Journal Club

Editor's Note: These short, critical reviews of recent papers in the *Journal*, written exclusively by graduate students or postdoctoral fellows, are intended to summarize the important findings of the paper and provide additional insight and commentary. For more information on the format and purpose of the Journal Club, please see http://www.jneurosci.org/misc/ifa_features.shtml.

The Spatiotemporal Structure of Population Coding in Monkey Parietal Cortex

Mark Stokes

Department of Experimental Psychology, Oxford University, Oxford OX1 3UD, United Kingdom
Review of Crowe et al.

Introduction

Conventionally, we think of a population code as a pattern of neural activity that represents information in the brain (DeCharms and Zador, 2000; Pouget et al., 2000). These patterns can be decoded upstream by detector units or an experimenter with an appropriate recording device. However, traditional accounts often assume that the mapping between the informational content and the spatial configuration that defines a unique pattern of activity is effectively stationary. The detector classifies output patterns according to the same rule, at least on the time scale of a single processing epoch.

It is becoming increasingly evident, however, that the same stimulus does not always elicit the same activation pattern (Buonomano and Maass, 2009), complicating the neural decoding problem. In the locust, for example, specific odors trigger a complex series of activity patterns within antennal lobe projection neurons (Mazor and Laurent, 2005). This complex but reliable trajectory through state-space could provide a more robust neural code than a more stable, or stationary, activation state. Indeed, the neurons

that receive output signals from these projection neurons respond maximally during the most dynamic phase of the olfactory response. Clearly, for the locust olfactory system the destination might be less important than the route traveled. Recent research by Crowe et al. (2010) demonstrates that similar dynamics might also underlie coding in monkey parietal cortex and explores how and why neural activation patterns evolve over time.

Crowe et al. (2010) used multielectrode arrays to record population responses in parietal area 7a while monkeys performed an object-construction task [Crowe et al. (2010), their Fig. 1]. Each experimental trial began with the presentation of a model stimulus, which consisted of a basic set of nine squares arranged to form an inverted “T” shape, plus an additional one or two squares extending either left or right of the main stem. Monkeys were trained to remember the location of these variable extensions in object-centered coordinates. After a short delay, a copy stimulus then was presented at a different screen position. The copy stimulus was identical to the original model, except one of the additional squares was removed. During a subsequent response phase, monkeys were presented sequentially with two choice stimuli that could complete the copy stimulus to match the model, and were rewarded for indicating the correct match with a key press response.

Using a cross-temporal approach to pattern classification of population activ-

ity, Crowe et al. (2010) demonstrate that coding of task-relevant spatial information (i.e., the object-centered location of an additional square on the model stimulus) in monkey parietal cortex varies across time, even when the sensory input remains constant. Importantly, activity pattern variability was not random but evolved along a systematic and reproducible trajectory. Moreover, patterns of neuronal activity at each time point were remarkably distinct: pattern classifiers trained on a specific time bin only performed optimally when tested on data from equivalent time bins [Crowe et al. (2010), their Fig. 5]. This temporal specificity suggests that time-dependent differences in the response profile do not simply reflect variable activation of a fixed, and more complete, population code. Rather, these data provide compelling evidence that task-relevant information is represented by a series of temporary activation patterns in parietal cortex.

Putative neural mechanisms: hidden states and internal interactions

The activation state of a network can be represented as a coordinate in multidimensional space, where each dimension is defined by the activation level (firing rate, integrated over some time window) of a different neuron (schematized in Fig. 1). This activation state is determined by the input, as well as the relative sensitivity of the population to that input. Sensitivity can be thought of as the “hidden state” of

Received Sept. 30, 2010; revised Oct. 29, 2010; accepted Nov. 2, 2010.

This research was supported by a Junior Research Fellowship, St. John's College, Oxford University. Thanks to A. C. Nobre and J. Duncan for helpful discussions and comments on the manuscript.

Correspondence should be addressed to Mark Stokes, Department of Experimental Psychology, South Parks Road, Oxford OX1 3UD, United Kingdom. E-mail: mark.stokes@sjc.ox.ac.uk.

DOI:10.1523/JNEUROSCI.5144-10.2011

Copyright © 2011 the authors 0270-6474/11/311167-03\$15.00/0

the system (Buonomano and Maass, 2009), which could arise from a myriad of possible biophysical mechanisms that influence the probability that a particular neuron will respond, given a certain input. At the population level, these hidden states will shape the pattern of activity that is elicited by a specific input pattern. For a paradigmatic example, consider synaptic weights. Over time, past experience modulates synaptic efficacy to transmit action potentials, thereby sculpting a pattern of weights that determines how downstream neurons will respond to specific input patterns.

In classic neural network architectures, hidden states are relatively stable. For example, synaptic weights may change only as a function of learning. However, what happens if hidden states vary at smaller time scales? In visual cortex of the anesthetized cat, Nikolić et al. (2009) found that the activation pattern elicited by a given stimulus also contains information regarding the identity of the previous stimulus. Presumably, the activation state elicited by the previous stimulus altered the sensitivity, or hidden state, of the neural population according to a specific pattern that was then expressed in the response to subsequent stimulation. The resultant activation state could then be read out to decode the identity of both stimuli. Similar mechanisms could operate during continuous stimulation. Constant input, X , to a network with an initial hidden state, H_t , will result in an activation state, A_t . If this state, in turn, triggers a temporary change in the hidden state of the network, now H_{t+1} , the same input times will now trigger a different activation state, A_{t+1} , which in turn could trigger another temporary change to the hidden states, H_{t+2} , and so on. These dependencies could prevent the system from settling into fixed-attractor states. Rather, such a system is more likely to enter temporary low-energy states, before hidden dynamics force the trajectory to a new point in state space (Rabinovich et al., 2008), as established in locust olfaction (Mazor and Laurent, 2005) and now in the parietal cortex of the behaving monkey (Crowe et al., 2010).

The results of Crowe et al. (2010) also point to another possible mechanism for dynamic population coding. They show that the secession of neural ensembles is predicted by intrinsic connectivity between neuron pairs, as indexed by correlated spontaneous activity. This relationship is consistent with the proposal that the trajectory through state space is, at least in part, deter-

Figure 1. Schematic of a dynamic population response profile. **A**, Color saturation (in red) represents activation level at time points, t_1 , t_2 , and t_3 . The inner cube surrounds three example neurons that are used to represent in **B**, the trajectory through three-dimensional activation state space.

mined by the synaptic strength between cells within the same local pool of neurons. In addition to generating an output response, local processing could also contribute to subsequent activation patterns within the same local circuit, which would in turn shape the evolution of the activation state.

Why code in four dimensions?

Evidence for nonstationary coding naturally raises the question: why code in time, as well as space? Increasing the available dimensionality for neural coding necessarily increases the representational capacity of the network. Specifically, the same points in activity space could be used to code unique representations if the sequence of activation states can inform decoding. The information capacity of a simplified network can be expressed in information theoretic terms as 2^n bits, where n is the number of binary coding units, or “neurons.” If we allow each unit to take a different state for each time point, t , the resultant capacity would scale accordingly, $2^{n \times t}$ bits.

However, the representational capacity of the brain is already enormous, without considering time—do we really need to take neural coding to this next dimension? Moreover, dynamic population coding could raise more problems than it solves. In particular, a spatiotemporal code requires a spatiotemporal decoder. Theoretically, pattern matching in space and time could be achieved by a downstream population with hidden states

that vary according to a complementary temporal profile. Whether or not this is neurobiologically plausible (especially concerning learning rules), such a decoding scheme would necessarily impose a lower-bound demand on temporal integration. Presumably, neural decoding always integrates over time to accumulate evidence, but if each time point essentially provides the same kind of information, then the network can “decide” at the shortest possible time after a criterion threshold has been reached. Therefore, the integration time is inherently flexible, depending on the strength of the signal, relative to noise, potentially resulting in very fast readout (Hung et al., 2005). In contrast, if a certain number of state transitions are required to disambiguate one signal from other potential signals, this imposes a fixed lower bound on the integration window even in the absence of noise, which would inevitably increase the minimum readout time.

Another possibility is that the evolving trajectory reflects functional differences in coding. For example, in parietal cortex, activity elicited by the initial input might progress through a series of processing steps to encode progressively more abstract, or task-relevant, features of the stimulus. As described above, these could evolve over time according to dynamic modulation of hidden states and/or local interactions defined by the inherent connectivity. Trajectories could also be shaped by time-dependent changes in the

input to parietal cortex. Dynamic codes from sensory-processing areas could provide time-varying parietal input, as could task-dependent feedback from prefrontal cortex. The complete trajectory could signal unique content (to a complementary spatiotemporal decoder), and/or multiple functionally distinct signals that can be decoded at multiple points along the trajectory (e.g., at temporarily stable attractor states). Within this framework, each differentiable state builds on the previous state, and therefore provides unique information about the current input, as well as recent input history (Nikolić et al., 2009). Importantly, a decoding algorithm need not necessarily consider the full trajectory: the recent activation history could be inferred from a given point in the multidimensional activation state space. In the case of constant stimulus input (Crowe et al., 2010), sequential, and distinct, activation states could reflect processing toward a more stable, and perhaps more abstract, representation for guiding behavior. In particular, given the mnemonic demands of the task in Crowe et al. (2010), such transformations could reflect processing toward a more time-durable representation, perhaps encoded across a sparse subset of neuron that is distinct from the neural population that is initially driven by the sensory input (Meyers et al., 2008). A sparse representation might be more durable, and resistant to interference from subsequent input, particularly if the later-stage neural ensembles are also less directly connected to sensory input.

Temporally distinct activation profiles in high-level brain areas, such a parietal

and prefrontal cortex, could also reflect distinct behavioral contexts per se, as they unfold sequentially throughout the experimental trial (Sigala et al., 2008). Moreover, the dissimilarity (i.e., distance in activation state space) of the neural code at different epochs would help differentiate information pertaining to specific task epochs, thus reducing cross talk between processing stages. In this scheme, the activity states code the behavioral context (or trial epoch/time), as well as stimulus-specific information. Crowe et al. (2010) were also able to decode the relative time point in the trial from dynamic parietal codes [Crowe et al. (2010), their Fig. 10], providing important evidence that the population dynamics in parietal cortex may also represent the sequential structure of the behavioral task.

Summary

Sophisticated pattern-analytic approaches to neurophysiology present an exciting new window into the temporal dynamics of population coding. Similar pattern-analytic approaches have also become widely used in functional magnetic resonance imaging (fMRI) studies; however, the coarse temporal resolution of fMRI can only provide a time-averaged window into the population response. This static picture belies the evident complexity of neural dynamics revealed by multielectrode recordings, as in Crowe et al. (2010). No doubt, future studies using real-time population measures will be able to address many of the questions that so far remain unanswered. In particular, how does oscillatory behavior shape the temporal profile of dynamic population coding? Multisite recording could also explore how dynamic patterns of activity are

read out by other brain areas to determine whether time-varying patterns represent stable content or multiple functionally distinct neural codes.

References

- Buonomano DV, Maass W (2009) State-dependent computations: Spatiotemporal processing in cortical networks. *Nat Rev Neurosci* 10:113–125.
- Crowe DA, Averbeck BB, Chafee MV (2010) Rapid sequences of population activity patterns dynamically encode task-critical spatial information in parietal cortex. *J Neurosci* 30:11640–11653.
- DeCharms RC, Zador A (2000) Neural representation and the cortical code. *Ann Rev Neurosci* 23:613–647.
- Hung CP, Kreiman G, Poggio T, DiCarlo JJ (2005) Fast readout of object identity from macaque inferior temporal cortex. *Science* 310:863–866.
- Mazor O, Laurent G (2005) Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* 48:661–673.
- Meyers EM, Freedman DJ, Kreiman G, Miller EK, Poggio T (2008) Dynamic population coding of category information in inferior temporal and prefrontal cortex. *J Neurophysiol* 100:1407–1419.
- Nikolić D, Häusler S, Singer W, Maass W (2009) Distributed fading memory for stimulus properties in the primary visual cortex. *PLoS Biol* 7:e1000260.
- Pouget A, Dayan P, Zemel R (2000) Information processing with population codes. *Nat Rev Neurosci* 1:125–132.
- Rabinovich M, Huerta R, Laurent G (2008) Neuroscience: transient dynamics for neural processing. *Science* 321:48–50.
- Sigala N, Kusunoki M, Nimmo-Smith I, Gaffan D, Duncan J (2008) Hierarchical coding for sequential task events in the monkey prefrontal cortex. *Proc Natl Acad Sci U S A* 105:11969–11974.