



Understanding exploration in humans and machines by formalizing the function of curiosity[☆]

Rachit Dubey and Thomas L Griffiths

Recent work in machine learning has demonstrated the benefits of providing artificial agents with a sense of curiosity — a form of intrinsic reward that supports exploration. Two strategies have emerged for defining these rewards: favoring novelty and pursuing prediction errors. Psychological theories of curiosity have also emphasized these two factors. We show how these two literatures can be connected by understanding the function of curiosity, which requires thinking about the abstract computational problem that both humans and machines face as they explore their world.

Address

Princeton University, USA

Corresponding author: Dubey, Rachit (rdubey@princeton.edu)

Current Opinion in Behavioral Sciences 2020, 35:118–124

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

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

For a complete overview see the [Issue](#) and the [Editorial](#)

Available online 5th October 2020

<https://doi.org/10.1016/j.cobeha.2020.07.008>

2352-1546/© 2020 Elsevier Ltd. All rights reserved.

Introduction

The last few years have seen significant progress in developing machine learning systems that approach or exceed human performance in specific tasks [1–4]. One of the areas in which substantial progress has been made is that of reinforcement learning [5,6]. In reinforcement learning, agents learn what actions to take based on rewards provided by their environment, with the goal of maximizing long-term reward [7]. Recent progress in reinforcement learning has included developing systems that are capable of learning to play computer games far better than humans (where the reward is the score in the game) and developing algorithms for robotic systems that are capable of learning how to move and perform basic tasks (where the reward reflects progress in these tasks).

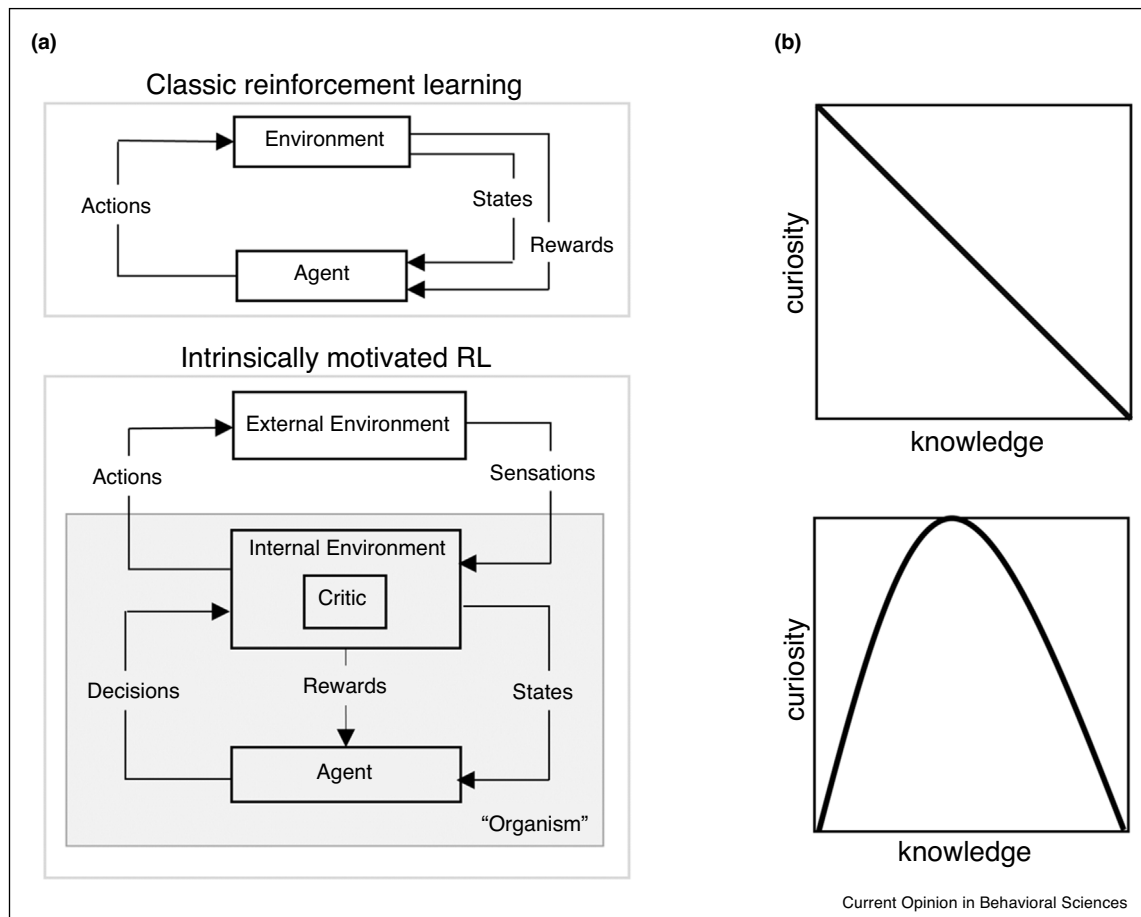
While the classic formulation of reinforcement learning focuses on problems that are defined in terms of *extrinsic* rewards (i.e. rewards provided by the designer or the environment to the agent), researchers quickly realized that effective algorithms for reinforcement learning often need to make use of *intrinsic* rewards (i.e. rewards generated by the agent itself [8*,9,10]; refer to [Figure 1a](#)). In particular, one of the key aspects of solving a reinforcement learning problem is effectively exploring the space of possible actions. This may require taking actions that currently seem less promising than actions that have already been found to be effective — the classic ‘explore/exploit’ problem. Early research in reinforcement learning discovered that appropriate exploration could be supported by providing agents with a reward for taking actions that led to parts of the space that had not previously been visited [11].

Recently, the insight that agents need appropriately structured intrinsic rewards for effective exploration has been incorporated into state-of-the-art reinforcement learning systems, resulting in significant improvements in performance [12,13,14**,15,16]. In some cases, it is even possible to define reinforcement learning systems that operate entirely in the absence of extrinsic rewards, with intrinsic rewards being sufficient to drive the agents to discover effective strategies [17,18]. The two forms of intrinsic rewards that have been found to be effective are one that focuses on rewarding novelty [19*,13], consistent with the earlier work mentioned above, and one that focuses on finding settings where the agent’s model of the world is wrong [20,12,14**]. These strategies both motivate uncertainty reduction, but do so with two different intrinsic rewards: either for actions that explore states with the greatest absolute uncertainty (novelty), or for actions that cause the greatest expected reduction in uncertainty (prediction error).

The two quantities that have been found to provide effective intrinsic rewards for machines — novelty and prediction error — have also been postulated as playing a key role in curiosity for humans ([Figure 1b](#)). In this paper, we argue that this should not come as a surprise: the convergent evolution of these two literatures reflect the fact that both humans and machines face the same underlying computational problem. By considering the function of curiosity, we can see why both of these factors are important and make predictions about when one is more important than the other. We begin by reviewing recent work on curiosity in machine learning before

[☆] This work was partially supported by grant number 61454 from the John Templeton Foundation.

Figure 1



(a) Intrinsically motivated reinforcement learning, adapted from [9]. While the standard RL view is that all rewards are extrinsic (top), intrinsically motivated RL considers the settings in which all rewards are intrinsically generated by the agent (bottom). Under this view, curiosity is an example of an intrinsic reward that helps the agent explore despite the absence of extrinsic rewards. (b) Curiosity in machine learning and psychology. In both psychology and machine learning, there are two broad classes of models of curiosity. Novelty-based models (top) posit that curiosity is a function of novelty (i.e. curiosity diminishes as knowledge increases). Prediction-error-based models (bottom) posit that curiosity is highest when predictions are always improving (i.e. curiosity is highest for moderately complex stimuli).

turning to psychological theories of curiosity and then considering the function of curiosity and its implications for both of these literatures.

Curiosity in modern machine learning

At the highest level, reinforcement learning (RL) solves control problems by identifying a *policy* that indicates how an agent should act in the world to maximize expected future rewards [7]. Recently, deep learning — representing policies in the weights of large artificial neural networks that are then learned through experience — has made it possible to scale up RL to problems that were previously intractable. One of the earliest success stories of deep RL was the Deep Q-Learning algorithm that solved a range of Atari games, some at a superhuman level [5]. This achievement was especially remarkable because this work was the first to demonstrate that deep learning

could be used to train RL agents to solve complex control tasks directly from raw image pixels. Following this, a large body of research has applied deep RL to achieve human-level performance on various other tasks [21,22]. However, these successes in RL require the reward function to be dense (i.e. with no long gaps between rewards) and to be well-aligned with the task. However, in many real-world settings, rewards extrinsic to the agent are either extremely sparse or absent altogether, thereby restricting the applications of RL to problems closer to the real-world.

An alternative to providing dense extrinsic rewards is to use ‘intrinsic’ rewards, that is, endow agents with the ability to generate rewards by themselves [10,9]. One popular form of such an intrinsic reward is endowing an agent with ‘curiosity’ — a self-generated reward that

enables the agent to continue exploring despite the absence of extrinsic rewards [20]. A number of algorithms have been proposed to endow curiosity to agents and they generally fall into two broad groups, being based on either novelty or prediction error.

Novelty-based algorithms pursue the idea that the brain is intrinsically motivated to seek out novel tasks. Generally, these algorithms estimate the distribution of the environmental state (by maintaining a ‘visitation count’ — a count of the number of times different states have been visited) and encourage the agent to explore by providing intrinsic rewards as an inverse function of the visitation count (i.e. by encouraging the agent to visit new states and discouraging the agent from revisiting the same states). Novelty-based algorithms have been successfully applied to solve various Atari games, navigation tasks, [23,19*,13,16,15] and have also gained popularity in the evolutionary computing literature where they have been shown to find solutions which algorithms relying on solely optimizing the fitness function could not find [24].

Algorithms based on prediction errors rely on the hypothesis that the brain is intrinsically motivated to pursue tasks in which one’s predictions are always improving [20,25,26,27*]. Algorithms that generate intrinsic rewards as a function of prediction errors predict the next state (given the current state and the executed action) and then minimize the error compared to the predicted state (after they execute the action to move to that state). As the agent explores more and the state predictions become more accurate, the error decreases leading to the agent becoming curious about other states. Since the rewards are generated based on how hard it is for the agent to predict the next state, prediction-error-based agents end up exploring the environment by pursuing tasks that are neither too easy or too difficult and instead by focusing on tasks on which their prediction errors are constantly reducing. This is because when the environmental complexity is too high for the agent, the agent would not improve its prediction error and such tasks would not generate rewards. When the environmental complexity is too low, then those tasks would be learned by the agent very quickly, leading to no further generation of intrinsic rewards (as the prediction error will quickly go down to zero). Curiosity thus seems to be a matter of finding the right balance so that the agent is constantly maximizing the rate of reducing prediction errors. Prediction-error-based algorithms have been also shown to be successful on a variety of RL tasks such as solving atari games and performing navigation [14*,18,28].

Psychological theories of curiosity

Some of the most influential studies in psychology on curiosity were conducted by the British and Canadian psychologist Daniel Berlyne. In one paper, Berlyne

conducted a series of experiments that demonstrated that rats spend more time exploring a novel stimulus compared to a familiar stimulus and as they become more familiar with the novel stimulus, they spend less time exploring it [29]. From these experiments, Berlyne proposed that novelty affects an organism’s receptor by making the organism more ‘curious’ and described novelty as a driving force that motivates an organism to seek out novel stimuli which diminishes as an organism gains more familiarity with a stimulus. Berlyne’s novelty-based theory of curiosity has since been supported by various empirical studies as well as neuroscientific studies that show that novel stimuli activate reward-responsive areas in the brain [30–33].

In another set of studies, Berlyne studied how stimulus complexity and incongruity affects curiosity. Through various experiments on human participants, Berlyne found that people spend more time exploring stimuli that are more complex or incongruous to them and that people are drawn towards optimally incongruent or surprising events [34,35]. These findings led to the development of an alternative perspective on curiosity which posits that curiosity is a function of stimulus complexity and is driven by stimuli that are neither too simple nor too complex [36,37]. A number of empirical studies have further supported this by showing that curiosity is an inverted U-shaped function of knowledge, with people showing the highest curiosity for topics for which they have moderate knowledge about [38,39]. This perspective is in line with prediction-error-based approaches as reducing prediction errors are similar to reducing incongruity. Researchers have also studied the influence of task difficulty on curiosity and shown that, similar to the effect of complexity, people spend most time exploring tasks that are moderately difficult [40,41]. More recently, studies have shown that stimuli that trigger prediction errors not only evoke greater curiosity, but are also remembered and recalled more better, suggesting possible pathways through which prediction errors and curiosity enhance memory encoding [42–45].

The function of curiosity

Both novelty-error-based and prediction-error-based approaches have found support in the machine learning as well as the psychology literature. However, in both these literatures, it is not yet clear under which situations which of the two distinct approaches are more advantageous. We recently [46**] showed that different psychological theories of curiosity can be reconciled by focusing on the abstract computational problem underlying curiosity and the form that the optimal solution to that problem should take. We believe that the same framework helps to understand why (and when) novelty and prediction-error are also effective intrinsic rewards for RL.

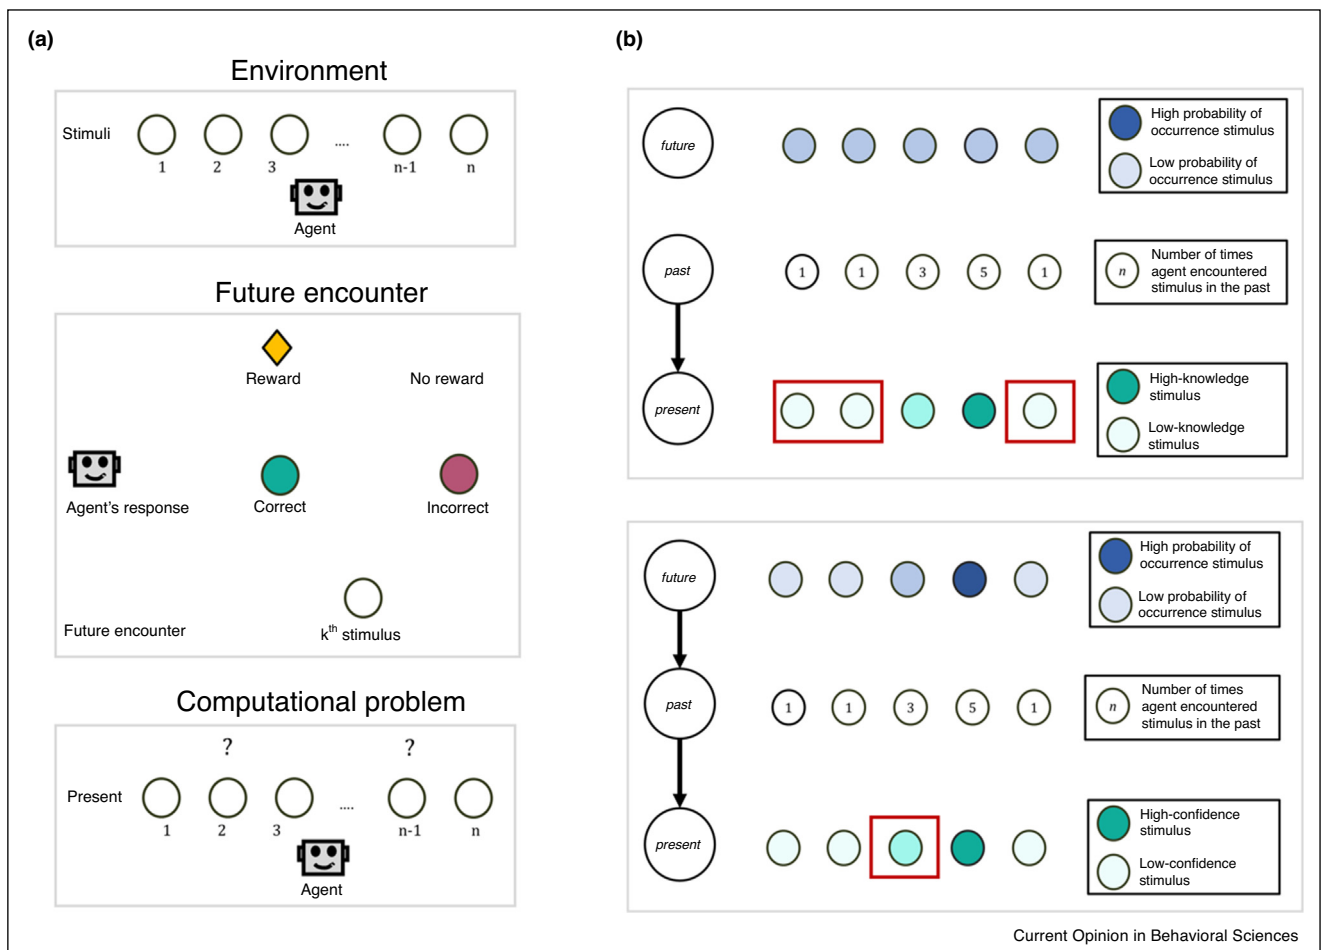
Our work posits that the function of curiosity is to help the agent take actions to maximally increase the usefulness of its knowledge. Depending on the environment, this problem can be solved by focusing on either novelty or prediction errors.

As an illustration, consider an agent in an environment where rewards are far into the future and are dependent on how the agent explores the environment in the present (also refer to Figure 2a). More specifically, with some probability, stimuli occur again in the future and the agent receives a reward upon producing the correct response (or equivalently the right action/s) to those stimuli in that future encounter (a stimulus here can be considered anything that has to be learnt or solved such as a puzzle,

game, task, etc.). Given this, the goal of the agent is to learn the correct response to as many stimuli as possible, as knowing more about the environment in the present maximizes its chances of obtaining rewards in the future. Because the agent has limited computational resources, it then faces the problem of deciding what to explore in the environment in the present to maximize future rewards.

The optimal solution to this problem posits that a rational agent should take actions that maximally increase it is the value of its knowledge, where this value is a function of how much the agent knows about the environment in the present and how probable it is to encounter various stimuli again in the future. For instance, if the agent comes across a stimulus that is completely novel, then the

Figure 2



(a) Computational problem underlying curiosity. *Top panel:* A toy environment with n stimuli and a reward-maximizing agent. *Middle panel:* In this environment, rewards can only be attained in the future. More specifically, some stimuli can occur again in the future (the k th stimulus in this example) and the agent receives a reward for producing the correct response (or equivalently the 'right' action) to that stimulus and no reward for producing the incorrect response to that stimulus. *Bottom panel:* Given this, the agent faces the problem of deciding which stimulus to explore in the present to maximize future rewards. **(b)** The causal environment determines when curiosity is driven by novelty or prediction errors. *Top panel:* An environment where future is independent of past and present. Here, it is optimal for the agent to be curious about novel stimuli in the present to maximize future rewards (highlighted in the red box, bottom row). *Bottom panel:* An environment where future is related to past and present. Here, it is optimal for the agent to be curious about moderately complex stimuli (highlighted in the red box, bottom row).

agent should explore that stimulus. However, if the agent believes that this stimulus will never occur again in the future, then it is perfectly rational for the agent to not be curious about this stimulus because knowing more about this stimulus does not help the agent to maximize its rewards in the future. From this perspective, the function of curiosity is to simply help the agent take actions to maximally increase its value of knowledge (or in other words prioritize exploration in the present to maximize rewards in the future).

One interesting implication of this analysis is that it shows that the causal structure of the environment determines the form that this optimal solution takes — whether curiosity is driven by novelty or prediction errors (refer to Figure 2b). More specifically, when the agent is in an environment where the past and future occurrences of stimuli are independent of each other, then it is rational for the agent to explore novel stimuli (akin to novelty based approaches). This is because, in this environment, all stimuli are equally likely to occur again in the future and the expected value is increased most rapidly by learning more about the stimuli for which the agent has least knowledge. On the other hand, when the agent is in an environment where past and future occurrences are related to each other, that is, the stimuli for which the agent has had more experience with are more likely to occur again in the future, then it is rational for the agent to explore moderately complex stimuli (akin to prediction-error approaches). This is because completely novel stimuli have little chance of occurring again in the future and therefore the expected value is increased the most by learning more about moderately complex stimuli (as they have a moderate chance of occurring again in the future). From this perspective, novelty-error-based and prediction-error-based algorithms are not competing approaches, but rather different solutions to the same computational problem albeit in different environmental structures.

This perspective provides insights relevant to both machine learning and psychology. First, it demonstrates that both novelty-error-based and prediction-error-based approaches have at their heart the same computational problem. This could be especially relevant to artificial intelligence, as it suggests under which environmental settings novelty-based techniques should be preferred over prediction-error techniques. Second, this analysis shows that in addition to current knowledge, curiosity is also sensitive to future probability of occurrence, implying that people are not only sensitive to the match or mismatch between a given stimulus and their current or former beliefs, but also how they expect the resolution of uncertainty to affect the future. This is relevant to both human and artificial curiosity. For researchers in artificial intelligence, it suggests incorporating a form of ‘usefulness’ estimation to the agents, wherein the agents

prioritize their exploration based on which states are more likely to be encountered again. This would be especially valuable in domains where the state space is too large. For researchers in psychology, this offers new insights towards understanding maladaptive curiosity — because it is optimal to have a drive to explore stimuli that increase our value of knowledge, simple miscalculations of this value can lead people (and even animals) to become suboptimally curious about things that may not have any real-world advantage [47–50] and develop a preference for advance information-seeking even if it is costly to do so [51–55]. Subsequently, this also offers new ways to think about how curiosity in various learning settings can be directed towards relevant stimuli with strategic value-based interventions [56–58].

Conclusion

In situations where humans and machines face the same computational problems, we can expect that they will converge on similar solutions. We have argued that this is the case for curiosity: both humans and machines are seeking to maximize the value of the knowledge they acquire about the world. Recognizing this shared problem helps to explain why novelty and prediction-error are important determinants of intrinsic rewards for both systems, and when it is most appropriate to focus on each. We anticipate that analyzing other cases where humans and machines face the same computational problems — and recognizing when they do not — will be a valuable tool for both artificial intelligence and cognitive science as these fields move forward.

Conflict of interest statement

Nothing declared.

References and recommended reading

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

- of special interest
 - of special interest
1. Krizhevsky A, Sutskever I, Hinton GE: **Imagenet classification with deep convolutional neural networks**. *NIPS* 2012.
 2. Ciregan D, Meier U, Schmidhuber J: **Multi-column deep neural networks for image classification**. *2012 IEEE Conference on Computer Vision and Pattern Recognition; IEEE: 2012:3642-3649*.
 3. Lake BM, Salakhutdinov R, Tenenbaum JB: **Human-level concept learning through probabilistic program induction**. *Science* 2015, **350**:1332-1338.
 4. Bejnordi BE, Veta M, Van Diest PJ, Van Ginneken B, Karssemeijer N, Litjens G, Van Der Laak JA, Hermsen M, Manson QF, Balkenhol M *et al.*: **Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer**. *JAMA* 2017, **318**:2199-2210.
 5. Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G *et al.*: **Human-level control through deep reinforcement learning**. *Nature* 2015.
 6. Silver D, Schrittwieser J, Simonyan K, Antonoglou I, Huang A, Guez A, Hubert T, Baker L, Lai M, Bolton A *et al.*: **Mastering the**

- game of go without human knowledge.** *Nature* 2017, **550**:354-359.
7. Sutton RS, Barto AG: *Reinforcement Learning: An Introduction*. MIT Press; 2018.
 8. Barto AG, Singh S, Chentanez N: **Intrinsically motivated learning of hierarchical collections of skills.** *Proceedings of the 3rd International Conference on Development and Learning* 2004:112-119
- One of the first papers on intrinsically motivated reinforcement learning wherein intrinsic rewards are implemented analogous to novelty responses of dopamine neurons.
9. Singh S, Lewis RL, Barto AG, Sorg J: **Intrinsically motivated reinforcement learning: an evolutionary perspective.** *IEEE Trans Auton Mental Dev* 2010, **2**:70-82.
 10. Barto AG: **Intrinsic motivation and reinforcement learning.** *Intrinsically Motivated Learning in Natural and Artificial Systems*; Springer: 2013:17-47.
 11. Kaelbling LP: *Learning in Embedded Systems*. MIT Press; 1993.
 12. Houthoofd R, Chen X, Duan Y, Schulman J, De Turck F, Abbeel P: **VIME: Variational information maximizing exploration.** *NIPS* 2016.
 13. Fu J, Co-Reyes JD, Levine S: *Ex2: Exploration With Exemplar Models for Deep Reinforcement Learning*. 2017arXiv:1703.01260.
 14. Pathak D, Agrawal P, Efros AA, Darrell T: **Curiosity-driven exploration by self-supervised prediction.** *International Conference on Machine Learning (ICML), vol 2017* 2017:2017
- This paper implements curiosity as an intrinsic reward based on an agent's prediction errors to enable learning even in the complete absence of external rewards.
15. Savinov N, Raichuk A, Marinier R, Vincent D, Pollefeys M, Lillicrap T, Gelly S: *Episodic Curiosity Through Reachability*. 2018arXiv:1810.02274.
 16. Conti E, Madhavan V, Such FP, Lehman J, Stanley K, Clune J: **Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents.** *Advances in Neural Information Processing Systems* 2018:5027-5038.
 17. Pathak D, Agrawal P, Efros AA, Darrell T: **Curiosity-driven exploration by self-supervised prediction.** *International Conference on Machine Learning (ICML), vol 2017* 2017:2017.
 18. Burda Y, Edwards H, Pathak D, Storkey A, Darrell T, Efros AA: *Large-Scale Study of Curiosity-Driven Learning*. 2018arXiv:1808.04355.
 19. Tang H, Houthoofd R, Foote D, Stooke A, Chen OX, Duan Y, Schulman J, DeTurck F, Abbeel P: **# Exploration: a study of count-based exploration for deep reinforcement learning.** *Advances in Neural Information Processing Systems* 2017:2750-2759
- This paper proposes a generalization of classic count-based exploration to high dimensional spaces through hashing and demonstrates its effectiveness on various deep RL benchmark problems.
20. Schmidhuber J: **Curious model-building control systems.** *IEEE International Joint Conference on Neural Networks*; IEEE: 1991:1458-1463.
 21. Lillicrap TP, Hunt JJ, Pritzel A, Heess N, Erez T, Tassa Y, Silver D, Wierstra D: *Continuous Control With Deep Reinforcement Learning*. ICLR; 2016.
 22. Duan Y, Chen X, Houthoofd R, Schulman J, Abbeel P: **Benchmarking deep reinforcement learning for continuous control.** *International Conference on Machine Learning* 2016:1329-1338.
 23. Bellemare M, Srinivasan S, Ostrovski G, Schaul T, Saxton D, Munos R: **Unifying count-based exploration and intrinsic motivation.** *NIPS* 2016.
 24. Lehman J, Stanley KO: **Abandoning objectives: evolution through the search for novelty alone.** *Evol Comput* 2011, **19**:189-223.
 25. Schmidhuber J: **Formal theory of creativity, fun, and intrinsic motivation (1990-2010).** *IEEE Trans Auton Mental Dev* 2010.
 26. Oudeyer P-Y, Kaplan F, Hafner VV: **Intrinsic motivation systems for autonomous mental development.** *Evol Comput* 2007.
 27. Schmidhuber J: **Driven by compression progress: a simple principle explains essential aspects of subjective beauty, novelty, surprise, interestingness, attention, curiosity, creativity, art, science, music, jokes.** *Workshop on Anticipatory Behavior in Adaptive Learning Systems*; Springer: 2008:48-76
- This paper provides an overview and mathematical intuition behind the compression progress hypothesis which is the underlying motivation for many prediction error based curiosity algorithms.
28. Colas C, Sigaud O, Oudeyer P-Y: *GEP-PG: Decoupling Exploration and Exploitation in Deep Reinforcement Learning Algorithms*. 2018arXiv:1802.05054.
 29. Berlyne DE: **Novelty and curiosity as determinants of exploratory behaviour.** *Br J Psychol Gen Sect* 1950, **41**:68-80.
 30. Ranganath C, Rainer G: **Neural mechanisms for detecting and remembering novel events.** *Nat Rev Neurosci* 2003, **4**:193-202.
 31. Wittmann BC, Bunzeck N, Dolan RJ, Düzel E: **Anticipation of novelty recruits reward system and hippocampus while promoting recollection.** *Neuroimage* 2007, **38**:194-202.
 32. Wittmann BC, Daw ND, Seymour B, Dolan RJ: **Striatal activity underlies novelty-based choice in humans.** *Neuron* 2008, **58**:967-973.
 33. Düzel E, Bunzeck N, Guitart-Masip M, Düzel S: **Novelty-related motivation of anticipation and exploration by dopamine (NOMAD): implications for healthy aging.** *Neurosci Biobehav Rev* 2010, **34**:660-669.
 34. Berlyne DE: *Mcgraw-Hill Series in Psychology. Conflict, Arousal, and Curiosity*. 1960.
 35. Berlyne DE: **Complexity and incongruity variables as determinants of exploratory choice and evaluative ratings.** *Can J Psychol/Revue canadienne de psychologie* 1963, **17**:274.
 36. Loewenstein G: **The psychology of curiosity: a review and reinterpretation.** *Psychol Bull* 1994, **116**:75.
 37. Golman R, Loewenstein G: **Curiosity, information gaps, and the utility of knowledge.** *Inform Gaps Util Knowl* 2015:96-135.
 38. Kang MJ, Hsu M, Krajbich IM, Loewenstein G, McClure SM, Wang JTY, Camerer CF: **The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory.** *Psychol Sci* 2009, **20**:963-973.
 39. Baranes A, Oudeyer P-Y, Gottlieb J: **Eye movements reveal epistemic curiosity in human observers.** *Vis Res* 2015, **117**:81-90.
 40. Baranes A, Oudeyer P-Y, Gottlieb J: **The effects of task difficulty, novelty and the size of the search space on intrinsically motivated exploration.** *Front Neurosci* 2014, **8**:317.
 41. Geana A, Wilson R, Daw ND, Cohen J: **Boredom, information-seeking and exploration.** *Proceedings of the Annual Conference of the Cognitive Science Society* 2016.
 42. Gruber MJ, Gelman BD, Ranganath C: **States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit.** *Neuron* 2014, **84**:486-496.
 43. Marvin CB, Shohamy D: **Curiosity and reward: valence predicts choice and information prediction errors enhance learning.** *J Exp Psychol: Gen* 2016, **145**:266.
 44. Gruber MJ, Ranganath C: **How curiosity enhances hippocampus-dependent memory: the prediction, appraisal, curiosity, and exploration (pace) framework.** *Trends Cogn Sci* 2019, **23**:1014-1025.
 45. Gruber MJ, Valji A, Ranganath C: *Curiosity and Learning: A Neuroscientific Perspective*. 2019.
 46. Dubey R, Griffiths TL: **Reconciling novelty and complexity through a rational analysis of curiosity.** *Psychol Rev* 2020, **127**:455

This paper presents a rational analysis of curiosity by considering the computational problem underlying curiosity. In doing so, this analysis serves as a useful tool to understand the relationship between previous distinct theories of curiosity and suggests a way to unify them into a single framework.

47. Vasconcelos M, Monteiro T, Kacelnik A: **Irrational choice and the value of information.** *Sci Rep* 2015, **5**:1-12.
48. Cabrero JMR, Zhu J-Q, Ludvig EA: **Costly curiosity: people pay a price to resolve an uncertain gamble early.** *Behav Process* 2019, **160**:20-25.
49. Wang MZ, Hayden BY: **Monkeys are curious about counterfactual outcomes.** *Cognition* 2019, **189**:1-10.
50. Lau JKL, Ozono H, Kuratomi K, Komiya A, Murayama K: **Shared striatal activity in decisions to satisfy curiosity and hunger at the risk of electric shocks.** *Nat Hum Behav* 2020, **4**:531-543.
51. Bromberg-Martin ES, Hikosaka O: **Midbrain dopamine neurons signal preference for advance information about upcoming rewards.** *Neuron* 2009, **63**:119-126.
52. Blanchard TC, Hayden BY, Bromberg-Martin ES: **Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity.** *Neuron* 2015, **85**:602-614.
53. Iigaya K, Story GW, Kurth-Nelson Z, Dolan RJ, Dayan P: **The modulation of savouring by prediction error and its effects on choice.** *Elife* 2016, **5**:e13747.
54. Brydevall M, Bennett D, Murawski C, Bode S: **The neural encoding of information prediction errors during non-instrumental information seeking.** *Sci Rep* 2018, **8**:1-11.
55. Kobayashi K, Ravaioli S, Baranès A, Woodford M, Gottlieb J: **Diverse motives for human curiosity.** *Nat Hum Behav* 2019, **3**:587-595.
56. Dubey R, Griffiths TL, Lombrozo T: **If it's important, then I am curious: a value intervention to induce curiosity.** *Proceedings of the Annual Conference of the Cognitive Science Society* 2019.
57. Hulleman CS, Godes O, Hendricks BL, Harackiewicz JM: **Enhancing interest and performance with a utility value intervention.** *J Educ Psychol* 2010, **102**:880.
58. Harackiewicz JM, Rozek CS, Hulleman CS, Hyde JS: **Helping parents to motivate adolescents in mathematics and science: an experimental test of a utility-value intervention.** *Psychol Sci* 2012, **23**:899-906.