

Minireview

Visual adaptation: Neural, psychological and computational aspects

Colin W.G. Clifford ^{a,*}, Michael A. Webster ^b, Garrett B. Stanley ^c, Alan A. Stocker ^d,
Adam Kohn ^e, Tatyana O. Sharpee ^f, Odelia Schwartz ^g

^a School of Psychology, University of Sydney, Sydney, NSW 2006, Australia

^b Department of Psychology, University of Nevada, Reno, NV 89557, USA

^c Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA

^d Center for Neural Science, New York University, New York, NY 10003, USA

^e Department of Neuroscience, Albert Einstein College of Medicine, Bronx, NY 10461, USA

^f Crick-Jacobs Center for Theoretical Biology and Computational Neurobiology Lab, The Salk Institute for Biological Studies, La Jolla, CA 92037, USA

^g Computational Neurobiology Lab, The Salk Institute, La Jolla, CA 92037, USA

Received 13 March 2007; received in revised form 31 July 2007

Abstract

The term *visual adaptation* describes the processes by which the visual system alters its operating properties in response to changes in the environment. These continual adjustments in sensory processing are diagnostic as to the computational principles underlying the neural coding of information and can have profound consequences for our perceptual experience. New physiological and psychophysical data, along with emerging statistical and computational models, make this an opportune time to bring together experimental and theoretical perspectives. Here, we discuss functional ideas about adaptation in the light of recent data and identify exciting directions for future research.

© 2007 Elsevier Ltd. All rights reserved.

Keywords: Sensory coding; Physiology; Theoretical neuroscience; Perception; Information processing; Psychophysics

1. Encoding aspects of adaptation

Our sensory systems encode information about the environment as patterns of activity across populations of neurons. An appealing hypothesis states that the response properties of neurons in sensory areas of the brain are matched to the signals to which they are exposed (e.g. Attneave, 1954; Barlow, 1961; Chechik et al., 2006; Fairhall, Lewen, Bialek, & de Ruyter Van Steveninck, 2001). According to this hypothesis, adaptation serves to maintain that match across changes in the diet of sensory stimulation. Understanding the link between sensory processing and environmental stimuli requires a better grasp of the statistical structure of environmental stimuli.

This area has progressed rapidly over the past decade, largely due to stronger computing power and advances in signal processing methodologies (Simoncelli & Olshausen, 2001).

Images of the natural world are highly structured over both space and time. For instance, the intensity values of neighboring spatial regions of an image tend to be similar (Field, 1987; Ruderman & Bialek, 1994; Simoncelli & Olshausen, 2001). This means that information about any given point in the image is contained in the recent history of the image structure at that point (temporal context) and the structure of the surrounding image (spatial context). Adaptation and surround modulation at the neuronal level might then exploit these regularities in image structure to optimize sensory coding (Schwartz, Hsu, & Dayan, 2007). The statistical analysis of natural images has been used to derive neural models using principles of efficient coding of the input signals such as redundancy

* Corresponding author. Fax +61 (0) 2 9351 2603.

E-mail address: colinc@psych.usyd.edu.au (C.W.G. Clifford).

reduction so that the outputs of neurons in the model tend towards statistical independence (e.g. Li & Atick, 1994; Olshausen & Field, 1996). Image statistics approaches have led to models of contextual modulation that can be directly tested against neuronal data from visual cortex (e.g. Rao & Ballard, 1999; Schwartz & Simoncelli, 2001). While such approaches have largely been applied to neuronal data on simultaneous spatial context, there have been only preliminary studies on adaptation (Buiatti & van Vreeswijk, 2003; Wainwright, Schwartz, & Simoncelli, 2002). There thus remains extensive scope to consider models of adaptation informed by the statistical structure of natural image sequences.

2. Decoding aspects of adaptation

To support decisions or behaviours, the information encoded by our sensory systems must be decoded into an appropriate form. Increasing evidence suggests that humans can combine stored knowledge and noisy sensory observations in an optimal fashion when performing a perceptual task. The Bayesian formulation of this problem defines the optimal strategy and provides a principled yet simple computational framework for perception. The Bayesian approach can account for a large range of existing perceptual data (see Knill & Richards, 1996, for a collection of early work) and has also found resonance in image statistics frameworks (e.g. Geisler, Perry, Super, & Gallogly, 2001).

In a recent study of visual motion perception, Stocker and Simoncelli (2006a) showed that human performance on a *trial-by-trial* basis is well explained by a Bayesian observer model that is optimal with respect to uncertainty reflected in the encoded sensory signals, where uncertainty includes external (stimulus) as well as internal (neural) noise. If the Bayesian observer model provides a valid and general computational explanation of perception then it ought to account for the known aftereffects of adaptation. Adaptation results in two fundamental changes on subsequent perception of suprathreshold stimuli (Clifford, 2002). Firstly and most saliently, it leads to a perceptual bias (Levinson & Sekuler, 1976) where the estimate of the stimulus parameter is repelled from the value of the adaptor (Fig. 1). Secondly, it changes the ability of the observer to discriminate stimulus parameters, where discrimination thresholds typically decrease for parameter values very similar to the adaptor, yet increase at values rather different from the adaptor (e.g. Phinney, Bowd, & Patterson, 1997).

An intuitive hypothesis is that adaptation changes the prior distribution in the Bayesian observer model. The prior is meant to reflect the distribution of values of the perceptual variable in the world. If the prior is estimated from the recent history of values sampled by the organism then repeated presentation of the same value should increase the prior probability in its vicinity. A qualitative examination, however, reveals such a change in the prior distribution would lead to a perceptual bias that is opposite




Fig. 1. Schematic of the repulsive direction aftereffect (Levinson and Sekuler, 1976). (a) Prior to adaptation, the direction of a pattern of dots moving vertically upwards is perceived veridically. (b) Adaptation is induced by prolonged exposure to a pattern of dots moving in an oblique direction. (c) Following adaptation, the perceived direction of a pattern of dots moving vertically upwards is repelled away from the adapting direction of motion.

to the repulsive aftereffect (Fig. 2). Rather, a sensible change in the likelihood function within the Bayesian observer model can account both for aftereffects and for changes in discrimination threshold induced by adaptation (Stocker & Simoncelli, 2006b). The likelihood function changes according to an increase in signal-to-noise ratio of the sensory signal in the vicinity of the adaptor value. Such an increase in sensory signal quality seems in qualitative agreement with ideas of efficient coding. This raises the promise that adaptation may be understood on the basis of a single, principled computational framework combining bottom-up efficient sensory encoding and top-down Bayesian decoding.

3. Adaptation at the single cell level

At the neuronal level, various forms of adaptation have been identified at multiple stages of processing in early sensory pathways. For example, biophysically distinct mechanisms of adaptation to the mean luminance and contrast of




Fig. 2. Hypothetical Bayesian model in which adaptation changes the prior distribution. Top: a Bayesian estimator optimally combines sensory evidence (given by the likelihood function—grey curve) with stored knowledge (represented as the prior probability distribution—bold curve) by multiplying both functions to obtain a posterior probability distribution. A sensible choice for the estimate $\hat{\theta}_{pre}$ is e.g. to select the θ with highest probability. Note, that for reasons of readability, we chose the prior distribution to be flat initially, in which case the posterior distribution becomes identical with the likelihood function (up to a scale factor). Bottom: Assuming that adaptation leads to an increase of the prior distribution around the adaptor θ_{adapt} , the posterior, and thus the estimate $\hat{\theta}_{post}$, would be shifted toward the adaptor. Thus, changing the prior distribution results in an *attractive shift* (red arrow) of the estimate toward the adaptor rather than the experimentally reported repulsive aftereffects (e.g. Levinson and Sekuler, 1976). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

sinusoidal visual stimuli have been shown to exist in the retina (Chander & Chichilnisky, 2001; Enroth-Cugell & Shapley, 1973; Hosoya, Baccus, & Meister, 2005; Shapley & Victor, 1978; Smirnakis, Berry, Warland, Bialek, & Meister, 1997), and similar functional effects have subsequently been observed in the thalamus and cortex. Given that sensory neurons are capable of diverse forms of adaptation, it is intriguing to hypothesize that neurons can adapt, at least partially, to any statistical parameter.

Sensory neurons are commonly characterized using the linear–nonlinear model (Rieke, Warland, de Ruyter van Steveninck, & Bialek, 1997). In this framework, the neural response for a given stimulus depends on linear filtering of the stimulus with a receptive field defined over some region of visual space and time. The nonlinear aspects of neural response are modeled by a function that relates the filtered stimulus values to response probability. Both components of the linear–nonlinear model—the linear filter and the nonlinear gain function—can be affected by adaptation (Marav-

all, Petersen, Fairhall, Arabzadeh, & Diamond, 2007; Nagel & Doupe, 2006). Adaptation to different statistical parameters of the stimulus can affect different aspects of neural responses. Neurons can accommodate changes in the mean luminance and contrast by adjusting the nonlinear function (Chander & Chichilnisky, 2001; Rieke, 2001). For example, if mean luminance is increased, neurons will typically adjust their threshold for firing to maintain their average firing rate. However a change in the distribution of image contrast across different spatial frequencies requires adjustments at the filtering stage if neurons are to maintain the optimality of coding (Smirnakis et al., 1997).

What could be the goal of such adaptive changes in neural coding? One theoretical idea is that adaptive behaviors at the early stages of processing serve to maximize the efficiency of encoding and, in particular, the amount of information about the stimulus conveyed by spike times (Brenner, Bialek, & de Ruyter van Steveninck, 2000; Rieke, Bodnar, & Bialek, 1995). If neurons performed only linear transformation of incoming signals, then knowledge of input statistical properties would permit exact description of the filtering operations that would maximize the amount of information transmission. Specifically, the transfer function of the optimal filter for a linear system is inversely proportional to the amplitude spectrum of the stimulus distribution up to some cut-off frequency beyond which neural noise dominates (Atick & Redlich, 1990, 1992; Barlow, 1990; Linsker, 1989; Srinivasan, Laughlin, & Dubs, 1982).

It is difficult however to generalize this argument to apply to a linear–nonlinear system. Instead, one can make predictions for how filtering properties should change after a change in stimulus statistics in order to maintain coding efficiency. Suppose the parameters of the linear–nonlinear system are optimally suited for encoding a stimulus distribution with some amplitude spectrum. Any change in the stimulus distribution could then be compensated for by adjusting neural filtering such that the *filtered* stimulus distribution remained unchanged. If such adjustments in filtering could be achieved then there would be no need for any adjustments at the nonlinear stage. Strictly speaking, this argument only applies to stimuli taken from a correlated Gaussian distribution. Despite this theoretical limitation, simple cells in the primary visual cortex (V1) appear to follow such an optimization principle in their response to low spatial frequencies (Sharpee et al., 2006). When probed with white noise and natural stimuli, neurons adjust their filtering of low spatial frequencies in such a way as to keep the spectral content of the filtered stimulus distribution constant. However, at mid-to-high spatial frequencies there are no changes in filtering so the filtered stimulus distribution covaries significantly with the stimulus distribution itself (Fig. 3).

Although the time scale of adaptation reported by Sharpee et al. (2006) is of the order of tens of seconds, this does not, of course, exclude the possibility of adaptation at shorter time scales. Several other forms of adaptation to changes in stimulus statistics appear to be very rapid,

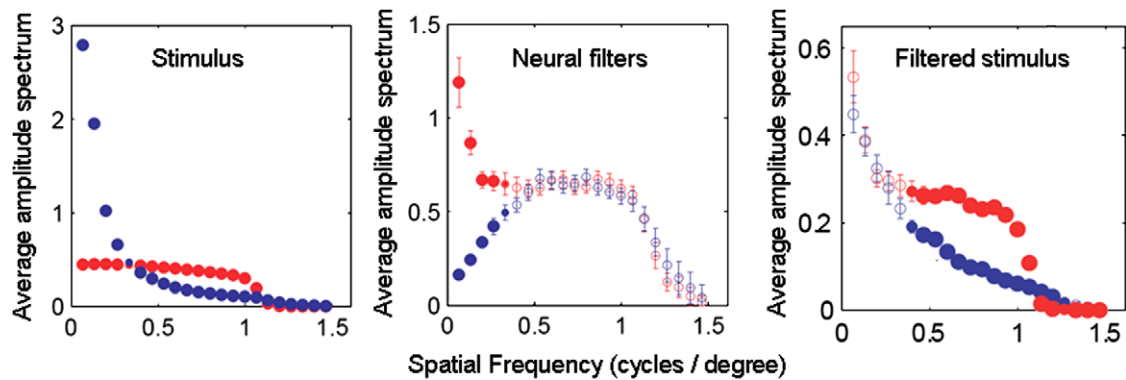


Fig. 3. Optimal adaptive filtering by simple cells in the primary visual cortex. Spectral content of natural (blue) and noise (red) stimuli is shown in the left column. Average spectral content of neural filters from the linear–nonlinear model is shown in the middle column; blue (red) when derived from natural (noise) stimuli. Neural filters were computed for 40 simple cells, using the method of maximally informative dimensions in order to compensate for effects of non-Gaussian correlations in the input statistics (Sharpee et al., 2006). The rightmost column shows the spectral content of stimuli after filtering. At low spatial frequencies, where neural filtering is stimulus-dependent, the spectral content of stimuli after filtering is the same. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

occurring nearly instantaneously and thus introducing the possibility of a static nonlinear mechanism rather than an active and ongoing process conventionally associated with adaptation (Brenner et al., 2000; Baccus & Meister, 2002; Fairhall et al., 2001; Maravall et al., 2007; McCormick, Connors, Lighthall, & Prince, 1985; Nagel & Doupe, 2006; Nowak, Azouz, Sanchez-Vives, Gray, & McCormick, 2003). A purely static nonlinear property would not normally be regarded as adaptation rather than, say, dynamic gain control. However, it is likely that different elements of adaptation operate on different time scales, reflecting a continuum from near instantaneous gain control (Chander & Chichilnisky, 2001; Fairhall et al., 2001; Nagel & Doupe, 2006; Nowak et al., 2003) to slowly developing modulations in encoding properties (Baccus & Meister, 2002; Hosoya et al., 2005; Sharpee et al., 2006).

4. Adaptation within a neural processing hierarchy

Adaptation has been investigated most extensively in the early visual system, but several recent studies have explored the adaptation of cortical areas linked more closely to perceptual experience. One such study, by Kohn and Movshon (2003), measured the spatial specificity of contrast adaptation in area MT, an extrastriate area in the macaque visual system important for motion processing (Dubner & Zeki, 1971). They found that adaptation could alter contrast sensitivity in one subregion of an MT receptive field without affecting sensitivity elsewhere, consistent with contrast adaptation occurring early in the processing stream and simply being inherited by MT. Subsequently, Kohn and Movshon (2004) reported that adaptation alters direction tuning in MT in a manner opposite to previously described effects on tuning in V1 (Dragoi, Rivadulla, & Sur, 2001; Dragoi, Sharma, & Sur, 2000; Dragoi, Sharma, Miller, & Sur, 2002). Specifically, adaptation on the flank of an MT tuning curve was found to cause the tuning curve to shift toward the adaptor, rather than to be repelled from

it as in V1 (see also Krekelberg, van Wezel, & Albright, 2006). With a simple population decoding model, these attractive shifts in tuning can explain the perceptual repulsion that follows adaptation. Such shifts also lead to a relative enhancement of the representation of frequently occurring stimuli, a potential neural basis for the changes in likelihood function required by Bayesian explanations for perceptual effects (Stocker & Simoncelli, 2006b).

To understand adaptation at higher levels of a sensory stream, it is important to know how neurons at each stage of processing adapt and how plasticity occurring early in the processing stream impacts downstream areas. In macaque IT cortex, for example, neuronal adaptation has been reported to show greater selectivity than the response of the corresponding neuron, suggesting that adaptation is occurring at or before the level of synapses onto the neuron (Sawamura, Orban, & Vogels, 2006). The effects observed by Kohn and Movshon (2003, 2004) can also be attributed to adaptation early in the visual system, with the novel changes in MT tuning arising from an interaction between adapted feedforward input and unaltered recurrent excitatory and inhibitory connections within MT. The complex nature of some perceptual aftereffects (see below) suggests, however, that higher levels of the visual system also adapt. Further work is needed to understand the relative importance of low-level vs. high-level effects. Tracing the effects of adaptation through the sensory system will undoubtedly also elucidate the functional organization of the visual system more generally.

Remarkably, adaptation to stimulus properties as sophisticated as orientation, spatio-temporal correlation (Hosoya et al., 2005) and relative motion can be observed as early as the ganglion cells of the retina even though neuronal selectivity for these properties is typically not observed outside visual cortex. These findings go against the long held view that pattern adaptation is evidence of a feature detecting neuron selective for that pattern (Mollon, 1977), suggesting that caution is required in interpret-

ing the results of psychophysical studies of pattern adaptation. Instead, Hosoya et al. (2005) demonstrate that a recurrent network of model amacrine cells can implement predictive coding (Srinivasan et al., 1982) dynamically through synaptic plasticity, such that the stimulus properties to which the network adapts are modified “on the fly” in response to the changing environment.

5. Perceptual consequences of adaptation

Behavioural studies with human observers have recently reported perceptual aftereffects to higher-level stimulus properties such as the identity, gender, ethnicity, emotion and attractiveness of faces (e.g. Rhodes, Jeffery, Watson, Clifford, & Nakayama, 2003; Webster, Kaping, Mizokami, & Duhamel, 2004). The robustness of these effects across variations in size, position and orientation between the adapting and test faces indicates that they are generated at a high-level of the visual processing hierarchy rather than inherited from adaptation at earlier stages of the visual pathway (e.g. Watson & Clifford, 2003). The existence of such high-level aftereffects suggests that adaptive coding is not only a property of the early stages of sensory processing but is instead a strategy employed throughout the perceptual processing hierarchy.

The dramatic nature of face aftereffects (Fig. 4) suggests that one of the primary functions of adaptation may be to regulate how things look (Webster, Werner, & Field, 2005). One way this is important is to provide perceptual constancy. Adapting to extraneous variations in the stimulus (e.g. a change in the color signal when the lighting changes) helps allow the visual system to maintain a stable perception of object properties (e.g. the surface reflectance). This process may also be important for maintaining constancy despite variations in the observer as sensitivity changes during the lifespan, or even for different points in the visual field. Thus, the object that appears white or in focus remains largely the same in old or young observers, and in the fovea or periphery, despite large differences in spectral sensitivity or spatial resolution.

In addition, this constancy may be important for synchronizing perception among observers by adapting them

to a common environment. Qualities like “white” or “in focus” represent perceptual norms. These norms play a special role in anchoring many perceptions, from color to faces, and adaptation may anchor these by normalizing visual coding relative to the stimulus distributions to which the observer is currently exposed, such as the average color or the average face (Webster et al., 2004). These norms also represent a prediction about the world, allowing higher-level mechanisms the same economy of explicitly coding only the “errors” in the predictions as is evident in the retina (Hosoya et al., 2005; Smirnakis et al., 1997; Srinivasan et al., 1982) and primary visual cortex (Sharpee et al., 2006). By discounting expected properties, a further functional consequence of adaptation may be to increase the salience of novel stimuli (Barlow, 1990). These novel stimuli are the very features of the world to which we are not currently adapted, so much of what we notice about the world arguably amounts to a visual ‘aftereffect’ (Webster et al., 2005). For example, changes in motion sensitivity are largely imperceptible while adapting to a waterfall but have striking consequences when we shift our gaze to the static surround. What is specifically striking is the aftereffect—how the scene differs from the adapting pattern. In the same way, at any moment we are under adaptation to certain recurring properties of the ambient environment. The ‘aftereffect’ of this adaptation may be that novel properties of the environment appear the most conspicuous.

6. Summary and future directions

Consideration of efficiency suggests that the goal of sensory coding and the goal of adaptation are intimately related: to match the response properties of our sensory systems to the prevailing environment. Establishing a computational theory of sensory adaptation then entails determining which properties of the response distribution tend to remain invariant across different stimulus distributions (e.g. Sharpee, Rust, & Bialek, 2004) and establishing the significance of potential constraints such as the metabolic cost of neuronal activity (Laughlin, de Ruyter van Steveninck, & Anderson, 1998; Lennie, 2003).




Fig. 4. Adapting to a particular expression (e.g. a happy face) induces an after-“affect” in a neutral face (e.g. so that it appears more angry).

The natural environment has shaped the evolution and development of our visual systems. Recent research has shown that neuronal responses to naturalistic stimuli cannot necessarily be inferred from the responses to conventional ‘laboratory’ stimuli such as sine wave gratings (Mante, Frazor, Bonin, Geisler, & Carandini, 2005; Sharpee et al., 2006). The difficulty of generalizing from simple to complex stimuli underscores the need for a better understanding of the statistical structure of the environment. Such an understanding will enable the development of artificial stimuli more closely matched to the statistical properties of naturalistic stimuli but whose properties can be controlled systematically. The use of naturalistic stimulus ensembles in turn requires new approaches for characterizing neuronal nonlinearities (e.g. Sharpee et al., 2004) if it is to enhance our understanding of the computational basis of adaptation.

As well as an appreciation of the statistics of natural images, a complete model of sensory adaptation will require detailed consideration of the uses to which the encoded sensory information is to be put. The search for a common computational framework in which to consider the adaptive encoding of sensory information and the corresponding decoding operations is of major theoretical significance. The integration of efficient sensory encoding with Bayesian decoding (Stocker & Simoncelli, 2006b) is an exciting innovation in this regard. Future work should investigate the implementation of such a computational framework in a biologically realizable architecture (e.g. Rao, 2004) and explore the behaviour of such networks under adaptation (e.g. Teich & Qian, 2003). Proposals of the functional benefit provided by altered neuronal responsiveness and tuning will need to consider how these changes affect downstream networks if any theory is to encompass the multiple levels of the visual hierarchy and the range of time scales over which adaptation occurs.

Finally, better characterization of the behavioural effects of adaptation is required if psychophysical data are to prove decisive between models. In domains such as orientation and motion perception, attractive aftereffects are observed under some conditions instead of repulsion. The ability to account for these phenomena may prove diagnostic of the underlying mechanisms (Clifford, Wenderoth, & Spehar, 2000). Application of techniques from signal detection theory such as equivalent noise analysis (Dao, Lu, & Doshier, 2006; Pelli & Farell, 1999) might also be profitably applied to tease apart the effects of adaptation on various determinants of psychophysical discrimination performance such as observer’s level of internal noise and sampling efficiency.

Acknowledgments

We are grateful to Peter Dayan for his role in organizing the workshop at Computational and Systems Neuroscience 2006 that allowed us to bring all the authors together and for helpful discussion.

References

- Atick, J. J., & Redlich, A. N. (1990). Towards a theory of early visual processing. *Neural Computation*, 2, 308–320.
- Atick, J. J., & Redlich, A. N. (1992). What does the retina know about natural scenes? *Neural Computation*, 4, 196–210.
- Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, 61, 183–193.
- Baccus, S. A., & Meister, M. (2002). Fast and slow contrast adaptation in retinal circuitry. *Neuron*, 36, 909–919.
- Barlow, H. B. (1961). Possible principles underlying the transformation of sensory messages. In W. A. Rosenblith (Ed.), *Sensory communication* (pp. 217–234). Cambridge, MA: MIT.
- Barlow, H. (1990). A theory about the functional role and synaptic mechanism of visual aftereffects. In C. Blakemore (Ed.), *Vision: Coding and efficiency*. Cambridge, UK: Cambridge University Press.
- Brenner, N., Bialek, W., & de Ruyter van Steveninck, R. (2000). Adaptive rescaling maximizes information transmission. *Neuron*, 26, 695–702.
- Buiatti, M., & van Vreeswijk, C. A. (2003). Variance renormalization: A key mechanism for temporal adaptation in neural vision? *Vision Research*, 43, 1895–1906.
- Chander, D., & Chichilnisky, E. J. (2001). Adaptation to temporal contrast in primate and salamander retina. *Journal of Neuroscience*, 21, 9904–9916.
- Chechik, G., Anderson, M. J., Bar-Yosef, O., Young, E. D., Tishby, N., & Nelken, I. (2006). Reduction of information redundancy in the ascending auditory pathway. *Neuron*, 51, 359–368.
- Clifford, C. W. G. (2002). Perceptual adaptation: Motion parallels orientation. *Trends in Cognitive Sciences*, 6, 136–143.
- Clifford, C. W. G., Wenderoth, P., & Spehar, B. (2000). A functional angle on adaptation in cortical vision. *Proceedings of the Royal Society of London Series B—Biological Science*, 267, 1705–1710.
- Dao, D. Y., Lu, Z. L., & Doshier, B. A. (2006). Adaptation to sine-wave gratings selectively reduces the contrast gain of the adapted stimuli. *Journal of Vision*, 6, 739–759.
- Dragoi, V., Sharma, J., & Sur, M. (2000). Adaptation-induced plasticity of orientation tuning in adult visual cortex. *Neuron*, 28, 287–298.
- Dragoi, V., Rivadulla, C., & Sur, M. (2001). Foci of orientation plasticity in visual cortex. *Nature*, 411, 80–86.
- Dragoi, V., Sharma, J., Miller, E. K., & Sur, M. (2002). Dynamics of neuronal sensitivity in visual cortex and local feature discrimination. *Nature Neuroscience*, 5, 883–891.
- Dubner, R., & Zeki, S. M. (1971). Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey. *Brain Research*, 35, 528–532.
- Enroth-Cugell, C., & Shapley, R. M. (1973). Adaptation and dynamics of cat retinal ganglion cells. *Journal of Physiology*, 233, 271–309.
- Fairhall, A. L., Lewen, G. D., Bialek, W., & de Ruyter Van Steveninck, R. R. (2001). Efficiency and ambiguity in an adaptive neural code. *Nature*, 412, 787–792.
- Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A*, 4, 2379–2394.
- Geisler, W. S., Perry, J. S., Super, B. J., & Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, 41, 711–724.
- Hosoya, T., Baccus, S. A., & Meister, M. (2005). Dynamic predictive coding by the retina. *Nature*, 436, 71–77.
- Knill, D. C., & Richards, W. (1996). *Perception as Bayesian inference*. Cambridge, UK: Cambridge University Press.
- Kohn, A., & Movshon, J. A. (2003). Neuronal adaptation to visual motion in area MT of the macaque. *Neuron*, 39, 681–691.
- Kohn, A., & Movshon, J. A. (2004). Adaptation changes the direction tuning of macaque MT neurons. *Nature Neuroscience*, 7, 764–772.
- Krekelberg, B., van Wezel, R. J., & Albright, T. D. (2006). Adaptation in macaque MT reduces perceived speed and improves speed discrimination. *Journal of Neurophysiology*, 95, 255–270.

- Laughlin, S. B., de Ruyter van Steveninck, R. R., & Anderson, J. C. (1998). The metabolic cost of neural information. *Nature Neuroscience*, *1*, 36–41.
- Lennie, P. (2003). The cost of cortical computation. *Current Biology*, *13*, 493–497.
- Levinson, E., & Sekuler, R. (1976). Adaptation alters perceived direction of motion. *Vision Research*, *16*, 779–781.
- Li, Z., & Atick, J. J. (1994). Towards a theory of striate cortex. *Neural Computation*, *6*, 127–146.
- Linsker, R. (1989). An application of the principle of maximum information preservation to linear systems. In D. S. Touretzky (Ed.), *Advance in neural information processing systems* (Vol. 1, pp. 86–194). Berlin: Springer.
- Mante, V., Frazor, R. A., Bonin, V., Geisler, W. S., & Carandini, M. (2005). Independence of luminance and contrast in natural scenes and in the early visual system. *Nature Neuroscience*, *8*, 1690–1697.
- Maravall, M., Petersen, R. S., Fairhall, A. L., Arabzadeh, E., & Diamond, M. E. (2007). Shifts in coding properties and maintenance of information transmission during adaptation in barrel cortex. *PLoS Biology*, *5*, e19.
- McCormick, D. A., Connors, B. W., Lighthall, J. W., & Prince, D. A. (1985). Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *Journal of Neurophysiology*, *54*, 782–806.
- Mollon, J. D. (1977). Neural analysis. In K. von Fienendt & I. K. Monstgaard (Eds.), *The perceptual world*. London: Academic Press.
- Nagel, K. I., & Doupe, A. J. (2006). Temporal processing and adaptation in the songbird auditory forebrain. *Neuron*, *51*, 845–859.
- Nowak, L. G., Azouz, R., Sanchez-Vives, M. V., Gray, C. M., & McCormick, D. A. (2003). Electrophysiological classes of cat primary visual cortical neurons in vivo as revealed by quantitative analyses. *Journal of Neurophysiology*, *89*, 1541–1566.
- Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse factorial code. *Nature*, *381*, 607–609.
- Pelli, D. G., & Farell, B. (1999). Why use noise? *Journal of the Optical Society of America A, Optics, Image Science, and Vision*, *16*, 647–653.
- Phinney, R. E., Bowd, C., & Patterson, R. (1997). Direction-selective coding of stereoscopic (Cyclopean) motion. *Vision Research*, *37*, 865–869.
- Rao, R. P. N. (2004). Bayesian computation in recurrent neural circuits. *Neural Computation*, *16*, 1–38.
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*, 79–87.
- Rhodes, G., Jeffery, L., Watson, T. L., Clifford, C. W. G., & Nakayama, K. (2003). Fitting the mind to the world: Face adaptation and attractiveness aftereffects. *Psychological Science*, *14*, 558–566.
- Rieke, F. (2001). Temporal contrast adaptation in salamander bipolar cells. *Journal of Neuroscience*, *21*, 9445–9454.
- Rieke, F., Bodnar, D. A., & Bialek, W. (1995). Naturalistic stimuli increase the rate and efficiency of information transmission by primary auditory afferents. *Proceedings of the Royal Society of London Series B—Biological Science*, *262*, 259–265.
- Rieke, F., Warland, D., de Ruyter van Steveninck, R., & Bialek, W. (1997). *Spikes: Exploring the neural code*. Cambridge, MA: MIT Press.
- Ruderman, D. L., & Bialek, W. (1994). Statistics of natural images: Scaling in the woods. *Physical Review Letters*, *73*, 814–817.
- Sawamura, H., Orban, G. A., & Vogels, R. (2006). Selectivity of neuronal adaptation does not match response selectivity: A single-cell study of the fMRI adaptation paradigm. *Neuron*, *49*, 307–318.
- Schwartz, O., & Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. *Nature Neuroscience*, *4*, 819–825.
- Schwartz, O., Hsu, A., & Dayan, P. (2007). Space and time in visual context. *Nature Reviews: Neuroscience*, *8*, 522–535.
- Shapley, R. M., & Victor, J. D. (1978). The effect of contrast on the transfer properties of cat retinal ganglion cells. *Journal of Physiology*, *285*, 275–298.
- Sharpee, T., Rust, N. C., & Bialek, W. (2004). Analyzing neural responses to natural signals: Maximally informative dimensions. *Neural Computation*, *16*, 223–250.
- Sharpee, T. O., Sugihara, H., Kurgansky, A. V., Rebrik, S. P., Stryker, M. P., & Miller, K. D. (2006). Adaptive filtering enhances information transmission in visual cortex. *Nature*, *439*, 936–942.
- Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual Review of Neurosciences*, *24*, 1193–1216.
- Smirnakis, S. M., Berry, M. J., Warland, D. K., Bialek, W., & Meister, M. (1997). Adaptation of retinal processing to image contrast and spatial scale. *Nature*, *386*, 69–73.
- Srinivasan, M. V., Laughlin, S. B., & Dubs, A. (1982). Predictive coding: A fresh view of inhibition in the retina. *Proceedings of the Royal Society of London Series B—Biological Science*, *216*, 427–459.
- Stocker, A. A., & Simoncelli, E. P. (2006a). Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience*, *9*, 578–585.
- Stocker, A. A., & Simoncelli, E. P. (2006b). Sensory adaptation within a Bayesian framework for perception. In Y. Weiss, B. Schoelkopf, & J. Platt (Eds.), *Advance in neural information processing systems* (Vol. 18, pp. 1291–1298). London: Springer.
- Teich, A. F., & Qian, N. (2003). Learning and adaptation in a recurrent model of V1 orientation selectivity. *Journal of Neurophysiology*, *9*, 2086–2100.
- Wainwright, M. J., Schwartz, O., & Simoncelli, E. P. (2002). Natural image statistics and divisive normalization: Modeling nonlinearity and adaptation in cortical neurons. In R. Rao, B. Olshausen, & M. Lewicki (Eds.), *Probabilistic models of the brain: Perception and neural function*. Cambridge, MA: MIT.
- Watson, T. L., & Clifford, C. W. G. (2003). Pulling faces: Investigating the face distortion aftereffect. *Perception*, *32*, 1109–1116.
- Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature*, *428*, 558–561.
- Webster, M. A., Werner, J. S., & Field, D. J. (2005). Adaptation and the phenomenology of perception. In C. W. G. Clifford & G. Rhodes (Eds.), *Fitting the mind to the world: Adaptation and aftereffects in high level vision* (pp. 241–277). Oxford: Oxford University Press.