

Imaging orientation selectivity: decoding conscious perception in V1

Geoffrey M Boynton

In V1, neurons preferring similar orientations are grouped in columns too small to be resolved by conventional fMRI. Two studies circumvent this limitation by using algorithms to recognize patterns of activation across a large area. This new trick allows the authors to distinguish responses to different orientations in human V1 and to study its contribution to conscious perception.

Functional MRI is currently the best non-invasive tool for measuring human brain activity at below-centimeter resolution¹. This spatial resolution is ideal for detecting and studying entire cortical maps, such as those in primary visual cortex (V1), but it is still a long way from measuring the responses of individual neurons. Electrophysiological recording and optical imaging in animals have shown that V1 neurons preferring similar orientations form columns about 500 μm across. Measuring the response from these homogenous clusters would be an important step toward increasing the spatial resolution of functional imaging, but even these have been too fine to be routinely resolved by fMRI—until now. In this issue, two studies^{2,3} show that it is possible to estimate the orientation of a stimulus from the pattern of fMRI responses it produces in V1. This enables us for the first time to study how this fundamental form of visual information is represented in human cortex.

Over the past decade, the spatial resolution of fMRI has been gradually improving through technical advances such as increased magnetic field strength, better receiving coils and more reliable gradients and amplifiers. But direct measurements below the resolution of a millimeter can be obtained only with massive amounts of signal averaging from a carefully selected group of subjects⁴. Attempts to study orientation selectivity with fMRI using indirect methods such as adaptation have also met with difficulty⁵. In the two current papers^{2,3}, the authors took an alternative approach to measuring orientation selectivity through a clever data analysis

trick. Remarkably, both groups used traditional high-field (3-T) fMRI data acquisition methods, which means that evidence of signals at the columnar level may already be available in all of our existing fMRI data sets.

How did they measure such a fine spatial structure without special equipment? In a simulated orientation 'pinwheel' map (Fig. 1a), different colors indicate the preference of orientation columns on a 9×9 mm region of the cortical surface⁶. Each of the nine 3×3 -mm fMRI voxels, shown as black squares, contains a broad range of orientation preferences. On closer inspection, however, some voxels contain more columns of one orientation preference than another (Fig. 1b). Although all voxels respond to all orientations, voxels clearly have a variable response across orientations. This variability is evidence of high spatial frequency information, even if the measurement tool is sampling at a lower frequency, a phenomenon known as 'aliasing' in the signal processing literature.

How can these weak biases in fMRI responses be used to predict the orientation of a subsequently viewed stimulus? The trick is to first associate a range of test stimulus orientations with their patterns of fMRI responses. There are a variety of ways of doing this. Kamitani and Tong² use a 'linear support vector machine' that creates 'classifiers' for each stimulus orientation by summing weighted responses across voxels, obtaining optimal weights during a training period. When oriented stimuli were presented after training, the response for each classifier to the fMRI image was calculated, and the actual stimulus was estimated from the classifier with the largest response estimates. Haynes and Rees³ used a 'linear discriminant analysis' method in which a Bayesian calcula-

tion found the orientation that was most likely to have induced a given pattern of responses.

Because fMRI responses are noisy, many voxels must be incorporated to obtain reason-




Figure 1 Patterns of orientation-selective responses measured with fMRI. (a) Synthetic orientation tuning data generated by band-pass filtering random orientation values⁶. The black squares represent 3×3 mm fMRI voxels. (b) Histograms showing the proportion of selectivity inside each voxel to each of the eight orientations shown below. This shows how different stimulus orientations produce slightly different patterns of responses in V1. Algorithms such as those used in the current studies^{2,3} can estimate from these responses the orientation of a subsequently presented stimulus.

Geoffrey M. Boynton is at the Salk Institute, 10010 North Torrey Pines Road, La Jolla, California, USA. email: boynton@salk.edu

able accuracy. Kamitani and Tong² show that in V1, the pattern of responses across a few hundred voxels can predict with nearly perfect accuracy which orientation was shown. Haynes and Rees³ found an impressive average accuracy of 80% from the response to a single two-second stimulus presentation.

It is not known how far we can go with this method. If too many columns fit inside the average voxel, then voxels should have nearly equal responses to all orientations. On the other hand, increasing voxel size dramatically increases the reliability of the signal from each voxel, so results may not be highly dependent on voxel size. This is not the first time that this method has been used to measure stimulus selectivity at a sub-voxel level. Haxby and colleagues⁷ used a similar method to show that the pattern of fMRI responses in the ventral temporal cortex could predict the category of an object (such as faces, cats, houses and shoes) that was being shown. It remains to be seen if this method will be successful at studying columnar structures in other sensory areas, such as those in auditory cortex, or even motor areas of the human brain.

Although orientation selectivity is ubiquitous in mammalian V1, the ability to examine it in humans provides the opportunity to study the role of this area in conscious visual perception. For example, if V1 is so early in the processing stream, does orientation selectively occur automatically, or can it be affected by the will of the subject? If it does occur automatically, does it always lead to a conscious visual percept? Both papers went well beyond verifying that orientation-selectivity can be measured in human V1 and have given important insights into these fundamental questions.

Kamitani and Tong² were able to predict the orientation that a subject was thinking about. Subjects were instructed to attend to one of two orthogonal orientations forming a 'plaid' stimulus. The physical stimulus did not

change across trials; only the instructions to the subjects did. The authors found that they could predict which orientation the subject was attending with 80% accuracy on a trial-by-trial basis. This suggests that attending to one orientation and ignoring the other changed the pattern of fMRI responses enough to look like only the attended orientation was presented. This is similar to the effects seen in single-neuron recordings when multiple stimuli are presented within the receptive field of a cell^{8,9}. But this is the first such evidence in humans, and the first in the primary visual cortex.

Haynes and Rees³, on the other hand, were able to predict the orientation of a stimulus that subjects could not see. Rapidly alternating a stimulus of co-oriented lines with a multi-oriented 'masking' stimulus renders the oriented stimulus invisible¹⁰. Subjects can clearly see the alternation between the oriented and the masking stimulus, but cannot determine the angle of the oriented stimulus. However, Haynes and Rees³ used the fMRI response to this alternating stimulus to predict the masked stimulus even though the subjects could not tell which orientation was being shown.

These two studies have interesting implications about the role of V1 in consciousness. Being just two synapses away from the eye, V1 is usually considered an early visual area. Early visual areas tend to represent properties of the physical stimulus, whereas visual areas later in the processing stream seem to hold our conscious percept, or our brain's interpretation of the stimulus^{5,11}. The finding by Haynes and Rees³ is consistent with this idea, and supports the theory¹² that we are not consciously aware of all of the processing going on in V1.

But is V1 a passive feed-forward image processing machine that is unaffected by what the observer was thinking or doing? It seems not. Allocating attention to a particular location in space (without moving the eyes) can affect fMRI and electrophysiological responses in V1

(refs. 13–15). Kamitani and Tong² show that allocation to the feature of an attended stimulus can affect V1 responses as well. Far from being unavailable to consciousness, V1 responses can be examined as if part of a 'mind-reading' exercise on their subjects, the authors suggest².

These two papers show that V1 appears to be neither at the beginning nor at the end of visual processing. Considering the array of feedback connections from higher visual areas to V1 and the feed-forward loops back through the geniculate to V1, perhaps it is unwise to consider any cortical visual area as 'early'. Instead, the distinction between early and late processing may all be in the timing. It is likely that the response to the 'invisible' stimulus by Haynes and Rees³ occurs early in the temporal response to the stimulus, and the modulations of the orientation-selective response by attention found by Kamitani and Tong² occur 100–200 ms later, as has been seen for spatial attention¹⁴. The answer may already be there in the data, waiting for another clever algorithm to tease it out.

- Engel, S.A., Glover, G.H. & Wandell, B.A. *Cereb. Cortex* **7**, 181–192 (1997).
- Kamitani, Y. & Tong, F. *Nat. Neurosci.* **8**, 679–685 (2005).
- Haynes, J. & Rees, G. *Nat. Neurosci.* **8**, 686–691 (2005).
- Cheng, K., Waggoner, R.A. & Tanaka, K. *Neuron* **32**, 359–374 (2001).
- Boynton, G.M. & Finney, E.M. *J. Neurosci.* **23**, 8781–8787 (2003).
- Rojer, A.S. & Schwartz, E.L. *Biol. Cybern.* **62**, 381–391 (1990).
- Haxby, J.V. *et al. Science* **293**, 2425–2430 (2001).
- Moran, J. & Desimone, R. *Science* **229**, 782–784 (1985).
- Reynolds, J.H., Chelazzi, L. & Desimone, R. *J. Neurosci.* **19**, 1736–1753 (1999).
- Macknik, S.L. & Livingstone, M.S. *Nat. Neurosci.* **1**, 144–149 (1998).
- Treue, S. *Trends Neurosci.* **24**, 295–300 (2001).
- Crick, F. & Koch, C. *Nature* **375**, 121–123 (1995).
- Gandhi, S.P., Heeger, D.J. & Boynton, G.M. *Proc. Natl. Acad. Sci. USA* **96**, 3314–3319 (1999).
- Martinez, A. *et al. Nat. Neurosci.* **2**, 364–369 (1999).
- Somers, D.C., Dale, A.M., Seiffert, A.E. & Tootell, R.B. *Proc. Natl. Acad. Sci. USA* **96**, 1663–1668 (1999).

Finding the G spot on fusion machinery

Jane Sullivan

Activation of G protein-coupled receptors can inhibit secretion of neurotransmitters and hormones. Two recent reports in *Nature Neuroscience* show that this inhibition is due to G $\beta\gamma$ binding to SNAP-25, directly blocking the vesicle fusion machinery.

Jane Sullivan is at the Department of Physiology and Biophysics, University of Washington School of Medicine, Box 357290, Seattle, Washington 98195, USA.
e-mail: jmsull@u.washington.edu

Activation of presynaptic G protein-coupled receptors (GPCRs) by ligands such as GABA, glutamate, serotonin or adenosine is a powerful negative feedback mechanism for modulating transmission at synapses throughout

the brain^{1,2}. For example, glutamate released from hippocampal mossy fibers during high-frequency trains of action potentials activates presynaptic G protein-coupled metabotropic glutamate receptors and inhibits subsequent