

Attention, Reward, and Information Seeking

 Jacqueline Gottlieb,^{1,2} Mary Hayhoe,³ Okihide Hikosaka,⁶ and Antonio Rangel^{4,5}

¹Department of Neuroscience and ²Kavli Institute for Brain Science, Columbia University, New York, New York 10032, ³University of Texas at Austin, Austin, Texas 78712, ⁴Division of Humanities and Social Sciences and ⁵Computational and Neural Systems, California Institute of Technology, Pasadena, California 91125, and ⁶Laboratory of Sensorimotor Research, National Eye Institute, National Institutes of Health, Bethesda, Maryland 20892

Decision making is thought to be guided by the values of alternative options and involve the accumulation of evidence to an internal bound. However, in natural behavior, evidence accumulation is an active process whereby subjects decide when and which sensory stimulus to sample. These sampling decisions are naturally served by attention and rapid eye movements (saccades), but little is known about how saccades are controlled to guide future actions. Here we review evidence that was discussed at a recent symposium, which suggests that information selection involves basal ganglia and cortical mechanisms and that, across different contexts, it is guided by two central factors: the gains in reward and gains in information (uncertainty reduction) associated with sensory cues.

Introduction

Over the past decades, many studies have probed the single-cell basis of simple decisions, taking as a model system the oculomotor (eye movement) system of nonhuman primates. In these investigations, monkeys are trained to choose among several options based on a simple rule, and report their decision by making a rapid eye movement (saccade). In general, these studies support the idea that decisions are guided by the values of alternative options and involve the accumulation of evidence to an internal bound (Gold and Shadlen, 2007; Kable and Glimcher, 2009).

The progress made by these investigations sets the stage for addressing more complex questions about eye movement decisions in natural behavior. Unlike in laboratory paradigms where the subject's goal is simply to make a saccade, in natural behavior eye movements are tightly coordinated with subsequent actions and serve a specific purpose: to accumulate sensory evidence relevant to those actions (Tatler et al., 2011). Thus, evidence accumulation is not a passive process as it is currently portrayed. Instead, it requires active decisions as to when and which source of information to sample.

Therefore, a central open question is how are saccades guided so as to sample information? Although the answers to this question are just beginning to be investigated, the studies presented at a recent symposium suggest that this guidance rests on three factors. One factor is the direct association between stimuli and rewards. These associations are learned and retained on long time scales, involve basal ganglia and cortical structures, and can automatically bias attention even when they are irrelevant to a current action. A second factor is the operant value of an action that is, what a subject expects to gain from a sensorimotor task or an economic choice. And a final factor is the uncertainty involved in

the task, or the gains in information that a saccade is expected to bring.

The lateral intraparietal area (LIP) encodes visual salience based on novelty and reward

In monkeys, saccade generation involves the transmission of sensory information from the visual system to two cortical oculomotor areas, the frontal eye field (FEF) and the LIP, which in turn send descending projections to subcortical motor structures, including the basal ganglia, superior colliculus, and brainstem motor nuclei (Fig. 1A). The FEF and LIP are of special interest because they play key roles in saccadic decisions. Both areas contain populations of target-selective cells that have visual responses and spatial receptive fields (RFs) and respond very selectively to stimuli that attract attention or gaze (Thompson and Bichot, 2005; Bisley and Goldberg, 2010). These neurons seem to convey a consistent message in a variety of tasks — the location of an attention-worthy object in the visual scene (Thompson and Bichot, 2005; Bisley and Goldberg, 2010). Experiments using microstimulation or reversible inactivation suggest that the LIP and FEF exert remote effects both on motor structures, biasing these structures to generate a saccade to the selected location, and on the visual system, generating top-down attentional enhancement of the selected item (Awh et al., 2006; Squire et al., 2013).

In contrast with this extensive knowledge of their properties and functional effects, very little is known about the genesis of the LIP and FEF target selection responses. How do the neurons in these areas acquire their selective responses, and how do they “know” which stimulus to sample?

A result that has been influential in addressing this question is the finding that target selection responses in LIP cells scale monotonically with the value of a planned saccade, when value is defined according to the magnitude, probability, or delay of a juice reward delivered for making that saccade (Sugrue et al., 2005; Kable and Glimcher, 2009). This result, which has been replicated in multiple tasks, supports the idea that saccade decisions can be described in economic or reinforcement terms. However, this

Received Aug. 7, 2014; revised Oct. 8, 2014; accepted Oct. 15, 2014.

The authors declare no competing financial interests.

Correspondence should be addressed to Dr. Jacqueline Gottlieb, Department of Neuroscience, Columbia University, 1051 Riverside Drive, Kolb Research Annex, Room 569, New York, NY 10032. E-mail: jg2141@columbia.edu
DOI:10.1523/JNEUROSCI.3270-14.2014

Copyright © 2014 the authors 0270-6474/14/3415497-08\$15.00/0

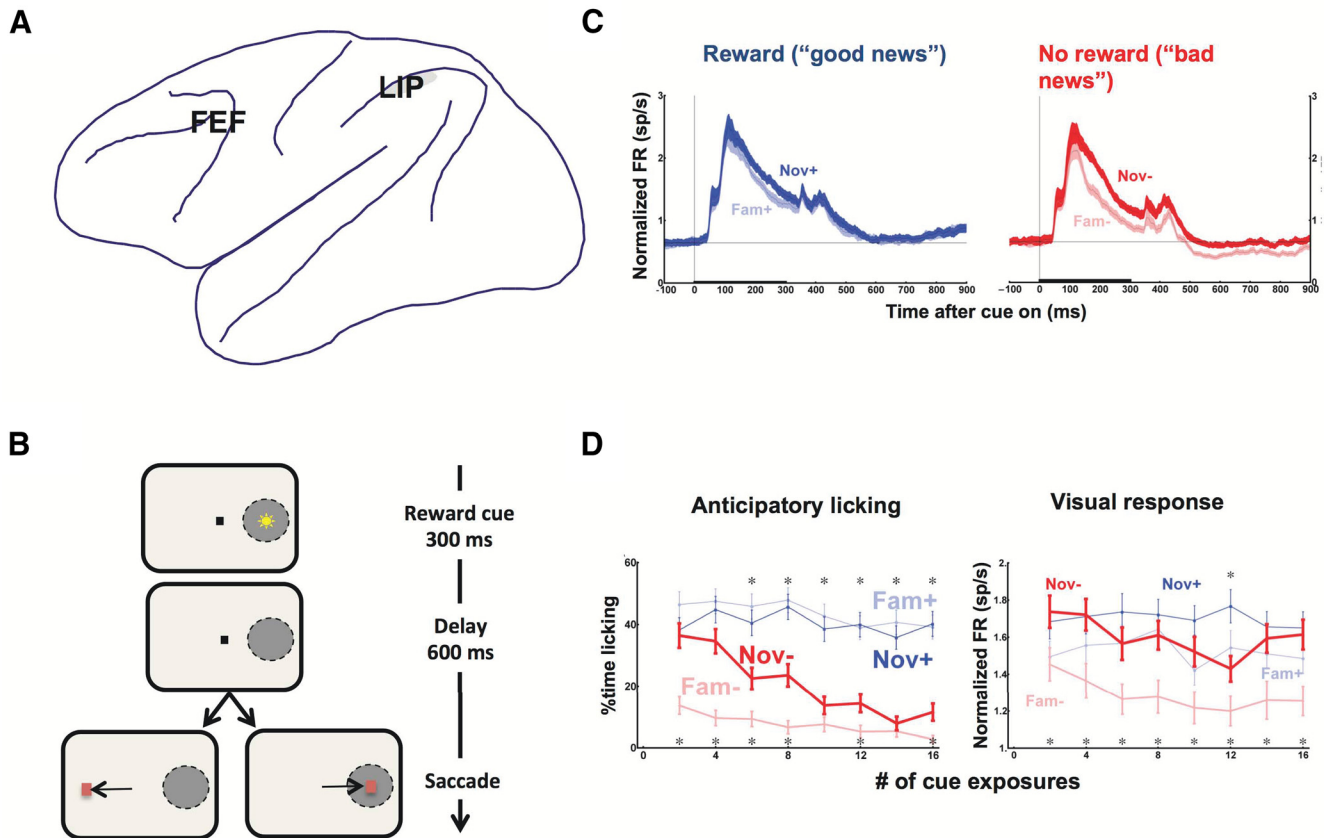


Figure 1. *A*, Cortical oculomotor areas. Lateral view of the macaque brain showing the approximate locations of the FEF and LIP. *B*, Task design. A trial began when the monkeys fixated a central fixation point (small black dot). A reward cue (yellow star) was then presented for 300 ms at a randomly selected location that could fall inside the RF of an LIP cell (gray oval) or at the opposite location (for simplicity, only the RF location is illustrated). The cue could fall into one of four categories depending on whether it was familiar (Fam) or novel (Nov) and signaled a positive (+) or a negative (−) outcome. The cue presentation was followed by a 600 ms delay period during which the monkeys had to maintain fixation (“Delay”), and then by the presentation of a saccade target at the same or opposite location relative to the cue. If the monkeys made a correct saccade to the target, they received the outcome predicted by the cue (a reward on Nov+ and Fam+ trials, but no reward on Nov− and Fam− trials). Trials with incorrect saccades were immediately repeated. *C*, LIP neurons are modulated by reward and novelty. Normalized activity (mean ± SEM) in a population of LIP cells, elicited by cues that appeared in the RF and that could be familiar or newly learned and bring “good news” (predicting a reward; Nov+ and Fam+) or bring “bad news” (predicting a lack of reward; Nov− and Fam−). The cues appeared for 300 ms (thick horizontal bar) and were followed by a 600 ms delay period during which the monkeys maintained fixation. The familiar cues showed strong reward modulations, with Fam− cues evoking a lower visual response and sustained delay period suppression that was not seen for Fam+ cues. However, newly learned cues elicited stronger overall responses and weaker reward modulations. In particular, Nov− cues did not elicit the sustained suppression seen for the Fam− cues. *D*, Learning of cue–reward associations as a function of the number of cue exposures during a session. The points show the duration of anticipatory licking and the normalized visual response (during the visual epoch, 150–300 ms after cue onset) as a function of the number of cue exposures during the session. Error bars indicate SEM. Anticipatory licking for the Nov− cues declined rapidly, but the visual response elicited by the Nov− cue remained high throughout the session. Although the monkeys rapidly learn negative cue–reward associations, they are slower to reduce the salience of a “bad news” cue. *B*, Reproduced with permission from Peck et al. (2009). *C, D*, Reproduced with permission from Foley et al. (2014).

idea has significant limitations because it introduces irresolvable confounds between reward and attention (Maunsell, 2004) and cannot explain a range of empirical observations that, in addition to their reward modulations, target selection cells respond strongly to unrewarded cues [including salient distractors (Gottlieb et al., 1998; Bisley and Goldberg, 2003; Suzuki and Gottlieb, 2013), novel items (Foley et al., 2014), and even punishment-predicting cues that are actively avoided in choice paradigms (Leathers and Olson, 2012)]. These discrepancies beg the question of how a value-based interpretation can be reconciled with the natural role of saccades in sampling information (Gottlieb, 2012).

To examine this question, we tested LIP responses in an “informational” context where visual stimuli brought information about a reward but did not support a decision based on that reward (Peck et al., 2009). As shown in Figure 1*B*, at the onset of a trial, the monkeys had uncertainty about the trial’s reward and viewed a brief visual stimulus that resolved this uncertainty, signaling whether the trial will end in a reward or a lack of reward

(i.e., brought “good” or “bad” news). After presentation of the cue inside or opposite the RF of an LIP cell, the monkeys were trained to maintain fixation for a brief delay and then make a saccade to a separate target that could appear either at the same or at the opposite location. Therefore, the cues did not allow the monkeys to make an active choice, but they could automatically bias attention toward or away from their visual field location. We detected these biases by comparing the postcue saccades that were spatially congruent or incongruent with the cue location.

We found that both positive and negative cues evoked a similar visual response shortly after their onset, suggesting that they both transiently attracted attention (Fig. 1*C*). However, at slightly longer delays, cues that brought “good news” produced sustained neural excitation, whereas cues that brought “bad news” produced sustained suppression in the neurons encoding their location. Consistent with these neuronal responses, saccades following a positive cue were facilitated if they were directed toward the cue-congruent location, suggesting that “good news” automatically attracted attention. By contrast, saccades

following negative cues were impaired if they were directed toward the cue-congruent location, suggesting that “bad news” automatically inhibited attention.

While these results are consistent with previously described reward modulations, they differ in that they were not based on operant rewards (the reward delivered by making a saccade) but on mere stimulus–reward associations. In other words, learnt reward associations modify salience, or the ability of a stimulus automatically to bias attention. This conclusion is supported by ample psychophysical evidence in human observers (see, e.g., Della Libera, 2009, #7519; Hickey, 2010, #8224; Anderson et al., 2011, #8168).

This conclusion however, raises an important question: if saccades are automatically biased by task-irrelevant rewards, how can they serve the goal of sampling information? To address this question, we examined how the neurons encode informational properties, such as the novelty of a visual cue. Using the same paradigm shown in Figure 2*B*, we introduced, along with the highly familiar reward cues, cues that were novel on their first presentation and were shown several times to allow the monkeys to learn their reward significance (Foley et al., 2014). The monkeys rapidly learnt when a novel cue signaled “bad news” and extinguished their anticipatory licking within the first few presentations (Fig. 1*D*, left). Strikingly however, the newly learned cues continued to produce enhanced salience and LIP responses for dozens of presentations (Fig. 1*D*, right). Therefore, salience is enhanced by novelty independently of reward associations, suggesting that it is also sensitive to the informational properties of visual cues.

Gaze orienting based on long-term value memory: basal ganglia mechanism

A challenge that animals face in natural behavior is the complexity of the visual scenes, and the need to memorize multiple potentially rewarding items. This is difficult to accomplish using only short-term memory because of its low capacity (Cowan, 2001). Recent evidence suggests that the basal ganglia implement a high-capacity mechanism that relies on long-term object-value memories and can automatically orient gaze to high-valued objects. We propose that this process is analogous to an “object skill” (Hikosaka et al., 2013) and allows animals and humans to rapidly reach good objects and obtain rewards.

The basal ganglia play an important role in decision-making by inhibiting unnecessary actions and disinhibiting necessary actions (Hikosaka et al., 2000; Gurney et al., 2001), and this mechanism can be used for gaze control. Indeed, experimental manipulations of the basal ganglia in macaque monkeys lead to deficits of gaze control (Hikosaka and Wurtz, 1985; Kori et al., 1995). Among basal ganglia nuclei, the caudate nucleus (CD) acts as a main gateway for gaze control by receiving inputs from cortical areas (including FEF and LIP) (Selemon and Goldman-Rakic, 1985; Parthasarathy et al., 1992), and relaying the signals to the superior colliculus (SC) via the substantia nigra pars reticulata (SNr) through direct and indirect pathways (Hikosaka et al., 1993, 2000) (Fig. 2*A*).

Notably, many neurons composing the CD–SNr–SC circuit in the monkey are strongly affected by expected reward (Hikosaka et al., 2006). This was shown using a saccade task in which the amount of reward varied depending on the direction of the saccade. The saccade to the target occurred more quickly with a higher speed when a larger reward was expected (Takikawa et al., 2002). Correspondingly, the visual and presaccadic responses of CD, SNr, and SC neurons were generally enhanced when a larger

reward was expected (Kawagoe et al., 1998; Sato and Hikosaka, 2002; Ikeda and Hikosaka, 2003). The behavioral and neuronal responses to a given target changed quickly (in several trials) when the associated reward amount changed blockwise (small-large-small, etc.). In short, the CD–SNr–SC circuit encodes reward values flexibly, relying on short-term memory.

However, recent studies show that neurons in a specific sector of the CD–SNr–SC circuit encode reward values stably, relying on long-term memory. Whereas previous studies were focused on the head/body of CD (CDh/CDb), the CD has a caudal-ventral extension called the “tail” (CDt), which is mostly unique to primates and receives inputs mainly from the inferotemporal cortex (IT) (Saint-Cyr et al., 1990) (Fig. 2*A*). A majority of CDt neurons show differential responses to visual objects (similar to IT neurons) (Caan et al., 1984; Brown et al., 1995; Yamamoto et al., 2012), but often with contralateral RFs (unlike IT neurons) (Yamamoto et al., 2012). Their outputs readily induce saccades (Yamamoto et al., 2012) by disinhibition through their connection to the caudal-lateral SNr (Kim et al., 2014) and then to SC (Yasuda et al., 2012).

Unexpectedly, unlike the neurons in CDh/CDb, neurons in the CDt were hardly influenced by rapidly changing reward values (Yamamoto et al., 2013). Instead, they changed their visual responses with prolonged experience across days, only when each object was associated with a high or low reward value consistently. Such stable value coding culminated at the caudal-lateral SNr where most neurons were inhibited by stably high-valued objects and excited by stably low-valued objects (Yasuda et al., 2012) (Fig. 2*B,C*). The neurons showed such categorical responses to as many objects as the monkey had experienced (so far, >300 objects), even after retention periods longer than 100 d. As expected from the robust inhibitory SNr–SC connection, all the monkeys tested (>10) acquired a preference for making saccades to high-valued objects among low-valued objects quickly and automatically, even with no immediate reward outcome (Fig. 2*D*). Importantly, this preference was disrupted by the reversible inactivation of CDt, but only for objects presented in the contralateral visual field (Kim and Hikosaka, 2013). Our findings are consistent with recent findings in humans (Anderson et al., 2011; Theeuwes, 2012; Chelazzi et al., 2013) and monkeys (Peck et al., 2009), showing that attention and gaze are captured automatically by objects with reward associations, and suggest that the CDt–SNr (caudal-lateral)–SC circuit contributes to the automatic gaze/attention capture.

Gaze control in active behavior

In natural behavior, one of the most important roles of saccades is to sample information to assist ongoing actions. For example, while navigating through natural environments, humans must gather visual information to control foot placement, avoid obstacles, and control heading (Fig. 3*A*). Despite these multiple demands on attention, subjects manage to select gaze targets to sample the relevant information at the appropriate time. How is this done, apparently so effortlessly, yet so reliably?

To address this problem, we made the simplifying assumption that complex behavior can be broken down into simpler and largely independent modules or subtasks (Hayhoe and Ballard, 2014), and gaze is allocated based on competition between the subtasks. For instance, in a walking task, one subtask involves heading toward a goal and another avoiding obstacles. The problem of choosing the right sequence of gaze targets then reduces to one of choosing which subtask should be attended at any moment (e.g., look toward the goal or look for obstacles).

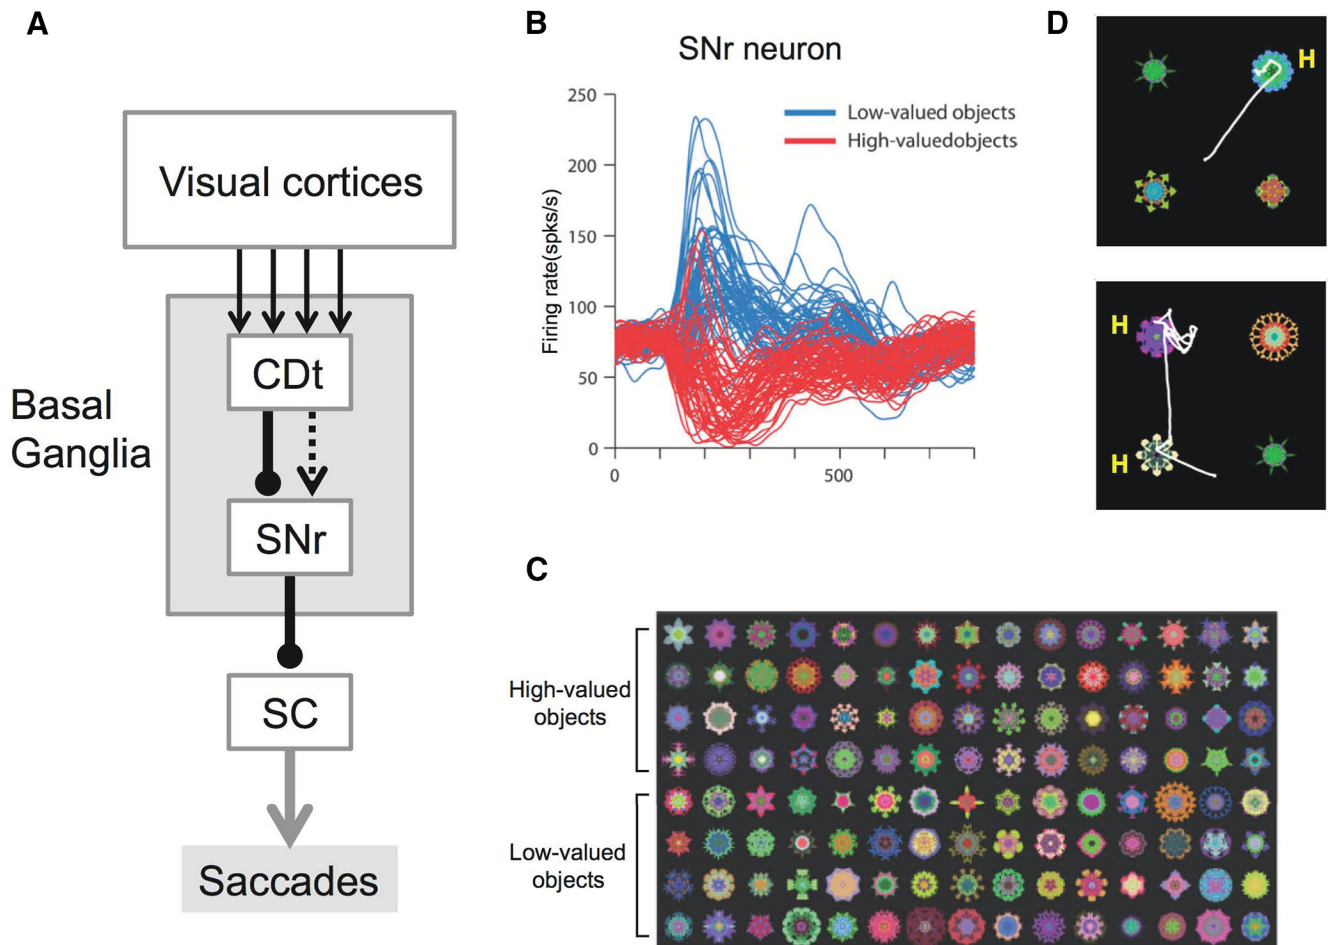


Figure 2. Basal ganglia mechanism for automatic gaze orienting to stably high-valued objects. **A**, Basal ganglia circuit controlling gaze orienting. This scheme represents the circuit mediated by the caudate tail (CDt), but equivalent circuits are present for the caudate head (CDh) or caudate body (CDb). CDt receives inputs from visual cortical areas, whereas CDh and CDb receive inputs from the frontal and parietal cortical areas. Arrows indicate excitatory connections (or effects). Lines with circular dots indicate inhibitory connections. Unbroken and broken lines indicate direct and indirect connections, respectively. **B**, An SNr neuron encoding stable values of visual objects. Shown superimposed are the neuron's responses to 60 high-valued objects (red) and 60 low-valued objects (blue). These objects (i.e., fractals) are shown in **C**. Before the recording, the monkey had experienced these objects with a large or small reward consistently for >5 daily learning sessions but had not seen them for >3 d. Firing rates (shown by spike density functions) are aligned on the onset of the object (time 0). The object disappeared at 400 ms. The neuron was located in the caudal-lateral part of SNr, which receives concentrated inputs from CDt (not CDh or CDb) and projected its axon to SC (as shown by antidromic activation). **D**, Free viewing task. On each trial, four fractal objects were presented simultaneously and the monkey freely looked at them. Examples of saccade trajectories are shown by white lines. The monkey tended to look at stably high-valued objects (denoted as H). No reward was delivered during or after free viewing. **A**, Reproduced with permission from Hikosaka et al. (2013). **B**, **C**, Reproduced from Yasuda et al. (2012).

As described in the previous sections, neural and behavioral evidence supports the idea that gaze is influenced by value considerations and uncertainty, or the need to acquire new information (Gottlieb, 2012). The role of uncertainty is less well understood but is critical in a natural task. For example, when walking on a well-known path, gaze choice will be less critical because much of the relevant information is available in memory, while in dynamic and unfamiliar environments subjects have higher uncertainty and hence greater need for information. Uncertainty arises from several sources. For example, peripheral objects may not be accurately localized or attended (Jovanovic et al., 2006), meaning that fixation is needed to specify their precise location. In turn, after fixating an obstacle, uncertainty about the location of that obstacle is likely to grow with time because of changes in the environment or memory decay, and another fixation might be required to obtain updated information. Therefore, a key role of gaze is to choose targets to reduce uncertainty about task-relevant visual information.

This principle has been captured in theoretical work by Sprague et al. (2007) who developed a reinforcement learning

model for gaze allocation that incorporates reward and uncertainty-based competition between modular subtasks. In support of this model, we gathered empirical evidence that both implicit reward (operationally defined by task goals) and uncertainty determine where and when a subject fixates in a driving task (Sullivan et al., 2012). In a more recent experiment, we investigated a walking task (Tong and Hayhoe, 2014) as illustrated in Figure 3. Subjects walked in a virtual environment seen through a head-mounted display and were instructed to stay close to the narrow path, to avoid a set of floating blue obstacles, and to intercept the brown targets. In different trials, subjects had different instructions about what subtasks to give priority to (“follow the path,” “follow the path and avoid obstacles,” or “follow the path, avoid obstacles, and intercept targets,” etc.). We also manipulated uncertainty about obstacle position by adding different amounts of random motion to the floating objects. Subjects spent more time looking at the obstacles when they were specifically instructed to avoid them (Fig. 3C, “task-relevant”), and this increased gaze allocation resulted in more effective avoidance. Adding uncertainty about obstacle position also in-

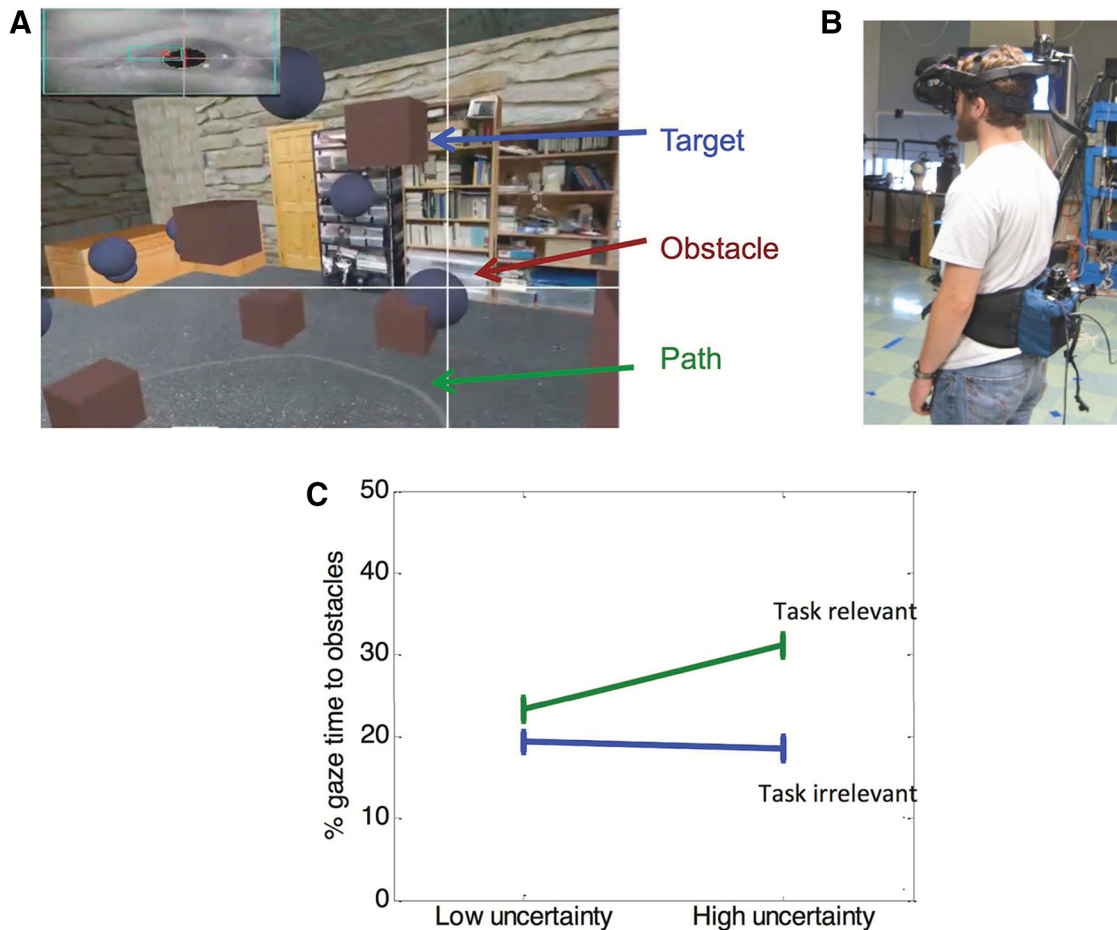


Figure 3. Virtual reality experiments for measuring gaze behavior. *A*, The view seen by a subject walking in a virtual environment with floating blue obstacles that must be avoided. The brown targets are to be intercepted, and the subject is instructed to stay close to the path indicated by the gray line. The white crosshair shows the direction of gaze. Inset, Eye image from an Arrington eye tracker mounted in the head-mounted display. *B*, View of the subject in the helmet. Both head and body position are tracked with a Hi-Ball system. *C*, Percentage of time spent fixating on the obstacles when the obstacles are stationary (low uncertainty) or moving (high uncertainty). Top points indicate the condition when subjects are explicitly instructed to avoid the randomly placed obstacles while following the path. In this case, fixations on obstacles provide task-relevant information. Bottom points indicate the condition when the subjects are instructed only to stay on the path and collect targets. In this case, the obstacles are irrelevant. Data are mean \pm SEM across 12 subjects.

creased the time spent looking at obstacles, but only when they were task relevant (Fig. 3C).

These results indicate that gaze control during complex sensorimotor tasks is sensitive both to momentary goals, or rewards, associated with the task and to the uncertainty about environmental states, consistent with the theoretical work of Sprague et al. (2007). Uncertainty is an essential factor that needs to be taken into account for a comprehensive theory of target selection, and understanding its neurophysiological underpinnings and interaction with reward is a key question for future investigations.

Visual attention during simple choice

Casual observation suggests that visual attention plays an important role in many economic decisions. Consider a shopper choosing a cereal box in the supermarket aisle. Confronted with dozens of options, his eyes move from one option to another until he is able to settle on a choice a few seconds later. This example motivates two basic questions: How do fixations affect choices? And, conversely, what determines what we fixate on while making a choice?

A series of recent papers have proposed and tested a model of how fixations affect simple choices, called the attentional drift diffusion model (aDDM) (Krajbich et al., 2010, 2011, 2012). Fig-

ure 4A summarizes the model for the case of a binary choice, where subjects make choices between pairs of snack food items (Fig. 4B). The model assumes that a choice is made by computing dynamically a stochastic relative decision value (RDV) that provides a measure of the relative attractiveness of left versus the right options. A choice is made the first time the RDV signal crosses a preestablished barrier for either the left or right options. Critically, the evolution of the RDV depends on the fixations, such that at every instant it assigns greater weight to the value of the item that is being fixated. Previous studies have shown that this model can provide a highly accurate quantitative account of the correlation between fixations and choices for the case of binary and trinary choice (Krajbich et al., 2010, 2011, 2012; Towal et al., 2013), that the size of the attentional bias is sizable, and, surprisingly, that the effect of attention is causal (Armel and Rangel, 2008; Milosavljevic et al., 2012). Human fMRI studies have shown that value signals computed in ventromedial prefrontal cortex at the time of choice, and widely thought to drive choices, are attentionally modulated in a way that could support the implementation of the aDDM (Lim et al., 2011).

The fact that fixations seem to have a sizable impact on choices motivates a deeper understanding of what determines the pattern of fixations during simple choices. The aDDM assumes that fix-

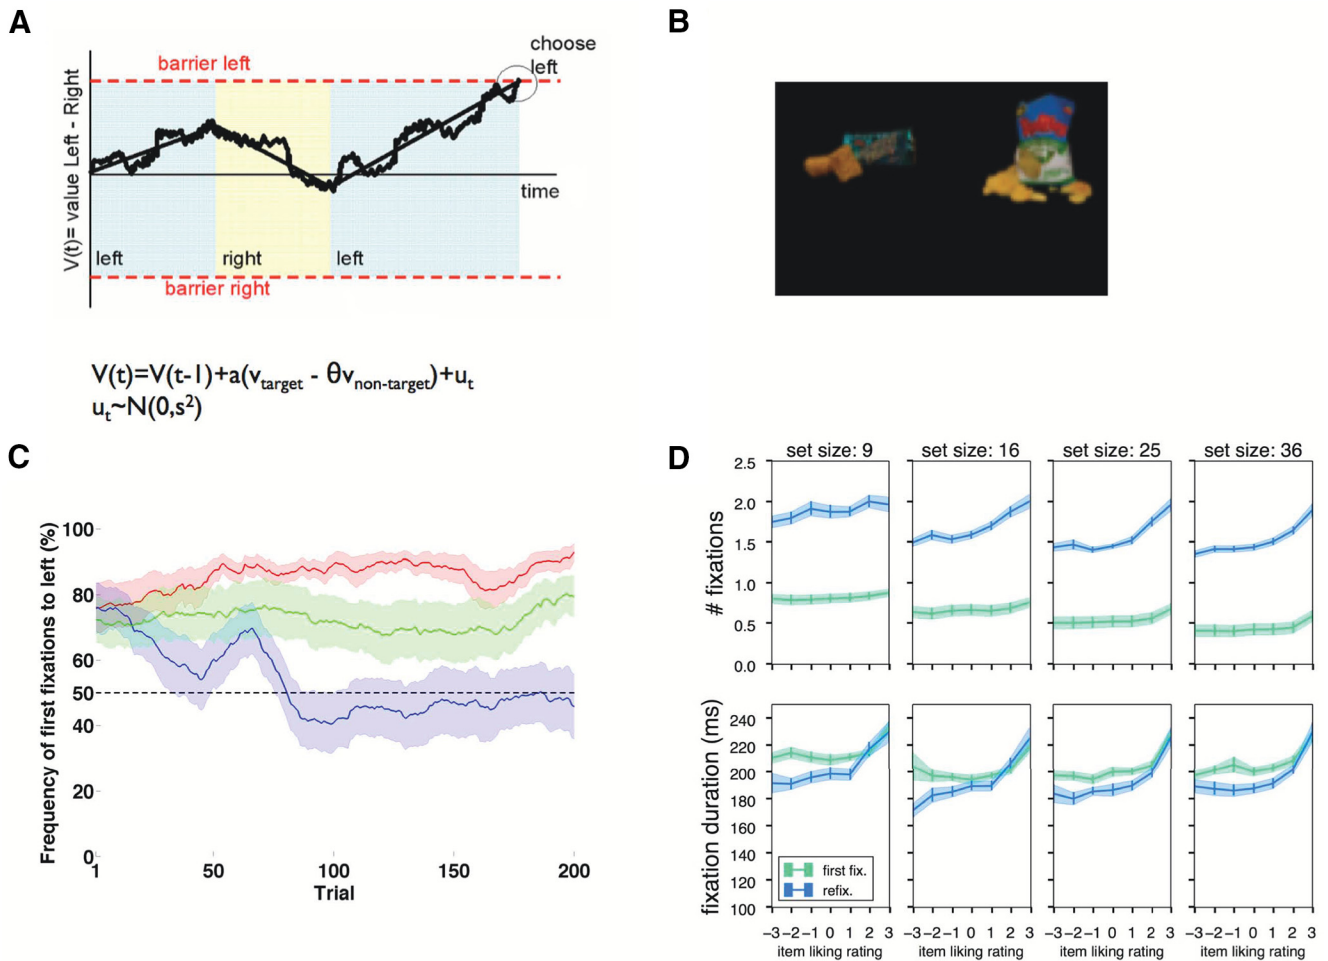


Figure 4. *A*, Summary of the aDDM for binary choices. *V* denotes the RDV signal that determines the choice. Left/right refer to the location of the fixations. *B*, Sample binary choice used in experiments. *C*, Probability that first fixation is to the left item in a modified binary choice task, as a function of condition and trial number. In green condition, best item is equally likely to be shown on left or right. In red condition, best item is shown on left in 70% of trials. In blue condition, best item is shown on right in 70% of trials. The identity of the best item changes over trials. *D*, Number of fixations and mean fixation duration as a function of choice set size and the value of the item. The green curve represents initial fixations. The blue curve represents refixations.

ation locations and durations are orthogonal to the value of the stimuli and are not affected by the value of the items or the state of the choice process. This is important because it assumes that fixations affect choices but that neither the choice process nor the value of the items affects fixations. This assumption is important because it implies that any feature of the environment that affects fixations can affect choices, even if it is unrelated to the underlying desirability of the stimuli. This stark assumption seems to be largely consistent with the data. Importantly, however, the model allows for a systematic influence of low-level visual features (e.g., relative brightness) on fixations, and through them on choices, which has also been observed in experiments (Towal et al., 2013).

These results raise two important questions. First, can visual attention be deployed in more sophisticated ways in more naturalistic choice settings? In particular, a natural interpretation of the experiments described above is that for simple choice tasks with relative short reaction times (1–2 s), visual attention must be guided by low-level features and does not have access to the state of the choice process. An important feature of the previous tests of the aDDM is that they were designed so that there was no correlation between low-level features and any information about values. But in many naturalistic settings, such a correlation might be present. For example, it might be the case that your local convenience store always places your favorite items on the left, in

which case your visual system could learn to place disproportionately its attention toward the left to improve choices. We have carried preliminary tests of this possibility and have found strong support for it, as summarized in Figure 4*C*. Another important feature of these experiments is that the stimuli are relatively novel for the subjects, but in many choice situations subjects might have developed strong stimulus–reward associations. The studies described in the previous sections (Peck et al., 2009; Yasuda et al., 2013) strongly suggest that these associations might bias attention to the best previously experienced products, which would improve the performance of the aDDM.

Second, is visual attention influenced by the value of choice options, and by the choice process, in more complex situations? To see why this is important, consider again a shopper facing a large number of options, and taking much longer to make the decision. It is natural to hypothesize that the fixation pattern in this task will change over the course of the decision, with random fixations early on, but a top-down narrowing of attention to a subset of the best options later on. Ongoing tests of this hypothesis suggest that this is indeed the case. For example, as summarized in Figure 4*D*, we found that fixation location and duration were more sensitive to stimulus value during refixations (which take place later in the trial) than during fixations (which take place earlier).

Conclusions

The studies we reviewed in the previous sections highlight the fact that saccades are specialized for sampling information or, in other words, accumulating evidence for ongoing actions. The fact that saccades do so in a goal-directed fashion implies that the brain has mechanisms for assigning priority to sources of information based on the demands of the task, and these mechanisms are only beginning to be investigated.

The first set of studies we discussed were conducted in monkeys, and suggest that one mechanism of saccade target selection is based on prior learning of stimulus–reward (Pavlovian) associations. Saccade-related responses in LIP and the basal ganglia are enhanced for reward-associated visual cues. This learning is slow and enduring and importantly, modifies salience automatically even when visual stimuli are irrelevant for actions, consistent with the phenomenon of value-driven attention capture reported in human observers (Della Libera and Chelazzi, 2009; Hickey et al., 2010; Anderson et al., 2011).

The second set of studies we reviewed were conducted in human observers, and examine saccades in active paradigms, including naturalistic sensorimotor behaviors and economic decisions. Under these conditions, saccades have a bidirectional relationship with the ongoing task. On one hand, saccades are influenced by the nature and values of the ongoing actions. On the other hand, saccades influence the task by selecting the sensory information that most strongly impacts the observer's actions.

A third and final factor, which is least understood, is the uncertainty or informational demand of the task. Analysis of gaze patterns in sensorimotor behaviors shows that gaze is strongly sensitive to the relative uncertainty of competing subtasks, and a central question for future investigations is how uncertainty and reward-based control interact. In some circumstances, these factors may be closely aligned. For instance, when subjects are performing a task, reducing uncertainty by definition increases the future expected rewards. However, in other cases they can be distinct. For instance, the value-based attentional capture described in humans and monkeys may be useful for rapidly detecting potential rewards in complex natural scenes, but it can also introduce suboptimal biases by prioritizing desirable over accurate information. Such “optimistic” biases are found in learning paradigms and were implicated in risky behaviors due to underestimation of unpleasant information (Sharot, 2011). Thus, parsing out the relative roles of reward and information gains, and the neural mechanisms by which both reach the oculomotor system, is a central question for understanding how the brain achieves top-down control of attention and gaze.

References

- Anderson BA, Laurent PA, Yantis S (2011) Value-driven attentional capture. *Proc Natl Acad Sci U S A* 108:10367–10371. [CrossRef Medline](#)
- Armell C, Rangel A (2008) Biasing simple choices by manipulating relative visual attention. *Judgm Decis Mak* 3:396–403.
- Awh E, Armstrong KM, Moore T (2006) Visual and oculomotor selection: links, causes and implications for spatial attention. *Trends Cogn Sci* 10:124–130. [CrossRef Medline](#)
- Bisley JW, Goldberg ME (2003) Neuronal activity in the lateral intraparietal area and spatial attention. *Science* 299:81–86. [CrossRef Medline](#)
- Bisley JW, Goldberg ME (2010) Attention, intention, and priority in the parietal lobe. *Annu Rev Neurosci* 33:1–21. [CrossRef Medline](#)
- Brown VJ, Desimone R, Mishkin M (1995) Responses of cells in the tail of the caudate nucleus during visual discrimination learning. *J Neurophysiol* 74:1083–1094. [Medline](#)
- Caan W, Perrett DI, Rolls ET (1984) Responses of striatal neurons in the behaving monkey: 2. Visual processing in the caudal neostriatum. *Brain Res* 290:53–65. [CrossRef Medline](#)
- Chelazzi L, Perlatto A, Santandrea E, Della Libera C (2013) Rewards teach visual selective attention. *Vision Res* 85:58–72. [CrossRef Medline](#)
- Cowan N (2001) The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behav Brain Sci* 24:87–114, discussion 114–185. [Medline](#)
- Della Libera C, Chelazzi L (2009) Learning to attend and to ignore is a matter of gains and losses. *Psychol Sci* 20:778–784. [CrossRef Medline](#)
- Foley NC, Jangraw DC, Peck C, Gottlieb J (2014) Novelty enhances visual salience independently of reward in the parietal lobe. *J Neurosci* 34:7947–7957. [CrossRef Medline](#)
- Gold JL, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574. [CrossRef Medline](#)
- Gottlieb J (2012) Attention, learning, and the value of information. *Neuron* 76:281–295. [CrossRef Medline](#)
- Gottlieb JP, Kusunoki M, Goldberg ME (1998) The representation of visual salience in monkey parietal cortex. *Nature* 391:481–484. [CrossRef Medline](#)
- Gurney K, Prescott TJ, Redgrave P (2001) A computational model of action selection in the basal ganglia: I. A new functional anatomy. *Biol Cybern* 84:401–410. [CrossRef Medline](#)
- Hayhoe M, Ballard D (2014) Modeling task control of eye movements. *Curr Biol* 24:622–628. [CrossRef Medline](#)
- Hickey C, Chelazzi L, Theeuwes J (2010) Reward guides vision when it's your thing: trait reward-seeking in reward-mediated visual priming. *PLoS One* 5:e14087. [CrossRef Medline](#)
- Hikosaka O, Wurtz RH (1985) Modification of saccadic eye movements by GABA-related substances: II. Effects of muscimol in monkey substantia nigra pars reticulata. *J Neurophysiol* 53:292–308. [Medline](#)
- Hikosaka O, Sakamoto M, Miyashita N (1993) Effects of caudate nucleus stimulation on substantia nigra cell activity in monkey. *Exp Brain Res* 95:457–472. [Medline](#)
- Hikosaka O, Takikawa Y, Kawagoe R (2000) Role of the basal ganglia in the control of purposive saccadic eye movements. *Physiol Rev* 80:953–978. [Medline](#)
- Hikosaka O, Nakamura K, Nakahara H (2006) Basal ganglia orient eyes to reward. *J Neurophysiol* 95:567–584. [CrossRef Medline](#)
- Hikosaka O, Yamamoto S, Yasuda M, Kim HF (2013) Why skill matters. *Trends Cogn Sci* 17:434–441. [CrossRef Medline](#)
- Ikeda T, Hikosaka O (2003) Reward-dependent gain and bias of visual responses in primate superior colliculus. *Neuron* 39:693–700. [CrossRef Medline](#)
- Jovancevic J, Sullivan B, Hayhoe M (2006) Control of attention and gaze in complex environments. *J Vis* 6:1431–1450. [CrossRef Medline](#)
- Kable JW, Glimcher PW (2009) The neurobiology of decision: consensus and controversy. *Neuron* 63:733–745. [CrossRef Medline](#)
- Kawagoe R, Takikawa Y, Hikosaka O (1998) Expectation of reward modulates cognitive signals in the basal ganglia. *Nat Neurosci* 1:411–416. [CrossRef Medline](#)
- Kim HF, Hikosaka O (2013) Distinct basal ganglia circuits controlling behaviors guided by flexible and stable values. *Neuron* 79:1001–1010. [CrossRef Medline](#)
- Kim HF, Ghazizadeh A, Hikosaka O (2014) Separate groups of dopamine neurons innervate caudate head and tail encoding flexible and stable value memories. *Front Neuroanat* 8:120. [CrossRef](#)
- Kori A, Miyashita N, Kato M, Hikosaka O, Usui S, Matsumura M (1995) Eye movements in monkeys with local dopamine depletion in the caudate nucleus: II. Deficits in voluntary saccades. *J Neurosci* 15:928–941. [Medline](#)
- Krajbich I, Armell C, Rangel A (2010) Visual fixations and the computation and comparison of value in simple choice. *Nat Neurosci* 13:1292–1298. [CrossRef Medline](#)
- Krajbich I, Rangel A (2011) Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proc Natl Acad Sci U S A* 108:13852–13857. [CrossRef Medline](#)
- Krajbich I, Lu D, Camerer C, Rangel A (2012) The attentional drift-diffusion model extends to simple purchasing decisions. *Front Cogn Sci* 3:193. [CrossRef Medline](#)
- Leathers ML, Olson CR (2012) In monkeys making value-based decisions, LIP neurons encode cue salience and not action value. *Science* 338:132–135. [CrossRef Medline](#)
- Lim SL, O'Doherty JP, Rangel A (2011) The decision value computations in

- the vmPFC and striatum use a relative value code that is guided by visual attention. *J Neurosci* 31:13214–13223. [CrossRef Medline](#)
- Maunsell JH (2004) Neuronal representations of cognitive state: reward or attention? *Trends Cogn Sci* 8:261–265. [CrossRef Medline](#)
- Milosavljevic M, Navalpakkam V, Koch C, Rangel A (2012) Relative visual saliency differences induce sizable bias in consumer choice. *J Consum Psychol* 22:67–74. [CrossRef](#)
- Parthasarathy HB, Schall JD, Graybiel AM (1992) Distributed but convergent ordering of corticostriatal projections: analysis of the frontal eye field and the supplementary eye field in the macaque monkey. *J Neurosci* 12:4468–4488. [Medline](#)
- Peck CJ, Jangraw DC, Suzuki M, Efem R, Gottlieb J (2009) Reward modulates attention independently of action value in posterior parietal cortex. *J Neurosci* 29:11182–11191. [CrossRef Medline](#)
- Saint-Cyr JA, Ungerleider LG, Desimone R (1990) Organization of visual cortical inputs to the striatum and subsequent outputs to the pallidum-nigral complex in the monkey. *J Comp Neurol* 298:129–156. [CrossRef Medline](#)
- Sato M, Hikosaka O (2002) Role of primate substantia nigra pars reticulata in reward-oriented saccadic eye movement. *J Neurosci* 22:2363–2373. [Medline](#)
- Selemon LD, Goldman-Rakic PS (1985) Longitudinal topography and interdigitation of corticostriatal projections in the rhesus monkey. *J Neurosci* 5:776–794. [Medline](#)
- Sharot T (2011) The optimism bias. *Curr Biol* 21:R941–R945. [CrossRef Medline](#)
- Sprague N, Ballard DH, Robinson A (2007) Modeling embodied visual behaviors. *ACM Trans Appl Percept* 4:11.
- Squire R, Noudoost B, Schafer R, Moore T (2013) Prefrontal contributions to visual selective attention. *Annu Rev Neurosci* 8:451–466. [CrossRef Medline](#)
- Sugrue LP, Corrado GS, Newsome WT (2005) Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat Rev Neurosci* 6:363–375. [CrossRef Medline](#)
- Sullivan BT, Johnson L, Rothkopf CA, Ballard D, Hayhoe M (2012) The role of uncertainty and reward on eye movements in a virtual driving task. *J Vis* 12:13. [CrossRef Medline](#)
- Suzuki M, Gottlieb J (2013) Distinct neural mechanisms of distractor suppression in the frontal and parietal lobe. *Nat Neurosci* 16:98–104. [CrossRef Medline](#)
- Takikawa Y, Kawagoe R, Itoh H, Nakahara H, Hikosaka O (2002) Modulation of saccadic eye movements by predicted reward outcome. *Exp Brain Res* 142:284–291. [CrossRef Medline](#)
- Tatler BW, Hayhoe MM, Land MF, Ballard DH (2011) Eye guidance in natural vision: reinterpreting salience. *J Vis* 11:5. [CrossRef Medline](#)
- Theeuwes J (2012) Automatic control of visual selection. *Nebr Symp Motiv* 59:23–62. [Medline](#)
- Thompson KG, Bichot NP (2005) A visual salience map in the primate frontal eye field. *Prog Brain Res* 147:251–262. [CrossRef Medline](#)
- Tong MH, Hayhoe M (2013) The effects of task and uncertainty on gaze while walking. *J Vis* 13:514. [CrossRef](#)
- Towal RB, Mormann M, Koch C (2013) Simultaneous modeling of visual saliency and value computation improves predictions of economic choice. *Proc Natl Acad Sci U S A* 110:E3858–E3867. [CrossRef Medline](#)
- Yamamoto S, Kim HF, Hikosaka O (2013) Reward value-contingent changes of visual responses in the primate caudate tail associated with a visuomotor skill. *J Neurosci* 33:11227–11238. [CrossRef Medline](#)
- Yamamoto S, Monosov IE, Yasuda M, Hikosaka O (2012) What and where information in the caudate tail guides saccades to visual objects. *J Neurosci* 32:11005–11016. [CrossRef Medline](#)
- Yasuda M, Yamamoto S, Hikosaka O (2012) Robust representation of stable object values in the oculomotor basal ganglia. *J Neurosci* 32:16917–16932. [CrossRef Medline](#)