

Representation of Concurrent Stimuli by Population Activity in Visual Cortex

Laura Busse,^{1,2,*} Alex R. Wade,^{2,3} and Matteo Carandini^{1,2}

¹UCL Institute of Ophthalmology, University College London, London EC1V 9EL, UK

²Smith-Kettlewell Eye Research Institute, San Francisco, CA 94115, USA

³Department of Neurology, University of California, San Francisco, San Francisco, CA 94117, USA

*Correspondence: laura.busse@cin.uni-tuebingen.de

DOI 10.1016/j.neuron.2009.11.004

SUMMARY

How do neuronal populations represent concurrent stimuli? We measured population responses in cat primary visual cortex (V1) using electrode arrays. Population responses to two superimposed gratings were weighted sums of the individual grating responses. The weights depended strongly on the relative contrasts of the gratings. When the contrasts were similar, the population performed an approximately equal summation. When the contrasts differed markedly, however, the population performed approximately a winner-take-all competition. Stimuli that were intermediate to these extremes elicited intermediate responses. This entire range of behaviors was explained by a single model of contrast normalization. Normalization captured both the spike responses and the local field potential responses; it even predicted visually evoked currents source-localized to V1 in human subjects. Normalization has profound effects on V1 population responses and is likely to shape the interpretation of these responses by higher cortical areas.

INTRODUCTION

Even the simplest sensory stimulus or planned movement causes a large pool of neurons to be active. The code for both sensory perception and motor output, therefore, is thought to lie in the collective activation profile of neuronal populations (Georgopoulos et al., 1982; McIlwain, 1986; Nicolelis et al., 1995; Pouget et al., 2000). This profile has been measured mostly when a population is faced with an individual sensory stimulus or a single motor output. In such cases, the profile of population activity is typically bell shaped, with the strength of each neuron's response depending on the match between the neuron's preferences and the sensory signal or the planned movement (Chen et al., 2006; Georgopoulos et al., 1986; Purushothaman and Bradley, 2005).

In nature, however, a sensory system is often confronted with the conjoint presence of multiple stimuli. Likewise, a motor system needs to represent a combination of multiple movements. In these circumstances, the population responses will not simply have a single peak centered on a particular stimulus

or movement, but rather a combination of multiple peaks (MacEvoy et al., 2009; Pasupathy and Connor, 2002; Treue et al., 2000). Understanding the rules of this combination is fundamental to our comprehension of population coding.

To investigate the representation of multiple stimuli in a neuronal population, we measured responses in primary visual cortex (V1) to sums of two component gratings. To control the strength of sensory stimulation, we varied grating contrast. To control the identity of the stimulated neurons, we varied grating orientation. Having control of both quantities provided a space that is ideal to test models of population coding.

We first investigated the impact of stimulus contrast and found that population responses are contrast-invariant: contrast scales the population response profiles without changing their shape. By comparing response properties of populations to those of the underlying single neurons, we explain how this invariance relates to the invariance of tuning curves seen in single units (Finn et al., 2007; Sclar and Freeman, 1982).

We then asked how population responses in V1 represent two superimposed component gratings (plaids), and we found that these responses could be well approximated by a weighted sum, in which the weights applied to the component gratings depend on contrast. The weights given to the component gratings of the plaid were always smaller than the sum of the weights to the individual components, consistent with the well-known phenomenon of cross-orientation suppression (DeAngelis et al., 1992; Morrone et al., 1982). We also found a profound effect of relative contrast: a gradual transition between two regimes. When the component contrasts are similar, the population gives sizeable weights to both components (equal summation). When the contrasts are dissimilar, however, the responses overwhelmingly favor the component with higher contrast (winner-take-all competition).

We were able to capture all these phenomena with a simple model based on contrast normalization. This model involves a division between a numerator that sums contributions of individual components and a denominator that grows with overall contrast. The model has been previously applied to individual neurons that are optimally tuned to one of the component gratings (Carandini et al., 1997; Freeman et al., 2002; Heeger, 1991; Heeger, 1992; Heuer and Britten, 2002). Here we extend it to entire populations and show that it can accurately describe their widely different regimes of operation.

To demonstrate the relevance of these findings to visual perception, we investigate them not only in the spike responses of anesthetized cats but also in visually evoked currents

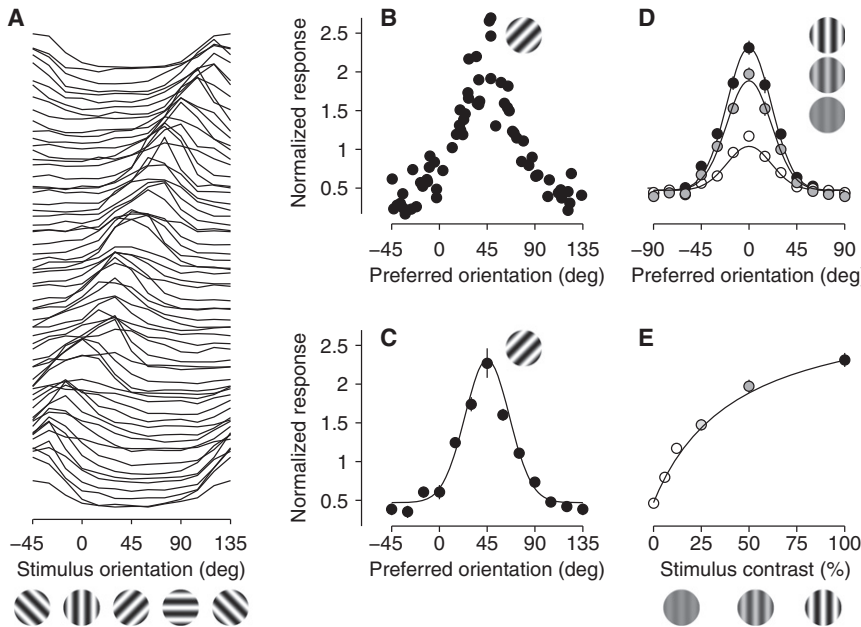


Figure 1. Tuning Curves and Population Responses to Single Orientation Stimuli

(A) Orientation tuning curves of all responsive sites (66 of 96), sorted according to preferred orientation. Each tuning curve is normalized by its mean across orientations.

(B) Population response to a 45° stimulus: responses of all sites (dots) as a function of preferred orientation of each site.

(C) The population response in (B) after binning sites with similar orientation preference (bin width: 15°). The curve is the best fitting circular Gaussian. Error bars indicate ±1 SE of responses across sites in each bin.

(D) Population response to a 0° stimulus for three contrasts: 12%, 50%, and 100%. The abscissa indicates preferred orientation relative to stimulus orientation. Data for stimuli of multiple orientations (0°, 30°, and 60°) are combined to obtain each population response. Error bars indicate ±1 SE of responses across sites in each bin. The curves fitting the data are circular Gaussians differing only in amplitude.

(E) Amplitude of the population responses as a function of stimulus contrast. The curve is the best-fitting hyperbolic ratio function ($c_{50} = 42.1\%$, $n = 1.0$). All fits are given by Equation (1). Experiment 84-12-16.

source-localized to V1 of human subjects. Moreover, we show with simple simulations how normalization in area V1 can have profound effects on the interpretation of V1 signals by higher cortical areas and, thus, shape perceptual judgments.

RESULTS

To measure the responses of a large population of neurons in area V1, we recorded from a 10 × 10-electrode array implanted in anesthetized cats (Figure 1A). The array covered an area of 16 mm² so that it included regions with a diversity of orientation preferences. As a result, spike responses measured at individual sites exhibited tuning curves whose preferences covered fairly uniformly the range of orientations (Figure 1A).

As expected, a stimulus of a given orientation evokes across the population a response whose profile peaks at the neurons that prefer the stimulus orientation (Figures 1B and 1C). This profile was well fit by a simple circular Gaussian function $G_{\varphi}(\theta)$ centered on stimulus orientation φ and varying with the preferred orientation of the neurons θ (Figure 1C).

Contrast-Invariance of Population Responses

We asked how these population responses are affected by stimulus contrast, and we found them to be invariant: changing contrast affected their profile in amplitude but not in width (Figures 1D and 1E). To test for invariance, we fitted the responses to an oriented stimulus (φ) by a separable model, the product of the Gaussian function of preferred orientation θ and a function of stimulus contrast c :

$$R_{\varphi}(\theta, c) = G_{\varphi}(\theta)f(c). \quad (1)$$

We tested this model by choosing a typical function for stimulus contrast (Figure 1E), the hyperbolic ratio that is commonly applied to individual neurons (Albrecht and Hamilton, 1982; Heeger, 1992):

$$f(c) = r_{\max} \frac{c^n}{c_{50}^n + c^n}, \quad (2)$$

where the parameters r_{\max} , c_{50} , and n determine the overall responsiveness, the semisaturation contrast, and the exponent of an accelerating nonlinearity related to spike threshold. The separable model provided excellent fits to the population responses. It explained 98.8% of the variance in the example data set (Figures 1D and 1E) and an average of 98.5% of the variance in all nine data sets. These results confirm earlier indications obtained by intrinsic imaging (Carandini and Sengpiel, 2004): population responses to individual orientations are contrast invariant.

A factor contributing to this invariance of population responses is surely the well-known invariance of tuning curves in individual neurons. Such tuning curves are scaled by stimulus contrast without changes in shape (Finn et al., 2007; Sclar and Freeman, 1982). The invariance seen in single neurons, however, is not sufficient to explain the invariance of population responses. For example, if neurons that are more sharply tuned responded only to higher contrasts, then increasing contrast would narrow the profile of population responses.

The invariance of population responses, indeed, makes a strong prediction that orientation tuning width and contrast sensitivity should be distributed independently across neurons. Increasing contrast, then, would not preferentially engage neurons that are more or less sharply tuned, thus leaving the

width of the population profile unchanged. This intuitive argument is formalized mathematically in the Appendix in the [Supplemental Data](#) (available online).

We tested this prediction and indeed found no systematic relationship between tuning width and contrast sensitivity across individual neurons (Figure 2). We measured orientation tuning curves (Figure 2A) and contrast responses (Figure 2B) in 75 well-isolated V1 neurons individually recorded in a separate set of experiments. We quantified contrast sensitivity of each neuron by the two parameters of the hyperbolic ratio: semisaturation contrast c_{50} , and exponent n . We then asked whether these parameters are independent of the neuron's tuning width (half-width at half height). We assessed statistical independence by calculating quartiles of the marginal distributions and testing for uniformity of the joint distribution. We saw no departure from independence between tuning width and semisaturation contrast c_{50} (Figure 2C, $\chi^2 = 4.57$; $df = 7$; $p = 0.71$). Similarly, there was no departure from independence between tuning width and exponent n (data not shown; $\chi^2 = 7.56$; $df = 7$; $p = 0.37$).

This analysis reveals that the properties of individual neurons and those of the population can be quite different. The mean tuning width of the individual neurons was $23^\circ \pm 14^\circ$ (mean \pm SD; $n = 75$; Figure 2A), comparable to the $22^\circ \pm 2^\circ$ measured in the population ($n = 9$ experiments; Figure 1D). The semisaturation contrast was also similar: $20\% \pm 11\%$ for individual neurons (median \pm MAD; Figure 2B) and $22\% \pm 9\%$ for the population (Figure 1E). The exponent, however, was much lower in the population than in most of the individual neurons: 2.2 ± 1.0 (median \pm MAD; Figure 2B) for individual neurons but only 1.0 ± 0.1 for the population. The contrast responses of the population (Figure 1E), therefore, are much shallower than those of most individual neurons.

These similarities and differences between neurons and population are simply explained. We built predicted population responses by averaging the fitted responses of the individual neurons after aligning their preferred orientation (Figures 2A and 2B, black curves). The predicted population tuning profile has a width of 20° , intermediate to that of individual neurons and similar to that of the actual population. The predicted population contrast response, instead, is much shallower than that of the average individual neuron: it has a similar semisaturation contrast (20%) but a lower exponent (1.4). These results echo the measurements made in the actual population. The shallow response derives from the fact that individual neurons have a broad range of semisaturation contrasts (Figure 2B).

Summation and Competition in Population Responses

We then asked how a population represents more than one stimulus and measured responses to plaids (Figure 3A). Plaids were obtained by superimposing two component gratings with different orientations; the component contrasts were independently varied.

The population responses to these plaids depended strongly on the relative contrasts of the components (Figure 3B). If the component contrasts were similar (e.g., 6% and 12%) or identical (e.g., both 12%), the population response to the plaid exhibited two peaks, one at each of the component orientations. However, if the component contrasts differed considerably

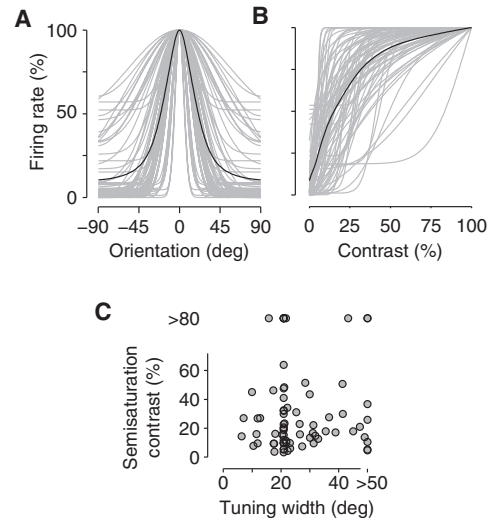


Figure 2. Relationship between Single Neuron Tuning Curves and Response Properties of the Population

(A) Orientation tuning curves of single neurons ($n = 75$, gray), normalized to the maximum response and centered on preferred orientation. Superimposed is the average orientation response profile of the population (black). (B) Contrast response functions. Same format as in (A). (C) Semisaturation contrast (obtained by fitting Equation 2) plotted as a function of tuning width (half-width at half-height) as obtained by fitting a circular Gaussian to each neuron's response. Tuning width and semisaturation contrast are independently distributed across neurons in the population.

(e.g., 50% and 12%), the response to the component of higher contrast (50% horizontal) dominated the population response to the plaid (Figure 3B, lower right). In these conditions, it is as if the 12% contrast vertical grating had almost disappeared. However, when this stimulus was presented alone it elicited a large response (Figure 3B, lower left).

To characterize these responses we described them with a weighted-sum model (Figure 3C). In this model, the response to the plaid with component orientations φ_1 and φ_2 and contrasts c_1 and c_2 is given by a linear combination of the responses to the component gratings:

$$\mathbf{R}_{1+2}(c_1, c_2) = w_1(c_1, c_2)\mathbf{R}_1(c_1) + w_2(c_1, c_2)\mathbf{R}_2(c_2), \quad (3)$$

where the responses to the individual components $\mathbf{R}_1(c_1)$ and $\mathbf{R}_2(c_2)$ are given by Equation (1) with Gaussians centered on φ_1 and φ_2 , and the scaling factors w_1 , w_2 depend on the combination of contrasts. Here and elsewhere, we express population responses (quantities in bold letters) as vectors: they are functions of stimulus orientation. These vectors are in turn affected by stimulus contrast.

Because the weighted-sum model provided good fits, its best-fitting weights w_1 , w_2 can be used to describe the rules of combination. Across contrast combinations, adding the two component gratings resulted in weights < 1 , a signature of cross-orientation suppression. Furthermore, the weights greatly depended on relative component contrasts. Weights were sizeable for both components if the two contrasts were similar ($w_1 = 0.99$, $w_2 = 0.24$) or identical ($w_1 = w_2 = 0.77$), resulting in two clear peaks in the activity profile (Figure 3C, left and middle). However,

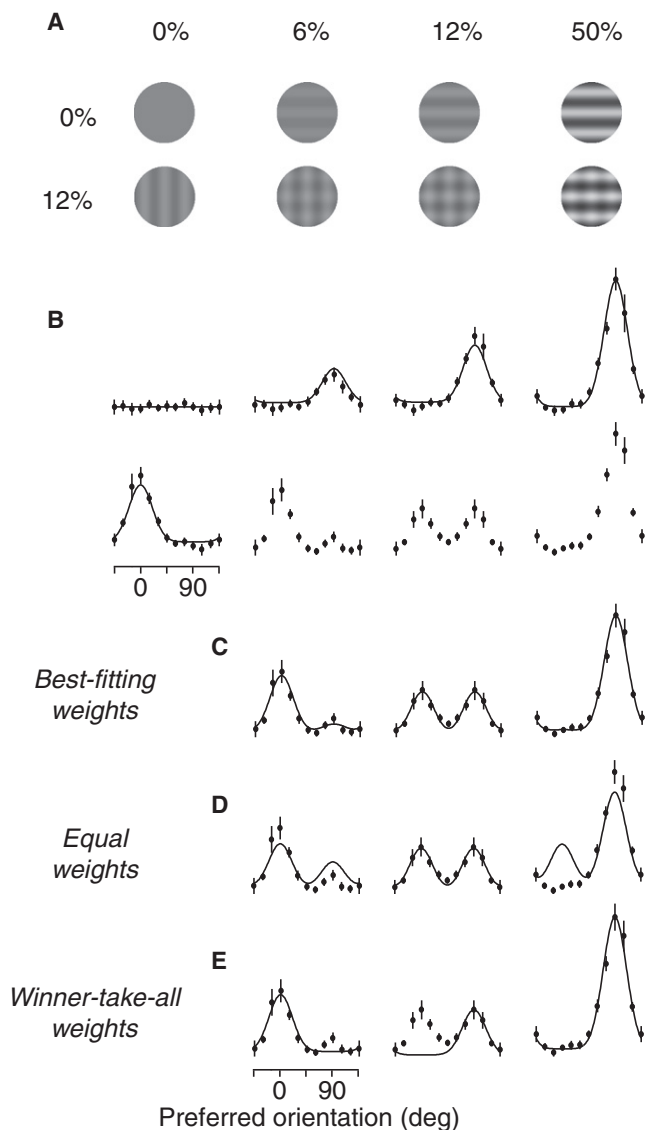


Figure 3. Population Responses to Plaids and Predictions of the Weighted-Sum Model

(A) Stimuli were gratings of two different orientations (first row and first column) and plaids obtained by summing the individual component gratings, for different combinations of component contrasts.

(B) Population responses to these stimuli. Error bars indicate ± 1 SE of responses across sites in each bin. Model predictions for the single grating responses are given by the separable model of orientation and contrast (traces).

(C–E) Predictions of the weighted-sum model for the population responses to plaids with best-fitting weights (C), equal weights (D), and winner-take-all weights (E). The data are replotted from (B) for comparison. Data in (B)–(E) are from experiment 83-10-15, plaid angle = 90° .

if the component contrasts differed considerably, the weight assigned to the lower-contrast component was close to zero ($w_1 = 0.02$, $w_2 = 0.92$) resulting in a single peak (Figure 3C, right).

To understand these values of the weights it helps to consider two extreme scenarios: equal weights and winner-take-all weights (Figures 3D and 3E). In the first scenario, the population

applies equal weights to both component responses: $w_1(c_1, c_2) = w_2(c_1, c_2)$. This scenario entails the presence of two peaks in population responses to all plaids, leading to reasonable or excellent fits if component contrasts are similar or identical (Figure 3D, left and middle), but unacceptable fits if component contrasts differ considerably (Figure 3D, right). In the second scenario, instead, the population response completely disregards the grating of lower contrast, $w_1 = 0$ if $c_1 < c_2$ and vice versa. Assuming winner-take-all weights leads to excellent predictions if component contrasts differ considerably (Figure 3E, right), but not if they are similar or equal (Figure 3E, left and middle). Indeed, winner-take-all weights can only predict a single peak in the activity profile. The sets of weights predicted by the two scenarios, therefore, have a complementary pattern of successes and failures.

To extend these qualitative findings we explored population responses obtained with a broad range of contrast combinations and with different plaid angles (Figures 4A–4C). As in the previous examples, these population responses were closely described by the weighted-sum model (Figure 4A, solid lines). We searched for the weights w_1 and w_2 that yielded the best fits and obtained a pair of weights for each contrast combination (Figure 4B). With these best-fitting parameters, the weighted-sum model made excellent predictions (mean fit quality index $q = 94.4\% \pm 0.5\%$ SE; $n = 9$; Figure 4C).

The best-fitting weights w_1 and w_2 depended markedly on component contrast (Figure 4B). First, the weights for plaids were consistently below unity, indicating that the response to a grating forming part of a plaid is smaller than the response to that grating presented alone (cross-orientation suppression). Second, with components of similar contrast both weights are considerable (Figure 4B, diagonal). Third, when component contrasts differ substantially, the weight given to the component of lower contrast is much weaker (and sometimes negligible) than the weight given to the component of higher contrast (Figure 4B, off the diagonal).

These data therefore reveal a range of behaviors extending from a regime of equal summation (at work when component contrasts are similar) to a regime of winner-take-all competition (at work when contrasts are very dissimilar). In intermediate contrast conditions the weights are intermediate between these extremes. To delineate the contrast conditions in which these two regimes operate, we return to the two extreme scenarios mentioned earlier: equal weights and winner-take-all weights (Figures 4D–4G). Assuming equal weights ($w_1 = w_2$, Figure 4D) leads to excellent predictions when component contrasts are similar ($q = 93.8\% \pm 0.8\%$ SE; Figure 4E, diagonal region), but results in poor predictions elsewhere ($q = 81.7\% \pm 1.1\%$). Assuming winner-take-all weights ($w_1 = 0$ if $c_1 < c_2$, Figure 4F) yields the opposite pattern of successes and failures (Figure 4G): poor fits when contrasts are similar ($q = 72.8\% \pm 1.2\%$), and good predictions elsewhere ($q = 91.9\% \pm 0.4\%$).

Normalization of Population Responses

The weighted-sum model describes accurately how the population integrates two superimposed inputs (Figures 4A–4C), but it does not capture explicitly the effects of component contrast. For each combination of component contrasts c_1 and c_2 , the

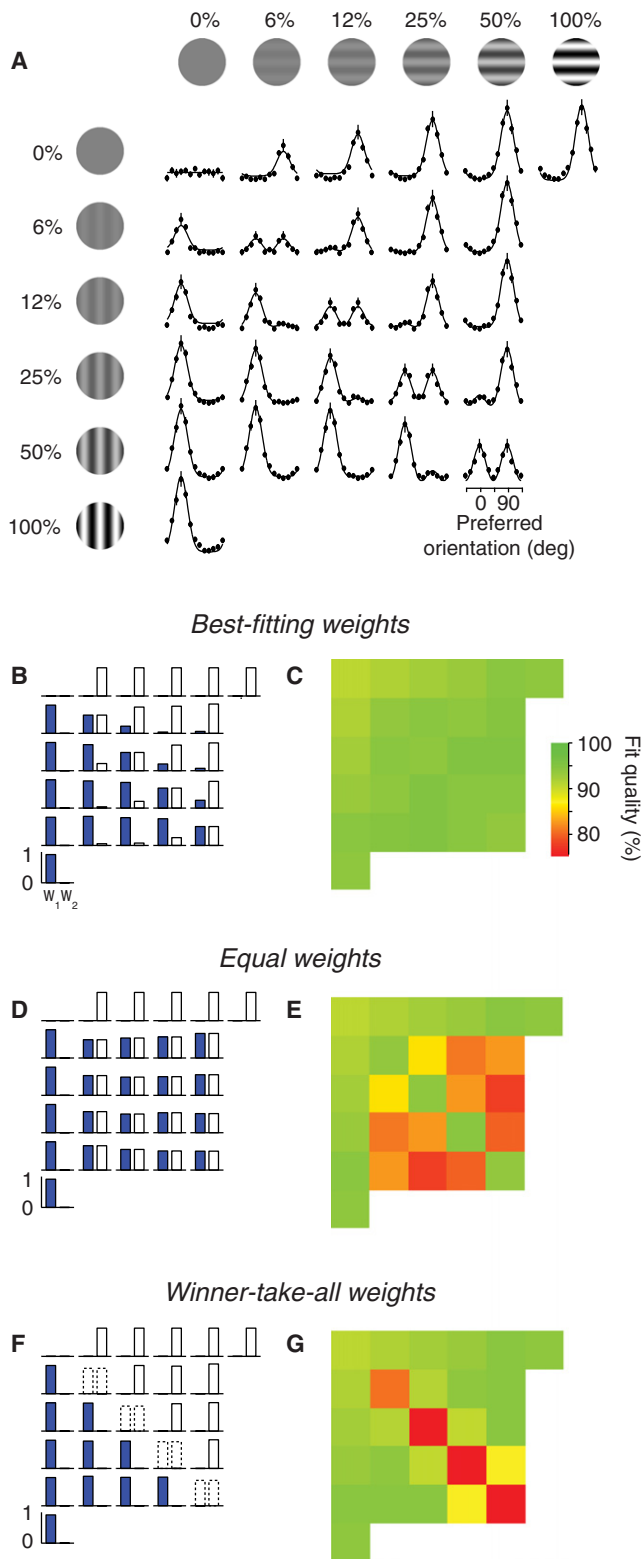


Figure 4. Predicted Weights and Fit Quality of the Weighted-Sum Model

(A) The complete stimulus set: gratings of different contrasts and their combinations (plaids). Population responses to plaids are fitted by the weighted sum

model requires two free parameters, the weights $w_1(c_1, c_2)$ and $w_2(c_1, c_2)$. If a new combination of contrasts c_1 and c_2 were to be tested, the model would make no prediction for the relevant weights.

What is needed, therefore, is a model that embodies the weighted-sum model while explicitly representing the role of contrast. Such a model should predict approximately equal weights when component contrasts are similar and approximately winner-take-all weights when component contrasts are dissimilar.

A promising candidate is the model of contrast normalization that has been developed for individual neurons. Normalization involves a ratio: the numerator sums the contributions of the different stimuli, weighted nonlinearly by component contrast, and the denominator scales these contributions on the basis of overall contrast (Albrecht and Geisler, 1991; Carandini et al., 1997; Heeger, 1991; Heeger, 1992; Kouh and Poggio, 2008). Normalization explains cross-orientation suppression for individual neurons, the ones that are selective for one of the two orientations in the stimulus (Carandini et al., 1997; Freeman et al., 2002). Can normalization predict the responses of the whole population, and can it predict the gradual transition between the regimes of equal summation and winner-take-all competition?

To apply the normalization model to our population responses, we express it as

$$R_{1+2}(C_1, C_2) = r_{\max} \frac{c_1^n G_1 + c_2^n G_2}{c_{50}^n + c_{rms}^n}, \quad (4)$$

where r_{\max} , c_{50} , and n are constants, and G_1, G_2 are the usual Gaussians centered on the component orientations ϕ_1 and ϕ_2 . The term $c_{rms} = \sqrt{c_1^2 + c_2^2}$ is the root-mean-square contrast of the stimulus. In the model, therefore, the component contrasts c_1 and c_2 appear both in the numerator and in the denominator. If one of these contrasts is zero (i.e., if the stimulus is a single grating), the model reduces to Equations (1) and (2), and hence incorporates contrast invariance. If instead both contrasts are positive, the model provides a closed-form prediction for how the two components of a plaid should combine to yield a single population response.

Crucially, the normalization model predicts a gradual transition of behaviors from equal summation to winner-take-all competition. As shown in the Appendix in the Supplemental Data, in two key cases the normalization model approximates the weighted sum model and we can derive expression for the weights: when the two component contrasts are similar and when they are very dissimilar. In the first case, the model approximates the scenario of equal weights, and in the second case it approximates the scenario of winner-take-all weights.

of the responses to the component gratings (weighted-sum model, curves). Error bars indicate ± 1 SE of responses across sites in each bin.

(B–G) Predicted weights (left) and average fit quality (right) for $n = 9$ experiments, for the weighted-sum model with best-fitting weights (B and C), equal weights (D and E), and winner-take-all weights (F and G) presented in the same format as (A). Plaid angle was 90° in 5 experiments, 45° in 2 experiments, and 30° in 2 experiments. Data in (A) are from experiment 82-6-3, plaid angle = 90° .

These behaviors arise rather intuitively from the combination of operations in the normalization model (Equation 4). Equal summation is a natural consequence of the sum in the numerator; winner-take-all competition, in turn, stems from the scaling of the summation terms by contrast; with the help of an exponent >1 , this exponent introduces imbalances in the effects of the two gratings when contrasts differ. Across all contrast conditions, the denominator scales the whole response, causing the effective weights to be <1 (cross-orientation suppression). This effect is of course most striking in the regime of winner-take-all competition, where the denominator reduces the weight applied to the weaker response to almost zero.

We tested the normalization model on the population responses and found that it provided good fits in both regimes (Figure 5). For the example experiment (Figure 5A) the model captures qualitatively all the fundamental behaviors of the population. Fits to other data sets, including those where the component gratings had different angles, were of comparable or higher quality (see Figure S1 available online). The good performance of the model is confirmed by an analysis of effective weights and fit quality across experiments (Figures 5B and 5C). The model captured the regime of equal summation seen when component contrasts are similar, predicting similar weights for the two component responses (Figure 5B, diagonal) and yielding good fits ($q = 91.5\% \pm 1.0\%$; Figure 5C). The model also captured the regime of winner-take-all competition, predicting progressively smaller weights for gratings whose contrast decreases relative to the other grating (Figure 5B, off-diagonal) and again yielding good fits ($q = 91.7\% \pm 0.6\%$ SE; Figure 5C). Median parameters of the normalization model were $n = 1.5 \pm 0.13$, $c_{50} = 13.1 \pm 4.2$, and $\sigma = 19.0 \pm 2.1$.

Importantly, in the normalization model the gradual transition from one regime to the other is explicitly controlled by the component contrasts, without requiring a change in model parameters. The normalization model, indeed, performs almost as well as the weighted-sum model with optimal weights. Yet it requires only five parameters: one for overall responsiveness, r_{\max} , two for contrast responses (c_{50} and n), and two for the circular Gaussian \mathbf{G} (width and vertical offset). The weighted-sum model, by comparison, requires the same five parameters plus two weights $w_1(c_1, c_2)$ and $w_2(c_1, c_2)$ for each combination of positive grating contrasts c_1 and c_2 . For our data sets that makes $5 + 16 = 21$ free parameters, more than four times the five free parameters in the normalization model.

Normalization in Subthreshold Responses

The physiological mechanisms underlying response normalization in individual V1 neurons are at the moment unclear and may rest on a combination of factors (Carandini et al., 1997; Carandini et al., 2002; Chance et al., 2002; Finn et al., 2007; Freeman et al., 2002; Priebe and Ferster, 2006). It is therefore of interest to know to what degree normalization is present in the subthreshold responses as opposed to spiking responses.

To measure subthreshold activity in populations of neurons, we analyzed local field potential (LFP) responses measured with the same 10×10 electrode array that measures the spike responses. The LFP comprises the combined subthreshold activity of the neurons surrounding the electrode (Katzner

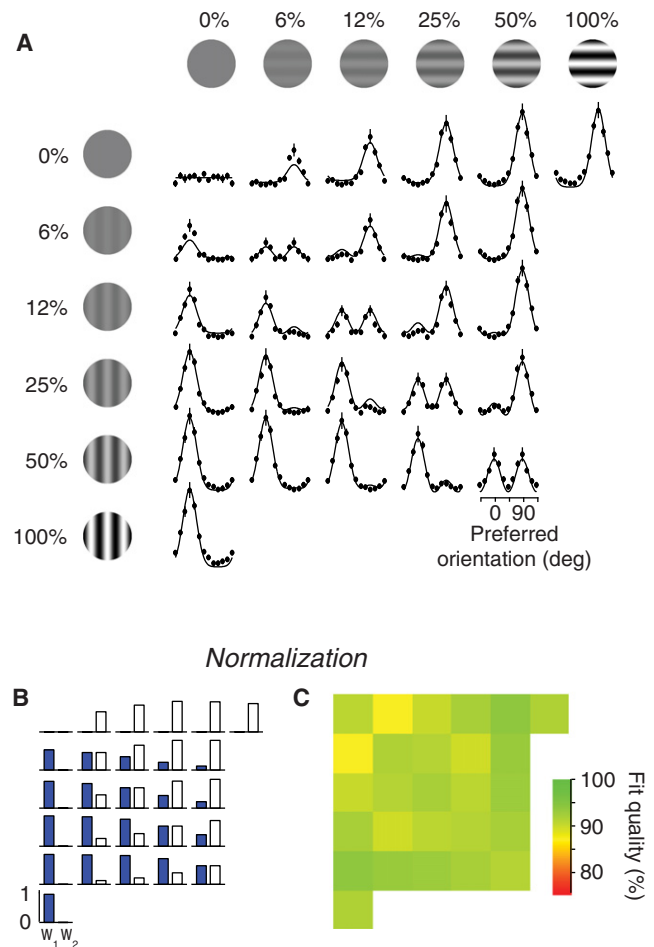


Figure 5. Predicted Population Responses, Weights, and Fit Quality of the Normalization Model

(A) Population responses are fitted by the normalization model with a single set of parameters. Error bars indicate ± 1 SE of responses across sites in each bin. (B) Weights predicted by the normalization model. (C) Average fit quality of the normalization model. Data in (A) are from experiment 82-6-3, plaid angle = 90° .

et al., 2009). To obtain the highest signal-to-noise ratios, we pooled these LFPs across all responsive sites of the array and across all animals ($n = 4$). To distinguish the responses to the test from the responses to the mask, we employed a frequency tagging method typically used in EEG research (Candy et al., 2001; Morrone and Burr, 1986; Regan, 1989): we made the test and mask contrast-reverse at different frequencies. These frequencies effectively tag the LFP responses, which oscillate at twice the frequency of reversal (Katzner et al., 2009).

This design provides an opportunity to test the same normalization model that we have applied to spike responses. Because of the pooling across sites, we can no longer study population responses as a function of preferred orientation. By tagging the two stimuli by frequency, though, we can distinguish the responses to test and mask. The predicted responses R_j , accordingly, depend on stimulus tag, $j = 1$ for the test and $j = 2$ for the mask:

$$R_j(C_1, C_2) = r_{\max} \frac{C_j^n}{C_{50}^n + C_{rms}^n} \quad (5)$$

Because adding a superimposed grating affects only the denominator, in this design cross-orientation suppression provides a frank reduction in tagged responses.

The LFP signals agreed closely with these predictions (Figures 6A and 6B). First, the LFPs evoked by the test grew and eventually saturated as the test contrast increased (Figure 6A, open symbols). The normalization model closely predicted this saturation (Figure 6A, curves). Second, these LFPs were reduced by adding a 25% mask grating (Figure 6A, closed symbols). The normalization model closely matched this effect, correctly predicting a rightward shift of the contrast response function (Figure 6A, curves). Third, the suppressive interactions between test and mask were mutual: increasing test contrast reduced the LFPs evoked by the mask (Figure 6B, closed symbols). The model did the same (Figure 6B, curves).

Because the predictions of the normalization model were accurate (98.2% explained variance), the model's parameters can be used to compare subthreshold responses to spike responses. We found that subthreshold population responses measured from LFPs saturated at lower contrasts ($c_{50} = 13\%$; $n = 1.5$) than those measured earlier from population spike responses ($c_{50} = 22\%$; $n = 1.0$). A similar difference between contrast responses of subthreshold potentials and spikes has been seen in individual neurons and may be a simple consequence of spike threshold (Finn et al., 2007).

Normalization in Human V1

Does normalization govern the population responses only under anesthesia or can it be used to explain responses in the awake cortex? Does it apply to humans? We sought an answer to these questions by measuring the electroencephalogram (EEG) of human subjects with a sensor net of 128 electrodes. From these signals, we obtained visual evoked potentials (VEPs). Human VEP responses to plaid stimuli exhibit suppressive interactions between the grating components (Burr and Morrone, 1987; Candy et al., 2001; Ross and Speed, 1991), and such interactions may be explained by normalization (Candy et al., 2001).

The VEP is a close approximation to the pooled LFP measured within the underlying cortex (Schroeder et al., 1991). To measure it, therefore, we used stimuli similar to those used to measure LFP in anesthetized cats: we made the test and mask gratings reverse in contrast at different frequencies. Test and mask thus caused distinct responses that oscillated at twice the reversal frequencies (Candy et al., 2001). To extract responses originating specifically from area V1, we estimated the current source density across the entire cortex and then defined a region of interest based on maps of retinotopy measured with fMRI (Appelbaum et al., 2006).

The amplitudes of the visually evoked currents source-localized to V1 in humans were remarkably similar to the pooled LFP signals in anesthetized cats (Figures 6C and 6D). First, increasing test contrast increased the responses to the test (Figure 6C, open symbols). Second, adding the mask to the test reduced these responses (Figure 6D, closed symbols). Third, increases in test contrast reduced the responses to the mask

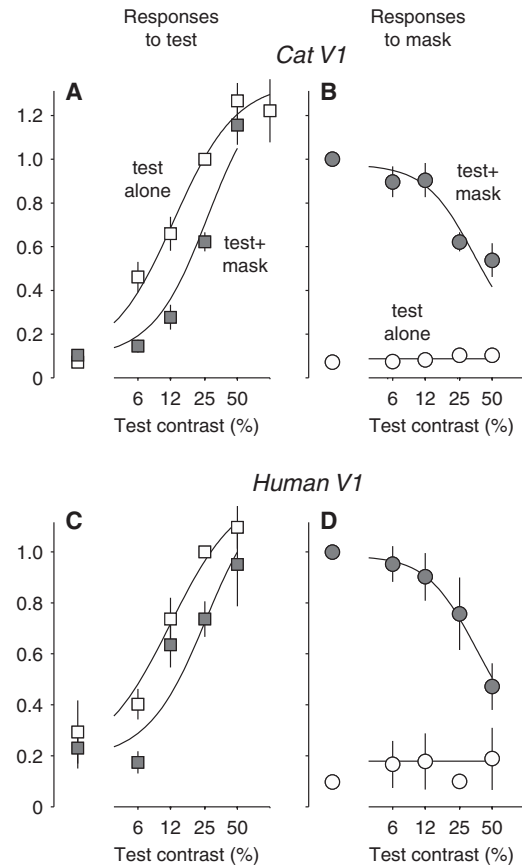


Figure 6. Normalization in Large-Scale Subthreshold Activity in Cats and Humans

Curves are fits of the normalization model. Error bars indicate standard error across subjects.

(A and B) LFPs obtained from area V1 in anesthetized cats in response to the test stimulus (A) and the mask stimulus (B). Response amplitudes were extracted at twice the stimulus frequency. Open symbols are conditions in which the test is presented alone, closed symbols indicate conditions in which a 25% contrast mask is added.

(C and D) Source-imaged VEP signals obtained from area V1 in human observers using current source density modeling and fMRI retinotopic mapping. Format as in (A) and (B). Responses in individual experiments were normalized to yield a value of 1 when the test stimulus was presented alone at 25% contrast.

stimulus (Figure 6D). Therefore, in the population responses measured in human subjects (just as in those measured from anesthetized animals), there is competition between the stimuli.

The human V1 population responses, again, were closely predicted by the normalization model (Equation 5) with fit parameters that were almost identical to those estimated for cat LFPs (Figures 6C and 6D, curves). As in anesthetized cats, the model (Equation 5) captures the competitive interactions between test and mask: the contrast response function to the test is shifted to the right in the presence of a mask (Figure 6C, curves), and responses to the mask decrease as the test contrast increases (Figure 6D, curves). As in cat LFPs, these predictions were accurate (96.2% explained variance). The model parameters for

fitting these human V1 responses ($c_{50} = 12\%$; $n = 1.4$) were similar to those for fitting pooled LFP responses in anesthetized cats ($c_{50} = 13\%$; $n = 1.5$). The similarity of subthreshold population responses in anesthetized cats and human observers indicates that V1 responses in both species and across behavioral states are governed by the same functional mechanism: contrast normalization.

Impact of V1 Normalization on Higher Areas

We finally asked what consequences our findings may have on the decoding of V1 signals by higher cortical areas. We have seen that normalization explains the regimes of equal summation and competition observed in V1 and that competition can be so extreme as to correspond to a winner-take-all effect. Such profound competition is bound to have consequences on downstream areas.

To investigate these consequences, we implemented a well-established model of a pattern-selective neuron in area MT (Figures 7A–7C). When stimulated with plaids, such neurons respond to the global direction of the plaid rather than the direction of motion of the individual components (Movshon et al., 1985). A successful model for these neurons (Rust et al., 2006; Simoncelli and Heeger, 1998) postulates that they sum the activity of a population of V1 neurons with weights appropriate to obtain the desired preferred direction (Figure 7A). We implemented this model for a model MT neuron whose direction tuning curve measured with drifting gratings peaks at 180° (Figure 7B). The neuron exhibits the same selectivity when the stimulus is a plaid composed of two gratings differing in direction by 120° : it responds most strongly when the global direction of the plaid is 180° (Figure 7C, black curve). The model incorporates contrast normalization in V1, but none of these phenomena rests on such normalization: neglecting normalization simply results in larger responses, without changes in tuning (Figure 7C, gray curves).

Normalization in area V1, however, can profoundly affect the responses of this MT neuron to a plaid with unequal component contrasts (Figure 7D). When component contrasts are 25% and 75%, normalization of population responses in area V1 leads to marked competition. The lower-contrast stimulus almost disappears from the population responses of V1, so the model MT neuron responds mostly to the higher contrast stimulus. Its tuning for plaid direction shifts toward 240° (i.e., the plaid direction in which the higher contrast grating drifts at 180°) (Figure 7D, black curve). The model MT neuron, in other words, has largely lost its ability to represent the overall motion of the plaid. This effect is almost entirely explained by normalization: neglecting normalization would lead to much smaller effects (Figure 7D, gray curve). This simple example, therefore, illustrates how normalization of population responses in area V1 can have profound consequences for the response properties of areas downstream in the visual hierarchy.

DISCUSSION

We investigated the representation of concurrent stimuli in the population responses of visual cortex. When the visual system is faced with multiple simultaneous orientations and contrasts, the population response can be described as a weighted sum

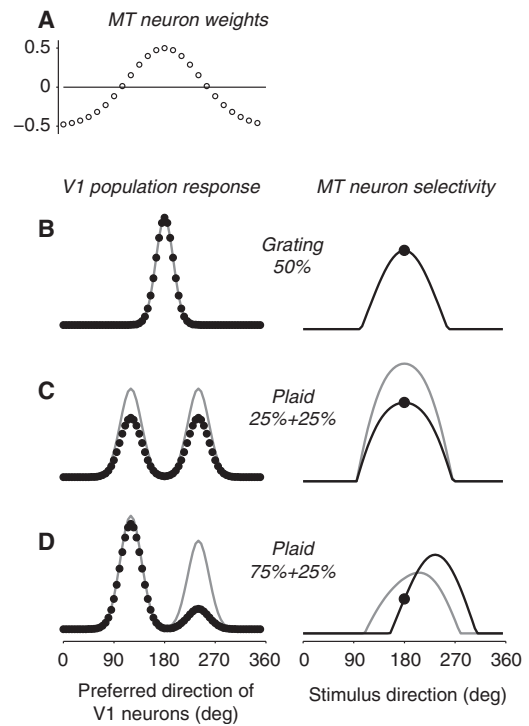


Figure 7. Model of the Effects of V1 Normalization on a Pattern-Selective MT Neuron

- (A) Weights of a model MT neuron as a function of preferred direction of V1 neurons.
- (B) V1 population response (left) and response of the MT neuron (right, filled circle) to a single grating of 50% contrast drifting in a direction of 180° . Gratings of different directions of motion will elicit shifted versions of the V1 population response (data not shown), which combine into the MT tuning curve (black line).
- (C) Responses to a plaid of equal component contrast (component directions are 120° and 240° , plaid motion direction is 180°). V1 population response (left), with normalization (black) and without normalization (gray). Response of the MT neuron to the plaid (right, filled circle). Plaids of different global motion directions will elicit shifted versions of the V1 population response (data not shown), which are combined into the MT tuning for plaids (black trace: with V1 normalization; gray trace: without V1 normalization). The MT neuron responds most strongly when the global motion direction of the plaid is 180° .
- (D) Responses to the same plaid, but with unequal component contrast. Same format as in (C). With normalization in V1, the MT neuron's tuning for plaid direction shifts toward 240° (i.e., toward the plaid) with directions of 180° (high contrast) and 300° (low contrast).

of the responses to the individual stimulus components. The weights applied to the stimulus components depend on the specific combination of stimulus strengths and lead to a range of behaviors from equal summation to winner-take-all competition.

Our results agree qualitatively but not precisely with a recent study that has explored population responses to a subset of our stimuli—namely, plaids whose components have identical contrast (MacEvoy et al., 2009). This study reported that population response to such plaids resembled the average of the responses to the component gratings, independent of overall

contrast. This behavior corresponds to equal-weights summation with weights of 0.5. Similar to their results, we find that the weights for same-contrast plaids do not depend on overall contrast (ANOVA, $F = 0.61$; $df = 3$; $p = 0.5$). However, we consistently found weights higher than 0.5 (average of 0.67 ± 0.14 SD; $n = 9$). The discrepancy may be due to differences in recorded signals (spikes versus blood related signals) and in species (cat versus tree shrew).

Our study, in fact, extends this previous finding by providing a more direct measure of neuronal responses, by probing the responses to plaids in which component contrasts differ, by showing that they obey a weighted-sum law across all contrast combinations, and by providing a unifying normalization model that accounts for the responses. The model has few free parameters, and exhibits a gradual transition between regimes of equal summation and winner-take-all competition that depends on the stimulus contrasts.

By applying and validating a normalization model to populations, moreover, we substantially extend prior studies of normalization in individual neurons. The normalization model has been shown to predict the responses of single neurons to stimuli that were tailored to the recorded neuron (Carandini et al., 1997; Freeman et al., 2002; Heeger, 1992). These prior results, however, concern only a subset of the neurons in a population: those that are stimulated optimally by either component of the plaids. Such neurons inhabit two of the 12 orientation preference bins in which we have divided the population. It was not known whether their collective action could be explained by a single normalization model. Moreover, the majority of neurons—those in the remaining 10 bins—are driven optimally by neither grating. It was unknown whether the normalization model would have predicted their responses.

While performing the apparently mundane task of contrast gain control, normalization provides a single mechanism that can gradually turn from equal summation to winner-take-all competition. Work in single neurons had well established that normalization could capture the effects of cross-orientation suppression and thus predict sublinear summation of responses to the individual components. As revealed by our measurements, however, the normalization model exhibits an unintuitive behavior: a strong population response to a single stimulus can be dramatically suppressed by adding a stimulus of higher contrast, so much that the population response to the compound stimulus seems to represent only the stimulus with higher contrast.

The ability of the normalization model to capture population activity in response regimes ranging from equal summation to winner-take-all competition resonates with a recent proposal of a canonical divisive computation (Kouh and Poggio, 2008). Although the mathematical expressions for the two divisive computations are similar, there are also important differences. First, the normalization model acts on the representation of a stimulus in the population activity; this representation is weighted by stimulus strength, not by synaptic strength. Second, the normalization model accounts for all responses with a single set of parameters, not by postulating different circuits with different sets of parameters. Hence, the normalization model predicts that the same neural circuitry can operate across all stimulus conditions.

The neural circuitry underlying response normalization is far from clear and may involve a combination of synaptic inhibition (Carandini et al., 1997), modulation of input noise (Chance et al., 2002; Finn et al., 2007), and nonlinearities in the input from the LGN (Carandini et al., 1997; Freeman et al., 2002; Priebe and Ferster, 2006). Our results of profound suppression of LFP responses in the presence of multiple stimuli indicate that spike threshold in area V1 may not play a major role in response normalization. This result certainly agrees with measurements of cross-orientation suppression in the membrane potential of V1 simple cells (Priebe and Ferster, 2006).

Similar to LFPs in anesthetized cats, normalization also affects population responses recorded from human V1. Our data are in general agreement with earlier measurements of cross-orientation suppression in the human VEP (Candy et al., 2001; Morrone and Burr, 1986). We extend the results of these studies in three ways. First, by using source-localization combined with fMRI retinotopic mapping, we concentrate on responses of area V1 rather than the whole occipital cortex. Second, we relate directly these population responses in humans to measurements of LFP obtained in cats with similar stimulation methods. Third, we demonstrate that both LFP and source-imaged VEP responses can be predicted quantitatively by a normalization model with similar parameter values.

The normalization observed in the human primary visual cortex should have behavioral consequences during the perception of superimposed orientations: perceptual thresholds for the test stimulus should be elevated in the presence of a mask. Such threshold elevations have indeed been found in a number of psychophysical experiments and are well characterized (Petrov et al., 2005). Also in agreement with our findings, psychophysical experiments exploiting reflexive ocular following movements suggest an intriguing combination of summation and winner-take-all mechanisms depending on the relative contrast of competing image motions (Sheliga et al., 2006).

With simple simulations, we show that normalization in the population responses of V1 can have profound effects on neurons in higher visual areas. Strikingly, normalization in V1 makes our model MT neuron lose its pattern-selectivity when the two components of a plaid have dissimilar contrast. This prediction may not be far from truth. First, it is consistent with a recent abstract reporting exactly such behavior in actual MT neurons (Kumbhani et al., 2008). Second, it is consistent with similar behaviors seen in the responses of MT neurons to multiple stimuli at different positions in their receptive field, which exhibit summation for stimuli with similar contrast and competition for stimuli with different contrast (Britten and Heuer, 1999; Heuer and Britten, 2002). Third, our prediction is consistent with the tendency of human observers to perceive plaids of different component contrast as moving in the direction of the component of higher contrast (Stone et al., 1990).

The normalization model for population responses could serve as a front-end of models of attentional selection. Attention enhances the processing of behaviorally relevant information and reduces the impact of irrelevant information; effects of attention have been likened to changes in effective stimulus contrast (Martinez-Trujillo and Treue, 2002; Reynolds et al., 2000). In our experiments, which do not engage attentional modulation, the

sensory representation of multiple stimuli overemphasizes the more salient object. Attention could act on top of these purely sensory effects: spatial attention could further accentuate the existing asymmetry of the sensory representation (Ghose and Maunsell, 2008; Moran and Desimone, 1985; Reynolds et al., 1999; Treue and Martinez-Trujillo, 1999), and feature-based attention could additionally sharpen the population response to the attended feature (Martinez-Trujillo and Treue, 2004). In line with these ideas, the effects of attention on psychophysical performance have been modeled as competitive interactions between visual filters (Lee et al., 1999). Furthermore, it has recently been proposed that attention might use existing normalization circuits to modulate neuronal responses to sensory stimuli (Reynolds and Heeger, 2009).

Normalization, in summary, makes population responses operate not only in a regime of equal summation, where the response to summed stimuli resembles the scaled sum of the responses, but also in a regime of winner-take-all competition, where the response to summed stimuli resembles the response to the stronger stimulus alone. Effects of this kind have been observed in single neurons of regions as diverse as cortical area MT (Britten and Heuer, 1999; Heuer and Britten, 2002) and inferior colliculus (Keller and Takahashi, 2005). They may be due to normalization operating in the population responses in those regions or in earlier stages. These results suggest that normalization may be a fundamental operation, one that shapes population responses more profoundly than might have been expected from earlier studies.

EXPERIMENTAL PROCEDURES

MUA and LFP Recordings

Experiments were conducted at the Smith-Kettlewell Eye Research Institute under the supervision of the Institutional Animal Care and Use Committee. Detailed methods for these experiments are published elsewhere (Katzner et al., 2009). Briefly, young adult cats were anesthetized with Sodium Pentothal (0.5–2 mg/kg/hr, i.v.) and Fentanyl (typically 10 µg/kg/hr, i.v.), supplemented by inhalation of O₂ mixed with N₂O (typically in a ratio of 70:30). Eye-movements were prevented by a neuromuscular blocker (Pancuronium Bromide, 0.15 mg/kg/hr, i.v.). A 10 × 10 electrode array (400 µm spacing, 1.5 mm electrode length) was implanted in area V1 to record multiunit activity (MUA) and local field potentials (LFP). Insertion depths were about 0.8–1 mm, resulting in recordings confined mostly to layers 2–3. A Cerebus 128-channel system (Blackrock, Utah) was used to sample the data.

Stimuli for MUA Analysis

Stimuli were contrast-reversing sinusoidal gratings presented monocularly on a CRT monitor (refresh rate 125 Hz, mean luminance 32 cd/m²). Gratings were modulated sinusoidally in contrast with a temporal frequency of 4 Hz (9 experiments). Spatial frequency was adjusted to optimally drive the majority of sites in the electrode array. Gratings were presented in a circular window (30° diameter) and lasted 2 s. The contrast of the gratings was 0%, 6%, 12%, 25%, and 50%. Plaids were obtained by summing two gratings. For each experiment, the angle between component orientations in the plaid was fixed. To reduce effects of adaptation, each experiment consisted of three pairs of orientations (e.g., 0°/90°, 30°/120°, 60°/150°). Data for different pairs were collapsed after adjusting for the difference between stimulus orientation and neuronal preference. In addition, we recorded responses to 100% contrast gratings of 12 different orientations to obtain orientation tuning curves for each site in the array. The stimuli were shown in random order in blocks presented at least eight times.

Analysis of MUA

In each channel of the multielectrode array, we set thresholds to ~4 SD of the background noise. Threshold crossing were considered multiunit activity and were pooled together for each site. For each site, orientation preference was defined as the vector average of firing rate responses to the 12 full-contrast tuning stimuli. Only responsive sites with at least minimal tuning (circular variance <0.85) were considered in all further analyses (Ringach et al., 2002). Using this criterion, on average 42 ± 13 sites (median ± MAD) were included. For the computation of population responses, sites were binned (15° bin width) according to their preferred orientation. Before averaging across sites, responses for each site were normalized to their average response to the tuning stimuli.

Stimuli for LFP Analysis

The stimuli used in the LFP experiments were identical to the stimuli used in the MUA experiments with the following exception: the two components were modulated with a temporal frequency of 3.5 and 5 Hz. Data were obtained in four experiments.

Analysis of LFPs

LFPs were sampled at 2 kHz with a wide front-end filter (0.3 Hz to 500 Hz). We further low-pass filtered the LFP below 90 Hz to exclude any contamination by multiunit responses. LFPs were averaged across all electrodes in the Utah array that were responsive (circular variance of multiunit responses to 12 different orientations <0.85). Amplitudes at twice the stimulus frequency were extracted from the averaged data using Fourier analysis.

Single Unit Recordings

The analysis in Figure 2 was performed on 75 well-isolated neurons recorded in previous experiments in nine cats. Details on stimuli and preparation are described elsewhere (Freeman et al., 2002).

EEG Recordings

Methods for VEP recordings and source localization are described elsewhere (Appelbaum et al., 2006). Briefly, the EEG was acquired with 128-channel HydroCell Sensor Nets and a Netstation digitization system (EGI, Eugene OR) in five subjects. For each subject, the 3D locations of the sensors were recorded using a “Fastrak” radio-frequency 3D digitizer (Polhemus, Colchester, VT) and coregistered to their T1-weighted anatomical magnetic resonance (MR) scans, from which a three-shell boundary element model of the skull and scalp was computed.

Stimuli for EEG Analysis

Stimuli were presented foveally on a CRT monitor (refresh rate 72 Hz, mean luminance 32 cd/m²). Stimuli were viewed binocularly and consisted of contrast-reversing gratings, modulated at temporal frequencies of 4.5 Hz (test) and 3.6 Hz (mask). Gratings were presented within a circular window spanning 5° of visual angle, and their spatial frequency was 1 cpd. The test gratings had 0%, 6%, 12%, 25%, or 50% contrast; the mask gratings had 0 or 25% contrast. Each combination of grating contrasts was presented at least 15 times. Each trial lasted 11.1 s. During the trial, subjects attended to a stream of simultaneously presented letters and were instructed to indicate the presence of a probe letter “T” among distracters “L.”

Analysis of EEG Signals

EEG signals were postprocessed using custom software to remove artifacts due to head motion and blinks. To localize the sources of the VEP activity, cortically constrained minimum norm source estimates (Hämäläinen and Ilmoniemi, 1994) were computed and related to each subject’s visual areas as defined by fMRI retinotopic mapping (Engel et al., 1997). We extracted the time course of activity from area V1 and obtained the responses to test and mask grating by Fourier analysis.

Modeling

The models used to fit MUA are given by Equations (1)–(4) in the main text. Before fitting the models, we subtracted spontaneous activity to the gray

screen from all responses. The model used to fit LFP and VEP responses is given by Equation (5). All models were fitted by a least-squares algorithm.

Fit quality was assessed by one of two measures. The first measure is the percentage of variance in the responses r explained by the model predictions m :

$$v = 1 - \frac{\sum_{i=1}^n (r_i - m_i)^2}{\sum_{i=1}^n (r_i - \bar{r})^2},$$

where the indices i indicate the bin of orientation preference and \bar{r} indicates the mean of the responses. The second measure (fit quality index) is the root mean square deviation between responses and model, normalized to the mean of the observed responses:

$$q = 1 - \frac{\sqrt{\frac{\sum_{i=1}^n (r_i - m_i)^2}{n}}}{\bar{r}}.$$

Simulation of MT Responses

A population of direction-selective V1 neurons was modeled using the average parameters obtained from fitting our population recordings with the normalization model (Equation 4, width, semisaturation contrast, and exponent), and assuming an equal spacing of preferred direction across neurons. The weighting profile of the MT neuron is given by a Gaussian with a sigma of 65° (inspired by the neuron in Figure 5E of Rust et al., 2006). V1 population responses to plaids of 120° angle were simulated using the normalization model (Equation 4) and multiplied by the MT weighting profile. Responses below zero were set to zero.

SUPPLEMENTAL DATA

Supplemental data includes an Appendix and one figure and may be found with this article online at [http://www.cell.com/neuron/supplemental/S0896-6273\(09\)00886-1](http://www.cell.com/neuron/supplemental/S0896-6273(09)00886-1).

ACKNOWLEDGMENTS

We are grateful to S. Katzner and A. Benucci for help with the experiments in cat, J. Rowland for help with the experiments in humans, and D. Fitzpatrick and S. MacEvoy for helpful discussions. This work was supported by National Institutes of Health (grants EY-17396 to M.C. and EY-18157 to A.R.W.), by a Project Grant from the UK charity Fight for Sight (to M.C.), by an Advanced Investigator Award to M.C. from the European Research Council, and by a postdoctoral fellowship of the German Academy of Sciences Leopoldina (BMBF-LPD9901/8-165 to L.B.). M.C. holds the GlaxoSmithKline/Fight for Sight Chair in Visual Neuroscience.

Accepted: October 28, 2009

Published: December 23, 2009

REFERENCES

Albrecht, D.G., and Geisler, W.S. (1991). Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Vis. Neurosci.* 7, 531–546.

Albrecht, D.G., and Hamilton, D.B. (1982). Striate cortex of monkey and cat: contrast response function. *J. Neurophysiol.* 48, 217–237.

Appelbaum, L.G., Wade, A.R., Vildavski, V.Y., Pettet, M.W., and Norcia, A.M. (2006). Cue-invariant networks for figure and background processing in human visual cortex. *J. Neurosci.* 26, 11695–11708.

Britten, K.H., and Heuer, H.W. (1999). Spatial summation in the receptive fields of MT neurons. *J. Neurosci.* 19, 5074–5084.

Burr, D.C., and Morrone, M.C. (1987). Inhibitory interactions in the human vision system revealed in pattern-evoked potentials. *J. Physiol.* 389, 1–21.

Candy, T.R., Skoczenski, A.M., and Norcia, A.M. (2001). Normalization models applied to orientation masking in the human infant. *J. Neurosci.* 21, 4530–4541.

Carandini, M., Heeger, D.J., and Movshon, J.A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.* 17, 8621–8644.

Carandini, M., Heeger, D.J., and Senn, W. (2002). A synaptic explanation of suppression in visual cortex. *J. Neurosci.* 22, 10053–10065.

Carandini, M., and Sengpiel, F. (2004). Contrast invariance of functional maps in cat primary visual cortex. *J. Vis.* 4, 130–143.

Chance, F.S., Abbott, L.F., and Reyes, A.D. (2002). Gain modulation from background synaptic input. *Neuron* 35, 773–782.

Chen, Y., Geisler, W.S., and Seidemann, E. (2006). Optimal decoding of correlated neural population responses in the primate visual cortex. *Nat. Neurosci.* 9, 1412–1420.

DeAngelis, G.C., Robson, J.G., Ohzawa, I., and Freeman, R.D. (1992). Organization of suppression in receptive fields of neurons in cat visual cortex. *J. Neurophysiol.* 68, 144–163.

Engel, S.A., Glover, G.H., and Wandell, B.A. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb. Cortex* 7, 181–192.

Finn, I.M., Priebe, N.J., and Ferster, D. (2007). The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. *Neuron* 54, 137–152.

Freeman, T.C.B., Durand, S., Kiper, D.C., and Carandini, M. (2002). Suppression without inhibition in visual cortex. *Neuron* 35, 759–771.

Georgopoulos, A.P., Kalaska, J.F., Caminiti, R., and Massey, J.T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J. Neurosci.* 2, 1527–1537.

Georgopoulos, A.P., Schwartz, A.B., and Kettner, R.E. (1986). Neuronal population coding of movement direction. *Science* 233, 1416–1419.

Ghose, G.M., and Maunsell, J.H.R. (2008). Spatial summation can explain the attentional modulation of neuronal responses to multiple stimuli in area V4. *J. Neurosci.* 28, 5115–5126.

Hämäläinen, M.S., and Ilmoniemi, R.J. (1994). Interpreting magnetic fields of the brain: minimum norm estimates. *Med. Biol. Eng. Comput.* 32, 35–42.

Heeger, D.J. (1991). Nonlinear model of neural responses in cat visual cortex. In *Computational Models of Visual Processing*, M.S. Landy and J.A. Movshon, eds. (Cambridge, MA: MIT Press), pp. 119–133.

Heeger, D.J. (1992). Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* 9, 181–197.

Heuer, H.W., and Britten, K.H. (2002). Contrast dependence of response normalization in area MT of the rhesus macaque. *J. Neurophysiol.* 88, 3398–3408.

Katzner, S., Nauhaus, I., Benucci, A., Bonin, V., Ringach, D.L., and Carandini, M. (2009). Local origin of field potentials in visual cortex. *Neuron* 61, 35–41.

Keller, C.H., and Takahashi, T.T. (2005). Localization and identification of concurrent sounds in the owl's auditory space map. *J. Neurosci.* 25, 10446–10461.

Kouh, M., and Poggio, T. (2008). A canonical neural circuit for cortical nonlinear operations. *Neural Comput.* 20, 1427–1451.

Kumbhani, R.D., Saber, G.T., Majaj, N.J., Tailby, C., and Movshon, J.A. (2008). Contrast affects pattern direction selectivity in macaque MT neurons. (Program No. 460.26. 2008 Neuroscience Meeting Planner. Washington, DC: Society for Neuroscience, 2008. Online.)

Lee, D.K., Itti, L., Koch, C., and Braun, J. (1999). Attention activates winner-take-all competition among visual filters. *Nat. Neurosci.* 2, 375–381.

MacEvoy, S.P., Tucker, T.R., and Fitzpatrick, D. (2009). A precise form of divisive suppression supports population coding in the primary visual cortex. *Nat. Neurosci.* 12, 637–645.

Martinez-Trujillo, J., and Treue, S. (2002). Attentional modulation strength in cortical area MT depends on stimulus contrast. *Neuron* 35, 365–370.

- Martinez-Trujillo, J.C., and Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr. Biol.* *14*, 744–751.
- McIlwain, J.T. (1986). Point images in the visual system: new interest in an old idea. *Trends Neurosci.* *9*, 354–358.
- Moran, J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science* *229*, 782–784.
- Morrone, M.C., and Burr, D.C. (1986). Evidence for the existence and development of visual inhibition in humans. *Nature* *321*, 235–237.
- Morrone, M.C., Burr, D.C., and Maffei, L. (1982). Functional implications of cross-orientation inhibition of cortical visual cells. I. Neurophysiological evidence. *Proc. R. Soc. Lond. B. Biol. Sci.* *276*, 335–354.
- Movshon, J.A., Adelson, E.H., Gizzi, M.S., and Newsome, W.T. (1985). The analysis of moving visual patterns. In *Study Group of Pattern Recognition Mechanisms*, C. Chagas, R. Gattass, and C. Gross, eds. (Vatican City: Pontificia Academia Scientiarum), pp. 117–151.
- Nicolelis, M.A., Baccala, L.A., Lin, R.C., and Chapin, J.K. (1995). Sensorimotor encoding by synchronous neural ensemble activity at multiple levels of the somatosensory system. *Science* *268*, 1353–1358.
- Pasupathy, A., and Connor, C.E. (2002). Population coding of shape in area V4. *Nat. Neurosci.* *5*, 1332–1338.
- Petrov, Y., Carandini, M., and McKee, S. (2005). Two distinct mechanisms of suppression in human vision. *J. Neurosci.* *25*, 8704–8707.
- Pouget, A., Dayan, P., and Zemel, R. (2000). Information processing with population codes. *Nat. Rev. Neurosci.* *1*, 125–132.
- Priebe, N.J., and Ferster, D. (2006). Mechanisms underlying cross-orientation suppression in cat visual cortex. *Nat. Neurosci.* *9*, 552–561.
- Purushothaman, G., and Bradley, D.C. (2005). Neural population code for fine perceptual decisions in area MT. *Nat. Neurosci.* *8*, 99–106.
- Regan, D. (1989). *Human Brain Electrophysiology: Evoked Potentials and Evoked Magnetic Fields in Science and Medicine* (New York, NJ: Elsevier), 70–112.
- Reynolds, J.H., Chelazzi, L., and Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *J. Neurosci.* *19*, 1736–1753.
- Reynolds, J.H., and Heeger, D.J. (2009). The normalization model of attention. *Neuron* *61*, 168–185.
- Reynolds, J.H., Pasternak, T., and Desimone, R. (2000). Attention increases sensitivity of V4 neurons. *Neuron* *26*, 703–714.
- Ringach, D.L., Shapley, R.M., and Hawken, M.J. (2002). Orientation selectivity in macaque V1: diversity and laminar dependence. *J. Neurosci.* *22*, 5639–5651.
- Ross, J., and Speed, H.D. (1991). Contrast adaptation and contrast masking in human vision. *Proc. Biol. Sci.* *246*, 61–69.
- Rust, N.C., Mante, V., Simoncelli, E.P., and Movshon, J.A. (2006). How MT cells analyze the motion of visual patterns. *Nat. Neurosci.* *9*, 1421–1431.
- Schroeder, C.E., Tenke, C.E., Givre, S.J., Arezzo, J.C., and Vaughan, H.G., Jr. (1991). Striate cortical contribution to the surface-recorded pattern-reversal VEP in the alert monkey. *Vision Res.* *31*, 1143–1157.
- Scial, G., and Freeman, R.D. (1982). Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Exp. Brain Res.* *46*, 457–461.
- Sheliga, B.M., Kodaka, Y., FitzGibbon, E.J., and Miles, F.A. (2006). Human ocular following initiated by competing image motions: Evidence for a winner-take-all mechanism. *Vision Res.* *46*, 2041–2060.
- Simoncelli, E.P., and Heeger, D.J. (1998). A model of neuronal responses in visual area MT. *Vision Res.* *38*, 743–761.
- Stone, L.S., Watson, A.B., and Mulligan, J.B. (1990). Effect of contrast on the perceived direction of a moving plaid. *Vision Res.* *30*, 1049–1067.
- Treue, S., Hol, K., and Rauber, H.J. (2000). Seeing multiple directions of motion—physiology and psychophysics. *Nat. Neurosci.* *3*, 270–276.
- Treue, S., and Martinez-Trujillo, J.C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* *399*, 575–579.