

UC Berkeley

UC Berkeley Previously Published Works

Title

Decoding Cognitive Processes from Neural Ensembles

Permalink

<https://escholarship.org/uc/item/6jz9713b>

Journal

Trends in Cognitive Sciences, 22(12)

ISSN

1364-6613

Author

Wallis, Joni D

Publication Date

2018-12-01

DOI

10.1016/j.tics.2018.09.002

Peer reviewed



HHS Public Access

Author manuscript

Trends Cogn Sci. Author manuscript; available in PMC 2019 December 01.

Published in final edited form as:

Trends Cogn Sci. 2018 December ; 22(12): 1091–1102. doi:10.1016/j.tics.2018.09.002.

Decoding cognitive processes from neural ensembles

Joni D. Wallis

Department of Psychology and Helen Wills Neuroscience Institute, U.C. Berkeley

Abstract

An intrinsic difficulty in studying cognitive processes is that they are unobservable states that exist in between observable responses to the sensory environment. Cognitive states must be inferred from indirect behavioral measures. Neuroscience potentially provides the tools necessary to measure cognitive processes directly, but it is challenged on two fronts. First, neuroscientific measures often lack the spatiotemporal resolution to identify the neural computations that underlie a cognitive process. Second, the activity of a single neuron, which is the fundamental building block of neural computation, is too noisy to provide accurate measurements of a cognitive process. In this paper, we will examine recent developments in neurophysiological recording and analysis methods that provide a potential solution to these problems.

Keywords

Neural decoding; neurophysiology; high-level cognition; orbitofrontal cortex; decision-making

Dealing with neuronal noise

One of the central tenets of neuroscience is that neurons represent information by changing the rate at which they fire action potentials ('spikes'). However, this rate code is noisy. The exact timing and number of spikes shows considerable variability from one neuronal response to another. Neurophysiologists, who are attempting to measure how the neuron codes information, must extract the meaningful signal from this noisy response. This can be achieved by recording the activity of the neuron over repeated trials of the same experimental condition and averaging the neuron's firing rate across these trials. A disadvantage of this approach is that the neural computations of interest may show considerable variability across trials. This is particularly problematic when studying cognition since the experimenter has little control over the cognitive process. An alternative approach is to record the activity of many single neurons simultaneously and then project the pattern of neural activity into a high dimensional space that can be used to classify the information represented by the neurons. This avoids the necessity of recording across multiple trials and enables the experimenter to 'read out' a cognitive process directly from

*Correspondence: wallis@berkeley.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

brain activity. This review will focus on the application of this methods to understanding decision-making, but the methods can equally be applied to studying any cognitive process.

From single units to neural ensembles

Neurophysiology has traditionally focused on understanding the information encoded by a single neuron. For example, a stimulus might be varied along some dimension and changes in the size of the neuronal response are measured. To deal with trial-to-trial variability, the neuron's response is averaged across multiple presentations of the stimulus. However, applying this same approach to understanding cognitive processes risks averaging away the very process one is interested in measuring. Consider a person choosing between two restaurants. Perhaps they may alternate between the two restaurants changing their mind several times, or maybe they slowly build the case for one restaurant over the other. The eventual choice might be the same and it may take the same amount of time, but the underlying dynamics of the decision were very different. A further problem is introduced when the investigator attempts to control the experiment. The repetition and tight control on learning and behaviors that is required to average across trials runs counter to the freedom necessary for the animal to cogitate. Both experimental and theoretical work have made a distinction between the kinds of choices we make habitually (e.g. selecting the same brand of soda in the grocery store) and those which are unique to a specific set of circumstances (e.g. deciding where to eat when on vacation) [1, 2].

Over the last couple of decades, there has been a dramatic increase in the number of neurons that are recorded in a single session. Whereas neurophysiologists used to use a single wire to record from one or two neurons simultaneously, many labs now use multicontact electrodes and routinely record from over a hundred neurons simultaneously. This has not just allowed data to be collected more quickly, but it has led to entirely new kinds of analysis methods, most importantly for the purposes of this review, the ability to extract or 'decode' information from the neural ensemble with single trial resolution. This approach has been most successfully employed by researchers studying motor control [3]. These methods have been used to construct brain machine interfaces for paralyzed patients [4, 5]. Recording arrays are implanted in motor and premotor cortical areas, and the pattern of neural activity can be decoded to determine the movement intended by the paralyzed patient. These decoded neural signals can then be used to control robot arms or computer cursors, thereby helping to restore activities of daily living. Recent work has begun to apply these same methods to decoding cognitive processes. Although the focus of this review is on neurophysiological recording, we note that considerable progress has also been made applying similar methods to other neural signals (Text Box 1).

A challenge in applying decoding to the problem of decision-making is in establishing an appropriate ground truth. To build a decoder, it is necessary to have a training set of data, in order to measure the neuronal response to a known parameter. For example, to decode stimulus A from stimulus B, we first need to present both stimuli many times to produce a distribution of neural responses that might be expected from one or other stimulus presentation. Establishing such a ground truth is relatively straightforward when we are dealing with concrete stimuli or responses, but we seem to have come full circle to the

problem we described in the beginning, since we are now faced with the difficulty of establishing a ground truth for an unobservable cognitive process. There have been at least two approaches to the problem, both of which have used well characterized neural responses to help establish a ground truth. The first approach has capitalized on spatial encoding in hippocampus, while the second approach has capitalized on value encoding in orbitofrontal cortex (OFC). We will consider each of these approaches in turn and how they have helped to uncover the cognitive processes underlying decision-making.

Hippocampal spatial trajectories and decision-making

Hippocampal neurons have place fields [6]. Each neuron has a preferred region of space where the neuron fires when the animal is in that location (Figure 1A). When the animal pauses, many of these neurons will fire synchronously, an event that is visible in the local field potential (LFP) as a sharp-wave-ripple (Figure 1B). Zooming in to a more fine grain time scale reveals that the neurons are not firing synchronously, but rather in the order in which their place fields were originally encountered [7]. It is as if the animal is replaying the experience while it rests (Figure 1C). Decoding algorithms can extract this information and construct the complete spatial trajectory that the animal is replaying (Figure 1D). This phenomenon can also be observed in the open-field. When the rat pauses as it explores an open-field, sequences of place cells will fire along one potential navigational trajectory and then another, after which the animal will move along one of the trajectories [8]. Furthermore, decoded trajectories are influenced by the location of rewards, with trajectories biased towards those heading towards a rewarded goal (Figure 2).

Information regarding upcoming trajectories can also be decoded from other aspects of the hippocampal signal. Although most spikes from a hippocampal neuron occurred when the animal was in the neuron's place field, there were additional scattered spikes that could occur outside of the neuron's place field [9]. These extrafield spikes were not distributed randomly, but rather tended to occur at choice points in the behavioral task. For example, at the choice point in a T-maze, hippocampal neurons would often encode future locations that the animal expected to traverse [9, 10]. This future encoding would sweep down one arm of the maze and then the other, as if the animal were simulating the two possible courses of action in order to make its choice (Figure 3). It was as if the neuron's place field had temporarily shifted to reflect the potential locations that the animal might traverse. This spiking was distinct from the activity that occurred during the sharp-wave ripples, as described above. The animal's behavior was also consistent with the possibility that the animal was simulating possible courses of action. When the rat reached the choice point it would often look backwards and forwards down either arm of the maze, a behavior referred to as vicarious trial-and-error learning [10]. Notably both the vicarious trial and error behavior and the hippocampal sweeps disappeared as the animals practiced the task and the behavior became more automatic. This emphasizes the importance of using less practiced and less constrained behaviors to ensure that an animal is engaging decision-making [9].

Taken together, these studies suggest that hippocampal neural activity may provide an insight into the decision-making process. The experimenters used decoders trained on established ground truths (hippocampal place fields during actual navigation) to measure

hidden cognitive states (navigational decisions). However, they lack a critical piece of the decision: what makes one option preferable to another?

Decoding value from orbitofrontal cortex

Flip-flopping value signals during decision-making

OFC has long been ascribed a central role in value-based decision-making [11]. This process involves assigning a value to potential outcomes and then selecting that option that leads to the most valuable outcome. It is distinct from perceptual decision-making, which involves making decisions about sensory stimuli, in that it requires value judgements. Unlike sensory stimuli, value judgments are frequently hidden states that must be inferred. For example, there is nothing intrinsic to a dark brown cuboid that predicts the taste of chocolate. Instead, we infer that taste from our prior experience of chocolate bars. Furthermore, different people assign different tastes different value: some people prefer milk chocolate, while others prefer dark. Patients with OFC damage have a specific difficulty in making value-based decisions [12, 13] and OFC neurons encode the predicted value associated with specific choice options [14–16], consistent with a role of OFC in value-based decision-making.

This is not to say that OFC is involved in all types of value decisions. As mentioned earlier, there is an important distinction between habitual and novel choices [1, 2], with OFC particularly implicated in the latter [17]. For example, both fMRI [18] and magnetoencephalography [19] experiments in humans have shown that OFC is only involved for relatively novel choices. OFC has also been described as implementing ‘goods-based’ decision-making [20]. One way to simplify the process of making a decision is to separate the valuation of things in the environment (‘goods’) from the actions necessary to obtain these goods. This avoids the combinatorial explosion that would occur if every combination of potential good and action were assigned a value. Consistent with a role in goods-based decision-making, OFC neurons are less likely to encode actions than other regions of the frontal lobe [21–23]. In addition, OFC damage appears to preferentially affect stimulus valuation compared to action valuation in both humans [24] and monkeys [25], although recent findings suggest that this distinction may be less clear cut than previously thought [26, 27].

Because value is a hidden state and inherently subjective, studying the mechanisms underpinning value-based decision-making is particularly difficult. Decoding provides a potential solution. In a recent study, we used decoding methods to probe how monkeys choose between different choice options [28]. Two monkeys learned a set of cues, each of which predicted a specific type and size of reward. We recorded from small ensembles of OFC neurons (usually between 10 and 20 neurons) from cytoarchitectonic areas 11 and 13 [29] and trained a classifier to identify the value of the reward predicted by the cue from the pattern of neural activity. We then used the classifier to decode neural activity when the animal was presented with a choice between a pair of the cues (Figure 4A). Although the output of the decoding was noisy, there were clear periods where the classifier was very confident that OFC neural activity was consistent with one of the picture values (posterior probability close to 1). Furthermore, there were periods where the decoder would flip-flop

between representing the value associated with one option to representing the value of the alternate option.

To show that this flip-flopping did not simply reflect noise in the output of the decoder, we examined whether the decoder output was predictive of choice behavior. There were more flip-flops on error trials, where the animal chose the less valuable option, compared to correct trials (Figure 4B) and stronger representation of the chosen option, relative to the unchosen option, produced faster choice responses (Figure 4C). Thus, flip-flopping was predictive of choice behavior. Like the hippocampal work, we were successful in using an established ground truth (OFC encoding of the value of a reward-predictive cue) to decode a cognitive hidden state, in this case, the dynamics underlying individual value-based decisions. Decision-making seems to consist of OFC alternately representing the value of the outcomes that are predicted by the two reward-predictive cues constituting the choice.

Neural mechanisms of flip-flopping

What might be the mechanism that underpins these dynamics? Recent work has begun to probe the properties of neurons in prefrontal cortex (of which OFC is a part) that make them particularly useful for cognitive processing. Cortical neurons possess temporal fields, which can be determined by calculating the autocorrelation of spiking during baseline periods, which measures the tendency of neurons to fire in repetitive patterns. As one moves from sensorimotor cortex to association cortex, the autocorrelation time constant increases, reaching its maximum in prefrontal cortex [30]. In other words, there is an increase in the timescale over which repeated patterns of firing evolve. This suggests that the firing properties of prefrontal neurons predispose them to integrating information over longer time intervals than sensorimotor cortical regions. Within OFC, neurons that maintain the predicted value of the reward-predictive cue until the time of the reward outcome show longer time constants than neurons that are not involved in this process [31]. Prefrontal cortex is also strongly implicated in the process of working memory [32–34], whereby information must be explicitly maintained by prefrontal neural ensembles across delays. Recent computational models of how this process might be accomplished suggest that the balance of excitation and inhibition within prefrontal neuronal ensembles can predispose them to certain attractor states (Text Box 2), which can be used to store information across temporal delays [35, 36].

Integrating these ideas, the flip-flopping during decision-making could consist of competition between OFC attractor states responsible for representing individual reward-predictive cues. This competition could be biased, similar to models that have been proposed for the competition for attention in posterior sensory regions [37]. This could be accomplished via synaptic strengthening from dopaminergic prediction error signals during the initial learning of reward predictions [38, 39] that could produce stronger ensembles and deeper attractor basins for more valuable cues. One prediction of this model is that the more valuable option should be represented either more frequently or for longer duration. Indeed, both predictions were true. Chosen states were significantly longer than unchosen states and more frequent (Figure 4D).

One challenge to this proposed mechanism is reconciling attractor dynamics with flip-flopping. If an attractor basin is in operation, once the neural network enters a stable state, it should not leave that state to enter the alternate state. A potential solution is the theta oscillation, which has been shown to affect flip-flopping in the hippocampus. Rats exposed to two different spatial contexts that are discriminated via light cues develop orthogonal neural representations in the hippocampus that encode each context. When the light cues are abruptly changed, indicating a shift of context, the pattern of hippocampal firing flip-flops between the two contextual representations, with the timing of the flip-flops aligned to theta [40]. Computational modeling has shown that the theta oscillation induces flip-flopping by altering the balance between external inputs and the internal recurrent neuronal connections that underpin the attractor state, which allows for a potential resetting of attractor states with each theta cycle [41]. Furthermore, synaptic plasticity can influence this process, biasing attractor states towards those with stronger underlying synaptic connections [42].

The theta signal in OFC is particularly strong during value-based decisions [43, 44]. More broadly, the LFP contains a good deal of information relating to the decision. Indeed, we have recently shown that there is little difference between the information encoded by OFC neurons and the information that is encoded by the OFC LFP [45]. Consequently, we were able to decode OFC flip-flopping solely from OFC LFP, albeit with somewhat reduced accuracy compared to decoding performance based on OFC spikes [28]. We have also shown that individual prefrontal neurons phase lock to different LFP frequencies [46]. This could enable frequency encoding whereby different brain areas are able to tune in to different OFC frequencies to obtain different kinds of information. Evidence for this comes from studies that have recorded simultaneously from OFC and other brain areas. For example, there is increased coherence between OFC and the hippocampus in the theta band during the learning of novel goal-directed behaviors [47].

Interaction between valuation and other cognitive processes

What is noticeably absent from the above discussion is how the choice is implemented. At some point, the value flip-flopping in OFC must be converted into a motor response. The evidence suggests that this implementation occurs in areas downstream of OFC. We found that there was not a direct relationship between the value flip-flopping and the choice response: there was no evidence that OFC needed to be in a particular value state in order for the choice response to occur [28]. In addition, many studies have shown that OFC encodes very little information about the choice motor response [21–23]. One possible candidate is lateral prefrontal cortex (LPFC), which has strong, reciprocal connections with OFC [48] and motor areas [49, 50]. During a decision-making task, LPFC neurons show a gradual evolution from encoding the value of the chosen option to encoding the motor response necessary to acquire that option [51]. Furthermore, the dynamics of the unfolding of value signals in OFC are correlated with the value-to-action transformation that occurs in LPFC [52], suggesting a close relationship between value coding in OFC and action selection in LPFC. It is likely that there are also other routes to action, perhaps through the cingulate cortex or striatum [53], particularly as lesions of LPFC do not necessarily disrupt value-based choices [54].

Decoding methods can also be applied to other cognitive processes that might interact with decision-making. Neural responses in LPFC are strongly driven by the locus of covert attention [55–58]. Single trial decoding of the attentional locus can be obtained by recording from ensembles of 50 or so LPFC neurons [59]. There has been substantial debate regarding the role that attention may play in decision-making [60, 61]. Computational modeling of choice behavior has suggested that value may act as a multiplier on the rate at which evidence is accumulated to support one or other option as those options are sequentially fixated during the decision [62]. Direct support for this model has come from single neuron studies, which show that OFC value coding depends on whether a reward-predictive cue is fixated [63]. However, a caveat with this study is that there was no behavioral task: the animal was simply free viewing the screen on which the reward-predictive cue was presented. The lack of value coding when the animal was not fixating the cue, may simply have reflected the fact that the animal was not engaged. Indeed, in our own studies, eye movements had no relationship to the flip-flopping of value representations during decision-making [28].

One possible explanation for the lack of relationship between eye movements and value representations is that the interaction and between value and attention may occur covertly and have only an indirect relationship with eye movements. Indeed, one of the posited reasons that we have working memory is so that we are not ‘stimulus-bound’: we can think about things we are not necessarily looking at. For example, if you are staring at the steak in the meat section of the grocery store, your OFC may be encoding your valuation of steak, but it could equally be encoding your valuation of chicken so that you can compare it to the steak. This brings us back to the problem outlined at the beginning of this paper: how do we access the contents of working memory, which is an unobservable, covert process? And the solution may be the same: decode the contents of working memory and compare it to the decoding of decision-making. One possibility is that it is shifts in covert attention that are controlling OFC flip-flopping. Alternatively, it may be that OFC flip-flopping guides shifts of covert attention. Future work could attempt to distinguish between these two possibilities by simultaneously decoding value information from OFC and spatial attention from LPFC.

Concluding remarks and future perspectives

Decoding enables neuroscientists to measure cognitive processes as they unfold. However, there are some caveats to keep in mind in interpreting the results of a decoding analysis. Most importantly, just because we can decode information from the activity of a neural ensemble, this does not mean that individual neurons are encoding that information. For example, it would be easy to build a decoder that could output the orientation of edges from retinal activity, but we also know that individual retinal neurons encode spots of lines rather than edges. In other words, we cannot infer the function of a neural ensemble from our ability to decode information from that ensemble. This problem becomes more pronounced as the sophistication of the decoder increases. For example, returning to our retina example, in principal we could build a decoder that would detect faces from retinal activity: after all, that is what the brain is doing. We also note that although we have presented ensemble analyses in opposition to single neuron analyses, the real power may come from combining the two approaches. For example, when averaging the activity of single neurons across trials,

investigators typically synchronize the neural activity to external events, such as the presentation of sensory stimuli or the animal's overt behavior. Combining single neuron and ensemble analyses raises the possibility of building encoding models of neural tuning that are synchronized to internal cognitive processes.

Finally, we note that decoding rests on the ability to record large numbers of neurons simultaneously, which enables the researcher to overcome the stochasticity that is inherent in the firing of single neurons. The last decade has seen a dramatic increase in the number of neurons that can be simultaneously recorded, with concomitant increases in the performance of decoders. Early efforts used arrays of microwires [64] or silicon probes [65] that were chronically implanted in the cortex. However, these arrays can typically only penetrate the brain for a few millimeters and so they cannot record from deep areas of the brain like OFC. More recently linear microarrays have been developed that are the same diameter as traditional recording electrodes but with many more contacts spaced along the shank. For example, the Neuropixels probe has 960 contacts and has been used to record from 700 neurons simultaneously [45]. Future experiments can capitalize on these technological developments to perform high resolution decoding of different cognitive processes from different brain areas, thereby providing a more complete picture of the computations underlying high-level cognition.

Acknowledgments

The author was supported by grants from the NIH (NIMH R01-MH117763, NIDA R21-DA041791, and NIMH R01-MH097990) during the writing of this review.

Glossary

Attractor state/basin

In a dynamical system, the attractor state or attractor basin, is the subset of the state space towards which objects will tend to move, irrespective of the starting conditions of the object.

Autocorrelation

A mathematical tool for finding repeating patterns. It is the correlation between serial observations as a function of the time lag between them.

Local field potential

The electrical potential recorded from an electrode positioned in neural tissue that reflects the summed electrical activity within $\approx 0.5 - 1$ mm of the electrode. It includes both action potentials and subthreshold membrane potentials.

Place fields

The spatial location that will cause a place neuron to fire whenever the animal is at that location

Sharp wave ripple

Large amplitude, short duration deflections in the local field potential, which are only found in the hippocampus and its neighboring areas.

Theta oscillation

A rhythmic oscillation between 4 and 10 Hz that is present in the local field potential throughout the brain. It is particularly prominent in the hippocampus.

References

1. Daw ND and Dayan P (2014) The algorithmic anatomy of model-based evaluation. *Philos Trans R Soc Lond B Biol Sci* 369 (1655).
2. Daw ND et al. (2005) Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci* 8 (12), 1704–11. [PubMed: 16286932]
3. Churchland MM et al. (2007) Techniques for extracting single-trial activity patterns from large-scale neural recordings. *Curr Opin Neurobiol* 17 (5), 609–18. [PubMed: 18093826]
4. Hochberg LR et al. (2012) Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature* 485 (7398), 372–5. [PubMed: 22596161]
5. Jarosiewicz B et al. (2015) Virtual typing by people with tetraplegia using a self-calibrating intracortical brain-computer interface. *Sci Transl Med* 7 (313), 313ra179.
6. O'Keefe J.a.N., L. (1978) *The hippocampus as a cognitive map*, Oxford University Press.
7. Foster DJ and Wilson MA (2006) Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* 440 (7084), 680–3. [PubMed: 16474382]
8. Pfeiffer BE and Foster DJ (2013) Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* 497 (7447), 74–9. [PubMed: 23594744]
9. Johnson A and Redish AD (2007) Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci* 27 (45), 12176–89. [PubMed: 17989284]
10. Redish AD (2016) Vicarious trial and error. *Nat Rev Neurosci* 17 (3), 147–59. [PubMed: 26891625]
11. Damasio AR (1996) The somatic marker hypothesis and the possible functions of the prefrontal cortex. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* 351 (1346), 1413–20.
12. Camille Net al.(2011) Ventromedial frontal lobe damage disrupts value maximization in humans. *J Neurosci* 31 (20), 7527–32. [PubMed: 21593337]
13. Vaidya AR and Fellows LK (2015) Testing necessary regional frontal contributions to value assessment and fixation-based updating. *Nat Commun* 6, 10120. [PubMed: 26658289]
14. Kennerley SW et al. (2011) Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nature neuroscience* 14 (12), 1581–9. [PubMed: 22037498]
15. Emrich SM et al. (2013) Distributed patterns of activity in sensory cortex reflect the precision of multiple items maintained in visual short-term memory. *J Neurosci* 33 (15), 6516–23. [PubMed: 23575849]
16. Padoa-Schioppa C and Conen KE (2017) Orbitofrontal Cortex: A Neural Circuit for Economic Decisions. *Neuron* 96 (4), 736–754. [PubMed: 29144973]
17. Glascher J et al. (2010) States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron* 66 (4), 585–95. [PubMed: 20510862]
18. Barron H . et al. (2013) Online evaluation of novel choices by simultaneous representation of multiple memories. *Nat Neurosci* 16 (10), 1492–8. [PubMed: 24013592]
19. Hunt LT et al. (2012) Mechanisms underlying cortical activity during value-guided choice. *Nature neuroscience* 15 (3), 470–6. [PubMed: 22231429]
20. Padoa-Schioppa C (2011) Neurobiology of economic choice: a good-based model. *Annual review of neuroscience* 34, 333–59.
21. Padoa-Schioppa C and Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. *Nature* 441 (7090), 223–6. [PubMed: 16633341]
22. Kennerley SW et al. (2009) Neurons in the frontal lobe encode the value of multiple decision variables. *J Cogn Neurosci* 21 (6), 1162–78. [PubMed: 18752411]

23. Grattan LE and Glimcher PW (2014) Absence of spatial tuning in the orbitofrontal cortex. *PLoS One* 9 (11), e112750. [PubMed: 25386837]
24. Camille N et al. (2011) Double dissociation of stimulus-value and action-value learning in humans with orbitofrontal or anterior cingulate cortex damage. *J Neurosci* 31 (42), 15048–52. [PubMed: 22016538]
25. Rudebeck PH et al. (2008) Frontal cortex subregions play distinct roles in choices between actions and stimuli. *J Neurosci* 28 (51), 13775–85. [PubMed: 19091968]
26. Gremel CM and Costa RM (2013) Orbitofrontal and striatal circuits dynamically encode the shift between goal-directed and habitual actions. *Nat Commun* 4, 2264. [PubMed: 23921250]
27. Rhodes SE and Murray EA (2013) Differential effects of amygdala, orbital prefrontal cortex, and prelimbic cortex lesions on goal-directed behavior in rhesus macaques. *J Neurosci* 33 (8), 3380–9. [PubMed: 23426666]
28. Rich EL and Wallis JD (2016) Decoding subjective decisions from orbitofrontal cortex. *Nat Neurosci* 19 (7), 973–80. [PubMed: 27273768]
29. Carmichael ST and Price JL (1994) Architectonic subdivision of the orbital and medial prefrontal cortex in the macaque monkey. *J Comp Neurol* 346 (3), 366–402. [PubMed: 7527805]
30. Murray JD et al. (2014) A hierarchy of intrinsic timescales across primate cortex. *Nat Neurosci* 17 (12), 1661–3. [PubMed: 25383900]
31. Cavanagh SE et al. (2016) Auto correlation structure at rest predicts value correlates of single neurons during reward-guided choice. *Elife* 5.
32. Goldman-Rakic PS (1987) Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In *Handbook of physiology, the nervous system, higher functions of the brain* (Plum F ed), pp. 373 – 417, American Physiological Society.
33. Lara AH and Wallis JD (2015) The Role of Prefrontal Cortex in Working Memory: A Mini Review. *Front Syst Neurosci* 9, 173. [PubMed: 26733825]
34. Chiang FK and Wallis JD (2018) Spatiotemporal encoding of search strategies by prefrontal neurons. *Proc Natl Acad Sci U S A* 115 (19), 5010–5015. [PubMed: 29686072]
35. Wimmer K et al. (2014) Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nat Neurosci* 17 (3), 431–9. [PubMed: 24487232]
36. Murray J. et al. (2017) Working Memory and Decision-Making in a Frontoparietal Circuit Model. *J Neurosci* 37 (50), 12167–12186. [PubMed: 29114071]
37. Desimone R and Duncan J (1995) Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18, 193–222. [PubMed: 7605061]
38. Takahashi YK et al. (2011) Expectancy-related changes in firing of dopamine neurons depend on orbitofrontal cortex. *Nature neuroscience* 14 (12), 1590–7. [PubMed: 22037501]
39. Schultz W et al. (1997) A neural substrate of prediction and reward. *Science* 275 (5306), 1593–9. [PubMed: 9054347]
40. Jezek K et al. (2011) Theta-paced flickering between place-cell maps in the hippocampus. *Nature* 478 (7368), 246–9. [PubMed: 21964339]
41. Stella F and Treves A (2011) Associative memory storage and retrieval: involvement of theta oscillations in hippocampal information processing. *Neural Plast* 2011, 683961. [PubMed: 21961072]
42. Mark S et al. (2017) Theta-paced flickering between place-cell maps in the hippocampus: A model based on short-term synaptic plasticity. *Hippocampus* 27 (9), 959–970. [PubMed: 28558154]
43. van Wingerden M et al. (2014) Phase-amplitude coupling in rat orbitofrontal cortex discriminates between correct and incorrect decisions during associative learning. *J Neurosci* 34 (2), 493–505. [PubMed: 24403149]
44. Pennartz CM et al. (2011) Population coding and neural rhythmicity in the orbitofrontal cortex. *Ann N Y Acad Sci* 1239, 149–61. [PubMed: 22145884]
45. Jun JJ et al. (2017) Fully integrated silicon probes for high-density recording of neural activity. *Nature* 551 (7679), 232–236. [PubMed: 29120427]

46. Canolty RT et al. (2010) Oscillatory phase coupling coordinates anatomically dispersed functional cell assemblies. *Proceedings of the National Academy of Sciences of the United States of America* 107 (40), 17356–61.
47. Young JJ and Shapiro ML (2011) Dynamic coding of goal-directed paths by orbital prefrontal cortex. *J Neurosci* 31 (16), 5989–6000. [PubMed: 21508224]
48. Barbas H and Pandya DN (1989) Architecture and intrinsic connections of the prefrontal cortex in the rhesus monkey. *J Comp Neurol* 286 (3), 353–75. [PubMed: 2768563]
49. Lu MT et al. (1994) Interconnections between the prefrontal cortex and the premotor areas in the frontal lobe. *J Comp Neurol* 341 (3), 375–92. [PubMed: 7515081]
50. Takada M et al. (2004) Organization of prefrontal outflow toward frontal motor-related areas in macaque monkeys. *Eur J Neurosci* 19 (12), 3328–42. [PubMed: 15217388]
51. Cai X and Padoa-Schioppa C (2014) Contributions of orbitofrontal and lateral prefrontal cortices to economic choice and the good-to-action transformation. *Neuron* 81 (5), 1140–51. [PubMed: 24529981]
52. Hunt LT et al. (2015) Capturing the temporal evolution of choice across prefrontal cortex. *Elife* 4.
53. Haber SN and Knutson B (2010) The reward circuit: linking primate anatomy and human imaging. *Neuropsychopharmacology* 35 (1), 4–26. [PubMed: 19812543]
54. Rudebeck PH et al. (2017) Specialized Representations of Value in the Orbital and Ventrolateral Prefrontal Cortex: Desirability versus Availability of Outcomes. *Neuron* 95 (5), 1208–1220 e5. [PubMed: 28858621]
55. Lara H and Wallis JD (2014) Executive control processes underlying multi-item working memory. *Nat Neurosci* 17 (6), 876–83. [PubMed: 24747574]
56. Messinger A et al. (2009) Multitasking of attention and memory functions in the primate prefrontal cortex. *J Neurosci* 29 (17), 5640–53. [PubMed: 19403830]
57. Schafer RJ and Moore T (2011) Selective attention from voluntary control of neurons in prefrontal cortex. *Science* 332 (6037), 1568–71. [PubMed: 21617042]
58. Buschman TJ and Miller EK (2009) Serial, covert shifts of attention during visual search are reflected by the frontal eye fields and correlated with population oscillations. *Neuron* 63 (3), 386–96. [PubMed: 19679077]
59. Tremblay S et al. (2015) Attentional filtering of visual information by neuronal ensembles in the primate lateral prefrontal cortex. *Neuron* 85 (1), 202–15. [PubMed: 25500502]
60. Gottlieb J (2012) Attention, learning, and the value of information. *Neuron* 76 (2), 281–95. [PubMed: 23083732]
61. Maunsell JH (2004) Neuronal representations of cognitive state: reward or attention? *Trends Cogn Sci* 8 (6), 261–5. [PubMed: 15165551]
62. Krajbich I et al. (2010) Visual fixations and the computation and comparison of value in simple choice. *Nat Neurosci* 13 (10), 1292–8. [PubMed: 20835253]
63. McGinty VB et al. (2016) Orbitofrontal Cortex Value Signals Depend on Fixation Location during Free Viewing. *Neuron* 90 (6), 1299–311. [PubMed: 27263972]
64. Carmena JM et al. (2003) Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol* 1 (2), E42. [PubMed: 14624244]
65. Donoghue JP (2002) Connecting cortex to machines: recent advances in brain interfaces. *Nat Neurosci* 5 Suppl, 1085–8. [PubMed: 12403992]
66. Naselaris T et al. (2009) Bayesian reconstruction of natural images from human brain activity. *Neuron* 63 (6), 902–15. [PubMed: 19778517]
67. Nishimoto S et al. (2011) Reconstructing visual experiences from brain activity evoked by natural movies. *Curr Biol* 21 (19), 1641–6. [PubMed: 21945275]
68. Huth AG et al. (2016) Decoding the Semantic Content of Natural Movies from Human Brain Activity. *Front Syst Neurosci* 10, 81. [PubMed: 27781035]
69. Horikawa T et al. (2013) Neural decoding of visual imagery during sleep. *Science* 340 (6132), 639–42. 15958. [PubMed: 23558170]
70. Bornstein AM et al. (2017) Reminders of past choices bias decisions for reward in humans. *Nat Commun* 8, 15958. [PubMed: 28653668]

71. Bornstein AM and Norman KA (2017) Reinstated episodic context guides sampling-based decisions for reward. *Nat Neurosci* 20 (7), 997–1003. [PubMed: 28581478]
72. Pasley BN et al. (2012) Reconstructing speech from human auditory cortex. *PLoS Biol* 10 (1), e1001251. [PubMed: 22303281]
73. Shenoy KV et al. (2013) Cortical control of arm movements: a dynamical systems perspective. *Annu Rev Neurosci* 36, 337–59. [PubMed: 23725001]
74. Foster DJ (2017) Replay Comes of Age. *Annu Rev Neurosci* 40, 581–602 [PubMed: 28772098]

Text Box 1.

Decoding other neural measures. Several studies have used decoding approaches to extract information from the BOLD signal measured via functional magnetic resonance imaging (fMRI). Researchers have been able to reconstruct images [66] and movies [67] that a subject is viewing from BOLD activity, as well as the semantic content of the movie [68]. It even seems possible to decode the visual content of dreams. Subjects were scanned while asleep in an fMRI, woken, and asked to report what they were dreaming. A classifier, that had previously been trained using the subjects' BOLD responses when viewing visual images, was able to decode the visual content of the dreams at levels significantly above chance [69]. Decoding methods have also been used to probe decision-making. Context can bias decision-making: if a cue had a low value in a given context, recalling that context decreases the cue's current value [70]. For example, recalling a previous episode where one had a poorly cooked steak, would decrease the likelihood of choosing steak. By decoding scenes from the parahippocampal BOLD response, researchers showed that the strength of the choice bias was directly proportional to how well the previous episodic memory was recalled [71]. These studies illustrate the power of decoding analyses to uncover otherwise unobservable states. However, the BOLD response takes several seconds and likely lacks the temporal resolution necessary to measure a cognitive process in real-time. More promising has been studies that have used intracranial electrocorticography signals from patients undergoing surgery for the treatment of epilepsy. For example, it has been possible to decode speech from neural activity in the superior temporal cortex evoked by speech [72]. To date, such methods have not tested cognitive processes, in part due to the difficulty of performing extensive cognitive testing in the surgical environment.

Text Box 2.

Dynamical perspectives of neural ensembles. Much of neuroscience has adopted a representational view of neural activity, whereby the goal is to ascertain what information is represented by neural firing. Although this approach has been extremely successful in explaining some regions of the brain (e.g. primary visual cortex), it has struggled when applied to other systems (e.g. primary motor cortex). An alternative, is to take a dynamical perspective, whereby the goal is to predict future states of neural activity based on the current state [73]. A key prediction of this perspective is that activity in the system will be only loosely tied to behavioral parameters, since many of the signals will instead be related to internal processes that help to compose the outputs. Consequently, the dynamical perspective has increased in influence as investigators have moved towards ensemble recordings and realized that behavioral parameters are frequently poorly encoded despite the decoder being able to extract information from the neural activity [3].

An important goal of the dynamical perspective is to identify the states through which the system traverses, particularly ‘attractor’ states. In a dynamical system, the attractor state or attractor basin, is the subset of the state space towards which the system will tend to move, irrespective of the starting conditions of the system. For example, a ball placed in a bathroom sink will move towards the plughole, irrespective of where it is initially placed: thus, the plughole is the attractor state. Consider the slightly more distorted sinks in Figure I that have multiple troughs. In sink A, if the ball is placed low, it will end up in the closest trough, but if it is placed higher, it will end up in the opposite trough. However, the symmetry of the sink ensures that across all starting positions there will be an equal number of balls in each trough. In contrast, in sink B, most of the balls starting on the right will end up at trough 2, as will some of the balls starting on the left. With respect to decision-making, the attractor states (troughs) could correspond to either option, with the depth of the troughs proportional to the value of the option.

Outstanding questions

- Is neural vacillation unique to decision-making or is it a broader mechanism used whenever multiple pieces of information must be processed?
- How are stimulus-reward associations learned in the first place and where are they stored?
- What is the role of neural oscillations in decision-making? Are they responsible for controlling vacillation? Are they responsible for regulating the interaction of different brain regions?
- How are decisions translated into actions? Parsing the valuation of goods in the environment from action selection simplifies the complexity of decision-making but raises the question of how the action can influence the valuation process. For example, one might favor a close, mediocre restaurant over a distant, good restaurant. How could such tradeoffs be resolved if goods and actions are separately valued?
- To what extent can the processes underlying value-based decision-making be extrapolated to other kinds of decision-making, such as perceptual decision-making (and vice versa!)?
- Is decision-making a unitary process and, if not, how are conflicts between competing systems resolved? For example, high calorie foods have a high value from a biological perspective, ensuring the organism's survival, but lower value if one is trying to lose weight. How are these separate values used to determine a choice?

Highlights

- Recent advances in analytic methods and high-channel count recordings have raised the possibility of reading out cognitive processes directly from the brain, as opposed to inferring cognitive processes indirectly from behavior.
- Decoding neural activity has been used to understand decision-making by using place cell activity in the hippocampus or value-selective neural responses in orbitofrontal cortex.
- Decoding could have broad applications for measuring other cognitive processes directly from neural activity, such as attention, working memory and reasoning.

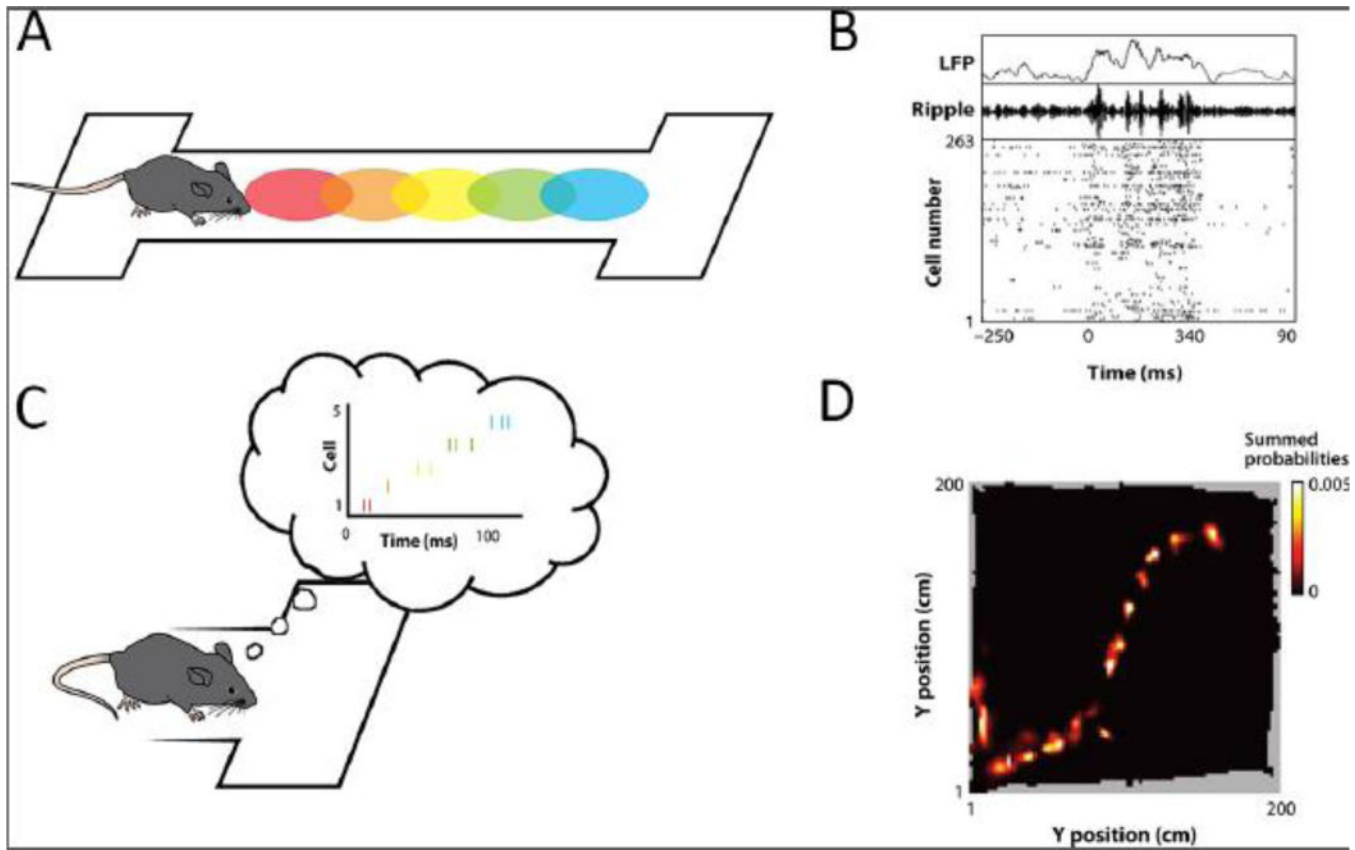


Figure 1. Replay events decoded from hippocampal place cells.

(A) As a rat walks along a linear track, different hippocampal neurons, with different place fields (colored ellipses) fire at different locations along the track. (B) *Top*: Raw LFP trace recorded from hippocampus. *Middle*: Ripple-filtered LFP trace. *Bottom*: Spike rasters from 263 simultaneously recorded hippocampal neurons, indicating the increased firing during ripple events. (C) When the rat pauses at the end of the track, the cells rapidly fire in the same order in which the place fields were encountered. (D) A Bayesian decoder allowed the spatial trajectory of the replay event to be constructed from the pattern of firing in the 263 neurons. Figure adapted with permission [74].

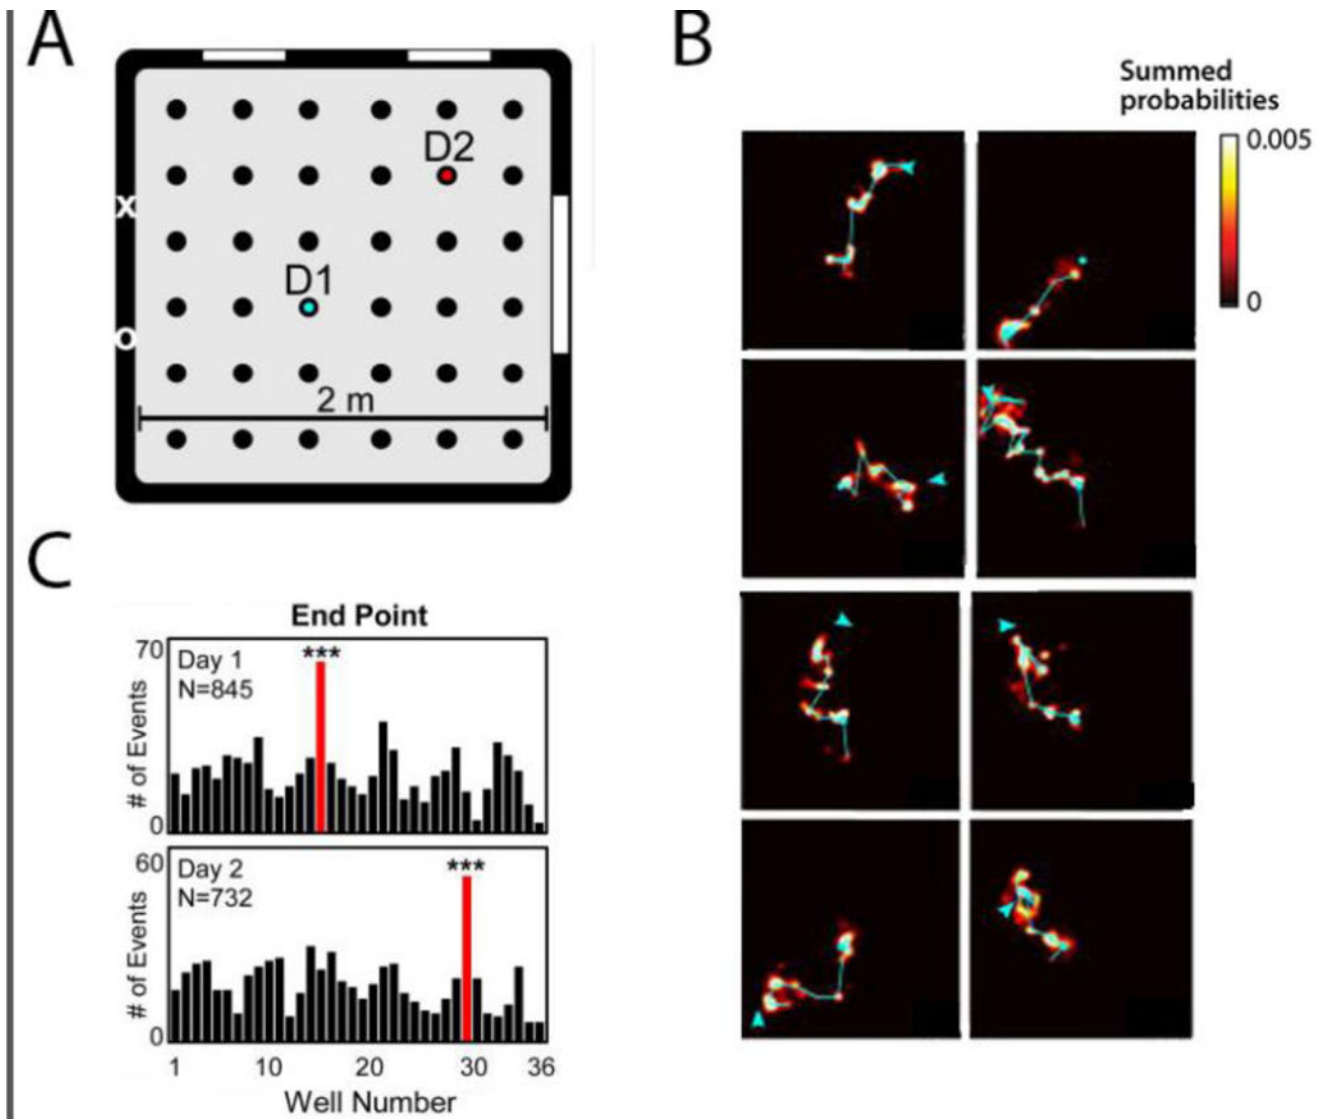


Figure 2. Decoding spatial trajectories in the open field.

(A) Rats explored an open-field, searching for a reward hidden in one of 36 wells. The reward was hidden at D1 and D2 on the first and second day, respectively. (B) Eight examples showing replay events decoded on the first day. The trajectories often traveled from the rat's current location (cyan triangle) to the reward location (cyan circle). (C) Frequency with which the end points of constructed replay trajectories occurred at different well locations. There was a bias for trajectories to end at the well in which the reward was hidden. Figure adapted with permission [8].

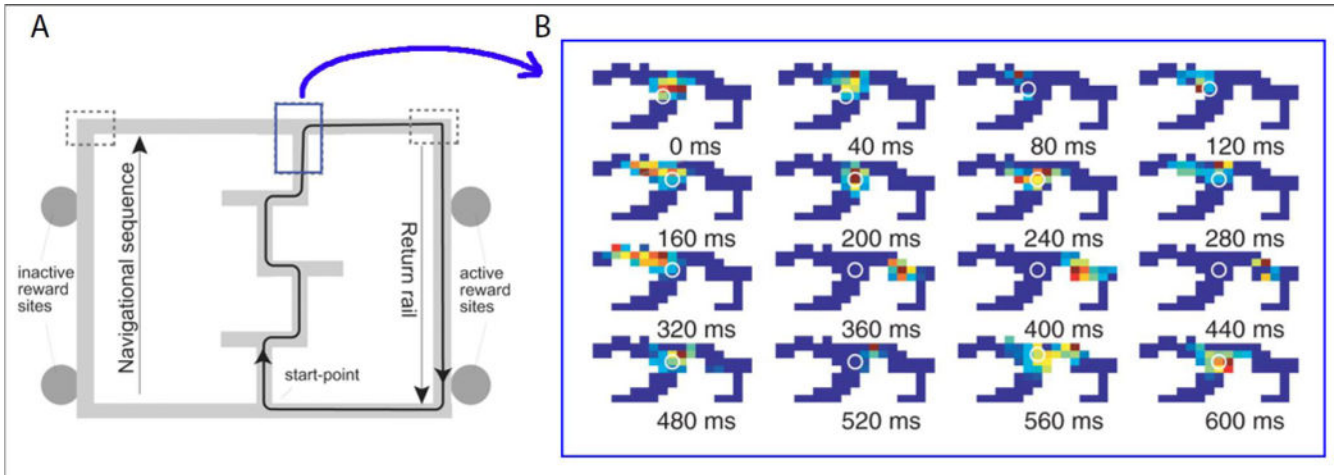


Figure 3. Decoding hippocampal activity on a single T-maze trial.

(A) A rat runs along a maze that requires them run down leftward or rightward arms on alternating trials. The choice point is outlined by the blue box. (B) The results of Bayesian decoding of spatial position from a hippocampal neural ensemble. The color scale indicates the location decoded with the highest probability and the white circle indicates the position of the rat. The hippocampal ensemble first represents locations in the leftward arm (from about 160 ms until 360 ms), followed by locations in the rightward arm (200 ms until 400 ms). Figure adapted with permission [9].




Figure 4. Decoding value from OFC neurons during decision-making.

(A) The plots show six trials, selected at random, in which the animal was choosing between pairs of the reward-predictive cues. The value of the two cues is indicated above the plots (Ch = chosen picture value, UnCh = unchosen picture value). The individual lines of the plots indicate the decoder output, which is the probability that the neural activity is consistent with the representation of one of the two cues. The lines are color-coded according to whether the value related to the picture that the animal ultimately chose, the picture he did not choose, or the unavailable pictures that were not presented on that trial. (B) There were significantly more flip-flops on error trials relative to correct trials (t-test, $p < 0.005$). (C) Strength of correlation between decoder output (posterior probability) and choice response time, for the chosen (red) and unchosen (blue) options. Stronger representation of the chosen option produced faster response times (negative correlation), while the opposite was true for the unchosen option. (D) The number of states per trial, averaged for each session. States were defined as a probability greater than 0.5 for longer than 80 ms. Chosen states were more prevalent than unchosen states.




Figure I. Representation of different attractor basins.

The color of the balls indicates whether they end up at trough 1 (blue) or trough 2 (red).