Contrast–response functions for multifocal visual evoked potentials: A test of a model relating V1 activity to multifocal visual evoked potentials activity

Donald C. Hood Department of Psychology, Columbia University, New York, NY, USA

Quraish Ghadiali Department of Psychology, Columbia University, New York, NY, USA

Jeanie C. Zhang Department of Psychology, Columbia University, New York, NY, USA

Norma V. Graham Department of Psychology, Columbia University, New York, NY, USA

S. Sabina Wolfson Department of Psychology, Columbia University, New York, NY, USA

Xian Zhang Department of Psychology, Columbia University, New York, NY, USA

The multifocal visual evoked potential (mfVEP) is largely generated in V1. To relate the electrical activity recorded from humans to recordings from single cells in nonhuman primate (V1) cortex, contrast–response functions for the human mfVEP were compared to predictions from a model of V1 activity (D. J. Heeger, A. C. Huk, W. S. Geisler, & D. G. Albrecht, 2000) based upon single-cell recordings from monkey V1 (e.g., D. G. Albrecht, 1995; D. G. Albrecht, W. S. Geisler, R. A. Frazor, & A. M. Crane, 2002; D. G. Albrecht & D. B. Hamilton, 1982; W. S. Geisler & D. G. Albrecht, 1997). A second purpose was to fully articulate the assumptions of this model to better understand the implications of this comparison. Finally, as the third purpose, one of these assumptions was tested. Monocular mfVEPs were obtained from normal subjects with a contrast-reversing dartboard pattern. The display contained 16 sectors each with a checkerboard. Both the sectors and the checks were scaled approximately for cortical magnification. In Experiment 1, there were 64 checks per sector. The contrast–response functions were fitted well up to 40% contrast by the theoretical population curve for V1 neurons; there was a systematic deviation for higher contrasts. The model, as articulated here, predicts that the contrast–response function should be the same and independent of the size of the elements in the display. Varying the size of the elements by varying the viewing distance in Experiment 2 produced similar results to those in Experiment 1. In Experiment 3, the viewing distance and sector size were held constant, but the size of the elements (and therefore the number of checks per sector) was varied. Changing check size by a factor of 16 had relatively little effect on the contrast–response function. In general, the mfVEP results were consistent with the model based upon the V1 neuron population. However, two aspects of the results require further exploration. First, there was a systematic deviation from the model's contrast–response function for higher contrasts. This deviation suggests that one or more of the model's assumptions may be violated. Second, the latency of the mfVEP changed far less than expected based upon single-cell data.

Keywords: V1, multifocal visual evoked potential, mfVEP, fMRI, cortical model, contrast, spatial frequency

Introduction

Over the last 40 years or so, a wealth of information has been accumulating about the cellular function of the primary (V1) visual cortex in nonhuman primates. For example, recordings have been made from thousands of V1 cells. It would be of considerable interest if the electrical activity recorded from the human scalp could be related systematically to these recordings. The obvious candidate for such comparisons is the visual evoked potential (VEP), a gross electrical potential generated by the cells in the occipital cortex.

Although the VEP is easily recorded with scalp electrodes, it is not a good choice for this comparison for two reasons.

doi: 10.1167/6.5.4 Received September 29, 2005; published April 24, 2006 ISSN 1534-7362

First, it is typically recorded in response to a relatively large stimulus. The response is therefore generated in wide regions of cortex and the waveform of the response complicated by the twists and turns of cortical folding. Second, the VEP has multiple sources, including those in regions beyond V1 (for a recent review, see Di Russo et al., 2005). These problems make the conventional VEP a poor choice for a comparison to single-cell activity.

A relatively new VEP method, the multifocal visual evoked potential (mfVEP) technique (Baseler, Sutter, Klein, & Carney, 1994; for a review, see Hood & Greenstein, 2003) circumvents these problems. With the mfVEP technique, many (typically 16–60) spatially local VEP responses can be recorded simultaneously, providing spatially localized measures of cortical activity. Further, three lines of evidence argue that the mfVEP is generated largely in V1. First, as originally pointed out by Baseler et al. (1994), the mfVEP waveforms reverse polarity as one crosses the horizontal meridian (see also Hood & Greenstein, 2003; Klistorner, Graham, Grigg, & Billson, 1998). Only potentials generated from inside the calcarine fissure should behave this way. The mfVEP from the upper visual field is reversed in polarity as compared with the lower one, whereas the conventional VEP recorded with the same electrode positions and same subjects can show the same polarity for upper and lower field stimulation (Fortune & Hood, 2003). Fortune and Hood (2003) showed that this difference was due to the fast sequence of mfVEP stimulation, which they speculated produces a response with a smaller extrastriate contribution than in the case of the conventional VEP. Second, dipole analysis suggests that most of the mfVEP signal is generated in V1 (Slotnick, Klein, Carney, Sutter, & Dastmalchi, 1999). Third, using an application of principal component analysis, Zhang and Hood (2004) provided evidence that the first principal component of the mfVEP was generated within the calcarine fissure and thus within V1.

One purpose of the current study was to record contrast–response functions for the mfVEP and to compare them to predictions based upon V1 single-cell activity. A predicted contrast–response function for V1 can be found in Heeger, Huk, Geisler, and Albrecht (2000), who fitted it to human fMRI BOLD contrast–response data (Boynton, Demb, Glover, & Heeger, 1999). This predicted function was based upon the analysis of contrast–response functions recorded from 333 V1 neurons in monkeys by Albrecht (1995), Albrecht, Geisler, Frazor, and Crane (2002), Albrecht and Hamilton (1982), and Geisler and Albrecht (1997). Although the predicted V1 function from Heeger et al. provides us with a powerful tool for relating V1 activity to mfVEP response functions, the assumptions underlying this predicted curve need to be specified. In their brief report in *Nature Neuroscience,* these assumptions were not delineated. A second purpose of this study was to fully articulate these assumptions. Finally, as the third purpose here, we test one of the key assumptions.

Methods: Experiments 1 and 2

Stimuli

The visual stimulus was designed and displayed using VERIS Software (EDI, San Mateo, CA). The display consisted of 16 sectors (Figure 1A), 6 in the outer ring, 6 in the middle ring, and 4 in the inner ring. It was viewed with the natural pupil. Each sector consisted of 64 checks, 32 at a higher luminance and 32 at a lower luminance. As Figure 1A indicates, the size of the checks increased as their distance from fixation increased. The scaling of the mfVEP display by Baseler et al. (1994) was designed to approximately correspond to the change in cortical magnification. During testing, each sector changed according to a pseudorandom reversal sequence in which the checks

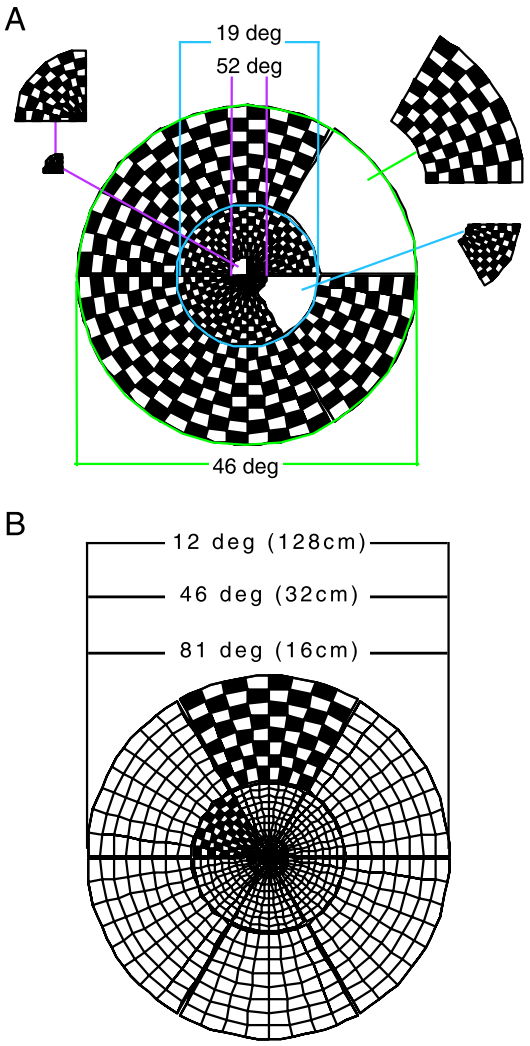

Figure 1. (A) The mfVEP display. For Experiment 1, the display was viewed at 32 cm. (B) Three viewing distances (16, 32, and 128 cm) were used in Experiment 2.

within the sector had a 50–50 chance of all remaining the same in luminance or switching to the other luminance every 13.3 ms (75 Hz frame rate). Technically, the pattern reversal stimuli were driven by a binary m-sequence. It should be noted that other multifocal modes of stimulation have been employed (James, 2003; Klistorner, Crewther, & Crewther, 1997; Maddess, James, & Bowman, 2005). For more details about the mfVEP, our procedures, or both, see Hood & Greenstein (2003).

Experiment 1

The visual stimulus was placed at an effective distance of 32 cm in front of the subject and was 46 deg in diameter (Figure 1A). Five contrast levels, 6.25%, 12.5%, 25%, 50%, and 99%, were presented first in order of increasing contrast, then in order of decreasing contrast, so that two recordings were obtained for each contrast level in a single session. Each recording was obtained in eight segments, each 30 s in length, with short breaks between segments. The contrast was the same for all eight segments. For all subjects, this session was repeated on three separate days.

Experiment 2

The display, the same as in Experiment 1, was viewed at three different distances, 16, 32, and 128 cm (see Figure 1B). The 32-cm distance was included to replicate Experiment 1. Six contrast levels were used: 8.7%, 16.8%, 35.6%, 70.3%, 89.9%, and 99.6%. As in Experiment 1, within a session, the contrast levels were presented first in order of increasing contrast, then in order of decreasing contrast, so that each contrast level was displayed twice. As in Experiment 1, each recording consisted of eight 30-s segments and each session was repeated on three separate days.

Subjects and recording

Three individuals, between 20 and 21 years, participated. All three had normal vision with 20/20 corrected acuity. All three participated in Experiment 1, whereas two of the subjects (S1 and S2) participated in Experiment 2. Procedures followed the tenets of the Declaration of Helsinki, and the protocol was approved by the Committee of the Institutional Board of Research of Columbia University.

The mfVEPs were recorded with gold cup electrodes placed at the inion, 4 cm above the inion, and 1 cm above and 4 cm to both the right and left of the inion (see Figure 3B in Hood & Greenstein, 2003). The forehead was used as the common ground. Differential recording was obtained with three channels using the inion as the reference, and three additional "channels" were derived offline. The continuous VEP record was amplified with the high- and low-frequency cutoffs set at 3 and 100 Hz (Grass PreAmplifier P511J, Quincy, MA).

A special purpose software, written in MATLAB (Mathworks Inc., MA), took as its input the 16 mfVEP records (the first slice of the second-order kernels for the 16 sectors of the display) from all three channels as derived by the VERIS software. The software filtered the records (low-pass cutoff at 35 Hz), displayed the results in the form of trace arrays (Figure 2), and performed quantitative analyses. Further, by taking the difference between pairs of channels, three additional "derived" channels were obtained, resulting in effectively six channels of recording. Except where noted, analyses were performed on the "best" responses, those with the largest signal-to-noise ratio (SNR), from the six "channels" as previously described (Hood & Greenstein, 2003; Hood, Greenstein, et al., 2002; Hood, Zhang, Hong, & Chen, 2002).

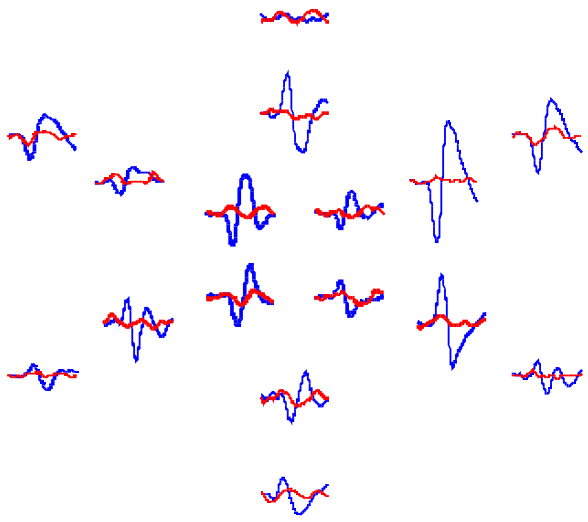

Figure 2. mfVEP responses from S1 for 6.25% (red) and 99% (blue) contrast.

Measurements of amplitude and timing

A typical measure of VEP amplitude is the root mean square (RMS) of the records taken over a particular time window. The amplitude plotted for Experiments 1 and 2 is the RMS calculated for the interval from 45 to 150 ms. Because the waveforms for the different contrast conditions were very similar (see Figure 6), varying this interval had little effect on the outcome.

The timing of the responses was measured for the two subjects (S1 and S2) and the condition, Experiments 1 and 2 (32 cm), for which we had the most data. To measure the relative latency, the responses from each of the 16 locations of each of the subjects were analyzed in the following way. First, the responses from the midline channel (i.e., between the inion and inion plus 4 cm) for the individual runs were averaged for each sector. The best channel was not used in this analysis because the waveform tends to vary with channel and this could confound our latency measure. Second, for each subject at each contrast, the 16 responses (one for each of the 16 sectors) were squared and summed, and the average summed curve was obtained (squaring the responses allows the responses from upper and lower fields,

which tend to be reversed in polarity, to be summed; Meigen & Kraemer, 2005). Records with little or no response were excluded from the analysis. In particular, based upon our experience with measurement of latency, only responses with SNR values greater than 1.7 were included in the sum (Hood et al., 2004). Third, the first peak (around 70 ms) was easily identified and its latency was measured.

Results: Experiments 1 and 2

Experiment 1

Figure 2 shows the individual mfVEP "best" responses recorded in one session from one individual. The responses to a contrast of 99% (blue) are clearly larger than those to a contrast of 6.25% (red). Figure 3A shows the mfVEP (RMS) amplitude versus contrast (contrast–response functions) for the three sessions of the same subject. For this analysis, the RMS amplitudes of the 16 mfVEP responses (one per sector) were averaged. In Figure 3B, the same data are normalized by dividing by the value at 25%. Although Figure 3A indicates that there is variability in amplitude from session to session, Figure 3B suggests that the shape of the curve remains approximately the same.

Figure 4A shows the average response curve for each subject, and Figure 4B shows the same data normalized to the value at 25% contrast. Although the overall amplitude varies among the subjects (Figure 4A), there is excellent agreement in the shape of the curves up to, and including, the 50% contrast point (Figure 4B). On average, the curves appear to reach a maximum by about 50% contrast.

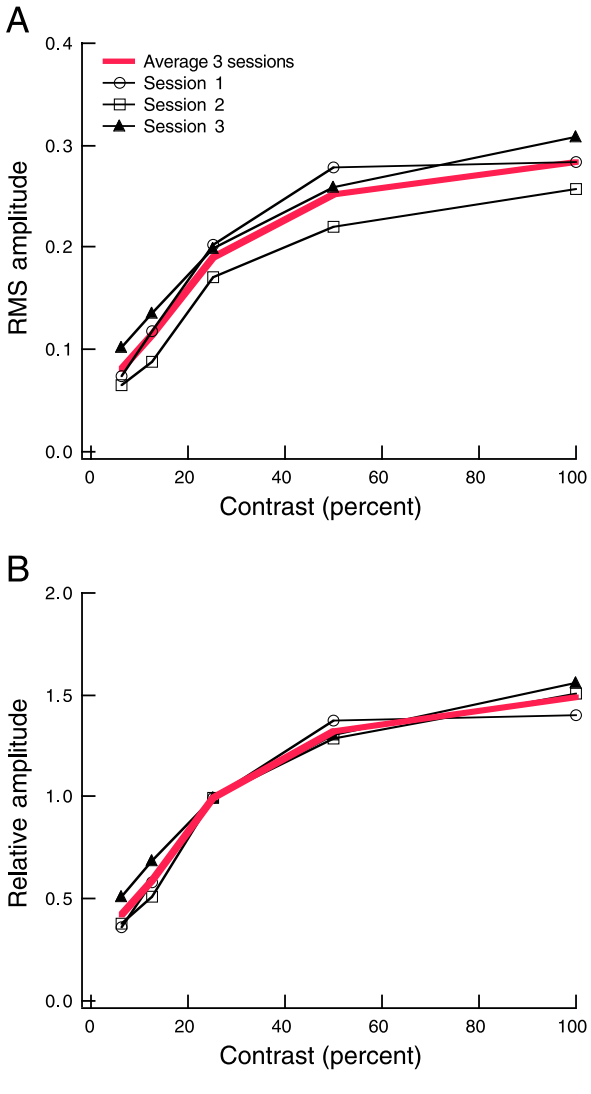

Figure 3. (A) Average RMS (in microvolts) amplitude (averaged across sectors) as a function of contrast for one subject (S1) and three separate sessions (symbols). The average of the three sessions is shown in red. (B) The data from panel A normalized to the value for 25% contrast.

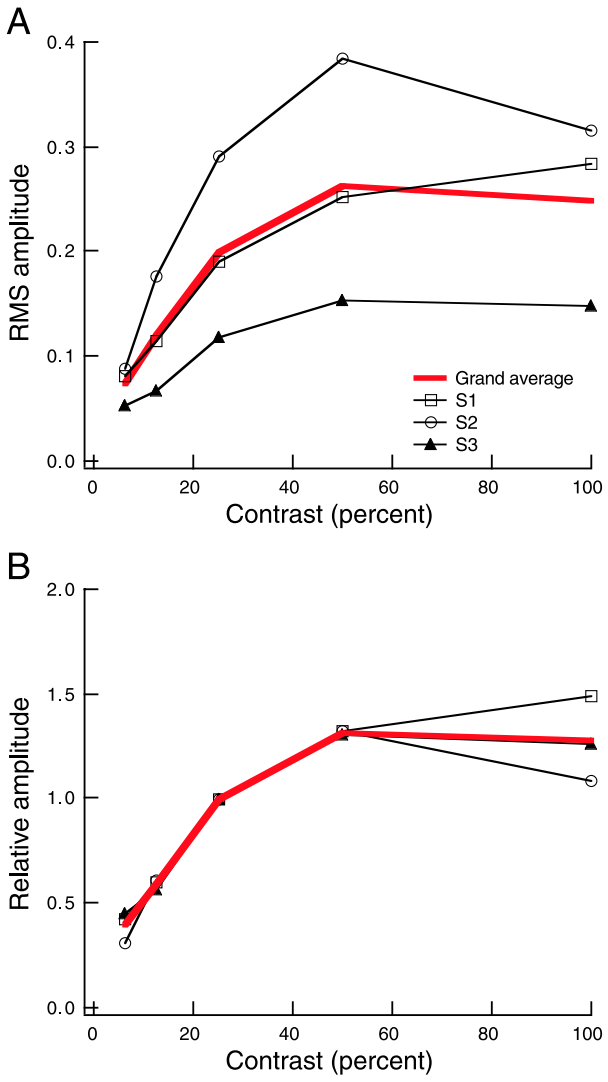

Figure 4. (A) Average RMS (in microvolts) amplitude (averaged across sectors) as a function of contrast for all three subjects. The average of the three subjects is shown in red. (B) The data from panel A normalized to the value for 25% contrast.

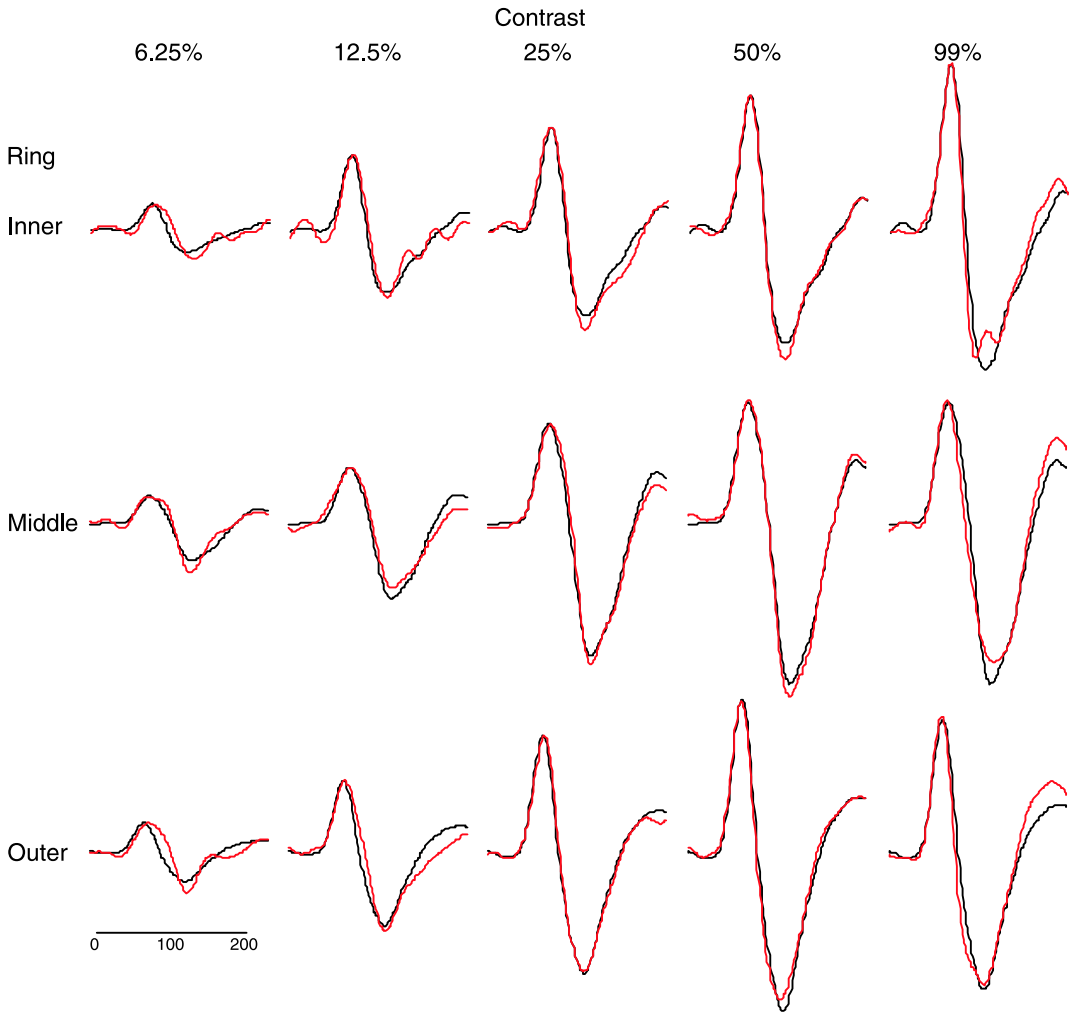

Figure 5. Sample mfVEP responses (red) from one subject (S1) and three locations, one from each of the three rings in the display. The black curve in each row is a template formed by averaging all the responses in that row, which were then scaled to fit each response.

If the waveform of the responses differed substantially as a function of contrast, then the window of analysis used to calculate the RMS amplitude could affect the shape of the contrast–response function. Figure 5 shows all the responses (red) from three locations for one of the subjects (S1). The three sectors, one from each ring, were chosen based upon the size (SNR) of the response. The black response in each row is the average of the responses (red) in that row for all contrasts scaled in amplitude by a multiplicative constant. This scaled template provides a reasonable fit to all responses. To a first approximation, the waveform for a given sector is the same at all contrasts.

Although, to a first approximation, the waveforms are similar for different contrasts, there are subtle differences. Notice in Figure 5 that there is a tendency for the responses at the lower contrasts to be slightly slower, and perhaps slightly broader, than the responses at high contrasts. As described in the Methods section, we measured the latency. The filled symbols in Figure 6 show the relative latency for S1 and S2 as a function of contrast. Latency, expressed

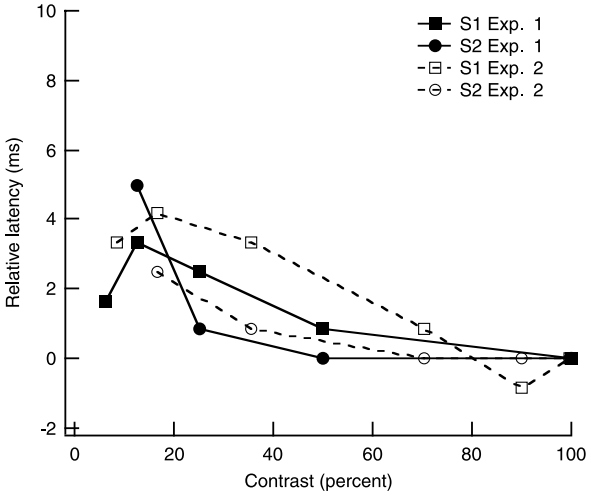

Figure 6. The relative latency of the mfVEP from Experiments 1 and 2 (32 cm) for two subjects.

relative to the value at full contrast, decreased as contrast increased, but only by about 4 ms. (Note that there is no point for S2 at the lowest contrast because only 2 of the 16 locations meet the SNR criterion for our latency measure.) The open symbols show the results for a comparable condition in Experiment 2.

Experiment 2

Figure 7A shows the average RMS amplitude of the two subjects for the three distances. As in Figures 3 and 4, the RMS amplitudes of the 16 mfVEP responses were averaged. The amplitudes were smaller for the closest/largest (open triangles: 16 cm distance) and furthest/smallest (open squares: 128 cm distance) displays than they were for the intermediate distance/size (open circles: 32 cm). The red circles are the average RMS amplitudes from Experiment 1 for the same two subjects. There is reasonably good agreement between the results from Experiment 1 at 32 cm (red circles) and the data for 32 cm (open circles) from Experiment 2. Figure 7B shows the data from Figure 7A normalized at the amplitude for 36% contrast. (Note that there were no data points for 36% in Experiment 1 so that these data were normalized to fall between the 17% and 36% points for the other conditions.) To a first approximation, the data for all conditions and both experiments fall along a common curve. The solid red curve is the predicted V1 function to be discussed in the next section along with panel C of Figure 7.

A model of the contrast–response function of V1 neurons

The model

The red curve in Figure 7B, supplied by Bill Geisler (personal communication), is based upon the contrast–response data from 333 V1 neurons of the monkey (macaque) recorded by Albrecht (1995), Albrecht et al. (2002), Albrecht and Hamilton (1982), and Geisler and Albrecht (1997). This theoretical curve provides a prediction of the average firing rate of V1 neurons (see caption to Figure 2 in Heeger et al., 2000). This curve is the same as the one published in Figure 2 of Heeger et al. (2000), which was reproduced here as Figure 8. Heeger et al. fitted this curve to the contrast–response data (filled symbols in Figure 8) from a human fMRI study (Boynton et al., 1999). To fit these data, they vertically scaled the V1 curve by a multiplicative constant (Heeger et al., 2000). As the stimulus conditions in the fMRI study and our mfVEP study differ from those in the single-unit studies, the comparison based upon only a vertical scaling needs justification. More generally, underlying comparisons such as those in Figures 7B and 8, there must be a model or a set of assumptions. What

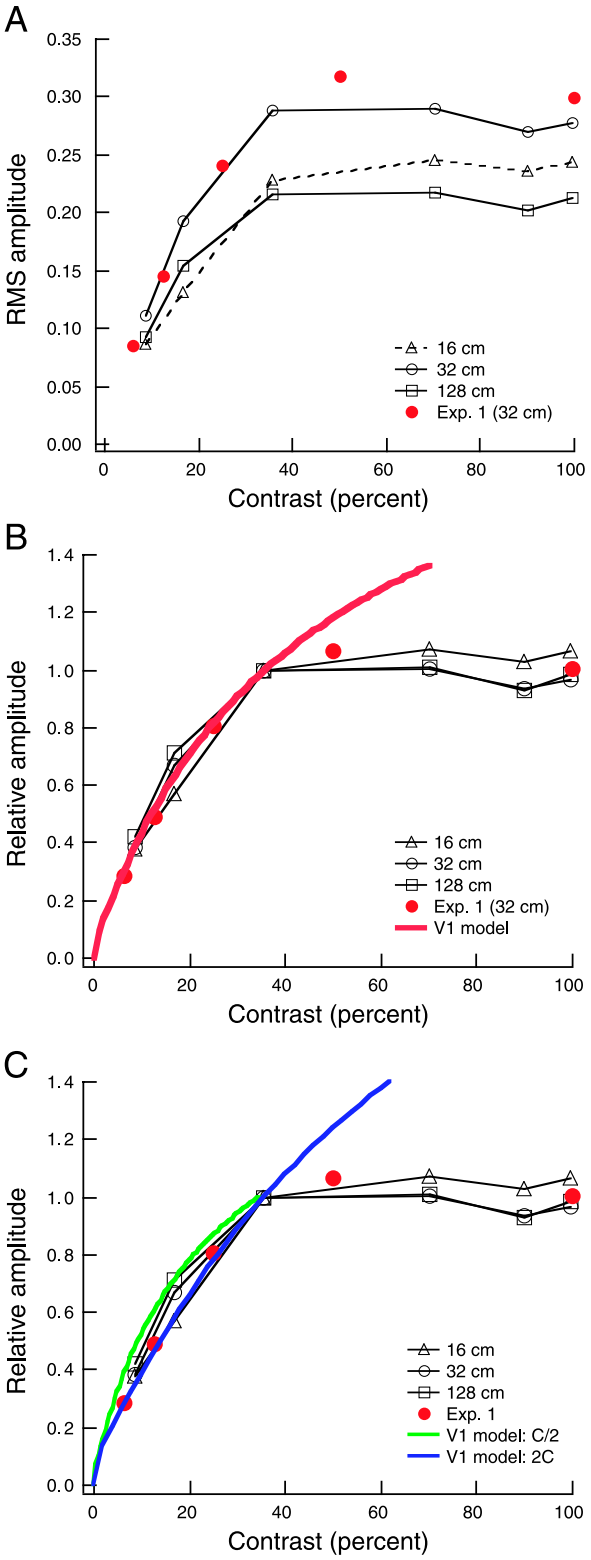

Figure 7. (A) Average RMS amplitude (in millivolts; averaged across sectors) as a function of contrast averaged for S1 and S2 and shown for Experiment 1 (red) and the three conditions of Experiment 2. (B) The data from panel A normalized as described in the text. The red curve is the prediction of a model described in the text. (C) The same plot as in panel B with the predictions for the model with C_{50} increased (blue) or decreased (green) by a factor of 2.

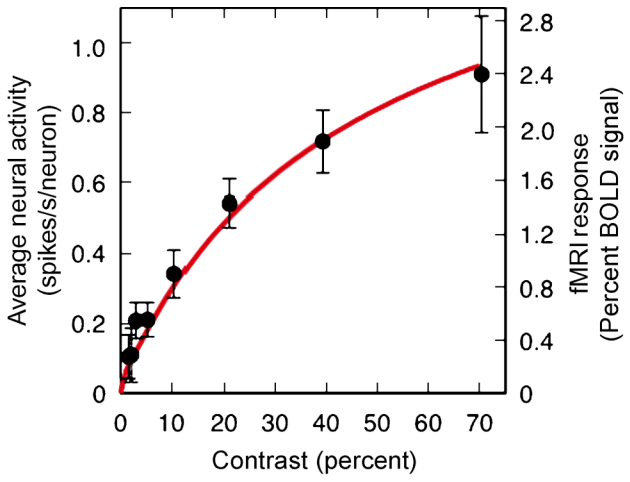

Figure 8. A comparison of fMRI data (symbols) to a model of V1 activity (red curve; from Heeger et al., 2000).

follows is our attempt to make the assumptions relating V1 physiology to mfVEP responses explicit. Something close to this model is implicit in the analysis of Heeger et al.

There are five assumptions:

1. The same equation,

$$R = [C^n/(C^n + C_{50}{}^n)]R_{\max}, \quad (1)$$

provides a reasonable fit to the contrast–response function of all cells (e.g., Albrecht & Hamilton, 1982; Geisler & Albrecht, 1997),

where R is the response of the cell and C is the contrast of the stimulus. C_{50}, n, and $R_{\max}$ are constants where n is the exponent and determines the slope of the initial portion of the curve, C_{50} is the contrast producing one-half maximum response, and $R_{\max}$ is the maximum response. Assumption 1 allows for the characterization of the individual contrast–response functions for the 333 cells so that they can be readily averaged. Support for this assumption can be found in Albrecht and Hamilton (1982), although in that study and in the work of others (e.g., Sclar, Maunsell, & Lennie, 1990), a range of shapes of the contrast–response function can be found.

2. The values of C_{50} and n are constant for each cell. Stimulating a cell with a nonpreferred stimulus, to a first approximation, only changes $R_{\max}$.

Assumption 2 asserts that varying the stimulus will affect only parameter $R_{\max}$ (for an example, see Figure 7 in Albrecht & Hamilton, 1982). In general, the evidence suggests that saturation is determined by the contrast of the stimulus and not by the amplitude of the responses (for reviews, see Albrecht, 1995; Albrecht et al., 2002; Carandini, Heeger, & Movshon, 1997; Geisler & Albrecht, 1997).

3. For a range of preferred spatial (and temporal) frequencies, the parameters n, $R_{\max}$, and C_{50} are approximately independent of preferred spatial (and temporal) frequency.

Changing the spatial or temporal characteristic of the stimulus will change the population of V1 cells dominating the response. According to Assumption 3, a change in the population of cells dominating the response will not change the predicted function. For example, if the cells preferring high spatial frequencies had, on average, larger values of C_{50}, then a change in the stimulus would yield a different theoretical curve. There is evidence that C_{50} is approximately the same for different preferred spatial frequencies (see Figure 6 in Albrecht & Hamilton, 1982), although as discussed below, these data show a small increase in C_{50} with increases in preferred frequency.

4. The contrast–response function is set shortly after the onset of a stimulus. Further, prolonged stimulation will not change this function.

This assumption is needed because the stimuli in the monkey experiments are brief, whereas those in the mfVEP are prolonged. Albrecht et al. (2002) provide evidence that the function is set within 10 ms of the onset of a stimulus. However, the possible effects of prolonged stimulation, such as those used in the mfVEP experiments, are less clear and will be discussed below.

5. The mfVEP is proportional to the sum of the responses of the V1 neurons.

An assumption is needed to link the mfVEP response to the theoretical V1 function. Heeger et al. (2000) used the agreement in Figure 8 between the BOLD signal of the fMRI and the V1 function to argue that the BOLD signal is a linear function of the neural activity.

The fit to the mfVEP data

The red curve in Figure 7B shows the fit of the V1 theoretical function. This function was fitted to the mfVEP data by normalizing (vertically scaling the curve by a multiplicative constant) to the mfVEP data from Experiment 1 at 25% contrast. The fit below 40% contrast is good. However, the data fall below the theoretical curve between 40% and 70% contrast. The theoretical curve does not extend beyond 70% contrast, as the monkey cells, in general, were not tested beyond this value.

The question naturally arises as to how sensitive the data are in distinguishing this curve from one with a different value of C_{50}. Figure 7C shows the effect of making C_{50} twice as large (blue curve) or twice as small (green curve). Roughly speaking, the data cannot distinguish between C_{50} values within a factor of ±2.

A test of Assumption 3

The visual system is less sensitive to high or very low spatial frequencies than it is to intermediate spatial frequencies. According to Assumption 3, these differences are due

at the cellular level to only a change in R_{max}; the value of C_{50} is independent of the preferred spatial frequency. However, there is evidence that the value of C_{50} may increase with increases in preferred spatial frequency. First, there is a suggestion in the monkey data that the value of C_{50} increases with increased spatial frequency (see Figure 6 in Albrecht & Hamilton, 1982). The value of C_{50} increased from about 19% contrast for cells with preferred frequencies below 1 cycle/deg to about 26% for cells with preferred frequencies above 10.7 cycles/deg. In addition, the contrast–response functions for the human BOLD fMRI signals showed a considerably higher value of the equivalent to C_{50} with a 2 cycle/deg plaid as compared with the 0.5 cycle/deg stimulus (Boynton et al., 1999), although the monkey data show relatively little change over this range.

The size (and spatial frequency content) of the stimuli was varied in Experiment 2. However, whereas these data supply one test of Assumption 3, the size of the individual sector and the total region of the retina stimulated were also varied. In Experiment 3, the viewing distance (retinal image) and sector size were held constant, but the size of the elements (and therefore, the number of checks per sector) was varied.

Experiment 3

Methods

Three subjects participated in this experiment, S1 and S2 from Experiments 1 and 2 and a third subject, S4, who was 19 years old. The display, as in Experiments 1 and 2, had 16 sectors and was viewed at 64 cm. In Experiments 1 and 2, each sector had 64 checks and the display was viewed at 32 cm (Experiment 1) or at one of three distances (16, 32, or 64 cm) in Experiment 2. In Experiment 3, the check size was varied by changing the number of checks in each sector. In particular, the checkerboard in each sector was either 2×2 (4 checks), 4×4 (16 checks), or 8×8 (64 checks), as shown in Figure 9. Six contrast levels were used: 8%, 16%, 32%, 50%, 70%, and 100%. For each of the three displays, four sessions per subject were run. Within each session, each of the six contrast conditions was run twice in a random order. With the exception of the correction applied to the RMS amplitudes, all other conditions were as described above in the Methods section for Experiments 1 and 2. The RMS amplitude of the record contains both noise and signal. As long as the noise is relatively small, the RMS amplitude is a reasonable measure of signal amplitude. However, in Experiment 3, the responses to the 2×2 stimulus were relatively small, especially at low contrasts. To get an estimate of the RMS amplitude of the signal, we subtracted the RMS of a (noise) window between 325 and 430 ms from the RMS of the signal window. We have previously shown that the noise window contains little or no signal (Hood, Zhang, et al., 2002; Zhang, Hood, Chen, & Hong, 2002). This "corrected RMS amplitude" provides a measure of the signal as long as the noise is random and there is no correlation between the noise in the signal window and the noise in the noise window. (A similar correction had relatively little effect on the data from Experiments 1 and 2).

Results

The data points in Figure 10A show the average RMS amplitude for S1 for the four runs; the smooth theoretical curves will be described below. In general, the 8×8 check display produces the largest response, whereas the 2×2 check display the smallest. Figure 10B shows the same data after the correction for noise as described in the Methods section. To obtain a qualitative comparison of the shape of these curves, they are shown normalized to their maximum value in Figure 11A. The normalized corrected amplitudes for the other two subjects are shown in the other panels of Figure 11. We were particularly interested here in whether

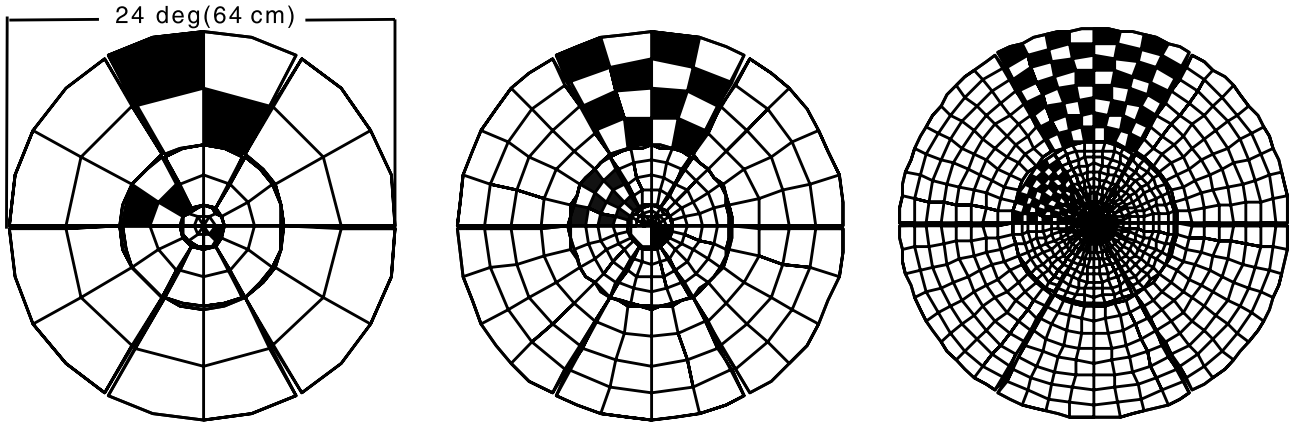

Figure 9. Schematic of stimuli used in Experiment 3 where the display was viewed at a distance of 64 cm but the number of checks per sector was set at 4 (left panel), 16 (middle panel), or 64 (right panel).

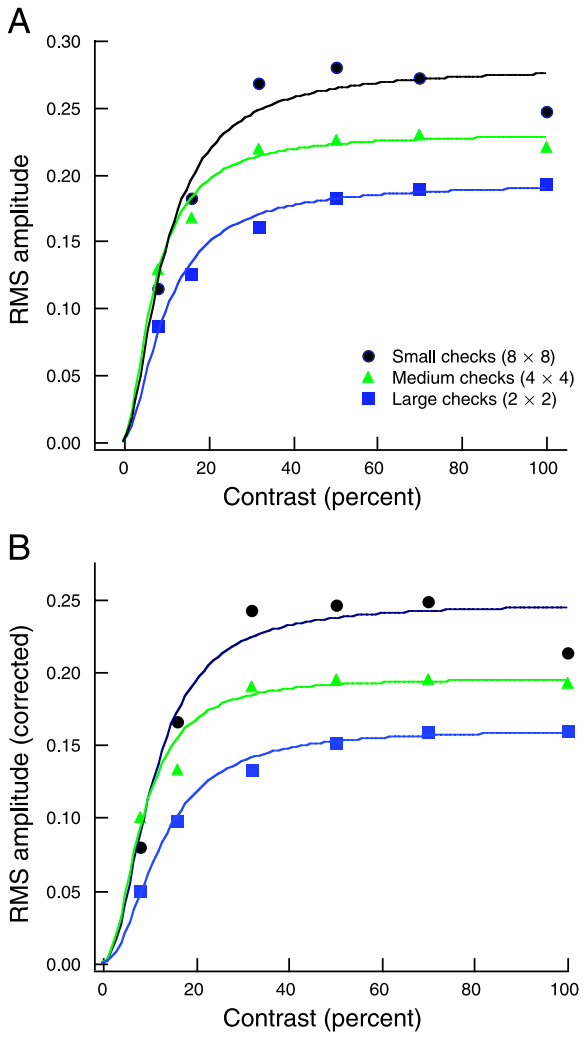

Figure 10. (A) Average RMS (in microvolts) amplitude (averaged across sectors) as a function of contrast for S1 and the three conditions (different symbols) of Experiment 3. The smooth curves show the fit of Equation 1 with *n* held constant. (B) The data from panel A corrected for noise as described in the text.

the curves appeared to be displaced relative to each other along the *X*-axis. The curves were normalized at their maximum values because, according to the model, only R_{max} should change; there should be no change in C_{50} and no displacement along the *X*-axis. Although the largest checks (blue squares) appear to fall to the right for two of the Ss, this trend is not found for the third. In general, there is no consistent trend for all three Ss.

To obtain a quantitative measure of the lateral placement of these data, Equation 1 was fitted to the data. Because we were interested in possible differences in C_{50}, n was held constant for each subject and only C_{50} and R_{max} were allowed to vary. The solid curves in both panels of Figure 10 show the fit. Table 1 contains the C_{50} values for the corrected data, the data shown in Figures 10B and 11. The fits were good with a goodness-of-fit statistic between 0.93 and 1.00. (The goodness-of-fit statistic was calculated as

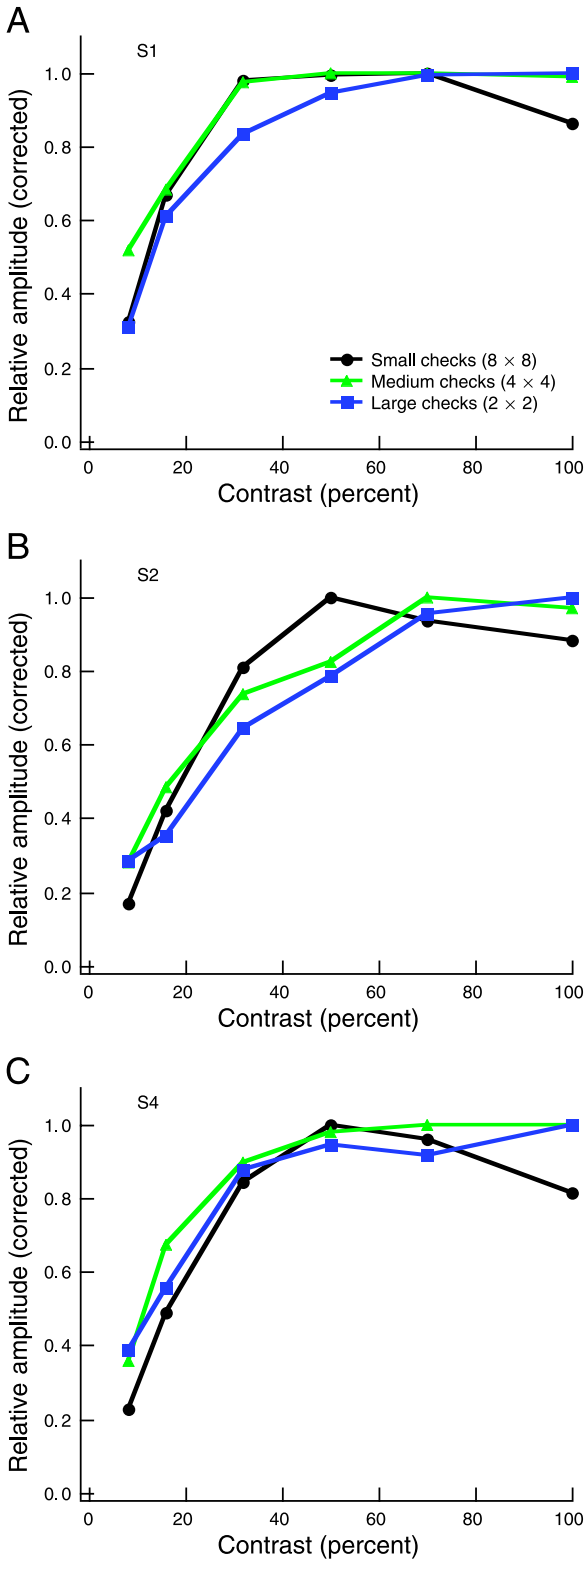

Figure 11. (A) Average RMS amplitude, corrected for noise and normalized to the maximum response, as a function of contrast for S1 (panel A), S2 (panel B), and S4 (panel C). The different symbols denote the three conditions of Experiment 3.

	2 × 2	4 × 4	8 × 8
Subject 1	12.6	8.7	10.9
Subject 2	21.6	15.8	17.2
Subject 3	11.2	10.9	14.2
Average	15.1	11.8	14.1

Table 1. C_{50} values for RMS corrected amplitude (n held constant).

$1 - \{[\Sigma(x_i - m_i)^2] / [\Sigma(x_i - \mu)^2]\}$, where x_i is the ith data point, m_i is the predicted value at that point, and μ is the mean of the data points.) In agreement with the qualitative appearance of Figure 11, there are only small, statistically nonsignificant differences in the computed C_{50} values for the three conditions, whether the corrected (Table 1) or uncorrected (Table 2) amplitudes are used.

Discussion

This study had three purposes: first, to compare the contrast–response functions for human mfVEPs to a model of V1 activity (Heeger et al., 2000); second, to articulate the assumptions of this model to better understand the implications of this comparison; and third, to test one of the key assumptions of the model.

The model does a reasonable job of describing the data up to a contrast of about 40%. However, above 40% the response of the model continues to increase, although, in general, the mfVEP data reach a maximum and remain relatively constant. The assumptions delineated above provide a basis for discussing this discrepancy between model and data.

Spatial frequency of the stimulus

According to the model, only the parameter R_{max} should be affected by manipulations of the spatial characteristics of the stimulus. Because we had reasons to believe that the parameter C_{50} varies with the size of the checks, this variable was manipulated in Experiments 2 and 3. For the range of check sizes used here, the variation in C_{50} was modest.

As mentioned above, the single-cell data from the monkey show only a small, nonsignificant increase (19–26% contrast or less than a factor of 1.4) in C_{50}, as the preferred frequency of the cell increased from less than 1 cycle/deg to greater than 10.7 cycles/deg. On the other hand, in their human fMRI study, Boynton et al. (1999) concluded that C_{50} differed for their two conditions, 0.5 and 2.0 cycles/deg (plaid/checkerboard composed of orthogonal sine wave gratings). In particular, C_{50} was larger for 2.0 cycles/deg by a factor of 3.2 and 1.7 for their two subjects when the fMRI and behavioral data were simultaneously fitted (Figures 3 and 4 in Boynton et al., 1999) and by an even larger factor when the fMRI data were fitted alone.

Due to the scaling of the mfVEP checkerboard pattern (see Figure 1), it is hard to make a quantitative comparison between our results and either the monkey V1 recordings or the human fMRI data. However, we can say that in Experiment 3, as in the Boynton et al. (1999) fMRI study, the range of check sizes varied by a factor of 4. Further, the stimulus in the Boynton et al. study was an annulus with an inner radius of 5 deg and an outer radius of 7 deg. At about 6 deg from the center of our display, the check size ranged from 0.3 cycles/deg (2 × 2) to 1.3 cycles/deg (8 × 8) in the general range of their study. Of course, the check sizes are larger in the outer part of the mfVEP display and considerably smaller in the center of the display, making a direct comparison to their stimulus impossible. In any case, what we can say at this point is that if the spatial frequency content of the stimulus affects the contrast–response function for the mfVEP data, the effects are relatively subtle for our conditions. To further test the effects of spatial frequency, as well as for ease of comparison to the single-cell data, we need to present sine wave stimuli in a multifocal paradigm. This will require specially designed multifocal software, which we are developing.

	2 × 2	4 × 4	8 × 8
Subject 1	10.5	7.4	9.8
Subject 2	6.1	14.4	12.7
Subject 3	6.9	8.7	11.3
Average	7.8	10.2	11.3

Table 2. C_{50} values for RMS amplitude (n held constant).

Other factors and possible violation of the model's assumptions

The relative modest effect of check size suggests that other factors must be involved in the discrepancy between the model and mfVEP results at higher contrast levels. Here, we consider five of the more obvious factors that may lead to a violation of one or more of the assumptions of the model. First, in the monkey experiments upon which the model's prediction was based, the stimulus was approximately 2 s in duration with 15-s breaks between presentations. For multifocal stimuli, reversals occurred on average at a rate of 37.5 Hz (one half of the frame rate of 75 Hz) for a period of 30 s followed by a brief break before the next stimulus. It is well known that V1 neurons decrease in responsiveness with prolonged stimulation (e.g., Albrecht, Farrar, & Hamilton, 1984; Movshon & Lennie, 1979; Sclar, Lennie, & DePriest, 1989). Although Albrecht et al. (2002) showed that the characteristic of the response function were set within 10 ms of the stimulus onset, they were also careful to distinguish between this fast-contrast gain control and a slow-contrast gain control. Unlike the fast type of gain control, the slow type does depend upon whether the stimulus is

optimal or not. That is, Assumptions 2 and 4 are violated. Thus, mfVEPs to repetitive stimulation, especially at the higher contrasts, may be attenuated by slow gain changes. Relevant here is the work of Maddess et al. (2005), who measured contrast–response functions for multifocal flashes with different temporal densities. These functions were well fitted by a power function with an exponent that varied with flash density. In fact, the response to stimuli with the lower flash density (1.3/s) continued to increase for contrasts above 40%. This finding suggests that, perhaps, Assumption 4 is violated and that the high-contrast, rapid mfVEP stimulus employed here is inducing a slow gain change. A recent fMRI study (Gardner et al., 2005) suggests that slow gain changes can affect the BOLD signal as well.

Second, the recordings in the monkey experiments were made largely within 5 deg of the fovea (Albrecht & Geisler, personal communication), whereas our display was considerably larger. Our mfVEP data within the central 5 deg were too variable to allow us to test the effects of eccentricity. Although there is no evidence that the contrast–response function of cortical neurons vary with eccentricity in ways that would explain our data, future mfVEP experiments will focus on the central region.

A third, and related, point concerns the region of cortex generating the mfVEPs measured with our particular temporal and spatial pattern of stimulation. Recently, Hall et al., (2005) presented magnetoencephalography (MEG) evidence that the contrast–response (MEG) function of the striate and extrastriate visual cortices differed. The function for the extrastriate more closely resembled our mfVEP results, whereas that for the striate cortex more closely resembled the V1 model. However, their stimuli were presented in a single quadrant and had very different temporal and spatial distributions. In any case, recall that our mfVEPs appear to be coming mainly, if not entirely, from V1. The arguments in favor of this conclusion are presented in the Introduction section. In addition, we (Zhang & Hood, 2004) have shown that the first component of a principal component analysis (PCA) is generated in the striate cortex. A reanalysis of the data using PCA produced essentially identical results to the RMS analysis used here. In short, the evidence argues that we are studying striate (V1) activity.

Fourth, the amplitude of the mfVEP may not be linearly related to underlying spike activity. Assumption 5 states that the mfVEP is proportional to the sum of the responses of the V1 neurons. This is equivalent to the assumption in Heeger et al. (2000) that the fMRI signal is proportional to the sum of V1 neurons. In fact, it is generally accepted that VEP responses are largely generated, not by spiking activity, but by synaptic activity, the interplay of EPSP and IPSP activity on pyramidal cells, and the resulting local field potentials (e.g., Schroeder, Tenke, & Givre, 1992). However, the current evidence from V1 suggests that the BOLD signal of the fMRI is also more closely related to synaptic activity and local field potentials than to spiking activity (Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001; Logothetis & Wandell, 2004). (Although for human auditory cortex, Mukamel et al., 2005, found that spiking activity was well correlated with BOLD signals and that this correlation was at least as high as for the local field potentials.) Thus, for similar reasons, neither the VEP nor the fMRI activity may be linearly related to spiking activity. That is, Assumption 5 may be violated for both the mfVEP and fMRI comparison to the model. The degree to which this violation affects the comparisons remains an open question. Although local field potentials and spiking activity are highly correlated (e.g., Logothetis et al., 2001; Mukamel et al., 2005), it is clear that they do not reflect the same aspects of V1 activity (Henrie & Shapley, 2005; Logothetis et al., 2001).

Fifth, to the extent that either the mfVEP or fMRI are dependant on local slow potentials, they will depend upon both excitatory and inhibitory potentials. Further, given that the mfVEP is due to an interplay of positive and negative components, perhaps the response to higher contrasts contains different underlying components than does the response to lower contrasts. This could explain why high-contrast mfVEP responses deviate from the predictions of the model. However, the similarity of the waveforms (Figure 6) suggests similar underlying components at all contrasts.

In any case, assuming for now that a linear relationship exists between spiking activity and the fMRI signal, as proposed by Heeger et al. (2000), then the comparison of fMRI results to the mfVEP recordings become relevant to the hypothesis linking the mfVEP amplitude to neural activity. Figure 12 shows the data from Experiments 1 and 2 and the theoretical function (red) from Figure 8. Recall that the fMRI data from Boynton et al. (1999) presented in Heeger et al. fall along the red curve (see Figure 8). Thus, there is a discrepancy between their fMRI data and our mfVEP data at higher contrasts. However, their fMRI data

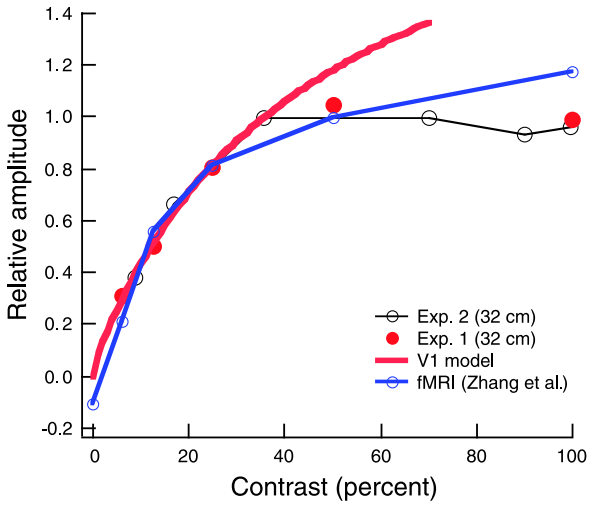

Figure 12. The relative RMS amplitudes from Experiments 1 and 2 (from Figure 7) shown with the theoretical curve (red) as in Figure 7 and the results (blue) from a fMRI experiment that used a similar display.

were obtained with a stimulus quite different from the mfVEP display used here. When our mfVEP results are compared to recent fMRI data obtained with a multifocal stimulus (Zhang, Hood, Ferrera, & Hirsch, 2005), similar to that used here, the fMRI results (blue curve) more closely approximate the mfVEP data. Although the mfVEP and fMRI experiments need repeating with identical stimuli, now there is no reason to believe that the function relating mfVEP to neural activity is necessarily different than the function relating fMRI BOLD signal to neural activity, although each may differ from the model for the reasons delineated above.

Latency and the model

Perhaps the most surprising aspect of our data is the relatively small change in latency seen with variations in contrast. The latency of the mfVEP changed by less than 5 ms (see Figure 6) with an increase in contrast. Baseler and Sutter (1997) also measured latency of mfVEP as a function of contrast. It is not easy to directly compare our results to their data as they did not report the latency of the mfVEP records per se. Rather, they measured the latency of two components extracted from the mfVEP using a complex algorithm. In any case, they also reported a modest increase in latency (about 10 ms) between 4%, the lowest contrast used, and 13% contrast with little change for higher contrasts.

The relatively modest change in latency for the mfVEP is in contrast to the latency changes, on the order of 40 ms, reported for primate V1 cells. For example, Albrecht (1995) reported a mean decrease in latency of 49.4 ms (n = 19) between threshold contrasts and 90% contrast or about 37 ms for the range from 5% to 90% contrast. In agreement, Gawne, Kjaer, and Richmond (1996) found a mean increase of 39.4 ± 2.3 ms (n = 37) in latency between their lowest contrast, 5%, and higher contrasts. Other studies report comparable changes (e.g., Albrecht et al., 2002; Carandini & Heeger, 1994; Carandini, Heeger, & Movshon, 1997). Whereas Albrecht's data show an increase of greater than 25 ms from 12.5% to 100% contrast, our data show an increase of less than 4–5 ms over this range.

We do not know why the single-cell data differ from the mfVEP data. However, we can offer three possible reasons for the difference. First, as mentioned above, the mfVEP is the sum of negative and positive contributions; hence, it is possible that a shift in the relative weighting of different component may give a false impression of latency. However, again, the striking similarity of waveform across contrasts argues against this possibility. The second possibility is based upon the observation that LGN cells appear to show smaller changes in latency as a function of contrast as compared with cortical cells. For example, Albrecht (1995) calculated that the decrease in latency (phase advance) with increased contrast of LGN X cells (Y cells) was smaller by a factor of 2.0 (1.5) as compared with the decrease in V1 cells. Because of these findings, it has been suggested to us that the mfVEP may reflect the LGN input to V1 neurons (Tony Movshon, personal communication). Although we cannot rule this out, subcortical (LGN) contributions to the conventional VEP are very small and appear very early in the time course (Schroeder et al., 1992). We would expect the same to be true for the mfVEP. Third, perhaps, cells whose responses are near their maximum amplitude and minimum latency dominate the mfVEP. V1 cells have a range of semisaturation values for latency, and these values are correlated with the semisaturation values for amplitude (Albrecht, 1995). To test this notion, we need a formal model for latency similar to the one provided for amplitude in Heeger et al. (2000).

Conclusion

The mfVEP contrast–response function is in general agreement with a model based upon the V1 neuron population. Consistent with this model, we found little change in the shape of the contrast–response function with change in the size of the elements of the stimulus. However, two aspects of the results require further exploration. First, there is the systematic deviation at higher contrasts between the model's predictions and the observed contrast–response function. This deviation suggests that one or more of the model's assumptions may be violated. Second, the latency of the mfVEP changed far less than expected based upon single-cell data.

Acknowledgments

We thank Duane Albrecht, Bill Geisler, and David Heeger for helpful discussion during the planning of this project. This work was supported in part by NIH/NEI grants EY02115 and EY0845.

Commercial relationships: none.
Corresponding author: Donald C. Hood.
Email: dch3@columbia.edu.
Address: Department of Psychology, Columbia University, New York, NY, USA.

References

Albrecht, D. G. (1995). Visual cortex neurons in monkey and cat: Effect of contrast on the spatial and temporal phase transfer functions. *Visual Neuroscience, 12,* 1191–1210. [PubMed]

Albrecht, D. G., Farrar, S. B., & Hamilton, D. B. (1984). Spatial contrast adaptation characteristics of neurones

recorded in the cat's visual cortex. *The Journal of Physiology, 347,* 713–739. [PubMed]

Albrecht, D. G., Geisler, W. S., Frazor, R. A., & Crane, A. M. (2002). Visual cortex neurons of monkeys and cats: Temporal dynamics of the contrast response function. *Journal of Neurophysiology, 88,* 888–913. [PubMed] [Article]

Albrecht, D. G., & Hamilton, D. B. (1982). Striate cortex of monkey and cat: Contrast response function. *Journal of Neurophysiology, 48,* 217–237. [PubMed]

Baseler, H. A., & Sutter, E. E. (1997). M and P components of the VEP and their visual field distribution. *Vision Research, 37,* 675–690. [PubMed]

Baseler, H. A., Sutter, E. E., Klein, S. A., & Carney, T. (1994). The topography of visual evoked response properties across the visual field. *Electroencephalography and Clinical Neurophysiology, 90,* 65–81. [PubMed]

Boynton, G. M., Demb, J. B., Glover, G. H., & Heeger, D. J. (1999). Neuronal basis of contrast discrimination. *Vision Research, 39,* 257–269. [PubMed]

Carandini, M., & Heeger, D. J. (1994). Summation and division by neurons in primate visual cortex. *Science, 264,* 1333–1336. [PubMed]

Carandini, M., Heeger, D. J., & Movshon, J. A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *The Journal of Neuroscience, 17,* 8621–8644. [PubMed] [Article]

Di Russo, F., Pitzalis, S., Spitoni, G., Aprile, T., Patria, F., Spinelli, D., et al. (2005). Identification of the neural sources of the pattern-reversal VEP. *Neuroimage, 24,* 874–886. [PubMed]

Fortune, B., & Hood, D. C. (2003). Conventional pattern-reversal VEPs are not equivalent to summed multifocal VEPs. *Investigative Ophthalmology and Visual Science, 44,* 1364–1375. [PubMed] [Article]

Gardner, J. L., Sun, P., Waggoner, R. A., Ueno, K., Tanaka, K., & Cheng, K. (2005). Contrast adaptation and representation in human early visual cortex. *Neuron, 47,* 607–620. [PubMed]

Gawne, T. J., Kjaer, T. W., & Richmond, B. J. (1996). Latency: Another potential code for feature binding in striate cortex. *Journal of Neurophysiology, 76,* 1356–1360. [PubMed]

Geisler, W. S., & Albrecht, D. G. (1997). Visual cortex neurons in monkeys and cats: Detection, discrimination, and identification. *Visual Neuroscience, 14,* 897–919. [PubMed]

Hall, S. D., Holliday, I. E., Hillebrand, A., Furlong, P. L., Singh, K. D., & Barnes, G. R. (2005). Distinct contrast response functions in striate and extra-striate regions of visual cortex revealed with magnetoencephalography (MEG). *Clinical Neurophysiology, 116,* 1716–1722. [PubMed]

Heeger, D. J., Huk, A. C., Geisler, W. S., & Albrecht, D. G. (2000). Spikes versus BOLD: What does neuroimaging tell us about neuronal activity? *Nature Neuroscience, 3,* 631–633. [PubMed] [Article]

Henrie, J. A., & Shapley, R. (2005). LFP power spectra in V1 cortex: The graded effect of stimulus contrast. *Journal of Neurophysiology, 94,* 479–490. [PubMed]

Hood, D. C., & Greenstein, V. C. (2003). Multifocal VEP and ganglion cell damage: Applications and limitations for the study of glaucoma. *Progress in Retinal and Eye Research, 22,* 201–251. [PubMed]

Hood, D. C., Greenstein, V. C., Odel, J. G., Zhang, X., Ritch, R., Liebmann, J. M., et al. (2002). Visual field defects and multifocal visual evoked potentials: Evidence of a linear relationship. *Archives of Ophthalmology, 120,* 1672–1681. [PubMed]

Hood, D. C., Ohri, N., Yang, E. B., Rodarte, C., Zhang, X., Fortune, B., et al. (2004). Determining abnormal latencies of multifocal visual evoked potentials: A monocular analysis. *Documenta Ophthalmologica: Advances in Opthalmology, 109,* 189–199. [PubMed]

Hood, D. C., Zhang, X., Greenstein, V. C., Kangovi, S., Odel, J. G., Liebmann, J. M., et al. (2000). An interocular comparison of the multifocal VEP: A possible technique for detecting local damage to the optic nerve. *Investigative Ophthalmology and Visual Science, 41,* 1580–1587. [PubMed] [Article]

Hood, D. C., Zhang, X., Hong, J. E., & Chen, C. S. (2002). Quantifying the benefits of additional channels of multifocal VEP recording. *Documenta Ophthalmologica: Advances in Opthalmology, 104,* 303–320. [PubMed]

James, A. C. (2003). The pattern-pulse multifocal visual evoked potential. *Investigative Ophthalmology and Visual Science, 44,* 879–890. [PubMed] [Article]

Klistorner, A., Crewther, D. P., & Crewther, S. G. (1997). Separate magnocellular and parvocellular contributions from temporal analysis of the multifocal VEP. *Vision Research, 37,* 2161–2169. [PubMed]

Klistorner, A. I., Graham, S. L., Grigg, J. R., & Billson, F. A. (1998). Multifocal topographic visual evoked potential: Improving objective detection of local visual field defects. *Investigative Ophthalmology and Visual Science, 39,* 937–950. [PubMed]

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature, 412,* 150–157. [PubMed]

Logothetis, N. K., & Wandell, B. A. (2004). Interpreting the BOLD signal. *Annual Review of Physiology, 66,* 735–769. [PubMed]

Maddess, T., James, A. C., & Bowman, E. A. (2005). Contrast response of temporally sparse dichoptic multifocal

visual evoked potentials. *Visual Neuroscience, 22,* 153–162. [PubMed]

Meigen, T., & Kraemer, M. (2005). Evaluation of the transmission time of visual evoked potentials to the visual cortex by multifocal recordings. *Investigative Ophthalmology and Visual Science, 46,* E-Abstract 3601.

Movshon, J. A., & Lennie, P. (1979). Pattern-selective adaptation in visual cortical neurones. *Nature, 278,* 850–852. [PubMed]

Mukamel, R., Gelbard, H., Arieli, A., Hasson, U., Fried, I., & Malach, R. (2005). Coupling between neuronal firing, field potentials, and fMRI in human auditory cortex. *Science, 309,* 951–954. [PubMed]

Schroeder, C. E., Tenke, C. E., & Givre, S. J. (1992). Subcortical contributions to the surface-recorded flash-VEP in the awake macaque. *Electroencephalography and Clinical Neurophysiology, 84,* 219–231. [PubMed]

Sclar, G., Lennie, P., & DePriest, D. D. (1989). Contrast adaptation in striate cortex of macaque. *Vision Research, 29,* 747–755. [PubMed]

Sclar, G., Maunsell, J. H., & Lennie, P. (1990). Coding of image contrast in central visual pathways of the macaque monkey. *Vision Research, 30,* 1–10. [PubMed]

Slotnick, S. D., Klein, S. A., Carney T., Sutter, E., & Dastmalchi, S. (1999). Using multi-stimulus VEP source localization to obtain a retinotopic map of human primary visual cortex. *Clinical Neurophysiology, 110,* 1793–1800. [PubMed]

Zhang, X., & Hood, D. C. (2004). A principal component analysis of multifocal pattern reversal VEP. *Journal of Vision, 4*(1), 32–43, http://journalofvision.org/4/1/4/, doi:10.1167/4.1.4. [PubMed] [Article]

Zhang, X., Hood, D. C., Chen, C. S., & Hong, J. E. (2002). A signal-to-noise analysis of multifocal VEP responses: An objective definition for poor records. *Documenta Ophthalmologica: Advances in Opthalmology, 104,* 287–302. [PubMed]

Zhang, X., Hood, D. C., Ferrera, J., & Hirsch, J. (2005). *The effect of attention and contrast on the BOLD response in V1 and beyond.* Manuscript submitted for publication.