

## Space and time in visual context

Odelia Schwartz<sup>\*‡</sup>, Anne Hsu<sup>§</sup> and Peter Dayan<sup>§</sup>

**Abstract** | No sensory stimulus is an island unto itself; rather, it can only properly be interpreted in light of the stimuli that surround it in space and time. This can result in entertaining illusions and puzzling results in psychological and neurophysiological experiments. We concentrate on perhaps the best studied test case, namely orientation or tilt, which gives rise to the notorious tilt illusion and the adaptation tilt after-effect. We review the empirical literature and discuss the computational and statistical ideas that are battling to explain these conundrums, and thereby gain favour as more general accounts of cortical processing.

### Saccade

A rapid eye movement (with speeds of up to 800° per second) that brings the point of maximal visual acuity — the fovea — to the image of interest.

“No man”, and concomitantly no sensory stimulus, “is an island.” That is, the perception of, and neurophysiological responses to, a target input depend strongly on both its spatial context (what surrounds a given object or feature) and its temporal context (what has been observed in the recent past). Psychophysical evidence for contextual effects is particularly widespread in vision<sup>1</sup>, including motion<sup>2</sup>, brightness<sup>3,4</sup>, orientation<sup>5</sup>, blur<sup>6</sup> and faces<sup>7,8</sup>. However, contextual influences also extend to other modalities, such as audition<sup>9</sup> and somatosensory processing<sup>10</sup>. Neurophysiological evidence for contextual influences is most extensive in the early visual processing of orientation and motion<sup>11–14</sup>, and in the whisking activity in the rodent somatosensory cortex<sup>15</sup>; but context is also likely to influence the neural processing of many other attributes, including colour<sup>16</sup> and border ownership<sup>17</sup>.

Here, we focus on the contextual modulation of visual orientation (see the examples in FIG. 1), a topic of broad interest that has given rise to many somewhat contentious results and theories. With some recent exceptions<sup>18–25</sup>, spatial and temporal contextual modulation of orientation processing have typically been treated separately, even being referred to by different names (for example, perceptual illusions and after-effects; non-classical receptive fields and adaptation, respectively). Nevertheless, despite quite different demands on the neural substrate (for example, horizontal intra-areal interactions for spatial context, and memory for temporal context), they are closely related both functionally and in terms of their impact on vision. We describe some of the many results that show the apparent psychophysical and electrophysiological parallels between the effects of spatial and temporal context (FIG. 2), and provide an overview of computational analyses of contextual processing that bear on this likeness.

A central issue for this Review is whether there is a functionally important commonality between spatial and temporal context that underlies the similarities between them that are evident in FIG. 2, or whether these similarities are mere coincidence. Many, perhaps even most, functional characterizations of cortical visual areas suggest that their processing either implicitly or explicitly reflects the statistical structure of the visual inputs. One obvious source of commonality would be if these statistics themselves exhibited similar properties. In fact spatial and temporal context can be seen as orthogonal slices through the full spatio-temporal structure of natural scenes and, indeed, it turns out that these slices do have a crucial shared property arising from the fact that objects are typically smooth and move or change slowly — inputs that are nearby in either space or time are typically similar (in some cases even taking saccades into account). We will see that many treatments of context make explicit or implicit appeal to these properties. Despite this common underlying foundation, the functional goals of contextual effects are still hotly debated. The understanding of contextual effects will have important implications for understanding normal processing of the spatio-temporal world, as well as for developing engineering and medical applications tuned to a complex changing sensory environment, such as adaptive visual and hearing aids.

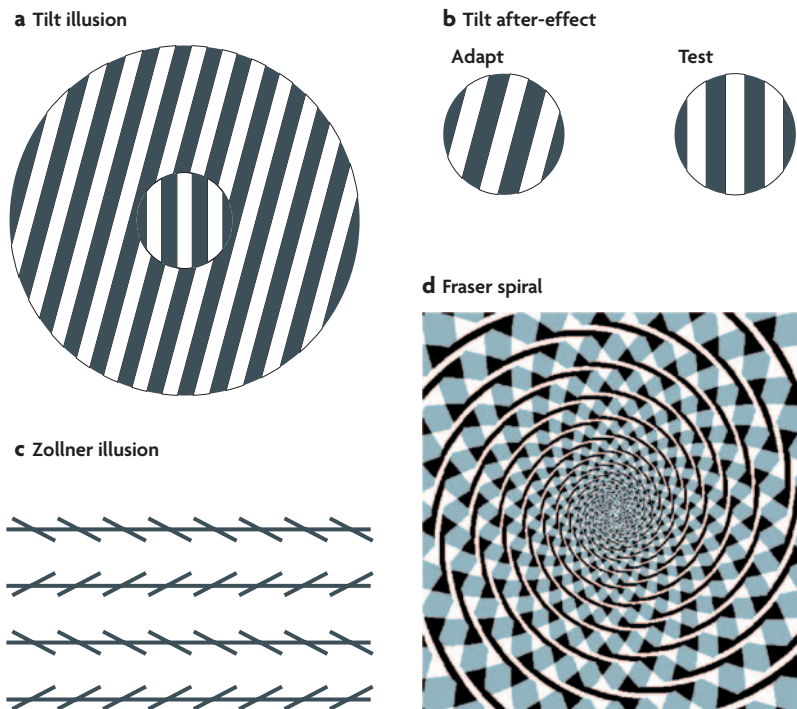
In the next section, we review experimental data describing how single cells and perceptions are affected by spatial and temporal context, and describe models that link these two datasets. We then consider the relationship between these changes and the statistical properties of natural scenes. We end by discussing some of the many gaps in the data and in our theoretical understanding of contextual effects.

<sup>\*</sup>Albert Einstein College of Medicine, Jack and Pearl Resnick Campus, 1300 Morris Park Avenue, Bronx, New York 10461 (718) 430–2000, USA.

<sup>‡</sup>Howard Hughes Medical Institute, and The Salk Institute, 10010 North Torrey Pines Road, La Jolla, California 92037, USA.

<sup>§</sup>Gatsby Computational Neuroscience Unit, UCL, London, UK.

Correspondence to O.S. email: odelia@salk.edu doi:10.1038/nrn2155



**Figure 1 | Examples of contextual tilt.** **a** | Tilt illusion. The presence of a surround spatial context tilted 15° clockwise from vertical causes the vertical centre target to appear repulsed away from the context, that is, tilted counter clockwise. The context and the target are separate in space but overlap in time. **b** | Tilt after-effect. The temporal context (adapt) is presented before the target (test) and this causes the vertical target to appear repulsed away from the context orientation. To induce the tilt after-effect, gaze at the left frame for at least 30 seconds and then quickly fixate upon the right frame. The context and test overlap in space but not in time. **c, d** | Two well-known illusions, which partly depend on contextual interactions of orientation. **c** | Zollner illusion. The four horizontal lines are parallel but appear tilted, owing to the shorter lines that are overlaid at an angle on top of each horizontal line and act as a context. **d** | Fraser spiral illusion<sup>200</sup>. The arc segments appear to form a spiral, but they are actually a series of concentric circles. The illusion is due to the tilted context pattern in the circles, which is augmented by the checkerboard pattern.

**Contextual effects on tilt**

We can approach the effect of context on tilt by measuring how the electrophysiological responses of populations of neurons to local orientation stimuli are changed by context (FIG. 2c,d), and by investigating the effects of context on the perception of orientation (FIG. 2a,b). Below we use mechanistic models of orientation tuning (FIG. 3) to assess whether the changes seen in population responses are consistent with experimental findings (FIG. 2). Many experiments indicate marked similarities between spatial and temporal context; however, although a good proportion are generally accepted, controversy remains even within each of the domains. Some of the controversy could arise from experimental differences, such as: species; anesthetized versus awake subjects; stimulus configurations; timescales of adaptation; or the possibility that certain of the stimuli actually impinge on the classical receptive field, thus rendering the effects not strictly contextual.

**Orientation tuning**

The property of visual neurons to only respond to stimuli (images) with a certain orientation or tilt (for example, vertically orientated bars).

**Tuning curve**

A tuning curve to a feature (such as orientation) is the curve describing the average response of a neuron as a function of the feature values.

**Rotation invariance**

When each input angle is treated in the same way; that is, when the input rotates, the output rotates.


**Electrophysiology.** Many neurons in the primary visual cortex (V1) are selective for local orientation within spatial receptive fields<sup>26</sup>, having mean firing rates that are conventionally modelled as coming from uni-modal, Gaussian-like tuning curves that are defined by three parameters: height, width and the ‘preferred orientations’ at which they peak. Local orientation is represented by a population of such neurons<sup>27–29</sup>, the orientation preferences of which are evenly distributed across the range, from –90° to 90°. The first two columns in FIG. 3a show the tuning curves of such a neuron population and their mean activities induced by presenting a bar at an angle of 20°. Note that in the absence of context, we assume homogenous processing of orientation (rotation invariance) and do not incorporate preference for absolute cardinal orientation axes<sup>30</sup>.

Contextual stimuli lead to suppression or enhancement of firing rates, broadening or sharpening of tuning widths, and repulsive or attractive shifts in preferred orientation, all depending on the relationship between the orientation of the contextual stimuli and the preferred orientation of the neurons. The first column in FIG. 3b–d illustrates archetypal changes to the population tuning curves that occur in response to a context (either spatial or temporal) orientated at 0°. In reality, the whole collection of tuning curves is rarely recorded in single experiments and many factors can affect the results.

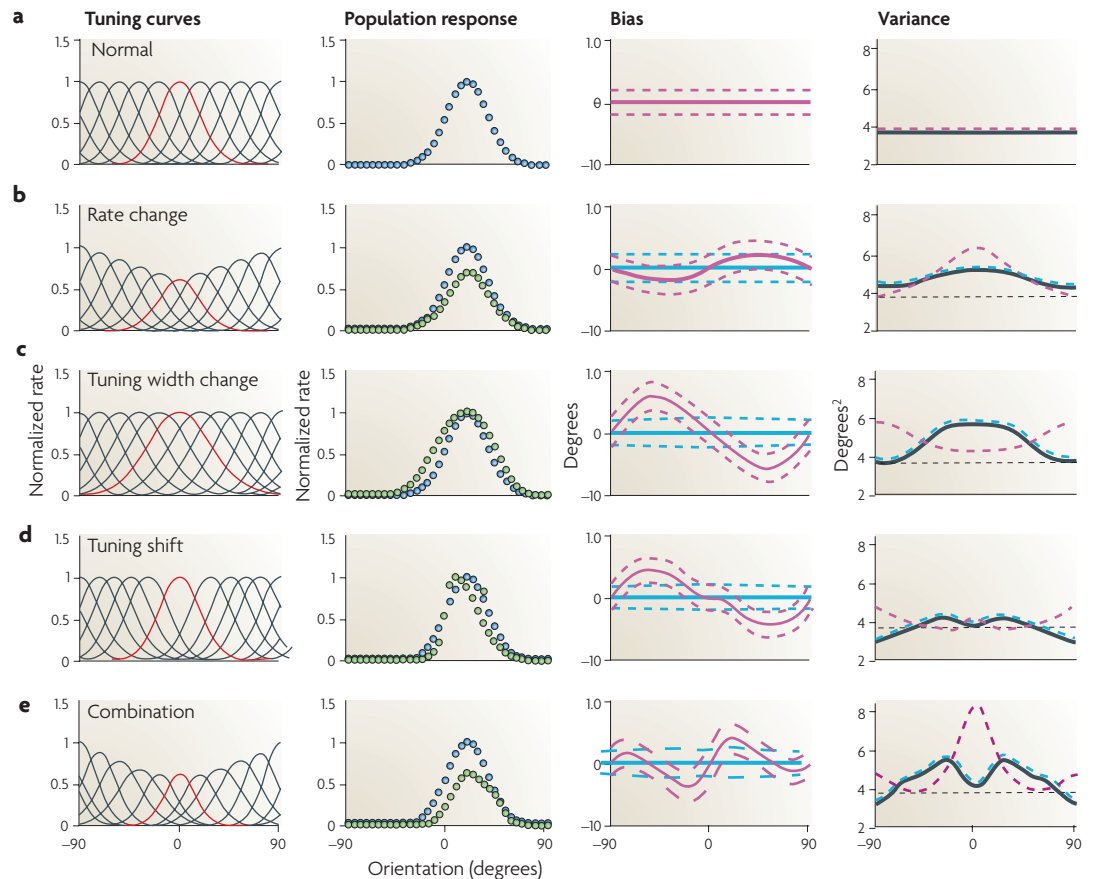
FIGURE 3b models tuning curve suppression, which is maximal for tuning curves that peak at the context orientation 0°. Suppression of neural responses has been widely observed in spatial contextual studies in the monkey<sup>31–36</sup> and cat<sup>24,37</sup>, and in temporal contextual studies in the monkey<sup>38</sup> and cat<sup>24,39,40</sup>. Contextual stimuli can also facilitate responses in neurons for which the preferred stimuli are orthogonal to the spatial or temporal context in both monkey and cat<sup>31,40,41</sup>, although there is variability across the population of cells<sup>31,40</sup>. FIGURE 2c shows example experimental data comparing suppression as a function of orientation difference between context and target, for spatial<sup>37</sup> and temporal<sup>39</sup> cases. Note that the data, unlike the model, do not specify the full tuning curves, but rather the amount of suppression at a particular location on the tuning curve for each orientation difference. Empirically, suppression is sometimes asymmetrical, affecting only the half of the tuning curve that contains the contextual orientation<sup>42</sup>, rather than being symmetrical as modelled here.

FIGURE 3c models increases in tuning curve widths (by increasing the Gaussian standard deviations), which are greatest for the neuron with preferred orientation at 0°, and decreases in tuning curve widths, as the difference between the preferred and contextual orientations grows. Spatial context has not been shown to produce systematic tuning width changes<sup>43</sup>. However, in (temporal) adaptation experiments, cells near the adapted orientation showed broadening of this sort<sup>39</sup>.

FIGURE 3d shows repulsive shifts of population tuning curves in the model away from 0°. The shifts are greatest for tuning curves that peak in the vicinity of the contextual orientation. Indeed, many experiments have shown that the preferred orientation of tuning curves



**Figure 2 | Example experimental data for spatial and temporal context effects.** The two columns show the similarity of spatial and temporal context effects on perception and neuronal response. **a** | Tilt illusion (spatial, left) and tilt after-effect (temporal, right) biases as a function of the angle difference between context and target orientations. Positive bias indicates repulsion and negative bias indicates attraction. **b** | Orientation discrimination thresholds for spatial and temporal data. **c** | Response suppression in cortical neurons as a function of the angle difference between the target and its context. The spatial data apply to 9 out of 32 neurons which showed 'iso-oriented suppression'. The temporal data are for suppression on the near flank of the tuning curve, which is similar but not identical in behaviour to suppression at the original preferred orientation. **d** | Repulsive tuning shifts in primary visual cortex neurons for space and time, for context-target differences of 15° (spatial context) and 15° and 30° (temporal context). **e** | Contrast-response functions for space and time in cortical neurons, either without a context (circles), or with a context orientation equal to the test (triangles). Modified, with permission, from: **a** REF. 19 © (2000) Royal Society of London; **b** (right) REF. 54 © (1985) American Institute Of Physics; **c** (left) REF. 37 © (1997) Springer; **c** (right) REF. 39 © (2000) Cell Press; **d** REF. 24 © (2005) Institute Of Physics Publishing; **e** (left) REF. 31 © (2002) American Physiological Society; **e** (right) REF. 40 © (2006) American Physiological Society. Data for **b** (left) from REF. 146.



**Figure 3 | Mechanistic population model for spatial and temporal context. a** | Tuning curves in the absence of context. **b–e** | The tuning parameters of the neuronal population are modified, given a context at  $0^\circ$ . Red indicates the tuning curve at the orientation equal to that specified by the context. **b** | Tuning curve suppression occurs when the preferred orientation is equal to that of the context, and gradually changes back to baseline as the preferred orientation moves further away from that of the context. **c** | Tuning curve widths decrease for neurons whose preferred orientations are further away from the contextual orientation, with the greatest shifts for values starting closest to the context. **d** | The preferred tuning values of neurons shift away from the contextual orientation, with the greatest shifts for values starting closest to the context. **e** | Example combination of tuning curve changes, based on the model of Clifford *et al.* (REF. 19). Columns left to right: tuning curves; population response (original blue; contextually modulated green) to a target orientation of  $20^\circ$ ; bias, assuming maximum-likelihood decoding, either suffering from the ‘coding catastrophe’ (magenta) or the true maximum likelihood (cyan), correcting for the new tuning properties of the neurons. Dashed lines show the standard deviation of the two decoded estimates; and variability of the estimators (magenta and cyan) and the inverse of the Fisher information (black). Rows (**b–e**) also display the baseline for the inverse of the Fisher information without any tuning manipulations (black dotted lines).

of V1 neurons in cats and monkeys shifts away from the orientation of the spatial context<sup>24,43</sup> or from the adapted orientation<sup>24,38,39</sup>. In the case of adaptation, such shifts have been reported after lengths of time as short as milliseconds<sup>24,38</sup> and have been documented to increase in strength for time scales on the order of minutes<sup>39</sup>. Figure 2d shows sample experimental data<sup>24</sup> that compare the distributions of tuning curve shifts when contextual and target orientations differ by  $15^\circ$  (spatial case) and by  $15^\circ$  and  $30^\circ$  (temporal context). Repulsive shifts in speed selectivity have also been noted in the presence of spatial context stimuli<sup>44</sup>. However, repulsive shifts might not be ubiquitous. Indeed, attractive shifts towards the contextual stimulus have been found in monkeys for cortical areas beyond V1, specifically in the visual motion processing middle temporal area in response to adaptation to the direction of motion<sup>42</sup>.

**Psychophysics.** The most straightforward examples of perceptual effects in orientation are two biases, known respectively as the tilt illusion, which is induced by spatial surround stimuli (FIG. 1a), and the tilt after-effect, which is caused by temporal adaptation to uniformly orientated textures (FIG. 1b) and which has been studied for time-scales ranging from seconds to minutes<sup>45</sup>. Sample data for both spatial and temporal context<sup>19</sup> are shown in FIG. 2a. The two biases depend in strikingly similar ways on the relative orientations of the target and the contextual stimuli<sup>5,18,46</sup>. Tilts that are close to the contextual orientation are ‘repelled’, that is, perceived as tilting further away from the contextual orientation (the direct effect), whereas tilts that differ greatly from the contextual orientation are attracted (the indirect effect, which is weaker).

Such biases have been widely documented for other stimulus attributes: the indirect effect has been observed

for motion stimuli<sup>47</sup>, and the direct effect has been reported for motion, colour, faces and blur<sup>2,6–8,20,47–49</sup>. Repulsion and attraction are also affected by factors such as timescale and geometry, as discussed later.

Discrimination thresholds for orientation are also altered in the presence of spatial context stimuli<sup>33,36,50,51</sup>, and temporal context stimuli ranging from milliseconds<sup>52</sup> to seconds<sup>53</sup> and minutes<sup>1,54–56</sup>, although some of these results are controversial. A notable effect that is qualitatively similar for both spatial and temporal contexts is that discrimination is impaired if the orientation of context stimuli is different from the orientation of the test target by 7° to 45° (REFS 50,51,54,56 and FIG. 2b). Contextual stimuli that are orthogonal to the orientation of the target stimulus seem to produce little or no impairment in discrimination; they might even aid discrimination in the case of adaptation<sup>39,52,56</sup>. Some studies<sup>50,54,56</sup>, although not others<sup>36,51,57</sup>, have found a modest improvement in discrimination if context tilts were equal to the target orientation, for both spatial and temporal contexts.

**Decoding.** In order to link the electrophysiological and psychophysical results, we need to understand how neuronal population responses are decoded to produce perceptual effects. Perceptual decisions involve processing in an extensive and incompletely understood hierarchy of cortical areas. However, there are many models for decoding the sorts of population codes shown in FIG. 3, which are at least partly validated by psychological and neurophysiological observations.

We focus on maximum-likelihood decoding (although other methods yield similar results<sup>58</sup>): each stimulus is associated with a likelihood of responses, that is, the probability of the neural responses given the stimulus ( $P(\text{response}|\text{stimulus})$ ). The responses are based on the neurons' (context-free) tuning curves; the perceived stimulus is the one that maximizes this likelihood of responses. Operating under the Poisson spiking neuron model (we assume that neurons in the population are independent, although correlations between neurons might also affect the result; for a recent review see REF. 59), the third column of FIG. 3a, which averages over  $10^5$  trials, shows that decoding is normally bias-free. Statistical considerations dictate the potential quality of a representation of a stimulus. This can be quantified by the Fisher information<sup>27–29</sup>, which, in regular cases such as this, is inversely related to the minimum threshold for discriminating two orientations or the minimum variance of the decoder. Column four of FIG. 3a shows the inverse of the Fisher information, which, in the absence of contextual stimuli, does not depend on the target orientations (and therefore neither does the threshold for discrimination).

Electrophysiological and perceptual data on contextual effects are conventionally linked by what we call a coding 'catastrophe' or decoding ambiguity<sup>60</sup>: downstream decoding mechanisms are assumed to be unaware of the changes in tuning caused by contextual stimuli. That is, they simply receive the changed input

from the previous layer, and do not 'know' that changes have occurred due to the context and therefore err (for instance, in reporting absolute orientation) when such changes arise. We will later see that functional advantages might be associated with these errors.

Assuming that such a coding catastrophe takes place, the third column in FIG. 3b–d shows the biases in perceived orientation as a function of the presented orientation for the anomalous maximum-likelihood decoder (which fails to correct for the changes). Each of the physiological phenomena of suppression, broadening and shifting can lead to perceptual repulsion and/or attraction. All of these effects could be combined in various ways to lead to the observed perceptual phenomena<sup>19,42,58,61</sup>. For example, FIG. 3e shows one possible combination<sup>19</sup> that gives rise to both direct and indirect biases (compare with FIG. 2a).

The dashed lines in the graphs in the third column show the standard deviations (over  $10^5$  trials) of the maximum likelihood estimator, which suffers from the coding catastrophe, and of the true maximum likelihood estimator; the equivalent lines in the fourth column show the variances. The variability of the true maximum likelihood decoder always follows the Fisher information. The different tuning curve manipulations have different impacts on this variability; note how the inverse of the Fisher information in the last columns of FIG. 3d,e bear some qualitative resemblance to the threshold data in FIG. 2b.

One reason for biased orientation inference might be that it is accompanied by improved discriminability. Indeed, it is a standard statistical manoeuvre to trade off bias for variance in estimation, and this can lead to a lower overall mean square estimation error. However, it appears that for the maximum likelihood estimator, which suffers from the coding catastrophe, all the tuning curve manipulations due to context (FIG. 3b–e) result in a larger mean square estimation error.

**Computation.** This rather conventional review of the electrophysiological and perceptual data presents two challenges to functional views of visual processing, namely why contextual stimuli should induce perceptual biases and altered discriminability, and why these perceptual effects, as well as the more general wealth of experimental results exemplified in FIG. 2, should be so similar for spatial and temporal contexts. To answer these questions, we have to elucidate computational frameworks for understanding visual processing, and the role of context in these frameworks. The coding catastrophe is one key to the puzzle — why the visual system might organize an adaptive change to tuning curves only to then seemingly ignore those changes in downstream processing. We discuss the range of functional approaches for contextual processing, and argue that they variously address, embrace and evade the coding catastrophe, although none of the answers they provide is yet completely compelling.

It is possible that estimating the orientations in the way they appear in the experiments might not be a relevant goal. For instance, it has been suggested that

**Discrimination threshold**

The smallest difference between two visual stimuli (for example, vertical versus tilted bars) that can reliably (that is, with a given probability of error) be reported by an observer.

**Population code**

Sensory events that are encoded by neuronal populations rather than by individual neurons.

**Poisson spiking neuron**

A simple model neuron for which the number of spikes emitted in a given time is Poisson distributed about a mean firing rate. Spikes are assumed to be independent both in time and across neurons.

**Fisher information**

Measures how quickly the likelihood of the population responses changes with stimulus parameters, and thereby provides a decoder-independent quantification of the potential accuracy of decoding (the Cramér-Rao lower bound).

**Mean square estimation error**


Estimation error can be quantified by the squared difference between the (population-based) estimate and its true value. The mean of this over trials is one measure of the accuracy of an estimate.

**Efficient coding**

When information is coded in an efficient and non-redundant manner, for instance, when the outputs of neurons in the population are statistically independent.

**Bayes (Bayesian approach)**

A statistical method that allows the use of prior information to evaluate the posterior probabilities of different hypotheses.



**Figure 4 | Statistics of natural movies for time and space.** **a** | Gathering image statistics<sup>24</sup>. Magenta corresponds to the target patch and cyan to context patches. For spatial statistics, the context is given by patches for eight surrounding spatial positions (statistics are averaged over all spatial positions in each frame). For temporal statistics, a single context patch of the neighbouring time frame of the movie in the same spatial location is collected. Actual statistics were gathered from the Catcam movie database<sup>201</sup>, using 6x6 pixel patches, and model-orientated filters<sup>202</sup>. Images of each frame were pre-processed by subtracting the mean and dividing by the standard deviation. **b** | Joint conditional statistics for target ‘complex cell’ horizontal filter responses, given a context filter response<sup>79,203</sup>. Intensity is proportional to the bin counts, except that each column is independently re-scaled to fill the range of intensities. Note that when the context response is near zero, the target response is also near zero; for larger values of the context response, the width of distribution of the target response increases. **c** | Orientation statistics (also see REF. 24). For all patches of context that were horizontal, the graphs depict the probability that the target patch has a given orientation. Orientation of a patch is determined by computing the sum of squares of even and odd responses for 16 orientated filters, and taking the maximum over all orientations.

the key computation for the primary visual cortex is in fact the allocation of salience to image locations based on statistical inhomogeneities in space<sup>62</sup> or time (for example, favouring novelty<sup>63</sup>), licensing estimation bias in favour of boosted salience. Equally, it could be argued that estimation is less important than discrimination, and therefore biases in estimation are of no importance compared with improvements in discriminability. However, the improvements in discriminability are arguably modest, and for particular contextual orientations there is even a deterioration in performance. Compared with the sheer magnitude of the perceptual biases, it has therefore been suggested that discriminability may not be the main goal<sup>64</sup>.

After describing the statistical core that is common to many functional treatments of visual processing, we discuss in detail the two collections of accounts, namely efficient coding and Bayesian modelling, that have been most popular as models of contextual effects.

**Image statistics.** Crudely speaking, the job for the areas of the brain that process the visual world is to take the complete spatio-temporal input  $I_s(t)$  up to time  $t$ ; that is, the activation of all pixels (or small patches, or orientations)  $s$  at every spatial location up to this time, and extract information about the nature and significant contents of the scene at time  $t$ . Most computational models start from the premise that the way the cortex decomposes images is influenced by the regularities inherent in the statistics of  $I_s(t)$  associated with the input scenes.

The statistical properties of typical natural scenes or movies have been extensively analysed<sup>65–70</sup>. Although the full distribution over natural movies  $P(I)$  eludes exact description, it is known that most objects in the visual world have large spatial and temporal footprints<sup>71,72</sup>, which implies that there are substantial low- and high-order correlations between  $I_{s_1}(t_1)$  and  $I_{s_2}(t_2)$  for inputs  $s_1$  and  $s_2$  that are spatially near to each other and times  $t_1$  and  $t_2$  that are close. This means that the temporal context ( $I(t_1)$  for  $t_1 < t_2$ , summing over space), and/or spatial context ( $\{I_{s_1}\}_{s_1 \in S_1}$  for some subset of pixels  $S_1$  in the neighbourhood of  $s_2$ , summing over time) will induce strong (and similar) expectations about the value  $I_{s_2}(t_2)$  (REFS 73–75), with information accruing both dynamically, as in a Kalman filter<sup>76</sup> and statically, as in a Markov random field<sup>77</sup>. Subjects’ knowledge of these correlations is evident in their ability to replace missing pixels in digital images<sup>78</sup> based only on neighbourhood information.

It has also proved possible to characterize local orientation in small patches of natural scenes (FIG. 4a). For instance, if natural scenes are filtered through receptive fields that resemble the orientation tuning of cortical simple cells, then the filter outputs at neighbouring spatial or temporal locations and for similar orientation preferences (FIG. 4b) are correlated, and exhibit higher order statistical coordination<sup>79–84</sup>.

Alternatively, we can examine the dominant orientations in two small patches of images that are neighbours in space or time, for instance, by considering a

## Kalman filter

A recursive formulation that estimates the present outcome dynamically in time, based on prior information and noisy measurements.

## Markov random field

An undirected graphical model that represents statistical dependencies between a set of variables. The Markov property is that a variable associated with one location in the image is only directly influenced by variables associated with neighbouring locations.

## Linear (or second order) de-correlation

Random variables are de-correlated if the off-diagonal elements of their covariance matrix (representing the second order statistics) are equal to zero. De-correlation is in general a weaker requirement than independence, because higher order statistics may still exhibit dependencies.

## Divisive normalization

Strictly speaking, when (for example) sum output across a population is kept constant by dividing each response by a (trial-dependent) quantity. Looser versions model gain control mechanisms in V1 and elsewhere.

## Gain control

When the (for example) sum output across a population is used to adjust the gain to an appropriate level for a range of input signal levels, with higher signal levels resulting in higher gain and reduced response. Stricter versions are denoted divisive normalization.

## Anti-Hebbian learning

A learning rule whereby whenever two units or neurons are active simultaneously, the effective connection between them becomes less excitatory or more inhibitory.

## Bayesian inference

Inference according to the standard laws of probability, notably including Bayes theorem. Conclusions are based on posterior distributions arising from combining observations (as probabilistic likelihoods) with prior information.

population of orientated filters as above, but choosing the physical orientation that corresponds to maximal activity in the population for each patch<sup>24,85–88</sup>. Similar orientations are notably more likely than different orientations in nearby spatial or temporal patches (FIG. 4c and REF. 24). These statistics will be affected by saccades and other eye movements. Unfortunately, there has only been limited investigation of this<sup>88</sup>, with the data suggesting that movements are likely to be made to image patches with both similar and opposite orientations. Although most adaptation experiments are performed under conditions of fixation, it is interesting to note that subsequent eye movements may affect the strengths of the after-effects<sup>89,90</sup>.

The essential kinship between spatial and temporal statistics answers the question about similarity; we argue below that its form might also explain the biases. Of course, there are also some key differences between spatial and temporal contexts, which arise from the larger number of spatial than temporal dimensions, the constraints of temporal causality and the effects of spatial geometry.

**Computational models.** There are two main classes of computational accounts, described in detail in the next sections, which invoke principles of either efficient coding or Bayesian probabilistic inference and modelling. In efficient coding terms, the context-mediated expectations imply that only some values of  $I_{s_2}(t_2)$  are likely, and these are the values on which coding and representational resources should be concentrated. The strong expectations arise from substantial redundancy in the input; efficient coding suggests that the contextual effects are a strategy for mitigating this inefficiency through de-correlation. This can be formalized using information theory, based on the conditional distribution  $P(I_{s_2}(t_2)|I_{s_1}(t_1))$ .

In Bayesian terms, one view is that the expectations that are inherent in this conditional distribution should be treated as a form of prior evidence, to be combined with information associated with  $I_{s_2}(t_2)$  itself to make inferences about the aspects of the visual world that determine  $I_{s_2}(t_2)$ . In such inference models, noisy and ill-posed inferences about  $I_{s_2}(t_2)$  can be partially resolved by incorporating relevant data from the context. A richer but more complicated view is provided by generative models. These suggest that the cortex builds a form of surrogate model of the way in which images are generated, and uses it to extract structure in the input. As image structures are spatially and temporally extended, the entities in such cortical models have to mix target and contextual stimuli in a statistically appropriate manner. Contextual effects then arise from the nature of this mixing.

## Efficient coding

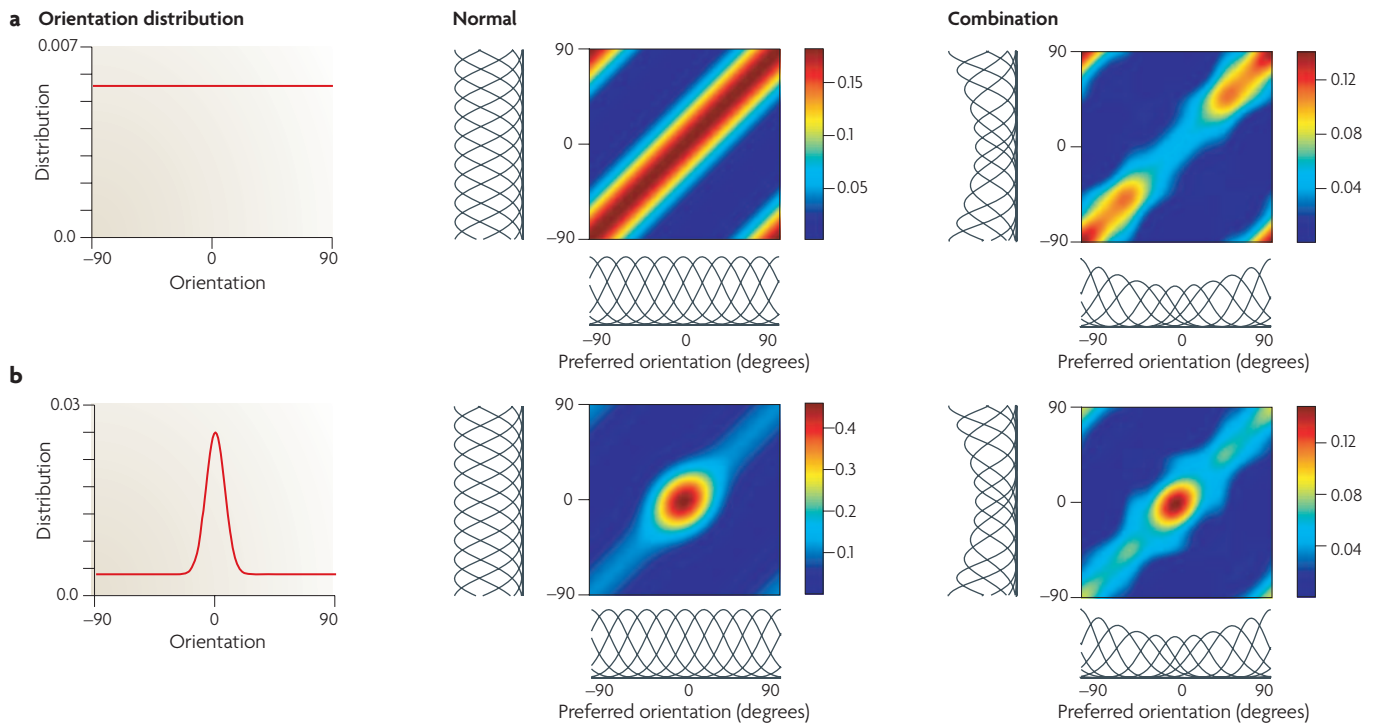
The efficient coding hypothesis<sup>91–92</sup> (influenced by information theory<sup>93</sup>) posits that neurons should use information-efficient, non-redundant codes for stimuli. Thus, without contextual information, codes should be generally non-redundant (and each neuron

should be used equally on average). When the context provides extra constraints, codes should be reorganized so that information is not repeated. This idea can be expressed in different terms — for instance, codes can be reorganized to concentrate on the aspects of the signal that are likely in the conditional distribution of the stimulus given the context  $P(I_{s_2}(t_2)|I_{s_1}(t_1))$ ; or to maximize information about the stimulus propagated by the neurons, given this conditional distribution. Although some efficient coding models have focused on only temporal or spatial context, it would be straightforward to map them onto the other domain. In general, however, the status of efficient coding as a principle for understanding cortical processing is unresolved<sup>70,94,95</sup>.

Efficient coding models span many levels of detail; different models also aim for different degrees of efficiency, for instance, settling for linear (or second order) de-correlation, instead of the independence between neural responses demanded by non-redundancy. Such models have been widely suggested to capture neural processing at the sensory periphery<sup>63,67,70,96–99</sup>.

More stringent (although, to our knowledge, not yet applied to the contextual phenomena described above) are models that seek, but do not necessarily achieve, independence (for example, Independent Component Analysis<sup>69,100</sup>). Decreasing the statistical dependence in filter outputs depicted in FIG. 4b was the explicit goal of models that apply non-linear transformations<sup>79,101</sup> that result in a form of divisive normalization<sup>102,103</sup>. These models can explain contextual phenomena in cortical neurons, such as response suppression, contrast gain control and tuning property changes<sup>79,101</sup>. Efficient coding has also been cited as a functional explanation for physiological contextual data<sup>24,32,38,104</sup>, such as the repulsive changes in orientation tuning curves produced by spatial or temporal contexts<sup>24,32,38</sup>. Other notions of efficiency have also been put forth in the spike rather than the mean firing rate domain, suggesting that when conditioned on the number of spikes, the neural response to the target and its context may be treated independently<sup>105,106</sup>.

Clifford *et al.*<sup>19</sup> suggested an appealing way of analysing perceptual contextual effects in terms of efficient coding. They considered a summary statistic (called the population vector) of the mean population activity for each target angle. Normally the distribution of this statistic is uniform, and thus maximally efficient across the target angle; Clifford and colleagues considered the classes of tuning changes (similar to those in FIG. 3e) that are necessary to maintain this uniformity in the face of biases in the distribution of input angles. They related the resulting population model to both bias<sup>19</sup> and discriminability<sup>36</sup>. By assuming instead that neurons that prefer orientations near the orientation specified by the context have greater variance in their responses, Wainwright<sup>107</sup> adopted the methods of efficient coding to suggest that these neurons should have comparatively suppressed responses, as in FIG. 3b, leading to the biases and changes in discriminability evident in the figure.



**Figure 5 | Correlation within a tuning curve population model.** The coloured plots show the covariances between all neurons averaged over distributions for target orientations that are either flat (**a**, left column; in the absence of contextual information) or when the contextual angle of 0° is favoured (**b**); and for either the unadapted tuning curves (normal, middle column) or tuning curves adapted according to the combined scheme (combination, right column; using the adaptation in FIG. 3e). Maximal efficiency mandates only diagonal covariances and equal use of all neurons; however, the peaked orientation distribution leads to substantial off-diagonal covariances and under-utilization of some neurons. The adaptations to the tuning curves improve these characteristics, but not completely.

We can directly examine how changes in tuning curves due to context affect the correlation statistics of the neural population model of FIG. 3. In the example in FIG. 5, tuning curve manipulations due to context reduce but do not eliminate the correlation dependencies of the population model. More mechanistic accounts include the suggestion that anti-Hebbian learning could be used to realize de-correlation<sup>108</sup> and the self-organizing network<sup>109,110</sup>, which has extensive inhibitory long-range lateral connections that serve this function.

The reason for biased inference under all these schemes is the coding catastrophe. Furthermore, the changed population response results in altered signal-to-noise ratios, which affects discrimination. Efficient coding does not address why the coding catastrophe occurs, because it lacks specification as to the computational goal beyond representation; rather, it embraces it without further question.

### Bayes

The other main class of functional approaches is associated with Bayesian inference and generative models. Bayesian methods in general have become popular for understanding the functional basis of a wide range of psychophysical and neural phenomena<sup>111</sup>, and several different, and not necessarily mutually exclusive, Bayesian ideas about contextual processing have been advocated. We discuss inference and generative models in turn.

**Inference models.** Bayesian inference models<sup>112–118</sup>, which are popular treatments of illusions and biases, start by assuming what the visual system might infer from a scene, such as local orientation. Two sources of information bear on this inference: prior expectations, perhaps based on long-run observed statistics, and observations about the given visual input (the same likelihood used by maximum-likelihood decoding).

The most common inferential Bayesian treatment of perceptual biases suggests that they arise from the imposition of prior expectations that are inappropriate for a particular scene. If spatial and temporal smoothness are interpreted as implying a prior that favours the same orientation for the target as for the context, Stocker and Simoncelli<sup>119</sup> point out that this would lead, if anything, to a bias favouring attraction rather than repulsion. The prior would be particularly important when observations are weak (for example, under conditions of uncertainty, such as low contrast or noisy target bars), a condition that does not seem to have been extensively tested. However, repulsion would never normally be expected. Instead, Stocker and Simoncelli<sup>119</sup> consider a class of the changes made in efficient coding that lead to higher likelihoods for tilts near to the adapting orientation. In a Bayesian model that makes inferences based on the mean of the posterior distribution (assuming a flat prior), they show that it is possible to capture substantial biases and discriminability changes without assuming the coding catastrophe.

**Prior**  
A probability distribution that captures the belief or expectation about a variable, in the absence of observations or evidence. Here, priors are specified through personal or evolutionary experience of environmental statistics.

Two alternative inferential Bayesian treatments address the coding catastrophe head-on. It has recently been observed that apparent repulsive biases can arise naturally in models in which there is competition among different explanations<sup>10,120,121</sup>, perhaps in the face of particular decision-theoretic loss functions. As a simplified example, consider the competition between explanations of the target: that it is either exactly consistent with the context, but suffering from perceptual noise; or that it is different. The closer its orientation to that of the context, the more likely it is to be seen as being the same; so the average perceptual orientation if it is seen as being different will favour values farther from the context, that is, repulsion. This idea has yet to be tested systematically on the data shown in FIG. 2.

More venerable, but perhaps only applicable to a restricted range of data on temporal context<sup>122</sup>, is the notion of Gibson's normalization<sup>5</sup> (not to be confused with divisive normalization associated with gain control), error correction<sup>123</sup> and calibration<sup>64</sup>. Crudely speaking, the idea underlying normalization or calibration is that systematic deviations over time in the input should not imply that the image contains particular values of  $\mathbf{I}(t_1)$  that indicate a changed prior on  $\mathbf{I}(t_2)$  (through the conditional distribution  $P(\mathbf{I}(t_2)|(\mathbf{I}(t_1)))$ ), but rather that the subject's sensory apparatus has fatigued or changed in a way that should be corrected in order to make correct inferences. One could put into Bayesian inferential terms the Gibsonian correction<sup>5</sup> by assuming that the statistics of the inferred images should be the overall long-run prior, and this can lead to repulsion. For instance, if the prior distribution indicates that vertical and horizontal tilts are most likely<sup>5,30</sup>, but recent experience has favoured neural responses consistent with a slightly off-vertical tilt, then the neural responses might need re-calibrating to be consistent with the long-run norm, which can lead to perceptual biases. This idea is controversial for tilt, given the biases prevalent for non-cardinal contexts<sup>124</sup>, but it has been well aired in the face and colour adaptation literature<sup>1,125</sup>. Webster *et al.*<sup>64</sup> go further, and suggest that it can be part of the process of regularization in the context of temporally and socially consistent communication.

Finally, anomalies might also arise not as an interaction of priors with likelihoods, but because true Bayesian inference is just too computationally complicated, thus making approximations necessary. By itself, this argument does not seem convincing: the ubiquity of contextual effects and specifically repulsion for a whole range of contextual attributes<sup>1,7,20</sup> would make it strange if such adaptations were not amongst the computations to which the approximations inherent in visual processing were tailored.

**Generative models.** A more holistic class of Bayesian approaches is associated with analysis-by-synthesis or generative models<sup>126–129</sup>. These suggest that the cortex builds a hierarchical statistical generative model of its input, with populations of neurons capturing the statistical regularities of visual inputs and offering a coordinate system within which to decompose (represent) new examples as a process of the sort of posterior inference discussed above. Note that the most efficient code for

inputs is closely related to the best probability distribution over those inputs. In these generative models, contextual processing is built into the nature of the representation itself, coming directly from the statistics of natural scenes. That is, the correct coordinate system for inputs is sensitive to spatially and temporally extended visual entities; and so inference is directly influenced by contextual manipulations.

One way to evade, rather than address or embrace the coding catastrophe, is to consider that the coordinate system used by the Bayesian model does not decompose the scene in the way the experimenter did. This is related to the notion that straightforward orientation estimation is not the visual system's goal. That is, perhaps pure local orientation does not form part of the Bayesian model's representational substrate for scenes; so rather than the semantics of the model being erroneous, as implied by the coding catastrophe, perhaps the semantics are inherently different in the first place. This explanation might seem odd, but it is common in discussions of other contextual effects such as lightness and brightness illusions<sup>3,130,131</sup>. There, the idea is that the visual system does not act as a veridical photometer, reporting the photon flux from different patches of visual space, but rather that it is interested in extracting two key properties that are associated with the patches — illumination (a global property that tends to be statistically coordinated for nearby spatial and temporal patches in a scene) and surface reflectance (typically a more local property). The underlying generative model assumes that the observed luminance in a patch is multiplicatively generated by these properties. Crudely, brightness illusions arise when correct Bayesian inference about these properties is at variance with the manner in which particular scenes were generated.

It has also been suggested that population responses to tilt arise as the output of a Bayesian generative model in which signals representing local structure (akin to reflectance) are multiplied by signals representing global structure (akin to illumination)<sup>132–134</sup>. The global structure does not have a direct physical analogue as for illumination, but rather it emerges owing to coordination of the orientation statistics of the type shown in FIG. 4b, representing structures such as orientated textures or edges<sup>81,83</sup>. Inference about the local structure amounts to a form of divisive normalization or gain control, and contextual information about the global signal gives rise to perceptual repulsion<sup>135</sup>.

Finally, predictive coding<sup>136–138</sup> is a Bayesian account which offers another link between hierarchical generative Bayesian approaches and efficient coding. According to the model of Rao and Ballard<sup>136</sup>, and ideas rooted in Mackay<sup>139</sup>, units in a downstream layer use feedback connections to inform lower layers of their current states. Lower layers then reorganize their codes (that is, change their population representation), so that the feed-forward bandwidth is used more efficiently to send information to amend the high level state, given their own, bottom-up information. Embracing the coding catastrophe, this can lead to perceptual anomalies.

#### Decision-theoretic loss function

The loss (or cost) associated with a particular decision about a quantity as a function of its true values. Bayesian decision theory suggests that choices should be made by minimizing expected losses under posterior distributions.

#### Gibson's normalization

The hypothesis that repulsive biases arise in the orientation domain due to a long-run prior favouring absolute cardinal axes.

### Perspectives and future directions

We have argued that the time is ripe to unify the wealth of disparate work on the effects of spatial and temporal context on visual processing and perception. We have focused on the comparatively simple case of orientation, discussing the strong similarities between spatial and temporal domains in terms of perceptual biases and changes in discriminability (FIG. 2a, b), and tuning curve and firing rate changes (FIG. 2c, d). These domains were putatively linked by mechanistic models of orientation processing (FIG. 3), and justified in efficient coding and Bayesian accounts by the key statistical properties that are shared by spatial and temporal contexts in natural scenes (FIG. 4).

Spatial and temporal contexts induce strong and similar perceptual biases, even for low-level features of an image, such as tilt. The similarity renders less appealing explanations that work in one, but not the other, domain, such as neural fatigue or Gibsonian regularization in the case of temporal context, or the intriguing functional account of adaptation, which posits that it preserves societal norms by ensuring that people exposed to similar environments have common perceptual experiences<sup>64</sup>.

Many current computational accounts appeal to some sort of coding catastrophe to explain perceptual bias. Some nascent accounts avoid this catastrophe, instead suggesting, for example, that the experimenter's and subject's models of image creation are incommensurate.

**Psychophysical issues.** There is obviously an inordinate range of possible spatio-temporal contexts, allowing structurally rich spatial patterns that change in systematic ways over time. Current psychophysical data are thus far from complete, particularly in terms of the comparison between, and integration of, spatial and temporal contextual effects. Furthermore, some results remain controversial.

Four classes of stimuli are of particular interest. One involves importing cases that have been well studied in one domain to the other domain. For example, studies of spatial, but not temporal, context have systematically investigated the separation of the target and contextual stimuli<sup>22,31,46,51,140,141</sup>, whereas in adaptation, mostly simpler manipulations of adaptation timescale have been tested<sup>142–144</sup>. Equally, adaptation experiments generally use foveal presentation of both context and target stimuli (however, see REF. 145), whereas more is known about the effects of retinal position on spatial context<sup>146</sup>.

A second class involves stimuli that couple spatial and temporal processing — rich, fast-changing, spatial patterns<sup>22,23</sup>. In the 'landing light' stimuli<sup>21</sup>, a contour is presented sequentially, with interesting effects on orientation estimation<sup>147</sup>. This can be seen as a temporal generalization of a case of spatial contour integration<sup>148</sup>, or a spatial generalization of a simple form of temporal adaptation.

The third class of stimuli highlights the differences between spatial and temporal contexts. For instance, geometry has a key role in spatial inference, with iso-orientated and ortho-orientated contextual surrounds having quite different statistical relationships with a

target, and it lies at the heart of various contextual models (for example, see REFS 120, 149; and see REF. 120 for explaining bias and discriminability data). Geometry also has a role in contextual effects that are involved in linking signals across space in order to bind common edges or borders<sup>65</sup>. There is no clear temporal analogue to geometry. Conversely, target figures that are more complicated than simple lines, such as faces, may not have strong spatial statistical constraints, but they do have rich temporal relationships that can lead to powerful adaptation.

The final class includes stimuli that are intended to distinguish between, and enrich, the modelling frameworks. For the statistical models, elucidation of the contextual effects on discriminability, and tests of the effects of likelihood manipulations are most pressing. Under Bayesian accounts, inference under conditions of uncertainty, such as low-contrast stimuli<sup>116</sup> or stimuli with added noise<sup>150</sup>, are critical. Perception under these conditions should depend more heavily on the contextual prior, and so constitute a means for testing the models. However, the experimental data are controversial. Furthermore, understanding the effects of more complex orientation contingencies, such as spatially overlapping orientations, could be most informative for modelling frameworks<sup>151–153</sup>.

Finally, we have focused on orientation estimation. Data from other orientation-based tasks, such as contrast detection<sup>140,141,154–157</sup>, which concentrate on low contrast target stimuli, and contrast discrimination<sup>158–161</sup>, offer windows into different aspects of contextual processing. There is also a large body of literature on adaptation in the case of motion (the motion after-effect<sup>2,48,49</sup>), which closely parallels orientation<sup>162</sup>, and there has been work on illusions in which temporal context affects brightness estimation<sup>4</sup>.

**Physiological issues.** The neurophysiological data are even less comprehensive. Problems include diverse responsiveness of neurons and the difficulty of interpreting results taken at a small number of punctuated locations in a complex and hierarchical architecture; differences across studies in species and behavioural states (anaesthetized versus awake); timescales of adaptation; and for spatial context, underestimating the classical receptive field, which could mean that it is invaded directly by contextual stimuli<sup>163</sup>.

It would be ideal, although perhaps unrealistic, to be able to couple all existing psychophysical results with their physiological counterparts. At least some neurophysiological data are available from the psychophysical protocols mentioned above, such as extensive studies on contrast manipulation<sup>31,36,40,164–168</sup> (FIG. 2e); coupling spatial surround and temporal adaptation experiments<sup>168</sup>; geometry<sup>31,140</sup>; timescales<sup>39,60</sup>; and more complex orientation contingencies<sup>169,170</sup>.

Ultimately, contextual processing needs to be understood under natural viewing conditions. Experimental techniques such as spike-triggered analysis and maximally informative dimensions allow one to present random stimuli or natural images, and to characterize the

#### Power law synapses

A synaptic adaptation that is (time) scale invariant; for example, having the same response shape at multiple timescales. This is in contrast to an exponential adaptation process with a single time constant.



29. Pouget, A., Dayan, P. & Zemel, R. Information processing with population codes. *Nature Rev. Neurosci.* **1**, 125–132 (2000).
30. Coppola, D. M., Purves, H. R., McCoy, A. N. & Purves, D. The distribution of oriented contours in the real world. *Proc. Natl Acad. Sci. USA* **95**, 4002–4006 (1998).
31. Cavanaugh, J. R., Bair, W. & Movshon, J. A. Selectivity and spatial distribution of signals from the receptive field surround in macaque V1 neurons. *J. Neurophysiol.* **88**, 2547–2556 (2002).
32. Müller, J. R., Metha, A. B., Krauskopf, J. & Lennie, P. Local signals from beyond the receptive fields of striate cortical neurons. *J. Neurophysiol.* **90**, 822–831 (2003).
33. van der Smagt, M. J., Wehrhahn, C. & Albright, T. D. Contextual masking of oriented lines: interactions between surface segmentation cues. *J. Neurophysiol.* **94**, 576–589 (2005).
34. Knierim, J. J. & van Essen, D. C. Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *J. Neurophysiol.* **67**, 961–980 (1992).
35. Nothdurft, H. C., Gallant, J. L. & Essen, D. C. V. Response modulation by texture surround in primate area V1: correlates of “popout” under anesthesia. *Vis. Neurosci.* **16**, 15–34 (1999).
36. Li, W., Thier, P. & Wehrhahn, C. Contextual influence on orientation discrimination of humans and responses of neurons in V1 of alert monkeys. *J. Neurophysiol.* **83**, 941–954 (2000).
37. Sengpiel, F., Sen, A. & Blakemore, C. Characteristics of surround inhibition in cat area 17. *Exp. Brain Res.* **116**, 216–228 (1997).
38. Müller, J. R., Metha, A. B., Krauskopf, J. & Lennie, P. Rapid adaptation in visual cortex to the structure of images. *Science* **285**, 1405–1408 (1999).
39. Dragoi, V., Sharma, J. & Sur, M. Adaptation-induced plasticity of orientation tuning in adult visual cortex. *Neuron* **28**, 287–298 (2000).
40. Crowder, N. A. et al. Relationship between contrast adaptation and orientation tuning in V1 and V2 of cat visual cortex. *J. Neurophysiol.* **95**, 271–283 (2006).
41. Levitt, J. B. & Lund, J. S. Contrast dependence of contextual effects in primate visual cortex. *Nature* **387**, 73–76 (1997).
42. Kohn, A. & Movshon, J. A. Adaptation changes the direction tuning of macaque MT neurons. *Nature Neurosci.* **7**, 764–772 (2004).
- Demonstrates the importance of studying systems hierarchically, reporting that population tuning changes at one level (in the primary visual cortex) can differ markedly from tuning changes at the next level (in the motion processing middle temporal area).**
43. Gilbert, C. D. & Wiesel, T. N. The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Res.* **30**, 1689–1701 (1990).
44. Li, C. Y., Lei, J. J. & Yao, H. S. Shift in speed selectivity of visual cortical neurons: a neural basis of perceived motion contrast. *Proc. Natl Acad. Sci. USA* **96**, 4052–4056 (1999).
45. Greenlee, M. W. & Magnussen, S. Saturation of the tilt aftereffect. *Vision Res.* **27**, 1041–1043 (1987).
46. Wenderoth, R. & van der Zwan, R. The effects of exposure duration and surrounding frames on direct and indirect tilt aftereffects and illusions. *Percept. Psychophys.* **46**, 338–344 (1989).
47. Schrater, P. R. & Simoncelli, E. P. Local velocity representation: evidence from motion adaptation. *Vision Res.* **38**, 3899–3912 (1998).
48. Mather, G., Verstraten, F. & Anstis, S. (eds) *The Motion Aftereffect* (MIT press, Massachusetts, USA, 1998).
49. Morgan, M., Chubb, C. & Solomon, J. A. Predicting the motion after-effect from sensitivity loss. *Vision Res.* **46**, 2412–2420 (2006).
50. Solomon, J. A. & Morgan, M. J. Stochastic recalibration: contextual effects on perceived tilt. *Proc. Biol. Sci.* **273**, 2681–2686 (2006).
51. Mareschal, I., Sceniak, M. P. & Shapley, R. M. Contextual influences on orientation discrimination: binding local and global cues. *Vision Res.* **41**, 1915–1930 (2001).
52. Dragoi, V., Sharma, J., Miller, E. K. & Sur, M. Dynamics of neuronal sensitivity in primate V1 underlying local feature discrimination. *Nature Neurosci.* **883**–891 (2002).
53. Westheimer, G. & Ge, A. Opposing views on orthogonal adaptation: a response to Clifford, Arnold, Smith, and Pianta (2003). *Vision Res.* **43**, 721–722 (2003).
54. Regan, D. & Beverley, K. I. Postadaptation orientation discrimination. *J. Opt. Soc. Am.* **2**, 147–155 (1985).
55. Clifford, C. W., Arnold, D. H., Smith, S. T. & Pianta, M. Opposing views on orthogonal adaptation: a reply to Westheimer and Ge (2002). *Vision Res.* **43**, 717–719 (2003).
56. Clifford, C. W., Wyatt, A. M., Arnold, D. H., Smith, S. T. & Wenderoth, P. Orthogonal adaptation improves orientation discrimination. *Vision Res.* **41**, 151–159 (2001).
57. Barlow, H. B., Macleod, D. I. A. & van Meeteren, A. Adaptation to gratings: no compensatory advantages found. *Vision Res.* **16**, 1043–1045 (1976).
58. Jin, D. Z., Dragoi, V., Sur, M. & Seung, H. S. Tilt aftereffect and adaptation-induced changes in orientation tuning in visual cortex. *J. Neurophysiol.* **94**, 4038–4050 (2005).
- Demonstrates how changes in tuning curves observed in cortical adaptation data can give rise to perceptual repulsion and attraction within a population decoding model.**
59. Averbeck, B. B., Latham, P. E. & Pouget, A. Neural correlations, population coding and computation. *Nature Rev. Neurosci.* **7**, 358–366 (2006).
60. Fairhall, A. L., Lewen, G. D., Bialek, W. & de Ruyter Van Steveninck, R. R. Efficiency and ambiguity in an adaptive neural code. *Nature* **412**, 787–792 (2001).
- Studies adaptation to variance statistics in a velocity-sensing neuron in the fly, and suggests that it might evade the coding catastrophe by reporting aspects of its state of adaptation through its long-run average firing rate.**
61. Teich, A. F. & Qian, N. Learning and adaptation in a recurrent model of V1 orientation selectivity. *J. Neurophysiol.* **89**, 2086–2100 (2003).
- Proposes a mechanistic population decoding model and draws out the perceptual implications for two different timescales of temporal context, namely adaptation and learning.**
62. Li, Z. A saliency map in primary visual cortex. *Trends Cogn. Sci.* **6**, 9–16 (2002).
63. Barlow, H. B. in *Vision: Coding and Efficiency* (ed. Blakemore, C.) 363–375 (Cambridge University Press, New York, USA, 1990).
- Barlow is a pioneer in the application of computational principles to explaining neural and psychophysical phenomena; here he discusses his original notion that efficient coding ideas can explain adaptation at the synaptic and perceptual levels.**
64. Webster, M. A., Werner, J. S. & Field, D. J. in *Fitting the Mind to the World: Adaptation and Aftereffects in High-Level Vision, Advances in Visual Cognition Series* (eds Clifford, C. W. & Rhodes, G. L.) 241–277 (Oxford University Press, USA, 2005).
- Points out that the effects of adaptation on perceptual discriminability are quite modest compared with the more striking perceptual repulsion, and suggests alternative functional frameworks with a focus on societal norms.**
65. Field, D. J. What is the goal of sensory coding? *Neural Comput.* **6**, 559–601 (1994).
66. Simoncelli, E. P. & Olshausen, B. A. Natural image statistics and neural representation. *Annu. Rev. Neurosci.* **24**, 1193–1216 (2001).
67. Li, Z. & Atick, J. J. Towards a theory of the striate cortex. *Neural Comput.* **6**, 127–146 (1994).
68. Olshausen, B. A. & Field, D. J. Emergence of simple-cell receptive field properties by learning a sparse factorial code. *Nature* **381**, 607–609 (1996).
69. Bell, A. J. & Sejnowski, T. J. The ‘independent components’ of natural scenes are edge filters. *Vision Res.* **37**, 3327–3338 (1997).
70. Zhaoping, L. Theoretical understanding of the early visual processes by data compression and data selection. *Network* **17**, 301–334 (2006).
- A review of theoretical frameworks for understanding early visual processing, including the notion that intracortical connections in the primary visual cortex are responsible for computing a map of the salience of regions of visual space, and that this can explain properties that are opaque to efficient coding principles.**
71. Field, D. J., Hayes, A. & Hess, R. F. Contour integration by the human visual system: evidence for a local ‘association field’. *Vision Res.* **33**, 173–193 (1993).
72. Elder, J. H. & Goldberg, R. M. Ecological statistics of Gestalt laws for the perceptual organization of contours. *J. Vis.* **2**, 324–353 (2002).
73. Olshausen, B. A. & Field, D. J. Vision and the coding of natural images. *Am. Sci.* **88**, 238–244 (2000).
74. Baddeley, R. The correlational structure of natural images and the calibration of spatial representations. *Cogn. Sci.* **21**, 351–372 (1997).
75. Dong, D. W. & Atick, J. J. Statistics of natural time-varying images. *Network* **6**, 345–358 (1995).
76. Grzywacz, N. M. & de Juan, J. Sensory adaptation as Kalman filtering: theory and illustration with contrast adaptation. *Network* **14**, 465–482 (2003).
77. Geman, S. & Geman, D. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE, Trans. Pat. Anal. Mach. Intell.* **6**, 721–741 (1984).
78. Kersten, D. Predictability and redundancy of natural images. *J. Opt. Soc. Am. A* **4**, 2395–2400 (1987).
79. Schwartz, O. & Simoncelli, E. P. Natural signal statistics and sensory gain control. *Nature Neurosci.* **4**, 819–825 (2001).
80. Hoyer, P. & Hyvärinen, A. A multi-layer sparse coding network learns contour coding from natural images. *Vision Res.* **42**, 1593–1605 (2002).
81. Karklin, Y. & Lewicki, M. S. A hierarchical Bayesian model for learning nonlinear statistical regularities in nonstationary natural signals. *Neural Comput.* **17**, 397–423 (2005).
82. Zetsche, C. & Nuding, U. Nonlinear and higher-order approaches to the encoding of natural scenes. *Network* **16**, 191–221 (2005).
83. Schwartz, O., Sejnowski, T. J. & Dayan, P. Soft mixer assignment in a hierarchical generative model of natural scene statistics. *Neural Comput.* **18**, 2680–2718 (2006).
84. Hyvärinen, A., Hurri, J. & Väyrynen, J. Bubbles: a unifying framework for low-level statistical properties of natural image sequences. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **20**, 1237–1252 (2003).
85. Sigman, M., Cecchi, G. A., Gilbert, C. D. & Magnasco, M. O. On a common circle: natural scenes and Gestalt rules. *Proc. Natl Acad. Sci. USA* **98**, 1935–1940 (2001).
86. Geisler, W. S., Perry, J. S., Super, B. J. & Gallogly, D. P. Edge co-occurrence in natural images predicts contour grouping performance. *Vision Res.* **41**, 711–724 (2001).
- Examines both absolute and Bayesian orientation spatial context statistics in natural images, and derives a model of contour grouping that is compared to perception.**
87. Howe, C. Q. & Purves, D. Natural-scene geometry predicts the perception of angles and line orientation. *Proc. Natl Acad. Sci. USA* **102**, 1228–1233 (2005).
- Measures image statistics of overlapping orientations in space, and suggests how these can give rise to perceptual tilt illusions.**
88. Dragoi, V. & Sur, M. Image structure at the center of gaze during free viewing. *J. Cogn. Neurosci.* **18**, 737–748 (2006).
89. Melcher, D. Spatiotopic transfer of visual-form adaptation across saccadic eye movements. *Curr. Biol.* **15**, 1745–1748 (2005).
90. Burr, D. & Morrone, M. C. Eye movements: building a stable world from glance to glance. *Curr. Biol.* **15**, R839–R840 (2005).
91. Barlow, H. B. in *Sensory Communication* (ed. Rosenblith, W. A.) 217–234 (MIT press, Massachusetts, 1961).
92. Atneave, F. Some informational aspects of visual perception. *Psychol. Rev.* **61**, 183–193 (1954).
93. Shannon, C. A mathematical theory of communication. *Bell Sys. Tech. J.* **27**, 379–423 (1948).
94. Barlow, H. B. Redundancy reduction revisited. *Network* **12**, 241–253 (2001).
95. Simoncelli, E. P. Vision and the statistics of the visual environment. *Curr. Opin. Neurobiol.* **13**, 144–149 (2003).
- Reviews the efficient coding hypothesis in recent literature, covering aspects of experimental testing of these principles and computational models based on efficient coding of natural images.**
96. Barlow, H. B. & Foldiak, P. in *The Computing Neuron* (eds Miall, C., Durbin, R. M. & Mitchison, G. J.) 54–72 (Addison-Wesley, England, 1989).
97. Atick, J. J. & Redlich, A. N. Towards a theory of early visual processing. *Neural Comput.* **2**, 308–320 (1990).
98. Atick, J. J., Li, Z. & Redlich, A. N. What does post-adaptation color appearance reveal about cortical color representation? *Vision Res.* **33**, 123–129 (1993).
99. Dan, Y., Atick, J. J. & Reid, R. C. Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J. Neurosci.* **16**, 3351–3362 (1996).

100. Smith, E. & Lewicki, M. Efficient auditory coding. *Nature* **439**, 978–982 (2006).

101. Wainwright, M. J., Schwartz, O. & Simoncelli, E. P. in *Probabilistic Models of the Brain: Perception and Neural Function* (eds Rao, R., Olshausen, B. A. & Lewicki, M.) 203–222 (MIT Press, Massachusetts, USA, 2002).

102. Geisler, W. S. & Albrecht, D. G. Cortical neurons: isolation of contrast gain control. *Vision Res.* **8**, 1409–1410 (1992).

103. Heeger, D. J. Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* **9**, 181–198 (1992).

104. Vinje, W. E. & Gallant, J. L. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* **287**, 1273–1276 (2000).

**Suggests that spatial context increases the efficiency of cortical neural processing.**

105. Movellan, J. R., Wachtler, T., Albright, T. D. & Sejnowski, T. J. in *Neural Information Processing Systems* (eds Becker, S., Thrun, S. & Obermayer, K.) 205–212 (MIT Press, Massachusetts, USA, 2002).

106. Schwartz, O., Movellan, J. R., Wachtler, T., Albright, T. D. & Sejnowski, T. J. Spike count distributions, factorizability, and contextual effects in area V1. *Neurocomputing* **58–60** (2004).

107. Wainwright, M. J. Visual adaptation as optimal information transmission. *Vision Res.* **39**, 3960–3974 (1999).

**Proposes that adaptation serves to optimize information transmission in an efficient coding context, using the tilt after-effect as a key example.**

108. Foldiak, P. Forming sparse representations by local anti-Hebbian learning. *Biol. Cybern.* **64**, 165–170 (1990).

109. Bednar, J. A. & Miikkulainen, R. Tilt aftereffects in a self-organizing model of the primary visual cortex. *Neural Comput.* **12**, 1721–1740 (2000).

110. Sirosh, J. & Miikkulainen, R. Topographic receptive fields and patterned lateral interaction in a self-organizing model of the primary visual cortex. *Neural Comput.* **9**, 577–594 (1997).

111. Doya, K., Ishii, S., Pouget, A. & Rao, R. P. N. (eds) *Bayesian Brain: Probabilistic Approaches to Neural Coding* (MIT Press, Massachusetts, USA, 2007).

112. Yuille, A. & Bulthoff, H. H. in *Bayesian Decision Theory and Psychophysics* (eds Knill, D. and Richards, W.) 123–161 (Cambridge University Press, New York, USA, 1996).

113. Yuille, A. & Kersten, D. Vision as Bayesian inference: analysis by synthesis? *Trends Cogn. Sci.* **10**, 301–308 (2006).

114. Balboa, R. M. & Grzywacz, N. M. The minimal local-asperity hypothesis of early retinal lateral inhibition. *Neural Comput.* **12**, 1485–1517 (2000).

116. Grzywacz, N. M. & Balboa, R. M. A Bayesian framework for sensory adaptation. *Neural Comput.* **14**, 543–559 (2002).

116. Weiss, Y., Simoncelli, E. P. & Adelson, E. H. Motion illusions as optimal percepts. *Nature Neurosci.* **5**, 598–604 (2002).

117. Stocker, A. A. & Simoncelli, E. P. Noise characteristics and prior expectations in human visual speed perception. *Nature Neurosci.* **9**, 578–585 (2006).

118. Kersten, D., Mamassian, P. & Yuille, A. Object perception as Bayesian inference. *Annu. Rev. Psychol.* **55**, 271–304 (2004).

119. Stocker, A. A. & Simoncelli, E. P. in *NIPS Advances in Neural Information Processing Systems* (eds Weiss, Y., Schölkopf, B. & Platt, J.) 1291–1298 (MIT Press, Massachusetts, USA, 2006).

**Proposes a Bayesian model of visual adaption in terms of adjustments to the likelihood, on the important basis that changes to the prior are likely to lead to perceptual attraction rather than repulsion.**

120. Schwartz, O., Sejnowski, T. J. & Dayan, P. in *Advances in Neural Information Processing Systems* **18** (eds Weiss, Y., Schölkopf, B. & Platt, J.) 1201–1208 (MIT Press, Massachusetts, USA, 2006).

121. Körding, K. & Tenenbaum, J. B. in *Advances in Neural Information Processing Systems* (eds Schölkopf, B., Platt, J. & Hoffman, T.) 737–744 (The MIT Press, Massachusetts, USA, 2006).

122. Over, R. Comparison of normalization theory and neural enhancement explanation of negative aftereffects. *Psychol. Bull.* **75**, 225–243 (1971).

123. Andrews, D. P. Error-correcting-perceptual mechanisms. *Q. J. Exp. Psychol.* **16**, 104–115 (1964).

124. Mitchell, D. E. & Muir, D. W. Does the tilt after-effect occur in the oblique meridian? *Vision Res.* **16**, 609–613 (1976).

125. Leopold, D. A., Bondar, I. V. & Giese, M. A. Norm-based face encoding by single neurons in the monkey inferotemporal cortex. *Nature* **442**, 572–575 (2006).

126. Neisser, U. *Cognitive Psychology* (Prentice-Hall, New Jersey, USA, 1967).

127. Hinton, G. E. & Ghahramani, Z. Generative models for discovering sparse distributed representations. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **352**, 1177–1190 (1997).

128. Grenander, U. & Srivastava, A. Probability models for clutter in natural images. *IEEE. Trans. Patt. Anal. Mach. Intell.* **23**, 423–429 (2002).

129. Zhu, S. & Mumford, D. Prior learning and gibbs reaction-diffusion. *IEEE. Trans. Patt. Anal. Mach. Intell.* **19**, 1236–1250 (1997).

130. Nundy, S. & Purves, D. A probabilistic explanation of brightness scaling. *Proc. Natl Acad. Sci. USA* **99**, 14482–14487 (2002).

131. Gilchrist, A. L. et al. An anchoring theory of lightness perception. *Psychol. Rev.* **4**, 795–834 (1999).

132. Andrews, D. P. & Mallows, C. Scale mixtures of normal distributions. *J. R. Stat. Soc.* **36**, 99–102 (1974).

133. Wainwright, M. J. & Simoncelli, E. P. in *Advances in Neural Information Processing Systems* (eds Solla, S. A., Leen, T. K. & Müller, K. R.) 855–861 (MIT Press, Massachusetts, USA, 2000).

134. Portilla, J., Strela, V., Wainwright, M. & Simoncelli, E. P. Image denoising using a scale mixture of Gaussians in the wavelet domain. *IEEE. Trans. Image Process.* **12**, 1338–1351 (2003).

135. Schwartz, O., Sejnowski, T. J. & Dayan, P. The tilt illusion, population decoding, and natural scene statistics. *Computational and Systems Neuroscience (COSYNE) Abstract* 280 (2007).

136. Rao, R. P. & Ballard, D. H. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neurosci.* **2**, 79–87 (1999).

137. Srinivasan, M. V., Laughlin, S. B. & Dubs, A. Predictive coding: a fresh view of inhibition in the retina. *Proc. R. Soc. Lond. B Biol. Sci.* **216**, 427–459 (1982).

138. Hosoya, T., Baccus, S. A. & Meister, M. Dynamic predictive coding by the retina. *Nature* **436**, 71–77 (2005).

139. Mackay, D. M. in *Automata Studies* (eds Shannon, C. E. & McCarthy, J.) 235–251 (Princeton University Press, New Jersey, USA, 1956).

140. Kapadia, M. K., Ito, M., Gilbert, C. D. & Westheimer, G. Improvement in visual sensitivity by changes in local context: parallel studies in human observers and in V1 of alert monkeys. *Neuron* **15**, 843–856 (1995).

141. Polat, U. & Sagi, D. Lateral interactions between spatial channels: suppression and facilitation revealed by lateral masking experiments. *Vision Res.* **33**, 993–999 (1993).

142. Magnussen, S. & Johnsen, T. Temporal aspects of spatial adaptation. A study of the tilt aftereffect. *Vision Res.* **26**, 661–672 (1986).

143. Rose, D. A square root law for adaptation to contrast? *Vision Res.* **32**, 1781–1788 (1992).

144. Kanai, R. & Verstraten, F. A. Perceptual manifestations of fast neural plasticity: motion priming, rapid motion aftereffect and perceptual sensitization. *Vision Res.* **45**, 3109–3116 (2005).

145. Muir, D. & Over, R. Tilt aftereffects in central and peripheral vision. *J. Exp. Psychol.* **85**, 165–170 (1970).

146. Solomon, J. A., Felisberti, F. M. & Morgan, M. J. Crowding and the tilt illusion: toward a unified account. *J. Vis.* **4**, 500–508 (2004).

147. Roberts, M. J. & Thielle, A. Attention and contrast modulate the influence of spatio-temporal context in orientation discrimination of human subjects. *FENS abstract* **3**, A053.15 (2006).

148. Kapadia, M. K., Westheimer, G. & Gilbert, C. D. Spatial distribution of contextual interactions in primary visual cortex and in visual perception. *J. Neurophysiol.* **84**, 2048–2062 (2000).

149. Li, Z. Computational design and nonlinear dynamics of a recurrent network model of the primary visual cortex. *Neural Comput.* **13**, 1749–1780 (2001).

150. Pelli, D. G. & Farell, B. Why use noise? *J. Opt. Soc. Am. A* **16**, 647–653 (1999).

151. Meese, T. S. & Georgeson, M. A. The tilt aftereffect in plaids and gratings: channel codes, local signs and “patchwise” transforms. *Vision Res.* **36**, 1421–1437 (1996).

152. Georgeson, M. A. Human vision combines oriented filters to compute edges. *Proc. Biol. Sci.* **249**, 235–245 (1992).

153. Smith, S., Wenderoth, P. & van der Zwan, R. Orientation processing mechanisms revealed by the plaid tilt illusion. *Vision Res.* **41**, 483–494 (2001).

154. Foley, J. M. & Boynton, G. M. Forward pattern masking and adaptation: effects of duration, interstimulus interval, contrast, and spatial and temporal frequency. *Vision Res.* **33**, 959–980 (1993).

155. Foley, J. M. & Yang, Y. D. Forward pattern masking: effects of spatial frequency and contrast. *J. Opt. Soc. Am. A* **8**, 2026–2037 (1991).

156. Wehrhahn, C. & Dresch, B. Detection facilitation by collinear stimuli in humans: dependence on strength and sign of contrast. *Vision Res.* **38**, 423–428 (1998).

157. Tanaka, Y. & Sagi, D. Long-lasting, long-range detection facilitation. *Vision Res.* **38**, 2591–2599 (1998).

158. Zenger-Landolt, B. & Koch, C. Flanker effects in peripheral contrast discrimination—psychophysics and modeling. *Vision Res.* **41**, 3663–3675 (2001).

159. Chen, C. C. & Tyler, C. W. Lateral modulation of contrast discrimination: flanker orientation effects. *J. Vis.* **2**, 520–530 (2002).

160. Greenlee, M. W. & Heitger, F. The functional role of contrast adaptation. *Vision Res.* **28**, 791–797 (1988).

161. Maattanen, L. M. & Koenderink, J. J. Contrast adaptation and contrast gain control. *Exp. Brain Res.* **87**, 205–212 (1991).

162. Clifford, C. W. Perceptual adaptation: motion parallels orientation. *Trends Cogn. Sci.* **6**, 136–143 (2002).

**Demonstrates perceptual analogies in the adaptation to orientation and motion, suggesting that common computational principles may underlie contextual processing in both domains.**

163. 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).

164. Albrecht, D. G., Farrar, S. B. & Hamilton, D. B. Spatial contrast adaptation characteristics of neurons recorded in the cat’s visual cortex. *J. Physiol. (Lond.)* **347**, 713–739 (1984).

165. Ohzawa, I., Sclar, G. & Freeman, R. D. Contrast gain control in the cat’s visual system. *J. Neurophysiol.* **54**, 651–667 (1985).

166. Kapadia, M. K., Westheimer, G. & Gilbert, C. D. Dynamics of spatial summation in primary visual cortex of alert monkeys. *Proc. Natl Acad. Sci. USA* **21**, 12073–12078 (1999).

167. Sceniak, M. P., Ringach, D. L., Hawken, M. J. & Shapley, R. Contrast’s effect on spatial summation by macaque V1 neurons. *Nature Neurosci.* **2**, 733–739 (1999).

168. Webb, B. S., Dhruv, N. T., Solomon, S. G., Tailby, C. & Lennie, P. Early and late mechanisms of surround suppression in striate cortex of macaque. *J. Neurosci.* **25**, 11666–11675 (2005).

169. Walker, G. A., Ohzawa, I. & Freeman, R. D. Disinhibition outside receptive fields in the visual cortex. *J. Neurosci.* **22**, 5659–5668 (2002).

170. Carandini, M., Movshon, J. A. & Ferster, D. Pattern adaptation and cross-orientation interactions in the primary visual cortex. *Neuropharmacology* **37**, 501–511 (1998).

171. Sharpee, T. O. et al. Adaptive filtering enhances information transmission in visual cortex. *Nature* **439**, 936–942 (2006).

172. Fuster, J. M. *Memory in the Cerebral Cortex: An Empirical Approach to Neural Networks in the Human and Nonhuman Primate* (MIT Press, Massachusetts, USA, 1995).

173. Goldman-Rakic, P. S. Regional and cellular fractionation of working memory. *Proc. Natl Acad. Sci. USA* **93**, 13473–13480 (1996).

174. Miller, E. K. & Cohen, J. D. An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* **24**, 167–202 (2001).

175. Renart, A., Song, O. & Wang, X. J. Robust spatial working memory through homeostatic synaptic scaling in heterogeneous cortical networks. *Neuron* **38**, 473–485 (2003).

176. Abbott, L. F., Varela, J. A., Sen, K. & Nelson, S. B. Synaptic depression and cortical gain control. *Science* **275**, 220–224 (1997).

177. Markram, H. & Tsodyks, M. V. Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature* **382**, 807–809 (1996).

178. 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).

179. Fusi, S., Drew, P. J. & Abbott, L. F. Cascade models of synaptically stored memories. *Neuron* **45**, 599–611 (2005).
180. Drew, P. J. & Abbott, L. F. Models and properties of power-law adaptation in neural systems. *J. Neurophysiol.* **96**, 826–833 (2006).
181. Wang, X. J., Liu, Y., Sanchez-Vives, M. V. & McCormick, D. A. Adaptation and temporal decorrelation by single neurons in the primary visual cortex. *J. Neurophysiol.* **89**, 3279–3293 (2003).
182. Carandini, M. Visual cortex: fatigue and adaptation. *Curr. Biol.* **10**, R605–R607 (2000).
183. Sanchez-Vives, M. V., Nowak, L. G. & McCormick, D. A. Cellular mechanisms of long-lasting adaptation in visual cortical neurons. *J. Neurosci.* **20**, 4286–4299 (2000).
184. Nowak, L. G., Sanchez-Vives, M. V. & McCormick, D. A. Role of synaptic and intrinsic membrane properties in short-term receptive field dynamics in cat area 17. *J. Neurosci.* **25**, 1866–1880 (2005).
185. Fitzpatrick, D. The functional organization of local circuits in visual cortex: insights from the study of tree shrew striate cortex. *Cereb. Cortex* **6**, 329–341 (1996).
186. Angelucci, A. & Bullier, J. Reaching beyond the classical receptive field of V1 neurons: horizontal or feedback axons? *J. Physiol. Paris* **97**, 141–154 (2003).
187. Levitt, J. B. & Lund, J. S. The spatial extent over which neurons in macaque striate cortex pool visual signals. *Vis. Neurosci.* **19**, 439–452 (2002).
188. Bair, W., Cavanaugh, J. R. & Movshon, J. A. Time course and time-distance relationships for surround suppression in macaque V1 neurons. *J. Neurosci.* **23**, 7690–7601 (2003).
189. Stetter, M., Bartsch, H. & Obermayer, K. A mean-field model for orientation tuning, contrast saturation, and contextual effects in the primary visual cortex. *Biol. Cybern.* **82**, 291–304 (2000).
190. Bressloff, P. C. & Cowan, J. D. An amplitude equation approach to contextual effects in visual cortex. *Neural Comput.* **14**, 493–525 (2002).
191. Sullivan, T. J. & de Sa, V. R. A model of surround suppression through cortical feedback. *Neural Netw.* **19**, 564–572 (2006).
192. Schwabe, L., Obermayer, K., Angelucci, A. & Bressloff, P. C. The role of feedback in shaping the extra-classical receptive field of cortical neurons: a recurrent network model. *J. Neurosci.* **26**, 9117–9129 (2006).
193. Das, A. & Gilbert, C. D. Topography of contextual modulations mediated by short-range interactions in primary visual cortex. *Nature* **399**, 655–661 (1999).
194. Okamoto, T., Watanabe, M., Aihara, K. & Kondo, S. An explanation of contextual modulation by short-range isotropic connections and orientation map geometry in the primary visual cortex. *Biol. Cybern.* **91**, 396–407 (2004).
195. Wiesel, J. & Sapiro, P. Extraclassical receptive field phenomena and short-range connectivity in V1. *Cereb. Cortex* **16**, 1531–1545 (2006).
196. Dragoi, V., Rivadulla, C. & Sur, M. Foci of orientation plasticity in visual cortex. *Nature* **411**, 80–86 (2001).
197. Gelb, M. H., Calvert, J. E., Harris, J. P. & Phillipson, O. T. Modification of visual orientation illusions by drugs which influence dopamine and GABA neurones: differential effects on simultaneous and successive illusions. *Psychopharmacology (Berl.)* **90**, 379–383 (1986).
198. Borst, A., Flanagan, V. L. & Sompolinsky, H. Adaptation without parameter change: dynamic gain control in motion detection. *Proc. Natl Acad. Sci. USA* **102**, 6172–6176 (2005).
199. Boynton, G. M. & Finney, E. M. Orientation-specific adaptation in human visual cortex. *J. Neurosci.* **23**, 8781–8787 (2003).
200. Fraser, J. A new visual illusion of direction. *Brit. J. Psychol.* **2**, 307–320 (1908).
201. Kayser, W., Einhauser, C. & Konig, P. Temporal correlations of orientations in natural scenes. *Neurocomputing* **52–54**, 117–123 (2003).
202. Simoncelli, E. P., Freeman, W. T., Adelson, E. H. & Heeger, D. J. Shiftable multi-scale transforms. *IEEE Trans. Inform. Theory* **38**, 587–607 (1992).
203. Simoncelli, E. P. in Proc. 31st Asilomar conf. on *Signals, Systems and Computers* 673–678 (Pacific Grove, California, USA, 1997).

#### Competing interests statement

The authors declare no competing financial interests.

#### Acknowledgements

This work was funded by the Howard Hughes Medical Institute (O.S.), the Gatsby Charitable Foundation (A.H., P.D.), the Biotechnology and Biological Sciences Research Council, the Engineering and Physical Sciences Research Council and the Wellcome Trust (P.D.). We are very grateful to C. Clifford, A. Kohn, A. Stocker and J. Solomon for comments on the manuscript and discussion, and to T. Sejnowski and E. Simoncelli for discussion.

#### FURTHER INFORMATION

Odellia Schwartz's homepage: [http://neuroscience.aecom.yu.edu/faculty/primary\\_faculty\\_pages/schwartz.html](http://neuroscience.aecom.yu.edu/faculty/primary_faculty_pages/schwartz.html)

Peter Dayan's homepage: [www.gatsby.ucl.ac.uk/~dayan](http://www.gatsby.ucl.ac.uk/~dayan)

Access to this links box is available online.