



Curiosity, information demand and attentional priority

Jacqueline Gottlieb^{1,2,3}, Michael Cohanpour^{1,3}, Yvonne Li^{1,3},
Nicholas Singletary^{1,3} and Erfan Zabeh^{3,4}

To make adaptive decisions in multi-dimensional environments, animals must infer the relevant features and afford them priority for the control of learning and actions. Prioritizing sources of information is the role of executive control and attention, but very little is known about the mechanisms by which the brain computes relevance or priority. We present a new decision-theoretic view of priority as a cognitive state that implements information gathering policies. This response is distinct from representations of information utility and it is recruited when the information serves instrumental and non-instrumental motives, suggesting that curiosity and instrumental information demand are more closely related than is traditionally assumed. We discuss the broader relevance of these findings for decision making in realistic, complex situations.

Addresses

¹ Department of Neuroscience, Columbia University, United States

² Kavli Institute for Brain Science, Columbia University, United States

³ Mortimer B. Zuckerman Mind Brain Behavior Institute, Columbia University, United States

⁴ Department of Biomedical Engineering, Columbia University, United States

Corresponding author: Gottlieb, Jacqueline (jg2141@columbia.edu)

Current Opinion in Behavioral Sciences 2020, 35:83–91

This review comes from a themed issue on **Curiosity (Explore versus Exploit)**

Edited by **Daphna Shohamy** and **Ran Hassin**

<https://doi.org/10.1016/j.cobeha.2020.07.016>

2352-1546/© 2020 Published by Elsevier Ltd.

Introduction

The burgeoning literature on curiosity that is reviewed in this volume reflects a renewed interest among neuroscientists in core questions of motivated cognition. How do humans and animals become motivated to *know*? And, given the practically infinite set of things they could potentially know, how do they decide *which* information to know? Despite having been clearly formulated by Daniel Berlyne in the mid-20th century [1], these questions remained dormant in neuroscience and cognitive

psychology only to be invigorated in the recent literature on curiosity.

The emerging literature on curiosity has relied on novel tasks of ‘information demand’, in which participants can ask for advance information about a situation. Conceptually related to ‘process tracking’ tasks in cognitive psychology [2,3], tasks of information demand go beyond the traditional approach of providing participants with clearly defined, experimenter-selected stimuli — for example, attend to a Gabor patch and report its orientation. Instead, they allow participants (humans or animals) to decide when, how much, and what type of information to request about the situation.

Converging evidence from these studies sheds light on the first question raised by Berlyne regarding the *motives* for information demand [4,5]. The results show that humans and animals are motivated to *know* even when the knowledge is costly and non-instrumental (i.e. cannot be exploited for material gains), a desire that drives curiosity [6–13]. In addition, animals are motivated to demand ‘instrumental’ information that can be exploited for material gains (i.e., guides subsequent actions). Finally, animals seek information that signals desirable (but not undesirable) outcomes, an affective or hedonic drive that can interfere with theoretically optimal sampling strategies.

A parallel line of research sheds light on the second, closely related question raised by Berlyne regarding information gathering *strategies*. Given that one has a goal, how does one decide *which* source of information to sample in pursuit of this goal? If you want to understand the origins of the universe, which book do you read or which podcast do you listen to? And if you want to make an investment decision, which advisor do you contact or which newspaper do you read?

In this chapter, we focus on emerging evidence on this second question — the organization of information gathering strategies. As the discussion will show, these are essentially attentional questions, and their neural correlates emerge in the fronto-parietal network involved in executive control and attention. We will discuss evidence supporting a view of attentional priority as an internal cognitive state that facilitates information gathering and is distinct from economic utility. Moreover, priority serves information sampling in instrumental and non-instrumental conditions, suggesting that curiosity and

instrumental information demand are more closely related than is traditionally assumed. We end by discussing the significance of the findings for higher-level behaviors and how we detect *relevance* in complex, realistic situations.

Attentional priority as internal optimization

The classical questions posed by Berlyne about the motives and strategies for information demand probe how animals go beyond merely *reacting* to information that is given to them — and *proactively* seek information using ‘question and answer’ strategies. Perhaps our most obvious interrogative strategies are *active sensing* behaviors through which we orient sensory receptors to touch, listen to or look at selected sensory cues. ‘Attention’ refers to the cognitive mechanisms that accompany active sensing behaviors, identify relevant information and link it to learning and actions.

Converging neurophysiological data show that, in monkeys and humans, spatial attention and rapid eye movements (saccades) are controlled by ‘priority’ responses — sparse representations of rich visual scenes that selectively encode attention-worthy locations. Priority responses are found in several nodes of the fronto-parietal network, including the dorsolateral prefrontal cortex, the frontal eye field, inferior parietal area 7A and the lateral intraparietal area (LIP) [14–16]. These areas contain spatially tuned neurons that have visual receptive fields (RF) and selectively encode attention-worthy locations, and have been shown to contribute to attentional modulations of visual responses and the planning of eye movements in subcortical structures [17–19].

While priority responses are elicited in the fronto-parietal network in a range of conditions, a longstanding question has been how these responses arise. Traditional theories have revolved around the dichotomy between ‘bottom-up’ versus ‘top-down’ modes of control. These modes are related to, respectively, the visual salience (contrast) and task-relevance of stimuli and they strongly influence behavior and visual priority [14,15,19]. However, there is increased recognition that this is an insufficient dichotomy [20]. First, the ‘bottom-up’/‘top-down’ distinction does not sufficiently capture empirical observations because ‘bottom-up’ salience is strongly context dependent [21] and attention is also captured by novelty, habits and reward associations, which do not fit neatly into a ‘bottom-up’ or ‘top-down’ category [20–22]. Second and perhaps most important, both constructs are merely descriptive and fail to address the core computational question in attention control. Beyond merely naming the factors that influence our attention, what is the general principle that explains how a stimulus becomes attention worthy in a situation?

Because sensory stimuli convey *information*, this question is closely aligned with questions of information demand, suggesting that attention can be described in a decision-theoretic perspective. Decision research has typically considered how animals select pragmatic actions that exert *outward* physical force — for example, when one asks for an apple or presses a button to deliver more juice. However, this framework can be equally extended to cognitive actions that exert *inward* effects — for instance, prioritize internal representations, and ask how the brain estimates the costs and benefits of these actions [4,23,24]. This is the key insight underlying Bayesian and utility-based accounts of attention control, which we discuss next.

A Bayesian account of priority

Because the decision-theoretic definition of ‘information’ is in terms of a reduction of uncertainty, the most natural theoretical account of attention is in the Bayesian framework of reasoning under uncertainty [4,23,24]. Converging evidence indicates that humans approximate Bayesian strategies but also show sub-optimal departures from these strategies including, as we review below, in attention allocation [26]. Thus, the Bayesian prescription for attention is unlikely to provide a full account of behavior but it is a useful benchmark to evaluate empirical observations. Let us start by briefly reviewing this prescription followed by empirical observations.

Bayesian logic assumes that decision makers form expectations (beliefs) about probabilistic events and update these expectations when receiving new information. The optimal rule governing this updating, described by the Bayesian formula, is to derive *posterior* beliefs by combining *prior expectations* with *new information*, weighting each factor by its reliability. This logic has traditionally been applied to laboratory scenarios in which participants are *given* sensory signals; for instance, one may be shown a cloud of dots with stochastic motion, and be asked to judge the probability that the majority of dots move in a rightward direction. To describe *information demand*, we must consider that in natural settings, individuals are not simply given a stimulus but must first decide to attend to the stimulus (i.e. ‘demand its information’). Consider a simple example in which a pedestrian reaches an intersection and wants to know the traffic light color to decide whether to wait or proceed. The pedestrian cannot discriminate the traffic light color in her peripheral vision but must first decide to look at (or attend to) the light. The decision to look, therefore, cannot depend on whether the traffic light emits a red or green color. Instead, according to Bayesian logic, it depends on how much more *accurate* the pedestrian expects to become *if* she were to discriminate the traffic light color.

In the Bayesian framework, expectations are described as probability distributions over a set of alternatives (for

example, the probability that waiting or crossing is the better decision). Accuracy, in turn, is measured by the ‘flatness’ of these distributions. Distributions that are flatter assign similar probabilities to alternative hypotheses and have lower accuracy (or, equivalently, higher entropy or higher uncertainty). Distributions that are more peaked favor a small subset of alternatives and have higher accuracy (or, equivalently, lower uncertainty) [3]. Using Bayesian logic, one can also predict future belief states and compute the *expected information gain* (EIG) – the accuracy of the posterior relative to the prior belief states – for alternative observations. In our simple example, a pedestrian can compute how much more certain she would become about crossing or waiting, if she saw that the traffic light showed a red or green color (relative to making a different discrimination (for instance, whether a cloud is white or gray).

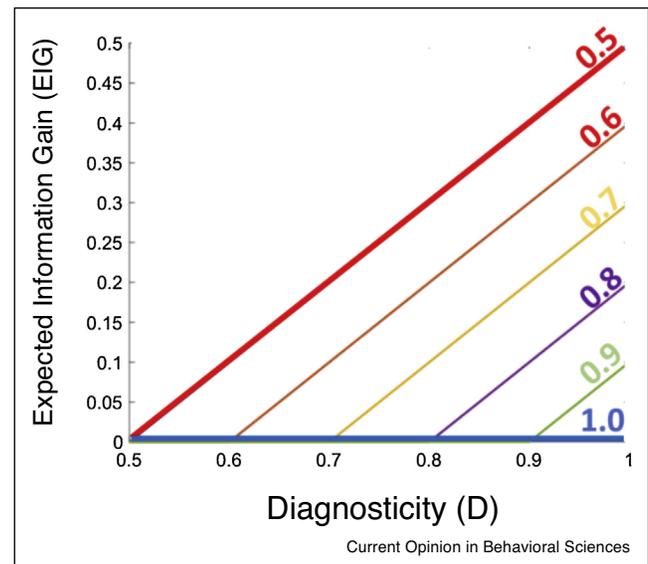
Applying the Bayesian formula shows that EIG depends on two factors: the decision maker’s *prior uncertainty* and the *diagnosticity* (*reliability*) of the available information. Both quantities are measured as prediction accuracy. A prior of 1.0 indicates that one has no uncertainty and can be 100% accurate *before* inspecting a cue; in contrast, a prior of 0.5 indicates maximal uncertainty (for a decision with 2 alternatives). Likewise, diagnosticity of 1.0 describes a perfectly accurate cue and diagnosticity of 0.5 describes a random, irrelevant cue (equivalent to % validity in traditional cueing tasks). Finally, EIG can be defined as the increase in accuracy of the posterior relative to prior expectations – i.e., the gains in accuracy expected after inspecting a cue [27].

As Figure 1 illustrates, EIG depends on a multiplicative interaction between diagnosticity and prior uncertainty. For each level of prior uncertainty, EIG increases as a function of D , but only if D exceeds prior accuracy — that is, if the information is more accurate than the decision maker’s prior beliefs. If one has high prior uncertainty (red line, 0.5) one can expect to gain information in proportion to diagnosticity. However, if one has no prior uncertainty (blue), even a perfectly reliable cue is merely redundant and has 0 EIG. Thus, a perfectly Bayesian actor should have maximal information demand if she has *high* prior uncertainty and information of high *reliability* (red, $D = 1.0$). In contrast, she should demand *no* information if she has perfect prior certainty (blue).

Uncertainty and diagnosticity are ingredients of attentional priority

While a Bayesian view of attention is common in theoretical models [23,28,29] and supported by studies of eye movement control [30,31,32], it has been relatively ignored in the neurophysiological literature, which has instead focused on the effect of rewards on attention and fronto-parietal activity (e.g. Refs. [25,33–35]). This focus is motivated by the fact that decision makers often exploit

Figure 1



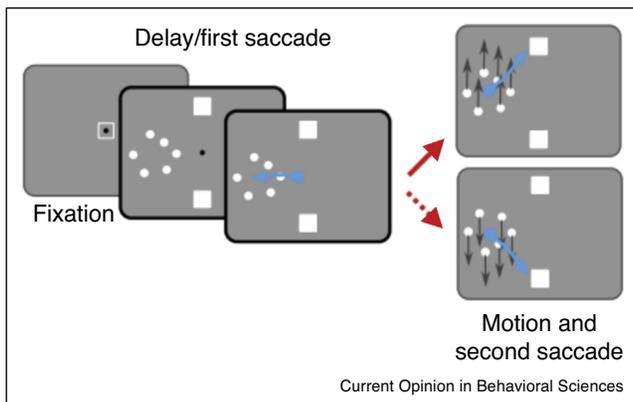
Expected Information Gains in a 2-alternative forced choice, as a function of diagnosticity (D , x-axis) and prior uncertainty (colors). Uncertainty is parametrized as decision accuracy: the maximum percent correct one can expect to obtain by guessing based on the priors and the colored traces show different levels of percent correct (numbers). EIG is defined as the accuracy of the posterior relative to the prior expectations. EIG is zero for $D < PC$ and proportional to D thereafter.

information for future reward gains. For example, a pedestrian is probably not curious about the traffic light color, but merely wants to use this information to further her goal of safely crossing the street. However, we recently showed that priority-sensitive neurons in the monkey parietal cortex – area LIP – encoded uncertainty and diagnosticity, the variables that define EIG, even when the animals exploited the information for instrumental incentives. Moreover, these variables were encoded independently of reward gains, suggesting that sensitivity to informational factors is ubiquitous even in tasks that do not explicitly probe curiosity.

We recorded LIP activity in a task in which monkeys made two contingent saccades: a first saccade to obtain information from a visual cue and a second saccade to report a decision based on the information (Figure 2). We focused on the neural responses related to the planning of the first, information-sampling saccade (Figure 2, ‘Delay/1st saccade’). Importantly, the monkeys obtained the cue’s information (motion toward one of the final decision alternatives) only after this saccade, allowing us to examine how parietal neurons encoded prior uncertainty and diagnosticity to contribute to information demand.

In one experiment using this task, we manipulated the monkeys’ prior uncertainty in alternate trial blocks [36**].

Figure 2



Stages of an instrumental information sampling task. The monkeys initiate a trial by achieving central fixation and view a display containing 2 targets (white squares) and a visual cue (stationary dots). After a 500 ms delay period, the monkeys make a saccade to the cue (blue arrow). The end of the saccade triggers the onset of motion in the cue, which has 100% coherence and is directed toward one of the targets, indicating the correct final decision. The diagnosticity (% validity) of the cue can be manipulated and signaled in advance to the monkeys. Upon motion onset, the monkeys are free to make a saccade to a target and receive a reward for a correct choice.

In high uncertainty blocks (U) the two targets for the final decision had equal probability of being correct; thus, the monkeys had high prior uncertainty and could expect the cue's information to resolve this uncertainty (equivalent to the red trace and $D = 1$ in Figure 1). In low uncertainty blocks (u), in contrast, a *single* alternative was correct on all trials; the monkeys had no prior uncertainty and the cue, while reliable, had 0 EIG (Figure 1, blue trace and $D = 1$). During neural recordings, we placed the cue in the RF of an LIP cell and examined how the neurons encoded the selection of the cue for the initial, information sampling saccade.

LIP neurons showed enhanced pre-saccadic responses on U relative to u blocks — when the cue had positive rather than zero EIG (Figure 3a). The encoding of uncertainty was significant in over 40% of individual cells and across the population and was independent of the metrics (latency, velocity, accuracy) of the cue-directed saccade.

Additional analyses showed that the enhancement by uncertainty did not reflect global arousal, but assignment of credit to the informative cue. Uncertainty did not modulate the neurons' responses to the target of the final decision, indicating that it specifically affected the visual cue. Moreover, the uncertainty-related enhancement correlated with the monkeys' ability to use the cue's information after the saccade. In high uncertainty (U) blocks, the accuracy of the monkeys' final decisions increased as a function of motion viewing time, and this increase was

steeper if LIP neurons showed stronger uncertainty modulations (Figure 3b, dark versus gray traces). Thus, the neurons' uncertainty enhancement before the saccade correlated with the monkeys' efficiency in using the information after the saccade.

In a second experiment, we investigated the role of diagnosticity [37]. We used a similar task except that monkeys always had high prior uncertainty and received cues with different diagnosticity (i.e. points along the red trace with $D = 1.0, 0.8$ and 0.55 in Figure 1). Moreover, rather than receiving a single cue as in the previous study, the monkeys received two cues of different diagnosticity and were free to choose which cue to inspect, allowing us to examine the prioritization of competing informative stimuli.

Monkeys chose to view the more reliable cue on $\sim 90\%$ of the trials, showing that they were highly sensitive to diagnosticity. LIP cells encoded diagnosticity-based prioritization (Figure 4). The neurons were selective for the direction of the saccade that the monkeys made to look at a cue, but this selectivity was stronger if the competing cues had clearly distinct rather than similar diagnosticity (Figure 4a, left versus center and right panels). This resulted in a robust early signal of diagnosticity (Figure 4b, red) that was distinct from the later saccade-related response (Figure 4b, green).

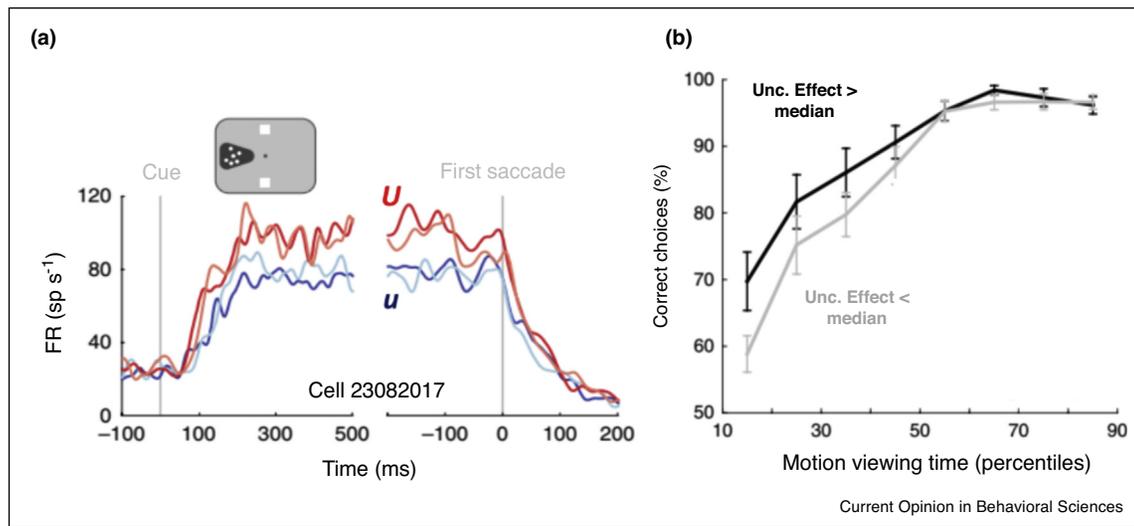
In sum, neurons encoding visual priority respond to uncertainty and diagnosticity — the two variables determining Bayesian EIG. These quantities are encoded a spatial fashion before information-sampling saccades, suggesting that they compute attentional policies that can minimize posterior uncertainty.

Priority maps do not consistently encode reward gains

The encoding of EIG in priority maps seems like a potential basis for curiosity-driven information demand — the prioritization of stimuli that allow one to *know* (enhance prediction accuracy). However, in the tasks we described, the monkeys used the information to obtain reward gains; can reward sensitivity explain our observations? Several further experiments refuted this hypothesis and showed that, although LIP neurons are sensitive to reward gains, they do not encode economic or information utility.

The Bayesian analysis described in the earlier section (Figure 1) can be used to compute the reward gains that one can expect from gaining accuracy in instrumental conditions. This quantity, known as the expected value of information (EVOI) can be defined as simply the product of EIG (the added probability of being correct) and the value at stake (the difference between the reward of a correct versus erroneous action).

Figure 3



Encoding of uncertainty by parietal oculomotor cells. **(a)** Uncertainty effect Responses of a representative LIP cell before a saccade to a cue that was expected to reduce uncertainty (U , red) or bring redundant information (u , blue). Firing rates (FR) are aligned on cue onset (left) and saccade onset (right). Pale and dark colors show, respectively, low and high reward magnitudes at each level of uncertainty. The cartoon shows the stimulus geometry, with the cue (dots) inside the RF of the cell (dark cone). **(b)** Neural sensitivity to uncertainty correlates with efficient use of the information. The accuracy for the final decision as a function of motion viewing time, split according to the strength of the neural uncertainty modulation (Unc. effect). For any viewing duration, accuracy was higher if the uncertainty effect in LIP was above relative to below the median for the sample.

To test whether LIP cells encoded EVOI, we used the task of Horan *et al.* (Figure 3) and manipulated the size of the reward that the monkeys expected for a correct response in each trial. Two reward sizes (large or small) were randomly interleaved in U and u blocks and were signaled to the monkeys through the fixation point color before the information sampling saccade. Figure 5a shows the EVOI that the monkeys experienced in the task (defined as the rewards the monkeys obtained by inspecting the cue relative to probe trials in which the cue motion was random and the monkeys merely guessed based on their priors). In u blocks, EVOI was minimal since the monkeys made correct choices based on their priors, whereas in U blocks EVOI was 0.5* reward size. This resulted in a positive interaction between reward size and prior uncertainty that clearly distinguishes EVOI from the more common metric of expected value (EV; Figure 5b, m defined as the product of reward size and probability with no reference to prior uncertainty).

Although LIP neurons were sensitive to reward size, they did not encode EVOI or EV (Figure 5c,d). Reward and uncertainty modulations were found in distinct classes of cells and combined additively rather multiplicatively in the average population response. Most strikingly, responses to reward and uncertainty had opposite polarity (Figure 5c). While uncertainty enhanced firing rates, reward had a suppressive effect with most neurons showing higher pre-saccadic responses for *lower* rewards

(Figure 5c,d). This negative reward modulation is incompatible with the theoretical prediction that reward and information interact positively, because information gains can only benefit, never reduce, reward gains (Figure 5a).

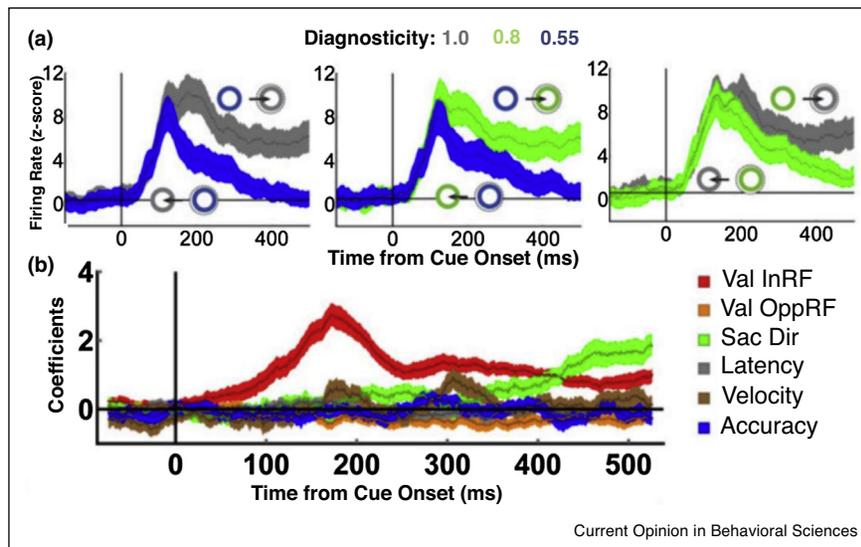
Additional evidence that the neurons did not encode subjective utility comes from the fact that they did not encode speed-accuracy tradeoffs — increases or decreases in post-saccadic viewing durations to adjust accuracy according to reward size [36**]. Moreover, a final control experiment showed that the neurons' reward sensitivity was inconsistent (the cells showed reward-related suppression and enhancement in different task contexts), refuting the idea that attentional priority equivalent to economic utility [4,38].

Hedonic effects of rewards

Before interpreting these results in a broader perspective, let us briefly mention a second puzzling set of results showing that rewards can have counterintuitive, apparently sub optimal effects on attention and information demand. Converging evidence, including from our laboratory, shows that humans and other animals are sensitive to the *affective content* of the information, and prioritize stimuli associated with positive or negative outcomes in ways that interfere with both epistemic and instrumental objectives [7,39–44].

Reward-based priority can be elicited by contextual manipulations. Monkeys and humans seek information

Figure 4



Priority responses are sensitive to diagnosticity. **(a)** Population responses Firing rates of LIP cells before saccades toward cues with different diagnosticity (colors). Diagnosticity was signaled by a colored border that could be discriminated in the visual periphery (legend) and the cues were positioned inside and opposite the RF (cartoons; RF denoted by faint gray circles). Neurons encoded the selection of the RF cue, but their directionality decreased as a function of the relative diagnosticity (left, medium and right panels). **(b)** Regression coefficients, time-resolved during the saccade decision epoch capturing the diagnosticity (%validity) of the cue in the RF (red, 'Val in RF'), after controlling for the diagnosticity of the opposite cue, saccade direction (green), and saccade velocity, endpoint accuracy, duration and latency (legend).

about a forthcoming outcome if they believe the outcome is likely to be positive, but *avoid* information (and sacrifice accuracy) if they believe it is likely to be negative (e.g. no reward or a bad medical diagnosis) [7,39,40]. Reward associations also induce humans to prioritize individual features that are not the most informative about a situation [41]. Finally, reward biases interfere with instrumental decisions, biasing the fidelity of visual information decoded from the temporal cortex [42], reducing the efficiency of instrumental information sampling [43] and distracting attention from task-relevant stimuli [44].

A study from our laboratory shows that LIP neurons encode reward-based priority [45]. The neurons had persistent activity that was biased toward the locations of cues predicting reward and away from the locations of cues predicting a lack of reward and interfered with the monkeys' ability to execute an instrumental saccade. Moreover, for highly familiar conditioned stimuli, these biases were evident from the earliest visual response, suggesting that they were associated with visual plasticity that allowed conditioned stimuli privileged bottom-up access to priority representations [45,46]. While a full discussion of reward-dependent prioritization is beyond the scope of this chapter, these results show that priority representations are recruited by multiple drives, some of which violate normative views of information utility.

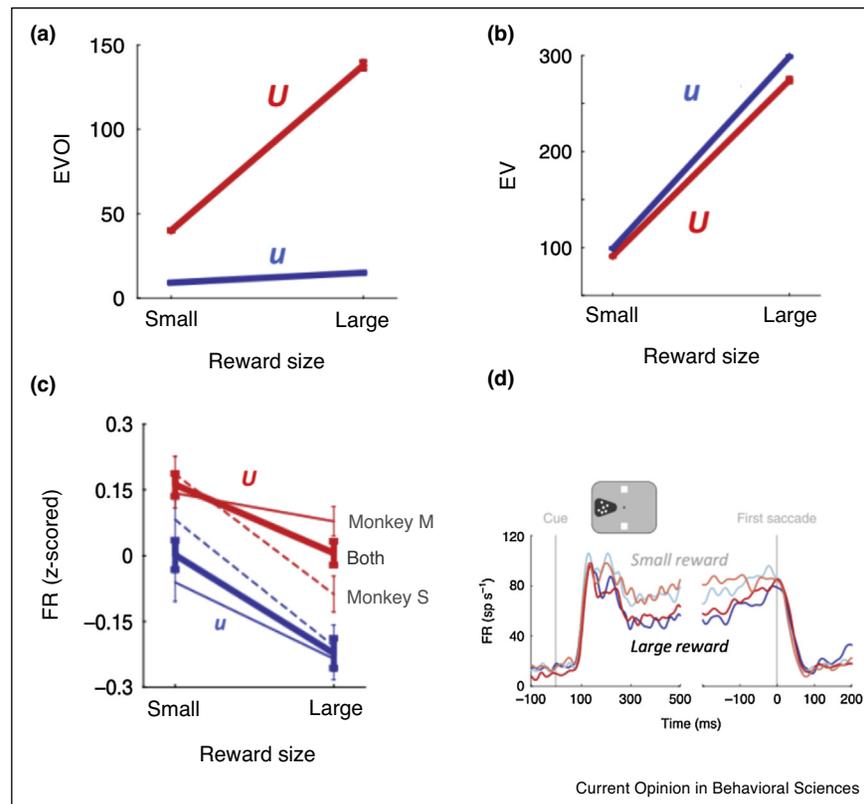
Value, cognitive state and uncertainty

Although our data are clear regarding the distinct representations of information and reward gains, this separation poses a puzzle. If rewards can specify optimal attention allocation, why do they exert apparently suboptimal effects — and why is EVOI not encoded in the priority map? The answer, we propose, requires us to think more deeply about two distinct computations: estimating the utility (cost and benefits) of an informational state versus implementing cognitive states for gathering information.

This distinction resonates with the expected value of control theory, which postulates that executive function involves distinct mechanisms for monitoring and regulation [16,47]. According to this theory, monitoring involves estimating the utility of a task and the level of control to engage in the task, and relies on the dorsal anterior cingulate cortex (dACC) and its interactions with valuation areas such as the ventromedial frontal cortex and midbrain dopamine cells [16,47]. In the context of information demand, the dACC may be critical for computing reward-dependent speed accuracy tradeoffs, which signal the subjective value of gathering information but were not encoded in LIP[48*].

Lateral fronto-parietal areas, in contrast, are proposed to play a regulatory role in *implementing* a desired cognitive policy. Mechanistic models of executive function postulate that, upon detecting a 'need for control', the dACC enhances

Figure 5



Rewards do not account for EIG responses. **(a)** EVOI The expected value of information (EVOI) that monkeys experienced in *U* and *u* blocks, defined as the expected value (EV) of trials with motion information relative to probe trials with random motion. **(b)** EV The expected value (product of reward magnitude and probability) of the same trials as in (a). EV is slightly higher in *u* blocks because of the higher accuracy in these blocks. **(c)** Average firing rates for the LIP populations were higher for *U* relative to *u* blocks (red versus blue) and lower for large relative to small reward sizes (x-axis), on average and individually in each monkey. **(d)** Enhancement by smaller rewards (with no uncertainty modulation) for a representative cell. Same format as in Figure 3a.

sensory gains perhaps acting through neuromodulators such as norepinephrine (NE) [49^{••},50]. If, in information demand, the dACC detects uncertainty and generates a boost of NE, this may explain the uncertainty-dependent enhancement we found in LIP cells (Figure 3b).

Another critical question concerns cognitive effort. While most decision research has focused on time and opportunity costs, a boost in NE is proposed to underlie mental effort — an internal *cognitive cost* — raising important questions about the neural mechanisms underlying cognitive effort [49^{••}]. Cognitive costs in information gathering are likely to arise from multiple factors including fatigue, computational complexity, and perceptual or memory load. Dissecting their underlying mechanisms will be essential for a full understanding of information sampling strategies.

Second, whereas dACC neurons and neuromodulators signal global uncertainty, LIP neurons are visually tuned

and select specific informative cues. This raises important questions about how uncertainty and diagnosticity are combined to assign credit to specific sources of information and generate spatially-organized sampling policies. Last but not least, the studies emphasize the role of uncertainty and its expected resolution in regulating information demand and the overlap of these mechanisms in instrumental and curiosity-based investigations.

Broader significance

Having reviewed the intricate neural mechanisms that prioritize information in the fronto-parietal network, let us end by considering the broader significance of this process for decisions in complex, multi-dimensional situations. Two particularly germane lines of research are the relevance of a representation (ROAR) framework proposed in social psychology [51,52] and models of multi-alternative economic decisions [53].

The ROAR framework of Eitam and Higgins examines how internal representations become ‘activated’ — transition from being implicit in memory to influencing immediate thoughts and actions. Consistent with studies of information demand, the ROAR framework proposes that activation depends on the motivational relevance of a representation, with the key motivational drives being *truth motivation* (establishing what’s real), *control motivation* (managing what happens) and *value motivation* (obtaining what’s desired). These motives seem to correspond loosely to the curiosity and instrumental motives we describe, suggesting that similar principles govern the prioritization of sensory and memory-derived information. Thus, an important question for future research concerns the mechanisms that prioritize information obtained — from without and within.

A second related line of research attempts to model realistic economic decisions and the deliberations that precede them (e.g. what to buy at a supermarket and where to look before making a choice; e.g. Ref. [54]). This research formalizes attention as a modulator of internal representations but, lacking a theory of attention control, adopts the unrealistic assumption that attention is randomly allocated. Studies of information demand can fill this gap by identifying factors that link attention with decision making such as hedonic value and EIG. But a crucial point, which has been implicit in our discussion, is that these high-level descriptors must be computed based on the decision makers’ knowledge of a situation. For instance, to determine the value of inspecting a package of butter or flour, one must know how to use butter and flour and their health implications. Thus, a complete understanding of information demand will require a description of the decision makers’ representations (‘internal models’) of a task and their use of these models to generate questions and obtain information [55].

Conflict of interest statement

Nothing declared.

Acknowledgements

The research described in this review was supported by a Human Frontiers Science Program (HFSP) Collaborative Research Grant to JG, a National Eye Institute RO1 (JG), and National Institute of Mental Health RO1 (JG).

References and recommended reading

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

- of special interest
 - of outstanding interest
1. Berlyne DE: **A theory of human curiosity.** *Br J Psychol Gen Sect* 1954, **45**:180-191.
 2. Johnson E: **Process models deserve process data: comment on Brandstatter, Gigerenzer, and Hertwig (2006).** *Psychol Rev* 2008, **115**:263-272.
 3. Coenen A, Nelson JD, Gureckis TM: **Asking the right questions about the psychology of human inquiry: nine open challenges.** *Psychon Bull Rev* 2018:1548-1587.
 4. Gottlieb J, Oudeyer PY: **Toward a neuroscience of active sampling and curiosity.** *Nat Rev Neurosci* 2018, **19**:758-770.
 5. Gottlieb J, Oudeyer PY, Lopes M, Baranes A: **Information seeking, curiosity and attention: computational and empirical mechanisms.** *Trends Cognit Sci* 2013, **17**:585-593.
 6. Zentall TR, Stagner JP: **Do pigeons prefer information in the absence of differential reinforcement?** *Learn Behav* 2012, **40**:465-475.
 7. Charpentier CJ, Bromberg-Martin ES, Sharot Tali: **Valuation of knowledge and ignorance in mesolimbic reward circuitry.** *Proc Natl Acad Sci U S A* 2018, **115**:E7255-E7264.
 8. Blanchard TC, Hayden BY, Bromberg-Martin ES: **Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity.** *Neuron* 2015, **85**:602-614.
 9. Bromberg-Martin ES, Hikosaka O: **Midbrain dopamine neurons signal preference for advance information about upcoming rewards.** *Neuron* 2009, **63**:119-126.
 10. Igaya K, Story GW, Kurth-Nelson Z, Dolan RJ, Dayan P: **The modulation of savouring by prediction error and its effects on choice.** *eLife* 2016, **21** pii: e13747.
 11. Kang MJ, Hsu M, Krajbich IM, Loewenstein G, McClure SM, Wang JT, Camerer CF: **The wick in the candle of learning: epistemic curiosity activates reward circuitry and enhances memory.** *Psychol Sci* 2009, **20**:963-973.
 12. Gruber MJ, Gelman BD, Ranganath C: **States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit.** *Neuron* 2014, **84**:486-496.
 13. Baranes AF, Oudeyer PY, Gottlieb J: **Eye movements encode epistemic curiosity in human observers.** *Vis Res* 2015, **117**:81-90.
 14. Katsuki F, Constantinidis C: **Early involvement of prefrontal cortex in visual bottom-up attention.** *Nat Neurosci* 2012, **15**:1160-1166.
 15. Suzuki M, Gottlieb J: **Distinct neural mechanisms of distractor suppression in the frontal and parietal lobe.** *Nat Neurosci* 2013, **16**:98-104.
 16. Shenhav A, Botvinick M, Cohen J: **The expected value of control: an integrative theory of anterior cingulate cortex function.** *Neuron* 2013, **79**:217-240.
 17. Gottlieb J, Snyder LH: **Spatial and cognitive functions of the parietal cortex.** *Curr Opin Neurobiol* 2010, **20**:731-740.
 18. Thompson KG, Bichot NP: **A visual salience map in the primate frontal eye field.** *Prog Brain Res* 2005, **147**:251-262.
 19. Bisley JW, Goldberg ME: **Attention, intention, and priority in the parietal lobe.** *Annu Rev Neurosci* 2010, **33**:1-21.
 20. Aw H, Belopolsky AV, Theeuwes J: **Top-down versus bottom-up attentional control: a failed theoretical dichotomy.** *Trends Cognit Sci* 2012, **16**:437-443.
 21. Tatler BW, Hayhoe MN, Land MF, Ballard DH: **Eye guidance in natural vision: reinterpreting salience.** *J Vis* 2011, **11**:5-25.
 22. Foley NC, Jangraw DC, Peck C, Gottlieb J: **Novelty enhances visual salience independently of reward in the parietal lobe.** *J Neurosci* 2014, **34**:7947-7957.
 23. Yu AJ, Dayan P: **Uncertainty, neuromodulation, and attention.** *Neuron* 2005, **46**:681-692.
 24. Kakade S, Dayan P: **Dopamine: generalization and bonuses.** *Neural Netw* 2002, **15**:549-559.
 25. Leong Y, Radulescu A, Daniel R, DeWoskin V, Niv Y: **Dynamic interaction between reinforcement learning and attention in multidimensional environments.** *Neuron* 2017, **93**:451-463.
 26. Benjamin D: **Errors in probabilistic reasoning.** *Handbook of Behavioral Economics.* Elsevier; 2019:69-186.

27. Nelson J, McKenzie C, Cottrell G, Sejnowski T: **Experience matters: information acquisition optimizes probability gain.** *Psychol Sci* 2010, **21**:960-969.
28. Dayan P, Kakade S, Montague PR: **Learning and selective attention.** *Nat Neurosci* 2000, **3(Suppl)**:1218-1223.
29. Pezzulo G, Rigoli F, Friston KJ: **Hierarchical active inference: a theory of motivated control.** *Trends Cogn Sci* 2018, **22**:294-306.
30. Yang SC, Lengyel M, Wolpert DM: **Active sensing in the categorization of visual patterns.** *eLife* 2016, **5**:e12215 <http://dx.doi.org/10.7554/eLife.12215>.
Clear demonstration of how Bayesian theory can describe viewing behavior in humans who sample visual information for a categorization decision.
31. Najemnik J, Geisler WS: **Eye movement statistics in humans are consistent with an optimal search strategy.** *J Vis* 2008, **8** 4 1-14.
32. Renninger LW, Verghese P, Coughlan J: **Where to look next? Eye movements reduce locatn uncertainty.** *J Vis* 2007, **7**:6.
33. Braver TS, Krug MK, Chiew KS, Kool W, Westbrook JA, Clement NJ, Adcock RA, Barch DM, Botvinick MM, Carver CS et al.: **Mechanisms of motivation-cognition interaction: challenges and opportunities.** *Cogn Affect Behav Neurosci* 2014, **14**:443-472.
34. Engemann JB, Damaraju E, Padmala S, Pessoa L: **Combined effects of attention and motivation on visual task performance: transient and sustained motivational effects.** *Front Hum Neurosci* 2009, **3**:4.
35. Sugrue LP, Corrado GS, Newsome WT: **Choosing the greater of two goods: neural currencies for valuation and decision making.** *Nat Rev Neurosci* 2005, **6**:363-375.
36. Horan M, Daddaoua N, Gottlieb J: **Parietal neurons encode information sampling based on decision uncertainty.** *Nat Neurosci* 2019, **22**:1327-1335.
The authors show that parietal oculomotor neurons are sensitive to prior decision uncertainty and assign credit to cues that resolve that uncertainty, and that these responses have surprising interactions with sensitivity to reward gains.
37. Foley NC, Kelley SP, Mhatre H, Lopes M, Gottlieb J: **Parietal neurons encode expected gains in instrumental information.** *Proc Natl Acad Sci U S A* 2017, **114**:E3315-E3323.
The authors show that parietal oculomotor neurons are sensitive to the diagnosticity of competing visual cues independently of the reward gains of these cues, providing a first demonstration that the neurons encode quantities relevant for Bayesian expected information gains.
38. Maunsell JH: **Neuronal representations of cognitive state: reward or attention?** *Trends Cogn Sci* 2004, **8**:261-265.
39. Daddaoua N, Lopes M, Gottlieb J: **Intrinsically motivated oculomotor exploration guided by uncertainty reduction and conditioned reinforcement in non-human primates.** *Sci Rep* 2016, **6**.
40. Loewenstein G, Molnar A: **The renaissance of belief-based utility in economics.** *Nat Hum Behav* 2018, **2**:166-167.
41. Kobayashi K, Ravaioli S, Baranès A, Woodford M, Gottlieb J: **Diverse motives for human curiosity.** *Nat Hum Behav* 2019, **3**:587-595.
42. Barbaro L, Peelen MV, Hickey C: **Valence, not utility, underlies reward-driven prioritization in human vision.** *J Neurosci* 2017, **37**:10438-10450.
43. Hunt LT, Rutledge RB, Malalasekera WM, Kennerley SW, Dolan RJ: **Approach-induced biases in human information sampling.** *PLoS Biol* 2016, **14**:e2000638.
44. Anderson B: **The attention habit: how reward learning shapes attentional selection.** *Ann N Y Acad Sci* 2016, **1369**:24-39.
45. Peck CJ, Jangraw DC, Suzuki M, Efem R, Gottlieb J: **Reward modulates attention independently of action value in posterior parietal cortex.** *J Neurosci* 2009, **29**:11182-11191.
46. Roelfsema PR, van Ooyen A, Watanabe T: **Perceptual learning rules based on reinforcers and attention.** *Trends Cogn Sci* 2010, **14**:64-71.
47. Shenhav A, Musslick S, Lieder F, Kool W, Griffiths TL, Cohen JD, Botvinick MM: **Toward a rational and mechanistic account of mental effort.** *Annu Rev Neurosci* 2017, **40**:99-124.
48. White JK, Bromberg-Martin ES, Heilbronner SR, Zhang K, Pai J, Haber SN, Monosov IE: **A neural network for information seeking.** *Nat Commun* 2019, **10**:5168.
Using a suite of information seeking and probabilistic reward tasks in monkeys, the authors detect responses related to uncertainty resolution in several basal ganglia and cortical areas including the dACC.
49. Silvetti M, Vassena E, Abrahamse E, Verguts T: **Dorsal anterior cingulate-brainstem ensemble as a reinforcement meta-learner.** *PLoS Comput Biol* 2018, **14**:e1006370.
The authors present a concrete computational model of how the dACC may optimize internal actions—such as attention and memory—to resolve cognitive conflict or uncertainty, and how it interacts with sensory and other fronto-parietal areas through the release of neuromodulators.
50. Eldar E, Cohen JD, Niv Y: **The effects of neural gain on attention and learning.** *Nat Neurosci* 2013, **16**:1146-1153.
51. Eitam B, Higgins T: **From reaction ('priming') to motivated selection: changing conceptualizations of accessibility.** *Curr Opin Psychol* 2016, **12**:58-62.
52. Eitam B, Higgins T: **Motivation in mental accessibility: relevance of a representation (roar) as a new framework.** *Soc Personal Psychol Comp* 2010, **4**:951-957.
53. Krajbich I, Rangel A: **Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions.** *Proc Natl Acad Sci U S A* 2011, **108**:13852-13857.
54. Krajbich I, Lu D, Camerer C, Rangel A: **The attentional drift-diffusion model extends to simple purchasing decisions.** *Front Psychol* 2012, **3**:193.
55. Wilson RC, Niv Y: **Inferring relevance in a changing world.** *Front Hum Neurosci* 2011, **5**:189.