

MOTIVATED COGNITION: NEURAL AND COMPUTATIONAL MECHANISMS OF CURIOSITY, ATTENTION, AND INTRINSIC MOTIVATION

Jacqueline Gottlieb, Manuel Lopes and Pierre-Yves Oudeyer

ABSTRACT

Based on a synthesis of findings from psychology, neuroscience, and machine learning, we propose a unified theory of curiosity as a form of motivated cognition. Curiosity, we propose, is comprised of a family of mechanisms that range in complexity from simple heuristics based on novelty, salience, or surprise, to drives based on reward and uncertainty reduction and finally, to self-directed metacognitive processes. These mechanisms, we propose, have evolved to allow agents to discover useful regularities in the world – steering them toward niches of maximal learning progress and away from both random and highly familiar tasks. We emphasize that curiosity arises organically in conjunction with cognition and motivation, being generated by cognitive processes and in turn,

Recent Developments in Neuroscience Research on Human Motivation

Advances in Motivation and Achievement, Volume 19, 149–172

Copyright © 2017 by Emerald Group Publishing Limited

All rights of reproduction in any form reserved

ISSN: 0749-7423/doi:10.1108/S0749-742320160000019017

motivating them. We hope that this view will spur the systematic study of curiosity as an integral aspect of cognition and decision making during development and adulthood.

Keywords: Intrinsic motivation; active learning; memory; attention; metacognition; development

INTRODUCTION

Countless studies in neuroscience and psychology have probed the neural basis of cognitive functions such as attention, memory, and mental representations. While these studies have traditionally remained independent from studies of decision making and motivation, this separation is beginning to change with the advent of evidence documenting strong effects of motivation on memory (reviewed in this book). These recent results suggest a more integrative view, whereby cognition and motivation are tightly intertwined. According to this view cognition is not a passive process that is simply “given” to us by the brain, but instead an active, motivated process – a mental act which, much like our physical acts, is proactively oriented toward a goal. Moreover, cognition is a source of motivation, and specific cognitive factors related to learning, memory or attention can causally impact motivational states.

This more integrative conception raises fundamental questions about the types of motivation that drive us to *think*. What are the factors that motivate us to learn, memorize or otherwise process new information? How do these intellectual drives serve our ability to thrive in the world, how do they control our actions and what are their neural substrates?

In this chapter, we will examine these questions with a focus on curiosity – a complex cognitive process that is defined as the intrinsic desire to learn or obtain information. Curiosity reaches its pinnacle in human beings in pursuits such as scientific research, and is arguably a key factor in the considerable success of our species. However, our understanding of curiosity is in its infancy, and its computational and neuroscientific basis are only beginning to be investigated.

We will review recent developments in neuroscience, cognitive psychology, and computational modeling and machine learning that pertain to these questions. We will start by reviewing fundamental properties of

curiosity and intrinsic motivation, followed by a survey of recent evidence that curiosity recruits motivational systems (including midbrain dopaminergic neurons and dopamine-recipient structures) and systems of selective attention including parietal areas involved in oculomotor control. Finally, we will review a number of factors that contribute to curiosity, including novelty, surprise, uncertainty, rewards and metacognitive control, and our current understanding of their neural mechanisms.

Throughout the discussion we will stress the idea that curiosity involves a family of mechanisms which, while highly sophisticated in humans, have their roots in more primitive motivational and information sampling systems that are found in many animal species. In addition, we emphasize the fact that mechanisms of curiosity-driven learning can be computationally modeled, and that such models are highly useful in formulating new hypotheses about the nature and function of curiosity in learning, development and behavior.

CURIOSITY AND INTRINSIC MOTIVATION

In Mark Twain's legendary book *Tom Sawyer*, Tom, a lovable and rambunctious teen, says this about intrinsic motivation: "*Work consists of whatever a body is obliged to do, and Play consists of whatever a body is not obliged to do.*" In the more formal (though perhaps less entertaining) language of science, an intrinsically motivated activity is one that is undertaken for no apparent reason except the behavior itself (Ryan & Deci, 2000a, 2000b).

Intrinsically motivated behaviors include behaviors that subjects are not obliged to do for survival and yet are highly motivating – such as children's play and adult leisure-time hobbies and creative pursuits. Intrinsically motivated activities are generally pleasurable, and can even cause special states of "flow" characterized by intense feelings of effortless control, concentration, enjoyment, and a contraction of the sense of time (Csikszentmihalyi, 1990).

From a computational perspective, intrinsically motivated behaviors can be characterized similarly to other goal-directed behaviors— as actions that seek to maximize an internal goal – formalized mathematically as a reward (value) function. A particular challenge however, is to understand what are the value functions that the agents seek to maximize. Whereas, in the vast majority of experiments in neuroscience and psychology, behavior is

shaped using easily measurable extrinsic rewards such as money, juice, food, or points, intrinsically motivated behaviors depend on internal factors that are much more difficult to characterize and are related to the individual's affective or cognitive structure. For instance, when creating a painting, an amateur individual may be motivated by the simple pleasure of gazing at the painting, the satisfaction of achieving personal goals (e.g., personal growth, learning, mastering a challenge), the satisfaction of achieving social goals (e.g., impressing her husband) or even the prospect of eventual material gains (being unexpectedly discovered and becoming a successful painter). These constructs can all be viewed as “rewards” in the widely accepted view of the term – as any factor that reinforces behavior and “makes you come back for more” (Thorndike, 1911). However, it remains a formidable challenge to identify which of these internal rewards come into play in any given context and how these motivations are computed by the brain.

Curiosity is a particular system of intrinsic motivation that drives agents to learn. The curious agent seems to have satisfied all his material needs – for food, social contact, safety, etc. As Tom Sawyer would say therefore, the agent has no more *Work* to do and a perfectly rational action that he may choose to take is simply conserve energy and do nothing at all – wait quietly until new primary needs arise. And yet, at least in some circumstances, the agent becomes intrinsically motivated to explore – and he expends time and effort to open a closed door (answer a question), and discover new parts of his environment that were not suspected before. What motivates the agent to carry out this extraordinary behavior? Why incur significant costs in energy, effort, and time, to pursue an uncertain (indeed, unknown!) reward when one's material needs have been satisfied?

A second remarkable aspect of curiosity, beyond its mere existence, is the fact that it generates not a random but a structured pattern of investigation. The agent in our cartoon is not interested indiscriminately in all the information that surrounds him, but becomes curious about specific items. Work in machine learning and robotics clearly shows that, because many of the tasks we may choose to explore in a natural environment are *unlearnable* or impossible to master, agents cannot assign “intrinsic” value to all sources of information as is sometimes claimed in the literature. An indiscriminate strategy of examining all the available information would result in collecting disparate pieces of information with nearly no discovery of useful structures, especially given the limited time and energy available over a biological life span. Therefore, a successful curiosity mechanism must assign value to possible endeavors very selectively, and in a way that

maximizes the agent's ability to predict his environment over vast portions of the learning space and on longer time scales. Computational studies show that, in environments that change quickly and/or continuously, curious individuals can gain an advantage by acquiring new skills and discovering new environmental structures (Singh et al., 2010; Barto, 2013). However, this long-term (evolutionary) advantage cannot specify the agent's actions on shorter time scale. Herein lies the second great mystery of curiosity: how do agents achieve structured exploration in a way that maximizes the long-term advantage, given that the specific goals are unknown when the exploration unfolds?

Emerging neuroscientific evidence, to which we turn next, suggests that implementing this system requires the concerted action of dopaminergic systems implicated in value and motivation, and cortical systems mediating cognitive processes of memory and attention.

DOPAMINERGIC SYSTEMS THAT PROCESS PRIMARY REWARDS ARE ACTIVATED BY CURIOSITY

To examine the motivational systems that are recruited by curiosity, Kang et al. used functional magnetic resonance imaging (fMRI) to monitor brain activity in human observers who pondered trivia questions (Kang et al., 2009). After reading a question, the subjects rated their curiosity and confidence regarding the question and, after a brief delay, were given the answer. The key analyses focused on activations during the *anticipatory* period – after the subjects had received the question but before they were given the answer.

Areas that showed activity related to curiosity during this epoch included the left caudate nucleus, bilateral inferior frontal gyrus (IFG), and loci in the putamen and globus pallidus (Fig. 1). In an additional behavioral task, the authors showed that subjects were willing to pay a higher price to obtain the answers to questions that they were more curious about – that is, could compare money and information on a common scale. They concluded that the value of the information, experienced as a feeling of curiosity, is encoded in some of the same structures that evaluate material gains.

Two recent studies extend this result in non-human primates by reporting that midbrain dopaminergic (DA) cells and cells in the orbitofrontal cortex (OFC), a pre-frontal area that receives DA innervation, encode the anticipation of obtaining reliable information from visual cues




Fig. 1. Brain regions that showed differential activity in high- versus low-curiosity trials during the first question presentation in Kang et al. (2009). Colored areas showed greater anticipatory activation on high-curiosity trials in experiment 1 ($p < .001$ uncorrected, $\text{prep} > .99$, extent threshold 5) using a median-split analysis (red), the modulator analysis (yellow), and the analysis of residual curiosity (green). See the original paper for a color rendition. The illustration at the right is a close-up view of the overlapping caudate activations. Ant, anterior; Pos, posterior; L, left; R, right; IFG, inferior frontal gyrus. *Source:* Reproduced with permission from Kang et al. (2009).

(Blanchard, Hayden, & Bromberg-Martin, 2015; Bromberg-Martin & Hikosaka, 2009).

In the study on DA cells, the subjects were trained on the so-called “observing paradigms,” where they had to choose between observing two cues that had equal physical rewards but differed in their offers of information (Bromberg-Martin & Hikosaka, 2009). Monkeys began each trial with a 50% probability of obtaining a large or a small reward and, before receiving the reward, had to choose to observe one of two visual items. If the monkeys chose the informative cue, this cue changed to one of two patterns that reliably predicted whether the trial will yield a large or small reward. If the monkeys chose the uninformative item, this item also changed to produce one of two patterns, but the patterns had only a random relation to the reward size.

The key feature of the behavioral task was that the extrinsic rewards that the monkeys received were equal for the two options (both targets had a 50% probability of delivering a large or small reward), and therefore there was no biological imperative for the monkeys to choose either option. Nevertheless, the monkeys developed a reliable and consistent preference for choosing the informative cue. A subsequent study of area OFC showed that the monkeys will choose the informative option even if its payoff is

slightly lower than that of the uninformative option – that is, monkeys are willing to sacrifice juice reward to view predictive cues (Blanchard et al., 2015). Finally, neuronal recordings showed that DA and OFC cells anticipated the receipt of reward information, and encoded the value that the monkeys placed on this information independent of the value of obtaining the reward itself.

Together, these investigations show that, in both humans and monkeys, the motivational systems that signal the value of primary rewards also signal an intrinsic desire to obtain information, raising important (and hotly debated) questions about the computations that relate values based on the two factors.

EYE MOVEMENTS AND ATTENTION

Although DA neurons and DA-recipient structures may enhance arousal and motivation in anticipation of information gains, they are not sufficient to explain the full scope of information seeking mechanisms. In addition to the motivational signals conveyed by these cells, the brain must construct a *full exploration policy* that specifies when, for how long, and from which source to gather information. Constructing such detailed policies requires input from cognitive processes that evaluate the semantic and informational properties of competing cues and their desirability as targets for exploration.

While little is currently known regarding these links, a recent study by Gruber and colleagues (Gruber, Gelman, & Ranganath, 2014) showed that a higher level of curiosity is associated with better memory performance, and enhanced activation of parahippocampal structures and its DA projections, providing valuable first evidence for a link between curiosity and memory. In addition, another recent study in our laboratory revealed a link between curiosity and attention (Baranes, Oudeyer, & Gottlieb, 2015).

In that latter study, we combined eye tracking with a trivia task similar to that used by Kang et al., in which we presented subjects with trivia questions and monitored their eye movements while they were waiting for and reading the answer. As shown in Fig. 2a, questions that were rated as eliciting higher curiosity were associated with faster anticipatory gaze shifts to the expected location of the answer. The subjects' eye position at the beginning of a trial was initially distant from the location where the answer was expected to appear (Fig. 2a, left panel: $-1,500$ ms), but gaze gradually

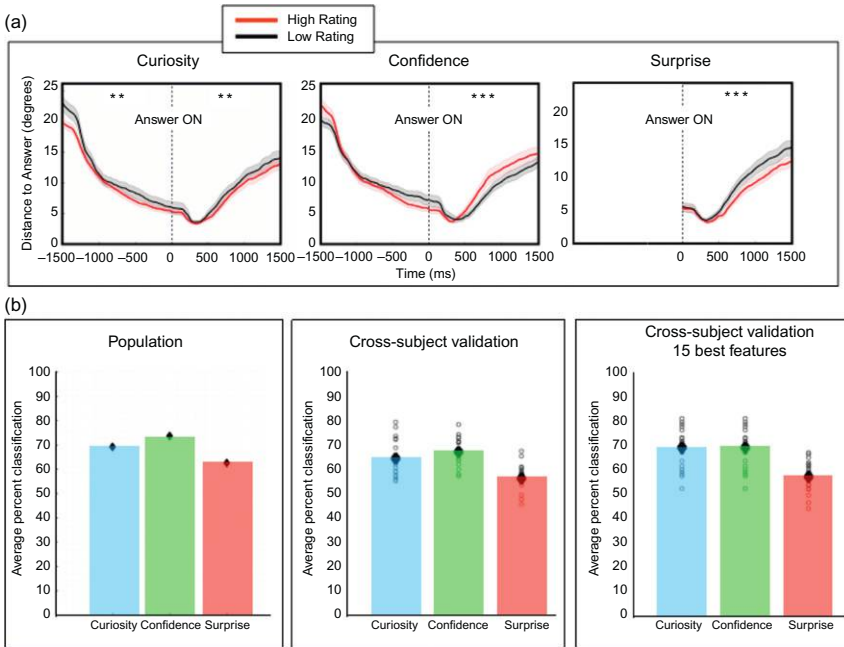


Fig. 2. Curiosity affects eye movements. (a) Eye movement patterns in relation to curiosity, confidence and surprise. For each trial with high or low ratings (black vs gray traces), we computed the distance between the eye position and the left edge of the answer box as a function of time. Distances were averaged for each subject, and we display the mean and SEM across subjects. Average distances before and after answer onset were compared with a one-way ANOVA; $**p < 10^{-45}$, $***p < 10^{-75}$. (b) Classification accuracy for different implementations of a machine learning algorithm. Left: Classification across the entire data set. Middle: Classification with cross-subject cross-validation. Right: Same as the middle panel but using only the 15 most informative features. In the middle and right panels, the open points show individual subject predictions and the black points and bars show average and SEM. *Source:* Reproduced with permission from (Baranes et al., 2015).

shifted to the answer location in an anticipatory fashion, well before the answer actually appeared (time 0 in the figure). Anticipatory eye movements were faster for questions that elicited high relative to low curiosity (black vs gray traces, left panel in Fig. 2a) and, across individuals, the magnitude of this acceleration was correlated with measures of curiosity-related personality traits (see also Risko, Anderson, Lanthier, & Kingstone, 2012).

As had also been shown by Kang et al., we found that curiosity had an inverted U-shaped relationship with the subjects' confidence that they knew the answer: subjects were most curious about questions for which they had intermediate confidence. In addition, ratings of curiosity were correlated with ratings of surprise: the more curious the subject were in a question, the more surprised they reported being when viewing the answer. Despite these correlations, the three ratings had dissociable eye movement signatures. While curiosity affected primarily anticipatory gaze, confidence and surprise primarily acted after the answer appeared (times 0–1,500 ms), when subjects lingered longer on an answer if they had reported having less confidence that they knew the answer (Fig. 2a, middle panel), and being more surprised after its presentation (Fig. 2a, right panel).

One intriguing possibility raised by these results is that, if eye movements reflect an individual's mental state, they may also be used to infer that state. To examine this possibility, we used a machine-learning algorithm with a cross-validation decoding technique (i.e., training the algorithm to match ratings with eye movement features in a subset of the data, and testing its predictions in a non-overlapping subset). As shown in Fig. 2b (left) the algorithm correctly predicted the subjects' curiosity, confidence, and surprise based on their eye movements alone with above chance accuracy of approximately 70%. Interestingly, classification accuracy was high even when the algorithm was trained and tested with data from different subjects (Fig. 2b, cross-subject validation) and even when we used a smaller set of eye movement features that were discovered to be most informative (Fig. 2b, 15 best features).

These results suggest that curiosity has powerful effects in recruiting memory and attention. While much remains to be learnt regarding these links, the possibility of using eye movements to read out epistemic states of curiosity, confidence or surprise, has exciting potential practical applications in a number of settings including educational situations.

WHAT MOTIVATES CURIOSITY?

Converging evidence suggests that some of the factors that generate curiosity include surprise, novelty, uncertainty about future rewards, and the probability of rewards for of specific items. We review each factor in turn.

Surprise

It has long been recognized that, far from being unbiased, the way in which we sample information from complex visual scenes depends strongly on our knowledge and expectations (Vo & Wolfe, 2015). In the hands of a professional magician, the manipulation of expectations can lead to spectacular misdirection and consequent surprise (Rieiro, Martinez-Conde, & Macknik, 2013). Predictive coding theories suggest that expectations play a key role in orienting attention by predicting away redundant information and freeing resources for detecting significant items (Friston et al., 2013). This systematic removal of information through active prediction may be critical for allowing us to see – and indeed, survive – as without it we may be overwhelmed by the sheer amount of information that our brains have to process.

Studies by Itti and Baldi have shown that surprise, defined in the domain of visual features, attracts human saccades during free-viewing exploration (Baldi & Itti; Itti & Baldi, 2009). Using a Bayesian algorithm combined with computational models of vision, the authors simulated the observers' beliefs about the expected distribution of pixel values at different visual locations, and defined surprise as the extent to which a visual input differed from these expectations. The authors showed that this quantitative metric could predict human free-viewing patterns with greater fidelity and flexibility relative to simpler intensity or contrast-based predictors. As the authors emphasize, surprise differs from standard measures of information in that it ascribes central importance to the observers' beliefs rather than being defined purely by the entropy (uncertainty) of a stimulus set. This makes it very clear that it is not the mere presence of information that attracts our attention, but the extent to which the information confirms or violates our prior expectations.

Novelty

Novelty, in contrast with surprise, is not context-specific but is defined by the total amount of exposure that observers had to a given observation. Novelty can be modeled mathematically as the dissimilarity between a stimulus and the representation of familiar stimuli encoded in the observer's memory (Barto, Mirolli, & Baldassare, 2013).

In a classical approach to reinforcement-learning (RL), novelty is thought to act as an internal reward that is equivalent to extrinsic rewards.

Consistent with this view, novel stimuli activate midbrain DA structures in humans and other animals (Horvitz, 2000; Laurent, 2008; Wittmann, Bunzeck, Dolan, & Düzel, 2007; Wittmann, Daw, Seymour, & Dolan, 2008), and provide a bonus for organizing reward-based exploration (Barto et al., 2013; Brafman & Tenenbholz, 2003; Kakade & Dayan, 2002; Laurent, 2008; Lopes & Oudeyer, 2012). A recent study in our laboratory complements the findings presented above (Fig. 2) and suggests that novelty also recruits attentional resources (Foley, Jangraw, Peck, & Gottlieb, 2014; Peck, Suzuki, Efem, & Gottlieb, 2009). However, in contrast with a traditional RL view, the attention effects of novelty were independent of those of reward.

To compare the ways in which novelty and reward impacted attention, we trained monkeys on a task in which they received visual cues that could be highly familiar or novel and could bring “good” or “bad” news – that is, signaled whether the trial will end with a reward or a lack of reward (Fig. 3a). After presentation of the cue at a peripheral location, the monkeys maintained fixation for a brief delay and then made a saccade to a *separate* target that could appear either at the same or at the opposite location as the cue. In this task therefore, the cues did not allow the monkeys to choose the trial’s outcome. Instead, they only brought information and could automatically bias attention toward or away from their visual field location.

We recorded the activity of visually responsive neurons in the lateral intraparietal area (LIP), a cortical area which, together with the frontal eye field (FEF), is implicated in the selection of targets for attention and gaze (Bisley & Goldberg, 2010; Fig. 3b). LIP neurons have visual receptive fields (RF), selectively encode the locations of attention-worthy items, and are thought to provide top-down signals for orienting attention and rapid eye movements (saccades) (Ibid.).

In our task using reward cues, LIP neurons had sharp visual responses if a reward cue appeared in their RF, suggesting that both positive and negative cues transiently attracted attention (Fig. 3c) (Foley et al., 2014; Peck, Jangraw, Suzuki, Efem, & Gottlieb, 2009). However at slightly longer delays, the orienting response in LIP changed according to the reward signaled by the cue. The neurons maintained slight excitation for a familiar cue that brought good news (Fam+), but developed sustained *inhibition* for a familiar cue that signaled bad news (a dip below baseline firing rates in the Fam– trace in Fig. 3c). Consistent with these neuronal responses, saccades were facilitated if they were directed toward the location of a positive cue (which was excited in the LIP representation) and impaired if they

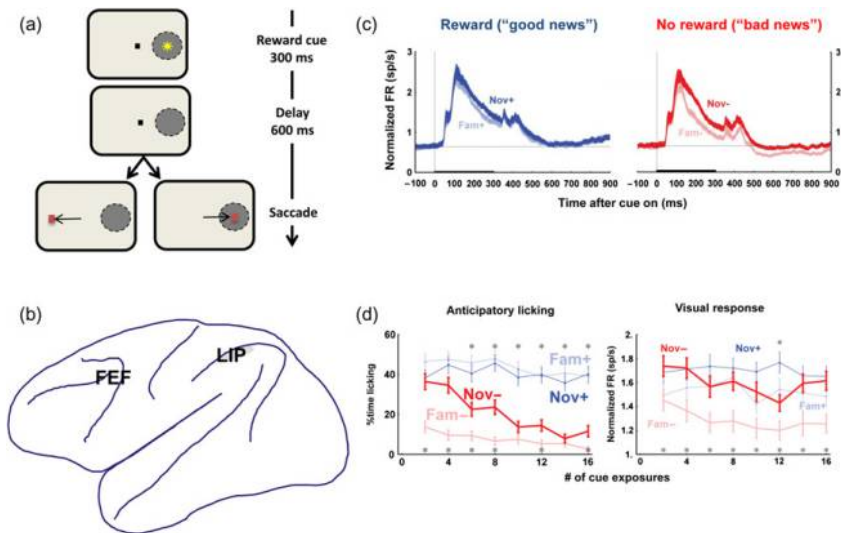


Fig. 3. Independent effects on reward and novelty on visual salience. (a) Task design. A trial began when the monkeys fixated a central fixation point (small black dot). A reward cue was then presented for 300 ms at a randomly selected location that could fall inside the RF of an LIP cell (gray oval) 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. (b) Cortical oculomotor areas. Lateral view of the macaque brain showing the approximate locations of the FEF and LIP. (c) LIP neurons are modulated by reward and novelty. Normalized activity (mean and standard error (SEM)) in a population of LIP cells, elicited by cues that appeared in the RF and which could be familiar or newly learn 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 responses and sustained delay period suppression that was not seen for Fam+ cues. However, newly learnt 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 show SEM. Stars show differences that are significant at $p < 0.01$. 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. *Source:* Reproduced with permission from Peck et al. (2009) (b) and Foley et al. (2014) (c and d).

were directed toward the location of a negative cue (which was suppressed in LIP). These neural and saccadic effects were spatially specific, occurring at the cue location but not at the opposite visual field location. That is, beyond producing global changes in arousal or motivation, the reward message conveyed by the cues modified the attentional processing of *specific* sources of information.

Comparison of novel and familiar cues showed that these reward-dependent attentional effects were much weaker or absent for newly learnt items, despite the fact that the monkeys quickly learnt the significance of those items. When they were first confronted with a novel cue, the monkeys showed anticipatory licking indicating that they expected to receive a reward following the cue, but this licking quickly extinguished if a cue turned out to signal a negative outcome (Foley et al., 2014) (Fig. 3d, left). However, LIP neurons showed enhanced responses to novel visual cues, and this enhancement persisted for dozens of presentations even for cues that signaled negative or positive outcomes (Fig. 3d, right). That is, newly learnt cues continued to produce enhanced salience and LIP responses on longer time scales, even when the cues signaled “bad news.”

As we discussed in the previous section therefore, novelty seems to exert dissociable effects through multiple pathways, both by activating motivational systems and through reward-independent visual/attentional effects. Understanding how these processes work in concert will be an important topic for future investigations.

Reward and Uncertainty

While the factors of novelty and surprise that we discussed above can engage attention independently of the observer’s task, attention can also be controlled in a top-down fashion – that is, tightly focused on achieving a goal. Since the early studies of Yarbus in the 1950s it has been appreciated that, when observers are engaged in a task, their eye movements are directed very selectively to task-relevant stimuli with very few glances to salient distractors, revealing the strength and importance of task-related control (Tatler, Hayhoe, Land, & Ballard, 2011).

Some insight into the computational basis of such task-related control comes from studies of naturalistic behaviors where subjects perform tasks such as driving in virtual reality settings (Gottlieb, Hayhoe, Hikosaka, & Rangel, 2014; Hayhoe & Ballard, 2014; Sullivan, Johnson, Rothkopf, Ballard, & Hayhoe, 2012). Behavior in such contexts was computationally

analyzed using RL models that partition the subjects' actions into discrete sub-tasks; for instance, while driving, one may have to coordinate between the sub-task of monitoring the speed and that of monitoring the road. These models suggest that gaze is allocated to competing sub-tasks based on two factors: the rewards and informational demands (uncertainty) of each individual task. This dual control mechanism allows subjects to direct gaze efficiently – to inform those actions that are not only valuable for achieving a goal but also have uncertainty and need for information (Hayhoe & Ballard, 2014; Sullivan et al., 2012; Tatler et al., 2011).

Remarkably, recent findings from our laboratory suggest that dual control by reward and uncertainty may operate even when subjects sample information spontaneously in intrinsically motivated fashion, without being obligated to perform an action (Daddaoua, Lopes, & Gottlieb, 2016). In the paradigm we used in that study, each trial had a different probability of delivering a reward, and monkeys received advance information about the probability that was in effect on a trial (e.g., 0%, 50%, or 100% reward probability; x -axis in Fig. 4). While waiting for reward delivery, the monkeys were given the *opportunity* to obtain additional information by searching for a reward cue that was hidden in a visual display, and which would provide perfect information about whether or not a reward will arrive. Importantly, the monkeys had no requirement to search for the additional cue; the reward arrived according to the signaled probability, whether the monkeys did or did not uncover the cue.

As shown in Fig. 4, the monkeys were intrinsically motivated to search, and their motivation depended on their prior knowledge of the reward probability. Search was almost completely absent if the monkeys expected to obtain no reward (0% likelihood), was maximal if the reward was uncertain (50% probability) and was also quite vigorous if the monkeys had certainty that they would receive a reward (100% likelihood). This pattern remained consistent over the experiment duration (Fig. 4, insets), suggesting that reward and uncertainty shape intrinsically motivated information gathering, similar to their role in instrumental (task-related) settings.

Along with these striking similarities, there are also critical differences between instrumental and non-instrumental settings. In an instrumental context, information guides ongoing actions, and the higher the quality of one's information, the higher the probability of success of those future actions. In contrast, in a non-instrumental context such as the one shown in Fig. 4, subjects cannot improve their chance of obtaining a reward by gathering information. The information the subjects gather can only regulate *their internal state*.

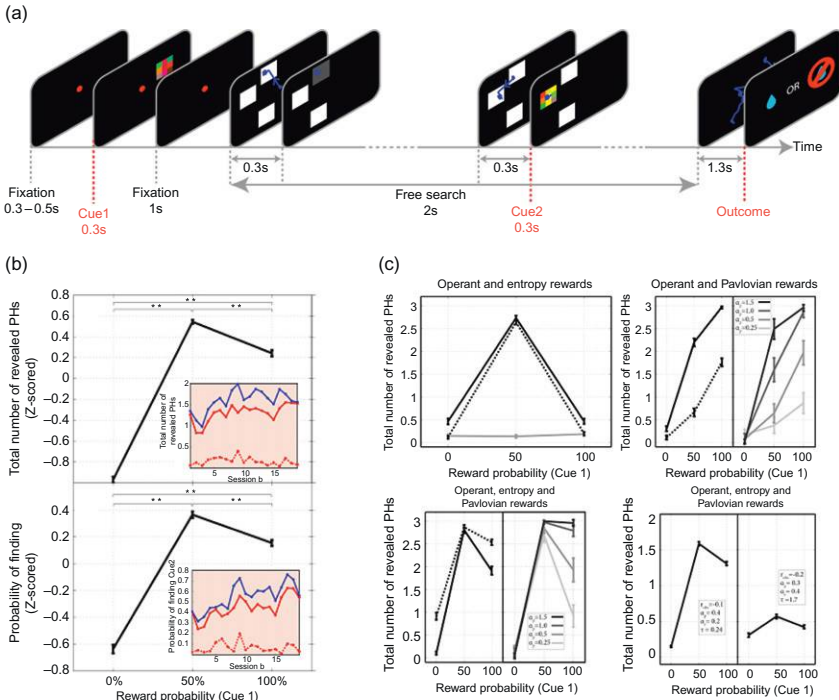


Fig. 4. Intrinsically motivated information search as a function of reward probability. (a) The layout of the task. See text for explanation. (b) The x-axis indicates the reward probability that was signaled to the monkey on each trial, before the information search began. The top panel shows the number of items sampled by the monkey as a function at each reward probability (mean and standard errors across all testing sessions (days), after z-scoring within individual sessions). The bottom panel shows the probability of finding the hidden cue (mean and standard errors, z-scored across all sessions). $*p < 0.025$ (Wilcoxon test). The insets in each panel show the average of the raw data in each daily session. The dotted, solid gray and solid black traces indicates 0%, 100%, and 50% reward probability. (c) Model simulations using different value functions. See text for details.

The fact that our subjects were interested in obtaining information when they had reward uncertainty (50% likelihood) suggests that, all things being equal, they prefer to resolve uncertainty as early as possible. The fact that the subjects were also interested in viewing positive cues when they had *no uncertainty* (100% likelihood) is quite remarkable and suggests that the mere act of viewing a positive cue has *intrinsic value*, even if the cue

is redundant and cannot be used to guide future actions. This conclusion is consistent with a rich literature showing that animals automatically approach and attend to positively conditioned Pavlovian cues (Castro & Berridge, 2014; Dayan, Niv, Seymour, & Daw, 2006; Flagel et al., 2011; Hickey, Chelazzi, & Theeuwes, 2010a, 2010b; Peck et al., 2009) and suggest that these automatic Pavlovian tendencies are important contributors to intrinsic motivation.

Especially importantly for our purpose, the dual role of reward and uncertainty has a striking resonance with dual-process psychological theories proposing that curiosity arises both from a desire to close “information gaps” (reduce uncertainty, or harvest information), and as a mere feeling of “interest” or “liking” of pleasurable items (Litman, 2007; Lowenstein, 1994). It remains a central question for future research to determine the precise ways in which curiosity is motivated by information gains and conditioned reinforcement from pleasurable cues, and to what extent this dual control may be inherited from more basic behaviors such as the task-related sampling of information.

Learning Progress and Metacognition

The four factors we reviewed above – novelty, surprise, reward and uncertainty – most likely act in combination and explain important aspects of curiosity mechanisms. However, several considerations suggest that, while these factors may be necessary, they are *not sufficient* to explain the full range of curiosity-based exploration.

Novelty and surprise are important heuristics for arousing curiosity, but they have the limitation that they do not necessarily signal significant or learnable environmental properties. A curiosity system that is based only on searching for novelty and surprise would only produce what early researchers called “diversive curiosity” (Lowenstein, 1994) – the type of transient curiosity we may show when we browse the internet with no specific aim – but cannot explain more deliberate, sustained investigative actions such as self-motivated study or scientific research.

Reward and uncertainty, on the other hand, can produce longer-lasting effects, but they are only well defined for highly practiced tasks, and are typically *not known* to agents when they embark on curiosity-driven exploration. For instance, in the tasks we described above the subjects were fully informed about the reward and uncertainty involved in a task (Figs. 3 and 4), but in more realistic settings a curious individual will begin

exploring with only vague estimates of these quantities. An agent cannot know the payoffs associated with a curiosity-driven action, and critically important, cannot even know the uncertainty associated with that action or indeed, assume that his investigations will *reduce uncertainty*.

Consistent with these intuitions, formal computational studies confirm that search strategies motivated by novelty, high or low uncertainty, or high or low entropy may be useful in well-delimited contexts, but are inefficient in acquiring knowledge and skills in unbounded large spaces that contain unlearnable tasks (Baranès & Oudeyer, 2009; Lopes & Montesano, 2014; Lopes & Oudeyer, 2010). Computational heuristics that motivate agents to explore states of high uncertainty may cause them to become trapped in random unlearnable tasks; conversely, heuristics that motivate agents to *minimize* uncertainty will cause them to focus exclusively on well-learned, predictable tasks. To understand the full range of our curiosity, therefore, we must account for the coexistence of two conflicting drives: the desire to reduce uncertainty on a short time scale versus the *intellectual risk taking* and will to increase uncertainty in order to learn on longer time scales.

To address these shortcomings, studies of artificial curiosity have developed computational strategies based on a metacognitive mechanism that assigns value to competing tasks based on the empirical learning progress (LP) related to each task (Baranes & Oudeyer, 2013; Lopes, Lang, Toussaint, & Oudeyer, 2012; Moulin-Frier, Nguyen, & Oudeyer, 2013; Oudeyer, Kaplan, & Hafner, 2007; Schmidhuber, 1991; Srivastava, Steunebrink, & Schmidhuber, 2013). Lopes et al. (2012) proved that exploration based on LP is equivalent to methods based solely on the number of visits (e.g., Brafman & Tennenholtz, 2003) but becomes more robust when encountering changing situations or having the wrong expectations.

LP-driven mechanisms prioritize competing tasks based on the rate of improvement – derivative – of the cost function that the learner is trying to maximize. LP can be defined based on the rate of improvement of predictions of a sensorimotor outcome or of the reward/success rate in a task. Compared to heuristics that search for high uncertainty, LP-driven mechanisms will motivate the learner to investigate situations that are initially uncertain and *keep exploring* them *only if* these situations lead to learning in practice. This can be formulated as an operational implementation of the information-theoretic framework of the free-energy principle (Friston et al., 2015).

An example of algorithmic architecture implementing an LP-driven curiosity process is the R-IAC architecture, detailed in Fig. 5b. In this

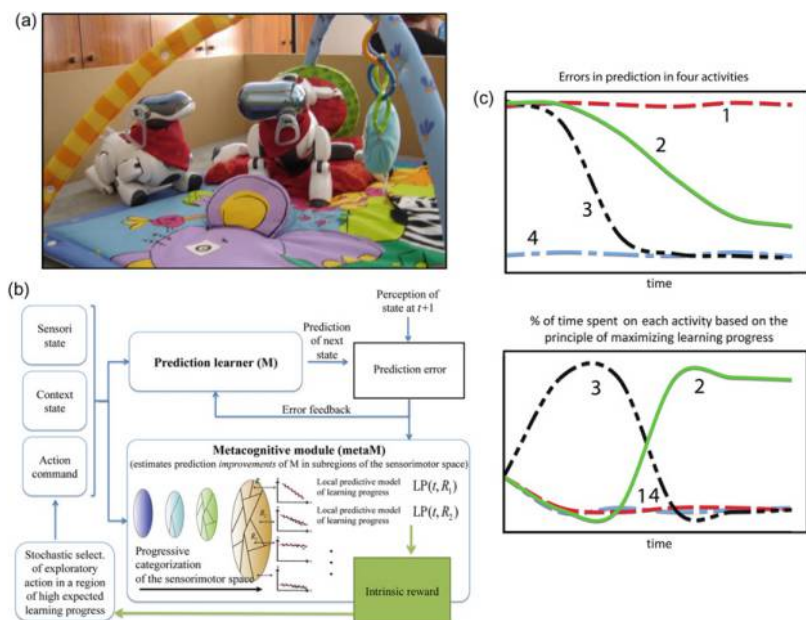


Fig. 5. Spontaneous curiosity-driven exploration can be efficiently driven by searching for situations that improve current predictions (LP), and self-organize a learning curriculum of increasing complexity. (a) The Playground Experiment: a quadruped robot placed on an infant play mat with a set of nearby objects, as well as an “adult” robot peer. The robot is equipped with a repertoire of motor primitives parameterized by several continuous numbers, which can be combined to form a large continuous space of possible actions. The robot learns how to use and tune them to affect various aspects of its surrounding environment, and exploration is driven by maximization of learning progress using the R-IAC architecture. We observe the self-organization of structured developmental trajectories, whereby the robot explores objects and actions in a progressively more complex stage-like manner while acquiring autonomously diverse affordances and skills that can be reused later on. The robot also discovers primitive vocal interaction as a result of the same process. (b) The R-IAC architecture implements this curiosity-driven process with several modules. A prediction machine (M) learns to predict the consequences of actions taken by the robot in given sensory states. A metacognitive module (metaM) estimates the evolution of errors in prediction of M in various subregions of the sensorimotor space, which in turn is used to compute learning progress as an intrinsic reward. Because the sensorimotor flow does not come pre-segmented into activities and tasks, a system that seeks to maximize differences in learnability is also used to progressively categorize the sensorimotor space into regions, which incrementally model the creation and refining of activities/tasks. Then an action selection system chooses activities to explore for which estimated learning progress is high. This choice is stochastic in order to monitor other activities for which learning progress might increase. (c) Confronted with four sensorimotor activities characterized by different learning profiles (i.e., evolution of prediction errors), exploration driven by maximization of learning progress results in avoidance of activities already predictable (curve 4) or too difficult to learn to predict (curve 1) to focus first on the activity with the fastest learning rate (curve 3) and eventually, when the latter starts to reach a plateau, to switch to the second most promising learning situation (curve 2). This allows the creation of an organized exploratory strategy necessary to engage in open-ended development. *Source:* Adapted with permission from (Kaplan & Oudeyer, 2007).

architecture, a robot learns to predict the consequence of its actions. Such predictive learning is made with statistical inference over the data collected when the robot carries out “experiments,” that is, tries an action and observes the results. The robot then chooses which task to perform based on a metacognitive module that monitors the evolution of prediction errors in various regions of the sensorimotor space: it selects regions to explore with a probability that is proportional to the rate of improvement in the past (such probabilistic scheme allows to continually search for new niches of progress).

In one study, we showed that such an architecture allows a robot to master hand-eye coordination much faster relative to strategies based on random exploration or a search for maximal uncertainty (Baranès & Oudeyer, 2009). Similar results were shown for the acquisition of other skills such as omnidirectional legged locomotion (Baranes & Oudeyer, 2013) or the manipulation of flexible objects (Nguyen & Oudeyer, 2013)

Interestingly, these analyses showed that, in addition to providing very efficient for acquiring new skills in large task spaces, LP-based algorithms produce exploration strategies that spontaneously progress from simple to more complex tasks in the absence of external instructions. For example, in the Playground Experiment (Fig. 5a) several behavioral and cognitive phases spontaneously formed during learning. After a phase of random body babbling, the robot focused on moving only certain body parts, and then focused on increasingly complex action-object affordances – beginning by learning how its leg can push or grasp objects, and ending up exploring how its vocalizations could produce reactions in another robot. Repeated runs of this experiment showed that in many cases similar developmental milestones appeared in a similar order while other robots showed deviations from these milestones or went through them in a different order, similar to the dual properties of universal tendencies and diversity seen in the development of infants (Oudeyer et al., 2007; Oudeyer & Smith, 2016).

In a related experiment on vocal development, robots used an LP-based algorithm to discover how to communicate with peers (Moulin-Frier et al., 2013). This experiment relied on a physical model of the vocal tract, its motor control and the auditory system, and showed how such a mechanism can explain the adaptive transition from vocal self-exploration with little sensitivity to the speech environment, to a later stage where vocal exploration becomes influenced by vocalizations of peers. Within the initial self-exploration phase, a sequence of vocal production stages self-organized, and shared properties with infant data: the vocal learner first discovered how to control phonation, then vocal variations of unarticulated sounds,

and finally articulated proto-syllables. As the vocal learner becomes more proficient at producing complex sounds, the imitating vocalizations of the teacher provide high LP resulting in a shift from self-exploration to vocal imitation.

One can apply such automatic organization of learning in intelligent tutoring systems. These systems' goal is to provide automatic assistance to learners based on their skills. As each student will have their own background, particular strengths, and weakness, a general model equal for all students will not accurately predict the behavior of any particular student. [Clement, Roy, Oudeyer, and Lopes \(2014\)](#) proposed the use of LP measures to allow an intelligent tutoring system to adapt to the particular learning progression of each individual student and showed that different paths provide a faster learning.

A salient property of LP-based curiosity algorithms is that they naturally give rise to a behavioral pattern that is a cornerstone of theories of intrinsic motivation – namely, the tendency of exploratory actions to progress from simpler to more complex tasks ([Ryan & Deci, 2000a, 2000b](#)). In a recent study, we replicated the latter effect in a laboratory setting by using a task where subjects were given a set of computer games of variable complexity and could freely choose the games they wished to play ([Baranes, Oudeyer, & Gottlieb, 2014](#)). A game lasted several seconds and required subjects to press a key as accurately as possible to intercept a series of dots that streamed past the center of the screen ([Fig. 6a](#)). Even though the subjects received no instruction about which game to select, they spontaneously organized their exploration in consistent patterns. Subjects did an initial survey of the entire space of the available games – including the most difficult games where dot speed were very high and performance was low – and then focused their exploration on games of intermediate complexity, where performance was 70–80% correct ([Fig. 6b](#)). This general trend was modulated by factors such as how much novelty could be found in games of a given complexity, and how the difficulty of tasks was spread along the game distribution.

Therefore, behavioral evidence is consistent with the idea that a self-organizing pattern based on task complexity shapes intrinsically motivated behaviors in a variety of contexts. However, more evidence is needed to establish whether this pattern indicates an LP-based mechanism. Many forms of learning are nonlinear in time (showing effects such as savings and consolidation) and it is unclear whether subjects can accurately track their LP or which aspects of progress determine intrinsic motivation. Addressing these questions will be critical for a better understanding of our most elaborate curiosity-based forms of exploration.

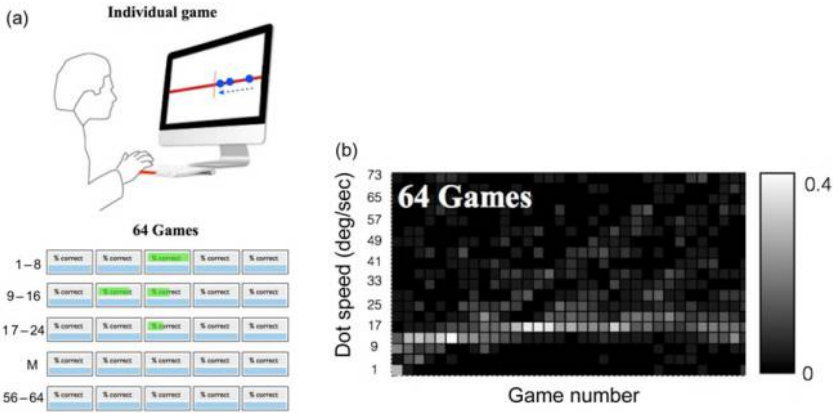


Fig. 6. Intrinsically motivated exploration in a laboratory game. (a) Task design. The top panel shows an individual game lasting ~30 seconds, in which subjects pressed a key to intercept a stream of moving dots (arrow) as they crossed the screen center. The bottom panel shows the selection screen with 64 games from which subjects could freely choose which game they wished to sample. (b) The speed (difficulty) of the selected games during a session. The color map indicates the probability of selection of a given speed, measured across all subjects in a sliding window over the session. *Source:* Adapted with permission from Baranes et al. (2014).

CONCLUSIONS

We reviewed a theory of curiosity that emerges from a synthesis of findings from neuroscience and machine-learning fields. A central theme that we stressed throughout the review is the fact that curiosity implies a tight interaction between cognitive and motivational systems. Rather than being an external process that acts *on* learning and cognition, curiosity arises organically in conjunction with these processes. When it is engaged in cognitive processing, the brain does more than simply discriminate, encode, and remember information – it also evaluates the epistemic and emotional qualities of these cognitive operations and uses them to generate the “interest” and intrinsic motivation that determine its future engagement with a task.

We have also stressed the fact that the mechanisms that generate curiosity would be ideally adapted to allow agents to discover new and useful regularities in large open-ended spaces that contain unlearnable tasks. This is a formidable challenge for which agents may have no optimal solution, and we proposed that organisms meet this challenge by combining a variety of

strategies. These strategies include simple heuristics such as exploration based on novelty, surprise, reward, and uncertainty that may have their roots in simpler active sensing behaviors. In addition, they may include more complex targeted investigations potentially based on metacognitive estimates of LP and information gain. These mechanisms may act in concert to autonomously organize exploration of vast unbounded spaces, steering agents away from overlearned (low uncertainty) tasks and away from unlearnable (high uncertainty) tasks, toward a middle range where the agent can make LP and discover new structures. While many of the views we outlined remain to be refined and substantiated through future research, we hope that they provide a useful roadmap to the key questions posed by that research.

REFERENCES

- Baldi, P., & Itti, L. Of bits and wows: A Bayesian theory of surprise with applications to attention. *Neural Networks*, 23, 649–666.
- Baranès, A., & Oudeyer, P.-Y. (2009). R-IAC: Robust intrinsically motivated exploration and active learning. *IEEE Transactions on Autonomous Mental Development*, 1, 155–169.
- Baranes, A., & Oudeyer, P. Y. (2013). Active learning of inverse models with intrinsically motivated goal exploration in robots. *Robotics and Autonomous Systems*, 61, 49–73.
- Baranes, A. F., Oudeyer, P. Y., & Gottlieb, J. (2014). The effects of task difficulty, novelty and the size of the search space on intrinsically motivated exploration. *Frontiers in Neuroscience*, 8, 1–9.
- Baranes, A. F., Oudeyer, P. Y., & Gottlieb, J. (2015). Eye movements encode epistemic curiosity in human observers. *Vision Research*, 117, 81–90.
- Barto, A., Mirolli, M., & Baldassarre, G. (2013). Novelty or surprise? *Frontiers in Psychology*, 4, 710.
- Barto, A. G. (2013). Intrinsic motivation and reinforcement learning. In G. Baldassarre & M. Mirolli (Eds.), *Intrinsically motivated learning in natural and artificial systems* (pp. 17–47). Berlin: Springer.
- Bisley, J., & Goldberg, M. (2010). Attention, intention, and priority in the parietal lobe. *Annual Review of Neuroscience*, 33, 1–21.
- Blanchard, T. C., Hayden, B. Y., & Bromberg-Martin, E. S. (2015). Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron*, 85, 602–614.
- Brafman, R. I., & Tennenholtz, M. (2003). R-max – A general polynomial time algorithm for near-optimal reinforcement learning. *The Journal of Machine Learning Research*, 3, 213–231.
- Bromberg-Martin, E. S., & Hikosaka, O. (2009). Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron*, 63, 119–126.
- Castro, D. C., & Berridge, K. C. (2014). Advances in the neurobiological basis for food “liking” versus “wanting”. *Physiology and Behavior*, 136, 22–30.
- Clement, B., Roy, D., Oudeyer, P. Y., & Lopes, M. (2014). Online optimization of teaching sequences with multi-armed bandits. *7th International conference on educational data mining*.

- Csikszentmihalyi. (1990). *Flow: The psychology of optimal experience*. New York, NY: Harper & Row.
- Daddaoua, N., Lopes, M., & Gottlieb, J. (2016). Intrinsically motivated oculomotor exploration guided by uncertainty reduction and conditioned reinforcement in non-human primates. *Scientific Reports*, 6.
- Dayan, P., Niv, Y., Seymour, B., & Daw, N. D. (2006). The misbehavior of value and the discipline of the will. *Neural Networks*, 19, 1153–1160.
- Flagel, S. B., Clark, J. J., Robinson, T. E., Mayo, L., Czul, A., Willuhn, I., ... Akil, H. (2011). A selective role for dopamine in stimulus-reward learning. *Nature*, 469, 53–57.
- Foley, N. C., Jangraw, D. C., Peck, C., & Gottlieb, J. (2014). Novelty enhances visual salience independently of reward in the parietal lobe. *Journal of Neuroscience*, 34, 7947–7957.
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, 6, 1–28.
- Friston, K., Schwartenbeck, P., Fitzgerald, T., Moutoussis, M., Behrens, T., & Dolan, R. J. (2013). The anatomy of choice: Active inference and agency. *Frontiers in Human Neuroscience*, 7, 598.
- Gottlieb, J., Hayhoe, M., Hikosaka, O., & Rangel, A. (2014). Attention, reward and information seeking. *Journal of Neuroscience*, 34, 15497–15504.
- Gruber, M. J., Gelman, B. D., & Ranganath, C. (2014). States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron*, 84, 486–496.
- Hayhoe, M., & Ballard, D. (2014). Modeling task control of eye movements. *Current Biology*, 24, 622–628.
- Hickey, C., Chelazzi, L., & Theeuwes, J. (2010a). Reward changes salience in human vision via the anterior cingulate. *Journal of Neuroscience*, 30, 11096–11103.
- Hickey, C., Chelazzi, L., & Theeuwes, J. (2010b). Reward guides vision when it's your thing: Trait reward-seeking in reward-mediated visual priming. *PLoS One*, 5, e14087.
- Horvitz, J. C. (2000). Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neuroscience*, 96, 651–656.
- Itti, L., & Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision Research*, 49, 1295–1306.
- Kakade, S., & Dayan, P. (2002). Dopamine: Generalization and bonuses. *Neural Networks*, 15, 549–559.
- Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T., & Camerer, C. F. (2009). The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychological Science*, 20, 963–973.
- Kaplan, F., & Oudeyer, P.-Y. (2007). In search of the neural circuits of intrinsic motivation. *Frontiers in Neuroscience*, 1, 225.
- Laurent, P. A. (2008). The emergence of saliency and novelty responses from Reinforcement Learning principles. *Neural Networks: The Official Journal of the International Neural Network Society*, 21, 1493–1499.
- Litman, J. A. (2007). Curiosity as a feeling of interest and feeling of deprivation: The I/D model of curiosity. In P. R. Zelig (Ed.), *Issues in the Psychology of Motivation*. Nova Science Publishers, Inc.
- Lopes, M., Lang, T., Toussaint, M., & Oudeyer, P.-Y. (2012). Exploration in model-based reinforcement learning by empirically estimating learning progress. Tahoe, USA: Neural Information Processing Systems (NIPS).
- Lopes, M., & Montesano, L. (2014). Active learning for autonomous intelligent agents: Exploration, curiosity and interaction, *arXiv: 1403.1497 [ca.AI]*.
- Lopes, M., & Oudeyer, P.-Y. (2010). Guest editorial active learning and intrinsically motivated exploration in robots: Advances and challenges. *IEEE Transactions on Autonomous Mental Development*, 2, 65–69.

- Lopes, M., & Oudeyer, P.-Y. (2012). The strategic student approach for life-long exploration and learning. *2012 IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL)* (pp. 1–8).
- Lowenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, *116*, 75–98.
- Moulin-Frier, C., Nguyen, S. M., & Oudeyer, P.-Y. (2013). Self-organization of early vocal development in infants and machines: The role of intrinsic motivation. *Frontiers in Psychology*, *4*.
- Nguyen, S. M., & Oudeyer, P.-Y. (2013). Socially guided intrinsic motivation for robot learning of motor skills. *Autonomous Robots*, *36*(3), 273–294.
- Oudeyer, P.-Y., Kaplan, F., & Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, *11*, 265–286.
- Oudeyer, P. Y., & Smith, L. (2016). How evolution may work through curiosity-driven developmental process. *Topics in Cognitive Science*, *8*(2), 492–502.
- Peck, C. J., Jangraw, D. C., Suzuki, M., Efem, R., & Gottlieb, J. (2009). Reward modulates attention independently of action value in posterior parietal cortex. *Journal of Neuroscience*, *29*, 11182–11191.
- Rieiro, H., Martinez-Conde, S., & Macknik, S. L. (2013). Perceptual elements in Penn & Teller's "Cups and Balls" magic trick. *PeerJ*, *1*, e19.
- Risko, E. F., Anderson, N. C., Lanthier, S., & Kingstone, A. (2012). Curious eyes: Individual differences in personality predict eye movement behavior in scene-viewing. *Cognition*, *122*, 86–90.
- Ryan, R. M., & Deci, E. L. (2000a). Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary Educational Psychology*, *25*, 54–67.
- Ryan, R. M., & Deci, E. L. (2000b). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychology*, *55*, 68–78.
- Schmidhuber, J. (1991). Curious model-building control systems. *IEEE International Joint Conference on Neural Networks* (Vol. 2, pp. 1458–1463).
- Singh, S., Lewis, R. L., Barto, A. G., & Sorg, J. (2010). Intrinsically motivated reinforcement learning: An evolutionary perspective. *IEEE Transactions on Autonomous Mental Development*, *2*(2), 70–82.
- Srivastava, R. K., Steunebrink, B. R., & Schmidhuber, J. (2013). First experiments with PowerPlay. *Neural Networks*, *41*, 130–136.
- Sullivan, B. T., Johnson, L., Rothkopf, C. A., Ballard, D., & Hayhoe, M. (2012). The role of uncertainty and reward on eye movements in a virtual driving task. *Journal of Visualization*, *12*.
- Tatler, B. W., Hayhoe, M. N., Land, M. F., & Ballard, D. H. (2011). Eye guidance in natural vision: Reinterpreting saliency. *Journal of Visualization*, *11*, 5–25.
- Thorndike, E. L. (1911). *Animal intelligence*. New York, NY: The Macmillan Company.
- Vo, M. L., & Wolfe, J. M. (2015). The role of memory for visual search in scenes. *Annals of New York Academy of Sciences*, *1339*, 72–81.
- Wittmann, B. C., Bunzeck, N., Dolan, R. J., & Düzel, E. (2007). Anticipation of novelty recruits reward system and hippocampus while promoting recollection. *NeuroImage*, *38*, 194–202.
- Wittmann, B. C., Daw, N. D., Seymour, B., & Dolan, R. J. (2008). Striatal activity underlies novelty-based choice in humans. *Neuron*, *58*, 967–973.