

In praise of artifice

Nicole C Rust & J Anthony Movshon

The visual system evolved to process natural images, and the goal of visual neuroscience is to understand the computations it uses to do this. Indeed the goal of any theory of visual function is a model that will predict responses to any stimulus, including natural scenes. It has, however, recently become common to take this fundamental principle one step further: trying to use photographic or cinematographic representations of natural scenes (natural stimuli) as primary probes to explore visual computations. This approach is both challenging and controversial, and we argue that this use of natural images is so fraught with difficulty that it is not useful. Traditional methods for exploring visual computations that use artificial stimuli with carefully selected properties have been and continue to be the most effective tools for visual neuroscience. The proper use of natural stimuli is to test models based on responses to these synthetic stimuli, not to replace them.

There are two important commonly held fallacies that drive the new fashion of using natural stimuli in visual neuroscience experiments. The first is that in 40 years of experimentation and modeling, we have failed to capture important aspects of the behavior of neurons in primary visual cortex. The second is that an important reason for this 'failure' is that we have been trapped in the world of 'simplistic' artificial stimuli, which lack the richness of natural stimuli and have therefore prevented us from uncovering crucial facts of cortical organization¹.

Standard models for cortical cells based on synthetic stimuli

We know a considerable amount about the response properties of neurons in primary visual cortex (V1). These neurons show stimulus selectivity simultaneously for a diverse set of parameters, such as the location, size, form and color of an object. An important goal of visual neurophysiology is to produce models of these neurons that describe how this stimulus selectivity arises. Ultimately, one hopes to integrate all these models into a single theory that can predict neuronal and population responses to any arbitrary stimulus.

In the classical tradition, a visual physiologist observes the response of a neuron to a particular stimulus or class of stimuli and considers whether the current generation of models captures that behavior or whether these models require elaboration. In other words, the physiologist treats the current model as a hypothesis and designs experiments to

test it. Naturally, these models are based on measurements made with stimuli that are simple and easily parameterized, such as bars, points of light and sinusoidal gratings, and are then tested with stimuli of increasing complexity.

The history of research on neurons in V1 illustrates the success of this approach. V1 neurons simultaneously represent information both about the spatial structure of a stimulus and where it is located in the visual field (for a review see ref. 2). This multidimensional representation is based on selectivity for such stimulus features as the position, orientation, size, binocular disparity and color of stationary stimuli, as well as the direction and speed of moving stimuli. The first widely accepted formal models of these cells (referred to here as the 'old standard models') captured these tuning properties by passing an image through one or more linear spatiotemporal filters (Fig. 1a). In these models, stimulus selectivity arises from the shape of these filters: stimuli that resemble the filters produce high firing rates whereas stimuli that differ produce negligible firing rates. To capture this behavior, model spatial filters based on Gabor or related wavelet functions are typically used^{3,4}. Within V1, neurons that are sensitive to the position of a stimulus within their receptive field ('simple cells') are commonly modeled as a single linear filter whose output is half-rectified by the inevitable threshold nonlinearity shared by all spiking neurons⁵. The position insensitivity of 'complex cells' is commonly modeled with two phase-shifted filters whose outputs are squared and summed (the energy model^{6,7}) before being passed through the nonlinear spiking threshold. The responses of both kinds of neurons are irregular, and this variability can be reasonably approximated by a Poisson spiking process.

These old standard models predict not only the selectivity of V1 neurons' responses to bars, edges and gratings, but they also provide a credible account of responses to a variety of more complicated targets, including checkerboards⁸, random dot textures and Glass patterns⁹, and photographs of natural scenes¹⁰. This is not to say, however, that the old standard models are completely satisfactory. Research has uncovered a number of interesting ways in which they fail, resulting in the continuing evolution of a more elaborate and comprehensive 'new standard model'.

There are five important elements of V1 receptive fields not captured by the old standard model. First, although the old model postulated one input filter for simple cells and two for complex cells, more sensitive spike-triggered analysis has shown that additional filters are often required to account fully for the dimensionality of the stimulus set to which these cells respond^{11,12}. Furthermore, the simple linear or quadratic transformations of filter outputs in the old model may need to be replaced with more general point nonlinearities. Second, there are three more or less distinct gain control mechanisms that change responses depending on the combination of stimuli being presented. One of these regulates luminance gain when the average illumination

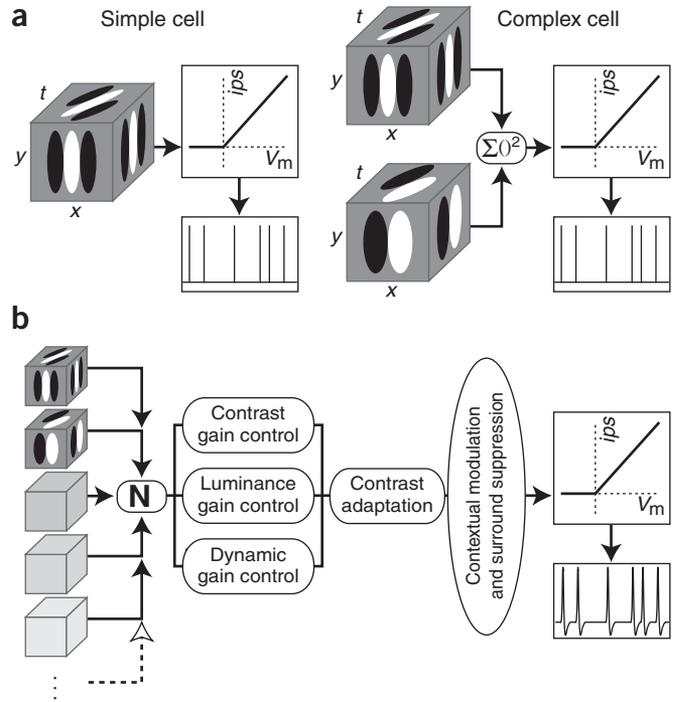
Nicole C. Rust is at the Howard Hughes Medical Institute and the Center for Neural Science, and J. Anthony Movshon is at the Center for Neural Science, New York University, New York, New York 10003, USA.
e-mail: rust@cns.nyu.edu

Published online 23 November 2005; doi:10.1038/nn1606

Figure 1 ‘Standard’ models of visual cortical cells, old and new. **(a)** The standard models of simple and complex cells in V1, circa 1985. The responses of simple cells are predicted by convolving the stimulus with a single linear spatiotemporal filter (represented by the cube; each face of the cube schematizes a section through the center of the filter), and passing the output through a spiking threshold nonlinearity obeying Poisson statistics that converts membrane voltage V_m into spikes at a rate ips (impulses per second). The responses of complex cells are predicted by summing the squared outputs of two linear spatiotemporal filters that have an approximate quadrature phase relationship in both space and time, and then passing the resulting signal through the same spiking nonlinearity. **(b)** The new standard model of visual cortical cells, circa 2005. The various deficiencies of the models in **a** have been corrected by allowing for multiple initial filters combined with an arbitrary nonlinearity N , and including several additional mechanisms: gain controls for luminance, contrast and temporal dynamics; contrast adaptation; and surround suppression and context effects. The output is passed through a spiking nonlinearity incorporating more realistic (e.g., Hodgkin-Huxley) dynamics. See text for details.

of the receptive field changes¹³. A second regulates contrast gain when the local average stimulus contrast varies^{14–17}. A third regulates the temporal dynamics of responses depending on the temporal character of concurrent stimulation¹⁸. These gain controls in combination capture such effects as the ability of an otherwise ineffective stimulus to reduce the response to an effective one¹⁴, as well as the saturation of signals and decreases in latency at high stimulus contrasts¹⁷. The three gain controls may not be completely independent of one another, and their basis in neural circuits is an active area of study. Third, it has long been known that V1 responses depend on the history of stimulation because of the phenomenon of contrast adaptation¹⁹. This important time-dependent nonlinearity seems to depend partly on a neuron’s history of activity^{20,21} and partly on changes in synaptic input^{22–24}, and it has a major influence on cortical responses. Fourth, the old standard model of a V1 neuron’s classical receptive field fails to deal with the existence of ‘hypercomplex’ cells²⁵ and, more generally, with the suppression of responses by stimuli presented in regions outside the classical receptive field²⁶. The addition of an inhibitory surround signal, originating in part from feedback signals from other cortical areas, extends the models to include this behavior^{27–29}. Fifth and finally, capturing the behavior of neurons on short time scales (<50 ms) requires that the old standard Poisson-spiking models be extended to include more realistic (e.g., Hodgkin-Huxley) spike generation^{30–32}. All of these elements can be combined to create a ‘new standard model’, schematically illustrated in **Figure 1b** as a synthesis and elaboration of the old standard models of **Figure 1a**.

All these extensions of the old standard model of V1 neurons were discovered and characterized using combinations of synthetic stimuli like bars and gratings; none of them was found using natural stimuli. Moreover, there is no case in which the response of V1 neurons to natural stimuli has been shown not to be captured by the new standard model. However, the components of this new standard model were for the most part discovered in relative isolation from one another, and an important challenge is to develop a set of measurement techniques to recover all the components of the new standard model for a single cell. Modern spike-triggered techniques are making a start on this problem (for example, see ref. 33), but they still fall short of achieving this goal. For the moment, then, the new standard model represents our accumulated understanding of the mechanisms in play in visual cortical processing, but it cannot be specified for individual cells. This creates a complex challenge for those who would use the new standard model to predict responses to complex stimuli, like natural images, that engage most or all of the mechanisms diagrammed in **Figure 1b**.



Uses and abuses of natural stimuli

There are two fundamentally different approaches that use natural stimuli. The principle behind the first approach is that one can deduce the properties of brain mechanisms for visual coding by reverse engineering: start with a set of natural scenes, and then infer the properties of the visual mechanisms that would best process those scenes. This is an appealing idea and seems both simple and direct, but the difficulty is that it is not clear what it means to say that a set of mechanisms are ‘best’ for processing natural scenes, because it is not clear what is to be optimized. One approach³⁴ is based on the hypothesis that the early visual system is designed to reduce redundancy in the neural code³⁵ and seeks the best set of mechanisms that satisfy a ‘sparse coding’ constraint. Others^{36,37} take a different approach and optimize their mechanisms for independence in the sense given by independent components analysis. A third approach finds mechanisms that satisfy a different criterion for independence³⁸. Each of these approaches is self-consistent and compelling in the terms the authors define, but the problem is that goals the brain satisfies in choosing neural codes are unknown. Because the specific coding models that emerge from these approaches depend in detail on the optimization chosen, the results can give only a qualitative impression of the true mechanisms. Another limitation of most of these efforts is that the computations usually generate mechanisms that only resemble the simple cells of the old standard model (**Fig. 1a**) and rarely incorporate more than one or two isolated features of the new standard model (**Fig. 1b**).

The second common use of natural images is more ambitious: to probe the visual system directly with natural images, motivated by the notion that the synthetic stimuli used in classic physiology experiments may not be sufficiently rich to uncover the full range of neuronal behavior¹. Implicit in this approach is the assumption that synthetic stimuli are in some way impoverished or ‘simplistic’ and therefore somehow miss important features of visual response. The main—and in our view, crippling—challenge is that the statistics of natural images are complex and poorly understood. Without understanding the constituents of natural images, it is imprudent to use them to develop a well-controlled hypothesis-driven experiment. Ironically, the only way to know what importance



to attach to different aspects of natural scene statistics is to have a comprehensive formal model of the neuron under study—but this is usually presented as the outcome rather than the foundation for the analysis.

The most popular way to try to solve this problem is to use ‘reverse correlation’ or spike-triggered approaches. In a traditional reverse correlation experiment, an experimenter constructs a stimulus-response model by presenting random stimuli (such as Gaussian white noise) and determining the characteristics of the subset of stimuli that elicit spikes (such as the spike-triggered average or spike-triggered covariance; for reviews, see refs. 33, 39). To yield an unbiased result, these techniques require that the stimuli are uncorrelated (‘white’), that their intensities are chosen from a Gaussian distribution and (in whatever descriptive space the experimenter chooses) that they span all dimensions of interest. A number of authors have applied similar approaches using natural images by modifying their analysis techniques to try to compensate for the strongly non-white and non-Gaussian nature of these stimuli^{40–47}. The problem is that when one extracts the subset of stimuli that are correlated with spikes, the correlations in the stimulus set make it difficult to determine why the stimulus-response correlation is present. Is it because of a mechanistic relationship between the stimulus and the neural circuit under study, or because of the correlations between the stimulus and other members of the stimulus set that may themselves be effective? Removing the effect of these correlations is straightforward if the system under study is simple. But if it has unknown architecture, it is necessary to make assumptions about the form of the underlying neural computation to distinguish the components of the response attributable to the neuron from those attributable to the correlations among the stimuli. The importance of these potential difficulties is difficult to determine with certainty, and it is certainly worthwhile to explore the question empirically. One ambitious attempt to use natural stimuli to study mechanism in cortical cells used a spike-triggered analysis and attempted to moderate the effects of stimulus correlations⁴¹. The data were fit to a model similar in form to the old standard model of **Figure 1a**, but this succeeded in capturing only a small fraction of the variance in the responses of V1 neurons to natural stimulus sequences, a result that falls well short of complete success.

So why do these analyses provide such disappointing results? Some have argued that this reflects deep limitations in our understanding of V1, and they suggest that we must refocus our efforts away from traditional, simple stimuli toward stimuli with more naturalistic characteristics¹ to overcome these limitations. But there is a simpler explanation. The significant and widely misunderstood limitation of spike-triggered approaches—using either synthetic or natural stimuli—is that they are not model-free; these techniques fit a specific model to the data, and in most cases this model is no more than a variant of the ‘old standard model’ shown in **Figure 1a**. Specifically, the data are used to fit a model in which the stimulus is first passed through one or more linear filters, the outputs of those filters are combined via an instantaneous nonlinearity, and noise is introduced into the system via a Poisson process. Without the elements of the new standard model (**Fig. 1b**), it is hardly surprising that the model performs poorly when put to a quantitative test. There is little doubt that the additional mechanisms represented in **Figure 1b** have a major role in natural scene responses. Yet the models evaluated in natural scene experiments lack these features; some work on subcortical processing suggests ways in which they might be incorporated⁴⁸. Another major limitation of the spike-triggered approaches is the absence of a realistic spike generator. These models are used to predict firing rate over the course of one or two frames (10–50 ms), yet within this time frame, deviations from Poisson spiking have a profound impact, especially at high firing rates. Techniques have been

proposed to incorporate realistic, non-Poisson spiking into spike-triggered characterizations^{30,31}, but these methods have not yet been generalized to models like the one in **Figure 1b**.

So the parsimonious interpretation of the ‘failure’ of the old standard model when faced with natural scenes is not very grand. Instead of reflecting some special feature of natural images that can reveal hitherto unsuspected neural machinery, it may be that the failure reflects only the limitations of the models used to evaluate the data, limitations that have been made very clear by numerous experiments using synthetic stimuli. In particular, until methods exist to fit the full model shown in **Figure 1b**, we cannot know whether the deficiencies in our ability to predict responses to natural stimuli from spike-triggered analyses—using either natural or synthetic stimulus sets—are due to known mechanisms or to novel ones.

Proper use of the natural and the artificial

Fitting and testing models that are general enough to predict the responses to arbitrary stimuli remains a central goal of visual neurophysiology. This process is useful because it provides a quantitative analysis of how close we are to reaching our goal of describing the behaviors of these neurons, and responses to natural scenes will always be the standard against which models are tested. But it seems to us that the limitations on using natural stimuli to build models rather than to test them are too important to ignore. Consider a scenario in which we are able to fit all known mechanisms with an integrated model but find that this model fails to accurately predict the response properties to natural scenes. What then? We would want to establish what machinery is missing from our description, and at this point, natural stimuli themselves are of no assistance. To determine mechanism, we must return to the classical approach of presenting artificial stimuli that are carefully designed as efficient and principled tests of specific hypotheses. In constructing these stimuli, we would be foolish to ignore the composition and structure of natural images, but only with synthetic stimuli could we carefully control and test model elements of computational importance. The proof of success will be found in predicting the responses to natural stimuli, but the predictions themselves will be made from artificial ingredients.

These disadvantages of natural stimuli do not invalidate their use in exploratory experiments. Indeed, for neurons with complex properties whose circuitry is unknown (such as those in higher cortical areas), these methods may be the best or even the only way to begin (for example, see ref. 49). But there are two phases to discovery. After the system is explored using stimuli that are chosen mostly for their effectiveness, model-building begins and our tools become the classical ones of hypothesis and test. In the study of primary visual cortex, we are fully engaged in that second stage. In our view, the proper use for studies of natural stimuli is to provide the benchmark against which success or failure can be measured. But success will require the model-guided use of artificial stimuli to uncover neuronal mechanism.

ACKNOWLEDGMENTS

The authors thank E.P. Simoncelli for discussions.

COMPETING INTERESTS STATEMENT

The authors declare that they have no competing financial interests.

Published online at <http://www.nature.com/natureneuroscience>

Reprints and permissions information is available online at <http://npg.nature.com/reprintsandpermissions/>

1. Olshausen, B.A. & Field, D.J. How close are we to understanding V1? *Neural Comput.* **17**, 1665–1699 (2005).
2. Lennie, P. & Movshon, J.A. Coding of color and form in the geniculostriate visual pathway. *J. Opt. Soc. Am. A* **22**, 1–21 (2005).

3. Jones, J.P. & Palmer, L.A. The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *J. Neurophysiol.* **58**, 1187–1211 (1987).
4. Ringach, D.L. Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *J. Neurophysiol.* **88**, 455–463 (2002).
5. Movshon, J.A., Thompson, I.D. & Tolhurst, D.J. Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *J. Physiol. (Lond.)* **283**, 53–77 (1978).
6. Adelson, E.H. & Bergen, J.R. Spatiotemporal energy models for the perception of motion. *J. Opt. Soc. Am. A* **2**, 284–299 (1985).
7. Spitzer, H. & Hochstein, S. A complex-cell receptive-field model. *J. Neurophysiol.* **53**, 1266–1286 (1985).
8. De Valois, K.K., De Valois, R.L. & Yund, E.W. Responses of striate cortex cells to grating and checkerboard patterns. *J. Physiol. (Lond.)* **291**, 483–505 (1979).
9. Smith, M.A., Bair, W. & Movshon, J.A. Signals in macaque striate cortical neurons that support the perception of glass patterns. *J. Neurosci.* **22**, 8334–8345 (2002).
10. Creutzfeldt, O.D. & Nothdurft, H.C. Representation of complex visual stimuli in the brain. *Naturwissenschaften* **65**, 307–318 (1978).
11. Touryan, J., Lau, B. & Dan, Y. Isolation of relevant visual features from random stimuli for cortical complex cells. *J. Neurosci.* **22**, 10811–10818 (2002).
12. Rust, N.C., Schwartz, O., Movshon, J.A. & Simoncelli, E.P. Spatiotemporal elements of macaque V1 receptive fields. *Neuron* **46**, 945–956 (2005).
13. Shapley, R.M. & Enroth-Cugell, C. Visual adaptation and retinal gain controls. *Prog. Retinal Res.* **3**, 263–346 (1984).
14. Bonds, A.B. Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Vis. Neurosci.* **2**, 41–55 (1989).
15. Geisler, W.S. & Albrecht, D.G. Cortical neurons: isolation of contrast gain control. *Vision Res.* **32**, 1409–1410 (1992).
16. Heeger, D.J. Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* **9**, 181–197 (1992).
17. Carandini, M., Heeger, D.J. & Movshon, J.A. Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.* **17**, 8621–8644 (1997).
18. Bair, W. & Movshon, J.A. Adaptive temporal integration of motion in direction-selective neurons in macaque visual cortex. *J. Neurosci.* **24**, 7305–7323 (2004).
19. Maffei, L., Fiorentini, A. & Bisti, S. Neural correlate of perceptual adaptation to gratings. *Science* **182**, 1036–1038 (1973).
20. Carandini, M., Movshon, J.A. & Ferster, D. Pattern adaptation and cross-orientation interactions in the primary visual cortex. *Neuropharmacology* **37**, 501–511 (1998).
21. Sanchez-Vives, M.V., Nowak, L.G. & McCormick, D.A. Membrane mechanisms underlying contrast adaptation in cat area 17 *in vivo*. *J. Neurosci.* **20**, 4267–4285 (2000).
22. Movshon, J.A. & Lennie, P. Pattern-selective adaptation in visual cortical neurones. *Nature* **278**, 850–852 (1979).
23. Chance, F.S., Nelson, S.B. & Abbott, L.F. Synaptic depression and the temporal response characteristics of V1 cells. *J. Neurosci.* **18**, 4785–4799 (1998).
24. Muller, J.R., Metha, A.B., Krauskopf, J. & Lennie, P. Rapid adaptation in visual cortex to the structure of images. *Science* **285**, 1405–1408 (1999).
25. Hubel, D.H. & Wiesel, T.N. Receptive fields and functional architecture of monkey striate cortex. *J. Physiol. (Lond.)* **195**, 215–243 (1968).
26. Blakemore, C. & Tobin, E.A. Lateral inhibition between orientation detectors in the cat's visual cortex. *Exp. Brain Res.* **15**, 439–440 (1972).
27. Angelucci, A., Levitt, J.B. & Lund, J.S. Anatomical origins of the classical receptive field and modulatory surround field of single neurons in macaque visual cortical area V1. *Prog. Brain Res.* **136**, 373–388 (2002).
28. Cavanaugh, J.R., Bair, W. & Movshon, J.A. Nature and interaction of signals from the receptive field center and surround in macaque V1 neurons. *J. Neurophysiol.* **88**, 2530–2546 (2002).
29. Sceniak, M.P., Ringach, D.L., Hawken, M.J. & Shapley, R. Contrast's effect on spatial summation by macaque V1 neurons. *Nat. Neurosci.* **2**, 733–739 (1999).
30. Keat, J., Reinagel, P., Reid, R.C. & Meister, M. Predicting every spike: a model for the responses of visual neurons. *Neuron* **30**, 803–817 (2001).
31. Paninski, L., Pillow, J.W. & Simoncelli, E.P. Maximum likelihood estimation of a stochastic integrate-and-fire neural encoding model. *Neural Comput.* **16**, 2533–2561 (2004).
32. Aguera y Arcas, B., Fairhall, A.L. & Bialek, W. Computation in a single neuron: Hodgkin and Huxley revisited. *Neural Comput.* **15**, 1715–1749 (2003).
33. Simoncelli, E.P., Pillow, J.W., Paninski, L. & Schwartz, O. Characterization of neural responses with stochastic stimuli. in *The Cognitive Neurosciences* (ed. Gazzaniga, M.) 327–338 (MIT Press, Cambridge, Massachusetts, 2004).
34. Olshausen, B.A. & Field, D.J. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381**, 607–609 (1996).
35. Barlow, H.B. Possible principles underlying the transformations of sensory messages. in *Sensory Communication* (ed. Rosenblith, W.A.) 217–234 (MIT Press, Cambridge, Massachusetts, 1961).
36. Bell, A.J. & Sejnowski, T.J. The independent components of natural scenes are edge filters. *Vision Res.* **37**, 3327–3338 (1997).
37. van Hateren, J.H. & van der Schaaf, A. Independent component filters of natural images compared with simple cells in primary visual cortex. *Proc. Royal Soc. Lond. B Biol. Sci.* **265**, 359–366 (1998).
38. Schwartz, O. & Simoncelli, E.P. Natural signal statistics and sensory gain control. *Nat. Neurosci.* **4**, 819–825 (2001).
39. Chichilnisky, E.J. A simple white noise analysis of neuronal light responses. *Network* **12**, 199–213 (2001).
40. Theunissen, F.E. *et al.* Estimating spatio-temporal receptive fields of auditory and visual neurons from their responses to natural stimuli. *Network* **12**, 289–316 (2001).
41. David, S.V., Vinje, W.E. & Gallant, J.L. Natural stimulus statistics alter the receptive field structure of V1 neurons. *J. Neurosci.* **24**, 6991–7006 (2004).
42. Touryan, J., Felsen, G. & Dan, Y. Spatial structure of complex cell receptive fields measured with natural images. *Neuron* **45**, 781–791 (2005).
43. Sharpee, T., Rust, N.C. & Bialek, W. Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comput.* **16**, 223–250 (2004).
44. Smyth, D., Willmore, B., Baker, G.E., Thompson, I.D. & Tolhurst, D.J. The receptive-field organization of simple cells in primary visual cortex of ferrets under natural scene stimulation. *J. Neurosci.* **23**, 4746–4759 (2003).
45. Ringach, D.L., Hawken, M.J. & Shapley, R. Receptive field structure of neurons in monkey primary visual cortex revealed by stimulation with natural image sequences. *J. Vis.* **2**, 12–24 (2002).
46. Prenger, R., Wu, M.C., David, S.V. & Gallant, J.L. Nonlinear V1 responses to natural scenes revealed by neural network analysis. *Neural Netw.* **17**, 663–679 (2004).
47. Felsen, G., Touryan, J., Han, F. & Dan, Y. Cortical sensitivity to visual features in natural scenes. *PLoS Biol.* **3**, e342 (2005).
48. Mante, V., Frazor, R.A., Bonin, V., Geisler, W.S. & Carandini, M. Independence of luminance and contrast in natural scenes and in the early visual system. *Nat. Neurosci.* **8**, 1690–1697 (2005).
49. Gross, C.G., Rocha-Miranda, C.E. & Bender, D.B. Visual properties of neurons in inferotemporal cortex of the macaque. *J. Neurophysiol.* **35**, 96–111 (1972).