



Free-energy and illusions: the Cornsweet effect

Harriet Brown* and Karl J. Friston

The Wellcome Trust Centre for Neuroimaging, University College London Queen Square, London, UK

Edited by:

Lars Muckli, University of Glasgow, UK

Reviewed by:

Gerrit W. Maus, University of California at Berkeley, USA
Peter De Weerd, Maastricht University, Netherlands

***Correspondence:**

Harriet Brown, Wellcome Trust Centre for Neuroimaging, Institute of Neurology Queen Square, London WC1N 3BG, UK.
e-mail: harriet.brown.09@ucl.ac.uk

In this paper, we review the nature of illusions using the free-energy formulation of Bayesian perception. We reiterate the notion that illusory percepts are, in fact, Bayes-optimal and represent the most likely explanation for ambiguous sensory input. This point is illustrated using perhaps the simplest of visual illusions; namely, the Cornsweet effect. By using plausible prior beliefs about the spatial gradients of illuminance and reflectance in visual scenes, we show that the Cornsweet effect emerges as a natural consequence of Bayes-optimal perception. Furthermore, we were able to simulate the appearance of secondary illusory percepts (Mach bands) as a function of stimulus contrast. The contrast-dependent emergence of the Cornsweet effect and subsequent appearance of Mach bands were simulated using a simple but plausible generative model. Because our generative model was inverted using a neurobiologically plausible scheme, we could use the inversion as a simulation of neuronal processing and implicit inference. Finally, we were able to verify the qualitative and quantitative predictions of this Bayes-optimal simulation psychophysically, using stimuli presented briefly to normal subjects at different contrast levels, in the context of a fixed alternative forced choice paradigm.

Keywords: free-energy, perception, Bayesian inference, illusions, Cornsweet effect, perceptual priors

INTRODUCTION

Illusions are often regarded as “failures” of perception; however, Bayesian considerations often provide a principled explanation for apparent failures of inference in terms of prior beliefs. This paper is about the nature of illusions and their relationship to Bayes-optimal perception. The main point made by this work is that illusory percepts are optimal in the sense of explaining sensations in terms of their most likely cause. In brief, illusions occur when the experimenter generates stimuli in an implausible or unlikely way. From the subject’s perspective, these stimuli are ambiguous and could be explained by different underlying causes. This ambiguity is resolved in a Bayesian setting, by choosing the most likely explanation, given prior beliefs about the hidden causes of the percept. This key point has been made by many authors (e.g., Purves et al., 1999). Here, we develop it under biologically realistic simulations of Bayes-optimal perception and try to make some quantitative predictions about how subjects should make perceptual decisions. We then try to establish the scheme’s validity by showing that these predictions are largely verified by experimental data from normal subjects viewing the same stimuli.

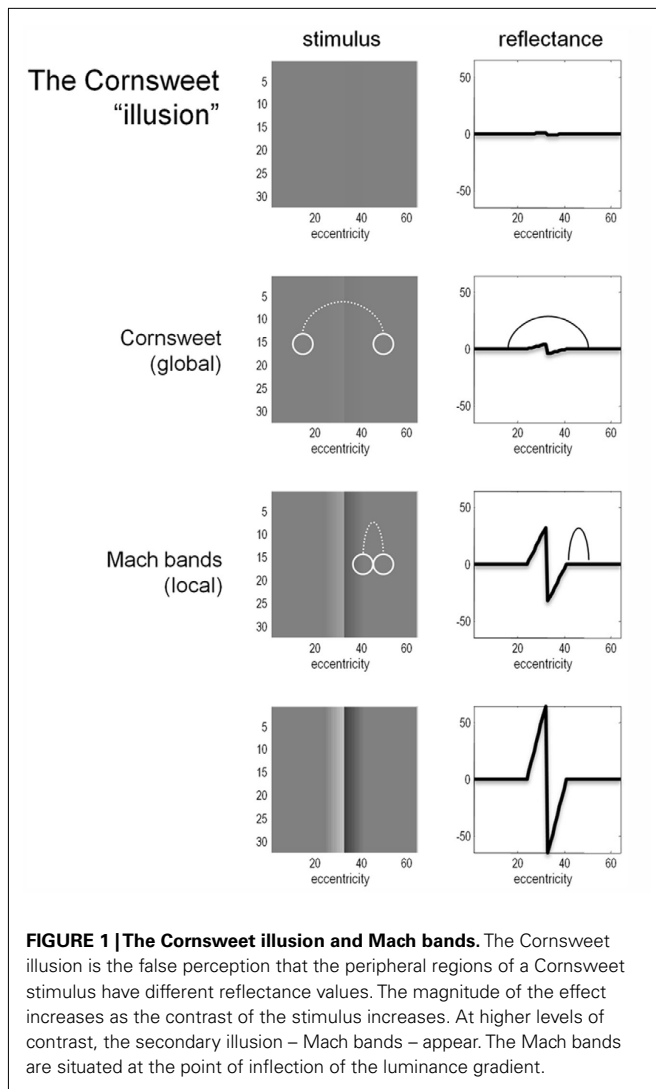
The example we have chosen is the Cornsweet illusion, which has a long history, dating back to the days of Helmholtz (Mach, 1865; O’Brien, 1959; Craik, 1966; Cornsweet, 1970). This is particularly relevant given our formulation of the Bayesian brain is based upon the idea that the brain is a Helmholtz or inference machine (von Helmholtz, 1866; Barlow, 1974; Dayan et al., 1995; Friston et al., 2006). In other words, the brain is trying to infer the hidden causes and states of the world generating sensory information, using predictions based upon a generative model that includes prior beliefs. We hoped to show that the Cornsweet effect can be explained in a parsimonious way by some simple prior beliefs

about the way that visual information is generated at different spatial and temporal scales.

THE CORNSWEET EFFECT AND THE NATURE OF ILLUSIONS

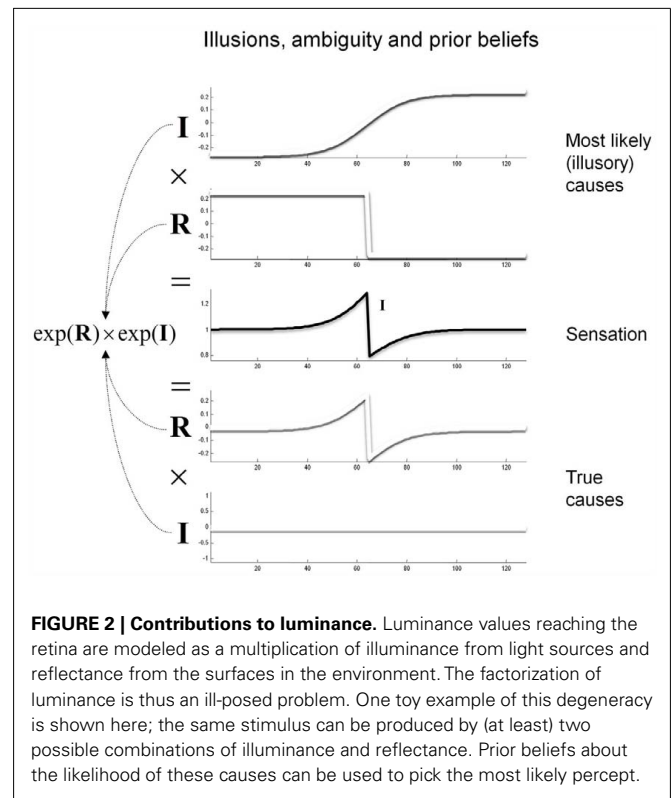
Figure 1 provides an illustration of the Cornsweet illusion. The illusion is the false percept that the peripheral regions of a stimulus have a different brightness, despite the fact they are physically iso-luminant. This illusion is induced by a biphasic luminance “edge” in the centre of the field of view (shown in the right hand column of **Figure 1**). The four rows of **Figure 1** show the Cornsweet effect increasing in magnitude as we increase the contrast of the stimulus. Interestingly, at high levels of contrast, secondary illusions – Mach bands (Mach, 1865; Lotto et al., 1999) – appear at the para-central points of inflection of the true luminance profile. It is this contrast-dependent emergence of the Cornsweet effect and subsequent Mach bands that we wanted to simulate, under the assumption that perception is Bayes-optimal.

The Bayesian aspect of perception becomes crucial when we consider the nature of illusions. Bayesian theories of perception describe how sensory data (that have a particular likelihood) are combined with prior beliefs (a prior distribution) to create a percept (a posterior distribution). One can regard illusory percepts as those that are induced by ambiguous stimuli, which can be caused in different ways – in other words, the probability of the data given different causes or explanations is the same. When faced with these stimuli, the prior distribution can be used to create a unimodal posterior and an unambiguous percept. If the percept or inference about the hidden causes of sensory information (the posterior distribution) is different from the true causes used to generate stimuli, the inference is said to be illusory or false. However, with illusory stimuli the mapping of hidden causes to their



sensory consequences is ill-posed (degenerate or many to one), such that a stimulus can have more than one cause. Thus, from the point of view of the observer, there can be no “false” inference unless the true causes are known. The perceptual inference can be optimal in a Bayesian sense, but is still illusory. However, not all possible causes of sensory input will be equally likely, so there will be an optimal inference in relation to prior beliefs about their causes. Prior beliefs can be learnt or innate: priors that are learnt depend upon experience while innate priors can be associated with architectural features of the visual brain, such as the complex arrangement of blobs, interblobs, and stripes in V1, that may reflect priors on the statistical structure of visual information selected by evolutionary pressure.

Prior beliefs are essential when resolving ambiguity or the ill-posed nature of perceptual inverse problems. Put simply, there will always be an optimal posterior estimate of what caused a sensation that rests upon prior beliefs. The example in **Figure 2** illustrates this: The central panel shows an ambiguous stimulus (luminance profile) that is formally similar to the sort of stimulus that induces the Cornsweet illusion. However, this stimulus can be caused in



an infinite number of ways. We have shown two plausible causes by assuming the stimulus is the product of (non-negative) illuminance and reflectance profiles. The lower two panels show the “true” causes generating stimuli for the Cornsweet illusion. Here, the stimulus has a reflectance profile that reproduces the Cornsweet stimulus and is illuminated with a uniform illuminant. An alternative explanation for exactly the same stimulus is provided in the upper two panels, in which two isoreflectant surfaces are viewed under a smooth gradient of ambient illumination. In this example, we have ensured that both the illuminant and reflectance are non-negative by applying an exponential transform before multiplying them to generate the stimulus.

The key point made by **Figure 2** is that there are many possible gradients of illuminance and reflectance that can produce the same pattern of sensory input (luminance). These different explanations for a particular stimulus can only be distinguished by priors on the spatial and temporal characteristics of the reflectance and illuminance. In this example, the ambiguity about what caused the stimulus can be resolved if we believe, *a priori*, that the visual world is composed of isoreflectant surfaces, as opposed to surfaces that (implausibly) get brighter or darker nearer their edges or occlusions (as in the lower panels). Under this prior assumption, an observer who infers the presence of spatially extensive isoreflectant surfaces, and explains the edge at the centre with an illuminance gradient, would be inferring its most likely cause. The Cornsweet “illusion” is thus only an illusion because the experimenter has chosen an unlikely combination of illuminance and reflectance profiles. In what follows, we will exploit priors on the spatial composition and generation of

visual input to simulate the Cornsweet effect and the emergence of Mach bands.

The Bayesian approach to visual perception has been exploited in previous work (Yuille et al., 1991; Knill and Pouget, 2004). In addition, several other visual illusions have been explained using Bayesian principles, including motion illusions (Weiss et al., 2002), the sound-induced flash illusion (Shams et al., 2005), and the Chubb illusion (Lotto and Purves, 2001). Additionally, Purves et al. (1999) demonstrated the Bayesian nature of the Cornsweet illusion: when presented in a context implying an illuminance gradient and reflectance step, the Cornsweet illusion is elicited more easily.

In terms of the neuronal systems mediating the Cornsweet illusion; some authors have implicated subcortical structures: for example, Anderson et al. (2009) found that BOLD signal in the lateral geniculate nuclei (LGN) best correlated with perception of the Cornsweet illusion, although correlations were also seen in visual cortex. Furthermore, the illusion could be abolished if the stimulus was not presented binocularly, suggesting an origin before V1. Mach bands similarly have been attributed to retinal mechanisms (e.g., Ratliff, 1965); however, Lotto et al. (1999) have suggested a high-level contextual explanation for their appearance. Irrespective of the cortical or subcortical systems involved, we will assume, in this paper that the same Bayesian principles operate and, crucially, rest on a hierarchical generative model that necessarily implicates distributed neuronal processing at the subcortical and cortical levels.

OVERVIEW

This paper comprises three sections. The first describes a simple generative model of visual input that entails prior beliefs about how visual stimuli are generated and can be used to infer their causes. This model is used in the second section to simulate the perception of the Cornsweet illusion and contrast-dependent emergence of Mach bands. In the third section, we test the predictions of the simulations using a psychophysics study of normal subjects.

A GENERATIVE MODEL FOR THE CORNSWEET EFFECT

Our simulations are based upon the free-energy formulation of Bayes-optimal perception. Put briefly, this is based upon the notion that self-organizing agents minimize the average surprise (entropy) of sensory inputs through minimizing a free-energy bound on surprise. Here, surprise is just the improbability of sampling some sensory information, in relation to a (generative) model of how those sensations were produced. By adjusting the free parameters of the model, the sensory information can be explained or predicted and surprise minimized. Mathematically, surprise is the (negative) log evidence for a model of the world that comprises hidden variables (causes and states) that generate sensory information. We have described in many previous publications how this principle leads to active inference and Bayes-optimal perception (Friston et al., 2006; Friston, 2009; Feldman and Friston, 2010). Free-energy is a function of sensory samples and a probabilistic representation of what caused those samples. This representation can be cast in terms of the most likely or expected states of the world, under a generative model of how they conspire

to produce sensory inputs. In brief, once we know the agent's generative model, one can use the free-energy principle to predict its behavior and perception. In the present context, our focus will be on perception and the role of prior beliefs that are an inherent part of the generative model. In what follows, we describe the model and then use it to simulate perceptual inference and electrophysiological responses.

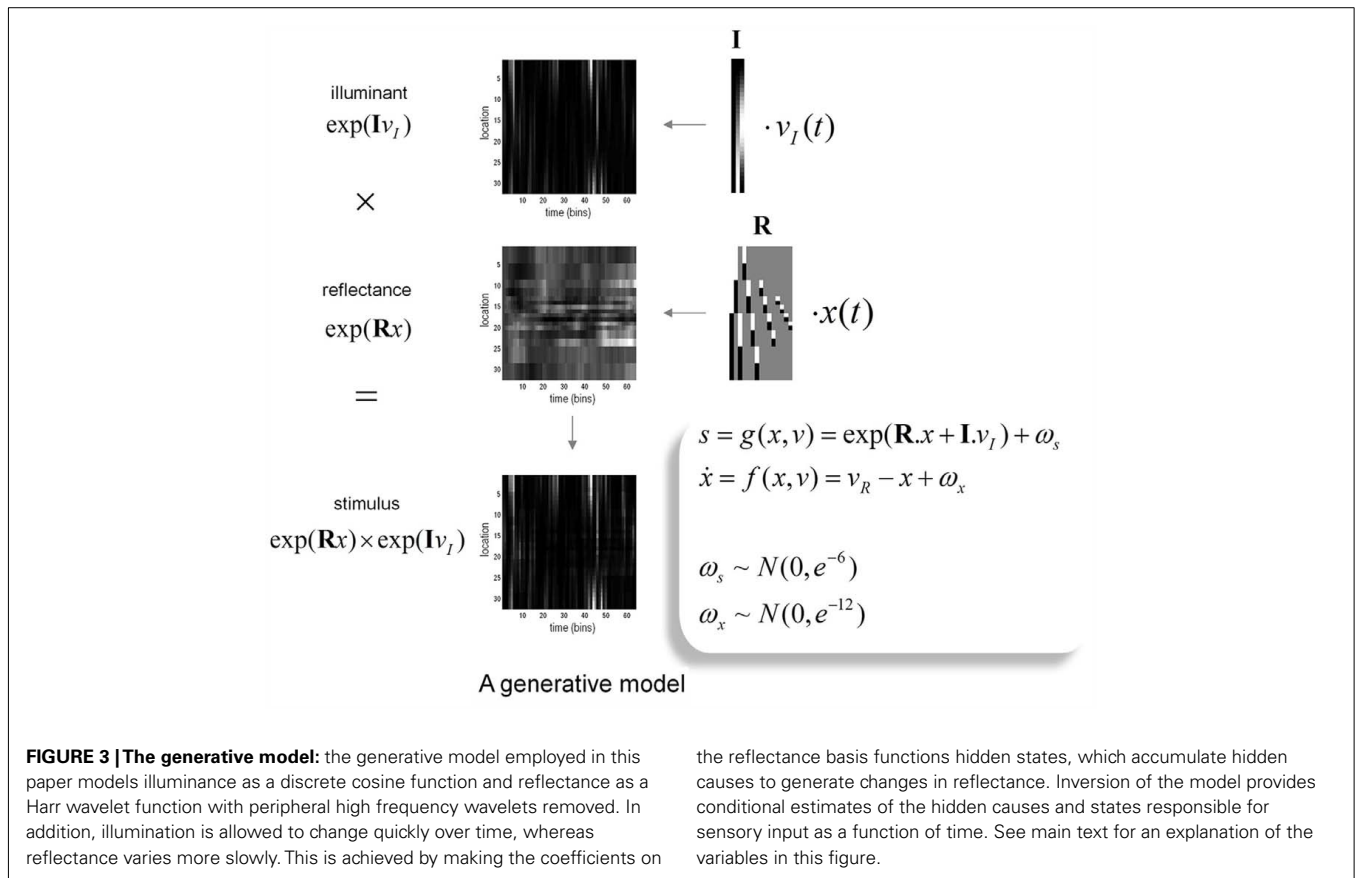
THE GENERATIVE MODEL

The generative model we used is straightforward: sensory input is the product of reflectance and illuminance, where illuminance varies smoothly over space but can fluctuate with a high frequency over time. Conversely, the reflectance profile of the visual world is caused by isoreflectant fields or surfaces that fluctuate smoothly in time. Crucially, the spatial scales over which these fluctuations occur have a scale-free nature, of the sort found in natural images (Burton and Moorhead, 1987; Field, 1987; Tolhurst et al., 1992; Ruderman and Bialek, 1994; Ruderman, 1997). To ensure positivity of the illuminant and reflectance we apply an exponential transform to the two factors before multiplying them (as in **Figure 2**). Equivalently, we can imagine the underlying causes (reflectance and illuminance) as being composed additively in log-space. This model is shown schematically in **Figure 3**, in terms of hidden causes and states. Mathematically, this model can be expressed as:

$$\begin{aligned} s &= g(x, v) = \exp(\mathbf{R} \cdot x + \mathbf{I} \cdot v_I) + \omega_s \\ \dot{x} &= f(x, v) = v_R - x + \omega_x \end{aligned} \quad (1)$$

Here, $s(t)$ are sensory signals generated from hidden states $x(t)$ and causes, $v(t)$ plus some random fluctuations $\omega(t)$. The difference between hidden causes and states is that states evolve dynamically, in response to perturbations by hidden causes – these dynamics are described by the equation of motion in the second equality above. Hidden causes $v(t) = (v_I, v_R)$ have been divided into those causing changes in luminance and those causing changes in hidden states that produce changes in reflectance. These hidden variables control the amplitude of spatial basis functions (\mathbf{R}, \mathbf{I}) encoding formal beliefs about the spatial scales of illuminance and reflectance. For the illuminant (\mathbf{I}) we use a low-order discrete cosine transform, while for the reflectant (\mathbf{R}) we use a low-order discrete wavelet transform.

The particular wavelet transform used here is a Haar wavelet set that has been thinned by removing high-order wavelets (with high spatial frequency) from the periphery of the visual field: This respects, roughly, the increasing size of classical receptive fields with retinotopic eccentricity. For simplicity (and ease of reporting the results), we restrict the simulations to a one-dimensional visual field. Because Haar wavelets afford local linear approximations to continuous reflectance profiles, the resulting reflectance has to be a mixture of isoreflectant surfaces at different spatial scales. To impose the scale-free aspect, we decrease the variance or, equivalently, increase the precision of the reflectance wavelet coefficients or hidden states in proportion to the order or spatial scale of their wavelet. This is implemented by placing a prior on the wavelet coefficients with the form $p(x_k) = N(0, e^{-3k})$, where k is the order of the wavelet. Neuronally, these basis functions could

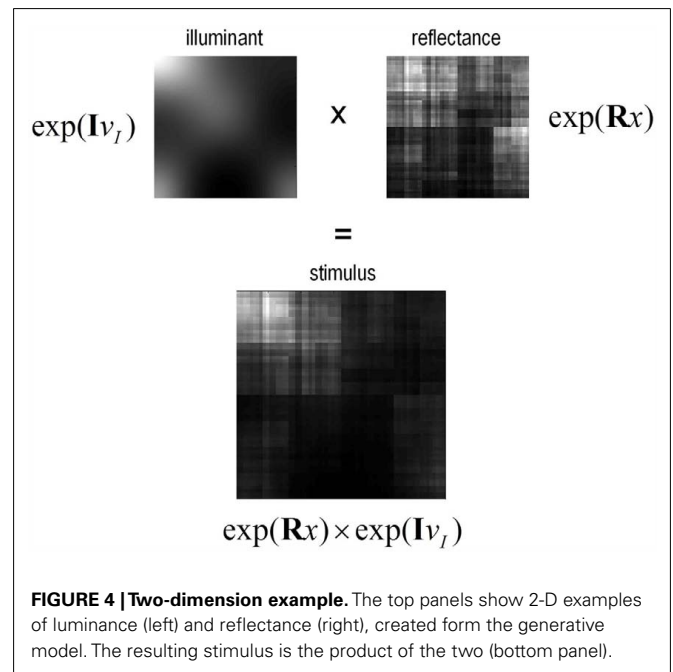


the reflectance basis functions hidden states, which accumulate hidden causes to generate changes in reflectance. Inversion of the model provides conditional estimates of the hidden causes and states responsible for sensory input as a function of time. See main text for an explanation of the variables in this figure.

stand in for a filling-in process such as that described by Grossberg and Hong (2006). Conversely, the illuminant is modeled as a mixture of smoothly varying cosine functions with a low spatial frequency. This is easily motivated by the fact that most sources of illumination are point sources, which results in smooth illuminance profiles. These were modeled here with the first three components of a discrete cosine transform (see Figure 4 for a graphical representation of the basis functions and how they are used to generate a stimulus).

By construction, this generative model of visual signals separates the spatial scales or frequencies of the illuminance and reflectance such that all the high frequency components are in the reflectance profile, while the low frequency components are in illuminance profile. Temporal persistence of reflectance is assured because the reflectance coefficients $x(t) \in \mathbb{R}^{16 \times 1}$ are hidden states that accumulate hidden causes $v_R(t) \in \mathbb{R}^{16 \times 1}$. This persistence reflects the prior belief that surfaces move in a continuous fashion. For simplicity, we mapped the hidden causes controlling illuminance $v_I(t) \in \mathbb{R}^{3 \times 1}$ directly to the stimulus (although this is not an important feature of our model). This can be thought of as accommodating rapid changes in illuminance of the sort that might be produced by a flickering candle.

Equation 1 defines our generative model in terms of the joint probability over sensory information and the hidden variables producing fluctuations in reflectance and illuminance. The fluctuations in the hidden causes are assumed to be Gaussian with a precision (inverse variance) of one, while the fluctuations in the



motion of the hidden states are assumed to have a log-precision of 12. Finally, we assume sensory fluctuations or noise with a log-precision of six. In the next section, we will manipulate the

log-precision of the sensory noise as a proxy for changing the contrast of the stimulus.

Figure 4 shows a snapshot of the sort of visual signals this generative model produces. Here, we have used the outer product of the discrete transforms above to generate a two-dimensional stimulus. We are not pretending that this is a veridical model of the real visual world. However, it is sufficient to explain the Cornsweet illusion and related effects by incorporating simple and plausible priors on the spatial scales over which illuminance and reflectance change. In the next section, we use this generative model to simulate perceptual and physiological responses to a stimulus, under the free-energy principle. This reduces to a Bayesian deconvolution of sensory input that tries to discover the most likely hidden causes and states generating that input.

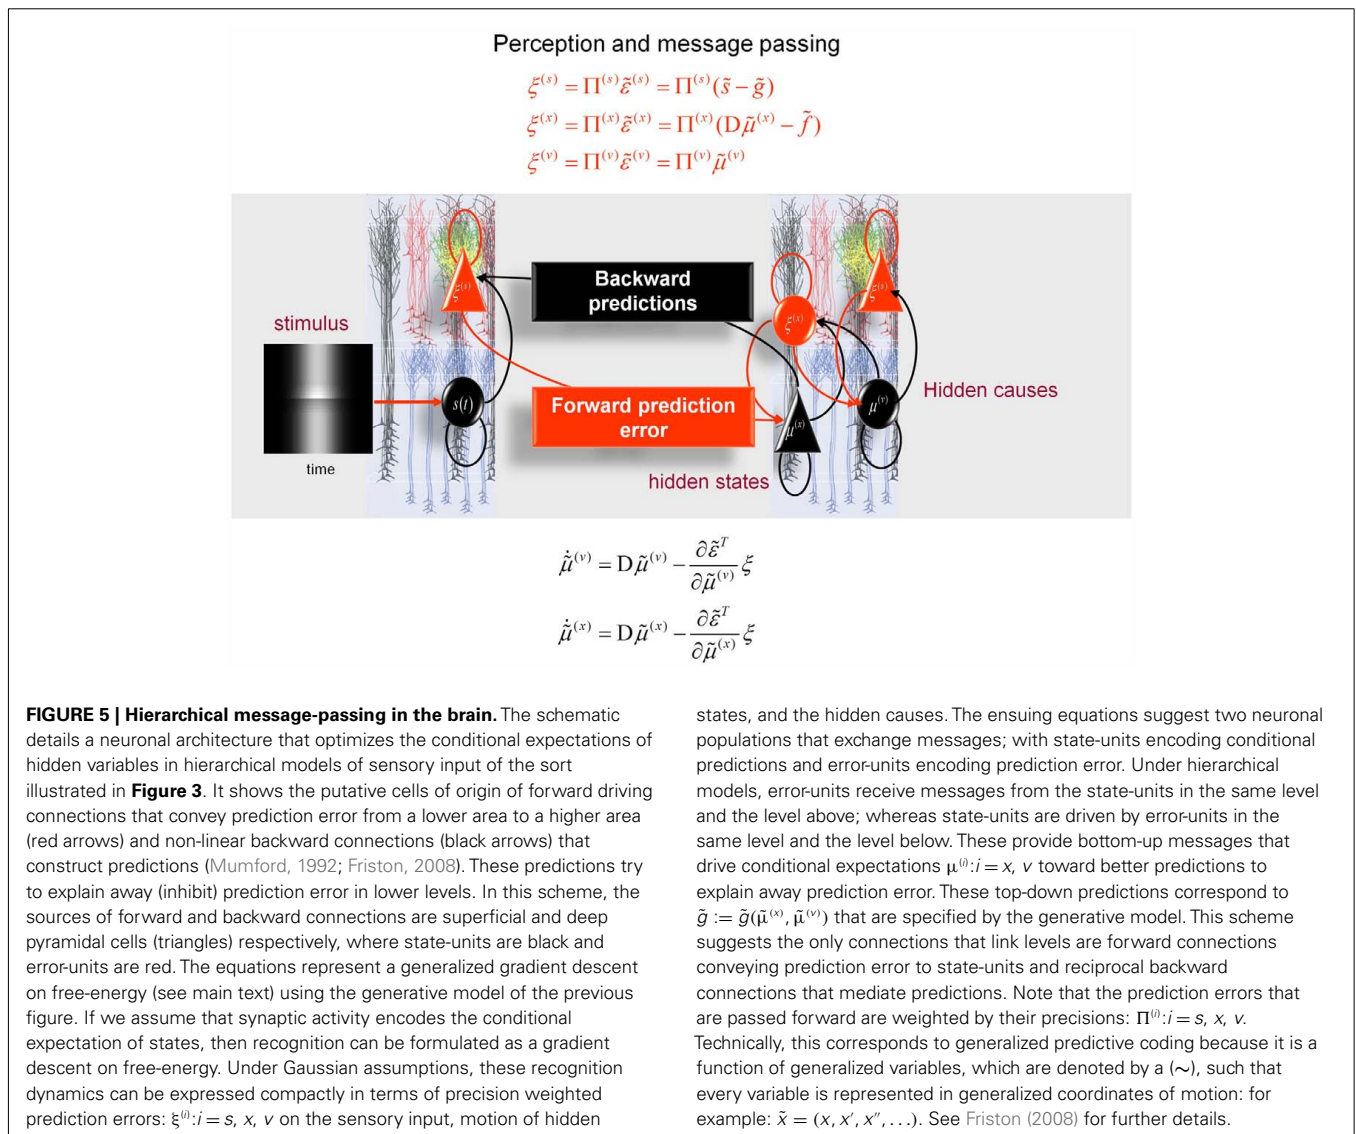
PERCEPTION AND PREDICTIVE CODING

This perceptual deconvolution can be regarded as the inversion of a generative model that maps from hidden causes (variables in the world) to sensory consequences. The inverse mapping corresponds

to inferring those variables by mapping back from the sensory consequences to the hidden causes and states. This can be implemented in a biologically plausible fashion using a generalized gradient descent on variational free-energy $F(\tilde{s}, \tilde{\mu})$; which is a function of (generalized) sensory states $\tilde{s}(t) = (s, s', s'', \dots)$ and the expected values $\tilde{\mu}(t) = (\mu_x, \mu_y)$ of hidden variables (see Friston, 2008 for details). In brief, this gradient descent corresponds to a Bayesian filtering, in which expected states of the world are continuously optimized using a prediction term and an update term:

$$\dot{\tilde{\mu}} = \mathcal{D}\tilde{\mu} - \frac{\partial}{\partial \tilde{\mu}} F(\tilde{s}, \tilde{\mu}) = \mathcal{D}\tilde{\mu} - \frac{\partial \tilde{\epsilon}^T}{\partial \tilde{\mu}} \xi \quad (2)$$

Under the simplifying assumption that probabilities are represented as Gaussian densities, this can be regarded as a generalized form of Kalman filtering, where the second (update or gradient) term can be expressed as a mixture of prediction errors (see the equations in **Figure 5**). This means that the generalized filtering



in Eq. 2 to corresponds to a generalized form of predictive coding. Predictive coding has become a popular metaphor for understanding perceptual inference in the visual system. For example, Rao and Ballard (1999) used predictive coding to provide a compelling explanation for extraclassical receptive field effects in striate cortex.

Put simply, in these simulations we assume that neural activity corresponds to the brain's representation of the most likely values of the hidden causes and states (hidden variables) and that these are continuously updated to minimize free-energy. The ensuing scheme has been discussed in terms of recurrent message-passing among different cell populations in hierarchical sensory cortex: see Figure 5 and Mumford (1992). This scheme rests upon the use of bottom-up prediction errors to optimize conditional estimates of hidden variables. These estimates are then used to produce top-down predictions that are compared with sensory input to form a bottom-up prediction error. In this context, the sum of squared prediction error can be regarded as free-energy. The recursive message-passing used in these schemes tries to minimize prediction error, such that the predictions approximate the true conditional or posterior estimates of the underlying hidden causes. It is this message-passing that we will stimulate in the next section and associate with neuronal responses, while using what they represent to predict how real subjects would respond behaviorally, in terms of their perceptual decisions or inference.

To simulate these responses we simply integrate or solve Eq. 2, using the functions $g(x, \nu)$ and $f(x, \nu)$ specified by a generative model in Eq. 1. These functions map hidden variables to sensory input and encode prior beliefs about the dynamics of hidden states. In short, by plugging the equations of our generative model in Figure 3 into the predictive coding scheme of Figure 5, we can simulate Bayes-optimal inference about the causes of sensations. Crucially, we can then reconstitute the posterior or conditional beliefs about these causes and associate these with percepts. In particular, we can take any mixture of the hidden variables and assess the posterior belief about that mixture. We will use this to quantify the Cornsweet and Mach band percepts, in terms of reflectance differences among different parts of the visual field. Note that the predictive coding scheme in Figure 5 weights the prediction errors by precision matrices. For example, the precision of sensory signals is $\Pi^{(s)} = I \cdot \exp(\gamma)$. These precisions are functions of log-precisions γ that encode the expected amplitude of random fluctuations.

SIMULATED RESPONSES

The simulated responses in Figure 6 were obtained by presenting the Cornsweet stimulus under uniform illumination. Here, the stimulus was presented transiently by modulating the illumination with a Gaussian envelope over time (see image inset). The resulting predictions are shown in Figure 6A as solid lines, while the red dotted lines correspond to the prediction error. These predictions are based upon the inferred hidden states and causes shown on the right and lower left respectively. The lines correspond to the posterior expectations and the gray regions correspond to 90% Bayesian confidence intervals. In terms of the underlying causes, the blue curve in Figure 6B is an estimate of the (log) amplitude of uniform illumination. This should have a roughly quadratic form (given the Gaussian envelope), peaking at around bin 30, which indeed

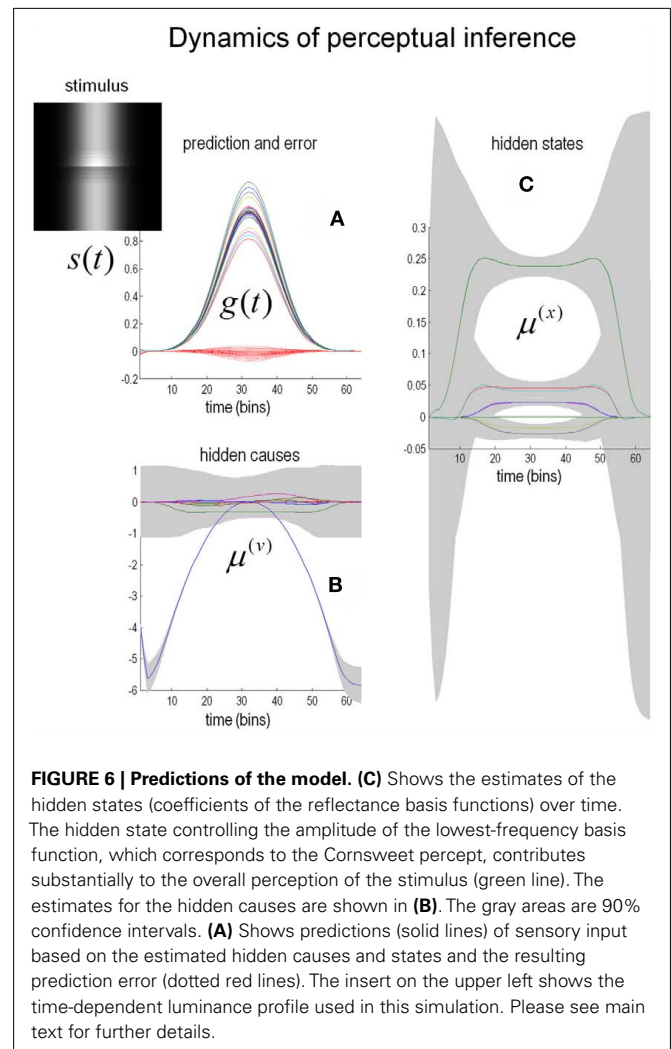
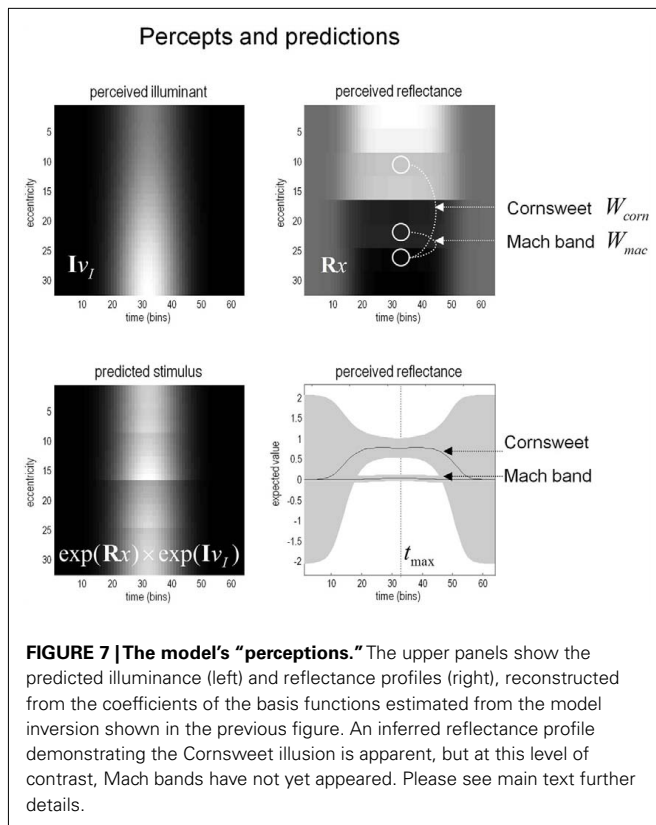


FIGURE 6 | Predictions of the model. (C) Shows the estimates of the hidden states (coefficients of the reflectance basis functions) over time. The hidden state controlling the amplitude of the lowest-frequency basis function, which corresponds to the Cornsweet percept, contributes substantially to the overall perception of the stimulus (green line). The estimates for the hidden causes are shown in (B). The gray areas are 90% confidence intervals. (A) Shows predictions (solid lines) of sensory input based on the estimated hidden causes and states and the resulting prediction error (dotted red lines). The insert on the upper left shows the time-dependent luminance profile used in this simulation. Please see main text for further details.

it does. The remaining causes that deviate from zero (Figure 6C) are the perturbations to the hidden states explaining or predicting changes in reflectance. These drive increases or decreases in the conditional expectations of the hidden states shown on the right. The green line is the coefficient of the second-order basis function splitting the visual field into an area of brightness on the left and darkness on the right. It can be seen that at the point of maximum illumination, there is an extremely high degree of confidence that this hidden state is bigger than zero. This is the Cornsweet percept.

The corresponding percepts in sensory space are shown in Figure 7 as a function of peristimulus time. The upper panels show the implicit reflectance and illuminance profiles encoded by the conditional expectations of hidden variables respectively. After an exponential transform (and multiplication) these produce the sensory predictions shown on the lower left. By taking a weighted mixture of the perceived reflectance in different regions of the visual field (shown by the white circles) one can estimate the conditional certainty about both the Cornsweet effect (differences in perceived reflectance on different sides) and the appearance of Mach bands (differences in perceived reflectance on the same side). The weights used to evaluate these mixtures are denoted as



W_{corn} and W_{mac} for the Cornsweet and Mach band effects respectively. The conditional expectation of these mixtures or effects $\mu_{\text{mac}} = W_{\text{mac}} \cdot \mu^{(x)}$ and their confidence intervals are shown on the lower right. At this level of visual contrast or precision (a log-precision of six), the Cornsweet effect is clearly evident with a high degree of certainty, while the confidence interval for the Mach band effect always contains zero. In other words, at this contrast (sensory precision) there is a Cornsweet effect but no Mach band effect. In the next section, we repeat the simulation above and record the conditional expectations (and confidences) about illusory effects at the point of maximum illumination for different levels of contrast.

CONTRAST OR PRECISION-DEPENDENT ILLUSORY PERCEPTS

Using the generative model and inversion scheme described above, we repeated the simulations over different levels of sensory precision. This can be regarded as a manipulation of contrast in the following sense: If we assume that the brain uses divisive normalization (Weber, 1846; Fechner, 1860; Craik, 1938; Geisler and Albrecht, 1992; Carandini and Heeger, 1994), the key change in sensory information, following an increase in contrast, is an increase in signal to noise; in other words, its precision increases (see the appendix of Feldman and Friston, 2010 for details). We use therefore a manipulation of the log-precision of sensory noise to emulate changes in visual contrast. It should be noted that we did not actually add sensory noise to the stimuli. The key quantity here is the level of precision assumed by the agent which, in these

simulations, we changed explicitly. In more realistic simulations, the log-precision would be itself optimized with respect to free-energy (see Feldman and Friston, 2010 for an example of this in the modeling of attention).

Figure 8 shows the results of perceptual inference under the Bayesian scheme described above. The only thing that we changed was the log-precision of the sensory input, from minus two (low) through to intermediate levels and ending with a very high log-precision of 16. The two graphs show the conditional expectation and 90% confidence intervals for the Cornsweet effect (upper panel) and Mach bands (lower panel) respectively, at the point of maximum illumination. It can be seen in both instances that under low levels of contrast (sensory precision) both effects are very small and inferred with a large degree of uncertainty. However, as contrast increases, conditional uncertainty reduces and, at a critical level, produces a confident inference that the effect is greater than zero (or some small threshold). Crucially, the point at which this happens for the Cornsweet effect is at a lower level of contrast than for the Mach bands. In other words, the Cornsweet illusion occurs first and then the Mach bands appear as contrast continues to rise. The explanation for this is straightforward; the Mach band illusion rests upon higher spatial frequencies in the generative model, which have a higher prior precision (encoding prior beliefs about the statistical – scale-free – structure of natural visual scenes). This means that there needs to be precise sensory evidence to change them from their prior expectation of zero. In short, at high levels of contrast or sensory precision, more and more fine detail in the posterior percept is recruited to provide the optimum explanation for the stimulus. Interestingly, as the contrast or sensory precision reaches very high levels, the veridical reflectance and illuminant profiles are inferred and, quantitatively, both the Cornsweet and mach band effects disappear. The three images show exemplar percepts, at low, intermediate and high levels of contrast respectively. The key difference in the spatial banding that underlies the Cornsweet and Mach band effects is evident in the difference between the intermediate and high levels of contrast.

The key prediction of these simulations is that we would expect subjects to categorize their percepts, following a brief exposure to a Cornsweet stimulus, differently at different levels of contrast. At low levels of contrast, we would expect them to categorize the stimulus as uniformly flat. At intermediate levels of contrast, we would expect them to categorize the stimulus as a Cornsweet percept, with isoluminant and uniform differences in the right and left hand parts of the visual field; while at higher levels of contrast one would expect the Mach bands to dominate and the stimuli would be categorized as possessing para-central bands. In principle, at very high levels of contrast, the subject should perceive the veridical stimulus. However, whether this level of contrast can be attained empirically is an open question. In the next section, we test these hypotheses psychophysically. We conclude this section by looking not at behavioral responses but at the neuronal responses implicit in the simulations.

SIMULATING NEURONAL RESPONSES

Figure 9 shows the prediction errors at low, intermediate and high levels of contrast. These are shown at the sensory level (upper row)

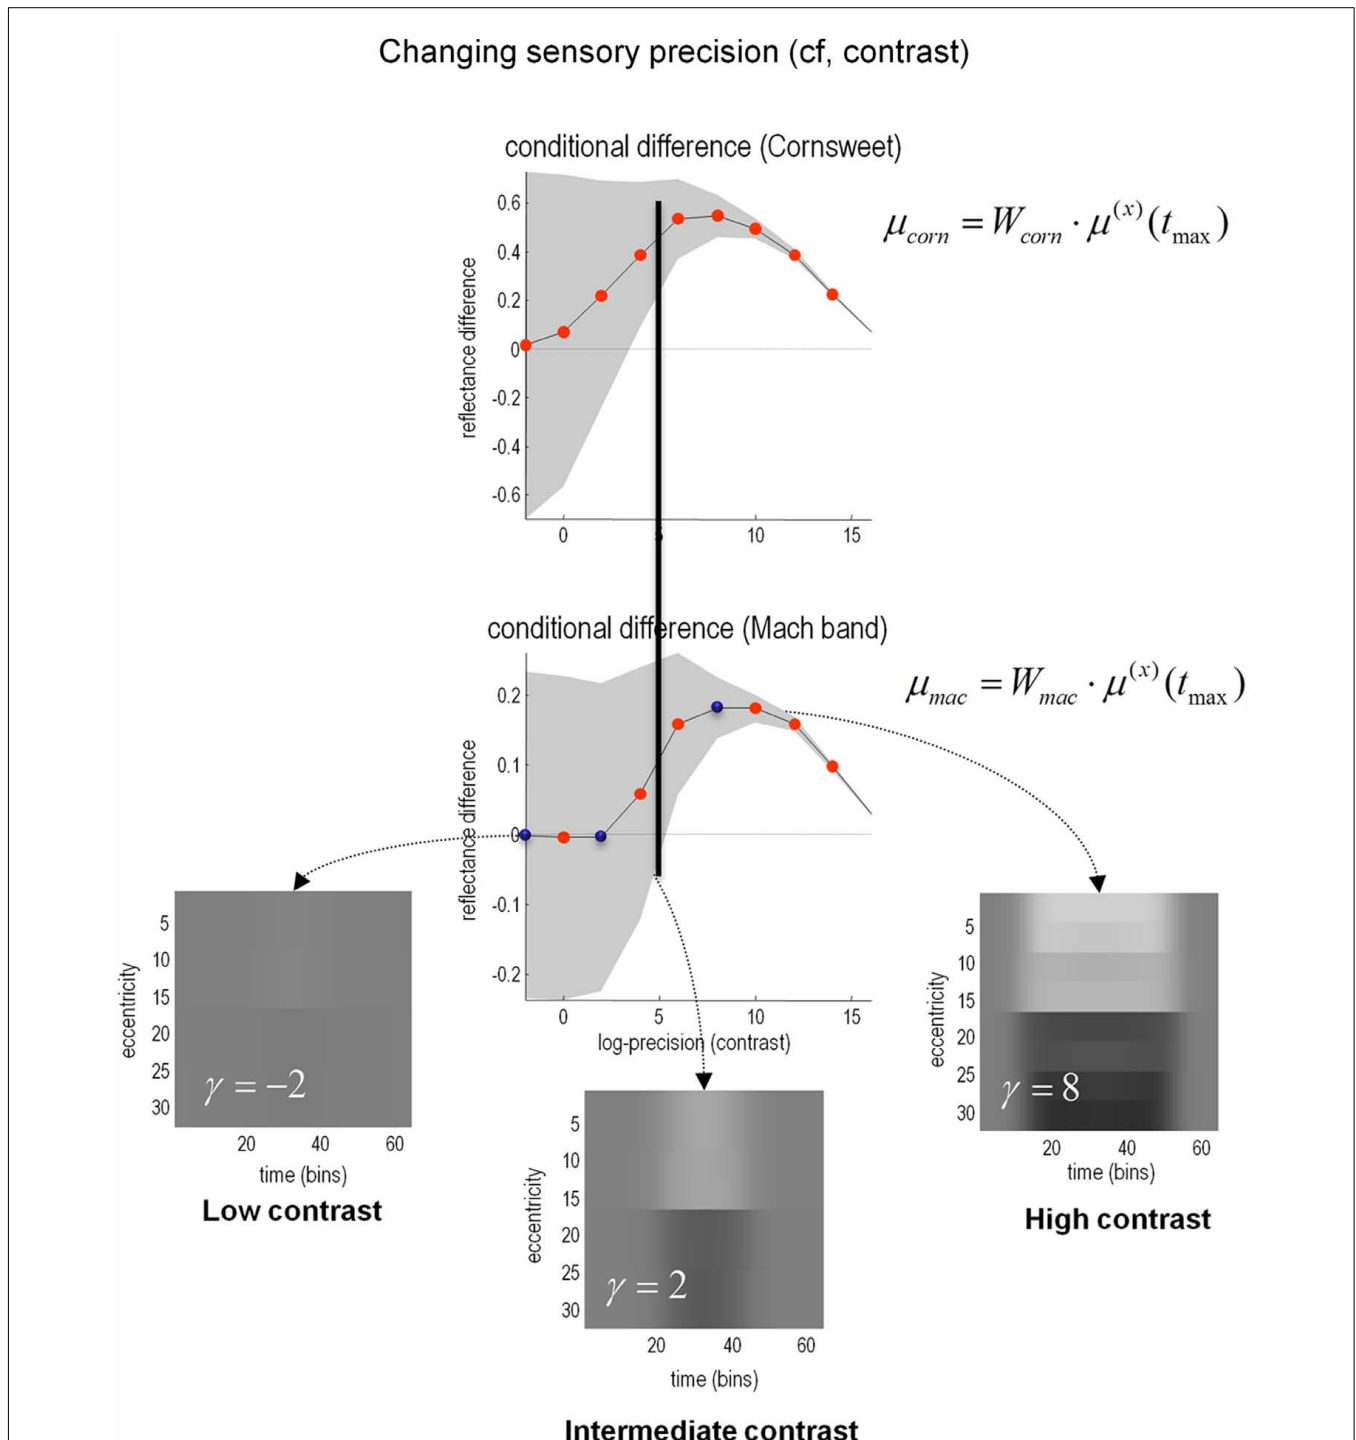
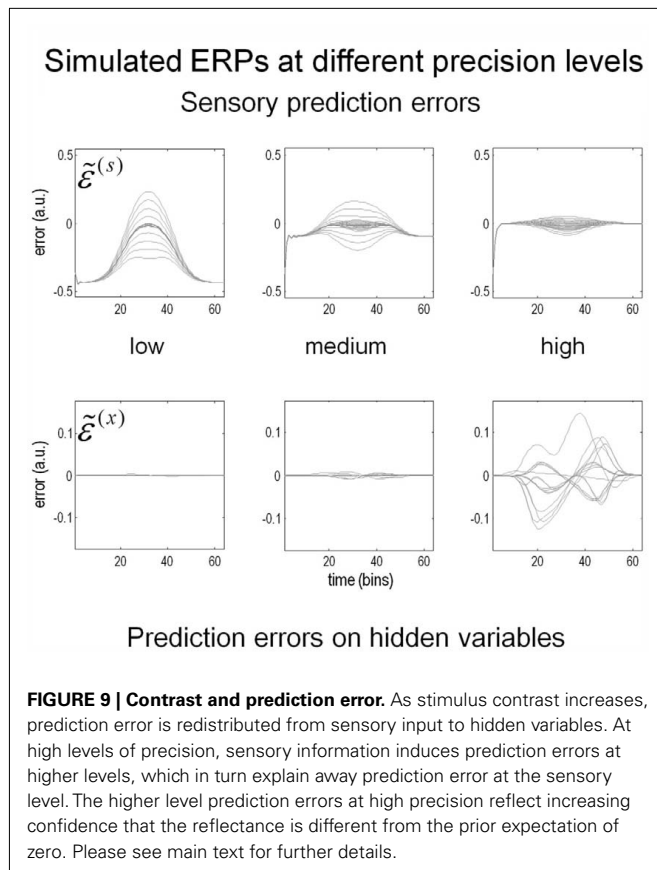


FIGURE 8 | The effect of contrast. In the results presented here the precision of the sensory data was used as a proxy for stimulus contrast. As precision increases, the strength of both the Cornsweet and Mach effects increases until, at very high levels of precision (contrast) the true luminance profile is perceived and the illusory percepts fade. Crucially, the Mach bands

appear at higher levels of contrast than the Cornsweet illusion. The inserts in the lower panels show the inferred reflectance is at different levels of contrast (indicated by the blue dots in the lower graph). The prediction errors associated with the processing of stimuli at these three levels are shown in the next figure. Please see main text for further details.

and at the higher levels of the hidden causes and states (lower row). The key thing to note here is that as contrast increases and the spatial detail of the posterior predictions increases, the sensory

prediction error falls. This is at the expense of inducing prediction errors at the higher level, which increase in proportion to the precision of sensory information. These higher prediction errors



are simply the difference between the posterior and prior expectations and reflect an increasing departure from a prior expectation of zero as contrast (the log-precision of sensory noise) increases. Although these results are interesting in themselves, they can also be regarded as a simulation of event related potentials. The reason that we can associate prediction error with observed electromagnetic brain responses is that it is usually assumed that prediction errors are encoded by the activity of superficial pyramidal cells (see **Figure 5**). It is these cells that are thought to contribute primarily to local field potentials and non-invasive EEG signals.

In the high-contrast condition, the prediction errors at the lower level are suppressed by the prediction of the presence of a Cornsweet stimulus. This sort of phenomenon has been demonstrated using fMRI (Alink et al., 2010; den Ouden et al., 2010); predictable stimuli cause less activation in stimulus-specific areas than unpredictable stimuli. However, the process simulated here is likely to produce more complicated neurophysiological correlates because of the confounding effect of precision; increased predictability (through increasing conditional confidence about the stimulus) will also increase estimates of precision. Since we believe that the prediction errors reported by superficial pyramidal cells are precision weighted, decreasing prediction error in lower sensory areas may be masked by the increasing precision of those errors. We will return to this and related issues in a subsequent paper looking at the neurophysiological correlates of contrast-dependent illusory effects. Here, we focus on psychophysical correlates:

A PSYCHOPHYSICAL TEST OF THEORETICAL PREDICTIONS

In this section, we report a psychophysics study of normal subjects exposed to the same stimuli used in the simulations above. We depart from the normal procedures for assessing illusions (which usually involve matching intensity differences) by using a forced choice paradigm. This is because we wanted to present stimuli briefly, for several reasons: First, brief presentation avoids the confounding effects of saccadic eye movements. Second, it allows us to prototype the paradigm for future use in electrophysiological (event related potential) studies, which require transient stimuli for trial averaging. Finally, a forced choice paradigm places constraints on the subject's choices that map directly to the model predictions. In what follows, we describe the paradigm and interpret our results quantitatively, in relation to the predictions of the simulations above.

SUBJECTS AND EXPERIMENTAL PARADIGM

We studied normal young subjects in accord with guidelines established by the local ethical committee and after obtaining informed consent. Eight participants (4 female) completed the Mach band paradigm; 19 (12 female) completed the Cornsweet paradigm.

EXPERIMENT 1 (CORNSWEET PARADIGM)

The Cornsweet illusion was assessed using a two-interval forced choice procedure. Subjects were presented with a (set contrast) Cornsweet stimulus and real luminance step for 200 ms (a Gaussian temporal envelope was not used), separated by an interval of 200 ms. One stimulus appeared to the left of fixation and one to the right; this was randomized across trials, as was the order of the stimuli. Subjects were asked to report the side on which the stimulus with the greatest contrast had appeared (**Figure 10**).

Six blocks were completed, using Cornsweet stimuli with Weber contrasts of 0.0073–0.734. A Quest procedure (Watson and Pelli, 1983) was used to select each step stimulus for comparison. The mean of the psychometric function was taken as the point of subjective equality between the Cornsweet stimulus and a real luminance. There were 200 presentations per block.

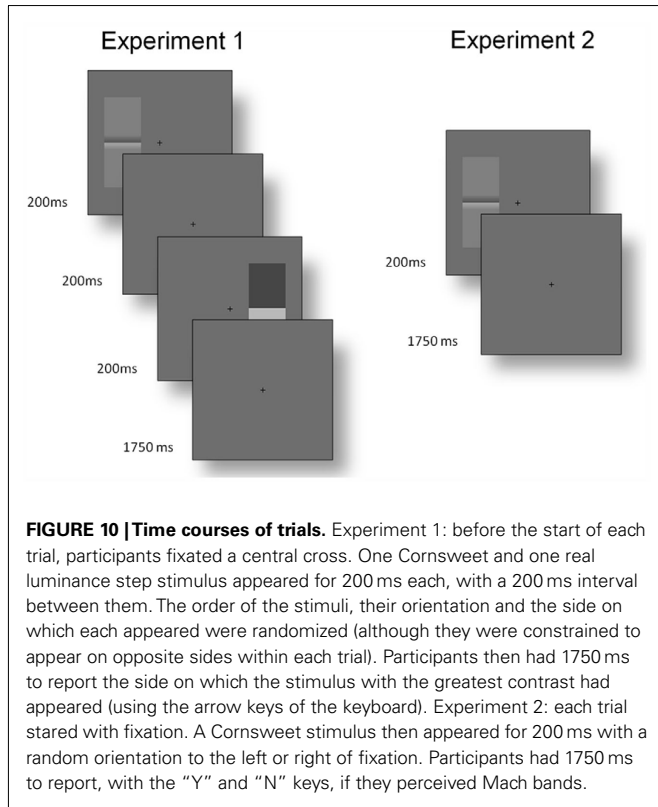
EXPERIMENT 2 (MACH BAND PARADIGM)

The same method could not be used to identify the strength of the Mach bands percept, as there is no non-illusory stimulus that can be used for matching. Consequently, we used a two-alternative forced choice paradigm: A single Cornsweet stimulus was displayed for 200 ms to the left or right of fixation and participants were asked to report if the stimulus contained Mach bands or not. Each participant completed six runs of 200 presentations at 10 Weber contrast levels from 0.0204 to 0.2038. The probability of reporting a Mach band was assessed as the relative frequency of reporting its presence over trials, within subject (**Figure 10**).

In both experiments, stimuli were displayed on an LCD monitor under ambient room lighting. Subjects were seated 60 cm away from the monitor, such that stimuli subtended an angle of 14.21° vertically and 6.10° horizontally, at 2.96°–7.06° eccentricity. The luminance ramp of the Cornsweet stimulus profile occupied 2.42°. Only the lowest levels of contrast the monitor was able to produce were employed; thus, luminance values were linearized *post hoc*.

RESULTS AND DISCUSSION

The results of the psychophysics experiments are shown in **Figure 11**, as a function of empirical (Weber) contrast levels. These results are expressed as the mean over all subjects and associated

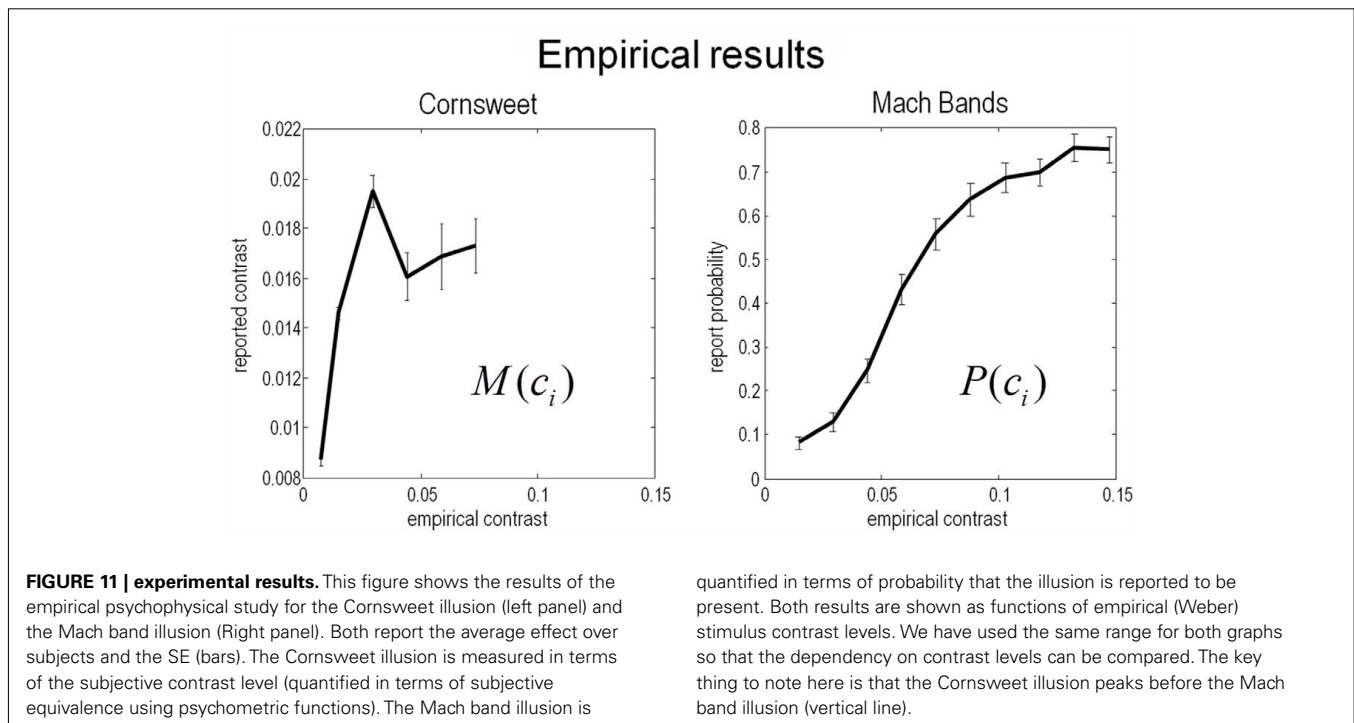


SE. The reported Cornsweet effect (as indexed by the point of subjective equality) peaked, on average over subjects, at a contrast of about 0.0025. At higher levels of contrast, as in the simulations, the effect fell quantitatively, plateauing at the highest contrast used in Experiment 1. Conversely, the probability of reporting a Mach band increased monotonically as a function of the empirical contrast, reaching about 75% at a Weber contrast of about 0.15.

Qualitatively, these empirical results compare well with the theoretical predictions shown in **Figure 9**: that is, the subjective or inferred Cornsweet effect emerged before the Mach bands, as contrast increases. We also see the characteristic “inverted U” dependency of the Cornsweet effect on contrast levels. The empirical profile is somewhat compromised by the small range of contrasts employed, however, this range was sufficient to disclose an unambiguous peak. Clearly, it would be nice to relate these empirical results quantitatively to the simulations shown in **Figure 9**. This presents an interesting challenge because the psychophysical data consist of reported levels of an effect and the probability of an effect for the Cornsweet and Mach Band illusions respectively. However, because our simulations provide a conditional or posterior probability over the effects reported, we can simulate both sorts of reports and see how well they explain the psychophysical data. This quantitative analysis is now considered in more detail.

A FORMAL BEHAVIORAL ANALYSIS

The simulations provide conditional expectations (and precisions) of both the Cornsweet and Mach band effects over a number of simulated (Weber) contrast levels, as modeled with the precision of sensory noise. This means that one can compute a psychometric function of contrast c that returns the behavioral predictions of both illusions respectively; namely, the level of the illusion and the probability of inferring a Mach band. To predict the reported

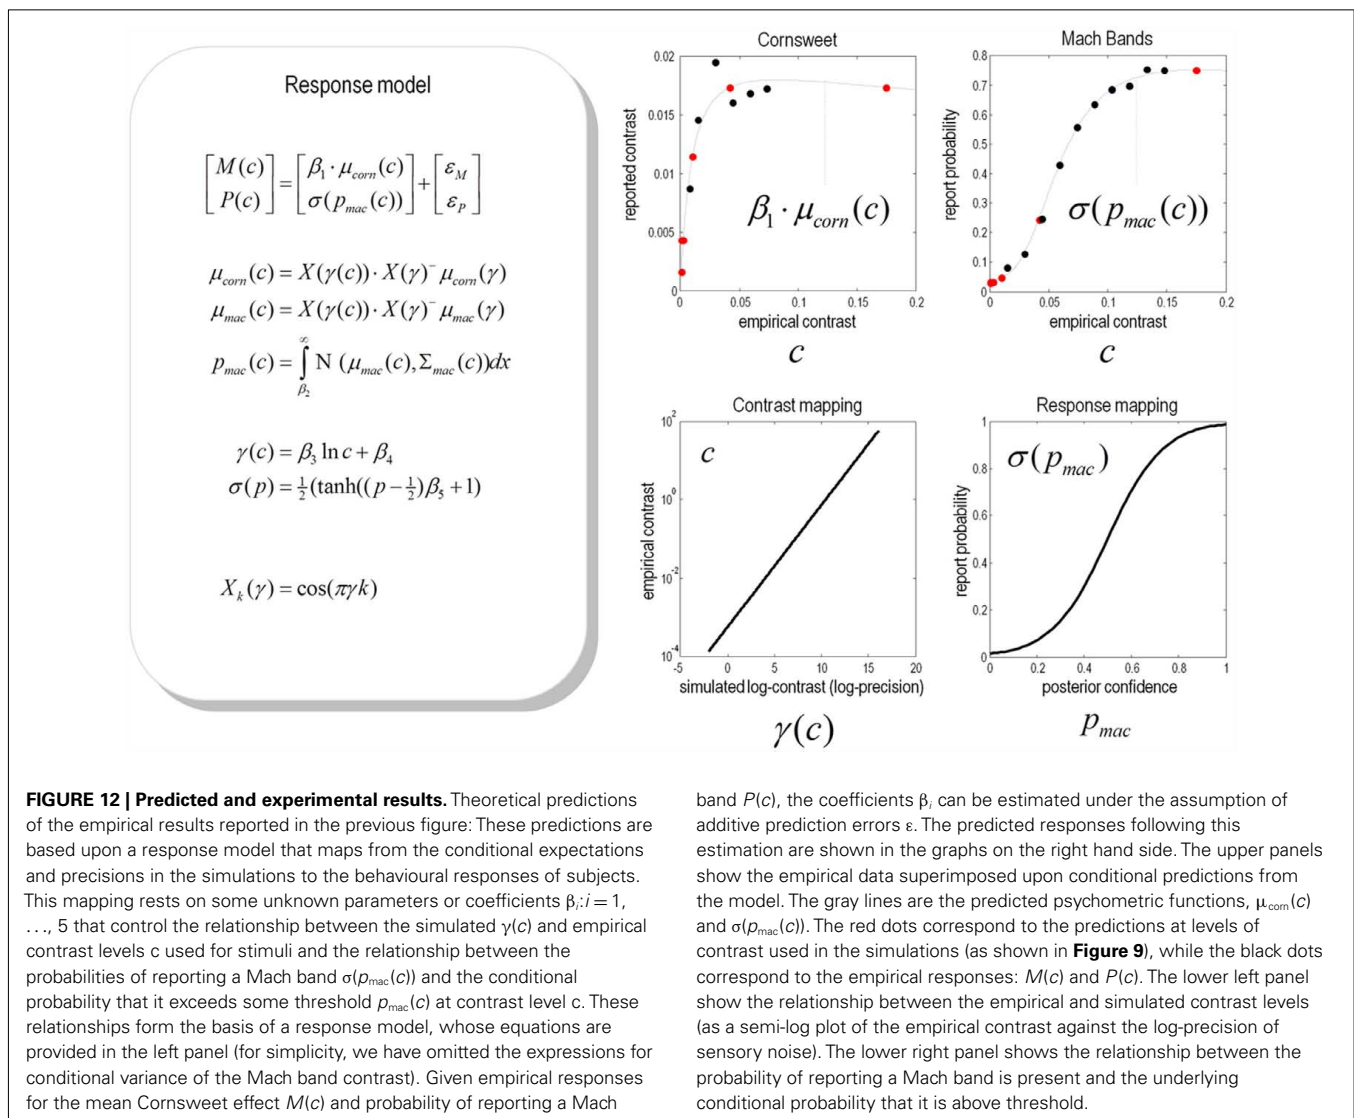


level of the Cornsweet illusion we can simply use the conditional expectation $\mu_{\text{corn}}(c)$ scaled by some (unknown) coefficient β_1 . To predict the probability of reporting the presence of Mach bands, one can integrate the conditional probability distribution over the Mach band effect above some (unknown) threshold β_2 . However, to do this, we need to know the relationship between the simulated and empirical contrasts:

As noted above, we used the log-precision of sensory noise γ to model log-contrast in accord with Weber's law. This means we can assume a linear relationship between the empirical log-contrast and simulated log-precision. This induces two further unknown coefficients – the slope and intercept (β_3 and β_4) that parameterise the relationship between the simulated and empirical contrasts. Finally, we need to relate the conditional probability of a suprathreshold Mach band effect to the probability of reporting its presence. Here, we assumed a simple, monotonic sigmoid relationship, under the constraint that when the conditional probability was 50:50, the report probability was also 50:50. The precise form of this mapping is provided in **Figure 11** (left panel) and has a

single (unknown) slope coefficient β_5 . These relationships provide a mapping between the results of the simulations and the observed responses averaged over subjects (under the assumptions of additive prediction errors). This is known as a response model and is detailed schematically in **Figure 11**. The predictions are based on the simulated responses in **Figure 9**, assuming a smooth psychometric function of contrast that was modeled as a linear mixture of cosine functions: $X_k(\gamma) = \cos(\pi\gamma k)$ for $k = 1, \dots, 6$). The coefficients of this discrete cosine set were estimated with ordinary least squares, using the responses of the model ($\mu_{\text{corn}}(\gamma), \mu_{\text{mac}}(\gamma)$) over different precision levels, at the time of maximum luminance.

Given the form of the relationships between the simulated and empirical contrasts and between the report probability and conditional probabilities for Mach bands, one can use the psychophysical data to estimate the unknown coefficients of these relationships: β_i for $i = 1, \dots, 5$. The results of this computationally informed response modeling are shown in the right panel of **Figure 12**. The upper panels show the same data as in **Figure 10** placed over the theoretical psychometric functions based on the



simulations of the previous section. These predictions are based on the mapping from simulated to empirical contrast levels (lower left) and the relationship between the probability of reporting a Mach band and the conditional confidence that it is present (lower right).

By construction, the relationship between the simulated and empirical contrasts is linear when plotted on a log-log scale. The slope of this plot suggests that the higher contrasts used in the simulations are, practically, not realizable in an empirical setting. This is because as the simulated contrast increases the corresponding empirical contrast increases much more quickly. The implication of this is that the contrasts employed in the psychophysics study correspond to the first few levels of the simulated contrasts. This means that it may be difficult to demonstrate the theoretically predicted attenuation of the Cornsweet illusion at very high levels of contrast.

The relationship between the report and conditional probabilities suggests that subjects have a tendency to “all or nothing” reporting; in the sense that a conditional confidence that the probability of reporting a Mach band is slightly greater than the conditional confidence it is above threshold. Conversely, subjects appear to report the absence of Mach bands with a probability that is slightly greater than the conditional probability it is below threshold. The resulting psychometric predictions (in the upper panels) show a remarkable agreement between the predicted and observed probabilities of reporting a Mach band. The correspondence between the predicted and empirical results for the Cornsweet illusion are less convincing but show that both asymptote to a peak level much more quickly than the probability of reporting a Mach band.

In summary, this analysis suggests that there is a reasonable quantitative agreement between the theoretical predictions and empirical results. Furthermore, in practical terms, it appears that the normal range of Weber contrasts that can be usefully employed corresponds to a relatively low level of sensory precision in the

simulations. This means that it may be difficult to demonstrate the “inverted U” behavior for the Cornsweet illusion seen in **Figure 9**. This is because it may be difficult to present stimuli at the ultra high levels of contrast required. Note that the empirical probability of reporting a Mach band does not decrease as a function of contrast level. This is to be anticipated from the theoretical predictions: increasing the contrast level increases the conditional precision about the inferred level of the Mach band effect, which means that the probability that is above threshold can still increase even if the conditional expectation decreases (as in **Figure 9**).

CONCLUSION

This paper has reviewed the nature of illusions, in the context of Bayes-optimal perception. We reiterate the notion that illusory percepts are optimal in that they may represent the most likely explanation for ambiguous sensory input. We have illustrated this using the Cornsweet illusion. By using simple and plausible prior expectations about the spatial deployment of illuminance and reflectance, we have shown that the Cornsweet effect emerges as a natural consequence of Bayes-optimal perception. Furthermore, we were able to simulate a contrast-dependent emergence of the Cornsweet effect and subsequent appearance of Mach bands that was verified psychophysically using a forced choice paradigm.

SOFTWARE NOTE

The simulations and graphics presented in this paper can be reproduced with the DEM toolbox distributed with the academic freeware SPM from <http://www.fil.ion.ucl.ac.uk/spm/>. The annotated files that implement the Cornsweet illusion simulations and the more general routines used for model inversion are provided as Matlab code.

ACKNOWLEDGMENTS

This work was funded by the Wellcome Trust. We thank Marcia Bennett for helping prepare this manuscript.

REFERENCES

- Alink, A., Schwiedrzik, C. M., Kohler, A., Singer, W., and Muckli, L. (2010). Stimulus predictability reduces responses in primary visual cortex. *J. Neurosci.* 30, 2960–2966.
- Anderson, E. J., Dakin, S. C., and Rees, G. (2009). Monocular signals in human lateral geniculate nucleus reflect the Craik-Cornsweet-O'Brien effect. *J. Vis.* 9, 1–18.
- Barlow, H. B. (1974). Inductive inference, coding, perception, and language. *Perception* 3, 123–134.
- Burton, G. J., and Moorhead, I. R. (1987). Color and spatial structure in natural scenes. *Appl. Opt.* 26, 157–170.
- Carandini, M., and Heeger, D. J. (1994). Summation and division by neurons in primate visual cortex. *Science* 264, 1333–1336.
- Cornsweet, T. (1970). *Visual Perception*. New York: Academic Press.
- Craik, K. J. (1938). The effect of adaptation on differential brightness discrimination. *J. Physiol. (Lond.)* 92, 406–421.
- Craik, K. J. W. (1966). *The Nature of Psychology*, ed. S. L. Sherwood. (Cambridge: Cambridge University Press).
- Dayan, P., Hinton, G. E., and Neal, R. M. (1995). The Helmholtz machine. *Neural Comput.* 7, 889–904.
- den Ouden, H. E., Daunizeau, J., Roiser, J., Friston, K. J., and Stephan, K. E. (2010). Striatal prediction error modulates cortical coupling. *J. Neurosci.* 30, 3210–3219.
- Fechner, G. T. (1860). “Spezielles zur Methode richtiger und falscher Fälle, in Anwendung auf die Gewichtsversuche,” in *Elemente der Psychophysik* (Leipzig: Breitkopf and Härtel).
- Feldman, H., and Friston, K. J. (2010). Attention, uncertainty, and free-energy. *Front. Hum. Neurosci.* 4:215. doi:10.3389/fnhum.2010.00215
- Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am. A* 12, 2379–2394.
- Friston, K. (2008). Hierarchical models in the brain. *PLoS Comput. Biol.* 4, e1000211. doi:10.1371/journal.pcbi.1000211
- Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends Cogn. Sci. (Regul. Ed.)* 13, 293–301.
- Friston, K., Kilner, J., and Harrison, L. (2006). A free energy principle for the brain. *J. Physiol. Paris* 100, 70–87.
- Geisler, W. S., and Albrecht, D. G. (1992). Cortical neurons: isolation of contrast gain control. *Vision Res.* 32, 1409–1410.
- Grossberg, S., and Hong, S. (2006). A neural model of surface perception: lightness, anchoring, and filling-in. *Spat. Vis.* 19, 263–321.
- Knill, D. C., and Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* 27, 712–719.
- Lotto, R. B., and Purves, D. (2001). An empirical explanation of the Chubb illusion. *J. Cogn. Neurosci.* 13, 547–555.
- Lotto, R. B., Williams, S. M., and Purves, D. (1999). Mach bands as empirically derived associations. *Proc. Natl. Acad. Sci. U.S.A.* 96, 5245–5250.
- Mach, E. (1865). Über die Wirkung der räumlichen Verteilung des Lichtreizes auf die Netzhaut. *Sitzungsber. Heidelb. Akad.*

- Wiss. Math. Naturwiss. Kl. 52, 303–322.
- Mumford, D. (1992). On the computational architecture of the neo-cortex. II. The role of cortico-cortical loops. *Biol. Cybern.* 66, 241–251.
- O'Brien, V. (1959). Contrast by contour-enhancement. *Am. J. Psychol.* 72, 299–300.
- Purves, D., Shimpi, A., and Lotto, R. B. (1999). An empirical explanation of the Cornsweet effect. *J. Neurosci.* 19, 8542–8551.
- Rao, R. P., and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2, 79–87.
- Ratliff, F. (1965). *Mach Bands: Quantitative Studies on Neural Networks in the Retina*. San Francisco: Holden-Day.
- Ruderman, D. L. (1997). Origins of scaling in natural images. *Vision Res.* 37, 3385–3398.
- Ruderman, D. L., and Bialek, W. (1994). Statistics of natural images: Scaling in the woods. *Phys. Rev. Lett.* 73, 814–817.
- Shams, L., Ma, W. J., and Beierholm, U. (2005). Sound-induced flash illusion as an optimal percept. *Neuroreport* 16, 1923–1927.
- Tolhurst, D. J., Tadmor, Y., and Chao, T. (1992). Amplitude spectra of natural images. *Ophthalmic Physiol. Opt.* 12, 229–232.
- von Helmholtz, H. (1866). *Treatise on Physiological Optics*, Vol. III, 3rd Edn (trans. by J. P. C. Southall 1925 Opt. Soc. Am. Section 26). New York: Dover.
- Watson, A. B., and Pelli, D. G. (1983). QUEST: a Bayesian adaptive psychometric method. *Percept. Psychophys.* 33, 113–120.
- Weber, E. (1846). “Der Tastsinn und das Gemeingefühl,” in *Handwörterbuch der Physiologie*, ed. R. Wagner (Leipzig: Wilhelm Engelmann).
- Weiss, Y., Simoncelli, E. P., and Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nat. Neurosci.* 5, 598–604.
- Yuille, A., Geiger, D., and Bulthoff, H. (1991). Stereo integration, mean field theory and psychophysics. *Network* 2, 423–442.
- Conflict of Interest Statement:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 30 November 2011; paper pending published: 21 December 2011; accepted: 07 February 2012; published online: 29 February 2012.

Citation: Brown H and Friston KJ (2012) Free-energy and illusions: the Cornsweet effect. *Front. Psychology* 3:43. doi: 10.3389/fpsyg.2012.00043

This article was submitted to *Frontiers in Perception Science*, a specialty of *Frontiers in Psychology*.

Copyright © 2012 Brown and Friston. This is an open-access article distributed under the terms of the Creative Commons Attribution Non Commercial License, which permits non-commercial use, distribution, and reproduction in other forums, provided the original authors and source are credited.