The Coordinated Mapping of Visual Space and Response Features in Visual Cortex

Hongbo Yu,[1,3] Brandon J. Farley,[1,3] Dezhe Z. Jin,[2] and Mriganka Sur[1,*]
[1]Department of Brain and Cognitive Sciences
Picower Institute for Learning and Memory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139
[2]Department of Physics
The Pennsylvania State University
University Park, Pennsylvania 16802

Summary

Whether general principles can explain the layouts of cortical maps remains unresolved. In primary visual cortex of ferret, the relationships between the maps of visual space and response features are predicted by a "dimension-reduction" model. The representation of visual space is anisotropic, with the elevation and azimuth axes having different magnification. This anisotropy is reflected in the orientation, ocular dominance, and spatial frequency domains, which are elongated such that their directions of rapid change, or high-gradient axes, are orthogonal to the high-gradient axis of the visual map. The feature maps are also strongly interdependent—their high-gradient regions avoid one another and intersect orthogonally where essential, so that overlap is minimized. Our results demonstrate a clear influence of the visual map on each feature map. In turn, the local representation of visual space is smooth, as predicted when many features are mapped within a cortical area.

Introduction

One of the keys to understanding cortical function is defining the rules that govern the organization of neurons according to their response properties. As some cortical areas are traversed, neuronal response properties change in a regular manner, forming systematic maps with highly stereotyped structures. Additionally, more than one response property can be mapped simultaneously within an area, as has been shown in visual, somatosensory, and auditory cortex (Hubel and Wiesel, 1963; Sur et al., 1981; Linden and Schreiner, 2003; Friedman et al., 2004). What determines the forms that these maps take? One theoretical proposal is that the structures reflect constraints imposed by "dimension-reduction," or a need to smoothly map several (more than two) response properties onto a two-dimensional cortical surface (Kohonen, 1982a; Durbin and Mitchison, 1990; Obermayer et al., 1990; Swindale, 1991). An explicit goal of dimension-reduction models is to minimize the connection length between neurons that share similar response properties, an approach also used in "elastic net" models (Goodhill and Will-

shaw, 1990) and "wire-length minimization" algorithms (Koulakov and Chklovskii, 2001) for explaining the structure of cortical columns.

Primary visual cortex (V1) is a model area for examining this issue since neurons in the region are mapped according to several well-understood response properties (i.e., features), including receptive field position, ocular dominance, orientation, and spatial frequency. The columnar organization of the area implies that the multiple feature maps are superimposed over a single two-dimensional cortical sheet. To accommodate the representation of all the features, dimension-reduction models predict specific spatial relationships between feature maps (Durbin and Mitchison, 1990; Obermayer et al., 1990). However, recent studies have questioned the relevance of such principles to cortical maps, in particular the map of visual space (Das and Gilbert, 1997; Bosking et al., 2002; Buzas et al., 2003).

A major prediction of dimension-reduction models is that, at locations where the value of a given feature changes rapidly across cortex, that of other independent features should change slowly; that is, there is a negative correlation between the gradient magnitudes of different maps (Obermayer et al., 1992). A corollary to this principle is that, at locations where the gradient magnitudes of two features are necessarily high, the contours of the two maps, and hence their gradient directions, are orthogonal. Experimental studies have demonstrated relationships for some pairs of functional maps consistent with these predictions. For example, locations where the preferred orientation map gradient is high, i.e., orientation pinwheels, are situated toward the center of ocular dominance and spatial frequency columns, where the latter maps have low gradients. Additionally, ocular dominance and orientation contours intersect near-orthogonally at the edges of ocular dominance columns, where the ocular dominance gradient is high (Blasdel and Salama, 1986; Bartfeld and Grinvald, 1992; Obermayer and Blasdel, 1993; Hubener et al., 1997).

In contrast, results from experiments that have examined the retinotopic and orientation maps appear inconsistent with the model predictions: whereas the original models (Durbin and Mitchison, 1990; Obermayer et al., 1990) predicted a negative correlation between the gradients of these two maps, such that receptive field locations should change slowly in locations where orientation preferences change rapidly, experimental results show either a positive correlation (Das and Gilbert, 1997) or no correlation (Bosking et al., 2002; Buzas et al., 2003). But the original models included only retinotopy and orientation as mapped features; more recent studies have demonstrated that the predicted relationships between the retinotopic and other maps are dependent on the number of features that are included in the model simulations (Swindale, 2004). It remains to be examined how experimentally measured relationships compare to model predictions, when the model simulations include all of the features known to be mapped within the cortical region of question.

*Correspondence: msur@mit.edu
[3]These authors contributed equally to this work.

We reasoned that an examination of the relationships between the retinotopic and other maps in V1 should address three related questions: (1) does the overall structure of the retinotopic map influence the mapping of other features, in a manner consistent with dimension-reduction principles; (2) are other feature maps related to each other by these principles, and (3) is the mapping of other features in turn reflected in the local structure of the retinotopic map? Since the relationships between the retinotopic and other maps have proven elusive, we examined a case where the dimension-reduction model makes a clear prediction: when the retinotopic map magnification differs for the two axes of visual space, the model predicts clear consequences for the remaining maps. The examination of maps in V1 of ferrets provided a unique opportunity to test this prediction. We found that, indeed, the retinotopic map powerfully influences the layout of other maps, as the model predicts, and that the other maps maintain specific local relationships with one another. Finally, the retinotopic map is predicted to be locally smooth in the case where three other variables are mapped within a cortical area. Our results demonstrate a strong interdependence between each visual cortex functional map, including the retinotopic map.

Results

The Dimension-Reduction Model

To model how several features could be mapped across the two-dimensional cortical surface, we used the self-organizing feature map algorithm (Kohonen, 1982b; Obermayer et al., 1992). In this algorithm, each possible feature combination within a multidimensional response feature space (representing x and y axis retinotopic location, orientation, ocular dominance, and spatial frequency) is represented at a location in cortex. In other words, the cortex achieves complete "coverage" of the feature space. In addition, the algorithm has a mechanism to promote "continuity" in the cortical representations; i.e., nearby locations on the cortical surface are forced to represent similar features so that orderly feature maps are formed. The constraints of coverage and continuity act in opposition and are commonly employed in some form in many computational models of visual cortex (Erwin et al., 1995; Swindale, 1996). We asked what form the feature mapping was predicted to take given these constraints and, in particular, how a global anisotropy in the retinotopic map was predicted to influence the layouts of the other maps. We reasoned that a match between the simulation predictions and in vivo experimental observations would suggest that the mapping of visual space and other features occurs in conformity with a dimension-reduction strategy.

Predicted Influence of the Retinotopic Map on the Structures of Other Maps

To determine the predicted influence of retinotopy on the mapping of other features, we compared the results of simulations in which visual space was represented either anisotropically or isotropically. To model an anisotropy, we used a retinotopic map whose magnifica-

tion was four times greater along the elevation axis of visual space representation than along the azimuth axis (Figure 1A). This degree of anisotropy was chosen because it resembles that found in ferret V1 (see Figure 3). We observed a strong influence of the visual map anisotropy on the layouts of the remaining maps formed in the simulation: domains in the orientation (Figure 1A), ocular dominance (Figure 1B), and spatial frequency (see Figure S1A in the Supplemental Data available online) maps were highly elongated. Furthermore, the domains from the latter maps were elongated along a common axis; this axis was parallel to the azimuth (low-magnification) axis of the retinotopic map and orthogonal to the elevation (high-magnification) axis. It follows from these observations that the high-gradient axis of retinotopy is orthogonal to the high-gradient axis of the other maps.

To quantify this observation, we calculated the gradient vector at every pixel of each functional map (see Experimental Procedures). The direction of this vector indicates the axis along which the mapped variable changes maximally. We found a bias toward orthogonal intersection angles between the visual map gradient and the gradients of either the orientation, ocular dominance, or spatial frequency maps (Figure 1C, dotted lines). Further, in those locations where the latter maps had a high-gradient magnitude, there was an even stronger tendency for their gradient direction to be orthogonal to that of the visual map (Figure 1C, solid lines).

In contrast to the above results, when an isotropic visual map was used in the simulation, domains in the resulting orientation (Figure 1D), ocular dominance (Figure 1E), and spatial frequency (Figure S1E) maps were not elongated. As a result, these domains did not align along a particular axis of the retinotopic map (Figure 1F). Taking these results together, our model predicts that a global anisotropy in visual space representation will strongly and specifically influence the layouts of other maps.

Predicted Gradient Relationships between Orientation, Ocular Dominance, and Spatial Frequency Maps

We next examined the relationships between the model's orientation, ocular dominance, and spatial frequency maps. A fundamental prediction of dimension-reduction models, which has previously been described, is that the highest gradient regions of multiple maps should avoid one another. We asked whether this could still occur with a strong anisotropy in the retinotopic map. Orientation, ocular dominance, and spatial frequency gradient maps from the same region of the model cortex were obtained. We found that the anisotropy caused the highest-gradient regions of these three maps to align along a common axis of the cortex (Figure 2A; cf. Figure 2F). Nevertheless, the highest-gradient regions were precisely interleaved and appeared to occupy mostly nonoverlapping regions of cortex (Figure 2A).

To quantify this relationship, all pixels within each gradient map were assigned a value between 1 and 10, where 1 indicated that the gradient was in the lowest 10th percentile of all gradient values within that map,
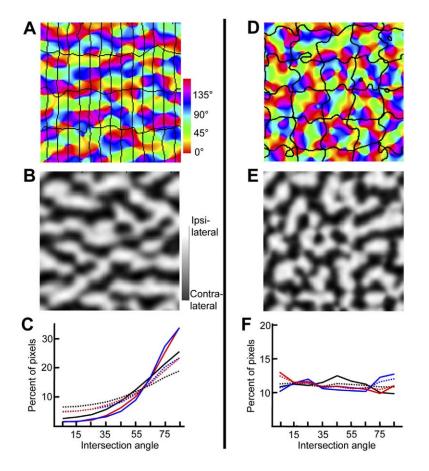
Figure 1. Effect of Visual Space Anisotropy on the Layouts of Other Feature Maps, as Predicted by a Dimension-Reduction Model

The retinotopic map was either anisotropic (A–C) or isotropic (D–F); it is illustrated in (A) and (D) by iso-elevation (horizontal in figure) and iso-azimuth (vertical in figure) contours; all contour intervals are identical. (A and D) The orientation map, denoted by color coding each pixel according to its preferred orientation. (B and E) The ocular dominance map. (C and F) Percent of pixels that have an intersection angle, within each 10° range whose center is indicated by the x axis tick mark, between the retinotopic gradient and the gradient of the orientation (black lines), ocular dominance (blue lines), or spatial frequency (red lines) map. Dotted lines indicate the calculation over all pixels of the cortex. Solid lines indicate the calculation over only those pixels whose orientation, ocular dominance, or spatial frequency gradients were within the highest 30th percentile.

and 10 the highest 10th percentile. The values from all three maps were averaged together at each pixel in the region, and the standard deviation of this average over the entire cortex was measured. We found that the standard deviation was smaller if the actual gradient maps were used (Figure 2B, red dotted bar) in the calculation, compared to when the pixels in each gradient map were randomly shuffled (Figure 2B, blue histogram). Thus, there is a tendency for the average gradient magnitude of all three maps to remain constant across the cortical surface. This demonstrates that the highest-gradient regions of the three maps mutually avoid one another.

To analyze directly whether the high-gradient regions of two maps avoid one another, we compared their gradients at corresponding pixels. Pixels were binned into ten groups according to their ocular dominance gradient percentile, and the average orientation gradient over all pixels within each group was determined and plotted (Figure 2C). A negative correlation can be seen, suggesting that the highest gradient regions of these two maps avoid one another. The same relationship exists between the orientation and spatial frequency maps (Figure S1B).

These gradient correlations among maps of orientation, ocular dominance, and spatial frequency also held when the initial retinotopic map was isotropic (Figures 2F–2H and Figure S1F). Thus, a fundamental prediction of dimension-reduction models is that high-gradient re-

gions of multiple maps avoid one another, and it holds whether the retinotopic map is anisotropic or isotropic.

Predicted Contour Relationships between Orientation, Ocular Dominance, and Spatial Frequency Maps

Above we described a fundamental prediction of dimension-reduction models, that high-gradient regions of multiple maps avoid one another. A second major prediction of this class of models, which is common to many cortical mapping algorithms, is that contours of different maps should cross at perpendicular angles (Erwin et al., 1995; Swindale, 1996). However, the constraint imposed by the anisotropy in the visual map challenges this prediction. With an anisotropic visual map, the orientation, ocular dominance, and spatial frequency domains are elongated along a parallel axis of cortex (see Figure 2A). This would seem to preclude the contours of these maps from crossing at perpendicular angles.

To examine this issue, we plotted the distribution of the intersection angles, over all map pixels, between the orientation and ocular dominance gradient vectors. When the visual map was isotropic, the intersection angles were orthogonal on average, as expected from previous studies (Figure 2J, red line). In contrast, when an anisotropic visual map was used, the intersection angles had no bias toward orthogonality (Figure 2E, red
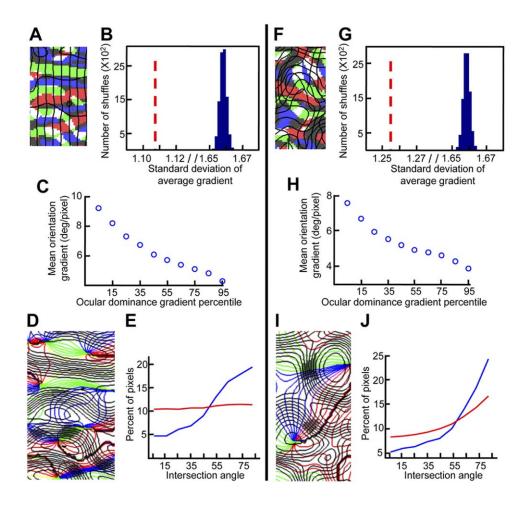
Figure 2. Predicted Effect of Visual Space Map Anisotropy on the Relationships between Other Feature Maps

The retinotopic map was either anisotropic (A–E) or isotropic (F–J). (A and F) High-gradient regions (pixels with gradients within the top 30th percentile) of orientation (blue), ocular dominance (green), spatial frequency (red), or two or more maps (white) are indicated. (B and G) Standard deviation across cortex of the average gradient magnitude of the orientation, ocular dominance, and spatial frequency maps (red dotted line). Histogram of this standard deviation calculation after pixels within each map were randomly shuffled, in 10,000 cases (blue histogram). (C and H) Pixels are grouped into ten bins according to their ocular dominance gradient percentile, and the mean orientation gradient for each group is indicated. (D and I) Orientation (color) contours superimposed on ocular dominance (black) contours. Gray pixels indicate locations where the gradient of both the orientation and ocular dominance maps are within the highest 30th percentile. (E and J) Percent of pixels that have an intersection angle, within each 10° range whose center is indicated, between the orientation and ocular dominance gradients. Red line shows the calculation over all pixels. Blue line indicates the calculation for those pixels where the gradients of both maps are within the highest 30th percentile.

line). Thus, an anisotropy in the visual map precludes an overall orthogonality between the remaining maps.

We next show, however, that the remaining maps have specific local regions where orthogonality is maintained. In the anisotropic case, a clear bias toward orthogonal intersection angles between the orientation and ocular dominance gradient vectors emerged specifically in those locations where the maps of these features had overlapping high gradients (Figure 2D, gray regions; Figure 2E, blue line). For the isotropic case, the bias toward intersection angles between orientation and ocular dominance gradients similarly became stronger in those locations where the two maps have high gradients (Figures 2I and 2J). The relationships between the contours of orientation and spatial frequency were very similar to those between orientation and ocu-

lar dominance (Figures S1C, S1D, S1G, and S1H). The results suggest that the overall orthogonality between orientation and ocular dominance (or spatial frequency) maps can vary, based on the nature of the retinotopic map. But the model predicts that, regardless of the nature of the retinotopic map, orthogonality will consistently occur between the remaining maps in their high-gradient overlap regions.

Together, the relationships we have described reflect the model's strategy for smoothly accommodating several features simultaneously on the two-dimensional cortical surface. In particular, they predict what influence the retinotopic map should have on other maps if the mapping of visual space is interdependent with that of other features. Whether the visual space and other maps in visual cortex demonstrate these relationships
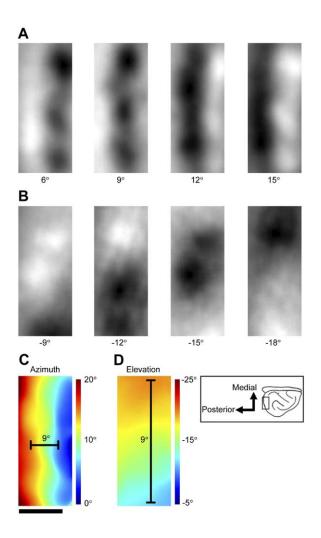
Figure 3. Optical Imaging Reveals that the Retinotopic Map in Ferret V1 Is Anisotropic

(Inset) Dorsal view of right hemisphere of ferret brain, indicating the orientation of the functional maps for all figures; outlined area indicates the region of cortex imaged. (A–B) "Single-condition" retinotopic maps. (A) Cortical activity pattern (activated pixels are darker) elicited by a flashing vertical bar. The stimulus position differed by 3° in horizontal space between successive panels, as indicated. (B) Activity elicited by a horizontal bar. The stimulus position differed by 3° in vertical space between successive panels. (C) Azimuth position preference map. Each pixel is color coded according to its preferred relative azimuth position in visual space. 9° of horizontal space was represented over 0.64 mm of cortex. (D) Elevation position preference map. 9° of vertical space was represented over 2.8 mm. Scale bar in (C) represents 1 mm and also applies to (A), (B), and (D).

requires experimental examination. We next compared these modeling predictions to the results we obtained experimentally in ferret V1.

Experimental Results from Ferret Visual Cortex: An Anisotropy in the V1 Retinotopic Map

We determined the layout of the retinotopic map in ferret V1 (Figure 3, inset) using optical imaging of intrinsic signals (see Experimental Procedures). The presenta-

tion of a vertical bar activated a narrow, mediolaterally elongated strip of cortex in V1, which represents an iso-azimuth contour (Figure 3A). As the vertical bar shifted from central to more peripheral regions of the visual field, the region of activation in V1 shifted posteriorly. A horizontal bar, on the other hand, activated a strip of cortex that was elongated along the anteroposterior axis of the cortex (Figure 3B). This strip represents an iso-elevation contour. Downward movement of the horizontal bar across the visual field caused the activated region to shift from lateral to medial across cortex (Figure 3B).

With our retinotopy stimulation paradigm, the stimulus passes over the receptive field of each cortical location periodically and causes a periodic activation pattern across time (data not shown) for each pixel. A pixel's relative receptive field location can be determined by its relative phase of activation within the stimulus cycle (Engel et al., 1994; Sereno et al., 1995; Kalatsky and Stryker, 2003; Mrsic-Flogel et al., 2003). We calculated the phase of the Fourier transform of each pixel's response time course, in response to stimulation by a horizontally or vertically drifting bar, and constructed maps for relative horizontal (Figure 3C) and vertical (Figure 3D) visual space, respectively.

While the general layout we show for the visual space map is in agreement with previous reports (Law et al., 1988; White et al., 1999), we additionally find a pronounced anisotropy in the representation of visual space (Figure 3). A much greater amount of cortex is devoted to the representation of a given amount of visual field elevation than for the equivalent amount of azimuth. In general, for the region of ferret V1 accessible to optical imaging, the magnification factor for the elevation axis of visual space was approximately three to five times greater than that for azimuth (Table S1).

The magnification anisotropy was confirmed by electrophysiological recordings. Receptive field positions and sizes were determined using the reverse correlation method of receptive field mapping. A representative case is shown in Figure 4. On- and off-subfields were summed to calculate the aggregate receptive field at each recording site (Figures 4A and 4B). The distribution of the receptive fields recorded at multiple sites (Figure 4A) revealed a wide range of azimuth position preferences (within a small anteroposterior extent of cortex) and a narrower range of elevations (Figure 4C). The composite retinotopic map confirmed that the magnification for the elevation axis of visual space was much greater than that for the azimuth axis (Figures 4D and 4E; note similarity to optical imaging-derived maps in Figures 3C and 3D).

Influence of the Global Visual Map Anisotropy on Feature Maps

Given the pronounced anisotropy of the retinotopic map in ferret visual cortex, we examined whether it would impact the layout of other functional maps, as the dimension-reduction model predicted. We first examined the map of orientation in V1, which was obtained by comparing the cortical activation pattern in response to grating stimuli presented at four different
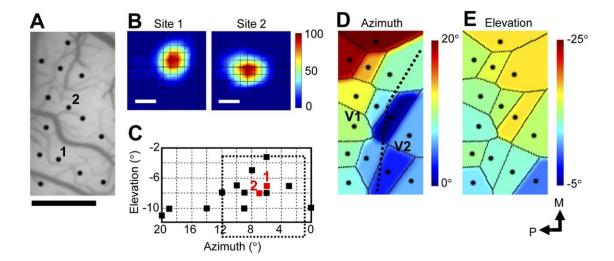
Figure 4. Multiunit Electrophysiological Recordings Confirm the Anisotropy of the Retinotopic Map in Ferret V1

(A) Image of cortical vasculature showing 14 recording sites. (B) Example multiunit receptive fields recorded at the two sites indicated in (A). The strength of the response to stimuli presented in each location of the stimulus grid (dotted box in [C]) is indicated (as percent maximum response). (C) Receptive field locations for all recording sites shown in (A). (D and E) Azimuth (D) and elevation (E) position preference map for the region of cortex shown in (A) (see Experimental Procedures). The dotted line in (D) indicates the approximate location of the V1/V2 border. Scale bar in (A) represents 1 mm and also applies to (D)–(E); bar in (B) is 3°.

orientations. Though this map has been previously described in ferret (Rao et al., 1997; White et al., 2001), a salient feature that we noticed was an anisotropy of the orientation domains. They were often elongated along the anteroposterior dimension of cortex (Figures 5A–5F). This was also true for the ocular dominance domains. The ocular dominance map in the binocular portion of V1, located at the most posterolateral edge of cortex (Law et al., 1988; Redies et al., 1990; White et al., 1999), contains ocular dominance columns that interleave similarly to those in V1 of the cat and macaque. In this part of V1, ocular dominance domains were anisotropic: they were elongated along the anteroposte-

rior dimension of cortex (Figure 5G), similarly to the iso-orientation domains. Finally, we examined whether spatial frequency is mapped in ferret visual cortex (see Experimental Procedures), and in what manner. We found that alternating high and low spatial frequency domains exist in V1 that are elongated along the anteroposterior dimension of cortex (Figure 5H and Figure S2A). The mapping of spatial frequency was verified in one animal by multiunit electrophysiological recordings (Figures S2B–S2D). Taken together, these results match those predicted by our model, suggesting that an anisotropy of the visual map is indeed accompanied by an anisotropy in the domains of the orientation, ocular
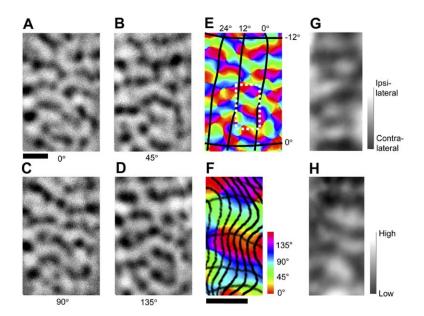

Figure 5. Orientation, Ocular Dominance, and Spatial Frequency Domains Are Anisotropic

(A–D) Cortical activity patterns elicited by the presentation of grating stimuli of the orientation indicated. (E) Orientation preference map with retinotopic contours superimposed. Vertical lines represent iso-azimuth contours, and horizontal lines are iso-elevation contours. (F) Detailed view of the region from (E) that is outlined. All contour intervals are 1.5°. (G) Ocular dominance preference map. (H) Spatial frequency preference map. Scale bars are 1 mm. Scale bar in (A) applies to (A)–(E), (G), and (H). Orientation, ocular dominance, and spatial frequency maps in this figure were obtained from different animals.
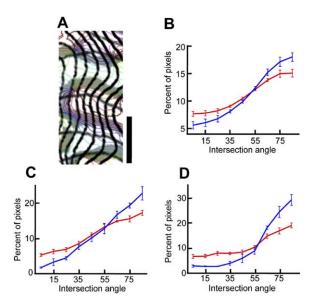
Figure 6. Relationships between the Map of Retinotopy and the Maps of Orientation, Ocular Dominance, and Spatial Frequency

(A) Orientation contours (color) superimposed with retinotopic contours (black), for the outlined region from Figure 5E. The gray regions are the high-gradient (top 30th percentile) regions of the orientation map. (B–D) Percent of pixels that have an intersection angle, within each 10° range whose center is indicated, between the retinotopic gradient and the gradient of the orientation (B), ocular dominance (C), or spatial frequency (D) map. Calculation over all pixels, red lines. Calculation restricted to pixels whose orientation (B), ocular dominance (C), or spatial frequency (D) gradient is within the highest 30th percentile, blue lines. Scale bar in (A) is 1 mm. Error bars denote standard error of the mean percentages over eight ferrets (B) or four ferrets (C and D).

dominance, and spatial frequency maps (cf. Figures 1A and 1B and Figure S1A).

We next directly compared the layout of the retinotopic map with that of other maps. By examining retinotopic contours superimposed on the orientation map (Figures 5E, 5F, and 6A), it can be seen that the high-gradient axes of these two maps are orthogonal. To quantify this relationship, we calculated the pixel-by-pixel intersection angles between the retinotopic gradient and the gradients of the other maps (see Experimental Procedures). Across the cortex, the percentage of pixels having a 60°–90° intersection angle would be 33.3% for a random distribution into the three bins of 0°–30°, 30°–60°, and 60°–90°. For the retinotopic and orientation maps, the number was 44.1% ± 4.5%, significantly higher than the chance level (p < 0.005, t test; n = 8 animals). For retinotopy and ocular dominance, the number was 48.0% ± 2.8% of pixels (p < 0.005, t test, n = 4 animals), and for retinotopy and spatial frequency, the number was 51.2% ± 4.9% of pixels (p < 0.005, t test, n = 4 animals). Thus there was a preponderance of near-orthogonal intersection angles between the retinotopic gradient and the gradients of the other maps, which reflects the global relationship between the anisotropy of the retinotopic map and the elongation of feature map domains.

Quantifying this in another way, we grouped pixels into nine 10° bins (0°–10°, 10°–20°, …, 80°–90°) according to the intersection angle between the retinotopic gradient and either the orientation, ocular dominance, or spatial frequency gradient. In all three cases, the percent of pixels per bin increased with intersection angle (orientation: Figure 6B, red line, r = 0.91, slope = 0.11 percent/deg; ocular dominance: Figure 6C, red line, r = 0.96, slope = 0.16 percent/deg; spatial frequency: Figure 6D, red line, r = 0.88, slope = 0.16 percent/deg).

The relationship between the retinotopic and orientation map gradients became even more robust if the analysis was performed only for those pixels of the cortex where the orientation gradient was within the highest 30th percentile of all values within the map (Figure 6A, gray regions). In this case, the percent of pixels per bin increased even more steeply with intersection angle (Figure 6B, blue line; r = 0.93, slope = 0.17 percent/deg). Similarly, the relationships between retinotopy and either ocular dominance or spatial frequency maps became stronger in the highest-gradient regions of the latter maps (ocular dominance: Figure 6C, blue line; r = 0.96, slope = 0.27 percent/deg; spatial frequency: Figure 6D, blue line; r = 0.89, slope = 0.35 percent/deg). These observations suggest a local influence of retinotopy on the layouts of other feature maps.

In summary, domains in the orientation, ocular dominance, and spatial frequency maps were anisotropic and were elongated specifically along the high-gradient axis of the retinotopic map. This implies that, on average, the high-gradient axis of the retinotopic map is orthogonal to that of other maps. The dimension-reduction model, which works under the assumption that the visual and other feature maps in visual cortex are interdependent, predicted just this effect (cf. Figure 1C).

Gradient Relationships between Orientation, Ocular Dominance, and Spatial Frequency Maps
The results above suggest that the visual space map constrains the mapping of other features. We next determined whether, given this constraint, the remaining feature maps were coupled by specific spatial relationships. We first examined the gradient relationships between orientation, ocular dominance, and spatial frequency. In Figure 7A, we superimpose the gradient maps of these three features. The figure shows that, while the high-gradient regions of each map are stretched along a similar axis of cortex, they interleave so as to avoid overlapping. To quantify this, each pixel from each gradient map was assigned a value between 1 and 10, where 1 indicated that the gradient was in the lowest 10th percentile of all gradient values within that map, and 10 the highest 10th percentile. The values from all three maps were averaged at each pixel in the region, and the standard deviation across the cortex of this averaged value was calculated. We found that the standard deviation was smaller if the actual gradient maps were used (Figure 7B, red dotted bar) compared to when the pixels in each gradient map were randomly shuffled (Figure 7B, blue histogram; p < 0.0001). This occurred in all three animals tested and indicates that the averaged gradient over all three maps stays relatively constant across the cortex.

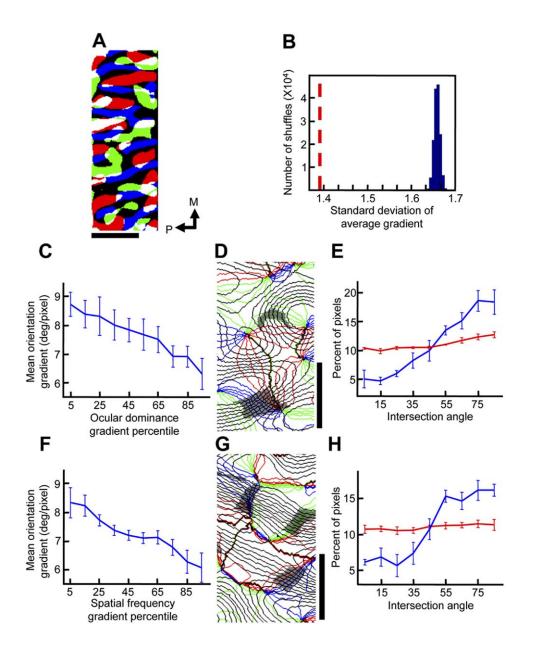We next compared the gradient magnitudes at corre-

Figure 7. Relationships between the Orientation, Ocular Dominance, and Spatial Frequency Maps

(A) High-gradient regions (top 30th percentile) of orientation (blue), ocular dominance (green), spatial frequency (red), or two or more maps (white). (B) Standard deviation across cortex of the average gradient magnitude of the orientation, ocular dominance, and spatial frequency maps (red dotted line). Histogram of standard deviations after pixels within each map were randomly shuffled, in 150,000 cases (blue histogram). (C and F) Pixels are grouped into ten bins according to their ocular dominance (C) or spatial frequency (F) gradient percentile, and the mean orientation gradient for each group is indicated. (D and G) Colored lines representing orientation contours are superimposed on black lines representing ocular dominance (D) or spatial frequency (G) contours. Gray regions indicate locations where high-gradient regions (top 30th percentile) of the orientation and ocular dominance (D) or spatial frequency (G) maps coincided. (E and H) Percent of pixels that have an intersection angle, within each 10° range whose center is indicated, between the orientation gradient and the ocular dominance (E) or spatial frequency (H) gradients. Average over all pixels, red line. Calculation for high orientation and ocular dominance (E) or spatial frequency (H) gradient overlap regions, blue line. Scale bar in (A) is 1 mm, bars in (D) and (G) are 0.5 mm. Error bars (C, E, F, and H) denote standard error of the mean values over five ferrets.

sponding pixels in orientation and ocular dominance maps. We found that the average orientation gradient calculated over all pixels whose ocular dominance gradient was within the highest 20th percentile was significantly lower than that calculated for pixels whose ocular dominance gradient was within the lowest 20th percentile (average difference of 1.9 ± 0.6 deg/pixel; means of 6.6 and 8.5 deg/pixel, respectively; paired t test, p < 0.001, n = 5 animals). In addition, all the pixels within a region were binned into ten groups according to their ocular dominance gradient percentile, and the average orientation gradient within each group was de-

termined. We found a strong negative correlation between the mean orientation gradient and the ocular dominance gradient percentile (r = –0.77, slope = –0.03 deg/pixel/percentile; Figure 7C).

The gradient relationships also held between orientation and spatial frequency maps. The average orientation gradient calculated for pixels whose spatial frequency gradient was within the highest 20th percentile was significantly lower than that calculated for pixels whose spatial frequency gradient was within the lowest 20th percentile (average difference of 2.1 ± 0.8 deg/pixel; means of 6.3 and 8.4 deg/pixel, respectively; paired t test, p < 0.001, n = 5 animals). There was a strong negative correlation between the mean orientation gradient and the spatial frequency gradient percentile, which held throughout the cortex (r = –0.71, slope = –0.02 deg/pixel/percentile; Figure 7F).

The results demonstrate that, in ferret V1, where the retinotopic map imposes a strong constraint on the remaining maps, the latter maintain specific relationships with one another. Their highest-gradient regions avoid one another, so that their combined gradient remains constant across the cortical surface. These relationships closely match the predictions of our dimension-reduction model (cf. Figures 2B and 2C and Figure S1B).

Contour Relationships between Orientation, Ocular Dominance, and Spatial Frequency Maps

We next examined whether the contours of orientation, ocular dominance, and spatial frequency maps tend to intersect at near-perpendicular angles. Previous studies have shown that this occurs in cat and monkey V1. But it may not occur in ferret V1 since (as noted above) the domains of orientation, ocular dominance, and spatial frequency maps are elongated along a parallel (rather than perpendicular) axis of cortex, reflecting the visual map anisotropy (see Figure 7A).

The contour lines of a superimposed orientation and ocular dominance map are illustrated in Figure 7D. Across the cortex, we found no significant tendency for the gradient vectors of the two maps to intersect at near-perpendicular angles. Quantification revealed that 36.2% ± 1.8% of pixels had gradients that intersected at an angle between 60° and 90°, which is not significantly different from the chance percentage of 33.3% (p > 0.1, t test, n = 5 animals). By grouping the pixels into nine 10° bins according to their intersection angles, the percent of pixels per bin did not increase greatly with intersection angle (Figure 7E, red line; r = 0.76, slope = 0.037 percent/deg). However, for the pixels of the cortex where both the orientation and ocular dominance gradients were within the highest 30th percentile, a clear tendency for orthogonal intersection angles between gradient vectors of the two maps emerged (Figure 7D, gray region; Figure 7E, blue line; r = 0.89, slope = 0.20 percent/deg). Considering this region, 49.3% ± 5.0% of pixels had gradient vector intersection angles within 60°–90° (p < 0.01, t test, n = 5 animals).

The relationships between the orientation and spatial frequency contours (Figure 7G) were similar. Across the cortex, there was only a weak tendency for orthogonal intersection angles between the gradients of the two maps, with 34.3% ± 2.8% of pixels having gradient in-

tersection angles within 60°–90° (p > 0.1, t test, n = 5 animals). By grouping the pixels into nine 10° bins according to their intersection angles, the percent of pixels per bin did not change greatly with intersection angle (Figure 7H, red line; r = 0.28, slope = 0.011 percent/deg). In contrast, for the pixels of the cortex where both the orientation and spatial frequency gradients were within the highest 30th percentile, there was a clear tendency for orthogonal intersection angles between the gradient vectors of the two maps (Figure 7G, gray regions; Figure 7H, blue line, r = 0.77, slope = 0.16 percent/deg). In that case, 47.1% ± 1.4% of gradient vectors had intersection angles within 60°–90° (p < 0.005, t test, n = 5 animals).

The results demonstrate, as our model predicted, that the visual map anisotropy decreases the overall tendency for perpendicular intersections between gradients of orientation and ocular dominance (and spatial frequency) maps. But the maps maintain strong orthogonal relationships in those locations where they have overlapping high gradients (cf. Figure 2 and Figure S1).

Influence of Feature Maps on Local Retinotopy: Model and Experiment

We showed above that the structure of the retinotopic map is reflected in the layouts of other feature maps. We lastly wanted to determine whether, in turn, the local structure of the retinotopic map is influenced by the mapping of these other features. Previous simulation studies predicted that the orientation map would have a strong influence on the retinotopic map (Durbin and Mitchison, 1990; Obermayer et al., 1990), but those simulations included only orientation and retinotopy as mapped features. We compared the predictions of simulations using only orientation and retinotopy to those that included additional features known to be mapped in ferret V1 and examined how the predictions in each case related to experimental relationships.

In the retinotopy-orientation simulation (the retinotopic map was anisotropic), we found that the retinotopic contours were distorted (Figure 8A). Iso-elevation lines were, in general, more widely spaced than average at pinwheel centers (dots in Figure 8A) and extremes of the orientation gradient map (although some counterexamples exist). We plotted the retinotopic gradient percentile as a function of the orientation gradient percentile (Figure 8C, red line) and found a clear negative correlation. Thus, in the retinotopy-orientation simulation, the detailed structure of the retinotopic map is visibly influenced by the orientation map, as has been shown previously with isotropic retinotopy (Durbin and Mitchison, 1990; Obermayer et al., 1990).

The results differed when the simulations included, besides orientation and retinotopy, the additional features of ocular dominance and spatial frequency (four components total, as in Figures 1 and 2). In this case, we did not find distortions in the retinotopic map that correlated with orientation pinwheels (Figure 8B), and there was no strong relationship between the gradient magnitudes of orientation and retinotopy (Figure 8C, blue line). Further, the retinotopy contours were smoother compared to the retinotopy-orientation simulation.

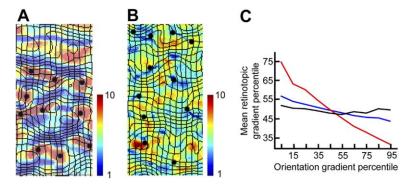We compared these predictions to experimentally

Figure 8. Local Gradient Relationships between the Retinotopic and Other Feature Maps

(A) Retinotopic contours are shown at high resolution from a two-component (retinotopy and orientation) simulation (retinotopic azimuth contour intervals represent four times the extent of visual space than elevation contour intervals). The background represents the normalized orientation gradient (10 indicates that a pixel is within the highest 10th percentile gradient, and 1 the lowest 10th percentile). Black dots represent orientation pinwheels. (B) Retinotopic contours are shown at high resolution from a four-component (retinotopy, orientation, ocular dominance, spatial frequency) simulation, as used throughout the manuscript (except in [A]). Retinotopic contour intervals same as in (A). Background represents the average normalized gradient over the orientation, ocular dominance, and spatial frequency maps. Black dots represent orientation pinwheels. (C) Pixels are grouped into ten bins according to their orientation gradient percentile, and the mean retinotopic gradient percentile for each group is indicated. Red line, from two-component simulation; blue line, from four-component simulation; black line, from optical-imaging experimental data.

measured relationships in ferret V1, where optical imaging was used to obtain retinotopic and orientation maps (see Figure 5). We found no strong correlation between the gradient magnitudes of retinotopy and orientation (Figure 8C, black line), or between the gradient magnitudes of retinotopy and either ocular dominance or spatial frequency (data not shown). We additionally measured the local relationships between the retinotopic and orientation maps using electrophysiological techniques (Figure S3) and found no strong correlation between the orientation gradient and either the receptive field gradient or the degree of receptive field overlap. These experimentally measured relationships between retinotopy and orientation were thus similar to the predictions of the four-component simulation and suggest that the retinotopic map may not be visibly distorted locally by the orientation map in the case where several features are mapped within a cortical area.

Discussion

Our study presents a comprehensive description of the mapping of visual space and of multiple other features within a cortical area. By directly comparing modeling and experimental results, we show that the spatial relationships between cortical maps, including those between visual space and other features, occur in conformity with a dimension-reduction strategy. This suggests that the constraints which drive map formation in the model, continuity (representing each feature smoothly across cortex) and coverage uniformity (representing each feature combination), may play a central role in determining the functional organization of visual cortex.

The Retinotopic Map and Other Feature Maps Are Interdependent

Recent experimental studies have come to conflicting conclusions regarding whether the structure of the retinotopic map is interdependent with that of other maps. In ferret V1, we find evidence for an interdependence: the retinotopic map in this species is strongly anisotropic, and the anisotropy is reflected in the layouts of other maps. Specifically, the gradient vectors of the ori-

entation, ocular dominance, and spatial frequency maps align along a specific axis of the retinotopic map, so that the highest-gradient axis of retinotopy is orthogonal to the highest-gradient axes of the remaining feature maps. Our model suggests that these relationships play a role in coordinating the mapping of visual space along with multiple additional features within a single cortical area.

While we suggest that retinotopy influences the layouts of other feature maps, some other factors have also been implicated. A role for the V1/V2 border is suggested by observations that ocular dominance columns in primates run perpendicular to this border (LeVay et al., 1985; Florence and Kaas, 1992). Since the retinotopic map and the V1/V2 border often have a specific relationship with one another, it can be difficult to distinguish between the influences of these two factors on the layouts of feature maps. However, in some cortical regions, retinotopy and the V1/V2 border are not aligned. In macaque V1, for example, they can deviate from each other, and here it was found that ocular dominance patterns follow local changes in the axis of the retinotopic anisotropy (Blasdel and Campbell, 2001). This favors the hypothesis that retinotopy itself has a direct relationship with other feature maps. Simulation studies have demonstrated that the shape of a cortical area is an additional factor that may influence the structures of feature maps (Bauer, 1995; Wolf et al., 1996). However, the shape of a cortical area also likely influences the degree of anisotropy of the retinotopic map; thus, cortical area shape may influence feature maps directly, or indirectly via its influence on the retinotopic map, which in turn influences the other feature maps.

The Local Pattern of Retinotopy
Whether local distortions exist in the retinotopic map (at a subcolumnar scale) that relate to the mapping of other features, such as pinwheel centers of the orientation map, has remained a question of particular interest. Early dimension-reduction models predicted specific local relationships between retinotopy and orientation (Durbin and Mitchison, 1990; Obermayer et al., 1990), whereas experimental measurements failed to confirm the predictions (Das and Gilbert, 1997; Bosking

et al., 2002; Buzas et al., 2003). However, a recent modeling study suggested that predictions of dimension-reduction models, regarding the relationships between retinotopy and other feature maps, are dependent on how many features are mapped within a cortical area (Swindale, 2004). In agreement with this, our studies suggest that the local relationships in ferret V1 between retinotopy and orientation maps are in close agreement with model predictions when a realistic number of features are simulated in the model.

For example, we show that low-gradient retinotopy regions are predicted to coincide with high-gradient orientation regions when these are the only two features mapped within an area. But if ocular dominance and spatial frequency maps also exist, then low retinotopic gradient regions should coincide with high-gradient regions of these additional maps as well. Given our observation that the high-gradient regions of orientation, ocular dominance, and spatial frequency maps occur in nonoverlapping locations of cortex, each of these feature maps should have opposing effects on the retinotopic map and smooth out the distorting effects of one another. In this case, the local retinotopic map is predicted to be smooth and to have no visible distortions that correlate with any other single map. These local relationship predictions are in agreement with our experimental results as well as those recently measured in other species (Bosking et al., 2002; Buzas et al., 2003). It remains possible that local distortions exist in the retinotopic map that cannot be detected by intrinsic-signal optical imaging or electrophysiological methods.

Contour Relationships between Multiple Maps
Mapping variables along orthogonal axes of cortex has long been suggested as a means to accommodate a smooth representation of multiple variables within a single cortical area (Hubel and Wiesel, 1977). However, since it is not possible for the domains of more than two maps to be mutually orthogonal, this cannot represent a general solution to the problem. In a modeling study (Obermayer et al., 1992), a relationship was suggested that could hold between more than two maps: that orthogonality occurs between any two maps when the high-gradient regions of those maps coincide. Our data directly confirm this prediction with experimental evidence. We find that although ocular dominance and orientation are not mapped along orthogonal axes (their domains are often elongated along a parallel axis), the contours of these maps intersect at perpendicular angles when high-gradient regions of the two maps coincide (and similarly for spatial frequency and orientation).

The general prediction that two features are mapped orthogonally in their high-gradient overlap regions may relate to previous observations in cat V1 that ocular dominance and orientation contours intersect at orthogonal angles in ocular dominance border regions (Bartfeld and Grinvald, 1992; Hubener et al., 1997). It may also relate to the finding in macaque V1 that orthogonality between orientation and ocular dominance contours increases in "linear" regions of the orientation map (Obermayer and Blasdel, 1993).

In some cases, two features can be mapped orthogonally throughout the cortex, not only locally. For example, orientation domains run along a specific axis of the retinotopic map in cat V2 (Cynader et al., 1987), and our results show that orientation, ocular dominance, and spatial frequency domains do the same in ferret V1. We suggest that this occurs due to the strong anisotropy in retinotopy within these two cortical areas; in these cases, the strong retinotopy gradient, which points along a constant axis of cortex, will cause the high-gradient regions and in turn the domains in other feature maps to be elongated along a constant axis.

While one recent study suggested that no orthogonal relationships exist between orientation and ocular dominance maps in ferret visual cortex (White et al., 2001), we found strong relationships. We suggest that the difference lies in the part of visual cortex examined: White and colleagues analyzed the V1/V2 border region, which has very large monocular domains. We analyzed the region containing interleaving ocular dominance columns, and we detect relationships consistent with those found in functionally similar regions of cat and primate.

Separability of Response Maps in V1
An assumption implicit in our study, that separable maps of multiple response properties exist in V1, is supported by numerous studies. It has previously been demonstrated that the orientation map (in binocular regions of cortex) is not dependent on the eye to which stimuli are presented; conversely, the map of ocular dominance is invariant to stimulus orientation (Blasdel, 1992). Further, the overall structure of the retinotopic map is not dependent on stimulus orientation (Blasdel and Campbell, 2001). While the receptive field of a neuron can differ for the two eyes, this relationship does not lead to systematic changes in retinotopic magnification, which implies that within binocular regions of cortex the overall pattern of retinotopy is not dependent on the eye to which stimuli are presented. Thus, we suggest that the maps of orientation, ocular dominance, and retinotopy are separable.

Whether orientation and spatial frequency maps can be approximated as separable from one another has also been examined. It has been shown, using sine-wave gratings, that the structure of the orientation map is not altered by the stimulus spatial frequency (Issa et al., 2000). On the other hand, the preferred spatial frequency of some neurons can change with orientation (Webster and De Valois, 1985). While one study suggests that the spatial frequency map has some dependence on stimulus orientation (Issa et al., 2000), our results suggest that the structure of the differential spatial frequency map in the region of ferret V1 we examined remains largely invariant to orientation.

Responses to texture stimuli are also consistent with separable maps of orientation and spatial frequency. An optical imaging study (Basole et al., 2003) showed, using texture stimuli composed of short line segments, that changing the bar length or axis of motion causes changes in the cortical population response. In particular, an activation pattern produced by long bars (gratings) moving orthogonally was reproduced by short

bars moving obliquely. Such equivalence has been previously described using interdigitating gratings composed of short line segments, which can elicit patterns of cortical activity that resemble activity patterns due to orthogonal gratings (Sheth et al., 1996). But compared to gratings, short line segments contain a broader range of spatial frequencies and orientations in their Fourier spectrum, and the population response to short segments in these studies is fully compatible with both the known spatiotemporal responses of V1 neurons (Mante and Carandini, 2003) and the existence of separable maps of orientation and spatial frequency (Baker and Issa, 2005; Mante and Carandini, 2005).

Experimental Procedures

Computational Model
To simulate the mapping of response features across the cortex we used the Kohonen self-organizing map algorithm (Kohonen, 1982b), as modified by Obermayer (Obermayer et al., 1992). A multidimensional feature space is defined, where each stimulus is represented as a multicomponent vector $V_s = (x_s, y_s, q_s \cos(2\phi_s), q_s \sin(2\phi_s), z_s, f_s)$ within this space. Here x and y correspond to azimuth and elevation retinotopic position, respectively, q is orientation selectivity, ϕ is orientation preference, z is ocular dominance, and f is spatial frequency preference. The feature x ranges from $(0, X)$, y from $(0, Y)$, q from $(0, Q)$, ϕ from $(0, \pi)$, z from $(0, Z)$, and f from $(0, F)$. The stimuli are mapped onto a cortical surface, which is represented as a two-dimensional grid of points with size $N \times N$. Each cortical point $r = (i, j)$ has a receptive field defined as $W_r = (x_r, y_r, q_r \cos(2\phi_r), q_r \sin(2\phi_r), z_r, f_r)$. At the beginning of the simulation, the maps are initialized as $x_r = i/N$, $y_r = j/N$, $q_r = Q/2$, $\phi_r = \pi/2$, $z_r = Z/2$, and $f_r = F/2$. The maps are formed through iterations (1.5 million) of three steps. (1) A stimulus V_s is chosen one at a time randomly from the complete feature space, assuming uniform distributions of each feature. (2) The cortical point $r_c = (i_c, j_c)$, whose preferred features are closest to those of the stimulus, is identified as the "winner." The closeness of the feature is measured with the Euclidian distance between the vectors $|V_s - W_r|^2$. (3) The preferred features of the cortical points are updated according to the equation $\Delta W_r = \alpha h(r)(V - W_r)$. Here, α is the learning rate, r is the cortical distance between a given cortical point (i, j) and the winner r_c, and $h(r) = \exp(-r^2/\sigma^2)$ is the neighborhood function. The neighborhood function restricts the changes in receptive fields to those cortical points nearby the winner (in cortical distance).

We used the following parameters for the simulations: $N = 513$, $\sigma = 5$, $\alpha = 0.02$, $X = \rho N$, $Y = N$, $Q = 40$, $Z = 60$, $F = 60$. Here, ρ is the elevation:azimuth magnification ratio of the retinotopic map and was chosen to match the magnification anisotropy that exists in the ferret or to simulate an isotropic map. Thus, the anisotropy was achieved by mapping different extents of azimuth and elevation visual space onto a square cortex. Alternately, a similar anisotropy could be obtained by simulating the mapping of an isotropic visual space onto an oval-shaped model cortex; similar map relationships resulted in both cases (data not shown). The maps displayed and analyzed were derived from a portion of the model cortex that did not include the boundary regions. We found that the gradient vector relationships between pairs of maps persisted for a range of simulation parameters (Q from 30 to 50, Z or F from 50 to 80, and σ from 5.0 to 5.5), while the degree or strength of these relationships systematically varied within these ranges. We were able to produce realistic orientation maps only when the orientation selectivity (q) was allowed to vary across the cortex; a number of parameter and annealing regimes were attempted to achieve a fixed orientation selectivity. The parameters used in the manuscript were chosen so that the relative wavelengths of multiple maps, as well as the orientation pinwheel density, matched between the simulation and experimental data. We found that while the map structures changed significantly during the initial simulation iterations, they did not change greatly between 1.5 million (the number used in this study) and 6 million (the maximum number tested) iterations.

Animals
Ten adult ferrets were used in these experiments. Animals were prepared for acute experiments according to protocols approved by MIT's Animal Care and Use Committee. Details have been described (Rao et al., 1997). Anesthesia was induced with ketamine (25 mg/kg) and xylazine (1.5 mg/kg) and maintained with isoflurane (1.0% to 1.5% in 70:30 mixture of N_2O/O_2) delivered through a tracheal cannula using artificial respiration. Fluid maintenance was achieved with a 50:50 mixture of 5% dextrose and lactated Ringer's solution, supplemented with Norcuron (0.25 mg/kg/hr) for paralysis. Expired CO_2 was maintained at 4%, and the anesthesia level was monitored continuously. A craniotomy and durotomy were performed to expose V1. A chamber was mounted on the skull around the exposed region and filled with agarose (1.5% in saline). This was covered by a cover glass and then silicone oil.

Optical Imaging
For a description of the optical imaging procedures, see the Supplemental Data. For orientation and spatial frequency maps, stimuli were presented binocularly and consisted of drifting, full-field square-wave gratings having one of four orientations (separated by 45°), one of four fundamental spatial frequencies (0.08, 0.125, 0.225, or 0.375 cycles/deg), and a temporal frequency of 1 Hz. Monocularly presented drifting gratings (four orientations, spatial frequency of 0.125 cycles/deg) were used to obtain ocular dominance maps. For further details on how orientation, ocular dominance, and spatial frequency maps were obtained, see the Supplemental Data.

To generate retinotopic maps, we used a periodic visual stimulation paradigm (presented to the contralateral eye) combined with continuous data acquisition optical imaging (Kalatsky and Stryker, 2003). For azimuth maps, the stimulus consisted of elongated vertical bars (1° × 30°) separated by 20°, which each flashed at 3 Hz and shifted their azimuthal location by 0.66° every second. A given location of space was thus stimulated every 30 s. Light-reflectance images were captured at 1 Hz. For elevation maps, the stimulus consisted of elongated horizontal bars (1° × 40°) separated by 15°, which drifted continuously at a rate of 1°/s in the elevation dimension. Each location of space was thus stimulated every 15 s. Images were captured at 3 Hz. Each retinotopic stimulus trial consisted of 8 to 24 cycles of stimulation, after which a blank screen was shown for 25 s, and each experiment consisted of 5 to 15 such trials. The light-reflectance data were averaged in-phase over all trials and cycles. Each frame of this averaged response (45 frames total for elevation, 30 for azimuth) thus consisted of the activation pattern resulting from stimulation during a restricted phase of the stimulus cycle. To obtain the retinotopic single-condition maps, each frame was subtracted from the mean of all the frames and then filtered (gaussian filter, standard deviation of 0.06 mm). To determine each pixel's preferred receptive field position, we calculated the phase of the Fast Fourier Transform (FFT) at the stimulation frequency, on the time course response of each pixel (Kalatsky and Stryker, 2003). Since the optical imaging signals follow the stimulation with an unknown lag time, this method provides maps of relative rather than absolute retinotopy. All of our analyses rely only on relative retinotopy values; for display purposes, we estimated the absolute values based on published maps of retinotopy in ferret (Law et al., 1988).

Analysis of Map Structures and Relationships
We used identical procedures for analyzing our computational and experimental data. Gradient maps were computed from the low-pass filtered maps of retinotopy, orientation, ocular dominance, or spatial frequency as the two-dimensional spatial derivative at each pixel. Let A(x, y) be the value at a pixel (x, y), dx = (A(x + 1, y) − A(x, y)) and dy = (A(x, y + 1) − A(x, y)). Then the gradient vector magnitude at (x, y) is

$$\sqrt{dx^2 + dy^2},$$

and the gradient vector angle is $(180/\pi)\text{atan}(dy/dx)$ and ranges from −90 to 90. For orientation, dx and dy were corrected to take into account circularity. The gradient magnitude describes how much a feature is changing around a given pixel in the map, and the gradi-

ent angle indicates the axis of cortex along which the feature is changing maximally (and is orthogonal to the map contour at that pixel).

The retinotopy gradient angle at a pixel was defined as being the direction orthogonal to the elevation gradient angle and approximates the cortical axis along which retinotopy changes maximally (see Supplemental Data). To obtain the retinotopy gradient magnitude at each pixel, we calculated the gradient magnitude separately for both the azimuth and elevation maps and summed them. The orientation gradient map was derived from the orientation angle map, i.e., the map of ϕ (see Supplemental Data), and the ocular dominance and spatial frequency gradient maps were derived from the ocular dominance and spatial frequency preference maps, respectively.

The gradient percentile indicates the percentage of pixels within a region whose gradient is at or below the gradient of the pixel in question. The gradient magnitude comparison of two maps, and the associated plots, were obtained by first grouping the pixels within a region into ten equal-sized bins according to the gradient magnitude percentile of one feature map (the x axis tick marks indicate the mean percentile for each bin). Then, for all pixels of each bin, the average gradient value (or percentile) of the second feature map was calculated. The gradient intersection angle comparisons of two maps, and the associated plots, were obtained by calculating the pixel-by-pixel difference in gradient directions for two feature maps (which range from 0° to 90°). The percent of total pixels from the region falling in each of nine 10° bins (0°–10°, 10°–20°, …, 80°–90°) is plotted.

Electrophysiology
For a description of the electrophysiology, see the Supplemental Data.

Supplemental Data
The Supplemental Data include four figures, one table, and Supplemental Experimental Procedures. They can be found with this article online at http://www.neuron.org/cgi/content/full/47/2/267/DC1/.

References

Baker, T.I., and Issa, N.P. (2005). Cortical maps of separable tuning properties predict population responses to complex visual stimuli. J Neurophysiol., in press. Published online March 9, 2005.

Bartfeld, E., and Grinvald, A. (1992). Relationships between orientation-preference pinwheels, cytochrome oxidase blobs, and ocular-dominance columns in primate striate cortex. Proc. Natl. Acad. Sci. USA 89, 11905–11909.

Basole, A., White, L.E., and Fitzpatrick, D. (2003). Mapping multiple features in the population response of visual cortex. Nature 424, 986–990.

Bauer, H.U. (1995). Development of oriented ocular dominance bands as a consequence of areal geometry. Neural Comput. 7, 36–50.

Blasdel, G.G. (1992). Differential imaging of ocular dominance and orientation selectivity in monkey striate cortex. J. Neurosci. 12, 3115–3138.

Blasdel, G., and Campbell, D. (2001). Functional retinotopy of monkey visual cortex. J. Neurosci. 21, 8286–8301.

Blasdel, G.G., and Salama, G. (1986). Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. Nature 321, 579–585.

Bosking, W.H., Crowley, J.C., and Fitzpatrick, D. (2002). Spatial coding of position and orientation in primary visual cortex. Nat. Neurosci. 5, 874–882.

Buzas, P., Volgushev, M., Eysel, U.T., and Kisvarday, Z.F. (2003). Independence of visuotopic representation and orientation map in the visual cortex of the cat. Eur. J. Neurosci. 18, 957–968.

Cynader, M.S., Swindale, N.V., and Matsubara, J.A. (1987). Functional topography in cat area 18. J. Neurosci. 7, 1401–1413.

Das, A., and Gilbert, C.D. (1997). Distortions of visuotopic map match orientation singularities in primary visual cortex. Nature 387, 594–598.

Durbin, R., and Mitchison, G. (1990). A dimension reduction framework for understanding cortical maps. Nature 343, 644–647.

Engel, S.A., Rumelhart, D.E., Wandell, B.A., Lee, A.T., Glover, G.H., Chichilnisky, E.J., and Shadlen, M.N. (1994). fMRI of human visual cortex. Nature 369, 525.

Erwin, E., Obermayer, K., and Schulten, K. (1995). Models of orientation and ocular dominance columns in the visual cortex: a critical comparison. Neural Comput. 7, 425–468.

Florence, S.L., and Kaas, J.H. (1992). Ocular dominance columns in area 17 of Old World macaque and talapoin monkeys: complete reconstructions and quantitative analyses. Vis. Neurosci. 8, 449–462.

Friedman, R.M., Chen, L.M., and Roe, A.W. (2004). Modality maps within primate somatosensory cortex. Proc. Natl. Acad. Sci. USA 101, 12724–12729.

Goodhill, G.J., and Willshaw, D.J. (1990). Application of the elastic net algorithm to the formation of ocular dominance stripes. Network-Comp. Neural. 1, 41–59.

Hubel, D.H., and Wiesel, T.N. (1963). Shape and arrangement of columns in cat's striate cortex. J. Physiol. 165, 559–568.

Hubel, D.H., and Wiesel, T.N. (1977). Ferrier lecture. Functional architecture of macaque monkey visual cortex. Proc. R. Soc. Lond. B. Biol. Sci. 198, 1–59.

Hubener, M., Shoham, D., Grinvald, A., and Bonhoeffer, T. (1997). Spatial relationships among three columnar systems in cat area 17. J. Neurosci. 17, 9270–9284.

Issa, N.P., Trepel, C., and Stryker, M.P. (2000). Spatial frequency maps in cat visual cortex. J. Neurosci. 20, 8504–8514.

Kalatsky, V.A., and Stryker, M.P. (2003). New paradigm for optical imaging: temporally encoded maps of intrinsic signal. Neuron 38, 529–545.

Kohonen, T. (1982a). Analysis of a simple self-organizing process. Biol. Cybern. 44, 135–140.

Kohonen, T. (1982b). Self-organized formation of topologically correct feature maps. Biol. Cybern. 43, 59–69.

Koulakov, A.A., and Chklovskii, D.B. (2001). Orientation preference patterns in mammalian visual cortex: a wire length minimization approach. Neuron 29, 519–527.

Law, M.I., Zahs, K.R., and Stryker, M.P. (1988). Organization of primary visual cortex (area 17) in the ferret. J. Comp. Neurol. 278, 157–180.

LeVay, S., Connolly, M., Houde, J., and Van Essen, D.C. (1985). The complete pattern of ocular dominance stripes in the striate cortex and visual field of the macaque monkey. J. Neurosci. 5, 486–501.

Linden, J.F., and Schreiner, C.E. (2003). Columnar transformations in auditory cortex? A comparison to visual and somatosensory cortices. Cereb. Cortex 13, 83–89.

Mante, V., and Carandini, M. (2003). Visual cortex: seeing motion. Curr. Biol. 13, R906–R908.

Mante, V., and Carandini, M. (2005). Mapping of stimulus energy in primary visual cortex. J. Neurophysiol., in press. Published online March 9, 2005.

Mrsic-Flogel, T., Hubener, M., and Bonhoeffer, T. (2003). Brain mapping: new wave optical imaging. Curr. Biol. *13*, R778–R780.

Obermayer, K., and Blasdel, G.G. (1993). Geometry of orientation and ocular dominance columns in monkey striate cortex. J. Neurosci. *13*, 4114–4129.

Obermayer, K., Ritter, H., and Schulten, K. (1990). A principle for the formation of the spatial structure of cortical feature maps. Proc. Natl. Acad. Sci. USA *87*, 8345–8349.

Obermayer, K., Blasdel, G.G., and Schulten, K. (1992). Statistical-mechanical analysis of self-organization and pattern formation during the development of visual maps. Phys. Rev. A. *45*, 7568–7589.

Rao, S.C., Toth, L.J., and Sur, M. (1997). Optically imaged maps of orientation preference in primary visual cortex of cats and ferrets. J. Comp. Neurol. *387*, 358–370.

Redies, C., Diksic, M., and Riml, H. (1990). Functional organization in the ferret visual cortex: a double-label 2-deoxyglucose study. J. Neurosci. *10*, 2791–2803.

Sereno, M.I., Dale, A.M., Reppas, J.B., Kwong, K.K., Belliveau, J.W., Brady, T.J., Rosen, B.R., and Tootell, R.B. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. Science *268*, 889–893.

Sheth, B.R., Sharma, J., Rao, S.C., and Sur, M. (1996). Orientation maps of subjective contours in visual cortex. Science *274*, 2110–2115.

Sur, M., Wall, J.T., and Kaas, J.H. (1981). Modular segregation of functional cell classes within the postcentral somatosensory cortex of monkeys. Science *212*, 1059–1061.

Swindale, N.V. (1991). Coverage and the design of striate cortex. Biol. Cybern. *65*, 415–424.

Swindale, N.V. (1996). The development of topography in the visual cortex: a review of models. Network-Comp. Neural. *7*, 161–247.

Swindale, N.V. (2004). How different feature spaces may be represented in cortical maps. Network *15*, 217–242.

Webster, M.A., and De Valois, R.L. (1985). Relationship between spatial-frequency and orientation tuning of striate-cortex cells. J. Opt. Soc. Am. A *2*, 1124–1132.

White, L.E., Bosking, W.H., Williams, S.M., and Fitzpatrick, D. (1999). Maps of central visual space in ferret V1 and V2 lack matching inputs from the two eyes. J. Neurosci. *19*, 7089–7099.

White, L.E., Bosking, W.H., and Fitzpatrick, D. (2001). Consistent mapping of orientation preference across irregular functional domains in ferret visual cortex. Vis. Neurosci. *18*, 65–76.

Wolf, F., Bauer, H.U., Pawelzik, K., and Geisel, T. (1996). Organization of the visual cortex. Nature *382*, 306–307.