Neural computations associated with goal-directed choice

Antonio Rangel and Todd Hare

In goal-directed decision-making, animals choose between actions that are associated with different reward outcomes (e.g., foods) and with different costs (e.g., effort). Rapid advances have been made over the past few years in our understanding of the computations associated with goal-directed choices, and of how those computations are implemented in the brain. We review some important findings, with an emphasis on computational models, human fMRI, and monkey neurophysiology studies.

Address
HSS & CNS, California Institute of Technology, 1200 E. California Boulevard, Pasadena, CA, United States

Corresponding author: Rangel, Antonio (rangel@hss.caltech.edu)

Introduction

Consider a canonical decision-making problem. Every day a hungry animal is placed at the bottom of a Y-maze and is allowed to run towards the upper left or right to collect a reward. The left arm leads to a highly liked food, but is also associated with a high cost since the animal is required to swim to reach it. The right end leads to a less desirable outcome, but does not require swimming. The foods randomly change each day. How does the animal decide which course to take?

A growing body of work has shown that this problem can be solved using two very different approaches [1–4]. In one approach animals learn the value of each action through trial-and-error using reinforcement learning, and then take the action with the highest learned value [2,4–9]. This strategy requires little knowledge on the part of the subject and can account for multiple aspects of behavior in many domains, but is only able to pick the optimal action on average. In another approach, animals estimate the value associated with each action in every trial using knowledge about their costs and benefits. With sufficient knowledge this approach, often called 'goal-directed' or 'model-based' decision-making [7,8,10], can

do much better since it is able to pick the optimal action in every trial [10].

Over the past decade significant advances have been made in understanding how the brain makes goal-directed choices. We review important findings from the past few years, as well as some of the most pressing open questions. Owing to space limitations we do not attempt to be comprehensive.

Computational framework

Models from psychology and economics suggest that goal-directed choice requires the following computations. First, the brain computes *stimulus values* that measure the value of the outcomes generated by each action. Second, it computes *action costs* that measure the costs associated with each course of action. Third, it integrates them into *action values* given by

$$Action\,Value = Stimulus\,Value - Action\,Cost.$$

Finally, the action values are compared in order to make a choice.

We now describe what is known about the computational and neural basis of these processes.

How are stimulus values encoded?

Several human fMRI studies have placed individuals in simple choice situations and have found that BOLD activity in the medial orbitofrontal cortex (mOFC) correlates with behavioral measures of stimulus values [11,12,13•]. These findings are consistent with monkey neurophysiology studies that have found stimulus value coding in OFC neurons during choice tasks [14,15••,16•,17] (Figure 1). Note, however, that we must be cautious when comparing OFC findings across species owing to potential connectivity differences [18–23].

In many circumstances the mapping from actions to outcomes is probabilistic, and the outcomes arise only after a delay. The above-model is easily extended to incorporate these complications by setting

$$Action\,Value = E[Discounted\,Stimulus\,Value|Action] - $$
$$E[Discounted\,Action\,Cost|Action],$$

where $E[\,]$ denotes the expectation operator, and both stimulus values and action costs are discounted owing to the temporal delay with which they occur. Human fMRI papers have investigated the encoding of stimulus values

in this more complex setting. Tom *et al.* [24] studied choices between random and certain monetary prizes and found that activity in mOFC correlated with stimulus values that were consistent with the predictions of Prospect Theory [25], a successful behavioral theory of how people evaluate risky payoffs. Levy *et al.* [26] have extended these findings by showing that the mOFC encodes stimulus values when there is ambiguity (i.e., incomplete knowledge) about the likelihood of different outcomes. Kable and Glimcher [27] studied choices between immediate and delayed monetary payoffs and found stimulus value signals for the delayed payoffs in an area of ventromedial prefrontal cortex adjacent to mOFC. The signals were a good match with the behavioral data: individuals with stimulus value signals that discounted the future more steeply were also less likely to delay gratification.

Several studies have begun to characterize the code used by the OFC to encode values. Two important findings have come from these investigations. First, there does not seem to be an anatomical dissociation between appetitive and aversive stimulus value coding. In related fMRI studies, Tom *et al.* [24] found that the same area of mOFC correlated positively with potential monetary gains and negatively with potential losses, and Plassmann *et al.* [28] found that mOFC activity correlated positively with the appetitiveness of foods and negatively with their aversiveness. Second, stimulus value signals are context independent. Padoa-Schioppa and Assad [16°] found that the stimulus value signals satisfy 'menu-independence', so that the value assigned to one stimulus does not depend on which other stimuli are in the choice set. This property is important because it guarantees consistency across decision situations. More recently Padoa-Schioppa [15°°] has found that stimulus value neurons in the OFC exhibit range adaptation, so that a given increase in value induces a larger change in firing rates when the range of potential values for a stimulus is small. This property is important because it allows single neurons to encode stimulus values in contexts where the range of potential values is small (e.g., $0–$2) or large (e.g., $0–$100 000).

How are stimulus values computed?
A popular theory states that stimulus values are learned through reinforcement learning and retrieved in OFC at the time of choice [2,9,29]. Although some evidence suggests that this process is at work in settings where animals repeatedly face a small number of stimuli [30], it cannot account for all observed behavior because humans are able to evaluate novel stimuli. We propose an alternative theory of stimulus value computation that takes advantage of the fact that most stimuli are complex bundles of more basic attributes (e.g., foods can be described by a list of perceptual properties such as size, color, and texture). Animals can evaluate any stimulus by learning the value of basic attributes, and then integrating

over them. This mechanism is more efficient because less information needs to be learned.

The conditioning theory has been extensively reviewed elsewhere [2,9,18,31], therefore we focus on what is known about the validity of the integration hypothesis. If it is correct, there should be neurons encoding the attributes associated with the stimulus being evaluated. For example, in the case of choices among lotteries, the model predicts the existence of neurons encoding statistics such as the expected value and variance.

Evidence consistent with the integration hypothesis comes from Kennerley *et al.* [32]. They found neurons in ACC, OFC, and LPFC neurons encoding the probability, magnitude, and effort associated with different options. Additional evidence is provided by the fMRI experiments discussed below.

Stimulus valuation in complex decision situations
Two recent human fMRI studies provide clues about how the brain has adapted to solve more sophisticated choice problems, such as dietary decisions with long-term consequences, or complex social decisions. In order to make good choices in these domains, the brain needs to compute the value of attributes such as the impact of the choice on future health, or on others' well-being. Hare *et al.* [33°°] studied dietary choices that involve self-control. Subjects made choices between stimuli that varied in their taste and health properties, which were measured independently. The study found that activity in the OFC encoded stimulus values regardless of the extent to which health or taste considerations drove the choices. However, health information had a greater influence on the OFC value signals (and choices) when a region of left DLPFC was activated. A functional connectivity analysis suggested that DLPFC might modulate the weight placed on different attributes during value computation in OFC (Figure 1).

More evidence that OFC integrates separate attributes into a stimulus value comes from an fMRI study of charitable decision-making. Hare *et al.* [34°] found that activity in OFC correlated with behavioral measures of the value that subject assigned to the charities. Moreover, functional connectivity analyses suggest that the OFC value signal integrated inputs from anterior insula and pSTC, areas that are thought to be crucial for social cognition. These findings suggest that these other variables are computed outside OFC, and that the information is then passed there to be integrated into stimulus values.

Interestingly, in both of the studies the effect operated by modulating activity in inferior frontal gyrus (IFG, BA 9), which might serve as a conduit of 'cognitive information'

Figure 1
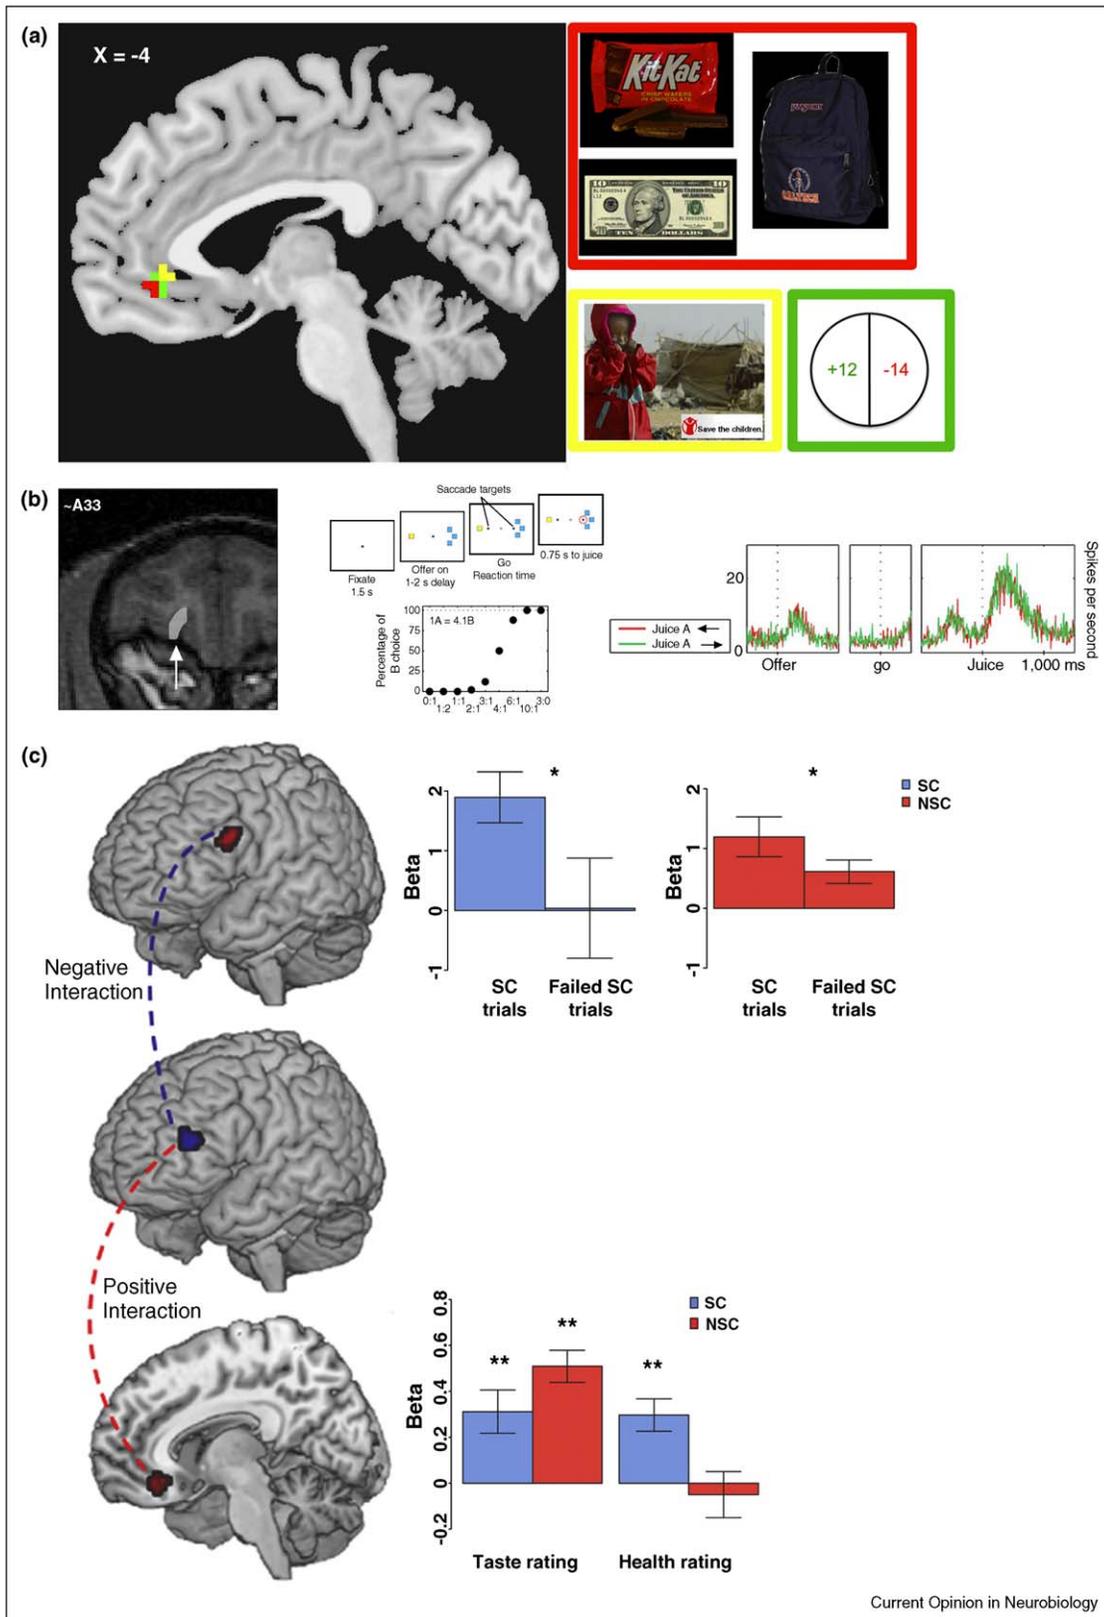
Stimulus value is reflected in VMPFC activity. **(a)** The overlay map shows the peak activations in mOFC/ACC for three fMRI studies of goal-directed decision-making. The peak from a study by Chib *et al.* [72] investigating decisions using consumer goods, food, and monetary rewards is shown in red.

into the OFC. Consistent with this hypothesis, the application of slow frequency repetitive transcranial magnetic stimulation over this area has been shown to affect purchasing [35], gambling [36], and social decisions [37,38].

How are action costs encoded and computed?

Almost every choice we make has costs associated with it. These costs come in two types. First are the costs of the actions required to obtain the stimuli, such as effort. Second are aversive stimuli that are bundled with the desired outcome. For example, purchasing a book requires giving up money. The key distinction between them is whether the cost is tied to the action or to the outcome. The distinction is meaningful because, for example, one can decrease the effort costs associated with purchasing an item (say by changing the physical distance to the store) without changing the price that has to be paid for the item.

Both types of costs have been investigated in the literature. Hare et al. [13•] studied purchasing decisions using fMRI and found that whereas the mOFC encoded the value of the items regardless of price, a more lateral area of OFC encoded the net value of the purchases, and thus reflected the costs of the stimuli. Talmi et al. [39•] studied stimulus value coding in a setting in which earning a monetary reward required getting electric shocks and found that the stimulus value signals in mOFC decreased with the size of the shocks.

Rudebeck et al. [40] investigated the difference between stimulus and effort-based action value coding. Macaques were asked to make choices between actions, without any visual stimulus associated with them, or between visual stimuli that did not have a fixed relationship with the actions used to indicate choices. The reward contingencies changed probabilistically over time. They found that lesions to the ACC sulcus, but not to the OFC, impaired action based choices, and that the opposite was true for stimulus based choices.

The results in this section suggest that there might be a dissociation between areas involved in encoding stimulus costs, predominantly in the OFC, and areas involved in encoding action costs, predominantly in ACC. This statement might seem at odds with previous findings

[19,41], but the apparent puzzle has to do with a misunderstanding in the existing literature that has treated delays in reward delivery as costs, and not as outcome attributes.

How are action values encoded and computed?

The computational model described above predicts that there should be neurons encoding each of the action values, regardless of whether the action is taken or not. Samejima et al. [42] recorded from striatal neurons during a probabilistic binary choice task and found neurons encoding the action values. In a closely related study Lau and Glimcher [43] found that about 60% of phasically active neurons in the caudate encoded either the value of particular actions (early in the trial) or the value of the chosen action (later in the trial). Kennerley et al. [32] found that most of the neurons that responded to both stimulus values and action costs were in the ACC, as opposed to OFC and LPFC. A related human fMRI study [44] found that BOLD activity in the dorsal ACC correlated with the action values during an effortful reward harvesting task.

Action value signals have been found in other areas besides the ACC. Kim et al. [45] recorded in DLPFC while monkeys chose either a left or a right eye movement associated with a larger delay reward or a smaller earlier reward. They found that a significant fraction of DLPFC neurons encoded the discounted value of the prize associated with one of the actions, but not the other, and they did before the choice was indicated by the animal. A recent human fMRI study [46] in which subjects had to chose between a button press or an eye movement in every trial found action value representations in supplementary motor areas. A related monkey neurophysiology study [47] also found signals consistent with action value coding in single SMA neurons. These findings suggest a dissociation between the OFC and areas of the ACC and SMA in goal-directed choice, with the former being specialized in the encoding of stimulus values, and the latter in the encoding of action values [19,32].

An important open question is how are the stimulus values and action costs integrated into action values. One possibility that deserves further investigation is that

---

Peak activity for choices over gambles representing both monetary gain and loss from Tom et al. [24] is shown in green. Yellow voxels represent the peak for decisions about charitable donations from Hare et al. [34•]. Examples of the stimuli associated with each peak are shown on the right inside a box of the corresponding color. **(b)** The MRI image shows the placement of electrodes in area 13 from Padoa-Schioppa and Assad [14]. A diagram of the task structure is shown in the upper middle and below that a choice curve showing the relative preference for juice A compared to juice B. The graph on the far right shows the firing rate for a stimulus selective neuron. Firing increases with the value of juice A regardless of the action required to select it. Images in part B were adapted from Ref. [14]. **(c)** The renderings on the left illustrate a potential pathway through which DLPFC activity might modulate value computations in mOFC as reported by Hare et al. [33••]. By inhibiting activity in a region of BA 46, DLPFC might bring down the value of unhealthy food items in the dietary self-control task. Dieters who successfully exercised self-control in this task had greater activity in left DLPFC compared to those subjects who did not use self-control. However, the top graph shows that within each group there was greater DLPFC activity when self-control was successful than on trials where self-control failed. The bottom graph shows that successful self-controllers incorporated both taste and health attributes into value signals computed in mOFC, whereas non-self-controllers computed values based on tastes alone.

Figure 2
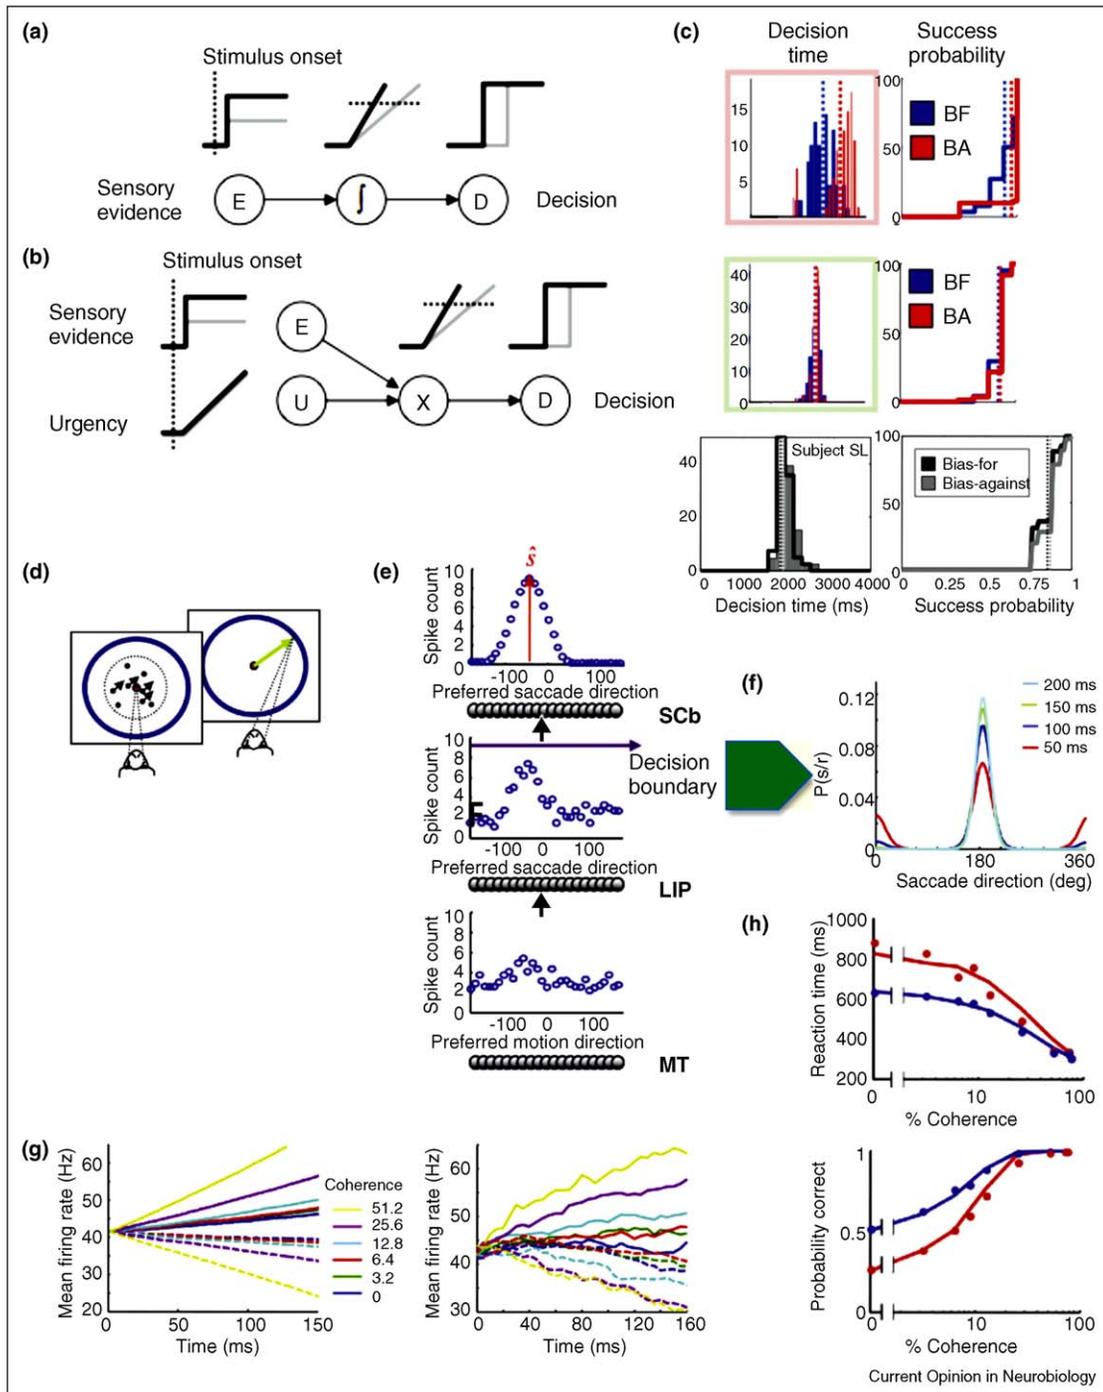
(a) Illustration of the main components of the diffusion model of perceptual decision-making. Evidence *E* in favor of a decision can be strong (black) or weak (gray) and is integrated over time. A decision is made when a common threshold is reached. **(b)** Illustration of the main components of the urgency-gating model of Cisek [60]. Now evidence is not integrated over time. Instead it is multiplied by an urgency signal *u* and a decision is made when *x* × *u* reaches a common decision threshold. **(c)** Data generated by the two models (top and middle) and actual behavior for a typical subject (bottom) in trials in which the first few pieces of evidence were biased for or against the decision. As can be seen in the bottom panel, the urgency model correctly predicts the absence of differences between the two models, but the diffusion model does not. **(d)** Continuous version of the Newsome–Shadlen perceptual discrimination task in which a subset of otherwise randomly moving dots move coherently in some direction. The animal indicates its guess about the direction of coherent movement at any time through an eye movement. Correct responses are rewarded. **(e)** Basic architecture of the Bayesian population model of Beck *et al.* [63[••]]. Each graph summarizes neural activity in an area at a particular instant of the decision task. SCb: bursting neurons in the superior colliculus. The *x*-axis denotes the preferred direction of movement for a given neuron. The *y*-axis denotes the local firing rate for that neuron. The model consists of a network with three interconnected layers of neurons with Gaussian tuning curves.

this might be done by separate ACC-striatal and OFC-striatal loops that converge in the caudate.

How are action values compared to make a choice?

The final stage in making a decision involves the comparison of the action values in order to make a choice. A significant amount of behavioral evidence suggests that the mapping from action values to choices is stochastic and follows a soft-max (or logistic) functional form, and that there is a speed-accuracy tradeoff. Two of the most important open questions in the field have to do with how the stochastic choice process is implemented: What exactly is the algorithm used by the brain to compare the action values? How is the comparison process implemented by the brain?

The first question has been addressed using a combination of modeling and psychometrics, mostly in the realm of perceptual decision-making (see Refs. [48–50] for outstanding recent reviews). A first class of models has approached the problem at the systems level by specifying the dynamic processes through which action values are compared. Most of the models that have been proposed are variations of the diffusion model. They assume that information about the value of the different actions is not directly accessible to the brain, but instead it needs to be computed on the basis of sequential Gaussian random samples of the underlying true values. Decisions are made by dynamically and optimally integrating the samples into a relative action value signal (e.g., $V_{\text{action 1}} - V_{\text{action 2}}$). The process terminates when the relative value action becomes sufficiently biased towards one of the choices [48,51–55]. Although this model has been quite successful in providing qualitative and quantitative explanations in many domains, it also has limitations that have begun to be addressed. First, the model does not generalize easily to the case of more than two alternatives. Bogacz and Gurney have proposed an extension in which decision thresholds depend on the amount of conflict between the alternatives (see also Refs. [56,57]). Second, a recent work by Cisek [58] and Ditterich [59] has shown models with urgency signals that increase with reaction time can account for a wider range of psychometric data than equivalent parameterizations of the diffusion model (Figure 2a–c). Finally, whereas diffusion models are silent about how the comparison process eventually triggers a motor response, Cisek [60] has extended the logic of these models to show how this could be accomplished, thus providing a joint model of action selection and motor planning.

The second class of models has addressed the same problem using neuronal models of the comparison process [50,61,62]. Although these models are significantly more complicated, they have been able to provide a better match to the behavioral and the neurometric data.

A new and highly promising alternative approach to these models are the Bayesian decision-making models based on probabilistic population codes of Pouget and collaborators (Figure 2d–h). In a series of recent papers [63[••],64] they have shown that as long as neurons follow approximately a Poisson-like distribution of spike counts, it is possible to make fully optimal Bayesian perceptual decisions using simple linear integration of neural activity, even in the presence of a large number of options, or with noise that changes within and across trials. Extending and testing these ideas in the realm of value based decision-making promises to be a profitable line of inquiry.

The best available evidence regarding potential substrates of the comparison process comes from the study by Kim et al. [45] described above. They found neurons in DLPFC that encoded the value of one of the actions dynamically, by ramping up their activity until a choice was made if the action was associated with the best prize, and ramping down activity otherwise. The dynamics of such neurons are consistent with the models discussed above. On the basis of human fMRI data Wunderlich et al. [46] have recently argued on that parts of the ACC might play a crucial role in the comparison process. Another hypothesis that has received a significant amount of attention is that the choice process might not be implemented in a single area, but instead might reside in the dynamics of cortico-basal-ganglia-thalamic loops [55,65–67].

Other clues about the neural basis of the comparator process come from the existence of wide spread chosen value signals reflecting the output of the comparison process. Padoa-Schioppa and Assad [14] found OFC neurons encoding the value of the chosen stimulus (often called 'chosen value' neurons). Kepecs et al. [68] found rat OFC neurons that encode either the 'chosen value' or a measure of 'uncertainty' on having selected the best stimulus. Several human fMRI choice studies have found that BOLD activity in the OFC correlates with chosen values [46,69,70]. Neurons encoding chosen values have also been found in the ACC [32,71] and the caudate [43]. At a minimum, these findings suggest that the output of the comparison process is passed to multiple areas, probably for the purpose of learning the value of actions via reinforcement learning.

MT neurons encode the instantaneous direction of motion. LIP and SCb neurons encode potential directions of motion. The population code in SCb represents the most likely direction of motion at any instant. **(f)** Activity in LIP can be passed through a Bayesian decoder to compute posterior probabilities $P(s|r)$ over the coherent direction of motion given the current levels of activity r. Note that the posteriors become sharper as time progresses. **(g)** The activity in LIP can also be used to predict firing rates as a function of coherence and time (left), which match well the predictions from Roitman and Shadlen [73] (right). **(h)** The model is also able to account for the choice and reaction time data in the multi-option choice experiment of Churchland et al. [57]. Red: four-choice experiment. Blue: two-choice experiment. **a–c** are from Ref. [60]. **d–h** are from Ref. [63[••]].

Importantly, none of the existing models has been systematically compared against neural activity during goal-directed choices. This exercise is one of the most pressing open questions for the field since very little is known about the neural substrates of action value comparisons.

Conclusions

Throughout the review we have emphasized a multitude of important and pressing open questions. However, it is important not to lose sight of the progress that has been made. We now know that OFC neurons encode stimulus values in a wide variety of contexts and that values are sensitive to internal physiological and cognitive states. We know that stimulus value signals respond to variables such as delay and risk in ways that are consistent with theories from behavioral economics. We know that during complex decisions other cortical areas (such as DLPFC, insula, and temporal cortex) are able to influence decisions by modulating the computation of stimulus values in OFC. We know that there is a dissociation between ACC, which specializes in action cost and value coding, and OFC, which specializes in stimulus value coding. We have begun to characterize some of the key computational properties of the processes though which stimulus values are compared to generate choices. Furthermore, the growing combination of computational models with sophisticated neuroscientific methods makes it likely that many of the open questions listed here will be resolved in the near future [4].

Acknowledgements

References and recommended reading

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

- of special interest
- •• of outstanding interest

1. Balleine BW: **Neural bases of food-seeking: affect, arousal and reward in corticostriatolimbic circuits**. *Physiol Behav* 2005, **86**:717-730.

2. Balleine BW, Daw N, O'Doherty J: **Multiple forms of value learning and the function of dopamine**. In *Neuroeconomics: Decision-Making and the Brain*. Edited by Glimcher PW, Fehr E, Camerer C, Poldrack RA. Elsevier; 2008.

3. Balleine BW, Dickinson A: **Goal-directed instrumental action: contingency and incentive learning and their cortical substrates**. *Neuropharmacology* 1998, **37**:407-419.

4. Rangel A, Camerer C, Montague PR: **A framework for studying the neurobiology of value-based decision making**. *Nat Rev Neurosci* 2008, **9**:545-556.

5. Dayan P, Niv Y, Seymour B, Daw ND: **The misbehavior of value and the discipline of the will**. *Neural Netw* 2006, **19**:1153-1160.

6. Dayan P: **The role of value systems in decision making**. In *Better Than Conscious? Implications for Performance and Institutional Analysis*. Edited by Engel C, Singer W. MIT Press; 2008

7. Rescola RA, Wagner AR: **A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and non-reinforcement**. In *Classical Conditioning II: Current Research and Theory*. Edited by Black AH, Prokasy WF. Appleton Century Crofts; 1972:406-412.

8. Barto AG: **Adaptive critic in the basal ganglia**. In *Models of Information Processing in the Basal Ganglia*. Edited by Houk JC, Davis JL, Beiser DG. MIT Press; 1995.

9. Niv Y, Montague PR: **Theoretical and empirical studies of learning**. In *Neuroeconomics: Decision-Making and the Brain*. Edited by Glimcher PW, Fehr E, Camerer C, Poldrack RA. Elsevier; 2008.

10. Daw ND, Niv Y, Dayan P: **Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control**. *Nat Neurosci* 2005, **8**:1704-1711.

11. Plassmann H, O'Doherty J, Rangel A: **Orbitofrontal cortex encodes willingness to pay in everyday economic transactions**. *J Neurosci* 2007, **27**:9984-9988.

12. Valentin VV, Dickinson A, O'Doherty JP: **Determining the neural substrates of goal-directed learning in the human brain**. *J Neurosci* 2007, **27**:4019-4026.

13. Hare TA, O'Doherty J, Camerer CF, Schultz W, Rangel A:
• **Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors**. *J Neurosci* 2008, **28**:5623-5630.
This human fMRI study uses a novel decision-making paradigm to dissociate three basic decision-making computations: stimulus values, stimulus costs, and prediction errors. It finds that activity in the medial OFC correlates with stimulus values, but does not represent costs, that activity in the central OFC correlates with net values, and that activity in the ventral striatum correlates best with prediction errors.

14. Padoa-Schioppa C, Assad JA: **Neurons in the orbitofrontal cortex encode economic value**. *Nature* 2006, **441**:223-226.

15. Padoa-Schioppa C: **Range-adapting representation of**
•• **economic value in the orbitofrontal cortex**. *J Neurosci* 2009, **29**:14004-14014.
Sometimes we make choices among small stake stimuli (e.g., which fruit to have for dessert), whereas at other times the stakes are much larger (e.g., which car to purchase). As a result the OFC valuation circuitry needs to be able to encode a wide range of values, while maintaining the relative precision with which it encodes them. This paper shows that OFC neurons address this problem by adapting its firing rate to the local range of values that it needs to encode. In particular, the paper shows that the increase in neuronal activity produced by a one-unit increase in the stimulus value decreases with the range of potential values of the stimulus.

16. Padoa-Schioppa C, Assad JA: **The representation of economic**
• **value in the orbitofrontal cortex is invariant for changes of menu**. *Nat Neurosci* 2008, **11**:95-102.
This study shows that OFC neurons encode absolute stimulus values. This stands in sharp contrast to relative value coding in which the value assigned to a stimulus depends on what other options are being considered. This is an important finding because absolute value coding promotes choice transitivity, which is a property of optimal choice.

17. Wallis JD, Miller EK: **Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task**. *Eur J Neurosci* 2003, **18**:2069-2081.

18. Rushworth MF, Mars RB, Summerfield C: **General mechanisms for making decisions?** *Curr Opin Neurobiol* 2009, **19**:75-83.

19. Rushworth MF, Behrens TE: **Choice, uncertainty and value in prefrontal and cingulate cortex**. *Nat Neurosci* 2008, **11**:389-397.

20. Ongur D, Price JL: **The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans**. *Cereb Cortex* 2000, **10**:206-219.

21. Ongur D, Ferry AT, Price JL: **Architectonic subdivision of the human orbital and medial prefrontal cortex**. *J Comp Neurol* 2003, **460**:425-449.

22. Hsu DT, Price JL: **Midline and intralaminar thalamic connections with the orbital and medial prefrontal networks in macaque monkeys**. *J Comp Neurol* 2007, **504**:89-111.

23. Saleem KS, Kondo H, Price JL: **Complementary circuits connecting the orbital and medial prefrontal networks with the temporal, insular, and opercular cortex in the macaque monkey**. *J Comp Neurol* 2008, **506**:659-693.

24. Tom SM, Fox CR, Trepel C, Poldrack RA: **The neural basis of loss aversion in decision-making under risk**. *Science* 2007, **315**:515-518.

25. Kahneman D, Tversky A: **Prospect Theory: an analysis of decision under risk**. *Econometrica* 1979, **4**:263-291.

26. Levy I, Snell J, Nelson AJ, Rustichini A, Glimcher PW: **The neural representation of subjective value under risk and ambiguity**. *J Neurophysiol* 2010, **103**:1036-1047.

27. Kable JW, Glimcher PW: **The neural correlates of subjective value during intertemporal choice**. *Nat Neurosci* 2007, **10**:1625-1633.

28. Plassmann H, O'Doherty J, Rangel A: **Aversive goal values are negatively encoded in the medial orbitofrontal cortex at the time of decision making**; J Neurosci under review.

29. Balleine BW, Dickinson A: **The role of incentive learning in instrumental outcome revaluation by sensory-specific satiety**. *Anim Learn Behav* 1989, **26**:.

30. Burke KA, Franz TM, Miller DN, Schoenbaum G: **The role of the orbitofrontal cortex in the pursuit of happiness and more specific rewards**. *Nature* 2008, **454**:340-344.

31. Behrens TE, Woolrich MW, Walton ME, Rushworth MF: **Learning the value of information in an uncertain world**. *Nat Neurosci* 2007, **10**:1214-1221.

32. Kennerley SW, Dahmubed AF, Lara AH, Wallis JD: **Neurons in the frontal lobe encode the value of multiple decision variables**. *J Cogn Neurosci* 2009, **21**:1162-1178.

33. Hare T, Camerer C, Rangel A: **Self-control in decision-making
•• involves modulation of the vMPFC valuation system**. *Science* 2009, **324**:646-648.
This human fMRI study investigates how the brain makes dietary decisions that involve self-control. Subjects make choices between stimuli that vary in their taste and health properties. The study finds that regardless of the extent to which subjects deploy self-control (defined as giving a higher weight to health than to taste considerations), activity in the OFC encodes the value of stimuli at the time of choice. Another finding is that the extent to which health information is represented in the OFC depends on the activity of a region of left DLPFC, which might modulate the stimulus value signals to induce dietary self-control.

34. Hare T, Camerer C, Knoepfle D, O'Doherty J, Rangel A: **Value
• computations in VMPFC during charitable decision-making incorporate input from regions involved in social cognition**. *J Neurosci* 2010, **30(2)**:583-590.
This human fMRI study investigates stimulus value coding during charitable decisions. It finds that activity in OFC correlates with behavioral measures of the value that subject assign to the charities at the time of decision-making. These areas overlap with those that have been shown in other studies to encode the value of 'non-social' stimuli. The study also shows that the stimulus value signal in OFC is likely to be computed by integrating inputs from anterior insula and pSTC, areas that are thought to be crucial for social cognition.

35. Camus M, Shimojo S, Camerer CF, O'Doherty J, Rangel A: **rTMS of right dorsolateral prefrontal cortex disrupts the computation of goal-directed values**. *Eur J Neurosci* 2008, **30**:1980-1988.

36. Knoch D, Gianotti LR, Pascual-Leone A, Treyer V, Regard M, Hohmann M, Brugger P: **Disruption of right prefrontal cortex by low-frequency repetitive transcranial magnetic stimulation induces risk-taking behavior**. *J Neurosci* 2006, **26**:6469-6472.

37. Knoch D, Pascual-Leone A, Meyer K, Treyer V, Fehr E: **Diminishing reciprocal fairness by disrupting the right prefrontal cortex**. *Science* 2006, **314**:829-832.

38. Knoch D, Schneider F, Schunk D, Hohmann M, Fehr E: **Disrupting the prefrontal cortex diminishes the human ability to build a good reputation**. *Proc Natl Acad Sci U S A* 2009, **106(49)**:20895-20899.

39. Talmi D, Dayan P, Kiebel SJ, Frith CD, Dolan RJ: **How humans
• integrate the prospects of pain and reward during choice**. *J Neurosci* 2009, **29**:14617-14626.
This paper investigates the integration of stimulus values and costs at the time of decision-making. It finds that the stimulus value of a monetary

reward is encoded less strongly in vmPFC when obtaining them is accompanied by a probabilistic cost in the form of an electric shock.

40. Rudebeck PH, Behrens TE, Kennerley SW, Baxter MG, Buckley MJ, Walton ME, Rushworth MF: **Frontal cortex subregions play distinct roles in choices between actions and stimuli**. *J Neurosci* 2008, **28**:13775-13785.

41. Roesch MR, Olson CR: **Neuronal activity in primate orbitofrontal cortex reflects the value of time**. *J Neurophysiol* 2005, **94**:2457-2471.

42. Samejima K, Ueda Y, Doya K, Kimura M: **Representation of action-specific reward values in the striatrum**. *Science* 2005, **310**:1337-1340.

43. Lau B, Glimcher PW: **Value representations in the primate striatum during matching behavior**. *Neuron* 2008, **58**:451-463.

44. Croxson PL, Walton ME, O'Reilly JX, Behrens TE, Rushworth MF: **Effort-based cost-benefit valuation and the human brain**. *J Neurosci* 2009, **29**:4531-4541.

45. Kim S, Hwang J, Lee D: **Prefrontal coding of temporally discounted values during intertemporal choice**. *Neuron* 2008, **59**:161-172.

46. Wunderlich K, Rangel A, O'Doherty JP: **Neural computations underlying action-based decision making in the human brain**. *Proc Natl Acad Sci U S A* 2009, **106**:17199-17204.

47. Sohn JW, Lee D: **Order-dependent modulation of directional signals in the supplementary and presupplementary motor areas**. *J Neurosci* 2007, **27**:13655-13666.

48. Bogacz R: **Optimal decision-making theories: linking neurobiology with behaviour**. *Trends Cogn Sci* 2007, **11**:118-125.

49. Smith PL, Ratcliff R: **Psychology and neurobiology of simple decisions**. *Trends Neurosci* 2004, **27**:161-168.

50. Wang XJ: **Decision making in recurrent neuronal circuits**. *Neuron* 2008, **60**:215-234.

51. Ratcliff R, Smith P: **A comparison of sequential sampling modles for two-choice reaction time**. *Psychol Rev* 2004, **111**:333-367.

52. Ratcliff R, McKoon G: **The diffusion decision model: theory and data for two-choice decision tasks**. *Neural Comput* 2008, **20**:873-922.

53. Usher M, McClelland J: **The time course of perceptual choice: the leaky, competing accumulator model**. *Psychol Rev* 2001, **108**:550-592.

54. Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD: **The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced choice tasks**. *Psychol Rev* 2006, **113**:700-765.

55. Bogacz R, Gurney K: **The basal ganglia and cortex implement optimal decision making between alternative actions**. *Neural Comput* 2007, **19**:442-477.

56. Niwa M, Ditterich J: **Perceptual decisions between multiple directions of visual motion**. *J Neurosci* 2008, **28**:4435-4445.

57. Churchland AK, Kiani R, Shadlen MN: **Decision-making with multiple alternatives**. *Nat Neurosci* 2008, **11**:693-702.

58. Cisek P, Puskas GA, El-Murr S: **Decisions in changing conditions: the urgency-gating model**. *J Neurosci* 2009, **29**:11560-11571.

59. Ditterich J: **Evidence for time-variant decision making**. *Eur J Neurosci* 2006, **24**:3628-3641.

60. Cisek P: **Integrated neural processes for defining potential actions and deciding between them: a computational model**. *J Neurosci* 2006, **26**:9761-9770.

61. Wang XJ: **Probabilistic decision making by slow reverberation in cortical circuits**. *Neuron* 2002, **36**:955-968.

62. Lo CC, Wang XJ: **Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks**. *Nat Neurosci* 2006, **9**:956-963.

63. Beck JM, Ma WJ, Kiani R, Hanks T, Churchland AK, Roitman J,
•• Shadlen MN, Latham PE, Pouget A: **Probabilistic population
    codes for Bayesian decision making**. *Neuron* 2008, **60**:1142-
    1152.
This paper provides a model of how fully optimal Bayesian perceptual
decisions can be made using population codes under plausible neuro-
biological assumptions.

64. Ma WJ, Beck JM, Latham PE, Pouget A: **Bayesian inference with
    probabilistic population codes**. *Nat Neurosci* 2006, **9**:1432-1438.

65. Graybiel AM: **The basal ganglia: learning new tricks and loving
    it**. *Curr Opin Neurobiol* 2005, **15**:638-644.

66. Mink JW: **The basal ganglia: focused selection and inhibition of
    competing motor programs**. *Prog Neurobiol* 1996, **50**:381-425.

67. Redgrave P, Prescott TJ, Gurney K: **The basal ganglia: a
    vertebrate solution to the selection problem?** *Neuroscience*
    1999, **89**:1009-1023.

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

69. Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ: **Cortical
    substrates for exploratory decisions in humans**. *Nature* 2006,
    441:876-879.

70. Hampton AN, Bossaerts P, O'Doherty JP: **The role of the
    ventromedial prefrontal cortex in abstract state-based
    inference during decision making in humans**. *J Neurosci* 2006,
    26:8360-8367.

71. Kennerley SW, Wallis JD: **Evaluating choices by single
    neurons in the frontal lobe: outcome value encoded
    across multiple decision variables**. *Eur J Neurosci* 2009,
    29:2061-2073.

72. Chib VS, Rangel A, Shimojo S, O'Doherty JP: **Evidence for a
    common representation of decision values for dissimilar
    goods in human ventromedial prefrontal cortex**. *J Neurosci*
    2009, **29**:12315-12320.

73. 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.

The goal-directed system learns about the contingency between actions and outcomes and ensures that behavior is appropriate given our current desire for these outcomes. When we learn new complex tasks, such as how to drive, our goal-directed system is actively engaged as we plan each maneuver. Conversely, model-based selection should be associated with greater goal-directed sensitivity to devaluation, due to the online incorpo-ration of outcome value information into choices between candidate actions (Daw et al., 2005; Dolan & Dayan, 2013). Reinforcement learning task On each trial, participants were presented with a choice between two fractals, each of which commonly (70 %; see Fig. 1A, white arrows) led to a particular second state displaying another fractal.