

Fig. 1. Generation of cholinergic neurons from primed fetal human neural stem cells in the medial septum of adult rat brain. **(a)** Unprimed cells remained largely undifferentiated, shown here as yellow-orange cells as a result of double labeling of red nestin and green GFP. ChAT staining was negative in these cells (not shown). **(b)** In contrast, primed hNSC cells did not stain for nestin (not shown) but did stain for ChAT, showing differentiation. Yellow-orange cells are double-labeled with ChAT (red) and GFP (green). Note that the hNSC-derived ChAT cells give rise to axons. Scale bar, 50 μ m.

retrophic factors. Alternatively, primed, but uncommitted, hNSCs could be instructed by the local environment to adopt specific fates. There could also be 'transdifferentiation' of committed precursors, at least with respect to neurotransmitter phenotype¹². Alternatively, some combination of all these mechanisms may be involved. Unfortunately, the present data do not allow us to distinguish between these possibilities

Laminin seems to be fundamental for the priming effect. Previous studies using a bFGF/heparin cocktail pretreatment before implantation generated far fewer neurons than observed in the current study^{9,10}. Wu *et al.*⁴ report that culturing on laminin spreads the hNSCs evenly out of the neurospheres, and speculate that this may expose the cells more evenly to the bFGF/heparin mix. The effect could also result from an

because implanted animals were only examined for a few weeks after implantation.

Another puzzle is the molecular nature of the priming. How exactly is this mix of bFGF, heparin and laminin causing such a dramatic and diverse induction of neural differentiation? Wu *et al.*⁴ suggest that heparin acts by potentiating FGF binding to its receptor, a well documented effect¹³. Heparin sulfate proteoglycans can also regulate a variety of other signaling pathways¹⁴. One way to determine if heparin's priming effect is mediated primarily by regulating the FGF receptor would be to ask whether the effects of heparin can be mimicked by substantially increasing the dose of bFGF.

inductive role of laminin-activated integrin signaling¹⁵. Future studies can address the molecular mechanism of priming by culturing on different adhesive substrates, and by directly manipulating integrin signaling using constitutively active versions of focal adhesion kinase, paxillin or c-src.

1. Cassidy, R. & Frisen, J. *Nature* **412**, 690–691 (2001).
2. Bjorklund, L. M. *et al. Proc. Natl. Acad. Sci. USA* **99**, 2344–2349 (2002).
3. Kim, J. H. *et al. Nature* **418**, 50–56 (2002).
4. Wu, P., Gu, Y., Huang, L.-Y. M. & Tarasenko, Y. *Nat. Neurosci.* **5**, 1271–1278 (2002).
5. Doering, L. C. & Snyder, E. Y. *J. Neurosci. Res.* **61**, 597–604 (2000).
6. Cao, Q. L. *et al. Exp. Neurol.* **167**, 48–58 (2001).
7. Rubio, F. J. *et al. Mol. Cell. Neurosci.* **16**, 1–13 (2000).
8. Englund, U. *et al. Exp. Neurol.* **173**, 1–21 (2002).
9. Fricker, R. A. *et al. J. Neurosci.* **19**, 5990–6005 (1999).
10. Ostefeld, T. *et al. Exp. Neurol.* **164**, 215–226 (2000).
11. Reubinoff, B. E. *et al. Nat. Biotechnol.* **19**, 1134–1140 (2001).
12. Patterson, P. H. & Chun, L. L. *Dev. Biol.* **56**, 263–280 (1977).
13. Yayon, A. *et al. Cell* **64**, 841–848 (1991).
14. Perrimon, N. & Bernfield, M. *Nature* **404**, 725–728 (2000).
15. Schwartz, M. A. & Ginsberg, M. H. *Nat. Cell Biol.* **4**, E65–E68 (2002).

Multidimensional space: the final frontier

Shimon Edelman

A new approach to analyzing the responses of V4 cortical neurons to objects suggests how the brain could represent a wide variety of shapes with a limited number of components.

The most general approach to understanding the challenges facing perceptual systems is to cast them as problems in data analysis. For theorists, this perspective immediately raises the question of the dimensionality of the data, which strongly influences the appropriate mathematical approaches to making sense of it.

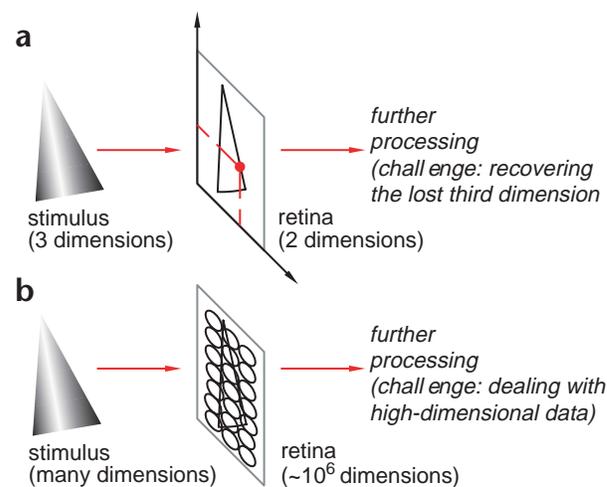
The author is in the Department of Psychology, 232 Uris Hall, Cornell University, Ithaca, New York 14853-7601, USA.
e-mail: se37@cornell.edu

Surprisingly, there is little agreement on this crucial issue in the cognitive sciences (Fig. 1). On the one hand, mathematicians and neuroanatomists know that the responses of N neurons span an N -dimensional measurement space^{1,2}. On the other hand, psychologists tend to frame the problem in terms of recovering the three-dimensional structure of the world from the retinal projection, which they describe as two-dimensional (for example, ref. 3, p. 146). Treating the retinal measurement space as two-dimensional only

makes sense if the stimulus is measured one point at a time—an assumption that is false because an image projected onto the retina is not coded in terms of discrete labeled points. Rather, the millions of photoreceptors all respond together, and their number is what determines the dimensionality of the signal with which the rest of the visual system must contend.

Neurophysiologists, who still collect most of their data by recording from one neuron at a time, face the dual challenge of resisting the temptation to oversimplify inherently multidimensional neural representations, and of transforming knowledge of single-cell responses into a coherent understanding of the behavior of groups of neurons. Thus, it is doubly satisfying to see how a study of visual shape selectivity in the monkey cortical area V4 by Pasupathy and Connor⁴ in this issue deftly rides out these troubled waters, avoiding both the Scylla of erroneous preconceptions and the Charybdis of muddled interpretation.

Fig. 1. The complex issue of dimensionality in vision. (a) A typical textbook notion of what happens at the front end of the human visual system stresses the loss of depth information and the need to recover it from the retinal representation. According to ref. 3, a computational approach to perception suggests that our brains compute 3D perceptual models of the environment, based on information from the 2D sensory receptors in our retinas. (b) The real challenge facing any mammalian visual system is how to deal with the extremely high nominal dimensionality of the retinal output¹¹. The actual dimensionality may be different, depending on the ensemble of data at hand; in fact, if the actual dimensionality of the visual signal were not much lower than its nominal dimensionality (which for humans runs in the millions, if one counts the number of axons in the optic nerve), fundamental tasks such as perceptual learning and generalization would be impossible¹².



The authors started by recording the responses of 109 neurons to parametrically controlled silhouette stimuli resembling raindrops, peanuts or tree leaves. Most researchers studying neural population coding would at this stage proceed to estimate the representation of some scalar characteristic of the stimulus, such as its size or orientation. In contrast, Pasupathy and Connor sought insight into how the ensemble of V4 cells represents the entire stimulus shape—all the bumps and dents of a silhouette, properly positioned. To that end, they first mapped, for each neuron, the shape tuning function defined over a curvature \times angular position parameter space. In this space, a raindrop shape, for example, is coded by a surface that has a single curvature peak at the 'north' angular position; complex silhouettes correspond, in general, to multi-peak surfaces. (Note that two dimensions of shape variation—curvature and position—are represented here simultaneously.) The authors then used basis-function decoding⁵ to estimate the tuning function of the entire population from those of the recorded neurons. In this method, a target function whose value is known only at a small number of sample points is approximated by a weighted average of several basis functions (think of approximating the shape of a hilly lawn by stretching a tarpaulin over several strategically placed airbags inflated to various degrees). When Gaussian basis function approximation is applied to the problem at hand, the potentially multi-peak population response surface emerges as a superposition of several Gaussian 'bumps' placed at the locations where data are available.

Knowledge of the full parametric response surface of an ensemble of neurons allowed the authors to predict the representation of novel silhouettes not included in the original set of 49 stimuli. As in any application of basis function approximation, the quality of the outcome depends on the number and 'spacing' of the data samples (shapes used as stimuli,

and neurons from which recordings are made) and on the suitability of the chosen basis functions (the fidelity with which the response of an individual neuron can be described by a Gaussian). Pasupathy and Connor are careful to separate those of their results that may be affected by these factors, such as the seemingly discrete peaks in the estimated response surface, from those that are not, such as the information that the ensemble response conveys about the geometry of the stimulus. They show how the shape of the stimulus can be reconstructed from the neural code as a superposition of continuous, parameterized components (stressing that visualizing the population code in this manner is not intended as a hypothesis of how the visual system uses such codes). This finding suggests how higher stages of cortical processing could, in principle, form an open-ended representation of a potentially infinite variety of structured silhouette shapes in terms of a small number of basic components, namely, curvature peaks at several angular positions.

The approach to structural representation outlined by Pasupathy and Connor disposes both with the assumption that the components of the neural code are alphabet-like, and with the requirement that they be represented independently of their spatial relations. Indeed, the curvature peaks in their ensemble code are graded and overlapping, and each of them is well-localized along the angular position dimension (and, on a larger scale, in retinotopic space, by virtue of the V4 neuron's receptive field). Consider this approach in the light of the commonly offered parallel between vision and language. In linguistics, structural representations are described as productive (addressing an infinite domain of objects by finite means) and systematic (capable of treating various rearrangements of the

same components as distinct yet related, rather than merely as different). Analogies between language and neural representations of shape, however, are problematic, for a variety of reasons ranging from philosophical to computational and implementational⁶. For example, positing the existence of alphabet-like (discrete and categorical) shape parts 'out there' in the world, which need merely to be detected and assembled into a relational description, is ontologically untenable⁷ and may be computationally infeasible⁸. Furthermore, abstract relational structures imported into neuroscience from linguistic theory (which itself traditionally shuns neurobiological issues; see ref. 9, p. 2) are ill suited to distributed neural implementation, where they give rise to the binding problem¹⁰, arguably unnecessarily⁶.

It is interesting to observe that the core of the method described by Pasupathy and Connor—transforming the high-dimensional ensemble response into the low-dimensional set of curvature values estimated at a small number of locations around the object's contour—is an excellent example of perception treated as a problem in dimensionality reduction. A low-dimensional representation, especially if it is tailored to the nature of the data, is a guarantee of good performance in behaviorally crucial tasks such as learning from examples and generalization to new stimuli (intuitively, a lower dimensional representation space can be 'covered' by fewer examples; the required number of examples grows exponentially with dimensionality). The approach of Pasupathy and Connor is thus a step in the right direction both methodologically and substantively, and it will be interesting to see if it can be extended from silhouettes to solid, shaded and textured shapes.

1. Braitenberg, V. *On the Texture of Brains* (Springer, New York, 1977).
2. Mumford, D. in *Large-Scale Neuronal Theories of the Brain* (eds. Koch, C. & Davis, J. L.) 125–152 (MIT Press, Cambridge, Massachusetts, 1994).
3. Sternberg, R. J. *Cognitive Psychology* 3rd edn. (Thomson Wadsworth, Belmont, California, 2002).
4. Pasupathy, A. & Connor, C. E. *Nat. Neurosci.* 5, 1332–1338 (2002).
5. Pouget, A., Zemel, R. S. & Dayan, P. *Nat. Rev. Neurosci.* 1, 125–132 (2000).
6. Edelman, S. *Trends Cognit. Sci.* 6, 125–131 (2002).
7. Smith, B. *Topoi* 20, 131–148 (2001).
8. Dickinson, S. et al. *Image Vis. Comput.* 15, 277–292 (1997).
9. Chomsky, N. *The Minimalist Program* (MIT Press, Cambridge, Massachusetts, 1995).
10. Roskies, A. L. *Neuron* 24, 7–9 (1999).
11. Edelman, S. *Representation and Recognition in Vision* (MIT Press, Cambridge, Massachusetts, 1999).
12. Edelman, S. & Intrator, N. in *Mechanisms of Perceptual Learning* (eds. Medin, D., Goldstone, R. & Schyns, P.) 353–380 (Academic, San Diego, 1997).

The cortical basis of motor planning: does it take two to tango?

René Marois

A new study using fMRI shows that the human frontal cortex—and not parietal cortex—is the primary locus of movement planning.

“I did everything Fred did, only backwards and in high heels,” quipped Ginger Rogers about her dance partnership with Fred Astaire. Watching them in perfect unison, one could be hard pressed to know who was leading and who was following. In brain function, the frontal and parietal cortex are to the cortical control of action what Fred and Ginger are to dance: a dynamic and seemingly inseparable double act¹. In fact, the recurrent co-activation of this frontoparietal network in functional imaging studies has made it difficult to tease apart their relative contributions to the control of action. Now, in this issue, Connolly and colleagues identify one stage of information processing that dissociates activity in the front and back of the brain: the bulk of movement planning is a property of the frontal, but not of the parietal, cortex².

The control of action has been extensively investigated in the frontal eye field (FEF) of the frontal lobe and the lateral intra-parietal (LIP) area of the parietal lobe, two key regions of the cortical network that controls where our eyes move and where our attention is directed³. Long studied in the monkey, putatively homologous brain regions have since been mapped in humans⁴ (Fig. 1). When

subjects rapidly shift their gaze from one object to another in a visual scene—an eye movement known as a saccade—or when they shift their attention to a different scene location from the one they are fixating, the FEF and LIP areas are invariably activated. Indeed, frontoparietal activation is ubiquitous in neuroimaging studies of attention and visual cognition⁵. Not surprisingly, the FEF and LIP are strongly interconnected and have similar physiological properties⁶. Should one thus conclude that the parietal and frontal cortex make equal contribution to the control of action? There is reason to believe that this may not be the case when we are preparing to act: although much evidence supports a role for the frontal cortex in the planning and preparation of movements^{7,8}, similar evidence for the parietal cortex is more equivocal^{9–11}. However, strong support for this notion had been lacking, primarily because most evidence is derived from single-neuron studies in non-human primates that have examined only the frontal or parietal cortex, but not both. In addition, few imaging studies of visuomotor cognition have aimed at dissociating the planning of a movement from the target of that movement.

Connolly and colleagues² addressed this issue by measuring brain activity in both frontal and parietal cortex with event-related fMRI while human subjects performed an eye movement planning task. At the onset of each trial, subjects

were instructed, via a central color cue, to execute an eye movement either toward (pro-saccade) or away from (anti-saccade) a peripheral target briefly presented either to the left or right visual field (Fig. 2). Crucially, a variable delay of 0, 2 or 4 seconds was introduced between the presentation of the cue and target. Subjects were faster at executing a saccade when there was a delay between the cue and target, demonstrating that some aspect of movement preparation beneficial to motor execution took place during the delay period. Connolly and colleagues found that for both pro- and anti-saccade trials, activation of the FEF ramped up during the cue and delay period, such that it was highest at the time of target presentation for the four-second delay and lowest at the zero-second delay. In other words, the delay between cue and target permitted a buildup of activity in the FEF before target presentation. In stark contrast, LIP showed no preparatory activity whatsoever during the delay period.

Any fMRI experiment that relies on a negative finding, such as the absence of parietal activation in this case, must safeguard against the trivial possibility

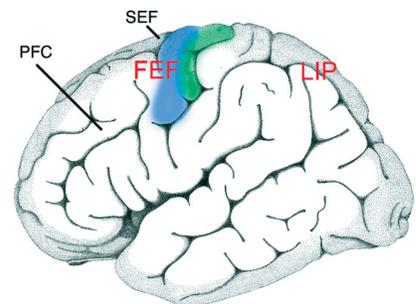


Fig. 1. Cortical centers of the oculomotor system. Connolly et al. measured brain activity in the FEF and presumptive LIP. Other cortical areas important in oculomotor control are the supplementary eye field (SEF) and the prefrontal cortex (PFC). The location of the primary motor cortex (green strip) and frontal premotor cortex (blue strip) is shown as a reference point.

The author is in the Department of Psychology, Vanderbilt University, 530 Wilson Hall, Nashville, Tennessee 37203, USA.
e-mail: rene.marois@vanderbilt.edu