lewicki, editors, *Probabilistic Models of the Brain*, chapter 2, pages 37–60. MIT Press, Cambridge, MA, 2002.

[10] P Mamassian, M Landy, and L T Maloney. **Bayesian modeling of visual perception**. In R P N Rao, B A Olshausen, and M S Lewicki, editors, *Probabilistic Models of the Brain*, chapter 1, pages 13–36. MIT Press, Cambridge, MA, 2002.

[11] L T Maloney. **Statistical decision theory and biological vision**. In D Heyer and R Mausfeld, editors, *Perception and the Physical World*, chapter 6, pages 145–189. Wiley, West Sussex, UK, 2002.

[12] T Cover and J Thomas. **Elements of Information Theory**. Wiley, New York, 1991.

[13] H B Barlow. **Possible principles underlying the transformation of sensory messages**. In W A Rosenblith, editor, *Sensory Communication*, pages 217–234. MIT Press, Cambridge, MA, 1961.

[14] S Panzeri, S R Schultz, A Treves, and E Rolls. **Correlations and encoding of informration in the nervous system**. *Proc Roy Soc B*, 266:1001–1012, 1999.

[15] P Dayan and L F Abbott. **Theoretical Neuroscience**. MIT Press, Cambridge, MA, 2001.

[16] F Attneave. **Some informational aspects of visual perception**. *Psych. Rev.*, 61:183–193, 1954.

[17] M V Srinivasan, S B Laughlin, and A Dubs. **Predictive coding: A fresh view of inhibition in the retina**. *J. R. Soc. Lond. B*, 216:427–459, 1982.

[18] J J Atick and A N Redlich. **Towards a theory of early visual processing**. *Neural Computation*, 2:308–320, 1990.

[19] J H van Hateren. **A theory of maximizing sensory information**. *Biol. Cybern.*, 68:23–29, 1992.

[20] D J Field. **What is the goal of sensory coding?** *Neural Computation*, 6:559–601, 1994.

[21] F Rieke, D A Bodnar, and W Bialek. **Naturalistic stimuli increase the rate and efficiency of information transmission by primary auditory afferents**. *Proc. R. Soc. Lond. B*, 262:259–265, 1995.

[22] F Rieke, D Warland, R R de Ruyter van Steveninck, and W Bialek. **Spikes: Exploring the Neural Code**. MIT Press, Cambridge, MA, 1997.

[23] W E Vinje and J L Gallant. **Sparse coding and decorrelation in primary visual cortex during natural vision**. *Science*, 287:1273–1276, Feb 2000.

[24] • S Nirenberg, S M Carcieri, A L Jacobs, and P E Latham. **Retinal ganglion cells act largely as independent encoders**. *Nature*, 411:698–701, June 2001.
An important paper, as it is one of the first to attempt to directly quantify the redundancy between pairs of neurons under naturalistic

stimulation conditions. The redundancy measure used is unusual and is motivated with respect to decoding (as opposed to encoding). This has generated a lively controversy (see the discussion on Peter Latham's web page), which should be resolved through careful discussion and comparison of different alternative measures.

[25] • W E Vinje and J L Gallant. **Natural stimulation of the nonclassical receptive field increases information transmission efficiency in V1**. *J. Neurosci.*, 22(7):2904–2915, Apr 2002.
Shows that the information rate, information per spike, and transmission efficiency increase as the spatial extent of a naturalistic stimulus is increased to extend beyond the classical receptive field.

[26] •• G Chechik, A Globerson, M J Anderson, E D Young, I Nelken, and N Tishby. **Group redundancy measures reveal redundancy reduction in the auditory pathway**. In T G Dietterich, S Becker, and Z Ghahramani, editors, *Adv. Neural Information Processing Systems (NIPS*01)*, volume 14, pages 173–180, Cambridge, MA, 2002. MIT Press.
Provides a direct experimental demonstration of redundancy reduction as one ascends the cat auditory pathway. Specifically, they show an increase in information content and reduction in mutual information in the primary auditory cortex as compared with the inferior colliculus during presentation of bird vocalization stimuli.

[27] S B Laughlin. **A simple coding procedure enhances a neuron's information capacity**. *Z. Naturforsch.*, 36c:910–912, 1981.

[28] Y Dan, J J Atick, and R C Reid. **Efficient coding of natural scenes in the lateral geniculate nucleus: Experimental test of a computational theory**. *J. Neuroscience*, 16:3351–3362, 1996.

[29] R Baddeley, L F Abbott, M C Booth, F Sengpiel, T Freeman, E A Wakeman, and E T Rolls. **Respones of neurons in primary and inferior temporal visual cortices to natural scenes**. *Proc. Roy. Soc. (Lond.)*, B264:1775–1783, 1998.

[30] R Linsker. **Self-organization in a perceptual network**. *Computer*, 21:105–117, 1988.

[31] H B Barlow, T P Kaushal, and G J Mitchison. **Finding minimum entropy codes**. *Neural Computation*, 1:412–423, 1989.

[32] T D Sanger. **Optimal unsupervised learning in a single-layer network**. *Neural Networks*, 2:459–473, 1989.

[33] P Foldiak. **Forming sparse representations by local anti-hebbian learning**. *Biol. Cybernetics*, 64:165–170, 1990.

[34] J J Atick. **Could information theory provide an ecological theory of sensory processing?** *Network: Computation in Neural Systems*, 3:213–251, 1992.

[35] B A Olshausen and D J Field. **Emergence of simple-cell receptive field properties by learning a sparse code for natural images**. *Nature*, 381:607–609, 1996.

[36] A J Bell and T J Sejnowski. **The 'independent components' of natural scenes are edge filters**. *Vision Research*, 37(23):3327–3338, 1997.

[37] J H van Hateren and D L Ruderman. **Independent component analysis of natural image sequences yields spatio-temporal filters similar to simple cells in primary visual cortex**. *Proc R Soc Lond B*, 265:2315–2320, 1998.

[38] E P Simoncelli and O Schwartz. **Image statistics and cortical normalization models**. In M. S. Kearns, S. A. Solla, and D. A. Cohn, editors, *Adv. Neural Information Processing Systems*, volume 11, pages 153–159, Cambridge, MA, 1999. MIT Press.

[39] P Reinagel. **How do visual neurons respond in the real world?** *Current Opinion*

in Neurobiology, 11:437–442, 2001.
A good overview of recent physiological experiments that utilize naturalistic stimuli.

[40] A Treves, S Panzeri, E T Rolls, M Booth, and E A Wakeman. **Firing rate distributions and efficiency of information transmission of inferior temporal cortex neurons to natural visual stimulation**. *Neural Computation*, 11:601–631, 1999.

[41] G G de Polavieja. **Errors drive the evolution of biological signaling to costly codes**. *J. Theoretical Biology*, 214:657–664, 2002.

[42] M Meister, L Lagnado, and D A Baylor. **Concerted signaling by retinal ganglion cells**. *Science*, 270(5239):1207–1210, Nov 1995.

[43]• D S Reich, F Mechler, and J D Victor. **Independent and redundant information of nearby cortical neurons**. *Science*, 294, Dec 2001.
Shows that responses of clusters of V1 neurons to white noise stimuli indicate a fairly low level of redundancy, for cluster of up to six neurons. Although this does not directly bear on the efficient coding hypothesis (because of the choice of stimuli), it seems surprising, given the results of Rieke95.

[44]• M C Wiener, M W Oram, Z Liu, and B J Richmond. **Consistency of encoding in monkey visual cortex**. *J. Neurosci.*, 21(20):8210–8221, October 2001.
The authors demonstrate that neurons in areas V1 transmit roughly equal amounts of information about a variety of stimuli (including bars, gratings, Walsh patterns, and photographic images). Their information measure is based on a spike count code taken over 400 msec intervals, and an assumption that spike count variability followed a Gaussian xo distribution.

[45]• G Lewen, W Bialek, and R de Ruyter van Steveninck. **Neural coding of natural motion stimuli**. *Network*, 12:317–329, 2001.
This technically difficult study of the motion-sensitive H1 neuron of the blowfly examines the information content in the neural response as a function of light level during naturalistic motion (achieved by rotating a fly on a high speed stepper motor. The authors show that the neural information about the motion continues does not saturate over a wide range of light levels.

[46] R Kern, C Petereit, and M Egelhaaf. **Neural processing of naturalistic optic flow**. *J. Neurosci*, 21(RC139), 2001.

[47] D Ringach, M Hawken, and R Shapley. **Receptive field structure of neurons in monkey primary visual cortex revealed by stimulation with natural image sequences**. *Journal of Vision*, 2:12–24, 2002.

[48] T Sharpee, N C Rust, and W Bialek. **Maximizing informative dimensions: Analyzing neural responses to natural signals**. In S Becker, S Thrun, and K Obermayer, editors, *Adv. Neural Information Processing Systems (NIPS*02)*, volume 15, Cambridge, MA, 2003. MIT Press.

[49] L Paninski. **Convergence properties of spike-triggered analysis techniques**. In S Becker, S Thrun, and K Obermayer, editors, *Adv. Neural Information Processing Systems (NIPS*02)*, volume 15, Cambridge, MA, 2003. MIT Press.

[50]• O Schwartz and E P Simoncelli. **Natural signal statistics and sensory gain control**. *Nature Neuroscience*, 4(8):819–825, August 2001.
Demonstrates that natural images and sounds exhibit higher order statistical dependencies that may be reduced using divisive normalization, a form of local contrast gain control. The paper also shows that when the parameters of such a model are optimized to reduce redundancy, the resulting model provides a good account of a number of nonlinear properties of early visual and auditory neurons.

[51] J H van Hateren and H P Snippe. **Information theoretical evaluation of parametric models of gain control in blowfly photoreceptor cells**. *Vision Research*, 41:1851–1865, 2001.
Derives a nearly optimal gain control model for blowfly photoreceptors based on information-theoretic measures in the context of naturalistic stimuli. It remains to be seen whether the resulting model is quantitatively consistent with physiological behaviors.

[52]• C Zetzsche and F Röhrbein. **Nonlinear and extra-classical receptive field properties and the statistics of natural scenes**. *Network: Computation in Neural Systems*, 12(3):331–350, 2001.
Examines a variety of higher order statistical properties of images, and their relationship to cortical models. In particular, the authors argue that gain control models and cascades of linear operators and nonlinear scalar operations are well-suited for natural image structures, as well as being consistent with nonlinear properties of cortical neurons.

[53]• P Hoyer and A Hyvärinen. **A multi-layer sparse coding network learns contour coding from natural images**. *Vision Research*, 42(12):1593–1605, 2002.
The authors use a sparse-coding model to derive an optimal basis set for processing the responses of model complex cells to natural images. The resulting basis encodes extended contours.

[54] W B Levy and R A Baxter. **Energy efficient neural codes**. *Neural Computation*, 8:531–543, 1996.

[55] S B Laughlin, R de Ruyter van Steveninck, and J C Anderson. **The metabolic cost of information**. *Nature neuroscience*, 1(1):36–41, May 1998.

[56] V Balasubramanian and M J Berry. **A test of metabolically efficient coding in the retina**. *Network: Comput. Neural Syst.*, 13:531–552, 2002.

[57] A Gottschalk. **Derivation of the visual contrast response function by maximizing information rate**. *Neural Computation*, 14:527–542, 2002.

[58] Horace B Barlow and P Foldiak. **Adaptation and decorrelation in the cortex**. In R Durbin, C Miall, and G Mitchinson, editors, *The Computing Neuron*, chapter 4, pages 54–72. Addison-Wellesley, New York, 1989.

[59] H B Barlow. **A theory about the functional role and synaptic mechanism of visual aftereffects**. In C Blakemore, editor, *Vision: Coding and Efficiency*. Cambridge University Press, 1990.

[60] M J Wainwright. **Visual adaptation as optimal information transmission**. *Vision Research*, 39:3960–3974, 1999.

[61] N Brenner, W Bialek, and R R de Ruyter van Steveninck. **Adaptive rescaling maximizes information transmission**. *Neuron*, 26:695–702, June 2000.

[62]• A Fairhall, G Lewen, W Bialek, and R de Ruyter van Steveninck. **Efficiency and ambiguity in an adaptive neural code**. *Nature*, 412:787–792, Aug 2001.
Elaborates on previous work examining changes in neural response due to changes in contrast (variance) of white noise stimuli. The paper shows that the dynamical properties of adaptation appear to be optimized for representation of both the statistics of the stimulus ensemble as well as the specific attributes of the most recently presented stimlus.

[63] A Hyvärinen and P Hoyer. **Emergence of topography and complex cell properties from natural images using extensions of ICA**. In S. A. Solla, T. K. Leen, and K.-R. Müller, editors, *Adv. Neural Information Processing Systems*, volume 12, pages 827–833, Cambridge, MA, May 2000. MIT Press.

[64] C Kayser, W Einhäuser, O Dümmer, P König, and K P Körding. **Extracting slow subspaces from natural videos leads to**

complex cells. In G Dorffner, H Bischof, and K Hornik, editors, *Proc. Int'l Conf. on Artificial Neural Networks (ICANN-01)*, pages 1075–1080, Vienna, Aug 2001. Springer-Verlag, Heidelberg.

[65] P Berkes and L Wiskott. **Applying slow feature analysis to image sequences yields a rich repertoire of complex cell properties**. In J R Dorronsoro, editor, *Proc. Int'l Conf. on Artificial Neural Networks (ICANN-02)*, pages 81–86, Madrid, Aug 2002.

[66] R M Balboa and N M Grzywacz. **The role of early lateral inhibition: More than maximizing luminance information**. *Vision Research*, 17:77–89, 2000.

[67] T von der Twer and D I A MacLeod. **Optimal nonlinear codes for the perception of natural colours**. *Network*, 12:395–407, 2001.