

synaptic deficits long before the formation amyloid deposits, possibly due to overexpression of APP or higher levels of soluble A β . Overexpression of human APP leads not only to increased levels of A β but also to an increase in a variety of APP C- and N-terminal fragments that can have biological effects. Whether soluble A β plays a role in these early deficits found in these models prior to A β aggregation has not been definitively answered. The authors report that their new model has no obvious behavioral abnormalities, although the behavioral abnormalities observed in other AD mouse models thus far require testing in complex learning and memory paradigms to fully appreciate (Janus and Westaway, 2001). Should the A β 42-overexpressing mice develop age- and amyloid-dependent cognitive decline, they will prove extremely valuable in further experiments to specifically test the role of soluble and insoluble forms of A β in both neurodegeneration and cognition. Those interested in A β metabolism should find these mice very useful, and we anxiously await the studies that will stem from the use of these new mouse models.

John D. Fryer and David M. Holtzman

Department of Neurology and Molecular Biology
and Pharmacology

Hope Center for Neurological Disorders
Washington University School of Medicine
St. Louis, Missouri 63110

Selected Reading

- Bales, K.R., Dodart, J.C., DeMattos, R.B., Holtzman, D.M., and Paul, S.M. (2002). *Mol. Interv.* 2, 363–375.
- Fryer, J.D., Simmons, K., Parsadanian, M., Bales, K.R., Paul, S.M., Sullivan, P.M., and Holtzman, D.M. (2005). *J. Neurosci.* 25, 2803–2810.
- Glenner, G.G., and Wong, C.W. (1984). *Biochem. Biophys. Res. Commun.* 120, 885–890.
- Gravina, S.A., Ho, L., Eckman, C.B., Long, K.E., Otvos, L., Younkin, L.H., Suzuki, N., and Younkin, S.G. (1995). *J. Biol. Chem.* 270, 7013–7016.
- Hardy, J., and Selkoe, D.J. (2002). *Science* 297, 353–356.
- Hsiao, K., Chapman, P., Nilsen, S., Eckman, C., Harigaya, Y., Younkin, S., Yang, F., and Cole, G. (1996). *Science* 274, 99–102.
- Iwatsubo, T., Odaka, A., Suzuki, N., Mizusawa, H., Nukina, N., and Ihara, Y. (1994). *Neuron* 13, 45–53.
- Janus, C., and Westaway, D. (2001). *Physiol. Behav.* 73, 873–886.
- Kim, S.H., Wang, R., Gordon, D.J., Bass, J., Steiner, D.F., Lynn, D.G., Thinakaran, G., Meredith, S.C., and Sisodia, S.S. (1999). *Nat. Neurosci.* 2, 984–988.
- McGowan, E., Pickford, F., Kim, J., Onstead, L., Eriksen, J., Yu, C., Skipper, L., Murphy, M.P., Beard, J., Das, P., et al. (2005). *Neuron* 47, this issue, 191–199.
- Strittmatter, W.J., and Roses, A.D. (1995). *Proc. Natl. Acad. Sci. USA* 92, 4725–4727.
- Vidal, R., Frangione, B., Rostagno, A., Mead, S., Revesz, T., Plant, G., and Ghiso, J. (1999). *Nature* 399, 776–781.

DOI 10.1016/j.neuron.2005.07.002

Cortical Maps: Where Theory Meets Experiments

Primary visual cortex (V1) has remarkably systematic functional maps. One commonly used class of computational models proposes that such maps are generated by a mechanism that projects the multiple dimensions of neuronal responses smoothly onto the two dimensions of cortex. In this issue of *Neuron*, Mriganka Sur and colleagues find a close match between such model predictions and measurements from ferret V1.

In most species of highly visual animals, neuronal response properties in primary visual cortex (V1)—the first cortical stage in the mammalian visual pathway—are laid out in remarkably systematic, periodic patterns on the cortical surface (Hubel and Wiesel, 1977; Hubener et al., 1997; but see Van Hooser et al., 2005). When tested with simple visual stimuli such as short line segments or small patches of grating-like textures, neurons at any given point on V1 can be characterized by their tuning for the different features that define these stimuli. Thus, a given neuron responds preferentially to stimuli appearing at a particular location in visual space, with a particular orientation, a particular spatial frequency (the spacing of the bars for a grating stimulus), ocular dominance (preference for stimuli being shown to the right versus the left eye), direction of movement, color contrast, etc. (The exact set of V1 tuning properties differs for different species.) As we move across the cortical surface, we find smooth interrelated maps of all these neuronal response properties. Neighboring points in V1 are tuned to nearby points in visual space, but they are also tuned to nearby values for all the other stimulus properties. While the mapped stimulus location changes monotonically as we move across V1, all the other stimulus properties cycle periodically over their full respective ranges, each map showing similar length scales of periodicity (Figure 1).

Over the years, this observation of the regularity of V1 maps has attracted a number of explanatory principles and theoretical models, most invoking patterned neuronal activity, but some also involving molecular messages (Swindale, 1996; Willshaw and Von Der Malsburg, 1979; Koulakov and Tsiganov, 2004; the last two deal with superior colliculus, but the principle is applicable to V1). In this issue of *Neuron*, Mriganka Sur and colleagues have made a valuable contribution to this literature. Their study (Yu et al., 2005) combines computational modeling and physiological measurements to provide a comprehensive test of one common class of purely activity-dependent cortical models—the so-called “dimension reduction” class, which proposes that V1 maps are formed through a mechanism that maps the multiple dimensions of neuronal response properties onto the two-dimensional cortical surface as smoothly as possible.

It has long been proposed that the organizing principle underlying the orderliness of V1 maps is one of trying to strike a balance between maximizing coverage

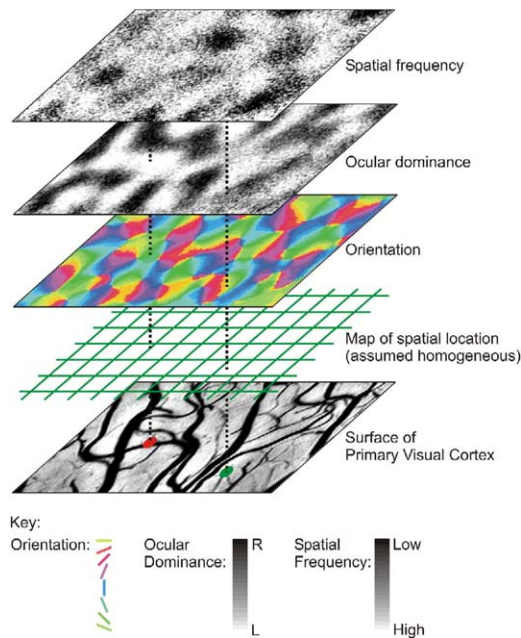


Figure 1. Schematic of Maps of Spatial Location, Orientation, Ocular Dominance, and Spatial Frequency over a Patch of Primary Visual Cortex

Neurons at the cortical position marked in red would be tuned to the particular stimulus location, orientation, ocular dominance, and spatial frequency identified by the dotted line leading up from the red spot. Neurons in the green spot would be tuned to the different set of values indicated by its dotted line. The key gives color codes identifying tunings for the different stimulus features.

and minimizing wiring length within cortex. Maximizing coverage means making sure that cortex can respond to all possible stimulus properties at every point in visual space. This quantifies our familiar perceptual experience that we seem to sample the visual world uniformly everywhere—it is not the case, for example, that we find ourselves more sensitive to red than to green in one patch of our visual world, or for detecting vertical versus horizontal lines in another patch. In modeling terms, this means that each patch of the cortical map of space should contain equal numbers of neurons tuned to every possible combination of stimulus properties. The other principle—of minimizing wiring length—is equivalent to demanding smoothness of mapping. Namely, that neurons with similar response properties should be placed as close to each other as possible on the cortical surface. A hypothetical example illustrates how these two principles could act in opposition to each other. For example, if all neuronal response properties were scattered randomly over V1 so that every patch of cortex had a fine salt-and-pepper mixture of neurons tuned to every possible stimulus property, the resultant map would have near-perfect coverage, but lousy smoothness.

With just two independent properties—orientation and ocular dominance—it is possible to have a V1 map that is both smooth and provides uniform coverage by having the two individual maps cross each other at right angles. This was proposed by Hubel and Wiesel




Figure 2. “Ice Cube” Model for the Joint Map of Orientation and Ocular Dominance on V1 Compared to a Model with Very Poor Coverage

(Left) Ice cube model: schematic maps of orientation and ocular dominance, intersecting each other at right angles, superimposed on a map of spatial location. The black lines on the spatial map indicate the boundaries of each repeating “module.” A shift of one module corresponds to a shift to a neighboring (partially overlapping) spatial position. Since every stripe of orientation crosses all stripes of ocular dominance equally, within a module, and every stripe of ocular dominance crosses every stripe of orientation equally, each module has a complete set of all possible combinations of values of orientation and ocular dominance. Thus, coverage is maximally uniform. (Right) If orientation and ocular dominance were to run parallel to each other, any given value of orientation would be tied to only a given value of ocular dominance. Thus, V1 would represent only a small set of combinations of orientation and ocular dominance would be very sparse.

(1977) in the “ice cube” model they posited for the structure of V1 (Figure 2). With more than two dimensions of neuronal response properties to map onto cortex, it is no longer possible to simultaneously have uniform coverage and perfect smoothness. A number of computational models have extended the ice cube idea of joint, intersecting maps to multiple dimensions by using “dimension-reduction” algorithms based on the principle of trying to map the multiple dimensions of neuronal response onto the two dimensions of cortex in the smoothest manner possible. This process can also be thought of as twisting and folding the two-dimensional model cortical surface multiple times into the multi-dimensional space of visual responses so as to touch all points in the higher dimensional space. These models have used variants of an elastic net or traveling salesman algorithm—or more biologically plausible “Hebbian” models of development.

The model V1 maps predicted by these dimension-reduction algorithms have properties that appear as logical extensions of the ice cube idea. One sees intuitively that, for any pair of maps, the most uniform coverage would result when the maps cross at right angles (Figure 2). With multiple superimposed response maps, it is of course geometrically impossible for all the maps to cross each other at right angles everywhere. But the constraint of trying to provide uniform coverage does predict that pairs of maps would be strongly biased toward being locally orthogonal, producing patches of ice cube-like structure, and that each map would develop similar scales of periodicity. The other constraint—trying to maximize smoothness—leads to a complemen-

tary set of predictions. With multiple maps it is not possible for all maps to be equally smooth everywhere and yet cross each other at right angles. So what happens is that each map shows regions of high smoothness, i.e., low gradient, or low rate of change, interspersed with regions of low smoothness, i.e., high gradient, or rapid change, or even singularities. Here, the mathematical constraint of trying to maximize the overall smoothness of the full V1 map makes another strong prediction. Namely, that individual maps will arrange themselves so as to avoid, as far as possible, each other's regions of high gradient. It is easy to see why this is the case. Since bumpiness in each individual map is unavoidable, what the algorithm does is to lay out the maps so that highly bumpy portions of one map are superimposed over the smoothest possible regions of other maps, thereby minimizing, to the extent possible, the average bumpiness of the joint multidimensional V1 map (Durbin and Mitchison, 1990). There have been very few experimental tests of these model predictions, however, and even these few appeared to contradict the models. In particular, where the models predicted that maps of space and of orientation would be bumpy, with their bumpy regions anticorrelated as explained above, experiments showed the two sets of bumpinesses to be either positively correlated (Das and Gilbert, 1997) or to bear no relation to each other (Buzas et al., 2003).

Now Mriganka Sur and colleagues have elegantly combined computational algorithms with physiological measurements to closely examine and, largely, validate these models of cortical organization (Yu et al., 2005). The authors considered the test case of ferret V1 where the map of space is sharply anisotropic, the cortical magnification factor (the cortical distance covered per degree of visual space) for the elevation ("vertical") axis of visual space being three to five times higher than for the azimuth ("horizontal"). The authors used a Kohonen algorithm (Kohonen, 2001) for generating self-organizing maps to map the six-dimensional space of V1 response properties, i.e., stimulus location (x , y), orientation (θ : the angle, and q : the strength of orientation preference), ocular dominance, and spatial frequency on to the two-dimensional model cortex. In a parallel physiological study of ferret V1 using optical imaging and electrode recordings, the authors found close agreement with their model predictions. In the model, anisotropy imposed in the spatial map led to orthogonal anisotropies in the maps of orientation, ocular dominance, and spatial frequency, a pattern reproduced in real ferret V1. Regions of high gradient magnitude ("bumpiness") in the different maps (other than of space) tended to avoid each other, as expected from the model constraint of trying to maximize smoothness. Having the enforced spatial anisotropy also led to a pattern of map crossings distinct from those in earlier theoretical models that also try to maximize coverage; the remaining maps no longer showed any overall tendency to cross each other at right angles (unlike the predictions from models where the spatial map was allowed to be isotropic), but only did so where patches of high gradient magnitude overlapped each other. Finally, the authors reproduced an earlier model result—backed up, here, with electrode recordings—showing

that adding extra response dimensions to the map (six, in this case) essentially eliminated any correlations between inhomogeneities in the maps of space and of orientation.

Why are these results important? First, of course, this work provides the most exhaustive test to date of the cortical mapping principle of attempting to maximize uniformity of coverage while at the same time maximizing local smoothness. And it is always pleasing to see an elegant biological theory being backed up with experimental data. The idea of uniform coverage, as already mentioned, is rooted in our perceptual experience of the uniformity of our visual abilities in space. The competing idea of maximal local smoothness feeds into another principle that is believed to underlie sensory processing, namely that much of neuronal processing is mediated through lateral interactions or lateral pooling between similar visual features. This includes the lateral suppression that generates center-surround receptive fields in the retina, the lateral pooling that generates oriented V1 receptive fields from their unoriented inputs, on through the lateral suppression leading to complex object selectivity in higher visual centers (Wang et al., 2000). Having cortical maps that are maximally smooth means that neurons that happen to be cortical neighbors are also neighbors in visual feature space; thus, any neural computation that depends on lateral comparison or lateral pooling among nearby visual features can be accomplished with just anatomically local comparisons mediated by short-range axons and dendrites. Such a process achieves the biologically important principle of minimizing the cortical wiring length involved in local computations, thus minimizing attendant metabolic costs and maximizing neuronal processing speed. Of course, for the full panoply of visual processing, such local computations are also complemented with specific long-range interactions as well as feedback from other cortical areas.

The current results also feed into long-standing debates on the relative importance of neural activity versus molecular mechanisms in determining cortical architecture. In this context, it is worth noting that the computational model used by Yu et al. did not explicitly use a cost function that attempted to maximize smoothness and uniformity of coverage. Rather, they used an activity-dependent or Hebbian algorithm where their model neurons started with a random set of response properties and were "shown" complete sets of visual stimuli; model synaptic strengths were updated each time using a rule that depended both on the activity simulated in the neurons and a model of smooth local interactions in cortex. This strategy has been used by many modelers, being both a biologically plausible simulacrum of development and much more computationally tractable than an explicit optimization of coverage and smoothness. Maps thus generated are known to be remarkably close to optimal in their coverage and smoothness (Durbin and Mitchison, 1990). But the model mechanism explicitly necessitates neural activity, with neuronal responses simultaneously driven by spatial location, orientation, ocular dominance, and spatial frequency so that the different maps "see" and interact with each other. It is noteworthy that the map

feature of spatial location is quite distinct from all other features in the authors' formulation, with an imposed anisotropy that then determines the patterns of all the other maps. This would suggest an underlying molecular marker rather than a purely activity-driven mechanism, although it is possible that constraints posed by the overall shape of V1 may lead to similar outcomes. Some recent work suggests that spatial mapping in V1, in particular, is driven by activity rather than molecular markers (Eglen et al., 2003). On the other hand, while a large body of literature suggests that ocular dominance is determined by neural activity, some recent work suggests that a molecular message may also be involved (Crowley and Katz, 2000). It would be very interesting to see whether insights from development could be used to inform models of cortical organization and whether formal, testable models can be developed that specifically discriminate between mechanisms that demand molecular markers, or activity alone, or some particular interaction between the two so as to gain further understanding of this important process.

Aniruddha Das

Columbia University Center for Neurobiology
and Behavior
1051 Riverside Drive, Unit 87
New York, New York 10032

Selected Reading

- Buzas, P., Volgushev, M., Eysel, U.T., and Kisvarday, Z.F. (2003). *Eur. J. Neurosci.* 18, 957–968.
- Crowley, J.C., and Katz, L.C. (2000). *Science* 290, 1321–1324.
- Das, A., and Gilbert, C.D. (1997). *Nature* 387, 594–598.
- Durbin, R., and Mitchison, G. (1990). *Nature* 343, 644–647.
- Eglen, S.J., Demas, J., and Wong, R.O.L. (2003). *Neuron* 40, 1053–1055.
- Hubel, D.H., and Wiesel, T.N. (1977). *Proc. R. Soc. Lond. B. Biol. Sci.* 198, 1–59.
- Hubener, M., Shoham, D., Grinvald, A., and Bonhoeffer, T. (1997). *J. Neurosci.* 17, 9270–9284.
- Kohonen, T. (2001). *Self-Organizing Maps* (New York: Springer).
- Koulakov, A.A., and Tsigankov, D.N. (2004). *BMC Neurosci.* 5, 30.
- Swindale, N.V. (1996). *Network. Comput. Neur. Syst.* 7, 161–247.
- Van Hooser, S.D., Heimel, J.A.F., Chung, S., Nelson, S.B., and Toth, L.J. (2005). *J. Neurosci.* 25, 19–28.
- Wang, Y., Fujita, I., and Murayama, Y. (2000). *Nat. Neurosci.* 3, 807–813.
- Willshaw, D.J., and Von Der Malsburg, C. (1979). *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 287, 203–243.
- Yu, H., Farley, B.J., Jin, D., and Sur, M. (2005). *Neuron* 47, this issue, 267–279.

DOI 10.1016/j.neuron.2005.07.004