Trends in
Cognitive Sciences

CellPress
OPEN ACCESS

Review

# Curiosity and the dynamics of optimal exploration

Francesco Poli [1,*], Jill X. O'Reilly [2], Rogier B. Mars [1,2] and Sabine Hunnius [1]

**What drives our curiosity remains an elusive and hotly debated issue, with multiple hypotheses proposed but a cohesive account yet to be established. This review discusses traditional and emergent theories that frame curiosity as a desire to know and a drive to learn, respectively. We adopt a model-based approach that maps the temporal dynamics of various factors underlying curiosity-based exploration, such as uncertainty, information gain, and learning progress. In so doing, we identify the limitations of past theories and posit an integrated account that harnesses their strengths in describing curiosity as a tool for optimal environmental exploration. In our unified account, curiosity serves as a 'common currency' for exploration, which must be balanced with other drives such as safety and hunger to achieve efficient action.**

## Highlights

Traditional theories frame curiosity as an intrinsic motivation to obtain information to resolve uncertainty.

Recent empirical breakthroughs challenge this perspective and reframe curiosity in terms of the learning process itself: curiosity is sustained when individuals perceive improvements in their performance and actively engage in the process of learning.

Unlike previous research, which remained agnostic about the change of curiosity over time, the learning progress framework provides insights into the temporal dynamics of curiosity, shedding light on its fluctuations based on the extent of learning progress individuals make.

By integrating the desire to obtain information with the drive to make learning progress, we offer a unified account of curiosity.

## What is curiosity?

Until recently, cognitive and neuroscientific theories of curiosity espoused a view in which curiosity is motivated by the goal of obtaining information or knowledge. When uncertainty arises (e.g., how will the book end? Who will win the lottery?), curiosity pushes us into seeking relevant information that resolves the uncertainty. In line with this idea, functional imaging studies showed that information activates the reward network, suggesting that receiving information is **intrinsically rewarding** (see Glossary) [1–4].

However, recent evidence from a variety of fields indicates that this is not the whole picture. Behavioral studies with human adults show that curiosity is motivated by the learning process itself. For instance, participants exploring probabilistic environments showed engagement with the task when they were learning the structure of the environment, but not when they either had learned the structure (low uncertainty) or when the structure was unlearnable (high uncertainty) [5,6]. Moreover, new studies on mice [7] show that midbrain dopamine, a key neurotransmitter for the experience of reward, is released depending on the rate of learning, indicating that the neural underpinnings of curiosity may be more closely tied to the learning process than to the mere resolution of uncertainty. This evolving body of evidence suggests that traditional theories, which primarily link curiosity to acquiring knowledge, are due for a substantive re-evaluation.

More broadly, the amount of experimental evidence on the cognitive and brain mechanisms underlying curiosity and exploration has grown exponentially in recent years, across cognitive (neuro)science [8], artificial intelligence [9], as well as educational [10], social [11], and organizational psychology [12,13]. Hence, the time is now ripe to translate this experimental progress into a robust, unitary theory of human curiosity. Here, we integrate the classic view of curiosity as motivated by the goal of obtaining knowledge with the newly emerging view of curiosity as motivated by **learning progress**.

[1]Donders Institute for Brain, Cognition, and Behaviour, Radboud University, Nijmegen, The Netherlands
[2]Wellcome Centre for Integrative Neuroimaging, Centre for Functional MRI of the Brain (FMRIB), Nuffield Department of Clinical Neurosciences, John Radcliffe Hospital, University of Oxford, Oxford, UK

*Correspondence: francesco.poli@donders.ru.nl (F. Poli).

We will first introduce the main views of curiosity put forward in cognitive neuroscience and their relationship with the concepts of **prediction error**, uncertainty, and **information gain**. We frame these aspects of curiosity in the predictive processing framework, a popular account for explaining how humans perceive and act upon the world [14], and analyze their limitations in supporting optimal curiosity-driven exploration. Then, we show how the account of curiosity as a drive to make learning progress can identify aspects of exploratory behavior that have been largely disregarded by previous research and how it can improve our understanding of curiosity.

## Deconstructing curiosity

Predictive processing offers a robust framework for modeling how we handle uncertainty, making it instrumental in understanding curiosity. It holds that our brains operate by generating predictions about the world and then adjusting those predictions based on the sensory evidence we receive. For example, entering a neighborhood where you have seen many cats before, you might expect to spot one (Figure 1). Yet, you unexpectedly notice an animal with a long prehensile tail and large teeth. This mismatch between the internal models of the world and the sensory evidence is defined as prediction error or surprise [15,16]. By adjusting your prior models of the world based on these error signals, the brain can improve its internal models and future predictions.

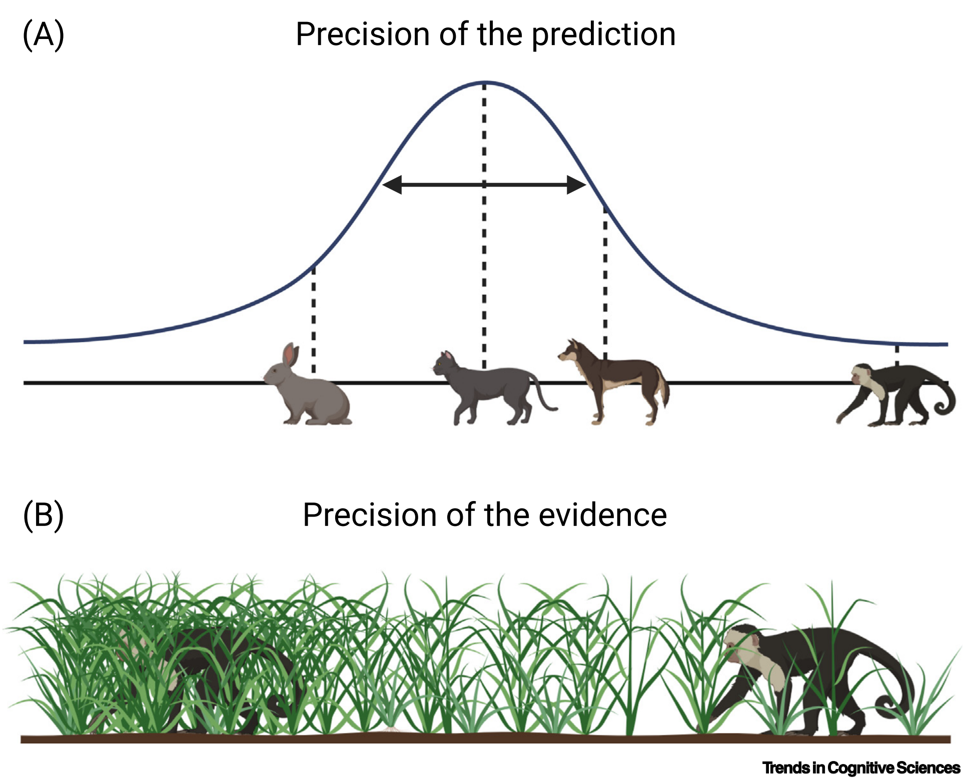


Figure 1. Predictions, evidence, and their precision. (A) Within predictive processing, predictions are probabilistic. The precision of the prediction is captured by the width of the distribution of possibilities (black arrow), where higher precision results in higher certainty of the prediction. Here, prediction is represented in a notional 1D space for illustration; however, in reality, the feature space for animal identification would be high dimensional. (B) The evidence also has varying degrees of precision from low (left) to high (right), depending on the amount of sensory noise (i.e., the density of the bushes). The unexpected sensory evidence (e.g., large teeth, prehensile tail) will travel through the brain in the form of prediction error signals, whereas the expected sensory evidence (e.g., black fur, four-legged) will not.

### Glossary

**Information gain:** the degree to which new evidence enhances the accuracy or completeness of existing knowledge or models.

**Intrinsic reward:** an internal incentive that drives behavior due to the inherent pleasure and interest in the activity itself, rather than any external benefits or penalties.

**Learning progress:** the rate of improvement in understanding or proficiency in a given domain. It can be quantified as the rate of change in prediction errors.

**Optimal exploration:** the most efficient strategy for seeking out and investigating new environments or ideas to gain knowledge or experience, while minimizing wasted physical or cognitive efforts.

**Precision:** a measure of the confidence or certainty of predictions and prediction errors, determining how much sensory feedback can alter the predicted state. It is the inverse of uncertainty.

**Prediction error:** the mismatch between the expected outcome of an action or event and the actual outcome, also referred to as surprise, novelty, or violation of expectations.

In this framework, predictions and evidence come with a fundamental feature: **precision** [16,17]. Predictions have a degree of precision in that, although one possibility is often more likely than others, there is uncertainty around it: the animal hiding is most likely a cat, but it might well be a dog, it could be a rabbit, and, in an extremely unlikely scenario, a monkey (Figure 1A). The evidence (and the prediction error resulting from it) has a degree of precision, too: how confident are you that your percept is accurate? The animal looked like a monkey, but if the bushes are thick and you cannot see well (Figure 1B), you should not trust perceptual information much. The final internal model must weigh in the precision of the prediction and the precision of the evidence to infer the most likely state of the world [18].

Some situations require us to change our internal models of the environment. Encountering high uncertainty, such as wondering what animal might be lurking in the bushes, or high prediction errors, such as being surprised upon seeing an unexpected animal, often prompts us to seek further information. Whereas the current article focuses on intrinsic exploration and curiosity, additional elicitors of information demand have been identified, including (monetary) value [19] or valence [20].

### Curiosity motivated by the goal of obtaining knowledge

There are two prominent perspectives on curiosity: one focuses on the role of uncertainty as a catalyst for curiosity, while the other emphasizes the pursuit of information gain as the primary mechanism underlying curiosity-driven exploration. The connection between uncertainty and curiosity has been substantiated by a considerable body of research [21–24]. Overall, there is empirical evidence that uncertainty is related to subjective reports of curiosity [22,25,26] as well as increased interest [27] and enhanced exploratory behavior [8,28,29], already from a very young age [30,31]. For example, recent findings indicate that curiosity rises as the uncertainty surrounding a lottery result increases [22]. Yet, curiosity had no substantial connection with the expected value of such lotteries, indicating that it was independent from extrinsic rewards such as monetary compensation [18]. The relation between curiosity and uncertainty holds even outside of controlled laboratory settings, expanding its validity to real-world contexts. Specifically, it was found that in classrooms, students were most curious when they experienced uncertain events, especially when this uncertainty was unexpected [32].

However, being drawn towards highly unpredictable situations can be maladaptive in the long run, in that if uncertainty cannot be reduced, a sustained focus would imply a waste of cognitive resources [33]. In accordance with this, it has been shown that during trivia questions, repeated exposure to unknown questions that elicit large prediction errors leads to the depletion of attentional resources, which results in participants disengaging from the task [27]. This indicates that the concepts of uncertainty alone cannot account for the temporal dynamics of curiosity-driven exploration. Specifically, it cannot explain why curiosity will not last if the uncertainty does not reduce [34].

A different line of research suggests that uncertainty is not directly related to curiosity [35], but that the main factor underlying curious behavior might be the opportunity to gain information that minimizes the uncertainty of the situation, rather than the uncertain situation in itself [35,36]. The link between curiosity and information gain has received substantial attention recently [35,37–39]. Humans and other animals are motivated to seek information [38,40], even when this information is not instrumental to acquiring extrinsic rewards [34,41,42]. Specifically, humans show a preference for information that is important to them, positive in valence, or salient [43,44]. Yet, we also feel curious about information even when it is unpleasant [45], costly to acquire [46], or when acquiring it is high-risk [1,47]. The intrinsic drive for information is also robust across development [48,49], starting from the first year of life [31,50,51]. This conceptualization of curiosity arising only when there is an opportunity to fill an information gap addresses the major limitation of theories that relate curiosity to uncertainty directly. Namely, it predicts that curiosity does not remain high when uncertainty is high but information cannot be obtained. This helps agents not to get stuck in unlearnable situations.

In Box 1, we use reinforcement learning as a tool to formalize the different mechanisms that might underlie curiosity. Reinforcement learning is commonly used to quantify the latent variables underlying learning processes [52,53]. As such, it offers a systematic way to represent the concepts of prediction error, uncertainty, and information gain. By combining insights from previous research with mathematical definitions of curiosity, we offer a more nuanced and comprehensive account of the mechanisms underlying curiosity. After introducing all mechanisms that might underlie curiosity, we will use these reinforcement-learning quantifications to illustrate the temporal dynamics of curiosity and propose a unified account.

Although a drive towards information gain successfully deters people from unlearnable situations, it is still suboptimal in driving efficient exploration. Some situations offer a large amount of

### Box 1. Implementing curiosity with a reinforcement learning model

Reinforcement learning models are widely used to learn a true value (e.g., the position of an object in space, or the reward associated with a certain choice) given continuous incoming evidence [52,53]. At each time-point $t$, an agent makes a prediction, observes the outcome, and updates their previous prediction based on the magnitude of the error. The update rule is:

$$P(s|m)_t = P_{t-1} + \alpha(r_{t-1} - P_{t-1}) \quad \text{[I]}$$

where the current prediction $P_t$ about a state of the world $s$ is obtained by updating the previous prediction $P_{t-1}$ by the prediction error. The prediction error is captured by the difference between the previous prediction $P_{t-1}$ and the evidence $r_{t-1}$. The learning rate $\alpha$ indicates how much the prediction error will change previous expectations. Note that predictions are conditioned on the model $m$, indicating that an agent will only generate predictions for states that are relevant to them [74].

Uncertainty is tightly coupled to precision, as uncertain situations arise from a lack of precision in the evidence or in the predictions. The precision of the prediction captures how certain we are about our prediction being correct. Hence, rather than a point estimate, our prediction is probabilistic. $u_{P_t}$ is the most likely outcome, but with uncertainty around it:

$$P_t \sim N\left(u_{P_t}, \sigma^2\right) \quad \text{[II]}$$

where the prediction $P_t$ comes from a normal probability distribution that contains uncertainty, expressed in terms of variance, or $\sigma^2$.

Prediction error is defined as the discrepancy between a prediction and the evidence:

$$PE_t = |r_{t-1} - P_{t-1}| \quad \text{[III]}$$

Prediction error is also referred to as surprise, as it quantifies how unexpected (i.e., surprising) the current state of the world is compared to what was predicted.

Information gain is defined as any change of the internal model $P_t$ from one time-point to the next. When internal models remain stable, no information is gained. A common way to operationalize the information gain between two predictions at consecutive time-points is the Kullback-Leibler divergence between the two [98]:

$$IG = \int P_t \log_2(P_t) - \int P_t \log_2(P_{t-1}) \quad \text{[IV]}$$

where the information gain captures the difference between the probability distribution of the current prediction $P_t$ and the previous prediction $P_{t-1}$. This construct is related to 'expected' information gain, which computes the information that is expected at time-point $t+1$ instead of the information that has already been received. These two constructs often produce similar estimates [51] and, for simplicity, they will not further be differentiated in this article.

information (e.g., a book about Bayesian statistics) but acquiring it might be challenging or time-consuming (the book is too complex or too long). In accordance with this, it was found that instead of seeking stimuli that are the most informative, people engage with stimuli that offer easily interpretable information [54]. Hence, as in the case of curiosity as a response to uncertainty, also a drive for information gain fails to fully capture this temporal aspect of curiosity, namely why curiosity wanes when the information is gained too slowly [55,56]. Next, we suggest that the learning progress account of curiosity offers a way to resolve this limitation.

### Curiosity motivated by the learning process itself

The learning progress account of curiosity holds that the brain is intrinsically motivated to engage in activities during which learning is occurring because this leads to an improvement of our internal models of the world [57,58]. Within this account, agents are expected to engage preferentially in activities and tasks that are just beyond their current level of skills or knowledge, with the goal of maximizing learning. At the implementation level, this entails that what we are curious about depends on the 'rate' at which we are decreasing our prediction errors [33,59] (Box 2). If prediction errors are not reduced, this might be because an activity is too difficult for the agent's current skillset or because the agent has already mastered the task. In these cases, curiosity about the activity should wane [31,60]. In this regard, the learning progress account gives a mechanistic explanation to the U-shape relation that is often found between human curiosity and uncertainty (i.e., a preference for intermediate levels of uncertainty or complexity) [4,30,61]: by maximizing learning progress, an agent will incidentally end up in environments that are neither too simple nor too complex.

This learning-maximization strategy is highly beneficial for guiding **optimal exploration**, as it allows us to improve internal models of the world at a high pace and avoids wasting cognitive resources on activities that cannot be learned. Indeed, multiple computational experiments [62,63] show the evolutionary advantages of an intrinsic motivation to learn, demonstrating that it is more efficient to evolve a general architecture that rewards learning itself rather than domain-specific modules for each cognitive ability.

Consistent with this computational work, recent studies suggest that learning progress may be experienced as intrinsically rewarding [57,64,65]. Studies conducted on mice indicate that the

#### Box 2. A formal definition of learning progress

Quantifying learning progress is not easy, as it requires tracking how quickly an agent's internal model of the world changes over time. Although many options exist [60,99,100], prediction errors are one of the simplest ways of capturing how well a model performs, as they tend to be high in a model that performs poorly. Hence, given the function of prediction error (i.e., how prediction errors change over time), learning progress is captured by its derivative:

$$LP = f(PE)' \qquad \text{[I]}$$

This value can be difficult to infer if the function of the prediction error is unknown. To solve this, previous research has captured curiosity with an approximate quantification of rate of change in prediction errors. For instance, learning progress was defined as the difference in prediction errors between two consecutive time-points [70]:

$$LP_t \cong PE_{t-1} - PE_t \qquad \text{[II]}$$

This approximation was employed successfully in research on artificial agents [101] and humans [5,6]. Regarding research on artificial agents, it was shown that endowing robots with an intrinsic drive to maximize learning progress led to faster and more robust learning [101]. Robots stayed focused on activities in which they were improving their predictions and avoided the ones that were either too easy or unlearnable. As concerns human behavior, it was shown that participants' learning progress in an open-ended exploration task predicted whether they remained engaged with the task or switched to a different one [2]. Moreover, similar to artificial agents, individuals who relied on learning progress to tailor their exploration had better learning outcomes [3].

release of midbrain dopamine, a crucial neurotransmitter associated with the sensation of reward, is influenced by the rate of learning [7] rather than the expected level of predictability of rewarding stimuli, contrary to previous assumptions [66]. Specifically, manipulations of mesolimbic dopamine using a neural network-based model showed the involvement of this neurotransmitter in adapting the rate of learning to improve efficient action over the environment [7]. These findings highlight the significant role of dopamine activity in regulating learning progress.

Although this conceptualization of curiosity has been used extensively to optimize the exploration strategies of artificial agents [60,67–69], evidence on whether humans display a pattern of behaviors consistent with such a drive was lacking until recently [5,6,70]. New studies now provide evidence for the idea that humans rely on learning progress to tailor their attention [31], task engagement [5,6,70,71], and sampling decisions [5]. For example, it was shown that how much learning progress adult participants make predicts whether they will remain engaged on a given task or switch to a different one instead [5]. Moreover, learning progress is associated with greater task engagement and subjective liking [70].

The learning progress account offers a simple heuristic mechanism to guide exploration, but its simplicity comes with a cost. If this is a generic drive, it cannot take into account how useful a given activity is to the agent's goal [72]. As such, every bit of learning progress is treated as equal, irrespective of its relevance. This might erroneously draw an agent towards an activity that allows for fast learning but holds little relevance for their goals. This contrasts with the predictive processing conceptualization of information gain, where information is specifically assessed as a function of its relevance for the agent [73]. In addition, the learning progress account cannot explain how individuals choose between different activities that offer a similar amount of learning progress [74]. We argue that overcoming the limitations of each account of curiosity in terms of their utility for exploration is only possible by integrating them.

## A unified account of curiosity

Here, we integrate the learning progress conceptualization of curiosity with the earlier described conceptualizations based on uncertainty and information gain, using the predictive processing framework as a guide. Specifically, we show that the maximization of learning progress should be considered an additional mechanism underlying curiosity, independent of but complementary to maximizing information gain (Figure 2). Conversely, we show that following uncertainty in itself is a suboptimal behavior.

### Curiosity as a tool for optimal exploration

Assessing the different accounts of curiosity requires a criterion for comparison, whether it be a theoretical ideal or an empirical benchmark. Without such a criterion, it is difficult to determine which mechanism underlying curiosity (i.e., seeking uncertainty, information gain, or learning progress) is more efficient. We will assess these mechanisms by their ability to foster optimal exploration of the environment. We define optimality in terms of an agent's success in improving their internal models of the world (i.e., improving the precision of their predictions) without spending too much time exploring in vain, thus wasting cognitive resources. In other words, optimal exploration optimizes the trade-off between an internal model's precision and the costs of exploration (e.g., time and effort). This viewpoint situates curiosity as an algorithmic mechanism that delivers a computationally demonstrable ecological advantage, promoting optimal environmental exploration. By defining a clear criterion for comparison, we are better positioned to critically evaluate the different mechanisms that might underlie curiosity and to integrate them to reach a more comprehensive operationalization of curiosity.

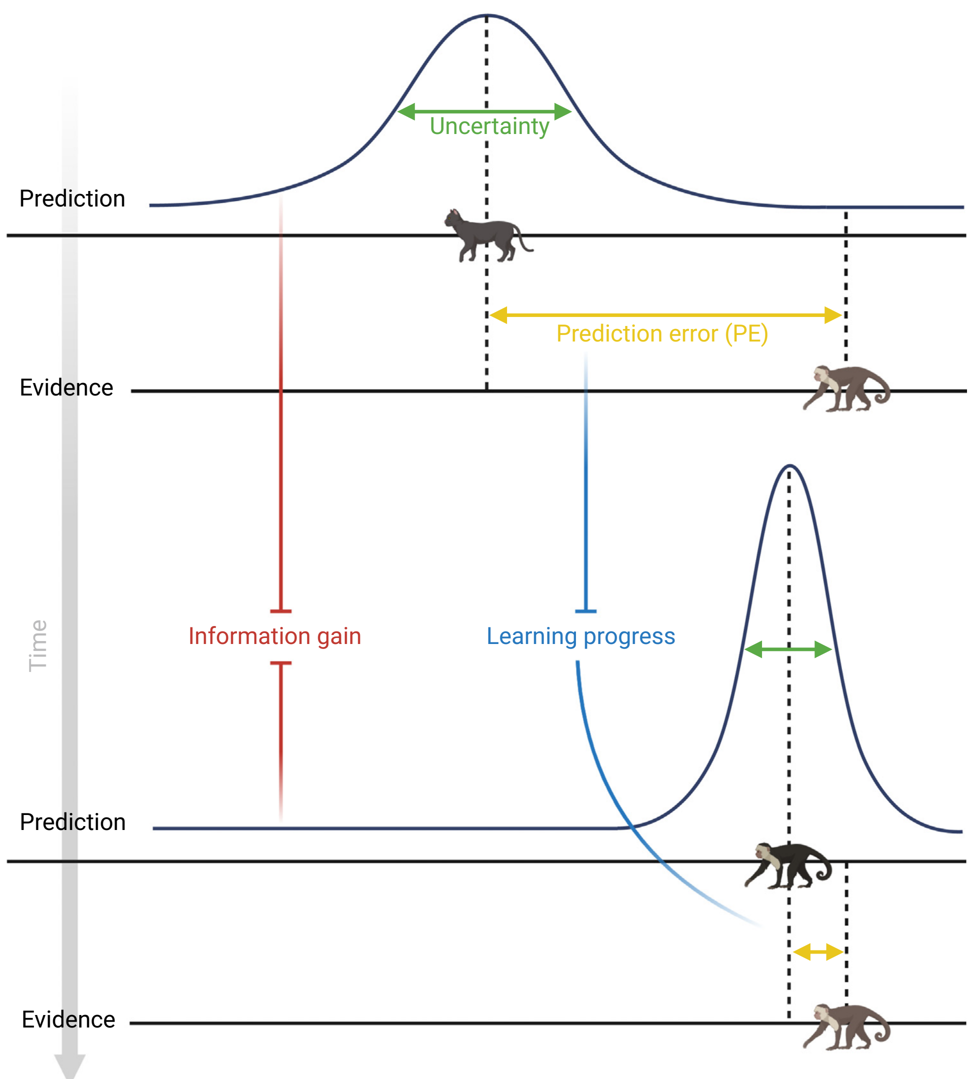


Figure 2. A schematic representation of the different factors underlying curiosity. Initially, an observer expects to see a cat with some degree of uncertainty (in green). However, they encounter a monkey instead, which leads to a prediction error (PE) (in yellow). Given the new evidence, the observer updates their internal model, now expecting to see a monkey. This expectation is confirmed by the incoming evidence (which is still the same monkey), thus leading to a small PE. The divergence between previous and current predictions (i.e., the difference in the location and width of the distributions) captures the information gain (in red). The rate of change in the PEs over time is a proxy of the learning progress (in blue). The observer's curiosity about the identity of the animal may be related to uncertainty, information gain, or learning progress.

We use the equations introduced in Boxes 1 and 2 to visualize the temporal dynamics of prediction error, uncertainty, information gain, and learning progress in a situation in which curiosity arises. As an example, we use the encounter with an animal, as introduced earlier. First, imagine the situation where after spotting an animal in the bushes, the animal slowly moves and becomes more visible (Figure 3A). From the initial belief that the animal is a cat, we update our internal model as increasingly more evidence becomes available. After 1 min, we are confident that we spotted a monkey (as indexed by the narrow probability distributions centered on the monkey in Figure 3C).

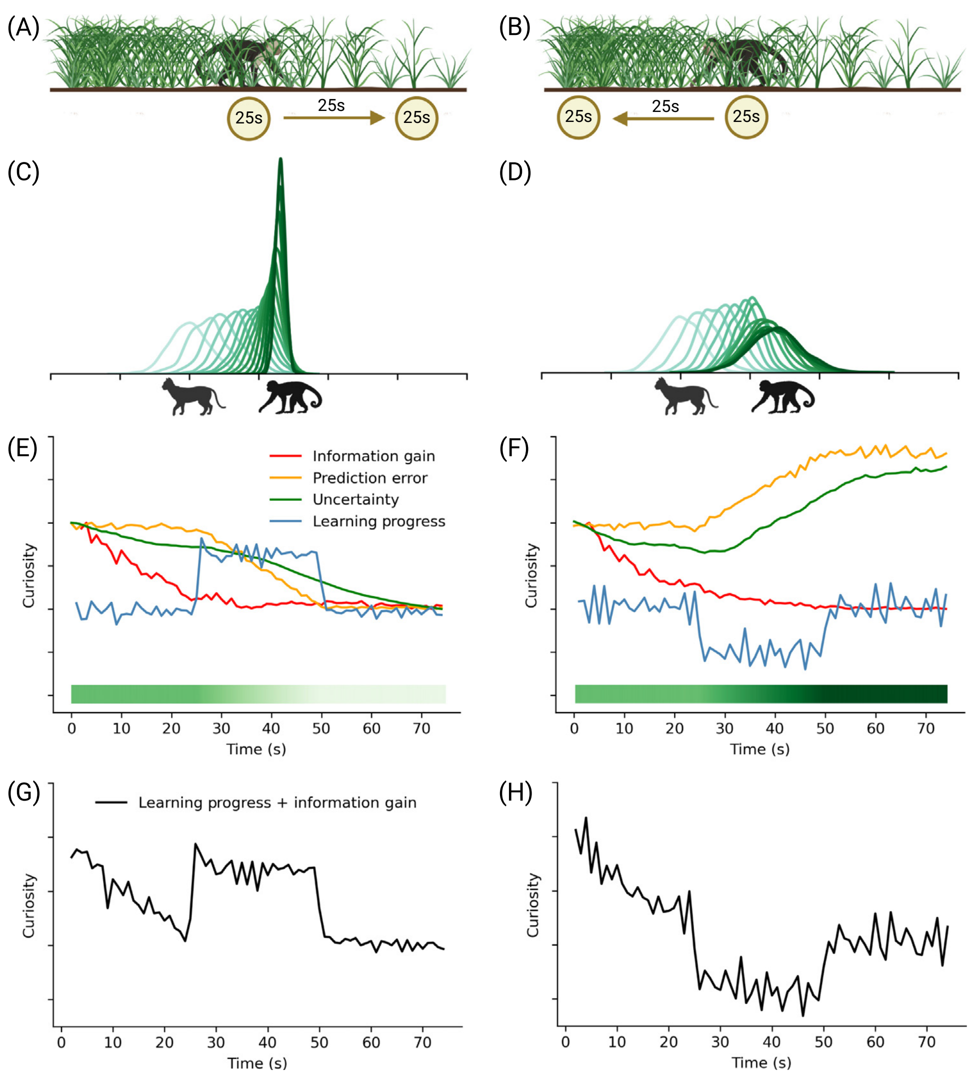


Figure 3. The temporal dynamics of curiosity. (A,B) An animal is hiding in the bushes for 25 s, then either starts moving out of the bushes (A) or deeper into the bushes (B) for 25 s. Finally, it stops moving. Computationally, the animal is represented in an arbitrary space as a vector of many features and every feature is learned independently through reinforcement learning. (C) As the animal gradually leaves the bushes, it delivers more information about its identity, which results in greater model change and greater confidence that the animal is a monkey. The bushes can be seen as Gaussian sensory noise that is added to the signal, where the level of the noise is proportional to the density of the bushes. (D) When the animal goes further into the bushes, it delivers less information about its identity, which results in reduced model change and reduced confidence in the identity of the animal. (E,F) The different accounts of curiosity are quantified using the equations in Boxes 1 and 2 to display their different temporal patterns. The green bar below the plots indicates the change in the level of noise across time. (G,H) Integrating learning progress and information gain allows to account for multiple aspects that are key to curiosity, such as increased interest when the animal is first spotted and persistence when learning is occurring.

To support optimal exploration, curiosity should be high at the beginning (when our model of the world is bad and we might have a chance to improve it) and when the animal starts moving (when our model of the world is finally getting better). As depicted in Figure 3E, uncertainty and information gain predict the initial high levels of curiosity, which is in line with optimal exploration, but not the latter increase. Conversely, learning progress predicts the latter increase, but not the former. Hence, all mechanisms are not fully optimal in promoting exploration, but all have unique strengths.

In an alternative situation, after initially being spotted, the animal starts moving deeper into the bushes (Figure 3B). We might realize it is unlikely to be a cat, but our certainty on whether it is truly a monkey will remain low (as depicted by wide distributions in Figure 3D). This results in an inability to optimally improve our internal model of the animal's identity. Contrary to the previous example, if curiosity was to support optimal exploration, it should peak at the start, but abruptly drop when the challenge of determining the animal's identity becomes impossible. As depicted in Figure 3F, information gain and uncertainty are again high at the start, thus capturing the initial curiosity that is instrumental to optimal exploration. However, only the values of learning progress quickly drop when the animal leaves, which captures a sudden reduction in curiosity that aids optimal exploration. Information gain reduces more slowly and uncertainty even keeps increasing. Hence, seeking uncertainty would lead us to invest resources in an unsolvable situation, which is in total contrast to optimal exploration.

### Integrating information gain and learning progress

Each account of curiosity has important limitations if taken in isolation. Following uncertainty can cause agents to become mired in unlearnable situations, seeking information gains fails to account for whether actual learning can occur, and focusing on learning progress does not satisfactorily address the initial state of curiosity. Hence, adhering to only one of these mechanisms will lead to suboptimal exploration. For example, only valuing information gain may lead to disengagement from activities where learning occurs but information is gained slowly, preferring the instantaneous acquisition of a different piece of knowledge instead. Conversely, only seeking learning progress might lead an agent to miss out on activities that are very informative to them, just because some other aspects of the world allow for greater learning.

Yet, merging the drives towards information gain and learning progress enables us to capture all essential aspects of curiosity as a tool for optimal exploration. Notably, an initial surge in curiosity, prompted by the need to gain information, is followed by a decrease. This diminishing curiosity becomes even more pronounced when learning is impossible (Figure 3H), yet it reverses temporarily when learning occurs (Figure 3G). These opposing temporal dynamics successfully capture the fact that engagement should remain high when learning takes place and decrease otherwise.

With this integration, the predictive processing framework is enriched by introducing the temporal dynamics of active learning, as a drive towards learning progress endows agents with a simple but powerful heuristic, which allows them to avoid wasting resources on situations where learning is not occurring. Moreover, this integration provides a 'common currency' for the value of exploration, underscoring the shared behavioral outcome of these diverse accounts: exploration. Establishing this common currency simplifies the calculation of the trade-offs with competing motivations and drives, such as avoiding danger or seeking rewards [75–77], thereby offering an efficient way to determine the optimal course of action.

Finally, the current work can approximate the actual computations that the brain may be performing to instantiate curiosity, moving away from vague and elusive definitions on one side [78] and from complex but biologically implausible implementations on the other. For example, previous computational work implemented curiosity as expected information gain [79], where thousands of hypothetical prospective scenarios are simulated and the one that is expected to minimize uncertainty the most is picked. Although effective, it is unlikely that the human brain implements such an effortful and costly strategy. Conversely, initial work relates the concepts presented in this article to their neural underpinnings [80,81]. For example, a subset of neurons in a network encompassing the anterior cingulate cortex, inferior dorsal striatum, and pallidum are involved in computing the trade-off between the utility of information and its accessibility, contingently prompting gaze behavior to gather this information [80].

In summary, understanding the mechanisms underlying curiosity is essential for comprehending human learning and exploration. By identifying and integrating various accounts of curiosity, we have developed a more precise model of human behavior. Next, we address how this improved model allows us to make more accurate predictions about how individuals explore their environment.

## From optimal to human exploration

In this article, we first reviewed empirical research on the cognitive mechanisms that have been proposed to underlie curiosity, highlighting their strengths and weaknesses. Then, we assessed these mechanisms in terms of their utility for optimal exploration. This led us to conclude that a mechanism integrating both learning progress and information gain would support optimal exploration. What remains to be determined is the extent to which human curiosity-driven exploration aligns with optimality [82,83].

Our review indicates that curiosity likely is the result of multiple underlying mechanisms. In this perspective, the extent to which each mechanism impacts behavior may differ between individuals. While some individuals might approximate optimality, others might be located in other portions of a ‘curiosity space’ that extends along the continuum of information gain, learning progress, and uncertainty (Figure 4). This space also allows for suboptimal curiosity mechanisms, which are more (e.g., information maximization) or less (e.g., information avoidance) efficient in driving exploration.

In line with this idea of curiosity space, existing empirical findings have reported individual differences in the extent to which information, learning progress, and uncertainty impact curiosity. For instance, some people were found to prefer acquiring new information [19] even to the point of incurring costs to obtain it [26,34], while others might even avoid information [84]. Similarly, while some individuals might seek learning progress [5,6], others might not [85,86]. Moreover, uncertainty is sometimes sought [22,86,87], but often avoided [88,89], especially by individuals with anxious traits [90]. Finally, irrespective of where they are situated in this curiosity space, individuals might vary in how strongly they experience curiosity [55,91]. For example, the personality trait of impulsivity has been shown to overlap with curiosity both behaviorally and in its neural substrates [55]. This might imply that enhanced impulsivity could lead to a heightened experience of curiosity. Focusing on individual differences is fundamental to come to an understanding of these complex interactions between multiple intertwined cognitive processes, which may be confounded at the group level [92].

Recent work started to develop model-based approaches to study individual differences in curiosity in adults [20,86], children [48,85,93], and infants [94,95]. Although these initial studies are promising, the end goal should be the successful measurement of stable individual differences in curiosity, which in turn would allow the prediction of people’s real-life choices and behaviors. Relatedly, we have considered the unfolding of curiosity on relatively short timescales. Recent work is starting to explore how these curiosity mechanisms operate in real-world contexts [96,97] and their functioning in longer time horizons or more complex environments is yet to be determined.

## Concluding remarks

In the current article, we defined curiosity as motivated by the goal of obtaining knowledge and as motivated by the learning process itself. Seeking uncertainty as well as maximizing information

### Outstanding questions

What is the exact interaction between information gain and learning progress in shaping curiosity? Does their importance differ across individuals?

How does individual variability in curiosity relate to differences in brain function and connectivity?

Do the cognitive mechanisms underlying curiosity change across phylogenesis and ontogenesis? Do they move from simpler but imperfect mechanisms to more optimal ones?

How does curiosity interact with other cognitive and affective processes, such as interest, reward processing, cognitive control, and emotion regulation?

Do the mechanisms underlying curiosity vary across clinical groups? What are the potential clinical implications of understanding curiosity and how can they be harnessed to enhance cognitive abilities and promote well-being?

Can we develop interventions or training programs to enhance curiosity? How can we leverage our understanding of the neural mechanisms of curiosity to promote lifelong learning and intellectual growth?

How can insights from cognitive neuroscience inform the design of educational strategies and environments that foster curiosity and optimal learning outcomes?

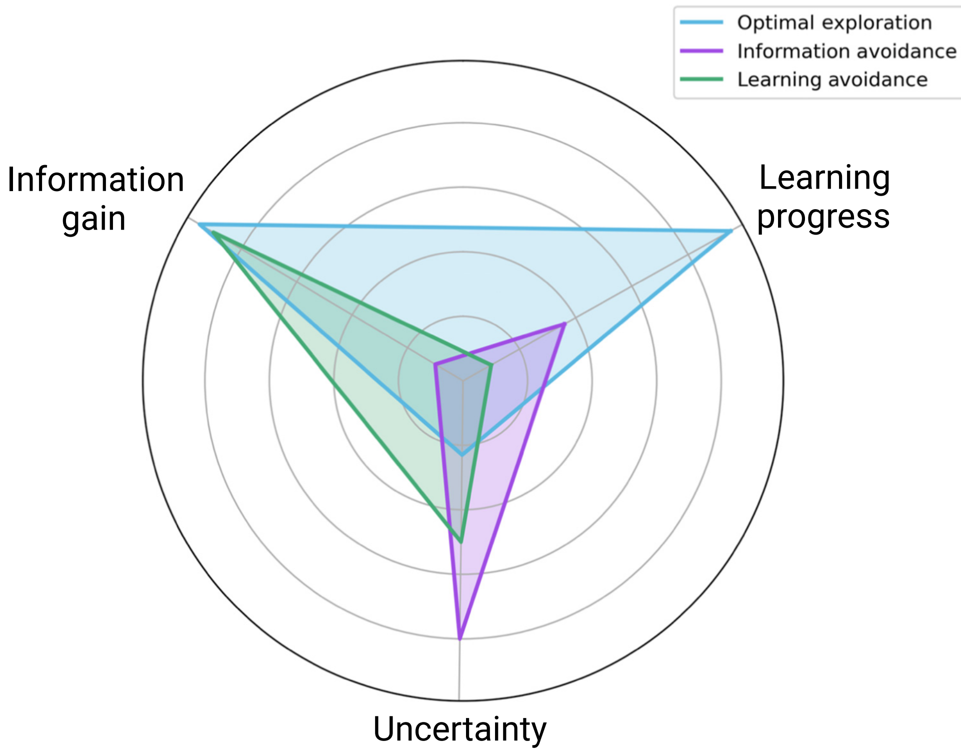


Figure 4. The curiosity space. Optimal exploration balances the maximization of information gain and learning progress (in blue). The specific weight of different curiosity mechanisms may vary across individuals, such that some individuals might be less inclined to gather new information (in purple). Conversely, other individuals may have a strong preference for obtaining instant information, taking into account, to a lesser extent, the aspect of learning progress (and the effort that comes with it) (in green).

gain or learning progress have all been proposed as different mechanisms to reach the same goal, namely a better understanding of the world around us. A comparison between these different mechanisms revealed fundamental differences in terms of their temporal dynamics and their utility for optimal exploration. This led us to conclude that only the concurrent tracking of information gain and learning progress offers a promising path for capturing curiosity. This theoretical work will allow us to test more precise hypotheses regarding the mechanisms underlying human curiosity. Although recent studies have advanced our understanding of these mechanisms, future work needs to increase its focus on the biological substrates that explain how curiosity originates and on the individual differences in the cognitive mechanisms that underlie curious behavior (see Outstanding questions). These advances will allow us to test biologically plausible theories of curiosity and to eventually explain and predict human behavior in the real world.

## Acknowledgments

This study was supported by a Donders Centre for Cognition internal grant to S.H. and R.B.M. ('Here's looking at you, kid.' A model-based approach to interindividual differences in infants' looking behavior and their relationship with cognitive performance and IQ; award/start date: 15 March 2018), a VICI grant from the Netherland Organization for Scientific Research NWO to S.H. ('Loving to learn - how curiosity drives cognitive development in young children'; serial number: VI.C.191.022), a Wellcome Trust center grant to R.B.M. ('Wellcome Centre for Integrative Neuroimaging'; serial number: 203139/Z/16/Z), an EPA Cephalosporin Fund and Biotechnology and Biological Sciences Research Council to R.B.M. (BB/N019814/1), a Netherlands Organization for Scientific Research VENI grant to R.B.M. (serial number: 016.195.197), and an Early Career Award of the Royal Netherlands Academy of Arts and Sciences to R.B.M.

## Declaration of interests

The authors declare no competing interests.

## References


1. Lau, J.K.L. *et al.* (2020) Shared striatal activity in decisions to satisfy curiosity and hunger at the risk of electric shocks. *Nat. Hum. Behav.* 4, 531–543
2. Dezza, I.C. *et al.* (2022) Independent and interacting value systems for reward and information in the human brain. *eLife* 11, e66358
3. Gruber, M.J. *et al.* (2014) States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron* 84, 486–496
4. Kang, M.J. *et al.* (2009) The wick in the candle of learning: epistemic curiosity activates reward circuitry and enhances memory. *Psychol. Sci.* 20, 963–973
5. Poli, F. *et al.* (2022) Contributions of expected learning progress and perceptual novelty to curiosity-driven exploration. *Cognition* 225, 105119
6. Ten, A. *et al.* (2021) Humans monitor learning progress in curiosity-driven exploration. *Nat. Commun.* 12, 5972
7. Coddington, L.T. *et al.* (2023) Mesolimbic dopamine adapts the rate of learning from action. *Nature* 614, 294–302
8. Gruber, M.J. and Ranganath, C. (2019) How curiosity enhances hippocampus-dependent memory: the prediction, appraisal, curiosity, and exploration (PACE) framework. *Trends Cogn. Sci.* 23, 1014–1025
9. Mazzaglia, P. *et al.* (2022) Curiosity-driven exploration via latent Bayesian surprise. In *Proceedings of the AAAI Conference on Artificial Intelligence* (36), pp. 7752–7760
10. Murayama, K. *et al.* (2019) Process account of curiosity and interest: a reward-learning perspective. *Educ. Psychol. Rev.* 31, 875–895
11. Kashdan, T.B. *et al.* (2013) How are curious people viewed and how do they behave in social situations? From the perspectives of self, friends, parents, and unacquainted observers. *J. Pers.* 81, 142–154
12. Wagstaff, M.F. *et al.* (2021) Measures of curiosity: a literature review. *Hum. Resour. Dev. Q.* 32, 363–389
13. Schweitzer, V.M. *et al.* (2023) (Don't) mind the gap? Information gaps compound curiosity yet also feed frustration at work. *Organ. Behav. Hum. Decis. Process.* 178, 104276
14. Clark, A. (2013) Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav. Brain Sci.* 36, 181–204
15. Egner, T. *et al.* (2010) Expectation and surprise determine neural population responses in the ventral visual stream. *J. Neurosci.* 30, 16601–16608
16. Kwisthout, J. *et al.* (2017) To be precise, the details don't matter: on predictive processing, precision, and level of detail of predictions. *Brain Cogn.* 112, 84–91
17. Friston, K. (2018) Does predictive coding have a future? *Nat. Neurosci.* 21, 1019–1021
18. O'Reilly, J.X. *et al.* (2012) How can a Bayesian approach inform neuroscience? *Eur. J. Neurosci.* 35, 1169–1179
19. Kobayashi, K. *et al.* (2019) Diverse motives for human curiosity. *Nat. Hum. Behav.* 3, 587–595
20. Kelly, C.A. and Sharot, T. (2021) Individual differences in information-seeking. *Nat. Commun.* 12, 7062
21. Berlyne, D.E. (1962) Uncertainty and epistemic curiosity. *Br. J. Psychol.* 53, 27–34
22. van Lieshout, L.L.F. *et al.* (2018) Induction and relief of curiosity elicit parietal and frontal activity. *J. Neurosci.* 38, 2579–2588
23. van Lieshout, L.L.F. *et al.* (2021) Curiosity or savouring? Information seeking is modulated by both uncertainty and valence. *PLoS One* 16, e0257011
24. Baer, C. and Kidd, C. (2022) Learning with certainty in childhood. *Trends Cogn. Sci.* 26, 887–896
25. Baranes, A. *et al.* (2015) Eye movements reveal epistemic curiosity in human observers. *Vision Res.* 117, 81–90
26. Iigaya, K. *et al.* (2016) The modulation of savouring by prediction error and its effects on choice. *eLife* 5, e13747
27. Silvia, P.J. (2005) What is interesting? Exploring the appraisal structure of interest. *Emotion* 5, 89
28. Cockburn, J. *et al.* (2022) Novelty and uncertainty regulate the balance between exploration and exploitation through distinct mechanisms in the human brain. *Neuron* 110, 2691–2702
29. Vogl, E. *et al.* (2020) Surprised–curious–confused: epistemic emotions and knowledge exploration. *Emotion* 20, 625–641
30. Kidd, C. *et al.* (2012) The Goldilocks effect: human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLoS One* 7, e36399
31. Poli, F. *et al.* (2020) Infants tailor their attention to maximize learning. *Sci. Adv.* 6, 5053–5076
32. Lamnina, M. and Chase, C.C. (2019) Developing a thirst for knowledge: how uncertainty in the classroom influences curiosity, affect, learning, and transfer. *Contemp. Educ. Psychol.* 59, 101785
33. Gottlieb, J. *et al.* (2013) Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends Cogn. Sci.* 17, 585–593
34. Rodriguez Cabrero, J.A.M. *et al.* (2019) Costly curiosity: people pay a price to resolve an uncertain gamble early. *Behav. Processes* 160, 20–25
35. Schwartenbeck, P. *et al.* (2019) Computational mechanisms of curiosity and goal-directed exploration. *eLife* 8, e41703
36. Friston, K.J. *et al.* (2017) Active inference, curiosity and insight. *Neural Comput.* 29, 2633–2683
37. Goupil, L. and Proust, J. (2023) Curiosity as a metacognitive feeling. *Cognition* 231, 105325
38. Marvin, C.B. and Shohamy, D. (2016) Curiosity and reward: valence predicts choice and information prediction errors enhance learning. *J. Exp. Psychol. Gen.* 145, 266–272
39. Wade, S. and Kidd, C. (2019) The role of prior knowledge and curiosity in learning. *Psychon. Bull. Rev.* 26, 1377–1387
40. Horan, M. *et al.* (2019) Parietal neurons encode information sampling based on decision uncertainty. *Nat. Neurosci.* 22, 1327–1335
41. Bennett, D. *et al.* (2016) Intrinsic valuation of information in decision making under uncertainty. *PLoS Comput. Biol.* 12, e1005020
42. Wang, M.Z. and Hayden, B.Y. (2019) Monkeys are curious about counterfactual outcomes. *Cognition* 189, 1–10
43. Golman, R. *et al.* (2022) The demand for, and avoidance of, information. *Manag. Sci.* 68, 6454–6476
44. van Lieshout, L.L. *et al.* (2019) Motives underlying human curiosity. *Nat. Hum. Behav.* 3, 550–551
45. Oosterwijk, S. (2017) Choosing the negative: a behavioral demonstration of morbid curiosity. *PLoS One* 12, e0178399
46. FitzGibbon, L. *et al.* (2021) The lure of counterfactual curiosity: people incur a cost to experience regret. *Psychol. Sci.* 32, 241–255
47. Hsee, C.K. and Ruan, B. (2016) The Pandora effect: the power and peril of curiosity. *Psychol. Sci.* 27, 659–666
48. Liquin, E.G. *et al.* (2021) Developmental change in what elicits curiosity. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (43), pp. 1360–1366
49. Aguirre, M. *et al.* (2022) Knowledge in sight: toddlers plan efficient epistemic actions by anticipating learning gains. *Cogn. Sci.* 46, e13103
50. Bazhydai, M. *et al.* (2020) "I don't know but I know who to ask": 12-month-olds actively seek information from knowledgeable adults. *Dev. Sci.* 23, e12938
51. Cao, A. *et al.* (2023) Habituation reflects optimal exploration over noisy perceptual samples. *Top. Cogn. Sci.* 15, 290–302
52. Sutton, R.S. and Barto, A.G. (2018) *Reinforcement Learning: An Introduction* (2nd edn), MIT Press
53. Sajid, N. *et al.* (2021) Active inference: demystified and compared. *Neural Comput.* 33, 674–712
54. Cheyette, S.J. *et al.* (2023) People seek easily interpretable information. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (45), pp. 913–919
55. Marvin, C.B. *et al.* (2020) Curiosity as the impulse to know: common behavioral and neural mechanisms underlying curiosity and impulsivity. *Curr. Opin. Behav. Sci.* 35, 92–98

56. Molnar, A. and Golman, R. (2023) Impatience for information: curiosity is here today, gone tomorrow. *J. Behav. Decis. Mak.* 37, e2360
57. Gottlieb, J. and Oudeyer, P.-Y. (2018) Towards a neuroscience of active sampling and curiosity. *Nat. Rev. Neurosci.* 19, 758–770
58. Twomey, K.E. and Westermann, G. (2018) Curiosity-based learning in infants: a neurocomputational approach. *Dev. Sci.* 21, e12629
59. Kaplan, F. and Oudeyer, P.-Y. (2004) Maximizing learning progress: an internal reward system for development. In *Embodied Artificial Intelligence* (Iiuda, F. *et al.*, eds), pp. 259–270, Springer
60. Oudeyer, P.-Y. and Kaplan, F. (2007) What is intrinsic motivation? A typology of computational approaches. *Front. Neurorobotics* 1, 6
61. Kidd, C. *et al.* (2014) The Goldilocks effect in infant auditory attention. *Child Dev.* 85, 1795–1804
62. Singh, S. *et al.* (2010) Intrinsically motivated reinforcement learning: an evolutionary perspective. *IEEE Trans. Auton. Ment. Dev.* 2, 70–82
63. Singh, S. *et al.* (2009) Where do rewards come from. In *Proceedings of the Annual Conference of the Cognitive Science Society* (31), pp. 2601–2606
64. Szumowska, E. and Kruglanski, A.W. (2020) Curiosity as end and means. *Curr. Opin. Behav. Sci.* 35, 35–39
65. Luciw, M. *et al.* (2013) An intrinsic value system for developing multiple invariant representations with incremental slowness learning. *Front. Neurorobotics* 7, 9
66. Schultz, W. *et al.* (1997) A neural substrate of prediction and reward. *Science* 275, 1593–1599
67. Forestier, S. *et al.* (2022) Intrinsically motivated goal exploration processes with automatic curriculum learning. *J. Mach. Learn. Res.* 23, 6818–6858
68. Lopes, M. *et al.* (2012) Exploration in model-based reinforcement learning by empirically estimating learning progress. In *Proceedings of the 25th International Conference on Neural Information Processing Systems* (1), pp. 206–214
69. Oudeyer, P.Y. *et al.* (2007) Intrinsic motivation systems for autonomous mental development. *IEEE Trans. Evol. Comput.* 11, 265–286
70. Sayalı, C. *et al.* (2023) Learning progress mediates the link between cognitive effort and task engagement. *Cognition* 236, 105418
71. Sakaki, M. *et al.* (2023) Motivated for near impossibility: how task type and reward modulate task enjoyment and the striatal activation for extremely difficult task. *Cogn. Affect. Behav. Neurosci.* 23, 30–41
72. Molinaro, G. and Collins, A.G.E. (2023) A goal-centric outlook on learning. *Trends Cogn. Sci.* 27, 1150–1164
73. Friston, K. *et al.* (2012) Free-energy minimization and the dark-room problem. *Front. Psychol.* 3, 130
74. Dubey, R. and Griffiths, T.L. (2020) Reconciling novelty and complexity through a rational analysis of curiosity. *Psychol. Rev.* 127, 455–476
75. Karayanni, M. and Nelken, I. (2022) Extrinsic rewards, intrinsic rewards, and non-optimal behavior. *J. Comput. Neurosci.* 50, 139–143
76. Leong, Y.C. *et al.* (2019) Neurocomputational mechanisms underlying motivated seeing. *Nat. Hum. Behav.* 3, 962–973
77. Juechems, K. *et al.* (2019) A network for computing value equilibrium in the human medial prefrontal cortex. *Neuron* 101, 977–987
78. Jirout, J. and Klahr, D. (2012) Children's scientific curiosity: in search of an operational definition of an elusive concept. *Dev. Rev.* 32, 125–160
79. Burda, Y. *et al.* (2018) Large-scale study of curiosity-driven learning. *arXiv*, Published online August 13, 2018. https://doi.org/10.48550/arXiv.1808.04355
80. Jezzini, A. *et al.* (2021) A prefrontal network integrates preferences for advance information about uncertain rewards and punishments. *Neuron* 109, 2339–2352
81. White, J.K. *et al.* (2019) A neural network for information seeking. *Nat. Commun.* 10, 5168
82. Lieder, F. and Griffiths, T.L. (2020) Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. *Behav. Brain Sci.* 43, e1
83. Modirshanechi, A. *et al.* (2023) The curse of optimism: a persistent distraction by novelty. *bioRxiv*, Published online July 7, 2023. https://doi.org/10.1101/2022.07.05.498835
84. Brouwers, M.C. and Sorrentino, R.M. (1993) Uncertainty orientation and protection motivation theory: the role of individual differences in health compliance. *J. Pers. Soc. Psychol.* 65, 102
85. Poli, F. *et al.* (2023) Exploration in 4-year-old children is guided by learning progress and novelty. *PsyArXiv*, Published online August 16, 2023. https://doi.org/10.31234/osf.io/f7u85
86. Poli, F. *et al.* (2023) Autistic traits foster effective curiosity-driven exploration. *PsyArXiv*, Published online September 13, 2023. https://doi.org/10.31234/osf.io/jnfdw
87. Zhu, J.-Q. *et al.* (2017) Information seeking as chasing anticipated prediction errors. In *Proceedings of the 39th Annual Meeting of the Cognitive Science Society*
88. Rosen, N.O. and Knäuper, B. (2009) A little uncertainty foes a long way: state and trait differences in uncertainty interact to increase information seeking but also increase worry. *Health Commun.* 24, 228–238
89. Grupe, D.W. and Nitschke, J.B. (2013) Uncertainty and anticipation in anxiety: an integrated neurobiological and psychological perspective. *Nat. Rev. Neurosci.* 14, 488–501
90. Lawrance, E.L. *et al.* (2022) The computational and neural substrates of ambiguity avoidance in anxiety. *Comput. Psychiatry Camb. Mass* 6, 8–33
91. Saragosa-Harris, N.M. *et al.* (2022) Real-world exploration increases across adolescence and relates to affect, risk taking, and social connectivity. *Psychol. Sci.* 33, 1664–1679
92. Kanai, R. and Rees, G. (2011) The structural basis of inter-individual differences in human behaviour and cognition. *Nat. Rev. Neurosci.* 12, 231–242
93. Molinaro, G. *et al.* (2023) Multifaceted information-seeking motives in children. *Nat. Commun.* 14, 5505
94. Poli, F. *et al.* (2023) Individual differences in processing speed and curiosity explain infant habituation and dishabituation performance. *Dev. Sci.*, Published online December 28, 2023. https://doi.org/10.1111/desc.13460
95. Altmann, E.C. *et al.* (2023) Curious choices: infants' moment-to-moment information sampling is driven by their exploration history. *PsyArXiv*, Published online November 27, 2023. https://doi.org/10.31234/osf.io/7d28k
96. Dubey, R. *et al.* (2021) Curiosity is contagious: a social influence intervention to induce curiosity. *Cogn. Sci.* 45, e12937
97. Abdelghani, R. *et al.* (2023) GPT-3-driven pedagogical agents to train children's curious question-asking skills. *Int. J. Artif. Intell. Educ.*, Published online June 30, 2023. https://doi.org/10.1007/s40593-023-00340-7
98. Kullback, S. and Leibler, R.A. (1951) On information and sufficiency. *Ann. Math. Stat.* 22, 79–86
99. Schmidhuber, J. (2009) Driven by compression progress: a simple principle explains essential aspects of subjective beauty, novelty, surprise, interestingness, attention, curiosity, creativity, art, science, music, jokes. In *Anticipatory Behavior in Adaptive Learning Systems*, pp. 48–76, Springer
100. Schillaci, G. *et al.* (2020) Tracking emotions: intrinsic motivation grounded on multi-level prediction error dynamics. In *2020 Joint IEEE 10th International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, pp. 1–8, IEEE
101. Gottlieb, J. *et al.* (2016) Motivated cognition: neural and computational mechanisms of curiosity, attention, and intrinsic motivation. *Adv. Motiv. Achiev.* 19, 149–172