

An unsupervised learning model of neural plasticity: Orientation selectivity in goggle-reared kittens

Anne S. Hsu ^{*}, Peter Dayan

Gatsby Computational Neuroscience Unit, University College London, 17 Queen Square, London WC1N 3AR, UK

Received 22 January 2007; received in revised form 13 June 2007

Abstract

The selectivities of neurons in primary visual cortex are often considered to be adapted to the statistics of natural images. Accordingly, simple cell-like tuning emerges when unsupervised learning models that seek sparse representations of input probabilities are trained on natural scenes. However, orientation tuning develops before structured vision starts, rendering these previous results moot as models of activity-dependent development. A more stringent examination of such models comes from experiments demonstrating altered neural response properties in goggle-reared kittens. We show that an unsupervised learning model of cortical responsivity accounts well for the dramatic effects of stimulus driven development during goggle-rearing.

© 2007 Elsevier Ltd. All rights reserved.

Keywords: Goggle-rearing; Unsupervised learning; Developmental models; Primary visual cortex; Plasticity

1. Introduction

Although the mammalian visual system undergoes substantial development before experiencing structured vision (Barlow, 1975; Blakemore & Van Sluysters, 1975; Buisseret & Imbert, 1976; Crair, Gillespie, & Stryker, 1998; Frégnac & Imbert, 1984; Godecke & Bonhoeffer, 1996; Hubel, 1988; Movshon & Blakemore, 1974), there is much evidence that visual neurons' response properties are significantly influenced by the statistics of visual inputs after eye opening (Blakemore & Cooper, 1970; Hirsch, Leventhal, McCall, & Tieman, 1983; Hirsch & Spinelli, 1970; Hubel, 1988; Hubel & Wiesel, 1963; Movshon & Blakemore, 1974; Sengpiel, Stawinski, & Bonhoeffer, 1999; Stryker & Sherk, 1975; Stryker, Sherk, Leventhal, & Hirsch, 1978; Tanaka, Ribot, & Tani, 2006; Wiesel, 1982). This poses important mechanistic questions about neural plasticity, and functional questions about how neural response properties are shaped by the visual environment. Indeed, many theoretic

cal studies have provided mechanistic (Barlow, 1975; Goodhill, 1993; Goodhill & Willshaw, 1990; Miller, Erwin, & Kayser, 1999; Miller, Keller, & Stryker, 1989; Swindale, 1996; Tanaka, Miyashita, & Ribot, 2004; von der Malsburg, 1973; Willshaw & von der Malsburg, 1979; Willshaw & von der Malsburg, 1976) and functional (Bell & Sejnowski, 1997; Hyvarinen & Hoyer, 2001; Karklin & Lewicki, 2003; Lewicki, 2002; Li & Atick, 1994; Olshausen & Field, 1997; Osindero, Welling, & Hinton, 2006; Rehn & Sommer, 2007) accounts.

Functional models aim to provide insight into the representational capacities and goals of the visual system rather than to offer mechanistic explanations of how neural responses are, or come to be, implemented. The most widely investigated class of functional models for visual receptive fields are based on informational (Barlow, 1981; Zhaoping, 2006; Li & Atick, 1994; Linsker, 1990) and probabilistic (Hinton & Ghahramani, 1997) notions, and involves variants of independent component analysis (ICA) or sparse/independent coding. ICA-like models posit that the functional goal of the visual cortex is to capture the statistics of its input with sparse and marginally- or conditionally-independent coding elements (Bell & Sejnowski,

^{*} Corresponding author.

E-mail addresses: ahsu@gatsby.ucl.ac.uk (A.S. Hsu), dayan@gatsby.ucl.ac.uk (P. Dayan).

1997; Hyvarinen & Hoyer, 2001; Olshausen & Field, 1997; Osindero et al., 2006; Rehn & Sommer, 2007; Teh, Welling, Osindero, & Hinton, 2003). When trained on natural images, these models result in linear or log-linear receptive fields (RFs) which are localized and orientation-tuned (i.e. Gabor-like) and also resemble other response properties of simple cells in V1.

Since most models have only been applied to normal development, much of which occurs without guidance from natural scene stimuli, it has remained unclear whether these descriptions of cortex apply to stimulus-driven cortical changes. An obvious way to test these ideas is to observe the consequences of changing the visual environment in a structured manner. Three broad classes of manipulation have been used, affecting the input from a single eye (e.g. monocular deprivation Hubel, Wiesel, & LeVay, 1977), the gross correlations between the two eyes (e.g. strabismus Hubel & Wiesel, 1965); or the fine correlations within both eyes (e.g. rearing with access to only restricted orientations Blakemore & Cooper, 1970; Hirsch & Spinelli, 1970; Sengpiel et al., 1999; Stryker & Sherk, 1975; Stryker et al., 1978; Tanaka et al., 2006). The last class is perhaps the most revealing, since the manipulation is less drastic than deprivation, but the results are much more dramatic than strabismus. However, except for Tanaka's mechanistic treatment (Tanaka et al., 2004), it has not so far been well addressed by modelling studies. In this paper we examine the validity of a functional description for visual cortical organization which has been previously widely applied only to normal visual development, by applying it to the case of development in the face of vision-distorting goggles worn for many months after eye-opening.

Against a backdrop of long-standing contention as to the neural consequences of rearing kittens in either striped cylinders or with goggles that limit the angle of view (Blakemore & Cooper, 1970; Hirsch & Spinelli, 1970; Stryker & Sherk, 1975; Stryker et al., 1978), recent studies have reinforced the view that restricted exposure to a single orientation significantly influences the development of orientation selectivity in primary visual cortex, biasing neurons towards representing the exposed orientation (Sengpiel et al., 1999; Tanaka et al., 2006). Tanaka and colleagues showed that severely restrictive striped goggle rearing exerts widespread influence on neural development (Tanaka et al., 2006) with the number of neurons in goggle-reared kittens preferring the exposed orientation being over three times that for normal kittens. There are also more subtle changes, such as to the shapes of the receptive fields and the proportion of oriented receptive fields. Tanaka and his colleagues used both optical imaging and electrophysiological methods to characterize the effects of their goggle-rearing regimen, and confirmed that these methods were mutually consistent. We show that the results of this stimulus driven cortical plasticity can be captured by an unsupervised learning model with the functional goals of capturing input statistics with sparsely distributed responses. To our knowledge this is the first

application of such models to the altered neural organization effected by major and sustained experimentally-controlled manipulations of developmental stimulus statistics.

2. Methods

It has long been observed that oriented, simple cell-like, RFs arise from applying to natural scenes, learning algorithms which are designed to capture the statistical structure of their input. This has motivated accounts of the functional goals of the cortical representation based on statistical unsupervised learning.

In particular, we consider the products-of-experts (POE) algorithm (Osindero et al., 2006; Teh et al., 2003), which is an extension of independent components analysis (ICA) to overcomplete representations. In the POE model, representational units respond to linear projections of the input, and capture structure by reporting unlikely stimuli, or those that can be described as violating constraints defined by the RFs. When trained on conventional natural scene input, the POE model produces model neuron units which have RFs resembling Gabor-like receptive fields, with a range of preferred spatial frequencies, orientations and shape characteristics that fairly well match those observed in cat and monkey cortex. Here, we study the consequences of training a POE with the unnatural scenes arising from a model of goggle-rearing.

In the POE model applied to natural scenes by (Osindero et al., 2006), the probability of the input is modelled as a product of generalized Student- t distributions. This model captures the statistics of its inputs, \mathbf{x} , as follows:

$$p_{\text{model}}(\mathbf{x}) = \frac{1}{Z(\theta)} \prod_{i=1}^n \frac{1}{(1 + \frac{1}{2}\mathbf{y}_i^2)^{\alpha_i}} \quad (1.1)$$

$$\mathbf{y}_i = \mathbf{J}_i \mathbf{x}$$

where \mathbf{x} contains the image data inputs, \mathbf{y}_i is the activity of the i th neuron, n is the total number of output units or model neurons and θ contains all the parameters of the model (\mathbf{J} and α). The linear mapping between input and output, is defined by matrix \mathbf{J} , where the i th row of \mathbf{J} is the response function, or in our case, the visual receptive field of the i th neuron α is a learned parameter that corresponds to the sparseness of the Student- t distribution. This decreases with increasing overcompleteness and ranges from 1 to 1.8. $Z(\theta)$ is the normalization constant. The parameters α and \mathbf{J} are adapted on the input statistics by maximizing the log-likelihood of this model given the input. That is, we seek:

$$\arg \max_{\mathbf{J}, \alpha} \{ \langle \log(p_{\text{model}}(\mathbf{x})) \rangle_{\mathbf{x} \sim \text{data}} \}$$

The gradient of the log likelihood is

$$\frac{\partial \langle \log(p(\mathbf{x})) \rangle}{\partial \theta} = \left\langle \frac{\partial}{\partial \theta} \left(-\alpha_i \sum_{i=1}^n \log \left(1 + \frac{1}{2} (\mathbf{J}_i \mathbf{x})^2 \right) - \log Z(\theta) \right) \right\rangle_{\mathbf{x} \sim \text{data}} \quad (1.2)$$

where $\langle \rangle_{\mathbf{x} \sim \text{data}}$ means an average over all image data and θ refers to either of the parameters α or \mathbf{J} .

In the complete case, for which the total number of neurons n equals the number of dimensions in the input \mathbf{x} , this model is exactly equivalent to standard ICA (Teh et al., 2003), and Eq. (1.2) can be written in a simple, closed form because $Z(\theta)$ can be shown to be equal to $|\text{inv}(\det(\mathbf{J}))|$. In general, ICA can be performed with a variety of functional forms other than that in the product of Eq. (1.1). (See Bell & Sejnowski, 1997 for further discussion). In these cases the exact gradient can be computed and the learning update rules are exactly equal to those proposed in traditional ICA (Teh et al., 2003).

In the over-complete cases, for which the number of neurons is greater than the dimensions of \mathbf{x} , the normalizing constant $Z(\theta)$ no longer has a simple analytic form and thus the exact gradient is computationally intractable. However, this gradient (Eq. (1.2)) can be shown to be equal to:

$$\frac{\partial(\log(p(\mathbf{x})))}{\partial\theta} = \left\langle \frac{\partial}{\partial\theta} \left(-\alpha_i \sum_{i=1}^n \log \left(1 + \frac{1}{2} (\mathbf{J}_i \mathbf{x})^2 \right) \right) \right\rangle_{\mathbf{x} \sim \text{data}} - \left\langle \frac{\partial}{\partial\theta} \left(-\alpha_i \sum_{i=1}^n \log \left(1 + \frac{1}{2} (\mathbf{J}_i \mathbf{x})^2 \right) \right) \right\rangle_{\mathbf{x} \sim \text{model}} \quad (1.3)$$

where $\langle \rangle_{\mathbf{x} \sim \text{data}}$ means averaged over samples of the input training data such that \mathbf{x} consists of samples from the input \mathbf{x} , and $\langle \rangle_{\mathbf{x} \sim \text{model}}$ means an average over samples where \mathbf{x} is drawn from $p_{\text{model}}(\mathbf{x})$ as given by Eq. (1.1), which is the distribution of data defined by the current model. In general $p_{\text{model}}(\mathbf{x})$ is not equal to the actual distribution of the image data. When the two do become equal, the gradient (1.3) becomes zero and (typically) a (global) maximum has been reached. In order to approximate the $\langle \rangle_{\mathbf{x} \sim \text{model}}$ term in the gradient, Eq. (1.3), Monte Carlo Markov Chain (MCMC) sampling methods must be used in order to sample from $p_{\text{model}}(\mathbf{x})$. An approximate scheme called Contrastive Divergence (CD) is available for avoiding part of the computational cost of this. Further details of CD learning and MCMC methods are provided in (Teh et al., 2003).

The inputs to the algorithm were derived from eleven natural images taken from the Van Hateren database (van Hateren & van der Schaaf, 1998). An input was either a 25×25 patch of a natural scene (each pixel has an angular resolution of approximately 2 min of arc), or a 25×25 patch of a natural scene that had been filtered in a way that simulates the effects of the goggles that Tanaka used. This was achieved by obtaining a 2D Fourier transform of the image and keeping the power for all components at the orientation selected by the goggles (GO) while performing Gaussian blurring on all other components (see Fig. 1). Our simulated filter had a bandwidth such that only a 7 degree range of orientations kept more than 30% of their maximum power.

Tanaka implemented what appears to be the most restrictive goggle-rearing regimen to date (Tanaka et al., 2006), fitting freely moving kittens with cylindrical lens goggles continuously from post-natal week three. For our purposes, the best way to match this regimen would not necessarily be to use exclusively goggle-filtered scenes as model inputs because of the many factors that may limit the instructive effects of goggle inputs. For instance, the presence of spontaneous activity whose structure is determined by developmental factors that precede the goggle rearing (Miyashita, Kim, & Tanaka, 1997; Miyashita & Tanaka, 1992; Ruthazer & Stryker, 1996; Tanaka et al., 2004; Thompson, 1997) could provide a counteracting bias. Indeed many studies have shown that relatively normal neural responses are found in animals raised in completely dark environments (Barlow, 1975; Blakemore & Van Sluyters, 1975; Buisseret & Imbert, 1976; Crair et al., 1998; Frégnac & Imbert, 1984; Godecke & Bonhoeffer, 1996; Movshon & Blake-more, 1974), which suggests that there is a spontaneous activity that may serve to restore normal tuning properties. Indeed, kittens which experienced longer periods of dark rearing showed more modest effects of goggle-rearing (Sengpiel et al., 1999; Tanaka et al., 2006). Since functional models trained on natural scenes have previously been shown to produce results which resemble the tuning properties of normal neurons (Bell & Sejnowski, 1997; Hyvarinen & Hoyer, 2001; Olshausen & Field, 1997; Osindero et al., 2006; Rehn & Sommer, 2007; van Hateren & van der Schaaf, 1998), we explore the effects of mixing different proportions

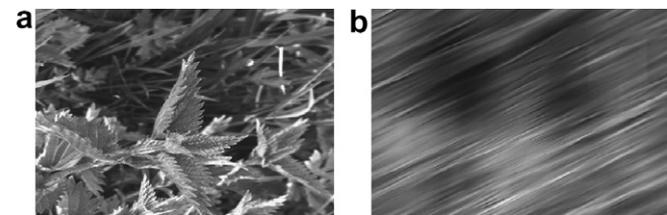


Fig. 1. Computer-simulated goggles. (a) An original natural image from the van Hateren database. (b) The image in (a) transformed to simulate goggles that restrict the image power to an input orientation of 56 degrees.

of goggle-filtered and un-filtered natural scene inputs. Differences in results induced by this, along with differences in rigor among striped rearing protocols, may help account for the variability in the outcomes of early goggle- and stripe-rearing experiments.

Here, we simulated the effects of goggle-rearing by training a (POE) model with inputs consisting of unadulterated natural scenes and/or natural scenes that had been filtered with (software-defined) goggles (see Fig. 1 and Supplementary Fig. 1). The cortical representation is substantially overcomplete in the sense that the visual cortex has roughly two orders of magnitude more cells than its thalamic input (Olshausen & Field, 1997) so we applied POE models with differing degrees of overcompleteness. As discussed above, different proportions of natural-scene inputs were also tested.

More formally, we created stimulus sets consisting of image patches from which 0, 50, 75, 90, 95, 98, and 100% were selected from goggle-filtered scenes, mixed randomly with others from un-altered natural scenes. Each set consisted of 150,000 samples in total. As is the convention, the input intensities were logarithmically transformed (Osindero et al., 2006; van Hateren & van der Schaaf, 1998), and were then whitened and reduced in dimension using principal components analysis to the most variable 196 dimensions. Whitening and the logarithmic transform speed up learning; they do not significantly change the results of learning. Learning proceeded in this whitened, reduced space using one-step contrastive divergence learning on batch sizes of 200 patches for 200 iterations. See Supplementary Methods for further learning details. See also (Osindero et al., 2006) for the details of the learning algorithm. The learned RFs were confirmed as being stable after about 100 iterations.

RFs were learned for 1 \times , 2 \times , and 4 \times overcomplete representations (quantifying the number of model neurons compared with the number of input dimensions) yielding sets of either 196, 392, or 784 RFs. All summary figures presented in the Results are averages of 4 runs with the same input and degree of model overcompleteness, using different initial random seeds for RF values. Following (Teh et al., 2003), α was set to 1.5 for the complete case and was initialized to 1 and learned for the overcomplete cases.

We used standard methods (Osindero et al., 2006; Ringach, 2002; Tanaka et al., 2006) to investigate the receptive field properties of the resulting individual RFs, such as their orientation selectivity, and also the properties of the population, such as the fraction tuned to the orientation favored by the goggles. In particular, we obtained tuning curves for each of our model neurons by averaging the absolute values of their responses to all phases of cosine gratings spaced 2 degrees apart. Absolute values were used in measuring tuning curves because ICA-like models do not differentiate between positive and negative model neuron responses (see the quadratic term in Eq. (1.1)). The frequencies of the cosine gratings were matched to the best-frequency of the Gabor functions that fit each of the model neurons most closely. RF parameters of aspect ratio and lengths were computed from the length/width and width, respectively, of the 2D Gaussian envelope corresponding to the best-fit Gabor functions of each neuron. Unless stated otherwise, histograms of preferred orientation for our model neurons are shown for 36 bins centered at 5–180 degrees. Best orientations of our model neurons were mapped circularly onto the range 2.5–182.5 degrees.

3. Results

The POE model, using regular natural scene input, has been shown (Osindero et al., 2006) to provide a reasonable fit to many receptive field (RF) properties of simple cells in the normal cat and macaque monkey. Model results compare well with simple cell properties of peak spatial frequency, spatial frequency bandwidth, aspect ratio, and shape (De Valois, Albrecht, & Thorell, 1982; Jones & Palmer, 1987; Parker & Hawken, 1988; Ringach, 2002). (See Supplementary Fig. 5.)

We studied the RF properties for the different degrees of overcompleteness and the different strengths of goggle-reared input (see Fig. 1). Except for the cases involving the most extreme percentages of goggle-filtered stimuli, many model neuron units developed localized, oriented gabor-like receptive fields (RFs) with well-formed, unimodal tuning curves (having a single peak and a circular variance as defined by Ringach (2002) of ≤ 0.6). Fig. 2a shows two examples.

As has also been observed with natural scene inputs (Osindero et al., 2006), the degree of overcompleteness has an effect on the total number of such localized model neurons. In particular, the absolute number rises with the total number of available units, but, as a percentage, falls (Fig. 2b; 0% bars). Fig. 2b further shows that the number of localized neurons also decreases when the proportion of goggle-filtered stimuli used in the training set is increased. This decrease can possibly be explained by the following: Localized RFs are necessary for maintaining sparse responses to inputs with natural scene statistics inputs. Since goggle-filtered stimuli occupy a much smaller volume of dimensional space than natural stimuli, the necessity for localized responses decreases. In Tanaka's data, two out of three goggle-reared kittens showed about 1.5–2 times as many non-oriented units as the normally-reared kitten (see Figs. 3 and 4 in Tanaka et al., 2006). In our model, the non-oriented neurons (Supplementary Fig. 2) also contribute to the appropriate representation of the statistics of input images. Unfortunately, the response properties of non-oriented neurons in goggle-reared kittens have not been documented, and their contribution to visual representation and behaviour is not clear. In view of the scope of the data in Tanaka et al. (2006), for the rest of the results, we report statistics only for model neurons that are uni-modal and pass a standard test for orientation-selectivity (i.e. circular variance ≤ 0.6).

The key consequence of training with stimuli containing an over-representation of the goggle-filtered orientation (GO) is an asymmetry in the distribution of cells which

favors that orientation. Different degrees of overcompleteness and over-representation lead to qualitative results that have different strengths, a finding that may have relevance to some of the early controversies in the experimental field about the relative strengths of the effects of different rearing protocols. We also find that for very high percentages of goggle-filtered stimuli, the preferred distribution histogram changes significantly in a way that has not been reported experimentally.

More formally, the degree of over-representation in each visual area can be defined by the over-representation index (ORI), which for optical imaging data, is the ratio of the number of pixels selective to GO, N_{GO} , to the average number of pixels per bin for all other orientations in the histogram, $N_{average}$.

$$ORI = \frac{N_{GO}}{N_{average}}$$

In Tanaka's study, area 17 and 18 of kittens reared with a strict goggle regimen showed ORI values ranging from 3.74 to 12. A kitten whose goggle rearing was interleaved with dark-rearing episodes showed lower ORI values (a result that was expected, since spontaneous activity in darkness was believed to aid normal development). In contrast, neurons in normally reared kittens showed no bias in orientation tuning preference (see Fig. 3g in Tanaka et al., 2006).

Fig. 3 shows the complete set of distributions for all levels of overcompleteness and all proportions of bias towards goggle-filtered inputs (here favoring the orientation 56 degrees). The plots that are highlighted in bold show the parameter range which produces results most comparable to Tanaka's data. In these cases, there is an over-representation of RFs oriented at GO and a roughly even distribution of neurons oriented at other orientations (with the exception that there is an absence of neurons immediately on either side of GO, a point we consider in more detail below). The over-representation of the GO, which scales with the proportion of goggle-filtered stimuli, can be over 3 times that of other orientations, as is also seen in the

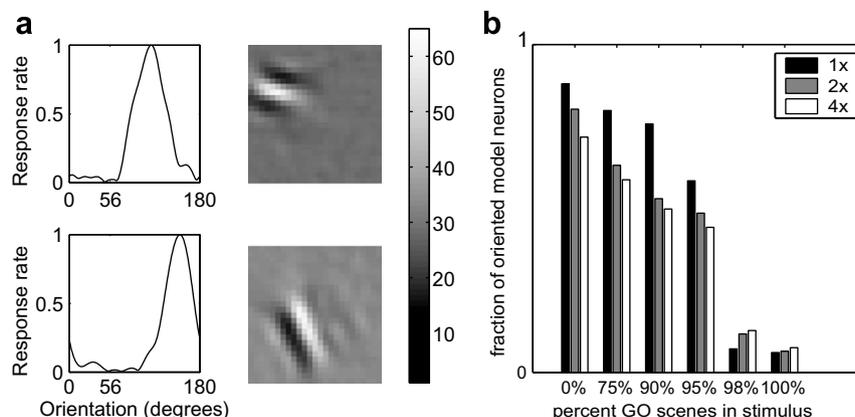
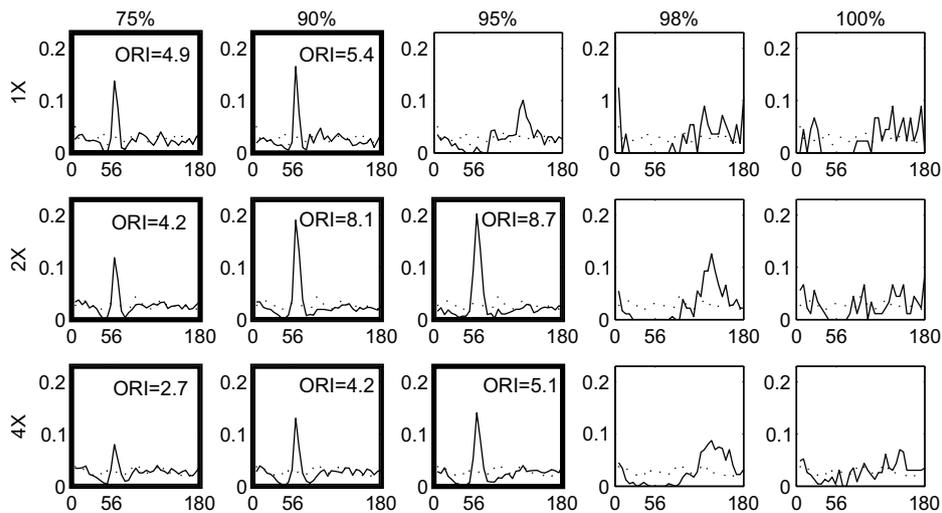


Fig. 2. Oriented localized tuning curves and their proportions. (a) Examples of tuning curves and respective localized RFs of our model neurons (taken from 2 \times overcomplete model trained on input consisting of 90% goggle-filtered stimuli). (b) Fraction of localized neurons for 1 \times , 2 \times , and 4 \times overcomplete models when trained with input with different percentages of goggle-filtered scenes mixed with natural scenes.

a) Histograms of model neurons' preferred orientation as a function of percentage of goggle-filtered stimuli and overcompleteness



b) Histogram of preferred orientation in goggle-reared kitten

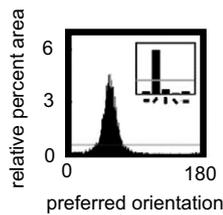


Fig. 3. Preferred orientation distributions. (a) Normalized histograms of model neurons' preferred orientations as a function overcompleteness (1 \times , 2 \times , 4 \times) and percentage of goggle-filtered scenes (75%, 90%, 95%, 98%, 100%) oriented at 56 degrees in the input. Dotted lines indicate the distribution when trained on purely natural scene input. In bold are results that are comparable to real neurons in goggle-reared kittens and for which over-representation indices (ORI) are shown. (b) Histogram of preferred orientation of pixels mapped using optical imaging of Area 17 for a goggle-reared kitten at 56 degrees with ORI value = 8.43 (Tanaka et al., 2006).

Tanaka data. Exact ORI values depend on the histogram bin width used for calculation. Tanaka's study used 1 degree bins. In our study, ORI was calculated from a preferred orientation histogram with 5 degree bins in order to obtain histograms with an adequate number of samples in each bin. The ORI values reported for our model neurons are averages over 4 runs. ORI values did not change by more than 10% when using histograms with 1, 2, 5 and 10 degree bins, and trends in the values remained the same.

We also found that the tuning curve widths, RF aspect ratios and lengths of our model neurons were similar to those of real neurons. In the parameter range that appears relevant to the experimental data, results with different percentages of goggle filtered stimuli are very similar so we only show results for 90% goggle-filtered stimuli. For our model neurons, the tuning widths (measured as full width at half maximum of the tuning curve, FWHM) are smaller for those neurons tuned to GO (Fig. 4a). Exactly the same phenomenon is apparent in Tanaka's data (Figs. 3, 4, Tanaka et al., 2006). Note that the tuning widths of RFs measured using unit recordings cannot be exactly compared to tuning widths measured by imaging (Tanaka et al., 2006). Fig. 4b shows the mean FWHM as a function

of orientation for imaged pixels of the same goggle-reared cat as Fig. 3b (Tanaka et al., 2006). Observe that while the black and white lines in Fig. 3a and b respectively both represent the mean values of model and data neurons' FWHM as a function of preferred orientation, the background on which the mean is overlaid differs. Results for our model neurons are overlaid on a 2D density histogram whereas the plot from Tanaka et al. (2006) is just a 2D plot of all imaged points, gray scale (in print)/color (online)-coded for preferred orientation.

Finally, in addition to having narrower tuning widths, our GO neurons are found to have receptive fields (RFs) with greater aspect ratios (length vs. width) and trends towards greater absolute spatial lengths than neurons at other orientations (see Fig. 5). Neurons with extremely long RFs which were elongated in parallel to the GO were also observed in the experimental study (Tanaka, personal communication, 2006). An earlier study involving a less severe form of stripe rearing (Hirsch & Spinelli, 1970), also found, using electrophysiology, that receptive fields at GO were markedly large. Additionally, our model GO neurons have RFs with higher spatial frequencies than neurons at other orientations (see Supplementary Fig. 3). These high spatial frequency RFs

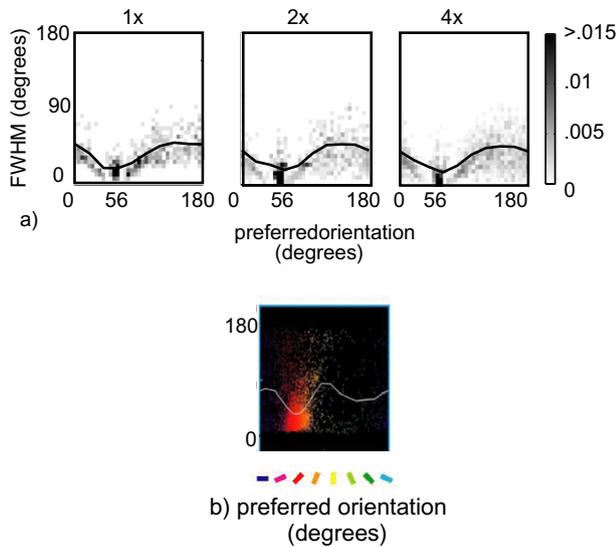


Fig. 4. Tuning width as a function of preferred orientation. (a) Full width half max (FWHM) as a function of preferred orientation for our model neurons trained on 90% goggle-filtered stimuli at 56 degrees with 1 \times , 2 \times , and 4 \times overcompleteness. For clarity we have replaced the scatter plot shown in Tanaka et al. (2006) by a 2D histogram of model neuron density with the mean value for all model neurons with a given preferred orientation plotted over in black. Preferred orientation bin widths are 5 degrees for the density plot and 20 degrees for the mean value plot. FWHM bin widths are 5 degrees. (b) FWHM from optical imaging of Area 17 of the same kitten as Fig. 3b, goggle-reared with input orientated at 56 degrees (Tanaka et al., 2006). The locations of gray scale (in print)/color (online) points on the plots indicate the FWHM and preferred orientation of a particular cortical location measured by imaging. The white line is the average FWHM for all imaged locations which had the given preferred orientation shown below.

are a result of the fact that goggle-filtered images have power in high spatial frequencies only at the goggle-orientation. This is a significant prediction of our model.

The most noticeable discrepancy between the output of our model using these parameters and the results from experimental goggle-rearing is the apparent absence of model neurons which prefer orientations that are less than 10° away from GO (see Fig. 3). There is no evidence for this in the kitten data. In our model, the pressure for this absence results from the strong competition among model neurons imposed by the sparseness constraint, not restricted (perhaps unlike cortex) by any constraint to the extent of receptive fields changes in response to the changed input statistics. This feature would likely be common to all models that prefer sparse outputs, for reasons that are discussed in the Supplementary Note. Furthermore, our model neurons oriented away from GO also contribute to the representation of goggle-filtered stimuli through a systematic dip in their tuning functions at GO (see Supplementary Fig. 4), a phenomenon which has also not been reported experimentally. It may not be present in real neurons due to the innate tendencies for natural tuning properties and limits on developmental plasticity (see Section 4).

The plots that are not highlighted in Fig. 3 show that the model can produce other classes of output when the bal-

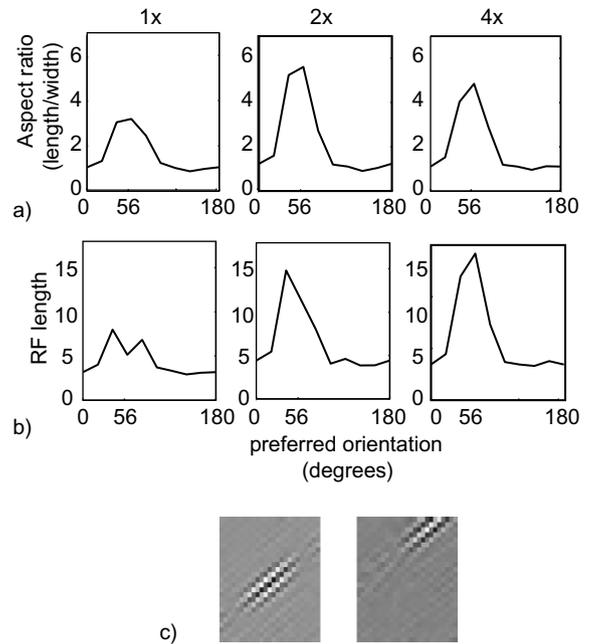


Fig. 5. Mean aspect ratios and RF lengths as functions of preferred orientation. (a) Aspect ratio (length/width of RF as defined by standard deviations of the envelope of the best-fit Gabor function), as a function of preferred orientation for our model neurons trained on 90% goggle-filtered stimuli at 56 degrees with 1 \times , 2 \times , and 4 \times overcompleteness. (b) Same as a, but for absolute spatial length measured in units of image pixels. Bin widths for aspect ratio and length are one unit. Orientation bin widths are 20 degrees. (c) Sample elongated RFs of model neurons oriented at GO.

ance between overcompleteness and input overrepresentation is different. At a critical percentage of goggle-filtered stimuli, the representation shifts to having no neurons preferring GO, along with a sharp increase of unoriented RFs. This critical percentage increases with the degree of overcompleteness (further simulations with results not shown were done to verify this trend). For the highest proportions of goggle-stimuli, most RFs were large and un-localized, and at 98% and 100%, few RFs even had obvious tuning characteristics. The unlocalized RFs appear because the lack of structure (i.e. the relatively small dimensional volume) in goggle-filtered stimuli no longer requires oriented, localized RFs for the production of sparse responses.

Compared with the experimental data, it may be that the percentage of goggle-filtered stimuli here is too large for there to be an experimental goggle-rearing regimen of comparable strength in the presence of innately driven spontaneous activity, especially in light of the ostensible extreme overcompleteness present in the cortex. Alternatively, these anomalous results could be a failing of this model for describing the most extreme regimens of goggle-rearing, revealing a limitation of this type of model in the face of extreme stimulus driven cortical alterations.

One dimension of variation that may be important in comparing the results of different experiments on goggle-rearing is the strength of the goggles used, which determines the severity to which visual input is restricted to a single orientation. To explore the effects of different

strengths, we simulated goggles for which a 20 degree range of orientations kept more than 30% of their maximum power. These goggles are significantly weaker than those used in the Tanaka study. The weaker goggles resulted in systematically smaller ORI values, consistent with the trend that rather weaker effects of GO over-representation were observed in previous experiments (Hirsch et al., 1983; Sengpiel et al., 1999). Additionally, the weaker goggles do not result in dips in the orientation histograms near but not at GO. That is, there is no longer an absence of model neurons tuned near to, but not at, GO. This absence reflects the fact that the weaker goggles allow through more power at a wider range of orientations, therefore weakening the competition for GO oriented RFs.

Another feature of model RFs learned with weaker goggles is the population representation at extreme proportions of GO stimuli (98–100%). Here, instead of the prevalence of un-localized RFs along with the complete absence of RFs aligned to GO that arise with the strong goggles, two characteristic types of localized RFs arise: one set which are elongated and aligned parallel with GO (as was also observed for lower percentages of GO), and one set which are fatter, and orthogonal to GO. Intuitively, the parallel RFs arise because the presence of power at orientations other than GO favors the same sort of solution as for lower percentages of GO stimuli; the orthogonal RFs arise to help represent the extreme proportions of GO.

In models such as ours, the whole collection of units evolves together, and is jointly responsible for representing inputs. Thus the development of one unit depends on the other units that are present. To explore the effects of limits to cortical plasticity, perhaps associated with a critical period, we simulated a case in which all the units were initialized with receptive fields that arise when the model is trained on natural-scenes data, and then a proportion of the units had their receptive fields clamped during the subsequent presentation of goggle-filtered images, so that they would preserve their original tuning. We observed (data not shown) that these clamped units did not significantly affect the pattern of results. Indeed, the distribution of RFs was similar to those seen before, and the GO oriented model neurons had the same elongated RFs and narrow tuning widths described above. Finally, we also ran our simulations without dimensionality reduction on smaller image patches (14×14 and 16×16). This produced similar results to our original dimension-reduced analysis with the only difference being that the presence of additional power at high frequencies had the effect of shifting the transition from representations with an over-representation of GO parallel RFs to an absence of GO parallel RFs to a lower percentage of GO-oriented stimuli (75%) for the $1 \times$ complete case.

4. Discussion

We have shown that a model based on particular principles of statistical unsupervised learning captures many features of cortical neural development under the drastic

environmental manipulation of rearing with goggles permitting exposure to just a single visual orientation (GO). The resulting model neurons share with the neurons of goggle-reared kittens an over-representation of RFs with preferred orientation at GO, a lower proportion of oriented-localized neural RFs relative to normally-reared kittens, narrower tuning widths for RFs at GO, and larger and more elongated shaped RFs at GO. We also found that the degree of over-representation of RFs at GO scales markedly with both the proportion of goggle-filtered stimuli used in the input and the strength of the simulated goggles, which is our attempt to replicate different degrees of severity in goggle-rearing protocols. This may explain the variable extents of over-representation observed across the many various experiments (Blakemore & Cooper, 1970; Hirsch & Spinelli, 1970; Sengpiel et al., 1999; Stryker & Sherk, 1975; Stryker et al., 1978; Tanaka et al., 2006).

We followed previous work and explored the behaviour of our model with overcomplete representations (Olshausen & Field, 1997; Osindero et al., 2006; Teh et al., 2003). While the representation in the cortex is typically assumed to be greatly overcomplete relative to its thalamic input, the (not unquestioned) estimate of ~ 100 times overcompleteness is too large to simulate. Therefore, it is important that our models exhibit reasonable behaviour as the degree of overcompleteness increases, licensing extrapolation. Indeed, we found that the model results still hold with increasing overcompleteness and also that the proportion of goggle-oriented stimuli that still gives rise to experimentally observed receptive-fields increases with overcompleteness.

The results we showed here are inevitably qualitative, because we could not simulate the actual degree of overcompleteness of the cortex, and we did not have a way of exactly simulating the balance between stimulus and spontaneous activity driven influences on the kitten's cortical development. Thus, we explored different levels of overcompleteness to show the robustness and general qualitative trends of our results within a broad range of parameters.

Another reason we emphasize the qualitative nature of our results is that the data in Tanaka et al. (2006) was collected from six goggle-reared kittens. While all kittens showed the systematic cortical changes mentioned above, there was no clear quantitative statistical relationships between different aspects of the changes. For example, the exact proportion of oriented vs. non-oriented neurons did not have an obvious monotonic relationship with ORI value. It remains unknown whether more data would establish a more consistent quantitative relationship between changes that is captured by our model (some of which are described below), or whether these relationships lie outside the scope of our model.

ICA-based unsupervised learning models embodying the same functional goals as ours have previously been applied to the development of normal cortical organization. Indeed, the $1 \times$ overcomplete version of our model is iden-

tical to traditional ICA (Teh et al., 2003). Our results show that they also well describe the plasticity engendered by (altered) stimuli. In fact, the conditions which we have considered, in which plasticity is driven by goggle-filtered natural scenes, is closer to the biological reality, since development in those previous accounts depends on input (i.e. natural scenes) that is not actually available during those phases of development.

One natural reaction to our study is that an ICA-based model would have no option but to produce the over-representation of RFs aligned with the goggle orientation. However, further consideration reveals a more counter-intuitive nature to these results. The task for all such models is to decompose the statistical distribution of inputs into separate pieces signalled by the output units. In the models, large neural responses signal the presence of rare (low probability) stimuli. Learning therefore captures the statistics of the input by finding RFs such that units would generally only be weakly activated by high probability images, leading to a representation that is sparse. Therefore, a reasonable naïve guess for the form of such RFs would be that they are aligned orthogonal to rather than parallel with GO.

We investigated the properties of such orthogonal RFs and found that while they do in fact produce sparse responses to goggle stimuli, they need to be fat and large in order to capture the statistics of goggle-stimuli structure. These large fat orthogonal RFs produce very un-sparse responses to natural scenes and consequently are poor at representing image data that even consists of a small proportion of natural scenes. Indeed, in our simulations such RFs are more common in the face of extreme proportions of goggle-filtered inputs.

Similarly, one might assume that model neurons with RFs orientated parallel with GO would respond all the time (very un-sparsely) to goggle-filtered stimuli. However, the elongated shape of the RFs parallel to GO actually make them very selective, responding with large activities only for particular classes of goggle-filtered input stimuli, thus preserving their sparse response distributions. Furthermore, RFs that are elongated parallel to GO also produce sparse responses to natural stimuli as well.

Aside from sparseness, the central task for unsupervised learning in this model is to find RFs that effectively represent the stimulus statistics. The various models differ according to the semantics of the units' activities, and the way that their outputs are combined to characterize the overall input distribution. In the POE model we used, the activity of a unit reports the violation of a set of constraints defined by its RFs. The improbability of the stimulus is determined by multiplying together the costs of all the violations. Stimuli rarely violate many common constraints simultaneously, and so the activity of the representational units is sparse. This view of RFs as setting constraints can help illuminate the reason for the characteristics of the model RFs that we find for goggle-filtered images: elongated GO RFs act together to constrain the stimulus to

particular striped patterns. Thus the over-representation of GO RFs serves to capture the striped statistics of the goggle-filtered scenes while maintaining sparse responses.

The POE model has some advantages over more conventional so-called causal generative models (Hyvarinen & Hoyer, 2001; Karklin & Lewicki, 2003; Olshausen & Field, 1997). Most particularly, it better captures the shapes and preferred frequencies of RFs of real neurons (Osindero et al., 2006; Ringach, 2002) than traditional sparse coding and ICA models (though recent work has shown that imposing hard sparseness constraints on sparse coding models can also result in a more realistic distribution of RF shapes; (Osindero et al., 2006; Rehn & Sommer, 2007)). Also, causal generative models are based on the assumption that natural images arise from independent causes. This assumption has its limitations because salient natural scene features often contain clear dependencies. The POE offers an alternative way of capturing input statistics that does not assume independent causes in its input.

Nevertheless, the extent to which the assumptions in different models accurately describe the organization of visual cortex remains unclear. It is likely that other overcomplete models also based on sparse characterizations of the input (Olshausen & Field, 1997; Osindero et al., 2006) will share at least some, if not all, of the results with the POE model discussed here. These results are not intended as a means of distinguishing between different ICA-based models. Instead, we hope our results will serve as an example of the applicability of this larger class of ICA-based models to cortical adaptation. Also, of course, the model is still highly simplified, for instance lacking a clear separation between excitatory and inhibitory units, or the sort of non-linear, complex-cell-like units whose functional roles are slowly becoming understood (Bell & Sejnowski, 1997; Hyvarinen & Hoyer, 2001; Karklin & Lewicki, 2003; Olshausen & Field, 1997; Rehn & Sommer, 2007). It has been suggested (Zhaoping, 2006) that these latter features play only a modest role in capturing input statistics, and so need not necessarily alter the essential structure of our findings.

The imperative for sparsity accounts for the two main aspects of our model that do not seem to have a clear parallel in the experimental data. First, our non-GO oriented model neurons have rather specific dips in tuning at GO in order to enhance sparseness (see [Supplementary Fig. 4](#)). Second, our model produces fewer neurons with best-orientations near, but not exactly, GO. This is because these neurons would not be able to maintain sparse selective responses to GO. One possible source of the difference is that we do not constrain the range of possible changes to the RFs of the model neurons, whereas it is quite likely that the real neurons are not so labile. Indeed, (Tanaka et al., 2006) suggests that two different functional mechanisms are responsible for the changes: instruction (shifts in preferred orientations of neurons towards GO) and selection (decreases in orientation selectivity of neurons that prefer orientations other than GO). One main rationale for selection is that neurons may not be arbitrarily plastic. This

requires them to shift gradually towards having peak tuning at GO and thus could explain why there are not fewer experimental neurons with preferred orientations near, but not at, GO. The requirement for gradual tuning shifts may be due to the significant lateral interactions present among experimental neurons with similar tuning. Constrained plasticity could also limit their ability to have tuning curves that dip specifically at GO. Understanding how best to characterize these constraints is an important direction for future work.

Additional studies involving larger samples of electrophysiological recordings would be the best way of further assessing the validity of this model. A first test would be to determine whether neural responses, including neurons with non-localized, oriented RFs, really do exhibit sparse responses to GO filtered input. An even more complete characterization of neurons' RFs and tuning properties would be needed in order to compare more precisely with the RF characteristics predicted by this model. Most notably, model neurons with RFs aligned to GO have higher peak spatial frequencies (see [Supplementary Fig. 3](#)). Also, the results of our simulation predict an extreme shift in neural representation in the case that goggle-stimuli represent very nearly all of the instructive force guiding development. Under these circumstances, our model predicts RFs that exhibit a marked absence of preferred orientations at goggle-orientation and that are mostly un-localized and un-oriented. These RFs may also have dips in tuning at GO. Unfortunately testing this would require such extreme manipulations that it may not be experimentally feasible.

Functional and mechanistic models are both important for understanding neural systems. Indeed, each has its own advantages and disadvantages. Our account inherits all the limitations of a purely functional account of cortical organization, in particular not suggesting a biologically testable mechanism that prescribes the step-by-step progress of cortical development. Thus, this type of model cannot resolve pertinent issues about how the cortex accomplishes its organization, as is addressed in the debate between instruction vs. selection. Furthermore, whereas mechanistic models can directly be refuted or at least refined by any failure in the face of very unusual inputs, it might seem that a functional model could be protected from this fate on the grounds that it is only intended to describe normal, not abnormal, function. However, one major intent for unsupervised learning models is to provide an account that can address the general capacity of different cortical regions to represent statistical structure in a variety of inputs. Therefore, if such models such as ours had failed in the face of altered input statistics, it would indeed have provided a route towards refinement.

We believe that the applicability of these models to extreme stimulus-driven cortical organization reinforces the insight that nervous system adaptation has functional goals similar to those of innate hard-wired organization. This expands on the domains in which these functional principles are able explain nervous system characteristics.

In the end, a fully coherent view of cortical organization will only be complete when functional models are incorporated and made compatible with more detailed mechanistic models of neural development (Miller et al., 1999, 1989). Additionally, the most challenging direction for future work is to extend both the experimental and theoretical paradigms to richer classes of manipulations to the input distribution. In particular, it is most compelling to treat the adaptations to input statistics shown over the course of months of goggle-rearing as being one end of a whole spectrum, at the other end of which are the millisecond scale changes evident in very short-term adaptation paradigms (Dragoi, Sharma, & Sur, 2000; Teh et al., 2003). Finding a functional explanation that accounts for the changes across this range of time scales is a pressing and important problem.

Acknowledgments

We would like to thank the Gatsby Charitable Foundation for funding and Shigeru Tanaka, Tadashi Yamazaki, Simon Osindero and also Yee Whye Teh for answering numerous questions and extensive discussions and help. We are also grateful to Richard Turner, Iain Murray, Pietro Berkes, Geoffrey Goodhill, Peter Latham, David Mackay and two anonymous reviewers for most helpful comments and advice.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at [doi:10.1016/j.visres.2007.07.023](https://doi.org/10.1016/j.visres.2007.07.023).

References

- Barlow, H. B. (1975). Visual experience and cortical development. *Nature*, 258, 199–204.
- Barlow, H. B. (1981). Critical limiting factors in the design of the eye and visual cortex. The Ferrier lecture, 1980. *Proceedings of the Royal Society, London B*, 212, 1–34.
- Bell, A. J., & Sejnowski, T. J. (1997). The “independent components” of natural scenes are edge filters. *Vision Research*, 37, 3327–3328.
- Blakemore, C., & Cooper, G. F. (1970). Development of the brain depends on visual environment. *Nature*, 228, 477–478.
- Blakemore, C., & Van Sluyters, R. C. (1975). Innate and environmental factors in the development of the kitten's visual cortex. *Journal of Physiology*, 248, 663–716.
- Buisseret, P., & Imbert, M. (1976). Visual cortical cells: Their developmental properties in normal and dark-reared kittens. *Journal of Neurophysiology*, 255, 511–525.
- Crair, M., Gillespie, D., & Stryker, M. (1998). The role of visual experience in the development of columns in cat visual cortex. *Science*, 279, 566–570.
- De Valois, R. L., Albrecht, D. G., & Thorell, L. G. (1982). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research*, 22, 545–559.
- Dragoi, V., Sharma, J., & Sur, M. (2000). Adaptation-induced plasticity of orientation tuning in adult visual cortex. *Neuron*, 28, 287–298.
- Frégnac, Y., & Imbert, M. (1984). Development of neuronal selectivity in primary visual cortex of cat. *Physiological Reviews*, 325–434.

- Godecke, I., & Bonhoeffer, T. (1996). Development of identical orientation maps for two eyes without common visual experience. *Nature*, *379*, 251–254.
- Goodhill, G. J. (1993). Topography and ocular dominance: A model exploring positive correlations. *Biological Cybernetics*, *69*, 109–118.
- Goodhill, G. J., & Willshaw, D. J. (1990). Application of the elastic net algorithm to the formation of ocular dominance stripes. *Network*, *1*, 41–59.
- Hinton, G. E., & Ghahramani, Z. (1997). Generative models for discovering sparse distributed representations. *Philosophical Transactions Royal Society B*, *352*, 1177–1190.
- Hirsch, H. V., Leventhal, A. G., McCall, M. A., & Tieman, D. G. (1983). Effects of exposure to lines of one or two orientations on different cell types in striate cortex of cat. *Journal of Physiology*, *337*, 241–255.
- Hirsch, H. V., & Spinelli, D. N. (1970). Visual experience modifies distribution of horizontally and vertically oriented receptive fields in cats. *Science*, *168*, 869–871.
- Hubel, D. H. (1988). *Eye, brain, and vision*. New York: WH Freeman.
- Hubel, D. H., & Wiesel, T. (1965). Binocular interaction in striate cortex of very young, visually inexperienced kittens. *Journal of Neurophysiology*, *28*, 1041–1059.
- Hubel, D. H., Wiesel, T., & LeVay, T. N. (1977). Plasticity of ocular dominance columns in monkey striate cortex. *Philosophical Transactions of the Royal Society of London [Biology]*, *278*, 377–409.
- Hubel, D. H., & Wiesel, T. N. (1963). Receptive fields of cells in striate cortex of very young, visually inexperienced kittens. *Journal of Neurophysiology*, *26*, 994–1002.
- Hyvarinen, A., & Hoyer, P. O. (2001). A two-layer sparse coding model learns simple and complex cell receptive fields and topography from natural images. *Vision Research*, *41*, 2413–2423.
- Jones, J., & Palmer, L. (1987). An evaluation of the two-dimensional gabor filter model of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, *58*, 1233–1258.
- Karklin, Y., & Lewicki, M. (2003). Learning higher-order structures in natural images. *Network-Computation in Neural Systems*, *14*, 483–499.
- Lewicki, M. (2002). Efficient coding of natural sounds. *Nature Neuroscience*, *5*, 356–363.
- Li, Z., & Atick, J. J. (1994). Towards a theory of striate cortex. *Neural Computation*, *6*, 127–146.
- Linsker, R. (1990). Perceptual neural organization: Some approaches based on network models and information theory. *Annals Review of Neuroscience*, *13*, 257–281.
- Miller, K. D., Erwin, E., & Kayser, A. (1999). Is the development of orientation selectivity instructed by activity. *Journal of Neurobiology*, *41*, 44–57.
- Miller, K. D., Keller, J. B., & Stryker, M. P. (1989). Ocular dominance column development: Analysis and simulation. *Science*, *245*, 605–615.
- Miyashita, M., Kim, D. S., & Tanaka, S. (1997). Cortical direction selectivity without directional experience. *Neuroreport*, *8*, 1187–1191.
- Miyashita, M., & Tanaka, S. (1992). A mathematical model for the self-organization of orientation columns in visual cortex. *Neuroreport*, *3*, 69–72.
- Movshon, J. A., & Blakemore, C. (1974). Functional reinnervation in kitten visual cortex. *Nature*, *251*, 504–505.
- Olshausen, B. A., & Field, D. J. (1997). Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Research*, *37*, 3311–3325.
- Osindero, S., Welling, M., & Hinton, G. (2006). Topographic product models applied to natural scene statistics. *Neural Computation*, *18*, 381–414.
- Parker, A. J., & Hawken, M. J. (1988). Two-dimensional spatial structure of receptive fields in monkey striate cortex. *Journal of the Optical Society of America A*, *5*, 598–605.
- Rehn, M., & Sommer, F. T. (2007). A network that uses few active neurons to code visual input predicts the diverse shapes of cortical receptive fields. *Journal of Computer Neuroscience*, *22*, 135–146.
- Ringach, D. L. (2002). Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *Journal of Neurophysiology*, *88*, 455–463.
- Ruthazer, E. S., & Stryker, M. P. (1996). The role of activity in the development of long-range horizontal connections in area 17 of the ferret. *Journal of Neuroscience*, *15*, 7253–7269.
- Sengpiel, F., Stawinski, P., & Bonhoeffer, T. (1999). Influence of experience on orientation maps in cat visual cortex. *Nature Neuroscience*, *2*, 732.
- Stryker, M., & Sherk, H. (1975). Modification of cortical orientation selectivity in the cat by restricted visual experience: A reexamination. *Science*, *190*, 904–906.
- Stryker, M., Sherk, H., Leventhal, A. G., & Hirsch, H. V. (1978). Physiological consequences for the cat's visual cortex of effectively restricting early visual experience with oriented contours. *Journal of Neurophysiology*, *41*, 896–909.
- Swindale, N. V. (1996). The development of topography in the visual cortex: A review of models. *Computation in Neural Systems*, *7*, 161–274.
- Tanaka, S., Miyashita, M., & Ribot, J. (2004). Roles of visual experience and intrinsic mechanism in the activity-dependent self-organization of orientation maps: Theory and experiment. *Neural Networks*, *17*, 1363–1375.
- Tanaka, S., Ribot, J., & Tani, T. (2006). Orientation-restricted continuous visual exposure induces marked reorganization of orientation maps in early life. *NeuroImage*, *30*, 462–477.
- Teh, Y., Welling, M., Osindero, S., & Hinton, G. (2003). Energy-based models for sparse overcomplete representations. *Journal of Machine Learning Research-Special Issue on ICA*, 1235–1260.
- Thompson, I. (1997). Cortical development: A role for spontaneous activity. *Current Biology*, *1*, R324–R326.
- van Hateren, J. H., & van der Schaaf, A. (1998). Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society of London Series B Sciences*, *265*, 359–366.
- von der Malsburg, C. (1973). Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik*, *14*, 85–100.
- Wiesel, T. N. (1982). Postnatal development of the visual cortex and the influence of environment. *Nature*, *299*, 583–591.
- Willshaw, D. J., & von der Malsburg, C. (1979). A marker induction mechanism for the establishment of ordered neural mappings: Its application to the retinotectal problem. *Philosophical Transactions of the Royal Society of London [Biology]*, *287*, 203–243.
- Willshaw, D. J., & von der Malsburg, C. (1976). How patterned neural connections can be set up by self-organization. *Philosophical Transactions of the Royal Society of London [Biology]*, *194*, 431–445.
- Zhao, L. (2006). Theoretical understanding of the early visual processes by data compression and data selection. *Network*, *17*, 301–334.