

Visual cortex maps are optimized for uniform coverage

Nicholas V. Swindale¹, Doron Shoham², Amiram Grinvald², Tobias Bonhoeffer³ and Mark Hübener³

¹ Dept. of Ophthalmology, University of British Columbia, 2550 Willow St., Vancouver, B.C., V5Z 3N9, Canada

² The Weizmann Institute of Science, Department of Neurobiology, Rehovot 76100, Israel

³ Max-Planck-Institut für Neurobiologie, Am-Klopferspitz 18A, D-82152 Martinsried, Germany

Correspondence should be addressed to N.V.S. (swindale@interchange.ubc.ca)

Cat visual cortex contains a topographic map of visual space, plus superimposed, spatially periodic maps of ocular dominance, spatial frequency and orientation. It is hypothesized that the layout of these maps is determined by two constraints: continuity or smooth mapping of stimulus properties across the cortical surface, and coverage uniformity or uniform representation of combinations of map features over visual space. Here we use a quantitative measure of coverage uniformity (c') to test the hypothesis that cortical maps are optimized for coverage. When we perturbed the spatial relationships between ocular dominance, spatial frequency and orientation maps obtained in single regions of cortex, we found that cortical maps are at a local minimum for c' . This suggests that coverage optimization is an important organizing principle governing cortical map development.

Mammalian primary visual cortex contains a single continuous representation of retinotopic visual space, on which orderly, periodic maps of several different visual stimulus properties are superimposed. These properties include ocular dominance and preferred orientation¹⁻⁴, direction of motion⁵⁻⁷, color⁸, spatial frequency⁹⁻¹¹ and possibly others¹². Although the total number of different maps that may be present in any one visual area of a particular species is unknown, the largest number so far described simultaneously in a single patch of cortex is three: in area 17 of the cat, optical imaging experiments have demonstrated functional maps of orientation, ocular dominance and spatial frequency¹³, each with a periodicity of roughly 1 mm (Fig. 1).

These maps are often structurally related. For example, in macaque monkey^{14,15}, iso-orientation lines tend to cross ocular dominance column borders at right angles, whereas orientation singularities (points where a single complete set of iso-orientation domains meet) tend to lie in the middle of left- or right-eye ocular dominance stripes. Similar relationships between orientation and ocular dominance maps, and between orientation and spatial frequency maps, have been observed in cat area 17 (ref. 13).

These findings are consistent with earlier suggestions⁴ that visual cortex maps develop according to a combination of continuity and completeness constraints, which act in opposition. The continuity constraint specifies that neighboring cortical locations should have similar receptive field properties, whereas the completeness constraint ensures that all combinations of the parameters represented in individual maps are distributed uniformly over visual space. A quantitative measure of completeness, known as coverage uniformity, or c' , has been devised¹⁶. It is calculated in the following way: for a given combination of features (for example, some unique combination of orientation, spatial frequency, eye and retinal location), the total neural response, A , in the cortex is calculated, taking into

account the spatial structure of the maps of these properties, and the receptive field tuning widths of individual cortical units. Uniform coverage means that A should be independent of the specific feature combination chosen. It is convenient to define c' as the standard deviation of A divided by its mean, taken over some representative set of stimuli. This makes it a dimensionless measure of 'noise' in the cortical representation of a particular feature space. If $c' = 0$, coverage is completely uniform; larger values correspond to an increasingly noisy representation: for example, if $c' = 1$, the standard deviation of the signal across the feature space is equal to the mean.

The hypothesis that cortical maps are organized so as to optimize (that is, minimize) c' was tested here by systematically perturbing the spatial relationships between maps of orientation, ocular dominance and spatial frequency obtained simultaneously in area 17 of the cat¹³ to see whether c' is at a local minimum. Two different methods were used to do this, both of which left continuity in the individual maps unchanged. In the first, the spatial relationships were altered by various combinations of flips (mirror inversions) about either the horizontal or vertical axes and/or 180° rotation (equivalent to a mirror inversion about one axis followed by a mirror inversion about the other). For three rectangular maps, there are a total of 16 transformations that disturb the point-to-point relationships between the maps in a unique way. (Note that some combinations of flips are equivalent: for example, flipping two of the maps about the vertical axis is equivalent to flipping just the third one.) The second method examined the possibility that map structure is close to a local optimum for coverage uniformity. To test this, a single map was displaced sideways by a given number of pixels relative to the other two, which remained fixed relative to each other. Coverage was then calculated for the region common to all three maps. This was done separately for each of the ocular dominance, spatial frequency and orientation maps, for a range of

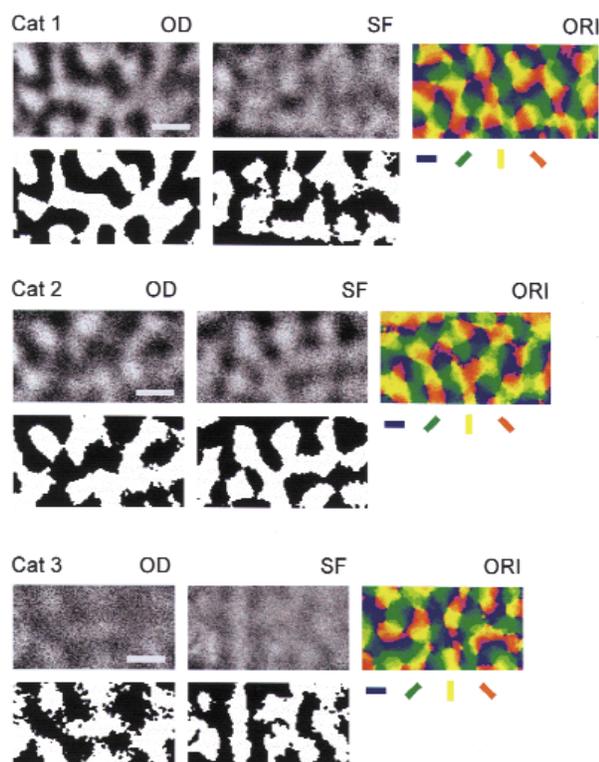


Fig. 1. Experimental data from three of the cats on which calculations were based. Top three images (OD, ocular dominance; SF, spatial frequency; ORI, orientation) show maps after rotating and trimming as described in the text. Bottom, ocular dominance and spatial frequency maps after smoothing and thresholding. Scale bar, 1 mm.

displacements of ± 10 pixels. The direction of the displacement was always along the long edge of the rectangular region, which corresponds roughly to the anterior–posterior axis of the brain. Flips and shifts were chosen as perturbations because they leave structure in the individual maps unchanged. Other types of perturbation, such as expansion, shrinkage, local distortions and

random changes in map values, were not used because in general they also change map continuity, that is, the degree of similarity in the receptive field properties of neighboring columns of cells in individual maps. We did not attempt to find local perturbations that leave continuity unchanged, although these, if they exist, would enlarge the range of tests that could be applied to the data.

RESULTS

Maps from five cats were analyzed. Two of these maps (Fig. 1; cats 1 and 2) showed strong, well-defined ocular-dominance and spatial-frequency patches. Cats 3, 4 and 5 showed somewhat less well-defined ocular-dominance and spatial-frequency patches. When maps were flipped or rotated relative to each other, in almost every case (78 of 80), coverage uniformity values got larger, that is, worse (Fig. 2). For cats 1 to 4, coverage values were larger, typically by about 10–25%, for every possible type of disarrangement. In cat 5, 14 of 16 manipulations produced increases of about 2–10% in c' , whereas 2 manipulations produced small, possibly insignificant, decreases.

In cats 1–4, displacement analysis indicated that the maps are close to a local minimum for coverage uniformity (Fig. 3). Moving any one of the three maps relative to the other two, in either direction, produced an increase in the value of c' , worsening coverage uniformity. The optima for the ocular dominance and spatial frequency maps were either at zero offset or ± 1 pixel from zero. Sometimes these optima seemed clearly misaligned, although only by one or two pixels. Because all the maps were obtained with interleaved stimulus presentations, it is unlikely that small movements of the camera and/or the cortex during the course of the experiment would cause systematic misalignments between the different maps. It is possible, perhaps, that the signals producing the different maps come from different depths in the cortex, and, depending on the angle of view, this might result in small systematic misalignments. Cat 5 showed local optima at zero displacement of the ocular dominance and the spatial frequency maps, although the change in coverage values was smaller compared to that observed in the other four animals. The orientation map in this animal, however, showed only a small, and arguably non-existent local optimum for zero displacement.

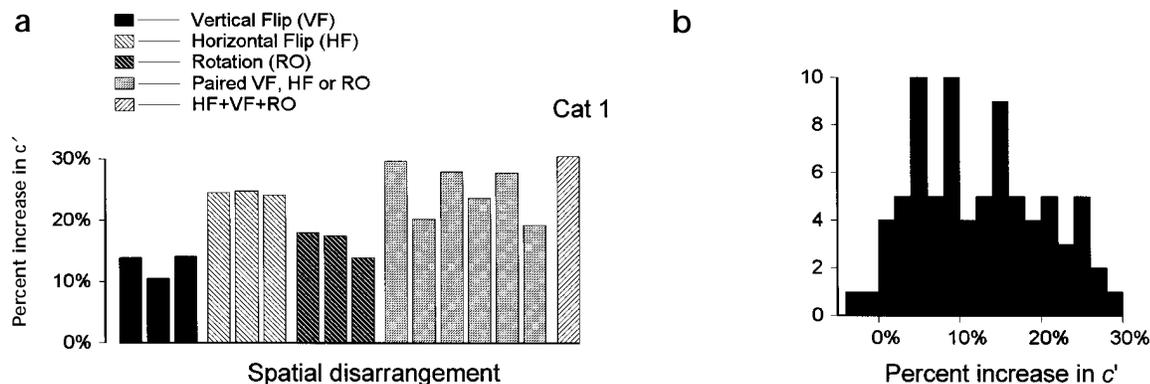


Fig. 2. Effects of various reflections and rotations of one or more maps, relative to each other, on coverage uniformity. The changes are expressed as the percentage change in c' (equation 3), relative to the value obtained with the maps in their original undisturbed positions. Sixteen possible types of perturbation were tested. (a) Values obtained in cat 1, arranged along the horizontal axis in roughly increasing order of disturbance. The first three values represent flips (mirror inversions) of one map about the horizontal axis. The next three values are from flips of a single map about the vertical axis; the next three values are rotations of a single map by 180° ; the next six values are combinations of flips or rotations applied to two maps, and the rightmost value is a combination of a horizontal flip, a vertical flip and a rotation applied to each of the three maps. (b) Distribution of changes in c' following these manipulations in all 5 experimental animals. The two cases where c' decreased were from cat 5.

articles

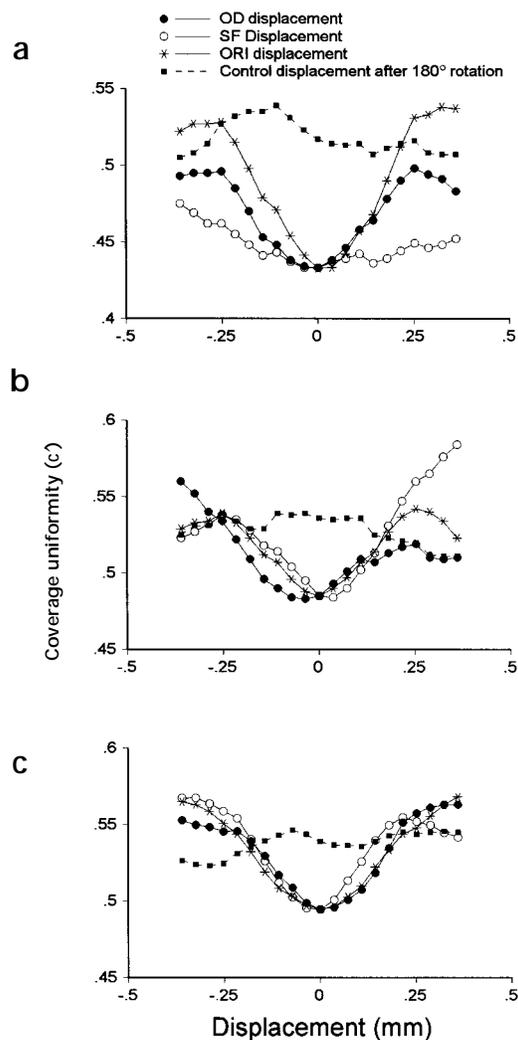
Fig. 3. Effect of displacing one of the maps sideways relative to the other two by a fixed amount (offsets are multiples of one pixel, 0.036 mm). The vertical axis shows coverage uniformity (c'), and the horizontal axis shows displacement in mm. Data from three cats are shown: cat 1 (a), cat 2 (b) and cat 4 (c). In the control conditions (squares), the map being moved was first rotated by 180°. In cats 1 and 4, this was the orientation map; in cat 2, the control was the ocular dominance map. The rotations have the expected effect of removing the local minimum in c' .

To check the program logic, and to control for the possibility that the local minima might be artifactual, the displacement analysis was repeated after carrying out several additional manipulations. First, the map that was to be moved relative to the other two was rotated by 180° (Fig. 3, squares). Second, Gaussian random noise values were substituted into the single-condition maps, which were then smoothed by convolution with a Gaussian filter with $\sigma = 3$ pixels to produce noise that was spatially correlated within, but not between, maps. Third, the images obtained in the second procedure were normalized by dividing by the mean of all of the images; this simulates the procedure of dividing by the cocktail blank, which is used to remove common artifacts (such as blood vessel noise) from single-condition optical images and might produce some form of spurious independence between the images. In all of these test cases, systematic changes in c' might be expected, but a well-defined minimum at zero should be absent. This was confirmed for all tests. In another set of tests, we examined the effects that might be produced by artifacts (such as blood vessels) common to all three maps. Common features (three thin vertical bars, 10 pixels wide and 40 pixels apart) were introduced into the spatial frequency and ocular dominance maps at corresponding locations. In one test, these features had the same polarity; in another, they were of opposite polarity. Both tests reduced or obliterated the local minimum at zero offset. This result would be expected on theoretical grounds, and it confirms that spatially coincident artifacts in the maps tend to reduce rather than accentuate the presence of local minima in c' .

DISCUSSION

Four of the five data sets tested provided unambiguous evidence that the spatial relationships between cortical maps of ocular dominance, spatial frequency and orientation optimize coverage uniformity, the uniform representation of all possible combinations of these three variables across visual space. Data from a fifth cat gave less clear-cut results, although the results from this animal were still substantially different from chance and point in the same direction.

A number of assumptions were necessary to do the analysis. These include estimates of orientation tuning width, average receptive field size and uniform retinotopy. The estimates of retinal magnification factor, orientation tuning width and receptive field size are within measured physiological ranges^{17,18}. They were made before doing the calculations and were not adjusted *post hoc* to give the most clear-cut results. We confirmed, nevertheless, that the coverage minima shown in Fig. 3 were qualitatively unaffected by changes in orientation tuning and receptive field width of $\pm 50\%$. An assumption of some kind about the nature of the retinotopic map was required because retinotopic data were not obtained in these experiments. Locally uniform retinotopy, such as has been demonstrated in the tree shrew (W.H. Bosking *et al.*, *Soc. Neurosci. Abstr.* 23, 755.7, 1997) is the simplest assumption in such circumstances. The possible effects of



a locally non-uniform retinotopic mapping on coverage have to be considered, but are hard to assess. There is evidence for discontinuities in the retinotopic map in the cat¹⁹, which has been interpreted to suggest that coverage must be non-optimal¹⁹. The present results do not support to this conclusion. It would be surprising if development went to significant lengths to ensure that the orientation, ocular dominance and spatial frequency maps were structured in such a way as to ensure good coverage, but the retinotopic map developed in such a way as to obliterate this effect. Nevertheless, the impact of the detailed fine structure of the retinotopic map on coverage in the cat remains to be demonstrated.

Relatively little weight can be attached to comparisons of the absolute values of c' across different animals (such as the five studied here). This is because the value depends on estimates of receptive field width, orientation tuning width and retinal magnification factor. Physiological noise, for example, spatially coincident artifacts in the maps, may also artifactually increase the measured value of c' . This is particularly relevant to the spatial-frequency maps, whose amplitude is small, rendering them more susceptible to vascular artifacts. Differences in c' between animals are therefore hard to interpret. Within individual animals, our manipulations suggest that random map superpositions would lead to a worsening of coverage by about 10–30%, which may not seem very large. Improvements of this amount in a sys-

tem whose performance is limited by noise may still be advantageous, however. It is also worth pointing out that our results only show that c' is at a local optimum for each map region studied. It is possible that solutions exist that are better in a global sense, and that they occur in other regions of the same cortical area, or that maps in some individual animals are, fortuitously or otherwise, globally better solutions. Some degree of interanimal variability in c' might therefore be expected, but this would not invalidate the significance of the existence of local optima in each individual patch of cortex.

It is possible that spatially organized maps of other visual features, such as line length, curvature, retinal disparity and so forth may be present in visual cortex. If so, then the calculation of coverage would need to be expanded to take into account these additional feature dimensions. Each additional dimension would degrade coverage²³ although it is not clear *a priori* what an acceptable upper limit for c' might be. This more complex situation, if it exists, does not invalidate the present results, because it is likely that a prerequisite for uniformity of coverage over the complete set of dimensions is uniformity over subsets of them.

Sliding and reflection perturbations were chosen here because they leave the structure of individual maps unchanged and thus allow the coverage constraint to be assessed independently of the continuity constraint. It is important that continuity should remain unchanged in these tests because it is possible to change map structure in ways that would improve coverage. For example, a salt-and-pepper mixing of receptive field properties would result in much better coverage uniformity than is observed in real maps, as would shrinking all the maps by the same amount. However both of these changes disrupt map continuity. A more complete test of the coverage hypothesis would include local rubber-sheet distortions (for example, movements of singularities) done in such a way as to leave continuity unchanged. We cannot rule out the possibility that coverage might be further improved by such perturbations. Devising them would be difficult, however, because there are potentially many ways to define continuity, and a distortion that left continuity unchanged under one measure would not necessarily do so under another. Why continuity is a constraint on cortical map structure is unclear. One possible explanation is that it minimizes the length of axons required to connect cells with similar properties²⁵.

Although the present results show that cortical maps seem to be designed to optimize coverage, they do not directly identify what types of structural relationships would produce this effect. Modeling studies^{16, 20–23} suggest that orthogonal spatial gradient relationships between pairs of maps are one of the ways that good coverage can be achieved. One way of demonstrating such a relationship is to calculate the distribution of intersection angles of the zero-crossings (that is, blob or stripe borders) in pairs of maps and show that it is skewed toward 90°. This has been demonstrated for orientation and spatial frequency, and for orientation and ocular-dominance column boundaries in the cat¹³ and for ocular dominance and orientation in the macaque monkey^{14,15}. Other pairs such as ocular dominance and spatial frequency have not been tested in this way, although it seems likely that such a relationship will be found. A second type of structural relationship that is likely to improve coverage is the placement of orientation singularities in the center of ocular dominance stripes¹⁶. This tendency has been observed in macaque monkey¹⁴ as well as in the cat¹³ and is likely to be one of the factors that cause coverage to depend on the structural relationships between the maps. The relative periodicities of the different maps also affect coverage uniformity¹⁶. However, as these are global properties of the

different maps, which are unchanged by reflection and sliding, their impact cannot be directly assessed here.

The present results lend support to the class of models of cortical development that implement the principle known as dimension reduction. This includes models based on the elastic net algorithm^{24,25} and the self-organizing feature map algorithm²⁶. These models assume, as a starting point, that the maps found in visual cortex are a solution to the problem of mapping a feature space with many dimensions as smoothly and completely as possible onto a two-dimensional surface (see ref. 27 for review). The implementation of completeness constraints in these models is likely to result in maps in which coverage is close to a local optimum. This is especially likely to be true for the elastic net model, which minimizes an energy term²⁴ that is formally almost identical to the coverage measure used here²². The observation that coverage seems to be close to a local optimum in cat area 17 therefore supports the completeness assumption on which these models are based.

METHODS

Standard optical imaging methods were used to obtain maps from area 17 of five anesthetized, paralyzed, young (2–3 month old) cats. Full details of the experimental methods used, and a more complete presentation of the experimental results from these animals, are given elsewhere¹³. In brief, functional topography maps of ocular dominance, spatial frequency and orientation were obtained using moving, high-contrast square-wave gratings with one of four orientations (0°, 45°, 90° and 135°) and two spatial frequencies (usually 0.2 c/deg and 0.6 c/deg) presented to either the contralateral or ipsilateral eye. To reduce artifacts unrelated to the stimuli, each image was divided by the 'cocktail blank', the sum of the images obtained with all the stimuli. Differential ocular dominance maps were obtained by adding all the single-condition maps (all orientations and spatial frequencies) obtained through the contralateral eye, and dividing the result by the sum of the maps obtained through the ipsilateral eye. Similarly, spatial frequency maps were obtained by dividing the sum of the single-condition maps obtained at one spatial frequency by the sum of the maps obtained at the second spatial frequency. Images were clipped at the 1.5 and 98.5% levels and scaled so that pixel values lay between 0 and 255. The maps were then high-pass filtered with a 30 × 30 kernel to remove hemodynamic noise.

Further preparation for the analysis presented here was as follows. Images were first rotated, if necessary, and then trimmed to select approximately the largest rectangular region falling within the map boundaries. This region was typically about 5 × 2.5 mm (approximately 140 × 70 pixels). Each image was smoothed by convolution with a Gaussian filter with $\sigma = 1$ pixel. Single-condition orientation maps were converted to angle maps by vector averaging. Following this, the ocular-dominance and spatial-frequency maps were segmented into two regions of equal area by application of a threshold equal to the median of the pixel values in the map. For the ocular-dominance map, pixel values falling below threshold were assigned (notional) ocular-dominance map values of -1 , and values above threshold were assigned values of $+1$. The spatial-frequency map was partitioned in the same way. The resulting images are shown in Fig. 1.

Coverage uniformity, denoted by c' , was defined^{16,23} as follows. Let $n_{ij} \in \{-1, 1\}$ be the ocular dominance value of cortical point (i, j) ; $m_{ij} \in \{-1, 1\}$ be the spatial frequency value; θ_{ij} the orientation preference in degrees and x_{ij}, y_{ij} the receptive field position in degrees of visual angle. Let the vector $\mathbf{v} = (n_s, m_s, \theta_s, x_s, y_s)$ represent a particular stimulus, and let the vector $\mathbf{w}_{ij} = (n_{ij}, m_{ij}, \theta_{ij}, x_{ij}, y_{ij})$ represent the center receptive field values of cortical point (i, j) . The total amount of activity, A , evoked in the cortex by a stimulus \mathbf{v} , can be written as

$$A(\mathbf{v}) = \sum_{ij \in C} f(\mathbf{v} - \mathbf{w}_{ij}), \quad (1)$$

where f defines the receptive field of point (i, j) , that is, the amount of activity evoked as a function of the difference between the stimulus value

and the center receptive field values, and the summation is taken over all, or a sufficient number, of the points in the map. Letting $n = n_s - n_{i,j}$, $m = m_s - m_{i,j}$, $\theta = \theta_s - \theta_{i,j}$, $x = x_s - x_{i,j}$ and $y = y_s - y_{i,j}$, the receptive field was defined as

$$f(n, m, \theta, x, y) = g(n)g(m)e^{-\theta^2/2\sigma_\theta^2}e^{-(x^2+y^2)/2\sigma_r^2}$$

$$\text{where } g(n) = \begin{cases} 1 & \text{if } n = 0 \\ 0 & \text{if } n \neq 0 \end{cases} \quad (2)$$

and σ_θ and σ_r define the widths of the orientation tuning curves and receptive field sizes, respectively. The Gaussian shapes of these tuning functions are a good approximation to the tuning functions observed experimentally.

$A(v)$ was calculated for a representative set of stimulus values and then coverage uniformity calculated as

$$c' = \text{standard deviation}(A)/\text{mean}(A) \quad (3)$$

If c' is small, it means that coverage is good, and that the total amount of activity evoked in the cortex by a given combination of stimulus parameters is relatively independent of the particular combination.

Parameter values and other simplifying assumptions required to do the calculations were as follows: orientation tuning width, σ_θ , was set equal to 25° ; a uniform and isotropic retinotopic map was assumed, so that $x_{i,j} = \alpha i$ and $y_{i,j} = \alpha j$. For the calculations, it was simplest to assume $\alpha = 1$ and then set σ_r equal to 12 pixels. This is a physiologically reasonable value for a receptive field width: in the maps, 1 pixel = 0.036 mm, and if a retinal magnification factor of $2^\circ = 1$ mm is assumed¹⁷, then 1° of visual field angle equals about 13.9 pixels, and σ_r scales to about 0.86° (ref. 18). To avoid an artifactual loss of responsiveness to stimuli falling close to the edges of the map, responses were normalized by edge weighting, that is, they were divided by the amount of the total receptive field volume falling within the boundaries of the map. Values of A were calculated for all combinations of $n_s \in \{-1, 1\}$, $m_s \in \{-1, 1\}$, $\theta_s \in \{0, 30, 60, 90, 120, 150\}$, $x_s \in \{1, 5, 9 \dots i_{max}\}$, $y_s \in \{1, 5, 9 \dots j_{max}\}$, where i_{max} and j_{max} are the size in pixels of the rectangularly trimmed map. For an average map, this was about 14,000 stimuli.

ACKNOWLEDGEMENTS

This work was supported by grants from the National Science and Engineering Research Council of Canada and the Medical Research Council of Canada to N.V.S. Additional support came from the Max-Planck-Gesellschaft and the EC Biotech Program (M.H. and T.B.) and from the Wolfson Foundation and German-Israeli Research Foundation to A.G.

RECEIVED 24 NOVEMBER 1999; ACCEPTED 6 JUNE 2000

- Hubel, D. H. & Wiesel, T. N. Receptive fields, binocular interaction, and functional architecture of cat striate cortex. *J. Physiol. (Lond.)* **160**, 106–154 (1962).
- Hubel, D. H. & Wiesel, T. N. Receptive fields and functional architecture of monkey striate cortex. *J. Physiol. (Lond.)* **195**, 215–243 (1968).

- Hubel, D. H. & Wiesel, T. N. Sequence regularity and geometry of orientation columns in the monkey striate cortex. *J. Comp. Neurol.* **158**, 267–294 (1974).
- Hubel, D. H. & Wiesel, T. N. Functional architecture of macaque monkey visual cortex. *Proc. R. Soc. Lond. B* **198**, 1–59 (1977).
- Swindale, N. V., Matsubara, J. A. & Cynader, M. S. Surface organization of orientation and direction selectivity in cat area 18. *J. Neurosci.* **7**, 1414–1427 (1987).
- Shmuel, A. & Grinvald, A. Functional organization for direction of motion and its relationship to orientation maps in area 18. *J. Neurosci.* **16**, 6945–6964 (1996).
- Weliky, M., Bosking, W. H. & Fitzpatrick, D. A systematic map of direction preference in primary visual cortex. *Nature* **379**, 725–728 (1996).
- Tootell, R. B. H., Silverman, M. S., Hamilton, S. L., De Valois, R. L. & Switkes, E. Functional anatomy of macaque striate cortex. III. Color. *J. Neurosci.* **8**, 1569–1593 (1988).
- Thompson, I. D. & Tolhurst, D. J. The representation of spatial frequency in cat visual cortex: a ¹⁴C-2-deoxyglucose study. *J. Physiol. (Lond.)* **300**, 58–59 (1981).
- Tootell, R. B. H., Silverman, M. S. & De Valois, R. L. Spatial frequency columns in primary visual cortex. *Science* **214**, 813–815 (1981).
- Shoham, D., Hübener, M., Schulze, S., Grinvald, A. & Bonhoeffer, T. Spatio-temporal frequency domains and their relation to cytochrome oxidase staining in cat visual cortex. *Nature* **385**, 529–533 (1997).
- DeAngelis, G. C., Ghose, G. M., Ohzawa, I. & Freeman, R. D. Functional micro-organization of primary visual cortex: receptive field analysis of nearby neurons. *J. Neurosci.* **19**, 4046–4064 (1999).
- Hübener, M., Shoham, D., Grinvald, A. & Bonhoeffer, T. Spatial relationships among three columnar systems in cat area 17. *J. Neurosci.* **17**, 9270–9284 (1997).
- Bartfeld, E. & Grinvald, A. Relationships between orientation-preference pinwheels, cytochrome oxidase blobs, and ocular dominance columns in primate striate cortex. *Proc. Natl. Acad. Sci. USA* **89**, 11905–11909 (1992).
- Obermayer, K. & Blasdel, G. G. Geometry of orientation and ocular dominance columns in monkey striate cortex. *J. Neurosci.* **13**, 4114–4129 (1993).
- Swindale, N. V. Coverage and the design of striate cortex. *Biol. Cybern.* **65**, 415–424 (1991).
- Albus, K. A quantitative study of the projection area of the central and paracentral visual field in area 17 of the cat. I. The precision of the topography. *Exp. Brain Res.* **24**, 159–179 (1975).
- Hetherington, P. A. & Swindale, N. V. Receptive field and orientation scatter studied by tetrode recording in cat area 17. *Visual Neurosci.* **16**, 637–652 (1999).
- Das, A. & Gilbert, C. D. Distortions of visuotopic map match orientation singularities in primary visual cortex. *Nature* **387**, 594–598 (1997).
- Obermayer, K., Ritter, H. & Schulten, K. A principle for the formation of the spatial structure of cortical feature maps. *Proc. Natl. Acad. Sci. USA* **87**, 8345–8349 (1990).
- Obermayer, K., Blasdel, G. G. & Schulten, K. Statistical-mechanical analysis of self-organization and pattern formation during the development of visual maps. *Phys. Rev. A* **45**, 7568–7589 (1992).
- Swindale, N. V. The development of topography in the visual cortex: a review of models. *Network* **7**, 161–247 (1996).
- Swindale, N. V. How many maps are there in visual cortex? *Cereb. Cortex* (in press).
- Durbin, R. & Willshaw, D. J. An analogue approach to the travelling salesman problem using an elastic net method. *Nature* **326**, 698–699 (1987).
- Durbin, R. & Mitchison, G. A dimension reduction framework for understanding cortical maps. *Nature* **343**, 644–647 (1990).
- Kohonen, T. Self-organized formation of topologically correct feature maps. *Biol. Cybern.* **43**, 59–69 (1982).
- Swindale, N. V. Elastic nets, travelling salesmen and cortical maps. *Curr. Biol.* **2**, 429–431 (1992).