Published in final edited form as:

Curr Opin Neurobiol. 2016 April ; 37: v–viii. doi:10.1016/j.conb.2016.03.003.

Editorial overview: Neurobiology of cognitive behavior: Complexity of neural computation and cognition

Alla Karpova¹ and Roozbeh Kiani²

¹Janelia Research Campus, Howard Hughes Medical Institute, 19700 Helix Drive, Ashburn, VA 20147, USA

²Center for Neural Science, New York University, 4 Washington Pl, Room 809, New York, NY 10003, USA

Abstract

We live in an age when our phones incorporate real-time updates on traffic and subway delays to help us navigate complex urban environments, vacuum cleaning robots map the layout of our apartments to optimize their cleaning strategies, and cars are beginning to drive themselves. But, amazing though today's artificial cognition systems may seem, the genuine mystery is the flexibility and adaptability with which their precursors and creators, brains, acquire and use knowledge. For at least two centuries, psychologists and cognitive scientists have studied human and animal behavior in an effort to better understand the faculties that support natural cognition: working memory, attention, multisensory integration, value-based decision-making, and analogical reasoning. They have helped organize our thinking about the likely algorithms the brain must use, and suggested central brain regions that are likely involved in executing them. Thus, it may seem like the stage is set for neurophysiologists to work out neural mechanisms underlying cognitive processes. However, despite tremendous progress over the past half century in understanding peripheral sensory and motor systems, the sheer complexity of central brain circuits has meant that cognition has retained most of its veils. For this issue, we asked contributors to define the experimental and theoretical challenges that cognition poses to neuroscience researchers, and to offer suggestions for how the field might productively tackle the complexity that comes with working on this frontier.

The classical approach to understanding the neural basis of cognition has relied on recording neural activity in animals performing well-controlled behavioral tasks, and searching for single neuron correlates of specific cognitive processes in brain regions thought to be responsible for those functions. A surprising degree of insight has been gleaned from this approach, and the chapter by **Brody and Hanks** provides an overview of progress its modernized version has enabled in understanding the neural underpinnings of one well-established cognitive function - perceptual evidence accumulation. Buried just underneath the apparently simple findings, however, were hints that they did not represent the entire picture. In one example, motor cortex was found to develop representations of specific sensory stimuli when those stimuli became contextually relevant for behavioral decisions (Zach et al., 2008). **Cumming and Nienborg** highlight a related finding in the field of perceptual decision-making: perplexingly, behavioral decisions appear to be far more correlated with fluctuations in the activity of one or a few sensory neurons than one would

expect. As these authors detail, an emerging view is that feedback-related entrainment of activity in sensory areas shapes these dynamics. The central role of cognitive feedback in shaping sensory representations is also highlighted by **Lehky and Tanaka**. They argue that in object recognition, this feedback represents high-level expectation of possible sensory inputs based on task information that can assist in object interpretation. More broadly, such unexpected mixing of sensory, motor, contextual and decision variables across cortical areas suggested that neural coding underlying even the simplest cognitive processes is much more complex and diverse than naïve models that rely on the segregation of such functions in the brain might assume. Rich task-relevant neural representations thus arise from the involvement of both local and distributed networks.

The recent ability to simultaneously monitor activity of large groups of neurons has given us an even deeper look at the richness and complexity of neural representations in cognitive tasks. Advances in high-density probe design and the increased sensitivity of genetically-encoded activity indicators have enabled population recordings from hundreds, and even thousands, of neurons simultaneously. Parallel advances in statistical analyses of multivariate time series data, such as approaches discussed by **Poorahmadi and Noorbaloochi**, have made it possible to extract meaning from this deluge of data. Such approaches have recapitulated and deepened the findings from seminal single neuron recordings that the activity of neuronal populations often displays many scales of spatial and temporal variability that go beyond what would be directly expected from the underlying task timescales (Druckmann and Chklovskii, 2012; Kaufman et al., 2014; Romo et al., 1999). It may seem then, that this added complexity of circuit dynamics will make the goal of deciphering the neural basis of cognitive function particularly vexing. However, it has been posited that some circuit dynamics have a negligible effect on a given computation. Indeed, when population activity is described as trajectories in activity state space, relatively low-dimensional dynamics are often found to be informative for task-specific computations (Mante et al., 2013; Mazor and Laurent, 2005). Although **Fetsch** correctly cautions us not to rush to conclude that any “interesting” dynamics constitute a mechanism for an actual biological neural computation, recent studies have lent support to the notion that certain modes of dynamics are specifically dedicated to encoding task-relevant information (Li et al., 2016; Sadtler et al., 2014). In one example, characterizing the effect of transient optogenetic suppression of network activity in a working memory task, Li et al. demonstrated that when the perturbation was performed early during the delay period, the detailed dynamics of activity recovered fully selectively along the hypothesized task-relevant mode (Li et al., 2016). What this emerging conceptual framework of activity mode-specific dynamics suggests is that one way to reign in the complexity of the neural code then is by separating those aspects of complex neural representations that can, or should, be explained given a particular task, away from the remaining – sometimes known as “null space” – dynamics.

A powerful theoretical avenue for identifying the explainable aspects of the neural code that can complement clever combinations of experimental and computational approaches is to model circuit dynamics in systems built to perform a particular computational building block. **Deneve** highlights the utility of this approach: by starting with a network designed for

efficient temporal processing of sensory stimuli, she demonstrates that some variability of neural responses can arise simply as a result of coordinated computation across the network. Taking a second approach, which seeks to exploit particular constraints of biological neural networks, **Wang and Kennedy**, in turn, detail how key properties of the structural organization of the primate cortex naturally constrain network dynamics in a way that the end result mirrors the hierarchy of timescales (from brief in sensory areas to more persistent in prefrontal cortex) found in the brain. **Fusi et al** combine both approaches by designing systems that can flexibly categorize inputs, for instance in a context dependent manner, but requiring the readout of such categorization to be linear, a constraint that confers greater biological plausibility. The central characteristic of such systems is that inputs need to be represented in high dimensions, and the authors demonstrate that this can be achieved through the mechanism of non-linear mixed selectivity – a frequently observed aspect of neural representations in higher order cortical areas. **Miller** takes a similarly dual approach by parsing out two central building blocks of more complex computations: neural selectivity and gain control, and developing simple models for how these computations can naturally occur in canonical cortical architecture. Overall, these theoretical approaches should greatly advance our understanding of complex neural computations by detailing the associated core neural responses and identifying surprising features of network dynamics.

The above-mentioned theoretical and experimental efforts are being applied in simpler, more tractable systems that may more readily provide insights into the neural mechanisms underlying some cognitive processes, and four chapters in this volume highlight several diverse approaches to attaining tractability. **Schneidman** summarizes insights derived from statistical models of the responses of small networks of sensory neurons and speculates about how these may extend to larger networks and brain areas deeper in the processing stream. **Gjoergjieva et al.** discuss the prospect of being able to model the dynamics of larger circuits on the basis of a comprehensive understanding of intrinsic and synaptic time scales for individual neurons and small circuits. **Golub et al.** describe an elegant way of bringing an entire sensorimotor system under direct experimental control through the clever use of brain computer interfaces (BCI). The authors argue that studying how an animal can alter the dynamics of even just the subset of neurons that can be observed in an experiment, we will eventually be able to get at the neural mechanisms behind many cognitive computations. Finally, **Haberkern and Jayaraman** discuss recent behavioral and physiological evidence in support of broadly relevant elements of cognitive processing in insects. Insect brains confer the advantages of stereotyped topography, small size, and, in the case of *Drosophila*, exquisite genetic access, which can make them more tractable for mechanistic studies of basic cognitive computations. In their diverse ways, these chapters highlight the advantages that these smaller systems offer in understanding the underlying neural circuit computations, and point to signs that the principles extracted from such research may be broadly applicable.

The universe of known complexity of neural computation is already large enough to occupy the field for many decades, so it may seem rash to simultaneously seek to understand the neural implementation of still more complex cognitive functions. **Koechlin** provides a blueprint of a comprehensive model for adaptive behavior in real-world open-ended

environments, and its algorithmic solution comprises a daunting set of interacting cognitive processes. If, as has been postulated recently (Gao and Ganguli, 2015), the low dimensionality of task-related neural dynamics observed in experiments to date is a reflection of the simplicity of the contemporary behavioral tasks, and yet we are only barely beginning to understand these dynamics, would we not drown in irreducible big data if we tackle ever something so complex? The answer to this question remains to be revealed, but one hopeful possibility is that if the added task complexity is more “naturally relevant”, then what now seems like unwieldy null space dynamics will take on a more interpretable form. Two articles in this volume provide hope that engaging animals in relevant, if more complex, ways can indeed bring about informative changes to neural representations. Drawing on recent findings, **Dudman and Krakauer** argue that neural dynamics in the basal ganglia track and predict kinematic parameters of movement execution when actions are under motivational control. They suggest that a more restrictive view of the basal ganglia’s role in the control of action selection stems from an uncoupling of these neural dynamics from movement kinematics in over-trained animals performing less demanding tasks. Along similar lines, **Kolling et al.** discuss how the conflicting interpretations of the neural representation in the anterior cingulate cortex are beginning to be disentangled in specifically-tailored behavioral tasks.

So how do we proceed with adding just the right type of cognitive complexity to behavioral tasks? **Tervo et al.** advocate one possible avenue. These authors propose that we build on the intuition widely adopted by the cognitive and computer science communities that animals approach complex environments by attempting to create hierarchical models of the environment’s latent structure. Developing behavioral tasks that have complex hierarchical relationships would thus be one way to add “natural” complexity. Of course, the ability to bring the underlying cognitive process under experimental control will depend not only on whether an animal’s solution to the added complexity is of the expected nature, but also on whether the experimenter can identify the exact structured abstraction the animal represents and uses to solve the task. Some ideas for how to best do that could come from fields with a rich tradition of rigorous behavioral control. **Fetsch** articulates the case for combining clever behavioral designs, careful analyses of behavioral data and normative behavioral modeling that enabled the fields of psychophysics and motor control to isolate and study specific behavioral component processes. The advantage of such an approach is made clear in the chapter by **Wolpert and Flanagan**, which outlines advances in the field of sensorimotor learning enabled by a combination of thoughtful behavioral approaches and computational modeling. However, **Barack and Gold** caution that many complex strategic factors that can contribute to task performance at times remain ignored even in psychophysics experiments. As the complexity of cognitive functions under study increases, it may become progressively more important to supplement the rigorous approaches to behavior discussed by these authors with experimental frameworks that permit direct behavioral access to the computations involved.

We hope that the diverse articles in this issue will give the reader a sense for ways, in which experimental and theoretical ways to probe the neural basis of complex computations have evolved over the past few decades of systems neuroscience. **Gomez-Marin and Mainen**

caution, however, that it may be too early to pick winning strategies and advocate for an open-ended and multi-pronged approach that draws freely from other fields as appropriate. **Yamins and DiCarlo** discuss just such an example of cross-pollination in their chapter on deep neural networks – a powerful approach from artificial intelligence that may help model and predict neural responses in visual cortex and beyond. In turn, **Daniels et al.** suggest that information theoretic analysis may help bridge the explanatory gap going from information processing at small scales, such as at the level of individual neurons, to more collective phenomena associated with, say, the interactions of neural populations and multiple brain regions.

Conclusion

The last half of the twentieth century witnessed a cognitive revolution, with insights from psychology, linguistics and neuroscience contributing to the recognition that many feats of natural intelligence required non-associationist explanations (Miller, 2003). This notion has fostered a robust debate on what cognition actually is, what its components might be, what the appropriate language to describe it is, and what it would take to link it to processes in the brain. While consensus about the central role of some cognitive faculties may be emerging in the cognitive science community, it is yet to be met by a comprehensive understanding of how the brain gives rise to these phenomena. Contemporary neuroscientists attempting to reverse engineer the brain and its ability to give rise to cognition are simultaneously equipped with enviable tools – dense multiunit recording; circuit perturbation methods; high dimensional quantification of behavior; a century of statistical tools to tease out insights from large datasets – and the bewildering requirement of wielding these tools to tackle one of the greatest scientific mysteries. The challenge we face though is not just in picking the right technical approaches to find answers to questions in cognition. In an era increasingly dominated by powerful genetic tools in model organisms, **Gomez-Marin and Mainen** warn against too restrictive a focus on mechanistic experiments when the boundaries of cognition have yet to be fully explored. Continuing the discussion about the advantages of, and limitations inherent in, the existing set of conceptual, experimental and analytical tools will be crucial for developing a roadmap for the next generation of experimental and theoretical studies in systems and cognitive neuroscience to tackle the complexity of cognition.

References

- Druckmann S, Chklovskii DB. Neuronal circuits underlying persistent representations despite time varying activity. *Current Biology*. 2012; 22:2095–2103. [PubMed: 23084992]
- Gao P, Ganguli S. On simplicity and complexity in the brave new world of large-scale neuroscience. *Current opinion in neurobiology*. 2015; 32:148–155. [PubMed: 25932978]
- Kaufman MT, Churchland MM, Ryu SI, Shenoy KV. Cortical activity in the null space: permitting preparation without movement. *Nature neuroscience*. 2014; 17:440–448. [PubMed: 24487233]
- Li N, Daie K, Svoboda K, Druckmann S. Robust neuronal dynamics in premotor cortex during motor planning. *Nature*. 2016 In press.
- Mante V, Sussillo D, Shenoy KV, Newsome WT. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*. 2013; 503:78–84. [PubMed: 24201281]
- Mazor O, Laurent G. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron*. 2005; 48:661–673. [PubMed: 16301181]

- Miller GA. The cognitive revolution: a historical perspective. *Trends in cognitive sciences*. 2003; 7:141–144. [PubMed: 12639696]
- Romo R, Brody CD, Hernández A, Lemus L. Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature*. 1999; 399:470–473. [PubMed: 10365959]
- Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, Byron MY, Batista AP. Neural constraints on learning. *Nature*. 2014; 512:423–426. [PubMed: 25164754]
- Zach N, Inbar D, Grinvald Y, Bergman H, Vaadia E. Emergence of novel representations in primary motor cortex and premotor neurons during associative learning. *The Journal of Neuroscience*. 2008; 28:9545–9556. [PubMed: 18799686]