Coding of the contrasts in natural images by visual cortex (V1) neurons: a Bayesian approach

Mazviita Chirimuuta, Philip L. Clatworthy, and David J. Tolhurst

Department of Physiology, University of Cambridge, Downing Street, Cambridge CB2 3EG, UK

Received September 30, 2002; revised manuscript received January 28, 2003; accepted February 12, 2003

Individual V1 neurons respond dynamically over only limited ranges of stimulus contrasts, yet we can discriminate contrasts over a wide range. Different V1 neurons cover different parts of the contrast range, and the information they provide must be pooled somehow. We describe a probabilistic pooling model that shows that populations of neurons with contrast responses like those in cat and monkey V1 would most accurately code contrasts in the range actually found in natural scenes. The pooling equation is similar to Bayes's equation; however, explicit inclusion of prior probabilities in the inference increases coding accuracy only slightly.

OCIS codes: 330.1800, 330.4060, 330.4270.

1. INTRODUCTION

Oriented contrast edges are widespread in natural scenes[1–4] and are thought to be important cues in object recognition.[5] Consistent with this, the primary visual cortex (V1) contains many neurons that respond well to contrast features of specific orientations and spatial frequencies.[6–11] One class of V1 neuron—the simple cell[6,7]—has receptive field structure[12–14] remarkably similar to idealized basis-function sets[2,3,15] that might encode the edges found in natural scenes. This paper is concerned with the contrast coding employed by simple cells, and we develop a model of the code for natural contrasts as a Bayesian maximum *a posteriori* estimator.

It is presumed that it is the activity of neurons in V1 in particular (both simple and complex cells) that codes for the magnitudes of stimulus contrasts (for some interesting arguments in favor, see Refs. 16–18). However, V1 neurons may seem unlikely components of a contrast code for a number of reasons. First, responses of neurons in cat and monkey V1 are noisy in the sense that, when the same stimulus is presented repeatedly, the number of action potentials elicited will vary markedly from trial to trial, with variance approximately twice the mean.[19–23] Second, V1 neurons are not usually responsive over the entire contrast range (from 0 to 100%); rather, the contrast-response function is a steep sigmoidal curve which can be fitted with the Naka–Rushton (hyperbolic ratio) equation[23–26] [see Eq. (2) below]. At high and low contrasts the curve will be flat, with the semisaturation constant c_{50} determining the contrasts where the curve is accelerating. Any single neuron will be capable of providing only a small amount of information about stimulus contrast.[27]

However, the responses of different neurons do cover different parts of the overall contrast range[19,20,23,24]: c_{50} varies between neurons. The brain must be able to combine individual responses in such a way as to limit the effects of noise and to derive information about the full contrast range (i.e., by pooling the responses of neurons with different c_{50} values). Bayesian theory provides a means of achieving this.[23,28] An observer can attain the highest accuracy in contrast identification tasks by choosing the contrast c, which is most probable given the noisy response r, i.e., the contrast for which $P(c|r)$ is maximum. This is the maximum *a posteriori* rule, and $P(c|r)$ can be calculated from the standard Bayes formula,

$$P(c|r) = \frac{P(r|c)P(c)}{P(r)}, \tag{1}$$

where $P(r)$ is a normalizing term, $P(r|c)$ is the likelihood distribution (the probabilities of the occurrence of each response given a particular stimulus contrast), and $P(c)$ represents *a priori* knowledge of the relative occurrences of the different contrasts. Note that the decision stage does not include any gain function,[29] which might be required to model behavioral contrast identification if, for example, it is more costly for an observer to misidentify high contrasts than low ones. Such constraints are not relevant to the current purposes, which are to explore the relationships between prior information about natural scenes and V1 coding strategies.

We present a simulation of the problem, measuring the combined contrast identification performance of sets of model neurons. The $P(r|c)$ distribution for each neuron is estimated from the output of the Naka–Rushton equation, corrupted by multiplicative noise over a large number of trials. The responses for the set of neurons are pooled to form joint distributions $P(\mathbf{r}|c)$ where $\mathbf{r}$ is the set of responses of a set of neurons to a single stimulus. $P(c)$ is calculated from natural scenes and is found to be a nonuniform distribution with most frequent contrasts at the midrange on a logarithmic scale.[30–33]

We explore the effects on coding accuracy of the Naka–Rushton parameters and of inclusion of the *a priori* information; we either include $P(c)$ calculated from natural scenes or replace it with a uniformly flat distribution. In the V1 physiology literature, c_{50} has been found to take a range of values, but with a bunched distribution in which

values around 0.1 are most common.[23–25,34] The c_{50} distribution determines the contrast for which identification is most accurate, and we show that the cat or monkey distributions result in maximum accuracy over the contrast range most frequent in natural images, even when the simulation does not include natural priors. Explicit inclusion of the prior probability distribution $P(c)$ adds little to the accuracy of contrast identification.

2. METHODS

A. Model Parameters

Various authors have fitted V1 neuron contrast-response functions (responses to sine-wave gratings of different contrasts) with the Naka–Rushton[35] equation [Eq. (2)]. The function can be viewed just as an empirical fit,[24] as a component of models of V1 circuitry,[36,37] or as derived from information-theoretic analysis.[38] For our purposes it is immaterial whether the function is simply empirical, though some of our findings may warrant discussion in the light of the principled interpretation of the equation

$$R = R_{\max}[c^q/(c_{50}^q + c^q)], \tag{2}$$

where R is a neuron's response (described as a spike count per trial) to contrast c of the grating, $R_{\max}$ is the neuron's maximum response, c_{50} is the contrast at half saturation, and q is an exponent (usually about 2).[24] Contrast is defined as

$$c = (L_{\max} - L_{\min})/(L_{\max} + L_{\min}), \tag{3}$$

where L is the luminance of a point on the stimulus display.

For simplicity, it is assumed that the stimulus presented is optimal for exciting the neuron, so that there is no uncertainty about stimulus parameters other than contrast. The parameter c_{50} governs the contrast sensitivity of the neuron, and has been measured in cat and monkey V1.[23–26] Values of c_{50} are not evenly distributed in V1, but instead the most common values are near 0.1. The c_{50} value governs the horizontal position of the response curve and therefore which particular contrast values will fall under the steep portion of the curve to give the greatest differential mean response. It is the value of c_{50}, therefore, that determines the contrast range over which a neuron can usefully encode contrast. Figure 1 shows histograms of the c_{50} distributions for cat (D. J. Tolhurst, unpublished data) and monkey[39]; for methods used to collect data, see Refs. 40 and 41. These two data sets are similar to those published by others, and we are very grateful for the use of the especially detailed and numerous monkey data. The mean and median of the monkey c_{50}'s are slightly higher than those of the cat; that is, the monkey neurons are on average less sensitive.[23,24]

Our model neuronal populations are each made up of 16 neurons whose individual contrast-response functions are defined by Eq. (2). For all simulations $R_{\max} = 10$ and $q = 2$, but the different model populations differ with respect to c_{50} values, which are indicated as symbols in Fig. 2. The cat and monkey populations have 16 different c_{50} values chosen to give a representative sample of the cat and monkey data, respectively, although we excluded the monkey neurons with c_{50} values above ~1.0 (see Section

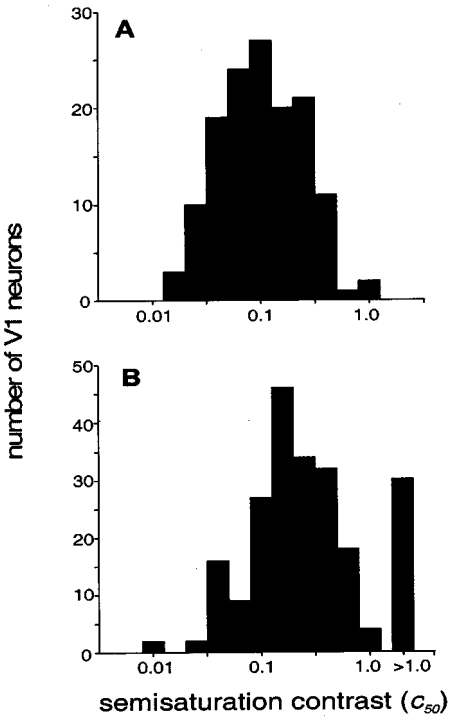

Fig. 1. Frequency of occurrence of V1 neurons with different c_{50} values: A, a sample of 138 neurons recorded in anaesthetized and paralyzed cats by D. J. Tolhurst in a number of studies (e.g., Ref. 19); B, a sample of 219 neurons recorded in anaesthetized and paralyzed monkeys (*Macaca fascicularis*); these data were very kindly provided by D. L. Ringach, M. J. Hawken, and R. Shapley.[39]

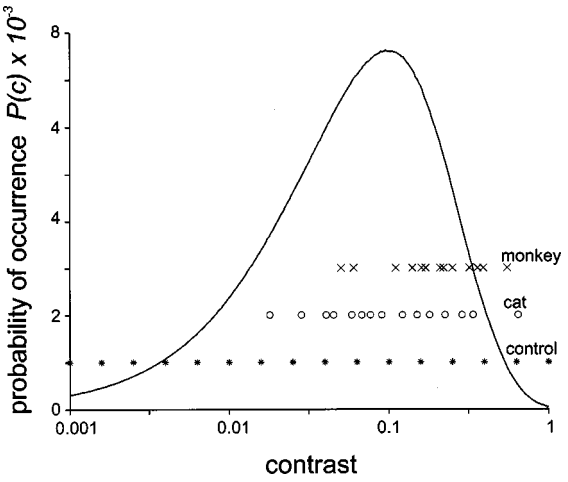

Fig. 2. Distribution of contrasts in natural scenes and c_{50} of models. Solid curve, estimation of the distribution of different contrasts in natural scenes. Gabor filters of eight different spatial frequencies at eight different orientations were convolved with 64 digitized images of natural scenes. 158×158 values were taken from the center of each 256×256 convolution, so that the distribution is based on more than 100×10^6 values. The number of occurrences of each contrast have been divided by the total number of occurrences of all contrasts to give the $P(c)$ distribution; there are 100 bins per log unit of contrast. The open circles show the 16 c_{50} values taken to represent the distribution of c_{50} measured in cats (Fig. 1B), the crosses show the 16 c_{50} values taken to represent the monkey distribution (Fig. 1A), the dots are 16 c_{50} values evenly spaced along the contrast axis.

4). The control distribution has 16 c_{50}'s that evenly span the contrast range from 0.001 to 1.0 in 15 logarithmic steps. Elsewhere we discuss the effects of changing $R_{\max}$, q, and the number of neurons in a model.[42]

B. Building the Contingency Table: $P(c/r)$

Equation (2) gives a deterministic response to contrast, whereas real V1 neurons respond with much trial-to-trial variability. In the model, noise is simulated by generating numbers from a Poisson probability distribution

$$P(x) = [\exp(-\mu)\mu^x]/x!, \tag{4}$$

where x is an integer and μ is the mean response for the contrast in question, calculated according to Eq. (2). However, these noisy responses will have variance equal to the mean rather than twice the mean (as observed in cat and monkey[20,22,23]). The greater noise is modeled by generating a second Poisson random number that uses the first noisy answer (x) as its mean to generate a new integer value r. This procedure is incorporated into the calculation of the response likelihood distribution:

$$P(r|c) = \sum_x \left[P(x) \frac{\exp(-\mu)x^r}{r!} \right]. \tag{5}$$

$P(r|c)$ is calculated for each of the neuron's possible responses given each of 311 contrasts from 0.001 to 1.26, in equal logarithmic steps (0.01 log units). This distribution is summed across contrast to give $P(r)$, the divisor of the Bayes equation:

$$P(r) = \sum_c P(r|c). \tag{6}$$

At this stage, the prior distribution of contrasts is taken to be flat, with each of the 311 contrasts equally likely. If $P(c)^*$ is the flat distribution, $P(c|r)^*$ is the individual posterior distribution without prior knowledge:

$$P(c|r)^* = \frac{P(r|c)P(c)^*}{P(r)}. \tag{7}$$

To combine the response cues generated by the 16 model neurons, the 16 individual *a posteriori* distributions are multiplied together to give the population distribution $P(c|\mathbf{r})^*$, the asterisk again signifying absence of prior knowledge. This pooling rule,[23] as described by Eq. (8), retains information about the differential firing rates of the neurons, in contrast to a pooling rule which simply sums action potentials[16,43]:

$$P(c|\mathbf{r})^* = \frac{\prod_n P(c|r_i)^*}{P(\mathbf{r})} = \frac{\prod_n P(c|r_i)^*}{\prod_n P(r_i)}, \tag{8}$$

where $\mathbf{r}$ is the set of n individual responses $\{r_1, r_2, r_3, \ldots, r_n\}$ and $P(c|r_i)^*$ is the probability of a given contrast given that the ith neuron gave a response r_i.

Finally, if we are to incorporate prior knowledge about the probability of occurrence of each contrast in natural scenes into the inference, the *a posteriori* distribution [Eq. (8)] is multiplied with the *a priori* distribution to give the resulting distribution $P(c|\mathbf{r})$:

$$P(c|\mathbf{r}) = P(c|\mathbf{r})^* P(c). \tag{9}$$

To estimate the natural priors distribution, we followed the method of Lauritzen *et al.*,[44] Peli,[45] and Tadmor and Tolhurst,[46] which measures the contrast of a region of a natural scene by convolving the region with a Gabor filter to represent a simple-cell receptive field,[47–49] dividing by the local mean luminance, and then finding the contrast of an optimal sinusoidal grating that elicits the same calculated response. We used odd-symmetric Gabor filters with a spatial-frequency bandwidth of 1.5 octaves[10,50] and with a variety of optimal orientations and spatial frequencies to determine the contrasts throughout 64 pictures of natural scenes digitized to 256×256 pixels and over 1000 gray levels.[51] The distribution is shown in Fig. 2.

C. Simulations

We ran simulations to evaluate the performance of the contrast coding of our three model neuron populations: cat c_{50}, monkey c_{50}, and control c_{50}. Performance was evaluated without [Eq. (8)] and with [Eq. (9)] inclusion of the prior probabilities.

We chose a set of 42 contrast values from 0.001 to 1.0, evenly spaced on a logarithmic scale, and simulated the presentation of 10,000 trials t at each contrast. On each trial, the responses of the 16 neurons in the model were chosen from a double Poisson distribution [see Eqs. (4) and (5)] with parameter μ set by the mean as calculated from the appropriate Naka-Rushton equation for that neuron. From the 16 responses $\mathbf{r}$, we find the most likely stimulus contrast $\hat{c}$ given that set of responses [c giving maximum $P(c|\mathbf{r})$]. The accuracy of coding at each contrast is

$$\text{accuracy} = \frac{t}{\sum_t [\log(\hat{c}) - \log(c)]^2}. \tag{10}$$

Accuracy is an inverse estimate of the variance of the model that acknowledges that the model responses will vary from trial to trial, and summarizes the range of estimated contrast responses in log space. It treats systematic errors in the same way as nonsystematic errors. In reality, a contrast code that produces systematic errors might be fine for certain tasks, such as contrast discrimination. The code might work well if it consistently and correctly identified which of two contrasts was the higher, even if it misidentified both. Accuracy was determined for each of the test stimulus contrasts.

We also calculated the mutual information $I(c; \hat{c})$ between the natural range of test contrasts c and the estimated contrasts $\hat{c}$ as an alternative metric for the performance of the models. When we calculated the mutual

information, each stimulus contrast between 0.001 and 1.0 was presented a number of times proportional to its estimated frequency of occurrence in natural scenes (Fig. 2). The total number of trials had to be more than 150,000 to minimize bias. The formula for mutual information is

$$I(c;\hat{c}) = \sum_{c}\sum_{\hat{c}} P(c;\hat{c})\log_2 \frac{P(c;\hat{c})}{P(c)P(\hat{c})}. \qquad (11)$$

3. RESULTS

The curve in Fig. 2 is the estimated distribution of contrast in natural scenes. The symbols below the curve represent the sets of c_{50}'s for each of our model populations (cat, open circles; monkey, crosses; control population, dots). Notice that most of the animal c_{50}'s fall under the peak of the natural contrast distribution, suggesting that the most frequently occurring contrasts will be encoded most reliably.[31]

The results of formal estimations of the accuracy of contrast coding are shown in Fig. 3 for A, cat; B, monkey; and C, control. Each plot shows accuracy (see Section 2) of the coding at each test contrast calculated without (open circles) and with (dots) inclusion of prior knowledge about the distribution of natural contrasts, i.e., using Eq. (8) and Eq. (9), respectively, to infer the contrast of a presented stimulus. The dashed curve (right ordinate) shows the distribution of natural contrasts $P(c)$ (or the "priors") redrawn from Fig. 2 but scaled to overlie the peak of the higher of the accuracy curves.

Note that the monkey and cat plots resemble the bell-shaped curves of the $P(c)$ distribution, with peaks very close to the peak of $P(c)$. On the other hand, the accuracy plots of the control population are flat and bear no resemblance to the $P(c)$ curve. This observation holds for simulations run with and without priors, suggesting that even before prior knowledge is included in the model, cat and monkey neurons best identify the contrasts most frequent in natural images. This optimization comes about just because of the bunched c_{50} distribution of the mammalian cortex (Fig. 1).

Surprisingly, there is not a dramatic difference between the accuracy results for the "priors" and "no-priors" simulations. For the monkey and cat populations, the peak accuracy is slightly higher when the priors are included, but with a reduction in accuracy at the highest contrasts so that the graph resembles more closely the right tail of the $P(c)$ curve. For the control population, the priors cause a bulge on the right-hand side of the plot so that maximum accuracy is near (but not on top of) the peak of $P(c)$. Again, there is a drop-off in accuracy at the highest contrasts.

Table 1 quantifies some of the above observations. It shows, for each neuronal population and for the priors and no-priors conditions, the area under the accuracy curve and the percentage of the accuracy area under the peak of the $P(c)$ curve (between contrasts 0.0186 and 0.295). The total area is an indication of the overall performance of the model, while percentage-under-peak indicates how well matched the neuronal population is to natural contrast stimuli. Observe that all three model populations have roughly the same total area, and that including natural priors in the simulation brings about an increase in area of about 10%, suggesting that knowledge

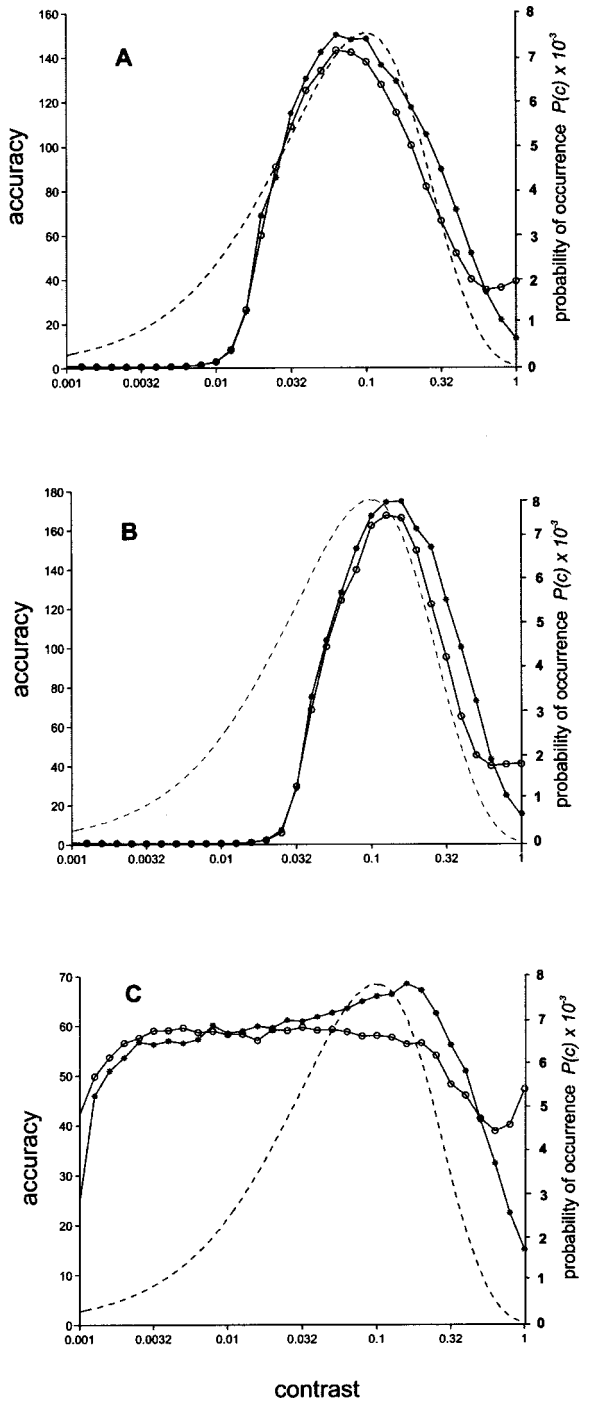

Fig. 3. Accuracy (left ordinate) of contrast coding for 16 model neurons with A, c_{50} sampled from the population of cat V1 neurons (circles in Fig. 2); B, c_{50} sampled from the population of monkey V1 (crosses in Fig. 2); C, c_{50} evenly spaced across the contrast axis from 0.001 to 1.0 (dots in Fig. 2). Note the difference in the scale of the left ordinate in C, control population, compared with the scales of A and B. Open circles, accuracy when stimulus contrast is inferred without including an *a priori* probability based on the frequency of occurrence of different contrasts in natural images; dots, accuracy when the inference is influenced by such prior knowledge; dashed curve (right ordinate) is the distribution of contrasts in natural images, the prior probability distribution (redrawn from Fig. 2).

Table 1. Summary of the Overall Performance of the Various Models of Contrast Coding

	Cat		Monkey		Control	
	No Priors	Priors	No Priors	Priors	No Priors	Priors
Total area[a]	1606	1769	1493	1674	1599	1642
Percentage area under $P(c)$ peak[b] (%)	85	87	85	85	44	49
Mutual information/bits[c]	2.40	2.35	2.20	2.19	2.12	2.12

[a] Area (in arbitrary units) under the accuracy curves of Fig. 3 in the contrast range 0.001 to 1.0.
[b] Percentage of the total area that falls in the contrast range 0.0186 to 0.295 (the range of contrasts for which probability of occurrence is greater than or equal to half of the maximum probability).
[c] Mutual information between inferred contrast and actual contrast when the models are stimulated with contrasts drawn from the distribution of contrasts in natural images.

of the priors does improve overall accuracy. The cat and monkey populations have a much greater proportion of the total accuracy concentrated around the peak of the distribution of contrasts in natural scenes, even in the no-priors condition (85% for both the cat and monkey compared with 44% for the control). This means that if, instead of presenting each stimulus contrast an equal number of times, one were to present stimulus contrasts with a frequency weighted by their frequency of occurrence in natural scenes, the cat and monkey populations would outperform the control population. This is the condition being tested when we calculate the mutual information between the natural scene contrasts and the model responses [Eq. (11)]. Table 1 shows the mutual information for the three model populations of neurons, with and without inclusion of the priors. In fact, these values show rather little difference in mutual information for the various models, although the rank order seems correct: The model of cat neurons gives the greatest accuracy and the highest mutual information, while the control model is least accurate and shows the lowest mutual information.

Adding priors to the simulation does not bring about a substantial change in the proportion of the accuracy curve under the $P(c)$ peak for any of the animal models, though this value does increase in the case of the control model. Priors do not make a great difference to the mutual information sum.

4. DISCUSSION

We have modeled how populations of VI neurons might encode the contrasts found in natural scenes and estimated how accurately each contrast would be encoded. Model populations of cat and monkey V1 neurons produce accuracy graphs that closely follow the estimated distribution of contrasts in natural images $P(c)$. This is largely because the steep portions of the neurons' contrast-response functions tend to occur around the contrasts most commonly encountered in natural scenes, rather than being evenly spaced across the contrast continuum. Most neurons have their semisaturation contrast (c_{50}) within half a log unit of the most frequently encountered contrast (0.1).

It is tempting to say, then, that the specific clustering of animal c_{50} values around a contrast of 0.1 is an adaptation to the task of encoding natural contrasts; or, in other words, that the selection of a particular c_{50} distribution might be a way in which the brain can build in prior knowledge of natural scenes when making decisions about contrast. This follows from the observation that including prior probabilities $[P(c)]$ in our control model (with evenly spaced c_{50}) shifts the accuracy peak closer to the $P(c)$ peak, whereas for the models of cat and monkey neurons, accuracy is maximum at the $P(c)$ peak even before explicit inclusion of prior probabilities. That is to say, there is a similarity between the effects of priors on the control model and the properties, without explicit priors, of the mammalian models. This matching of performance with stimulus is the desired effect of prior knowledge because, given finite coding resources, it results in a code which performs best in the conditions most common in the natural environment.[52] In the case of contrast coding, V1 neurons will provide the most reliable description of the contrast of edges or of other features when they are at their most usual contrast levels. This matching could arise genetically or developmentally through early cortical plasticity, as for other stimulus dimensions.[53,54]

One should, however, bear in mind the experimental finding[55,56] that the exact value of c_{50} adapts to the range of contrasts with which the neuron is being stimulated in the experiment. For example, presenting a series of low-contrast gratings will result in the measurement of a lower c_{50} than found if stimulating is with high contrasts. In this paper, we have used values of c_{50} taken from neurons stimulated over the whole contrast range. If neurons were presented with a set of gratings with contrasts chosen in proportion to their occurrence in natural scenes, we might expect that the c_{50}'s would show even more clustering around a contrast of 0.1 and an even more convincing match between V1 c_{50}'s and natural contrasts.

This conclusion as to the match between c_{50} and the distribution of contrasts in natural scenes follows from observing the shape of the accuracy curve, and further consideration should be given to the appropriateness of the accuracy metric that we have proposed. This particular accuracy measure [Eq. (10)] was devised to be equivalent to the inverse of the variance of the model's contrast guesses (in log space) around the correct value. The measure was thought to be appropriate for a stochastic model such as this, where the model's answers will form a distribution around the correct value and the vari-

ance is a measure of how far from this value the model answer tends to be.

An absolute measure, such as the percentage of trials in which the model gives the correct answer, would not be appropriate because the model is so rarely exactly right (at best, on 7% of trials for the animal c_{50} populations and on 5% for the control population), though it guesses with greater and lesser degrees of error information that is lost to the absolute measure. Even so, the percentage-correct-versus-contrast plot resembles the accuracy curve, with a flat curve for the control population and a bell curve peaking around 0.1 for the animal c_{50} populations. These similarities can be seen in other measures such as the inverse standard deviation, although the spread of the curve may change. Therefore, we argue, it is appropriate to use the shape of the accuracy curve (especially its peak) as an indication of fitness to natural-scene statistics.

Another means of assessing the fitness of the pooled neural responses to the distribution of contrasts in natural scenes is to plot the contrast-response function of the population, normalized to a maximum of 1, together with the cumulative distribution of contrasts in natural scenes.[30] If the responses are well matched, the curves should overlie. Elsewhere[42] we present the finding that, for our model, the match to the cumulative $P(c)$ curve is good if one plots the square root of population response (equivalent to the signal-to-noise ratio in our multiplicative noise model). This finding agrees with those inferred from the shape of the accuracy plot.

In this paper we present the mutual information between natural-scene contrasts as a quantitative measure of the neural code's fitness, and the model output is between 2.1 and 2.4 bits for all c_{50} populations. These values are concordant with the results of a psychophysical experiment performed in our laboratory (D. J. Tolhurst and J. Brown, unpublished observations) which found humans capable of correctly identifying the contrasts of 7.93-cycle/deg gratings to just four contrast levels (i.e., 2 bits).

The mutual information does not show a great difference between the mammalian and the control populations: 0.3 and 0.1 bits for the cat and monkey populations, respectively. This is a rather small improvement, even bearing in mind that information is a logarithmic measure, and it is probably due to the fact that the control population has quite good accuracy over virtually all of the contrast range. It therefore outperforms the cat and monkey models at all contrasts except for the critical 1-log-unit range centered about a contrast of 0.1. Some of these contrasts, though not at the peak of the $P(c)$ distribution, occur frequently enough to boost the mutual information sum in favor of the control model. Still, the mammalian populations do seem to have an advantage over the control, and it may well be the case that there is evolutionary benefit for improvements in coding efficiency that are actually quite small in terms of mutual information.[57]

It is interesting that explicitly including priors in the simulations makes little difference to the accuracy plots and virtually no difference to the mutual information figures. But note that the prior probabilities are never very small even at the contrast extremes; the tail at low contrasts is quite shallow. Then, because of the multiplicative pooling rule, the effect of the prior probabilities is dependent on the number of neurons. That is, priors become less 'effective' with a population of model neurons because the prior distribution becomes just one part (out of 17 in our models) of a long product of individual neuron likelihood distributions. In order to avoid the dependency of the effect of priors on n (the number of neurons) one might be tempted to use the following in place of Eq. (9):

$$P(c|\mathbf{r}) = P(c|\mathbf{r})*P(c)^n, \tag{12}$$

which follows from a more direct interpretation of Bayes's rule [Eq. (1)]: the cue provided by each neuron is weighted by the prior probabilities. However, this results in an accuracy plot which tends to spike at the peak in the priors distribution and, as the number of neurons increases, this plot becomes identical for all c_{50} populations, suggesting that the actual responses of the neurons are lost in favor of overreliance on the raised power of the prior information. As this is not a desirable outcome, a compromise might be to use

$$P(c|\mathbf{r}) = P(c|\mathbf{r})*P(c)^{n/a}, \tag{13}$$

where a is the fixed number of neurons to be multiplied by the priors, and n would have to be evenly divisible by a. This formulation could increase the efficacy of the priors but the choice of a would be totally arbitrary, as is the choice of n, which was taken to be 16 for all of our simulations.

We modeled the contrast-response functions of all the neurons as having the same maximum firing rate (10 action potentials) and the same steepness parameter ($q = 2$). We investigate the effects of these parameters elsewhere.[42] A different choice of maximum firing rate would change the accuracy and information values that we cite, but it would not change the comparative behavior of the models nor our conclusions about the effects of explicit inclusion of prior probabilities in the decision rule. We chose $q = 2$ because this is very close to the mean value actually found[24] and it is interesting theoretically.[36,38]

Figure 1B shows that there are significant numbers of neurons in monkey V1 with c_{50}'s from 1 to 12.[39] These contrast values are literally impossible to achieve with sinusoidal gratings (the experimental tools used to study the neurons). These neurons could not have had q as large as 2, or they would not have generated responses to achievable contrasts of less than 1.0 in the experiments. This reminds us that real neurons may differ quite markedly from each other and that it is a great simplification to model all the neurons as having the same "average" behavior.

It would be feasible to model a population of cat or monkey neurons in which all the parameters for the Naka–Rushton equation (2) are taken from neurophysiological data (i.e., for the case where each model neuron represents a particular unit whose contrast responses can be fitted by the Naka–Rushton). It would be interesting to see whether such a model has increased mutual information or a better fit to the natural-scene contrasts curve. The focus of this paper has been on the effects of priors

and the c_{50} parameter, and in this case including such extra detail would make the interpretation of the model results more difficult. For example, one would not know if the shape of the accuracy curve were due to a difference in c_{50} across neuronal populations or to the combined differences of a number of parameters.

Another of our simplifying assumptions is that all model neurons are optimally tuned to the stimuli, and that there is no uncertainty about the stimulus parameters other than contrast. This assumption is probably not relevant to the conclusions drawn about the animal c_{50} distribution because it has been shown experimentally that c_{50} remains constant under optimal and nonoptimal stimulus conditions.[24] That is, the contrast-response function of a neuron presented with a stimulus to which it is poorly tuned will begin to respond and will saturate at the same contrast as the contrast-response function under optimal stimulation, but at a lower maximum firing rate. To model such an effect, one would need only to add another set of neurons with lower $R_{\max}$ than the others without altering the c_{50} distribution. The resulting ambiguity as to whether a low firing rate is due to a low contrast or a suboptimal stimulus can be resolved by comparing firing rates across the groups of neurons.

When discussing these results one must bear in mind that contrast identification is one of many tasks in which V1 neurons will be engaged, and it may well be one of the less ecologically relevant ones. To illustrate: An edge feature will have an orientation and width as well as a contrast. If an animal's task is to recognize the object of which that feature is a part, accurate representation of orientation and spatial frequency will probably be more important than accurate contrast identification. So if there is a trade-off between accurate orientation and spatial-frequency coding versus contrast—which will arise given that contrast is best coded while the neuron is in the steep part of the contrast curve, whereas other features may be best coded when the neuron is saturating—contrast will lose out. The psychophysical finding mentioned above, that humans perform poorly in a contrast-labeling task, suggests that this is the case in human vision.

5. CONCLUSIONS

This paper has described a Bayesian approach to the study of the V1 contrast code. Such considerations as selection of model neuronal parameters and implementation of prior knowledge put into sharp focus the problem of the biological implementation of Bayesian analysis: How is it that the brain could retain knowledge of probability distributions or apply the maximum *a posteriori rule*? The prior-knowledge interpretation of c_{50}—the model parameter found to affect significantly the fit of model performance to natural-scene statistics—is interesting precisely because it suggests a way in which the brain can conduct one part of the Bayesian analysis.

ACKNOWLEDGMENTS

M. Chirimuuta was supported by a studentship from the Medical Research Council (UK). We are very grateful to D. L. Ringach, M. J. Hawken, and R. Shapley for allowing us to use their unpublished data and for performing the analyses on our behalf. We are grateful to the anonymous reviewers for the helpful comments on the original manuscript.

M. Chirimuuta may be reached by e-mail at mc325@cam.ac.uk.

REFERENCES

1. D. L. Ruderman, "The statistics of natural images," Network **5**, 517–548 (1994).
2. B. A. Olshausen and D. J. Field, "Sparse coding with an overcomplete basis set: a strategy employed by V1?" Vision Res. **37**, 3311–3325 (1997).
3. J. H. van Hateren and A. van der Schaaf, "Independent component filters of natural images compared with simple cells in primary visual cortex," Proc. R. Soc. London Ser. B **265**, 359–366 (1998).
4. W. S. Geisler, J. S. Perry, B. J. Super, and D. P. Gallogly, "Edge co-occurrence in natural images predicts contour grouping performance," Vision Res. **41**, 711–724 (2000).
5. D. Marr, *Vision* (Freeman, San Francisco, Calif., 1982).
6. D. H. Hubel and T. N. Wiesel, "Receptive fields of single neurones in the cat's striate cortex," J. Physiol. (London) **148**, 574–591 (1959).
7. D. H. Hubel and T. N. Wiesel, "Receptive fields, binocular interaction and functional architecture in the cat's visual cortex," J. Physiol. (London) **160**, 106–154 (1962).
8. J. A. Movshon, I. D. Thompson, and D. J. Tolhurst, "Spatial and temporal contrast sensitivity of neurons in areas 17 and 18 of the cat's visual cortex," J. Physiol. (London) **283**, 101–120 (1978).
9. R. L. De Valois, E. W. Yund, and N. Hepler, "The orientation and direction selectivity of cells in macaque visual cortex," Vision Res. **22**, 531–544 (1982).
10. R. L. De Valois, D. G. Albrecht, and L. G. Thorell, "Spatial frequency selectivity of cells in macaque visual cortex," Vision Res. **22**, 545–559 (1982).
11. J. P. Jones, A. Stepnowski, and L. A. Palmer, "The two-dimensional spectral structure of simple cell receptive fields in cat striate cortex," J. Neurophysiol. **58**, 1212–1232 (1987).
12. J. P. Jones and L. A. Palmer, "The two-dimensional spatial structure of simple receptive fields in cat striate cortex," J. Neurophysiol. **58**, 1187–1211 (1987).
13. G. C. De Angelis, J. G. Robson, I. Ohzawa, and R. D. Freeman, "Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex," J. Neurophysiol. **68**, 144–163 (1992).
14. D. L. Ringach, "Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex," J. Neurophysiol. **88**, 455–463 (2002).
15. R. J. Baddeley and P. J. Hancock, "A statistical analysis of natural images matches psychophysically derived orientation tuning curves," Proc. R. Soc. London Ser. B **246**, 219–223 (1991).
16. G. M. Boynton, J. B. Demb, G. H. Glover, and D. J. Heeger, "Neuronal basis of contrast discrimination," Vision Res. **39**, 257–269 (1999).
17. D. J. Heeger, A. Huk, W. S. Geisler, and D. G. Albrecht, "Spikes versus BOLD: What does neuroimaging tell us about neuronal activity?" Nat. Neurosci. **3**, 631–633 (2000).
18. L. L. Kontsevich, C. C. Chen, and C. W. Tyler, "Separating the effects of response nonlinearity and internal noise psychophysically," Vision Res. **42**, 1771–1784 (2002).
19. D. J. Tolhurst, J. A. Movshon, and I. D. Thompson, "The dependence of response amplitude and variance of cat visual cortical neuron on stimulus contrast," Exp. Brain Res. **41**, 414–419 (1981).
20. D. J. Tolhurst, J. A. Movshon, and A. F. Dean, "The statistical reliability of signals in single neurons in cat and monkey visual cortex," Vision Res. **23**, 775–785 (1983).

21. A. Bradley, B. C. Skottun, I. Ohzawa, G. Sclar, and R. D. Freeman, "Visual orientation and spatial frequency discrimination: a comparison of single neurons and behavior," J. Neurophysiol. **57**, 755–771 (1987).
22. R. Vogels, W. Spileers, and G. A. Orban, "The response variability of striate cortical neurons in the behaving monkey," Exp. Brain Res. **77**, 432–436 (1989).
23. W. S. Geisler and D. G. Albrecht, "Visual cortex neurons in monkeys and cats: detection, discrimination and identification," Visual Neurosci. **14**, 897–919 (1997).
24. D. G. Albrecht and D. B. Hamilton, "Striate cortex of monkey and cat: contrast response function," J. Neurophysiol. **48**, 217–237 (1982).
25. G. Sclar, J. H. R. Maunsell, and P. Lennie, "Coding of image contrast in central visual pathways of the macaque monkey," Vision Res. **30**, 1–10 (1990).
26. F. Sengpiel, R. J. Baddeley, T. C. B. Freeman, R. Harrad, and C. Blakemore, "Different mechanisms underlie three inhibitory phenomena in cat area 17," Vision Res. **38**, 2067–2080 (1998).
27. D. J. Tolhurst, "The amount of information transmitted about contrast by neurons in the cat's visual cortex," Visual Neurosci. **2**, 409–413 (1989).
28. W. S. Geisler and D. G. Albrecht, "Bayesian analysis of identification in monkey visual cortex: nonlinear mechanisms and stimulus certainty," Vision Res. **35**, 2723–2730 (1995).
29. P. Mamassian, M. Landy, and L. T. Maloney, "Bayesian modeling of visual perception," in *Probabilistic Models of the Brain*, R. P. N. Rao, B. A. Olshausen, and M. S. Lewicki, eds. (MIT Press, Cambridge, Mass., 2002).
30. S. Laughlin, "A Simple Coding Procedure Enhances a Neuron's Information Capacity," Z. Naturforsch. Teil C **36**, 910–912 (1981).
31. D. J. Tolhurst, "The limited contrast-response of single neurons in cat striate cortex and the distribution of contrasts in natural scenes," J. Physiol. (London) **497**, 64P (1996).
32. N. Brady and D. J. Field, "Local contrast in natural images: normalisation and coding efficiency," Perception **29**, 1041–1055 (2000).
33. Y. Tadmor and D. J. Tolhurst, "Calculating the contrasts that retinal ganglion cells and LGN neurons encounter in natural scenes," Vision Res. **40**, 3145–3157 (2000).
34. A. Anzai, M. A. Bearse, R. D. Freeman, and D. Q. Cai, "Contrast coding by cells in the cat's striate cortex," Visual Neurosci. **12**, 77–93 (1995).
35. K. I. Naka and W. A. H. Rushton, "S-potentials from colour units in the retina of fish (Cyprinidae)," J. Physiol. (London) **185**, 536–555 (1966).
36. D. J. Heeger, "Half-squaring in responses of cat striate cortex," Visual Neurosci. **9**, 427–443 (1992).
37. D. J. Heeger, "Normalization of cell responses in cat striate cortex," Visual Neurosci. **9**, 181–197 (1992).
38. A. Gottschalk, "Derivation of the visual contrast response function by maximizing information rate," Neural Comput. **14**, 527–542 (2002).
39. D. L. Ringach, Departments of Neurobiology and Psychology, Franz Hall, Room 8441B, University of California, Los Angeles, Los Angeles, California 90095-1563 (personal communication, 2002).
40. D. L. Ringach, M. J. Hawken, and R. Shapley, "Dynamics of orientation tuning in macaque primary visual cortex," Nature **387**, 281–284 (1997).
41. D. L. Ringach, R. Shapley, and M. J. Hawken, "Orientation selectivity in macaque V1: diversity and laminar dependence," J. Neurosci. **22**, 5639–5651 (2002).
42. P. L. Clatworthy, M. Chirimuuta, J. S. Lauritzen, and D. J. Tolhurst, "Coding of the contrasts in natural images by populations of neurons in striate visual cortex (V1)," Vision Res. (to be published).
43. M. N. Shadlen, K. H. Britten, W. T. Newsome, and J. A. Movshon, "A computational analysis of the relationship between neuronal and behavioural responses to visual motion," J. Neurosci. **16**, 1486–1510 (1996).
44. J. S. Lauritzen, A. Pelah, and D. J. Tolhurst, "Perceptual rules for watermarking images: A psychophysical study of the visual basis for digital pattern encryption," in *Human Vision and Electronic Imaging IV*, B. E. Rogowitz, and T. N. Pappas, eds., Proc. SPIE **3644**, 392–402 (1999).
45. E. Peli, "Contrast in complex images," J. Opt. Soc. Am. A **7**, 2032–2040 (1990).
46. Y. Tadmor and D. J. Tolhurst, "Discrimination of changes in the second-order statistics of natural and synthetic images," Vision Res. **34**, 541–554 (1994).
47. S. Marcelja, "Mathematical description of the responses of simple cortical cells," J. Opt. Soc. Am. A **70**, 1297–1300 (1980).
48. D. J. Field and D. J. Tolhurst, "The structure and symmetry of simple-cell receptive field profiles in the cat's visual cortex," Proc. R. Soc. London Ser. B **228**, 379–400 (1986).
49. J. P. Jones and L. A. Palmer, "An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex," J. Neurophysiol. **58**, 1233–1258 (1987).
50. D. J. Tolhurst and I. D. Thompson, "On the variety of spatial frequency selectivities shown by neurons in area 17 of the cat," Proc. R. Soc. London **B213**, 183–199 (1981).
51. D. J. Tolhurst, Y. Tadmor, and Tang Chao, "The amplitude spectra of natural images," Ophthalmic Physiol. Opt. **12**, 229–232 (1992).
52. M. V. Srinivasan, S. Laughlin, and A. Dubs, "Predictive coding: A fresh view of inhibition in the retina," Proc. R. Soc. London Ser. B **216**, 427–459 (1982).
53. C. Blakemore and G. F. Cooper, "Development of the brain depends on visual environment," Nature **228**, 477–478 (1970).
54. C. Blakemore and R. C. van Sluyters, "Innate and environmental factors in the development of the kitten's visual cortex," J. Physiol. (London) **248**, 663–716 (1975).
55. I. Ohzawa, G. Sclar, and R. D. Freeman, "Contrast gain control in the cat's visual system," J. Neurophysiol. **54**, 651–665 (1985).
56. A. B. Bonds, "Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex," Visual Neurosci. **2**, 41–55 (1989).
57. W. S. Geisler and R. L. Diehl, "Bayesian natural selection and the evolution of perceptual systems," Philos. Trans. R. Soc. London, Ser. B **357**, 419–448 (2002).