



HHS Public Access

Author manuscript

Curr Opin Neurobiol. Author manuscript; available in PMC 2017 April 01.

Published in final edited form as:

Curr Opin Neurobiol. 2016 April ; 37: 16–22. doi:10.1016/j.conb.2015.12.002.

The importance of task design and behavioral control for understanding the neural basis of cognitive functions

Christopher R Fetsch

Department of Neuroscience, Columbia University, USA

Abstract

The success of systems neuroscience depends on the ability to forge quantitative links between neural activity and behavior. Traditionally, this process has benefited from the rigorous development and testing of hypotheses using tools derived from classical psychophysics and computational motor control. As our capacity for measuring neural activity improves, accompanied by powerful new analysis strategies, it seems prudent to remember what these traditional approaches have to offer. Here I present a perspective on the merits of principled task design and tight behavioral control, along with some words of caution about interpretation in unguided, large-scale neural recording studies. I argue that a judicious combination of new and old approaches is the best way to advance our understanding of higher brain function in health and disease.

Introduction

It is an exciting time for systems and cognitive neuroscience. Methods for collecting and analyzing data are improving at a remarkable pace [1,2], and questions once limited to human and nonhuman primate studies are now being addressed in smaller animal models for which large-scale data collection and powerful circuit dissection techniques are more tractable [3,4]. However, the pursuit of advanced technology and ‘big data’ should not come at the expense of well-defined hypotheses and rigorous behavioral control. For although we wish to understand the inner workings of the mind, we only have access to a coarse distillation thereof, namely behavior. The degree of insight attainable in an experiment is therefore limited not by how many neurons we record, but by the quality of the mapping we can create between internal states and behavioral reports.

The first part of this review describes some principles of task design that originate in the quantitative study of perception or movement, yet should prove useful for probing processes that lie squarely in between. The second part consists of a brief commentary on recent approaches in the literature that might benefit from such ‘old fashioned’ behavioral tools. My goal is not to suppress enthusiasm for new methods or to criticize any particular

Corresponding author: Fetsch, Christopher R (crf2129@columbia.edu).

Conflict of interest statement

Nothing declared.

approach, but to encourage renewed emphasis on smart task design and careful quantification of behavior to help make the most of these advances.

What is meant by behavioral control, and why is it important?

If we want to understand how neural activity gives rise to our sensations, thoughts, memories, and decisions, we must come to grips with the fact that everything we know about such internal processes can only be inferred by observing external behavior. This is obviously true for animal models, but also applies in human subjects for whom introspection can be misleading and whose verbal reports are just another form of behavioral assay. Making things worse, this inference problem is ill posed: many internal states could lead to the same behavioral outcome.

So how do we make progress? The key is to constrain the space of possibilities as much as possible. Psychophysics does this by (a) carefully controlling sensory input, (b) measuring behavioral responses in a principled, quantitative fashion, and (c) accounting for response variability with an underlying statistical process, which then informs the search for neural mechanism. Similarly, research in computational motor control relies on monitoring the output of the system (i.e. eye or limb movements) with high resolution, and constructing normative models that explain the variability of behavior in terms of what is being optimized, and how [5–7]. In both cases, this kind of groundwork has been highly successful in advancing our understanding of sensory and motor processes. My contention is that the same degree of rigor will be needed to support inferences about the neural basis of cognitive functions.

Lessons from psychophysics

When vision scientists gained access to the physiology of visual neurons in the mid-20th century, there was great initial excitement but also caution [8,9]. Brindley [10] and Teller [11] developed the concept of a ‘linking hypothesis’ as way to formalize what one can and cannot conclude about how perception works based on properties of sensory neurons. This idea, together with the insights of Barlow [12], Marr [9], and others, began to codify the brain–mind relationship in a way that emphasized testable predictions and an understanding of the computations the system must perform. The resulting synthesis included a recognition that strong interpretive statements about neurophysiology should be predicated on a robust characterization of behavior (e.g. psychophysics) and a mathematical framework that connects this behavior to a postulated internal state.

Several authors have recently articulated the basic principles of psychophysics as they pertain to neuroscience [3**,13**,14,15]. Here are a few considerations to keep in mind when designing a behavioral task, and especially when training animal subjects to perform it.

Account for errors

One of the most critical benchmarks is to ensure that subjects are working ‘at threshold’ or nearly so — that is, trained to asymptotic performance and fully engaged in the task. To illustrate this point, imagine we want to understand the stochastic process that leads to errors

in a perceptual judgment. To do this we have to convince ourselves that the observed errors are not generated by a separate process, distinct from the one we wish to study. This is why, whenever possible, it is important to include multiple levels of difficulty that span the range from (near) perfect to chance performance, and to vary them randomly from trial to trial. If your subjects can perform near 100% correct on the easiest trials, even when they cannot predict trial difficulty in advance, you can rule out a large class of so-called ‘lapses’ — also termed guesses or random choices — for explaining behavior on more difficult trials.

It might seem like a maximum of 85–90% correct is sufficient, but consider that a lapse rate of 10–15% on a binary choice task — not uncommon in animal studies, particularly with rodents — implies that the subject is guessing on 20–30% of the trials (because half the guesses will be correct). Crucially, these trials are random from the experimenter’s point of view and cannot simply be detected and removed. No scientist would tolerate a software glitch that covertly replaced a substantial fraction of their data with random numbers, yet this is essentially what a high lapse rate permits, at least in principle. Although sophisticated behavioral modeling [16] can isolate the contribution of lapses from other sources of variability, this only applies to probabilities over ensembles of trials. It does not prevent lapses from corrupting attempts to relate single-trial neural activity to behavior. It is worth noting here that rodents are capable of achieving very low lapse rates, at least in certain tasks [17]. However, if high performance cannot be demonstrated, this may be a sign that the chosen species is not a good model for the task or process being studied.

In summary, experimenters should make every attempt to design tasks and training procedures such that subjects are capable of near-perfect performance on the easiest conditions, even though most analyses ought to focus on the difficult ones.¹ Along with selecting an appropriate task structure (e.g. 2-alternative forced choice) [3^{**},18] and normative modeling framework, this strategy will facilitate an accounting of how different types of errors arise. This issue seems particularly important for interpreting large-scale, exploratory studies, which may be more likely to zero in on features of neural activity that end up being red herrings due to uncontrolled behavioral variables.

Care about time

Time is fundamental to cognition, not only in the explicit way in which it supports abilities like learning and prediction, but more generally for structuring and regulating internal processes. The real world is not partitioned into discrete trials [13^{**}], so the brain has evolved to implement rules for terminating or switching between processes in the absence of external cues. In many cases, simply measuring the time it takes subjects to respond can be a powerful constraint for models of the internal process governing behavior [19–21]. Most importantly for the present topic, response-time (RT) tasks [22,23] are often the best way to identify the relevant time window(s) for analysis of neural data.

¹Indeed concentrating the bulk of trials on difficult conditions (e.g., varying discriminability on a logarithmic scale) can be an effective way to ensure sufficient behavioral variability for testing models of the underlying process. This strategy also maximizes one’s sensitivity for detecting differences between conditions — or effects of causal manipulations — that may be subtle but nonetheless provide key insights.

For example, many types of decisions involve the accumulation of evidence up to a threshold, or bound [24]. To study this process and its neural correlates it is essential to demonstrate behaviorally that the subject is indeed accumulating evidence. With well-trained subjects and an appropriate balance between reward and punishment (e.g. to discourage fast guessing), a choice-RT design allows the experimenter to analyze neural data restricted to the epoch of decision formation, barring sensory and motor delays.

In cases where RT measurements are not attainable or inappropriate for the task, other temporal constraints can be quite useful. One approach is to ‘compel’ subjects to initiate the motor response at a specific time [25], even — remarkably — before the sensory information for the decision becomes available [26]. These designs help to isolate the decision process from motor preparation and other variables such as the speed-accuracy trade-off. Alternatively, one can embrace this trade-off and simply vary the stimulus duration, using accuracy as a function of time to infer the presence of a bound and to estimate the time window of accumulation [27,28]. To promote engagement and discourage procrastination, it is often advisable to draw durations from a distribution with an early peak and a long tail, such as an exponential. In contrast, durations that are fixed [29,30] — or uniformly distributed with a relatively long minimum [31**] — may permit subjects to delay accumulation, or worse, vary its timing unpredictably across trials, thus complicating the interpretation of neural data.

These issues would be esoteric were it not for the ubiquity of putative integration or diffusion-like mechanisms in the literature. What is less widely recognized is that searching for their neural implementation must be accompanied by behavioral evidence that such a process is actually taking place, and in what time frame.

Lessons from computational motor control

So far I have argued that the rigorous measurement and modeling of behavior in the psychophysical tradition can benefit the study of complex cognitive functions as well. A similar payoff can accrue to those who borrow experimental methods and computational approaches from the field of motor control [5–7]. Furthermore, it turns out the motor system may offer a more direct glimpse of cognitive processing than is widely appreciated.

In a series of studies [32,33,34*], Cos, Cisek and colleagues analyzed arm kinematics and electromyographic (EMG) activity while human subjects decided between two reach targets in a free-choice task. The choice options differed in their biomechanical costs, and the authors also varied factors such as requirements for precision and stability. The results [32,33] suggested that the decision process takes these factors into account well before movement onset, and does so very rapidly after visual specification of the choice options [34*].

The latter study exploited the known relationship between evoked EMG potentials and corticospinal excitability, using transcranial magnetic stimulation (TMS) of motor cortex to probe the state of motor preparation at different times within a trial. Similarly, Selen *et al.* [35] built upon what was known about intermediate-latency and long-latency stretch reflexes

[36] to discover that the motor system receives a continuous flow of information related to a perceptual task. Subjects discriminated the direction of motion in a random-dot display and indicated their choice with a reaching movement while grasping a manipulandum. Selen *et al.* [35] then perturbed the manipulandum at random times to initiate a reflexive EMG response whose gain reflects the current preparatory state. They found that these reflex gains tracked the evolving decision variable (i.e. accumulated sensory evidence), as though a provisional plan of action was being continuously fed to motor-related structures in the brain during decision formation — an idea supported by neurophysiology in oculomotor tasks as well [24,37–39].

These and other findings [40] are consistent with the view that decisions, at least when the outcome is communicated with a motor act, arise as a competition between provisional behaviors or intentions [41,42], rather than being completed in a space of goods [43] or abstract categories [44] prior to action selection. But no matter where the reader stands on this particular issue, the key point is that insights into ostensibly covert cognitive processing can be made through careful measurement and manipulation of the motor system, guided by normative theory (e.g. optimal feedback control) [5,36]. Just as an accounting of the source of errors is an advantage of threshold psychophysics, so can a detailed understanding of movement and its variability help limit the range of explanations applied to neuronal datasets. It seems possible that better measurement and control of motor processes could help explain some of the ‘perplexing’ features of neuronal responses in higher brain areas, features that have recently motivated a reliance on complex dynamical structures and decoding schemes.

Pitfalls and limitations of large-scale population analyses

It is tempting to think that our understanding of the brain would immediately take a giant leap forward if only we could measure more of it. To be sure, major discoveries throughout history have often followed from technological advances that allowed greater access to the components of a system. We should not, however, allow the promise of big data to excuse obfuscation or sloppy thinking [1]. Neither should the sophistication of statistical methods dictate what sort of result constitutes an established finding about neural mechanism, as opposed to a description of the data or exploratory analysis for hypothesis generation [45,46].

The reverse inference problem

Naturally, larger and more complex datasets create new challenges for statistical inference. This has been a topic of great interest in the neuroimaging community, with several high-profile critiques highlighting the perils of improper correction for multiple comparisons [47,48] and the inflation of effect size from non-independence or circularity errors [49,50]. A more general problem is the inherent weakness of so-called reverse inference [51]: the attribution of specific mental processes based on brain activity patterns. With fMRI, the validity of reverse inference depends on how frequently a pattern of activity occurs across a variety of experimental contexts — the ‘base rate’ or prior probability of the observation. Certain regions or networks tend to be active in many studies across multiple cognitive

domains, thus it can be problematic to make strong claims about particular mental states based on such activation.

Although this problem is exacerbated by the coarse spatial and temporal resolution of fMRI, even cellular-resolution imaging and multielectrode physiology are not immune. Hypothetically, suppose that some interesting and novel feature of population activity is found to be associated with an experimental manipulation or task. This is a perfectly reasonable basis for forward inference (how the brain responds to the manipulation), but it does not imply that the feature underlies the cognitive process being investigated,² nor even that its presence is particularly diagnostic. Again it depends, among other things, on the base rate of the observation, which typically cannot be estimated in a single study using these techniques.

An example from the motor system may again be illustrative. Churchland *et al.* [52*] reported that the evolution of population activity in motor and premotor cortex exhibits a striking rotational structure during reaching movements, one that seems well suited for generating the measured EMG responses. Several data-shuffling controls indicated that rotational dynamics are not trivially observed in any dataset of similar size and complexity, nor were they observed when the animal was stationary. In a sense, this rules out a high base rate for such dynamics, but only within the actual recorded population. One might wonder how often rotations could be extracted from populations in other brain areas, including those unrelated to movement preparation. Indeed, partly in response to this concern, Churchland and colleagues (Churchland, Lara, and Cunningham, *Soc. Neurosci. Abs.* 2015; Russo, Perkins, London, and Churchland, *Soc. Neurosci. Abs.* 2015) recently reported an absence of strong dynamics or rotational structure, both upstream (supplementary motor area) and downstream (the muscles) of primary motor cortex, even when the movement itself was periodic. These findings strengthen the conclusions of their earlier work [52*], and provide a template for how to deal with the reverse inference issue as it applies to large-scale neurophysiology.

What we can and cannot learn from high-dimensional structures

Like Churchland *et al.* [52*], many recent studies (e.g. [30,53–56]) describe the activity of a large neural population as a coordinate in a high-dimensional space, where the number of dimensions is equal to the number of neurons. Typically some form of dimensionality-reduction is performed, and the positions or trajectories within the reduced state space (and/or performance of a statistical classifier) are used to infer something about the neural representation. This strategy is clearly worthwhile —indeed often necessary — for visualizing the data and furnishing a conceptual framework in which to generate new hypotheses [45,46,57]. It is generally unwarranted, however, to jump to the conclusion that state-space dynamics or advanced decoding algorithms constitute a mechanism for biological neural computation. Take for instance the observation that a linear decoder can be trained to read out task-related variables from a population whose constituent neurons respond with complex or ‘mixed’ selectivity [53,54]. This provides a useful description of

²Here I am not making the obvious point that correlation does not imply causation, but rather highlighting the difference between weak and strong inference from correlation.

the signals available in the data, but says very little about whether the brain takes advantage of these signals in a way remotely like the experimenter does.

In an analysis of prefrontal cortex activity during two different working memory tasks, Rigotti and colleagues [53] elegantly demonstrated the potential role of nonlinear mixed selectivity, including its relationship to dimensionality and how such representations confer an advantage for linear classifiers. Intriguingly, the dimensionality of the neural representation collapsed on error trials, suggesting a behavioral role for high dimensionality. However, this interpretation suffers from an incomplete understanding of the source of errors, that is, whether they indicate failures of working memory, behavioral response selection, or were simply lapses. The fact that stimulus identity could still be decoded with high accuracy on error trials was taken as support for the second possibility [53], but this presupposes the particular read-out strategy being evaluated.

Raposo *et al.* [54] analyzed posterior parietal cortex activity while rats performed a multisensory event-rate discrimination task. They found that, although individual neurons did not fall into clear functional categories, the population as a whole could be read out to predict the animal's choice and the stimulus modality (visual or auditory). However, the stronger claim [54] that activity is decoded in a particular way to support performance —or that exploration of different state space dimensions in different epochs plays a causal role — seems premature without tighter links to behavior and a biologically plausible model for combining sensory inputs and forming a decision. This brings up the related point that decoding analyses should take the perspective of the subject rather than the experimenter. The task of the subject (and thus the neural population) is to determine, from the evidence, which choice is the correct one — not which choice it is about to make, and not a task-irrelevant property of the trial such as stimulus modality [54].

Returning to fMRI, similar caveats — and other, more nuanced ones [58,59] — apply to the class of methods known as multivariate (or multi-voxel) pattern analysis (MVPA), which also uses machine-learning algorithms to extract information from a distributed pattern of activity. This type of analysis can reveal the representational content of the BOLD signal across a set of voxels, but not whether or how that information is used by the rest of the brain. Again the point here is deeper than just 'correlation is not causation'. It is that the way a classifier solves a problem — ignoring connectivity and limitations on synaptic plasticity — likely bears no resemblance to what the brain is actually doing, and may not say much about where it is doing it [58].³

Lastly, at the risk of being accused of Luddism, I would submit an appeal to the simplicity and directness of an earlier time. Imagine if Hubel and Wiesel had recorded many neurons simultaneously while animals viewed natural movies. Complex dynamics would surely have been observed, and a support vector machine might have been put forth as a candidate mechanism for vision. But one wonders what would have become of the simple and powerful insights arising from the selectivity of single cells for orientation, motion, color,

³Some skeptics of neuroimaging (and others) will find these points obvious, but may not have considered the parallels with recent large-scale neural recording studies that seem to elude such criticism.

and binocular disparity [60] — not to mention faces, toilet brushes and so on. The trajectory of systems neuroscience may have been much different.

Of course this is somewhat tongue-in-cheek, as cognition is not vision (although it seems unlikely that evolution would produce radically different neural mechanisms for the two functions [41]). Arguably, however, the framing of hypotheses in terms of computations reflected in the activity of single neurons [12,61–64] was not from a lack of imagination or statistical prowess on the part of the founders of systems neuroscience. It was based on careful observations and reasoning about the structure of the nervous system and the specific computations required of a given task.

It has been said that computations are the bridge between neural circuits and behavior [65], but not all computations are equally likely to serve this purpose. As the field continues to enumerate a set of fundamental or ‘canonical’ neural operations [13**,65–67], it is worth being mindful and transparent about the distinction between a tool employed by scientists and a mechanism employed by the brain. It may turn out to be prophetic that the statistics of population activity in higher brain areas are well suited for the application of machine learning classifiers. And perhaps it is true that neural networks wire themselves up in just the right way to permit computation with state space trajectories governed by recurrent dynamics. There are reasons to believe the answers will often be more straightforward than that (e.g., constraints of anatomy and biological noise [68], and elements of Barlow’s ‘neuron doctrine’ [12]), if only we ask the right questions. Either way, we cannot hope to find clarity unless neural data collection is coupled with well-motivated and controlled behavioral measurements.

Concluding remarks

The preceding commentary is offered in a spirit of goodwill, and with several caveats, first and foremost that the author is manifestly not an expert in many of the methods and topics discussed. I do not wish to imply that all experimental variables or computations should be observable at the level of single neurons, or that behavioral methods must be refined until they are. I am not saying that machine learning and dynamical systems approaches cannot have explanatory value on their own. And I would not suggest that all cognitive functions can be construed as a psychophysical task, or fully understood within a framework inspired by motor control. What I do hope to convey is that the philosophy and toolkit of traditional systems neuroscience can still contribute much in the way of rigor and parsimony, even (or especially) in the era of large-scale recording and sophisticated data analysis.

The BRAIN initiative and related ventures have focused much-needed attention on the development of tools for measuring and perturbing activity across large neural circuits. This effort will surely revolutionize our understanding of normal and diseased brain function. But if we wish to decipher the neural basis of cognition, we must strive for a better understanding of the cognitive states we wish to explain, states that can only be revealed through the lens of behavior.

Acknowledgments

The author's postdoctoral fellowship is supported by the Howard Hughes Medical Institute. I thank members of the Shadlen lab for discussions, especially Mehdi Sanayei, Shushruth, and Luke Woloszyn who contributed specific points and examples. I am especially grateful for the support of my mentor Michael Shadlen, whose influence pervades this essay.

References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest

- of outstanding interest

1. Sejnowski TJ, Churchland PS, Movshon JA. Putting big data to good use in neuroscience. *Nat Neurosci.* 2014; 17:1440–1441. [PubMed: 25349909]
2. Freeman J, Vladimirov N, Kawashima T, Mu Y, Sofroniew NJ, Bennett DV, Rosen J, Yang C-T, Looger LL, Ahrens MB. Mapping brain activity at scale with cluster computing. *Nat Meth.* 2014; 11:941–950.
- 3•• Carandini M, Churchland AK. Probing perceptual decisions in rodents. *Nat Neurosci.* 2013; 16:824–831. A timely and insightful review of recent efforts to bring the rigorous study of perceptual decision making to rodent models, which have several advantages including well developed molecular and circuit-manipulation tools. The authors cover many of the same points touched on here, plus other important considerations for experimental design, guided by concepts from psychophysics and signal detection theory. [PubMed: 23799475]
4. Ahrens MB, Engert F. Large-scale imaging in small brains. *Curr Opin Neurobiol.* 2015; 32:78–86. [PubMed: 25636154]
5. Franklin DW, Wolpert DM. Computational mechanisms of sensorimotor control. *Neuron.* 2011; 72:425–442. [PubMed: 22078503]
6. Scott SH. The computational and neural basis of voluntary motor control and planning. *Trends Cogn Sci.* 2012; 16:541–549. [PubMed: 23031541]
7. Lisberger SG, Medina JF. How and why neural and motor variation are related. *Curr Opin Neurobiol.* 2015; 33:110–116. [PubMed: 25845626]
8. Hubel D. Vision and the brain. *Bull Am Acad Arts Sci.* 1978; 31:17–28.
9. Marr, D. Vision. WH Freeman & Co; 1982.
10. Brindley, GS. Physiology of the Retina and the Visual Pathway. Edward Arnold Ltd; 1960.
11. Teller DY. Linking propositions. *Vis Res.* 1984; 24:1233–1246. [PubMed: 6395480]
12. Barlow HB. Single units and sensation: a neuron doctrine for perceptual psychology? *Perception.* 1972; 1:371–394. [PubMed: 4377168]
- 13•• Shadlen MN, Kiani R. Decision making as a window on cognition. *Neuron.* 2013; 80:791–806. This review summarizes a 25-year body of work, from the linking of visual neurons to psychophysical performance, to the quantitative study of higher aspects of decision-making, such as confidence. The authors walk us through the rationale behind studying simple perceptual decisions to expose fundamental computational principles that can be seen as building blocks for more complex cognitive functions. They also briefly revisit how considerations of E-I balance and noise in cortical circuits (see [68]) lend support to Barlow's notion of a 'direct code', as opposed to a complex representation distributed across a diverse neural population. [PubMed: 24183028]
14. Fetsch CR, Deangelis GC, Angelaki DE. Bridging the gap between theories of sensory cue integration and the physiology of multisensory neurons. *Nat Rev Neurosci.* 2013; 14:429–442. [PubMed: 23686172]
15. Read JCA. The place of human psychophysics in modern neuroscience. *Neuroscience.* 2015; 296:116–129. [PubMed: 24880153]

16. Erlich JC, Brunton BW, Duan CA, Hanks TD, Brody CD, Carandini M. Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *Elife*. 2015; 4:e05457.
- 17•. Histed MH, Carvalho LA, Maunsell JHR. Psychophysical measurement of contrast sensitivity in the behaving mouse. *J Neurophysiol*. 2012; 107:758–765. Mice were trained to perform a visual detection task, and about half achieved reliable and stable performance with very low lapse rates. In addition to establishing a genetically tractable model for studying the neural basis of visual perception, this study argues against a fundamental limit on rodent psychophysical performance — albeit with the caveat that such a limit may exist for more complex tasks for which nonhuman primates are often better suited. [PubMed: 22049334]
18. Morgan MJ, Melmoth D, Solomon JA. Linking hypotheses underlying Class A and Class B methods. *Vis Neurosci*. 2013; 30:197–206. [PubMed: 24476966]
19. Stone M. Models for choice-reaction time. *Psychometrika*. 1960; 25:251–260.
20. Link SW, Heath RA. A sequential theory of psychological discrimination. *Psychometrika*. 1975; 40:77–105.
21. Vickers, D. *Decision Processes in Visual Perception*. Academic Press; 1979.
22. Roitman JD, Shadlen MN. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci*. 2002; 22:9475–9489. [PubMed: 12417672]
23. Siegel M, Buschman TJ, Miller EK. Cortical information flow during flexible sensorimotor decisions. *Science*. 2015; 348:1352–1355. [PubMed: 26089513]
24. Gold JJ, Shadlen MN. The neural basis of decision making. *Annu Rev Neurosci*. 2007; 30:535–574. [PubMed: 17600525]
25. Resulaj A, Kiani R, Wolpert DM, Shadlen MN. Changes of mind in decision-making. *Nature*. 2009; 461:263–266. [PubMed: 19693010]
26. Stanford TR, Shankar S, Massoglia DP, Costello MG, Salinas E. Perceptual decision making in less than 30 milliseconds. *Nat Neurosci*. 2010; 13:379–385. [PubMed: 20098418]
27. Kiani R, Hanks T, Shadlen M. Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J Neurosci*. 2008; 28:3017–3029. [PubMed: 18354005]
28. de Lafuente V, Jazayeri M, Shadlen MN. Representation of accumulating evidence for a decision in two parietal areas. *J Neurosci*. 2015; 35:4306–4318. [PubMed: 25762677]
29. Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci*. 1996; 13:87–100. [PubMed: 8730992]
30. Mante V, Sussillo D, Shenoy KV, Newsome WT. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*. 2013; 503:78–84. [PubMed: 24201281]
- 31••. Latimer KW, Yates JL, Meister MLR, Huk AC, Pillow JW. Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science*. 2015; 349:184–187. This study used state-of-the-art statistical techniques to ask whether spike trains in the lateral intraparietal area (LIP) reflect a latent diffusion-like signal (i.e. gradual accumulation of evidence) or are more consistent with discrete steps to a stereotyped low or high state. The analysis favored the stepping model for the majority of neurons. However, the behavioral task was unable to identify the time window of decision formation, owing to the relatively long, experimenter-controlled viewing durations. This highlights the importance of behavioral control in the temporal dimension for isolating neural computations underlying cognition. [PubMed: 26160947]
32. Cos I, Bélanger N, Cisek P. The influence of predicted arm biomechanics on decision making. *J Neurophysiol*. 2011; 105:3022–3033. [PubMed: 21451055]
33. Cos I, Medleg F, Cisek P. The modulatory influence of end-point controllability on decisions between actions. *J Neurophysiol*. 2012; 108:1764–1780. [PubMed: 22773776]
- 34•. Cos I, Duque J, Cisek P. Rapid prediction of biomechanical costs during action decisions. *J Neurophysiol*. 2014; 112:1256–1266. This is the latest in a series of studies showing that the formation of a decision about where to move the arm is influenced by the biomechanical cost (effort) of the entire anticipated trajectory, even before the onset of movement. Here they provided subjects with a variable amount of time to view the configuration of potential movement

targets, and showed that 200 ms of viewing time was sufficient to bias choices based on predicted effort. They also perturbed primary motor cortex with TMS and found that the resulting EMG potentials (which reflect corticospinal excitability) indicated an influence of predicted effort within as little as 150 ms. The approach highlights the possibility of using quantitative measurement and manipulation of the motor system as a window onto fairly sophisticated cognitive processes, such as balancing reward versus cost in decision making. [PubMed: 24899673]

35. Selen LPJ, Shadlen MN, Wolpert DM. Deliberation in the motor system: reflex gains track evolving evidence leading to a decision. *J Neurosci.* 2012; 32:2276–2286. [PubMed: 22396403]
36. Pruszynski JA, Scott SH. Optimal feedback control and the long-latency stretch response. *Exp Brain Res.* 2012; 218:341–359. [PubMed: 22370742]
37. Kustov AA, Robinson DL. Shared neural control of attentional shifts and eye movements. *Nature.* 1996; 384:74–77. [PubMed: 8900281]
38. Kable JW, Glimcher PW. The neurobiology of decision: consensus and controversy. *Neuron.* 2009; 63:733–745. [PubMed: 19778504]
39. Ding L, Gold JJ. The basal ganglia's contributions to perceptual decision making. *Neuron.* 2013; 79:640–649. [PubMed: 23972593]
40. McKinstrey C, Dale R, Spivey MJ. Action dynamics reveal parallel competition in decision making. *Psychol Sci.* 2008; 19:22–24. [PubMed: 18181787]
41. Shadlen, MN.; Kiani, R.; Hanks, TD.; Churchland, AK. Neurobiology of decision making: an intentional framework. In: Engel, C.; Singer, W., editors. *Strüngmann Forum Report: Better Than Conscious? Decision Making, the Human Mind, and Implications for Institutions.* 2008. p. 1-32.
42. Cisek P, Kalaska JF. Neural mechanisms for interacting with a world full of action choices. *Annu Rev Neurosci.* 2010; 33:269–298. [PubMed: 20345247]
43. Padoa-Schioppa C. Neurobiology of economic choice: a good-based model. *Annu Rev Neurosci.* 2011; 34:333–359. [PubMed: 21456961]
44. Freedman DJ, Assad JA. A proposed common neural mechanism for categorization and perceptual decisions. *Nat Neurosci.* 2011; 14:143–146. [PubMed: 21270782]
45. Sussillo D. Neural circuits as computational dynamical systems. *Curr Opin Neurobiol.* 2014; 25:156–163. [PubMed: 24509098]
46. Cunningham JP, Yu BM. Dimensionality reduction for large-scale neural recordings. *Nat Neurosci.* 2014; 17:1500–1509. [PubMed: 25151264]
47. Bennett CM, Baird AA, Miller MB, Wolford GL. Neural correlates of interspecies perspective taking in the post-mortem Atlantic salmon: an argument for proper multiple comparisons correction. *J Serendip Unexpect Results.* 2011; 1:1–5.
48. Nichols T, Hayasaka S. Controlling the familywise error rate in functional neuroimaging: a comparative review. *Stat Methods Med Res.* 2003; 12:419–446. [PubMed: 14599004]
49. Vul E, Harris C, Winkielman P, Pashler H. Puzzlingly high correlations in fMRI studies of emotion, personality, and social cognition. *Perspect Psychol Sci.* 2009; 4:274–290. [PubMed: 26158964]
50. Kriegeskorte N, Simmons WK, Bellgowan PSF, Baker CI. Circular analysis in systems neuroscience: the dangers of double dipping. *Nat Neurosci.* 2009; 12:535–540. [PubMed: 19396166]
51. Poldrack RA. Inferring mental states from neuroimaging data: from reverse inference to large-scale decoding. *Neuron.* 2011; 72:692–697. [PubMed: 22153367]
52. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV. Neural population dynamics during reaching. *Nature.* 2012; 487:51–56. An analysis of premotor and motor cortex activity during reaching movements revealed unexpected patterns of neural coactivation, described in state space as quasi-oscillatory or rotational structure. Several controls and follow-up studies suggested that this structure is not trivially observed — that is, it does not have a high prior probability or ‘base rate’. Thus, it appears to indicate something unique about (pre)motor cortex, namely that it establishes specific internal dynamics to serve as a basis for outgoing motor commands. [PubMed: 22722855]

53. Rigotti M, Barak O, Warden MR, Wang X-J, Daw ND, Miller EK, Fusi S. The importance of mixed selectivity in complex cognitive tasks. *Nature*. 2013; 497:585–590. [PubMed: 23685452]
54. Raposo D, Kaufman MT, Churchland AK. A category-free neural population supports evolving demands during decision-making. *Nat Neurosci*. 2014; 17:1784–1792. [PubMed: 25383902]
55. Harvey CD, Coen P, Tank DW. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*. 2012; 484:62–68. [PubMed: 22419153]
56. Stokes MG, Kusunoki M, Sigala N, Nili H, Gaffan D, Duncan J. Dynamic coding for cognitive control in prefrontal cortex. *Neuron*. 2013; 78:364–375. [PubMed: 23562541]
57. Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV. Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron*. 2010; 68:387–400. [PubMed: 21040842]
58. Anderson, ML.; Oates, T. A critique of multi-voxel pattern analysis. In: Ohlsson, Catrambone R., editor. *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society*; 2010; p. 1511-1516.
59. Haynes J-D. A primer on pattern-based approaches to fMRI: principles, pitfalls, and perspectives. *Neuron*. 2015; 87:257–270. [PubMed: 26182413]
60. Rust NC, Movshon JA. In praise of artifice. *Nat Neurosci*. 2005; 8:1647–1650. [PubMed: 16306892]
61. Werner G, Mountcastle VB. The variability of central neural activity in a sensory system, and its implications for the central reflection of sensory events. *J Neurophysiol*. 1963; 26:958–977. [PubMed: 14084169]
62. Johnson KO. Sensory discrimination: decision process. *J Neurophysiol*. 1980; 43:1771–1792. [PubMed: 6251182]
63. Parker AJ, Newsome WT. Sense and the single neuron: probing the physiology of perception. *Annu Rev Neurosci*. 1998; 21:227–277. [PubMed: 9530497]
64. Movshon JA, Tolhurst DJ. Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *J Physiol (Lond)*. 1978; 283:53–77. [PubMed: 722589]
65. Carandini M. From circuits to behavior: a bridge too far? *Nat Neurosci*. 2012; 15:507–509. [PubMed: 22449960]
66. Carandini M, Heeger DJ. Normalization as a canonical neural computation. *Nat Rev Neurosci*. 2011; 13:51–62. [PubMed: 22108672]
67. Pack CC, Bensmaia SJ. Seeing and feeling motion: canonical computations in vision and touch. *PLoS Biol*. 2015; 13:e1002271. [PubMed: 26418156]
68. Shadlen MN, Newsome WT. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci*. 1998; 18:3870–3896. [PubMed: 9570816]