

Three challenges for connecting model to mechanism in decision-making

Anne K Churchland¹ and Roozbeh Kiani²



Recent years have seen a growing interest in understanding the neural mechanisms that support decision-making. The advent of new tools for measuring and manipulating neurons, alongside the inclusion of multiple new animal models and sensory systems has led to the generation of many novel datasets. The potential for these new approaches to constrain decision-making models is unprecedented. Here, we argue that to fully leverage these new approaches, three challenges must be met. First, experimenters must design well-controlled behavioral experiments that make it possible to distinguish competing behavioral strategies. Second, analyses of neural responses should think beyond single neurons, taking into account tradeoffs of single-trial versus trial-averaged approaches. Finally, quantitative model comparisons should be used, but must consider common obstacles.

Addresses

¹ Cold Spring Harbor Laboratory, United States

² Center for Neural Science, New York University, United States

Current Opinion in Behavioral Sciences 2016, 11:74–80

This review comes from a themed issue on **Computational modeling**

Edited by **Peter Dayan** and **Daniel Durstewitz**

<http://dx.doi.org/10.1016/j.cobeha.2016.06.008>

2352-1546/© 2016 Elsevier Ltd. All rights reserved.

Introduction

Major strides in our understanding of the neural mechanisms of decision-making were made by a powerful approach: studying visual decisions in human [1,2] and nonhuman [3] primates alongside single-neuron recording to evaluate potential underlying mechanisms. This approach generated key insights in the field, including an appreciation for the circumstances that lead subjects to integrate visual information over time and an opportunity to narrow down the neural mechanisms that might support such choices via carefully designed analyses of neural responses [4,5,6^{**},7–12].

In recent years, this approach has been augmented in a number of ways. First, many new animal models are used alongside primates, including rodents [13,14^{*},15] and invertebrates [16,17^{*}]. Further, the focus on visual stimuli has expanded; new studies include decisions guided by

olfactory [18,19], auditory [20,21], somatosensory [22–24], gustatory [25], and multisensory [26,27] stimuli. Finally, a wealth of new techniques for measuring and manipulating neurons has drastically changed the kind of data that is available to investigate decision-making mechanisms in the brain. These include the ability to monitor many neurons simultaneously [28,29^{*},30], and the opportunity to target neural populations defined by cell type or circuit [31,32,33^{**}]. These techniques provide a new view on neural activity during decision-making and have the potential to provide important new insights into underlying neural mechanisms.

The new animal models, modalities and techniques mean that the field is poised to make great strides in tackling unsolved problems in perceptual decision-making. However, the rapid changes necessitate a consideration of what aspects of experimental design are fundamental for advancing our understanding of decision-making. In this review, we argue that a shared understanding in three key areas is needed to fully leverage the tools and approaches that are in the field today. These are: designing behavioral experiments to afford insight into subject's strategy, analyzing population level neural activity and finally, avoiding obstacles when using these measurements to distinguish candidate models.

Well-controlled experiments to distinguish alternative behavioral strategies

Animals in laboratory tasks are skilled at developing strategies that lead to reward, but these do not always match the strategy that the experimenter had in mind. Determining how animals perform a task is challenging, but it is a necessity when the subject's strategy can influence the interpretation of results. Studies of the decision-making process are particularly susceptible to such misinterpretations. Animals may not uniformly adopt the best strategy because they misunderstand the task structure or because experimenters fail to constrain the solutions to the task. Special attention must be paid to an animal's training history and experimental interventions that shape the behavior. These can instill suboptimal strategies, or even worse, introduce complex reorganization of the neural circuits that furnish the behavior [34]. Experimenters should also employ appropriate analytical tools and control experiments to detect and verify strategies that underlie the behavior.

The need for these analytical tools is underscored by the fact that similar behavioral patterns could arise from

different strategies. For example, 10% lapse rate in a psychometric function could happen because the task is too difficult or because the subject elects to disengage and respond randomly on a large fraction of trials (20% in a 2-AFC task). A difficult task may be favored, especially for studying threshold-level behavior, but random behavior on 20% of trials can cause major problems for interpreting data, just as no experimenter would want a device that behaves randomly 20% of the time. A similar problem arises in value based decisions and foraging tasks where changes in the behavior can be attributed to either noisy integration of past choices and outcomes, or to random switches for further exploration [35–37]. Identifying the true strategy is critical for interpreting neural data. For some aspects of behavior, identification of strategy is extremely challenging (e.g., lapse rate in a trained animal). For some others, it is possible to distinguish different hypotheses using a combination of experimental design and targeted models [21,38,39,40**]. Two recent examples stand out. First, Gold and colleagues used these methods to show that in monkeys engaged in perceptual decisions, trial-to-trial variability of choice behavior stems from the influence of prior trials [41] (this has also been noted in mice; [42]). Further, the relative influence of prior trials and sensory evidence on a choice is shaped by training. Prior influences are strongest when perceptual sensitivity to the relevant sensory evidence is weakest and then decline steadily during training as sensitivity improves. Second, Scott and colleagues used a model based approach to interpret lapse rates on judgments about stimulus number [43]. Their model included noise that scaled with the number of stimuli; hence the high stimulus numbers that defined some easy trials were inherently error-prone.

Post hoc analyses can be a powerful tool for affording insight into an animal's strategy. A prominent class of such analyses borrows from a classic technique used to map receptive fields in visual areas using stimuli that fluctuate stochastically over time [44,45]. In behaving animals, experimenters can use stimuli that similarly fluctuate and track how these fluctuations relate to behavior. For example, when the strength of a stimulus (its motion energy, for example) fluctuates over time, experimenters can leverage those fluctuations to gain insight into which moments of a stimulus presentation influence an eventual decision. This analysis can distinguish strategies in which animals tend to favor early versus late evidence presented during decision formation [6**,12,21,46**,47**] (Figure 1). Similarly, in perceptual judgments about visual stimuli, a post hoc analysis of stimulus fluctuations can reveal an animal's internal estimate of the category boundary that separates one class of stimuli from another [48,49]. In some cases, this analysis uncovers that the animal's internal category boundary differs from that set by the experimenter, contributing to suboptimal performance.

Analysis of neural responses: thinking beyond single neurons


Neural measurements are inherently noisy. Cortical neurons elicit different patterns of spikes from trial-to-trial even when the incoming sensory stimulus is identical [50]. Appropriately handling this variability is an essential component of data analysis. Traditional data analyses typically average the responses of many trials together (trial averaging) to better estimate the single-neuron response to each trial. Often single neurons are then themselves averaged (neuron averaging), generating a population peristimulus time histogram. These averaging techniques allow experimenters to acquire a better estimate of the underlying mean, potentially affording insight into neural mechanism. Further, because this approach uses stimulus parameters optimized for each neuron, it can focus on a small population that may be most relevant for a decision (e.g., rightward-selective and leftward-selective MT neurons for discrimination of rightward and leftward motion). These studies have laid the foundation for understanding the neural mechanisms of decision-making and will continue to be influential in the future. In this section, we explain how recent work has highlighted some of the shortcomings of the traditional single-neuron approach and has provided alternatives. We also explain why alternatives to the traditional approach have their own shortcomings; these are tractable, but have yet to receive sufficient attention.

Trial averaging can obscure trial-to-trial dynamics

Averaging across trials can obscure important links between neural responses and behavior. For example, consider an experimenter who wished to understand how idiosyncratic decision biases are reflected in neural data. Because biases depend on recent reward history [41,42], averaging across many trials would remove the signal that is of interest to the experiment. Instead, measuring the responses of many neurons on a single trial can provide insight into how the network changes for biased versus unbiased decisions. A second example is changes of mind: subjects sometimes revise a decision mid-trial [51–54,55*,56]. A signature of this can be evident in the data, but because changes of mind take place at different moments on different trials, trial averaging will obscure the effect. Finally, trial averaging can obscure temporal dynamics, for example, by temporally blurring transitions which occur abruptly [57].

The advantages of single trial analyses are beginning to be accepted. Less often discussed is a consideration of how both single-trial analysis and traditional trial averaging involve tradeoffs. A shortcoming of single-trial analysis is that stochastic fluctuations in spikes could be interpreted as signal when they are in fact just related to the spike generation process [58,59]. A single spike train provides limited insight into the mean, variance, and moment-to-moment dynamics of a neural response.

Figure 1



Stochastic fluctuations in stimulus signal intensity can offer insight into behavioral strategy. *Top*: Schematic single-trial traces showing stimulus intensity (i.e., motion strength or repetition rate) fluctuating over time. Values above zero indicate evidence in favor of one decision category, described as 'right' because the subject might report the decision by making an eye or body movement to the right; values below zero indicate evidence in favor of the other decision category ('left'). *Left*: Examples that ultimately led a (hypothetical) subject to select 'left'. *Right*: Examples that ultimately led a (hypothetical) subject to select 'right'. *Bottom*: Schematic traces reflecting averages over many trials of the kind shown at *top*. Values close to zero (dashed line) indicate moments in time in which the stimulus had little impact on the eventual choice. Negative values indicate evidence at the corresponding time led to a leftwards choice. Positive values indicate evidence at the corresponding time led to a rightwards choice. Colors indicate two possible behavioral strategies. Red: support for a strategy in which subjects increase the weight assigned to evidence as it arrives over time. Early evidence (left side of red traces) is largely ignored (values are close to 0). Blue: support for an alternate strategy in which subjects decrease the weight assigned to evidence as it arrives over time. Late evidence (right side of blue traces) is largely ignored (values are close to 0). *Left*: computed from examples leading to a leftwards choice. *Right*: computed from examples leading to a rightwards choice.

Knowing about these inaccuracies and their magnitude is key to proper analysis of data. Old-fashioned averaging methods would reduce the influence of these inaccuracies on the final interpretation of the data but they do so at the cost of obscuring trial-to-trial variability and other important aspects of response dynamics.

Neuron averaging can obscure population heterogeneity

Averaging responses across neurons is an effective way to handle the reality that firing rates computed from individual neurons can be noisy. This is especially true for experiments in which the use of multiple stimulus strengths and/or multiple sensory modalities lead to a

large number of stimulus conditions, and an imperfect estimate of the underlying firing rate on each one at the level of single neurons.

A shortcoming of averaging neurons is that it relies on the assumption that the parameters of interest in the neurons are reflected uniformly across the population. An alternative possibility is that neurons reflect idiosyncratic combinations of either task parameters or response features [46,60]. If that's the case, averaging might hide response features in data that modulate neurons more sparsely, even if the modulation is consistent and can be easily decoded. Dimensionality reduction methods can reveal a

small number of parameters which, when linearly combined, can capture most of the response variability of each neuron in the population [61]. Targeted dimensionality reduction in which the dimensions largely correspond to user-specified parameters (such as time or stimulus strength) can further aid in such situations, allowing an experimenter to see the timecourse of modulation of a particular parameter, even if it accounts for a small amount of the overall variance [62,63]. Further, these methods can reveal order at the population level when single neurons appear bewilderingly complex [64,65].

Population-level analyses offer an alternative to current averaging approaches, but a shortcoming of such methods is that in their current instantiation, little consideration is given to the user's confidence in the firing rate estimate for each single neuron. In traditional population averaging, a number of methods were used for taking into account the standard error on the estimate of each neuron when combining them together [10]. Current population-level analyses can benefit from methods that adjust the influence of individual neurons based on the reliability of single neuron responses. Important strides are being taken in that direction [66]. Another challenge with population response analyses is that their complexity can make them unintuitive, even for experts. It is sometimes unclear what the expected outcome of the analyses is for alternative hypotheses and how susceptible the results are to measurement noise and neural response variability. Researchers can provide clarity by applying their analyses to synthetic data that are tailored for each hypothesis but share the noise properties of the recorded neural responses [46].

Obstacles to model comparison

Recent advances in computational and systems neuroscience have led to an increase in the number of quantitative models that one can use to explain cognitive and decision-making processes. At the same time, increased accessibility of powerful computers and specialized software has made model selection techniques exceedingly easy to implement. These approaches make it easier to quantitatively compare competing models, which seems, at first glance, to simplify the job of identifying the best ones. However, a number of pitfalls for model comparison mean that a deep understanding of these tools is required in order to avoid errors.

A common pitfall is overgeneralization, wherein researchers compare specific instances of two classes of models but generalize the outcome to all models in the two classes. The goal of model comparison in systems neuroscience is to make statements about specific neural mechanisms, which are often captured by a subset of model parameters. Individual models, however, often have additional parameters and implicit assumptions, the values of which can have a large impact on model

performance. For example, to test whether parietal neural responses represent accumulation of evidence through a gradual buildup or instantaneous change of firing rates, one must also make assumptions about starting time of accumulation, stopping criterion, and spiking statistics [67,68]. Inferential problems arise when the space of 'unimportant' model parameters and assumptions is not adequately explored (e.g., due to fixing some parameters) or when there are complex interactions within the model. Drawing broad conclusions about a neural mechanism based on comparison of specific instantiations of complex, multi-parameter models is susceptible to errors because variations of one parameter can change the model behavior and its fit to experimental data. In the above example, assuming that the starting time of the accumulation process is fixed can falsely reduce the likelihood of accumulation models and bias the conclusions because starting times could vary across parietal neurons [67]. As the complexity of tested models increases, unintended interactions of model parameters and overgeneralization errors become more problematic and deserve extra attention. It is critical to verify implicit model assumptions and understand interactions of all model parameters. Creating hierarchies of nested models and systematic tests of these models can alleviate the overgeneralization pitfall [40,69]. Unfortunately, however, tracking these errors may not be always practical as the space of testable models grows rapidly (often exponentially) with the number of model parameters.

A putative solution to this problem is to compare models in a principled way, such as through the use of Bayesian model comparisons. These leverage Bayes factors — the ratio of averaged likelihood of competing models — to inform model selection. Popular model comparison methods include the Bayesian information criterion (BIC), the Akaike information criterion (AIC), and the deviance information criterion (DIC), which is closely related to AIC. An appealing feature of these criteria is that they can make it possible, at least in theory, to compare models with different numbers of parameters by introducing a penalty term for the number of the parameters in the model [70]. These criteria are useful and have revealed, for example, that individual subjects can differ in their decision-making strategies [71].

Unfortunately, multiple pitfalls can arise from lack of knowledge about appropriate model comparison methods, error functions, and penalty for degrees of freedom. AIC, BIC, and DIC impose different penalties and may produce contradictory results. Lack of a clear understanding about which criterion is appropriate for a model comparison can lead to the selection of an incorrect model. For large sample sizes, AIC tends to penalize inadequately for the number of model parameters and, therefore, is susceptible to favoring complex models that overfit the data. In contrast, BIC tends to penalize excessively for the

number of parameters and favors models that underfit the data. For low sample sizes, the order reverses — AIC underfits. The safest practice is to use an array of model selection criteria and seek consensus among them. A lack of consensus across different criteria often indicates high model uncertainty, which should persuade researchers to revisit their model design. A second pitfall is related to priors for model parameters. Bayesian model fitting and comparisons rely on careful selection of priors [72,73], but information about priors is often lacking in experimental data. When reliable information about priors is lacking, one must ensure the results of model comparison are robust to changes of prior distributions within a biologically plausible range. However, like the last pitfall, this one does not have an easy cure: biologically plausible priors are rarely known and the set of possible distributions can be too big to search systematically. We recommend that researchers do not think about the calculation of BIC or other criteria as the end point of their model selection. Rather, they should use these criteria as a starting point and explore why a model is selected and what drives a superior fit to the data. Only through such a ‘mechanistic’ lens one may hope to generate true insights by employing Bayesian model comparisons.

For the last point in this section we focus on another common misconception about model selection. It is sometimes assumed that a model that passes a cross-validation test (i.e., explains the data it is not trained for or fit to) is exonerated from the abovementioned pitfalls. Unfortunately, that is not necessarily true. Although passing a cross-validation test is necessary for the suitability of a model, it is not sufficient. Further, cross-validation is often a phenomenological criterion, not a mechanistic one, and should be interpreted accordingly. The success of a cross-validation test does not imply that the neural mechanisms suggested by the model are correct. Despite these shortcomings, cross-validation is a useful tool and a good first step for establishing a model, especially when it subjects the model to a novel feature of the data (not just a group of randomly-chosen held-out trials). For instance, demonstrating that a model fit to reaction times can predict an animal’s choice or confidence about the choice is a good indicator that the neural mechanisms implied by the model are worth exploring [40^{••},74,75].

While we believe that quantitative model comparison techniques can advance our ability to distinguish candidate decision-making models, caution is clearly warranted. Overgeneralization must be protected against, and an exclusive reliance on Bayesian information criteria could lead to premature exclusion of candidate models. Instead, Bayesian methods can be used as a starting point, to identify key parameters, thus allowing experimenters to design the right experiments and analyses to robustly distinguish models. Fortunately, our recently-acquired

ability to record and manipulate large populations of neurons while animals are engaged in well-designed decision-making tasks have expanded our experimental repertoire and made incisive, hypothesis-driven experiments increasingly more accessible.

Conflict of interest

Nothing declared.

Acknowledgements

We thank Braden Purcell and Gouki Okazawa for useful discussions and insightful comments. This work was supported by the Simons Collaboration on the Global Brain (AKC & RK), R01EY022979 (AKC), R01MH109180 (RK), the Pew Foundation (AKC), and the Klingenstein Foundation (AKC).

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. Ratcliff R, Rouder JN: **Modeling response times for two-choice decisions.** *Psychol Sci* 1998, **9**:347-356.
2. Carpenter RH, Williams ML: **Neural computation of log likelihood in control of saccadic eye movements.** *Nature* 1995, **377**:59-62.
3. 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.
4. Ratcliff R, Hasegawa YT, Hasegawa RP, Smith PL, Segraves MA: **Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task.** *J Neurophysiol* 2007, **97**:1756-1774.
5. Law CT, Gold JI: **Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area.** *Nat Neurosci* 2008, **11**:505-513.
6. Nienborg H, Cumming BG: **Decision-related activity in sensory neurons reflects more than a neuron’s causal effect.** *Nature* 2009, **459**:89-92.
7. Huk AC, Shadlen MN: **Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making.** *J Neurosci* 2005, **25**:10420-10436.
8. Cohen MR, Newsome WT: **Context-dependent changes in functional circuitry in visual area MT.** *Neuron* 2008, **60**:162-173.
9. Rishel CA, Huang G, Freedman DJ: **Independent category and spatial encoding in parietal cortex.** *Neuron* 2013, **77**:969-979.
10. Churchland AK, Kiani R, Shadlen MN: **Decision-making with multiple alternatives.** *Nat Neurosci* 2008, **11**:693-702.
11. Kiani R, Shadlen MN: **Representation of confidence associated with a decision by neurons in the parietal cortex.** *Science* 2009, **324**:759-764.
12. Kiani R, Hanks TD, Shadlen MN: **Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment.** *J Neurosci* 2008, **28**:3017-3029.
13. Carandini M, Churchland AK: **Probing perceptual decisions in rodents.** *Nat Neurosci* 2013, **16**:824-831.
14. Harvey CD, Coen P, Tank DW: **Choice-specific sequences in parietal cortex during a virtual-navigation decision task.** *Nature* 2012, **484**:62-68.

The authors identified neurons in mouse posterior parietal cortex that responded during visually-guided decisions in virtual reality. Individual neurons tended to respond transiently. However, at the population level,

responses could distinguish left versus right and correct vs. incorrect decisions.

15. Glickfeld LL, Histed MH, Maunsell JH: **Mouse primary visual cortex is used to detect both orientation and contrast changes.** *J Neurosci: Off J Soc Neurosci* 2013, **33**:19416-19422.

16. Bendesky A, Tsunozaki M, Rockman MV, Kruglyak L, Bargmann CI: **Catecholamine receptor polymorphisms affect decision-making in *C. elegans*.** *Nature* 2011, **472**:313-318.

17. Ohyama T, Schneider-Mizell CM, Fetter RD, Aleman JV, Franconville R, Rivera-Alba M, Mensh BD, Branson KM, Simpson JH, Truman JW *et al.*: **A multilevel multimodal circuit enhances action selection in *Drosophila*.** *Nature* 2015, **520**:633-639.

This paper examined how nociceptive and mechanosensory cues are combined by drosophila larvae to select among two possible actions: fast locomotion and escape.

18. Kepecs A, Uchida N, Zariwala HA, Mainen ZF: **Neural correlates, computation and behavioural impact of decision confidence.** *Nature* 2008, **455**:227-231.

19. Smear M, Shusterman R, O'Connor R, Bozza T, Rinberg D: **Perception of sniff phase in mouse olfaction.** *Nature* 2011, **479**:397-400.

20. Znamenskiy P, Zador AM: **Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination.** *Nature* 2013, **497**:482-485.

21. Brunton BW, Botvinick MM, Brody CD: **Rats and humans can optimally accumulate evidence for decision-making.** *Science* 2013, **340**:95-98.

22. Guo ZV, Li N, Huber D, Ophir E, Gutnisky D, Ting JT, Feng G, Svoboda K: **Flow of cortical activity underlying a tactile decision in mice.** *Neuron* 2014, **81**:179-194.

23. Adibi M, Diamond ME, Arabzadeh E: **Behavioral study of whisker-mediated vibration sensation in rats.** *Proc Natl Acad Sci U S A* 2012, **109**:971-976.

24. Musall S, von der Behrens W, Mayrhofer JM, Weber B, Helmchen F, Haiss F: **Tactile frequency discrimination is enhanced by circumventing neocortical adaptation.** *Nat Neurosci* 2014, **17**:1567-1573.

25. Miller P, Katz DB: **Stochastic transitions between neural states in taste processing and decision-making.** *J Neurosci* 2010, **30**:2559-2570.

26. Drugowitsch J, DeAngelis GC, Angelaki DE, Pouget A: **Tuning the speed-accuracy trade-off to maximize reward rate in multisensory decision-making.** *Elife* 2015:4.

27. Raposo D, Sheppard JP, Schrater PR, Churchland AK: **Multisensory decision-making in rats and humans.** *J Neurosci* 2012, **32**:3726-3735.

28. Keller PJ, Ahrens MB: **Visualizing whole-brain activity and development at the single-cell level using light-sheet microscopy.** *Neuron* 2015, **85**:462-483.

29. Peron SP, Freeman J, Iyer V, Guo C, Svoboda K: **A cellular resolution map of barrel cortex activity during tactile behavior.** *Neuron* 2015, **86**:783-799.

Two-photon volumetric imaging was used to allow the authors to image ~12 000 neurons/mouse during a whisker object localization task. They observed sparse and spatially intermingled representations of multiple features of the tactile stimulus.

30. Ruff DA, Cohen MR: **Attention can either increase or decrease spike count correlations in visual cortex.** *Nat Neurosci* 2014, **17**:1591-1597.

31. Taniguchi H, He M, Wu P, Kim S, Paik R, Sugino K, Kvitsiani D, Fu Y, Lu J, Lin Y *et al.*: **A resource of Cre driver lines for genetic targeting of GABAergic neurons in cerebral cortex.** *Neuron* 2011, **71**:995-1013.

32. Scanziani M, Haussler M: **Electrophysiology in the age of light.** *Nature* 2009, **461**:930-939.

33. Yang H, Kwon SE, Severson KS, O'Connor DH: **Origins of choice-related activity in mouse somatosensory cortex.** *Nat Neurosci* 2016, **19**:127-134.

The authors identified cells that were members of specific anatomical pathways to pinpoint which circuits support choices about somatosensory stimuli. Axons in primary somatosensory cortex for neurons with cell bodies in secondary somatosensory cortex carried the strongest choice signals, pointing to this pathway as critical for the decision circuit.

34. Chowdhury SA, DeAngelis GC: **Fine discrimination training alters the causal contribution of macaque area MT to depth perception.** *Neuron* 2008, **60**:367-377.

35. Karlsson MP, Tervo DG, Karpova AY: **Network resets in medial prefrontal cortex mark the onset of behavioral uncertainty.** *Science* 2012, **338**:135-139.

36. Lau B, Glimcher PW: **Dynamic response-by-response models of matching behavior in rhesus monkeys.** *J Exp Anal Behav* 2005, **84**:555-579.

37. Donahue CH, Seo H, Lee D: **Cortical signals for rewarded actions and strategic exploration.** *Neuron* 2013, **80**:223-234.

38. Glaze CM, Kable JW, Gold JI: **Normative evidence accumulation in unpredictable environments.** *Elife* 2015:4.

39. Seo H, Cai X, Donahue CH, Lee D: **Neural correlates of strategic reasoning during competitive games.** *Science* 2014, **346**:340-343.

40. Purcell BA, Kiani R: **Neural mechanisms of post-error adjustments of decision policy in parietal cortex.** *Neuron* 2016, **89**:658-671.

The authors identified two factors that explain why subjects typically have slower reaction times following error trials. First, subjects have reduced sensitivity to stimuli on trials that follow errors. Second, subjects have a higher decision bound. These observations suggest decision parameters are updated dynamically depending on recent trial history.

41. Gold JI, Law CT, Connolly P, Bennur S: **The relative influences of priors and sensory evidence on an oculomotor decision variable during perceptual learning.** *J Neurophysiol* 2008, **100**:2653-2668.

42. Busse L, Ayaz A, Dhruv NT, Katzner S, Saleem AB, Scholvinck ML, Zaharia AD, Carandini M: **The detection of visual contrast in the behaving mouse.** *J Neurosci* 2011, **31**:11351-11361.

43. Scott BB, Constantinople CM, Erlich JC, Tank DW, Brody CD: **Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats.** *Elife* 2015, **4**:e11308.

44. Horwitz GD, Chichilnisky EJ, Albright TD: **Blue-yellow signals are enhanced by spatiotemporal luminance contrast in macaque V1.** *J Neurophysiol* 2005, **93**:2263-2278.

45. DeAngelis GC, Ohzawa I, Freeman RD: **Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. I. General characteristics and postnatal development.** *J Neurophysiol* 1993, **69**:1091-1117.

46. Raposo D, Kaufman MT, Churchland AK: **A category-free neural population supports evolving demands during decision-making.** *Nat Neurosci* 2014, **17**:1784-1792.

Recordings from rat posterior parietal cortex were used to demonstrate that neurons in decision making areas can reflect idiosyncratic combinations of multiple task parameters.

47. Sheppard JP, Raposo D, Churchland AK: **Dynamic weighting of multisensory stimuli shapes decision-making in rats and humans.** *J Vis* 2013:13.

The authors showed that rats and humans similarly integrate auditory and visual signals that bear on the same decision. Because auditory and visual stimulus strength fluctuated over time, the authors were able to demonstrate that both species used most of the evidence that arrived over the course of the 1000 ms decision.

48. Bondy AG, Cumming BG: **Monkeys behaving badly: probing macaque subjects' internal task strategies with psychophysical reverse correlation.** *Society for Neuroscience Abstracts* 2016.

49. Sha L, Kiani R: **Mechanisms of perceptual learning and learning interference in visual categorization.** *Society for Neuroscience Abstracts* 2015. 60.07.
50. Zohary E, Shadlen MN, Newsome WT: **Correlated neuronal discharge rate and its implications for psychophysical performance.** *Nature* 1994, **370**:140-143.
51. Berg Rvd, Anandalingam K, Zylberberg A, Kiani R, Shadlen MN, Wolpert DM: **A common mechanism underlies changes of mind about decisions and confidence.** *eLife* 2016. in press.
52. Kaufman MT, Churchland MM, Ryu SI, Shenoy KV: **Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex.** *Elife* 2015, **4**:e04677.
53. Bollimunta A, Totten D, Ditterich J: **Neural dynamics of choice: single-trial analysis of decision-related activity in parietal cortex.** *J Neurosci* 2012, **32**:12684-12701.
54. Resulaj A, Kiani R, Wolpert DM, Shadlen MN: **Changes of mind in decision-making.** *Nature* 2009, **461**:263-266.
55. Kiani R, Cueva CJ, Reppas JB, Newsome WT: **Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials.** *Curr Biol* 2014, **24**:1542-1547.
- The authors made simultaneous recordings from hundreds of neurons in prearcuate gyrus of monkeys during a motion discrimination task. A signature of the animal's changing its mind on some trials was identified in this population response. The properties of this signature conformed to expectations based on prior theoretical and behavioral studies.
56. Ponce-Alvarez A, Nacher V, Luna R, Riehle A, Romo R: **Dynamics of cortical neuronal ensembles transit from decision making to storage for later report.** *J Neurosci* 2012, **32**:11956-11969.
57. Durstewitz D, Vittoz NM, Floresco SB, Seamans JK: **Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning.** *Neuron* 2010, **66**:438-448.
58. Churchland AK, Kiani R, Chaudhuri R, Wang XJ, Pouget A, Shadlen MN: **Variance as a signature of neural computations during decision making.** *Neuron* 2011, **69**:818-831.
59. Goris RL, Movshon JA, Simoncelli EP: **Partitioning neuronal variability.** *Nat Neurosci* 2014, **17**:858-865.
60. Park IM, Meister ML, Huk AC, Pillow JW: **Encoding and decoding in parietal cortex during sensorimotor decision-making.** *Nat Neurosci* 2014, **17**:1395-1403.
61. Cunningham JP, Yu BM: **Dimensionality reduction for large-scale neural recordings.** *Nat Neurosci* 2014, **17**:1500-1509.
62. Brendel W, Romo R, Machens CK: *Demixed Principal Component Analysis.* NIPS; 2011.
- The authors developed a method for targeted dimensionality reduction that reveals the dependency of high dimensional data on individual task parameters. This allows visualization of data in a space that is more intuitive compared to more traditional methods of dimensionality reduction, such as PCA.
63. Machens CK, Romo R, Brody CD: **Functional, but not anatomical, separation of "what" and "when" in prefrontal cortex.** *J Neurosci* 2010, **30**:350-360.
64. Mante V, Sussillo D, Shenoy KV, Newsome WT: **Context-dependent computation by recurrent dynamics in prefrontal cortex.** *Nature* 2013, **503**:78-84.
65. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV: **Neural population dynamics during reaching.** *Nature* 2012, **487**:51-56.
66. Kobak D, Brendel W, Constantinidis C, Feierstein CE, Kepecs A, Mainen ZF, Romo R, Qi XL, Uchida N, Machens CK: **Demixed principal component analysis of neural population data.** *Elife* 2016:5.
67. Shadlen MN, Kiani R, Newsome WT, Gold JI, Wolpert DM, Zylberberg A, Ditterich J, de Lafuente V, Yang T, Roitman J: **Comment on "Single-trial spike trains in parietal cortex reveal discrete steps during decision-making".** *Science* 2016, **351**:1406.
68. Latimer KW, Yates JL, Meister ML, Huk AC, Pillow JW: **Response to Comment on "Single-trial spike trains in parietal cortex reveal discrete steps during decision-making".** *Science* 2016, **351**:1406.
69. Fetsch CR, Kiani R, Shadlen MN: **Predicting the accuracy of a decision: a neural mechanism of confidence.** *Cold Spring Harb Symp Quant Biol* 2015.
70. Lunn DJ, Best N, Thomas A, Wakefield J, Spiegelhalter D: **Bayesian analysis of population PK/PD models: general concepts and software.** *J Pharmacokinetic Pharmacodyn* 2002, **29**:271-307.
71. Kanitscheider I, Brown A, Pouget A, Churchland AK: **Multisensory decisions provide support for probabilistic number representations.** *J Neurophysiol* 2015. jn 00787 02014.
72. Lindley DV: **A statistical paradox.** *Biometrika* 1957, **44**:187-192.
73. Robert CP: **On the Jeffreys-Lindley paradox.** *Philos Sci* 2014, **81**:216-232.
74. Mazurek ME, Roitman JD, Ditterich J, Shadlen MN: **A role for neural integrators in perceptual decision making.** *Cereb Cortex* 2003, **13**:1257-1269.
75. Niwa M, Ditterich J: **Perceptual decisions between multiple directions of visual motion.** *J Neurosci* 2008, **28**:4435-4445.