A Scalar Neural Code for Categories in Parietal Cortex: Representing Cognitive Variables as “More” or “Less”

Matthew V. Chafee[1,2,3,*]
[1]Department of Neuroscience, University of Minnesota, Minneapolis, MN 55455, USA
[2]Brain Sciences Center, Minneapolis VA Medical Center, Minneapolis, MN 55417, USA
[3]Center for Cognitive Sciences, University of Minnesota, Minneapolis, MN 55455, USA
*Correspondence: chafe001@umn.edu
http://dx.doi.org/10.1016/j.neuron.2012.12.025

In this issue, Fitzgerald et al. (2013) show that LIP neurons in monkeys encode categorically distinct task conditions using a scalar code. Activity scales up or down to encode different categories, with neurons maintaining proportional levels of activity in relation to one another.

One of the primary goals of cognitive neuroscience is to effectively equate cognitive and neural processes. To achieve this, experimenters design behavioral paradigms in which determining the correct behavioral response is contingent upon a specific cognitive operation of interest. They then measure neural activity (through a variety of techniques in a variety of species) during task performance and relate the information encoded by that neural activity to the cognitive processing required. In studies of the neural correlates of cognition in monkeys, this approach has been applied to neural activity in the posterior parietal cortex with considerable success. For example, physiological signals in this area have been identified at the single-neuron level that carry information about decisions (Kiani and Shadlen, 2009), categories (Freedman and Assad, 2006; Goodwin et al., 2012; Swaminathan and Freedman, 2012), rules (Stoet and Snyder, 2004), numbers (Nieder and Miller, 2004), relationships (Chafee et al., 2007), and time (Leon and Shadlen, 2003).

The information reflected in this neural activity is relatively abstract, provides a logical basis for selecting the appropriate action in the task, and typically does not correlate with the attributes of the specific stimuli or movements involved—all characteristics one might expect of neural signals that support cognitive processing. However, demonstrating a correspondence between neural and cognitive representations necessarily involves making an inference as to the nature of the cognitive process taking place in the brain of a subject based on the pattern of behavior that one can observe. That inference is not necessarily straightforward, insofar as any single behavioral response could be the product of a large number of alternative cognitive operations. Further, although we often assume humans and monkeys share a core set of simple cognitive abilities, the degree to which these abilities differ between species is not precisely known. Finally, it is often necessary to first train monkeys to perform cognitive tasks before studying the neural correlates of cognitive processing (although, see Qi et al., 2012). Training can take many months, involves repeatedly rewarding successively more accurate behavioral responses, is likely to recruit reward-driven synaptic plasticity in the cortex, and therefore is likely to sculpt the computations performed by the cortical neurons and networks that are under study. It is not surprising then that at the end of training, neural signals are detected that reflect the cognitive processing the task in question was designed to recruit. However, it is not possible, a priori, to predict, before neural recording, *which* neural signals will emerge, the nature of the neural representation that will enable successful performance, the algorithm the brain will discover during training, or even the precise nature of the cognitive process that ultimately will result. Neural recording after training can help to answer all of these questions, and in some cases, the answers are surprising.

In this issue of *Neuron*, Fitzgerald et al. (2013) provide evidence that LIP neurons encode categorically different task conditions using a scalar neural code. In this representation, firing rates of neurons maintain a constant relation to one another but scale together in parallel to higher or lower levels to encode different task conditions. This is surprising because in prior studies, parietal cortex was found to represent categorical task conditions in cognitive tasks by selectively activating largely distinct subpopulations of neurons to encode each task condition, rather than scaling activity in a single group of neurons (Goodwin et al., 2012; Nieder and Miller, 2004; Swaminathan and Freedman, 2012). Fitzgerald et al. (2013) were led to the conclusion that LIP employed a scalar code by their discovery that the distribution of neurons in LIP preferring each of several categories or conditions in a task was strikingly biased. Given a task with two or three outcomes, they found that the majority of LIP neurons were most active for the same outcome and were similarly less active for the alternatives, leading to the idea that LIP employed a scalar code to meet the cognitive demands of the task. After eliminating trivial accounts for this neuronal bias based on the monkeys’ performance, Fitzgerald et al. (2013) arrived at the insight that the bias in population representation was predicted by the prior finding that network dynamics in area LIP were one-dimensional (Ganguli et al., 2008).

Ganguli and colleagues examined neural activity in LIP during the delay period of attention and perceptual decision tasks. They found that the firing rates of LIP neurons exhibited the same relative levels of activity with respect to one another during the delay period that they exhibited in their spontaneous activity, during the intertrial interval (Ganguli et al., 2008), suggesting that activity scaled up or down to code task-critical information. Representing the pattern of activity in the network as a vector in a rate space (with dimensions equal to the number of neurons), a network employing one-dimensional dynamics would represent information essentially by lengthening or shortening the population rate vector, rather than by changing its direction (as the direction of the vector is controlled by the distribution of firing rates over neurons and not the global level of activity). Fitzgerald et al. (2013) reasoned that if LIP neurons utilized one-dimensional dynamics and a scalar code to differentiate between stimulus categories and task conditions, then most neurons would be maximally activated for the same task condition, explaining the observed population bias. By providing strong evidence that population representations in area LIP are biased and neural dynamics are approximately one-dimensional under a variety of task conditions and cognitive demands, Fitzgerald et al. (2013) strengthen the evidence that this is a generalized characteristic of how LIP neurons code information in cognitive tasks.

In their report in this issue, Fitzgerald et al. (2013) trained monkeys to perform categorization, association, and perceptual decision tasks, and they recorded neural activity in an area of posterior parietal cortex (area LIP) known to play a role in sensorimotor control (Snyder et al., 1997), attention (Bisley and Goldberg, 2003; Gottlieb et al., 1998), and cognition (Freedman and Assad, 2006; Kiani and Shadlen, 2009). In the categorization task, monkeys viewed a patch of dots that moved in different directions and reported whether each stimulus belonged to one of two arbitrary motion categories (up-left or down-right) by imposing a diagonal category boundary based on direction. In the paired-associate task, monkeys learned to associate three pairs of visual objects and selected the object that was paired with a sample stimulus each trial. In the perceptual decision experiment, monkeys decided whether to saccade to a red or green target based on the direction of a visual motion stimulus. In each case, monkeys discriminated between task conditions based on a nominal cognitive variable (whether category, association, or decision) that could take on a small number of different values in the experiment. Many of us would have predicted that after training and during performance, each category or outcome in the task would activate comparable numbers of LIP neurons with comparable strength, because each was rewarded with comparable frequency and played an equivalent role in controlling behavior.

Instead, what these authors show is that the category preferences of LIP neurons were highly biased. Many more neurons were found to be selectively activated to encode motion category *A* than *B*, for example, in spite of their behavioral equivalence. A similarly skewed and biased population representation was observed in the paired associate and perceptual decision tasks. Most neurons in LIP preferred the same stimulus pair or direction of motion, with relatively few neurons preferring the others. What these data imply is that rather than encoding these categorically different task conditions as discrete cognitive representations, parietal neurons *rank them*, coding categories, associations, and decisions as points along a neural continuum, as "more" or "less." That suggests that LIP neurons collectively signal a scalar value, with different ranges of global activity corresponding to the different task conditions. For example, given a task with three conditions (*A*, *B*, and *C*), LIP neurons might produce 30 impulses per second to represent condition *A*, 20 impulses per second to encode condition *B*, and ten impulses per second to represent category *C* (hence the biased population preference for condition *A*). The readout of task condition would then be based on the overall firing rate in the population, rather than on the distribution of firing rates over neurons within the population. The fact that Fitzgerald et al. (2013) were able to show that this bias in population representation generalized across individual monkeys, experimental paradigms, and even laboratories suggests that it might reflect a computational principle rather than a task-specific effect.

An alternative explanation for the bias in the population representation of categories was that since visual stimuli in different categories varied in their visual features (for example direction of motion), a population bias in the preferred direction of LIP neurons could account for the bias in category preferences observed. Population biases among neurons tuned to direction of movement have been shown in area M1, for example (Naselaris et al., 2006). However, in the present case, the bias in population representation was much stronger during the delay period after the offset of visual stimuli than during the time that the stimuli were visible (a bias in visual motion processing would predict the reverse). In addition, and importantly, the bias in population representation was a function of training to perform the categorization task. The authors found that the distribution of preferred direction categories in area LIP was not similarly biased in a monkey viewing the same motion stimuli before it had been trained to categorize those stimuli. These observations provide compelling evidence that the bias is augmented in the delay period, is a consequence of training, and therefore does not reflect an intrinsic (pre-existing) asymmetry in the population distribution of preferred motion directions in area LIP. That is, the population bias arises during training to meet the cognitive processing demands of the task, rather than reflecting how LIP neurons intrinsically encode the visual stimuli involved.

Fitzgerald et al. (2013) found that spontaneous and evoked activity levels were correlated in single neurons, as did Ganguli and colleagues (Ganguli et al., 2008). One might wonder whether these data suggest that one-dimensional dynamics emerge as a consequence of neurons differing in their intrinsic excitability. Both groups observed that one-dimensional dynamics were less prevalent during periods of visual input, and this argues against that interpretation. If differences in excitability accounted for one-dimensional dynamics, these dynamics should be particularly evident when the network was excited by extrinsic input. It still seems possible that variability in the baseline level of activity across neurons could contribute to one-dimensional dynamics and scalar coding, in which case signals driven by extrinsic input would be superimposed on an initial and persistent difference in baseline firing rate. However, the present data provide convincing evidence that the population representation of task conditions is considerably biased, and one would not necessarily predict this on the basis of pre-existing differences in the baseline firing rate of neurons.

A strict one-dimensional coding scheme would predict that *all* neurons would have the same preferred category in categorization tasks or exhibit the same ranking of response magnitude over categorically distinct task conditions. Certainly, in this study and others (Chafee et al., 2007; Freedman and Assad, 2006; Goodwin et al., 2012; Nieder and Miller, 2004), this is not the case. Although there may be a bias for neurons to prefer some over other categories at the population level, this bias is not absolute, with some neurons preferring the more weakly represented category. This suggests that the population vector would change in direction as well as length to encode different categories, and it becomes a question of the degree to which scaling of activity versus redistribution of activity within the population carries information about cognitive variables. Fitzgerald et al. (2013) describe population dynamics in area LIP as being "low-dimensional" rather than one-dimensional in categorization tasks, and it is important to note that the mechanism of neural coding in LIP sits somewhere in between a strict scalar code, utilizing one-dimensional dynamics, and a more traditional population code in which different categories are at least partially encoded by the redistribution of activity over neurons. In addition, Fitzgerald et al. (2013) are careful to note that one-dimensional dynamics may emerge from network constraints that do not derive directly from the necessity to encode categories or categorically distinct task conditions. For example, to perform the behavioral tasks they employed, cortical networks had to sustain patterns of activity in the absence of continuing sensory input (to mediate working memory). The required patterns of connectivity between neurons could impose the proportionality in firing rates observed. At this point, therefore, the exact computational role of scalar coding in LIP and its relation to behavior remains an open question.

Still, the degree to which one-dimensional dynamics pertain to the neural representation of categories, as reflected in the strong population biases observed, is surprising. One would not necessarily guess that the brain would scale activity in the same group of neurons to represent different categories, which are discrete cognitive representations that differ in kind rather than degree. Training monkeys to perform a cognitive paradigm essentially guarantees that neural correlates of the cognitive processing required will be recovered after training. However, as these results show, it is not possible a priori to determine how the brain will solve the cognitive problem that has been set, or to determine the nature of the computations that neurons will perform to mediate successful performance. In some cases, the characteristics of the underlying neural representation are unanticipated.

REFERENCES

Bisley, J.W., and Goldberg, M.E. (2003). Science *299*, 81–86.

Chafee, M.V., Averbeck, B.B., and Crowe, D.A. (2007). Cereb. Cortex *17*, 2914–2932.

Fitzgerald, J.K., Freedman, D.J., Fanini, A., Bennur, S., Gold, J.I., and Assad, J.A. (2013). Neuron *77*, this issue, 180–191.

Freedman, D.J., and Assad, J.A. (2006). Nature *443*, 85–88.

Ganguli, S., Bisley, J.W., Roitman, J.D., Shadlen, M.N., Goldberg, M.E., and Miller, K.D. (2008). Neuron *58*, 15–25.

Goodwin, S.J., Blackman, R.K., Sakellaridi, S., and Chafee, M.V. (2012). J. Neurosci. *32*, 3499–3515.

Gottlieb, J.P., Kusunoki, M., and Goldberg, M.E. (1998). Nature *391*, 481–484.

Kiani, R., and Shadlen, M.N. (2009). Science *324*, 759–764.

Leon, M.I., and Shadlen, M.N. (2003). Neuron *38*, 317–327.

Naselaris, T., Merchant, H., Amirikian, B., and Georgopoulos, A.P. (2006). J. Neurophysiol. *96*, 3231–3236.

Nieder, A., and Miller, E.K. (2004). Proc. Natl. Acad. Sci. USA *101*, 7457–7462.

Qi, X.L., Meyer, T., Stanford, T.R., and Constantinidis, C. (2012). J. Neurosci. *32*, 6161–6169.

Snyder, L.H., Batista, A.P., and Andersen, R.A. (1997). Nature *386*, 167–170.

Stoet, G., and Snyder, L.H. (2004). Neuron *42*, 1003–1012.

Swaminathan, S.K., and Freedman, D.J. (2012). Nat. Neurosci. *15*, 315–320.