Balancing exploration and exploitation with information and randomization

Robert C. Wilson^{1,2,3,*}, Elizabeth Bonawitz⁴, Vincent D. Costa⁵, and R. Becket Ebitz⁶

¹Department of Psychology, University of Arizona, Tucson AZ USA

²Cognitive Science Program, University of Arizona, Tucson AZ USA

³Evelyn F. McKnight Brain Institute, University of Arizona, Tucson AZ
USA

⁴Department of Psychology, Rutgers University - Newark, Newark NJ
USA

⁵Department of Behavioral Neuroscience, Oregon Health and Science
University, Portland OR USA

⁶Department of Neuroscience, University of Montréal, Montréal, Québec,
Canada

*Correspondence to: bob@arizona.edu

Abstract

Explore-exploit decisions require us to trade off the benefits of exploring unknown options to learn more about them, with exploiting known options, for immediate reward. Such decisions are ubiquitous in nature, but from a computational perspective, they are notoriously hard. There is therefore much interest in how humans and animals make these decisions and recently there has been an explosion of research in this area. Here we provide a biased and incomplete snapshot of this field focusing on the major finding that many organisms use two distinct strategies to solve the explore-exploit dilemma: a bias for information ('directed exploration') and the randomization of choice ('random exploration'). We review evidence for the existence of these strategies, their computational properties, their neural implementations, as well as how directed and random exploration vary over the lifespan. We conclude by highlighting open questions in this field that are ripe to both explore and exploit.

Introduction

When you go to your favorite restaurant, do you always get the same thing (the pizza you know and love), or are you ever tempted by something else (the ravioli you’ve never tried)? Exploiting an old favorite guarantees the reward of a good meal, but is uninformative (you already know the pizza is good). On the other hand, exploring something new brings no guarantees about the quality of the meal, but does guarantee that you learn something (in this case, whether you like the ravioli or not). This is an example of the explore-exploit dilemma, a behavioral dilemma that occurs any time our desire for information conflicts with our need for reward [1, 2, 3, 4, 5, 6].

Outside of the contrived examples at the start of explore-exploit papers, the explore-exploit dilemma is ubiquitous. It is faced by birds in the sky as they forage for food [7] and fish in the sea as they decide where to hunt [8]. Monkeys face it when navigating changing environments [9] and rats face it all the time in the lab [10]. Even slime molds, without a brain but with a rudimentary ability to process information and a drive to find food, confront the explore-exploit dilemma [11]. As humans we face it when we travel to work [12], decide where to fish [13], or pick out food to order online [14]. As scientists we face it when choosing the next experiment to run. And as experts on explore-exploit decisions, the authors of this paper face it every time we introduce the dilemma to a new audience in our papers and talks: should we exploit the tried-and-trusted restaurant examples we have used before [15], or branch out and explore something new?

Given the ubiquity of the explore-exploit dilemma, an obvious question is: How do we solve it? Based on recent literature in psychology and neuroscience, we argue that humans and animals combine two major (but not mutually exclusive) strategies for exploration: information-seeking (directed exploration) and behavioral variability (random exploration). These strategies are associated with different neural correlates and the natural balance between them changes over the lifespan. Ultimately, we argue that a holistic model of human exploration must account for this duality and we review recent progress in developing such models. We end by highlighting some of the major open questions in the field.

Directed and random strategies for exploration

Optimal solutions to the explore-exploit dilemma are intractable in all but the simplest cases [16, 17, 18, 19]. The reason for this difficulty is that optimal solutions require us to perform massive simulations of the future — considering how choices impact future outcomes and how those outcomes will impact future choices. Because of this computational complexity, most theoretical work has turned to approximate methods. While there are literally hundreds of such approximate algorithms [20], across all of these approaches, two general strategies have proven to be especially effective: an explicit bias for information, ‘directed exploration’ [16, 17], and the randomization of choice, ‘random exploration’ [21, 22, 2] (Figure 1).

Figure 1: Illustration of a simple explore-exploit task and algorithms that attempt to solve it. (a) In this ‘bandits task,’ participants choose between three slot machines, or ‘one-armed bandits,’ loosely based on those in a casino. In this case the bandits pay out rewards sampled from Gaussian distributions whose mean is different for each machine. Participants are not told these mean payoffs and so must explore in order to learn which bandit is best. However, because they receive the reward from the bandit they choose, they also face pressure to exploit the bandit they currently believe to be best. After several plays, participants can have quite different information about each bandit, which can be summarized as a posterior distribution over the mean payoff from each option, (b). In this case, the participant has the most information about the red option, which they also believe to be best (narrow distribution with highest mean), some information about the blue distribution (wide distribution, lower mean), and no information about the yellow option (uniform distribution). (c - f) Different algorithms make use of the information in different ways. (c) A purely exploitative algorithm sets the value of each option, $Q(a)$, deterministically to the mean of the posterior, $r(a)$. This algorithm always exploits the

option with highest $r(a)$, in this case choosing the red option with 100% probability. (d) A purely directed algorithm (in this case the Upper Confidence Bound algorithm, [23]) adds a deterministic information bonus, set to twice the standard deviation of the posterior distribution for each bandit, to the value of each option. This information bonus is largest for the yellow option and the algorithm chooses this exploratory option with 100% probability. (e) A simple random strategy for exploration is to add random noise to the value of all options [22]. This algorithm is not deterministic, and any option can be chosen. However, because the values still include the expected payoff, this option is most likely to choose the option it believes to be best, the red option. (f) A more sophisticated random exploration strategy is Thompson sampling [21]. In this algorithm the value of each option is sampled randomly from the posterior distribution over the mean payoff from each option. The mean of this sampled value is still equal to the expected reward, but its variability - i.e. the decision noise - scales with the uncertainty in each option. Thus random exploration in Thompson sampling is highest when the algorithm is uncertain and reduces to zero when the payoff structure is fully known.

In directed exploration, exploration is ‘directed’ towards more informative options by a deterministic ‘information bonus,’ which increases the value of informative options. Mathematically, one way to implement this strategy is by increasing the value of informative options such that the value of taking action a

$$Q(a) = r(a) + IB(a) \tag{1}$$

where $r(a)$ is how good we expect option a to be and $IB(a)$ is the information bonus for each action. In this simple model, the choice is made by picking the option with the highest $Q(a)$. Different algorithms have different forms for the information bonus. One popular approach, epitomized by the ‘Upper Confidence Bound’ (or UCB) algorithm (Figure 1c), is to set the information bonus proportional to the *uncertainty* about the expected payoff from each option [23]. In some situations the performance of this strategy can be close to optimal despite having a fraction of the computational cost.

In random exploration, random decision noise drives exploration by chance. Mathematically, one way to implement random exploration is by adding noise to the computation of value

$$Q(a) = r(a) + \eta(a) \tag{2}$$

where $\eta(a)$ is zero-mean random noise sampled from some probability distribution and again the choice is the action with largest $Q(a)$. The exact flavor of random exploration is determined by the form of the noise distribution. For example, adding logistically-distributed random noise to values will yield the famous softmax choice function (Figure 1d). Random forms of exploration can perform quite well in many environments and other noise distributions can solve other problems. For example, annealing algorithms reduce the noise over time, while an algorithm called Thompson Sampling scales the noise with the agent’s uncertainty. Both approaches lead to high levels of random exploration

when the environment is unknown, but reduce random exploration later on allowing the agent to exploit what it has learned and leading to near optimal performance in some cases (Figure 1e) [21, 24].

Of course, directed and random exploration are not mutually exclusive and it is possible to use both strategies at once. In fact, there is strong evidence that humans and animals use both strategies; how exactly this holistic strategy could be implemented is a question we address towards the end.

The existence of directed and random exploration in behavior

While the success of directed and random exploration in artificial systems suggests that they may be good strategies for humans and animals, identifying these strategies in behavior proved difficult because of a number of psychological confounds. For example, directed exploration is often opposed by risk and ambiguity aversion [25, 26], while random exploration is difficult to disentangle from boredom or disengagement from the task. As a result, early work looking for directed exploration led to mixed results (e.g. [27, 28]), and the interpretation of choice variability as random exploration has been controversial [29, 30, 31].

A key advance in proving the existence of directed and random exploration was the design of tasks that manipulate how valuable it is to explore, independent of confounding factors. This approach builds on normative models, which identify the environmental features that *should* influence exploration because they alter the utility of new information. By manipulating these environmental features, we create conditions where exploration has value and conditions where it does not. We then identify directed and random exploration by looking for changes in information seeking and behavioral variability between these conditions.

Manipulations that increase the value of exploration include: increasing the time horizon [32, 15], increasing the uncertainty participants have about different options [33], and adding completely novel options that participants have not seen before [34, 35, 36, 37, 38]. Conversely, manipulations that increase the value of exploitation include raising the stakes of a one-off decision and decoupling information from choice [30]. In all of these cases, information seeking and behavioral variability increase when it is valuable to explore and decrease when it is valuable to exploit, providing strong evidence for the existence of both directed and random exploration in behavior.

The existence of directed and random exploration in the brain

Findings from neuroscience support the existence of directed and random exploration and further suggest that the neural processes underlying them may be dissociable.

Directed exploration, but not random exploration, has been associated with several prefrontal structures and systems, including activity in frontal pole [39, 40], mesocorticolimbic regions [36, 37], frontal theta oscillations [41], and prefrontal dopamine [28, 35, 42, 43]. A causal role for frontal regions in directed exploration is suggested by [44], who showed that continuous theta-burst transcranial magnetic stimulation to frontal pole abolishes horizon-dependent directed exploration, but leaves random exploration untouched. Building on a long literature in navigation and foraging [45], Johnson and colleagues [46] proposed that hippocampus should also play a key role in directed exploration.

The unique neural correlates of random exploration are less circumscribed. A number of studies have suggested that increased neural variability as a mechanism for generating exploratory noise. For example, random exploration is associated with a sudden increase in neural variability and a loss of choice tuning in decision-making circuits in monkeys [47]. Likewise, increased neural variability in motor related circuits in both humans [40], rodents [48], and birds [49, 50] is also related to the use of random exploration. However, choice tuning and changes in neural variability have also been associated with directed exploration [51] and could reflect other computations [31].

There is also some evidence that random exploration may be modulated by catecholamines norepinephrine and dopamine. Norepinephrine levels and related measures, like pupil size [52, 53], predict increased noise in behavior [54, 55, 56, 57] and the brain [58, 31]. However, results from systemic pharmacological studies tend to be more mixed [59, 60], and one study even found that blocking norepinephrine only affected exploration of completely novel options, a form of directed exploration [38], that is also modulated by dopamine [35]. Dopamine has been also been associated with random exploration, with decreased tonic dopamine associated increased behavioral variability [61] in rats, and dopaminergic neurons being found to modulate song variability in song birds [62].

Directed and random exploration develop differently over the lifespan

Multiple lines of evidence show that directed and random exploration develop differently over the lifespan. Even infants and preschoolers recognize uncertainty in their environment and make exploratory choices to deconfound variables [63], resolve belief violations [64, 65], and reduce ambiguity in intuitive theories [66]. Traditional bandits tasks reveal that the trade-off between exploration and exploitation shifts over development. Preschoolers are more likely to engage in directed exploration than 7-9 year-olds who are more directed than adults [67, 68].

In naturalistic tasks, preschoolers are sensitive to reward expectation and balance this against factors like whether information is decoupled from reward [69, 70]. However, in more cognitively-demanding tasks, factors like horizon appear to more weakly influence directed exploration for young teens (age 10-12), and grow in influence through adoles-

cence and adulthood (early 20s) [71], remaining stable into old age [72].

Conversely, random exploration as behavioral variability, (e.g. noise over choice behavior) is prevalent in preschoolers and decreases through adulthood [67, 68, 73]. However, horizon-dependent random exploration appears more constant over adolescence and young adulthood [71], though is also reduced in healthy older adults [72]. Taken together, these results suggest that some forms of directed and random exploration may be most prevalent in early childhood and diminish with age, but that sensitivity to factors like horizon may follow different trajectories, highlighting the dissociability of the different types of exploration.

Integrating directed and random exploration

Although directed and random exploration rely on dissociable cognitive and neural systems, it is clear that both are used in behavior. In this section we discuss holistic models for how directed and random exploration could be combined.

Perhaps the simplest holistic strategy is to combine the information bonus and decision noise into one equation to estimate the value of choosing a particular option [15, 6]

$$Q(a) = r(a) + IB(a) + n(a) \quad (3)$$

In this case, directed and random exploration have additive effects on value, with exploration occurring whenever the information bonus or decision noise tips the balance away from the exploitative choice.

Another approach would be to mix the distinct strategies for directed exploration, random exploration, and exploitation over time. In this case, a higher-order process decides *when* to explore (and perhaps also whether to use a directed or random strategy), while the decision of *what* to choose is driven by component processes of exploitation, and directed and random exploration. Consistent with this separation of what and when, several studies argue for dissociable exploratory and exploitative states, on the basis of qualitative differences in behavior and neural activity [74, 75, 53, 36]. Further, directed exploration can occur at random times, an observation that is difficult to reconcile with a unified mechanism for deciding when and what to explore [76].

A third approach to combining directed and random exploration, contends that the two strategies emerge as different behavioral facets of a unified algorithm known in the machine learning literature as Deep Exploration [77]. In this model [78], the explore-exploit choice is made by mental simulation of a small number of plausible futures that are ‘deep,’ in that they extend multiple time steps into the future, but narrow, in that the number of simulations used is small. Random exploration arises from this model because the simulations are stochastic. More subtly, directed exploration also arises from Deep Exploration if the mental simulations extend multiple time steps into the future and include lose-shift behavior, i.e. mentally simulating a switch away from an explore option if the simulated outcome is bad. Such lose-shift simulations limit the simulated downside of ex-

ploring, which boosts the simulated value of exploring and hence biases the exploration towards informative options.

Beyond the existence of directed and random exploration, Deep Exploration also accounts for the horizon dependence, uncertainty dependence and feedback dependence of directed and random exploration. Moreover, it predicts that there should be a tradeoff between directed and random exploration. As people use more simulations to make their decision they should exhibit more directed exploration and less random exploration, a prediction that holds both across the population and within subject [78].

Open questions

What should we continue to exploit?

There are outstanding issues in almost every section above and much future work will involve exploiting these ideas. Open questions include: What are the neural mechanisms of directed and random exploration? And how do they relate to other cognitive processes such as working memory [79] and motor control [47, 40, 48, 49, 50]? How does explore-exploit behavior in general (and directed and random exploration in particular) vary across species and cultures? Are individual differences in directed and random exploration stable traits? Do directed and random exploration change in mental illness [80, 81, 82, 83, 84, 85]? And is modifying explore-exploit, with neural stimulation or behavioral therapy, a potential therapeutic target [86, 44]?

In regard to the computations underlying explore-exploit decisions: Are they really made by mental simulation as implied by Deep Exploration? And if Deep Exploration is really at play, how can this unifying theory explain the dissociation between directed and random exploration? Are there separate decision processes for ‘when’ versus ‘what’ to explore? And, more generally, are exploration and exploitation distinct behavioral states or are they drives that work together in a more continuous fashion?

What is left to explore?

There are a number of more open-ended questions we should explore, although this is by no means an exhaustive list.

How does explore-exploit behavior relate to other constructs?

Because we still lack precise operational definitions, there is perennial confusion about the relationship between explore-exploit behavior and a number of other related constructs. These include ‘classic’ measures in behavioral economics such as risk taking, ambiguity attitude, and loss aversion, as well as measures from ecology such as foraging, stopping problems, and search [19]. The relationship with foraging and search has been questioned

empirically [87] and whether these are distinct behaviors or share some variance is a key question.

Outside of the physical world, explore-exploit problems may also occur in the mental domain. Indeed, mental search and theory change has framed learning as an explore-exploit problem in which we choose between exploiting our current beliefs about the world or mentally exploring for new hypotheses (see, e.g. [88, 89, 90, 91]). The degree to which the processes of mental search align with those in the physical world remain a topic for future work; additionally behavioral methods to tease apart directed versus random search in mental space remain elusive.

Finally, it is an open question how explore-exploit behavior relates to emotions such as curiosity, interest, and affect [92, 93, 94, 95, 96, 97]. While these concepts are thought to inspire or inhibit exploration or exploitation, the relationship between them and explore-exploit behavior has not been sufficiently explored. There is a conceptual bias that exploration is an intrinsically appetitive act, intended to maximize future rewards and minimize future losses. However, depending on the decision context and choice outcomes exploring could be construed as aversive and avoided. By examining directed and random exploration in both appetitive and aversive environments we can understand how explore-exploit decisions differ from simpler approach-avoidance behaviors.

How do we explore complex environments?

Many explore-exploit papers begin in the real-world, with a motivating example like the restaurant example in this paper, but they almost always end up in the lab, with participants performing highly controlled tasks with few options and simple probability distributions over rewards. The real world does not look like this, and the obvious question is whether that matters for explore-exploit behavior or not?

That is, are directed and random exploration enough for the complexity of real-world explore-exploit problems? Or are we missing something fundamental from the theory? A new type of exploration perhaps [38]? Or exploration based on stereotyped behaviors, that are not based on information or randomization, but are nevertheless effective [98]?

Relatedly, does the complexity of the real world, make implementing the more sophisticated algorithms for directed and random exploration too difficult? For example, Thompson Sampling and Deep Exploration require the ability to simulate future outcomes, a computation that gets more difficult as the complexity of those outcomes increases. Conversely, simpler strategies such as a fixed bias for novel stimuli [99] or adding decision noise that is not tuned to simulated outcomes may perform well enough in more complex settings that a more difficult computation is not justified.

Towards this end of understanding how directed and random exploration are implemented in more complex settings, is recent work by [100] who used a task with hundreds of options and a complex reward structure. While this increased complexity required a new model of *learning* to explain behavior, the decision process, based on an information bonus and decision noise, did not. Similarly, work on exploration using online food orders found evidence for uncertainty driven exploration in this real-world case [14]. As

impressive as this work is, we have only just begun to scratch the surface of the complexity of the real world and there is, quite literally, a whole world out there left to explore!

Acknowledgements

This work was supported by grants from the National Institute on Aging (R01AG061888 to RCW), the National Science Foundation (SES #1627971 to EB), the McDonnell Foundation (to EB and to RCW), the Jacobs Foundation (to EB), a Brain and Behavior Foundation Young Investigator Award (#27298 to RBE) and an Unfettered Research Grant from the Momental Foundation (to RBE).

Highlighted papers

- ** [30] Demonstrates that most noise in reinforcement learning tasks is learning noise that has autocorrelation over time. However, the noise used for random exploration is different and appears to have no autocorrelation.
- ** [47] Suggests a possible mechanism for generating random exploration with random neural activity
- ** [14] Leverages a large data set of food orders to show that real-world explore-exploit decisions are modulated by uncertainty
- ** [36] Monkeys use directed exploration to manage explore-exploit tradeoffs and these signals are coded in motivational brain regions (amygdala and ventral striatum).
- ** [40] Uses neuroimaging to identify different areas of the brain associated with directed and random exploration.
- ** [78] Proposes Deep Exploration as a unifying account of explore-exploit behavior.
- ** [67] Investigates the developmental profile of directed and random exploration in a bandits task with a large number of options.
- * [84] Schizophrenia is associated with changes in directed exploration, but not random exploration
- * [76] Directed exploration occurs even when exploration has no value, suggesting a random mechanism for deciding *when* to explore.
- * [100] Explore-exploit behavior probed using a bandits task with a large number of options. Also notable for the sophisticated computational model needed to analyze the data.

References

- [1] Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285, 1996.
- [2] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [3] Jonathan D Cohen, Samuel M McClure, and Angela J Yu. Should i stay or should i go? how the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481):933–942, 2007.
- [4] Thomas T Hills, Peter M Todd, David Lazer, A David Redish, Iain D Couzin, Cognitive Search Research Group, et al. Exploration versus exploitation in space, mind, and society. *Trends in cognitive sciences*, 19(1):46–54, 2015.
- [5] Katja Mehlhorn, Ben R Newell, Peter M Todd, Michael D Lee, Kate Morgan, Victoria A Braithwaite, Daniel Hausmann, Klaus Fiedler, and Cleotilde Gonzalez. Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, 2(3):191, 2015.
- [6] Eric Schulz and Samuel J Gershman. The algorithmic architecture of exploration in the human brain. *Current opinion in neurobiology*, 55:7–14, 2019.
- [7] John R Krebs, Alejandro Kacelnik, and Peter Taylor. Test of optimal sampling by foraging great tits. *Nature*, 275(5675):27–31, 1978.
- [8] David W Sims, Emily J Southall, Nicolas E Humphries, Graeme C Hays, Corey JA Bradshaw, Jonathan W Pitchford, Alex James, Mohammed Z Ahmed, Andrew S Brierley, Mark A Hindell, et al. Scaling laws of marine predator search behaviour. *Nature*, 451(7182):1098–1102, 2008.
- [9] Harriet R Thatcher, Colleen T Downs, and Nicola F Koyama. Anthropogenic influences on the time budgets of urban vervet monkeys. *Landscape and Urban Planning*, 181:38–44, 2019.
- [10] Brian J Jackson, Gusti Lulu Fatima, Sujean Oh, and David H Gire. Many paths to the same goal: balancing exploration and exploitation during probabilistic route planning. *Eneuro*, 7(3), 2020.
- [11] Chris R Reid, Tanya Latty, Audrey Dussutour, and Madeleine Beekman. Slime mold uses an externalized spatial “memory” to navigate in complex environments. *Proceedings of the National Academy of Sciences*, 109(43):17490–17494, 2012.

- [12] Shaun Larcom, Ferdinand Rauch, and Tim Willems. The benefits of forced experimentation: striking evidence from the london underground network. *The Quarterly Journal of Economics*, 132(4):2019–2055, 2017.
- [13] Shay O’Farrell, James N Sanchirico, Orr Spiegel, Maxime Depalle, Alan C Haynie, Steven A Murawski, Larry Perruso, and Andrew Strelcheck. Disturbance modifies payoffs in the explore-exploit trade-off. *Nature communications*, 10(1):1–9, 2019.
- [14] Eric Schulz, Rahul Bhui, Bradley C Love, Bastien Brier, Michael T Todd, and Samuel J Gershman. Structured, uncertainty-driven exploration in real-world consumer choice. *Proceedings of the National Academy of Sciences*, 116(28):13903–13908, 2019.
- [15] Robert C Wilson, Andra Geana, John M White, Elliot A Ludvig, and Jonathan D Cohen. Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, 143(6):2074, 2014.
- [16] Richard Bellman. A problem in the sequential design of experiments. *Sankhyā: The Indian Journal of Statistics (1933-1960)*, 16(3/4):221–229, 1956.
- [17] John C Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(2):148–164, 1979.
- [18] Shunan Zhang and J Yu Angela. Forgetful bayes and myopic planning: Human learning and decision-making in a bandit setting. In *Advances in neural information processing systems*, pages 2607–2615, 2013.
- [19] Bruno B Averbeck. Theory of choice in bandit, information sampling and foraging tasks. *PLoS computational biology*, 11(3), 2015.
- [20] Sébastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and non-stochastic multi-armed bandit problems. *arXiv preprint arXiv:1204.5721*, 2012.
- [21] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- [22] Christopher John Cornish Hellaby Watkins. *Learning from delayed rewards*. PhD thesis, King’s College, Cambridge, 1989.
- [23] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- [24] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on learning theory*, pages 39–1, 2012.
- [25] Daniel Ellsberg. Risk, ambiguity, and the savage axioms. *The quarterly journal of economics*, pages 643–669, 1961.

- [26] Colin Camerer and Martin Weber. Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of risk and uncertainty*, 5(4):325–370, 1992.
- [27] Nathaniel D Daw, John P O’doherly, Peter Dayan, Ben Seymour, and Raymond J Dolan. Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095):876–879, 2006.
- [28] Michael J Frank, Bradley B Doll, Jen Oas-Terpstra, and Francisco Moreno. Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. *Nature neuroscience*, 12(8):1062, 2009.
- [29] Matthew R Nassar and Michael J Frank. Taming the beast: extracting generalizable knowledge from computational models of cognition. *Current opinion in behavioral sciences*, 11:49–54, 2016.
- [30] Charles Findling, Vasilisa Skvortsova, Remi Dromnelle, Stefano Palminteri, and Valentin Wyart. Computational noise in reward-guided learning drives behavioral variability in volatile environments. *Nature neuroscience*, pages 1–12, 2019.
- [31] Timothy H Muller, Rogier B Mars, Timothy E Behrens, and Jill X O’Reilly. Control of entropy in neural models of environmental state. *Elife*, 8:e39404, 2019.
- [32] Alejandro Kacelnik. *Studies of foraging behaviour and time budgeting in great tits (Parus major)*. PhD thesis, University of Oxford, 1979.
- [33] Samuel J Gershman. Deconstructing the human algorithms for exploration. *Cognition*, 173:34–42, 2018.
- [34] Bianca C Wittmann, Nathaniel D Daw, Ben Seymour, and Raymond J Dolan. Striatal activity underlies novelty-based choice in humans. *Neuron*, 58(6):967–973, 2008.
- [35] Vincent D Costa, Valery L Tran, Janita Turchi, and Bruno B Averbeck. Dopamine modulates novelty seeking behavior during decision making. *Behavioral neuroscience*, 128(5):556, 2014.
- [36] Vincent D Costa, Andrew R Mitz, and Bruno B Averbeck. Subcortical substrates of explore-exploit decisions in primates. *Neuron*, 103(3):533–545, 2019.
- [37] Vincent D Costa and Bruno B Averbeck. Primate orbitofrontal cortex codes information relevant for managing explore–exploit tradeoffs. *Journal of Neuroscience*, 40(12):2553–2561, 2020.
- [38] Magda Dubois, Johanna Habicht, Jochen Michely, Rani Moran, Ray Dolan, and Tobias Hauser. Noradrenaline modulates tabula-rasa exploration. *bioRxiv*, 2020.
- [39] David Badre, Bradley B Doll, Nicole M Long, and Michael J Frank. Rostrolateral prefrontal cortex and individual differences in uncertainty-driven exploration. *Neuron*, 73(3):595–607, 2012.

- [40] Momchil S Tomov, Van Q Truong, Rohan A Hundia, and Samuel J Gershman. Dissociable neural correlates of uncertainty underlie different exploration strategies. *Nature communications*, 11(1):1–12, 2020.
- [41] James F Cavanagh, Christina M Figueroa, Michael X Cohen, and Michael J Frank. Frontal theta reflects uncertainty and unexpectedness during exploration and exploitation. *Cerebral cortex*, 22(11):2575–2586, 2012.
- [42] Samuel J Gershman and Bastian Greshake Tzovaras. Dopaminergic genes are associated with both directed and random exploration. *Neuropsychologia*, 120:97–104, 2018.
- [43] Karima Chakroun, David Mathar, Antonius Wiehler, Florian Ganzer, and Jan Peters. Dopaminergic modulation of the exploration/exploitation trade-off in human decision-making. *Elife*, 9:e51260, 2020.
- [44] Wojciech K Zajkowski, Malgorzata Kossut, and Robert C Wilson. A causal role for right frontopolar cortex in directed, but not random, exploration. *Elife*, 6:e27430, 2017.
- [45] John O’keefe and Lynn Nadel. *The hippocampus as a cognitive map*. Oxford: Clarendon Press, 1978.
- [46] Adam Johnson, Zachary Varberg, James Benhardus, Anthony Maahs, and Paul Schrater. The hippocampus and exploration: dynamically evolving behavior and neural representations. *Frontiers in human neuroscience*, 6:216, 2012.
- [47] R Becket Ebitz, Eddy Albarran, and Tirin Moore. Exploration disrupts choice-predictive signals and alters dynamics in prefrontal cortex. *Neuron*, 97(2):450–461, 2018.
- [48] Masayoshi Murakami, Hanan Shteingart, Yonatan Loewenstein, and Zachary F Mainen. Distinct sources of deterministic and stochastic components of action timing decisions in rodent frontal cortex. *Neuron*, 94(4):908–919, 2017.
- [49] Ashesh K Dhawale, Maurice A Smith, and Bence P Ölveczky. The role of variability in motor learning. *Annual review of neuroscience*, 40:479–498, 2017.
- [50] Satoshi Kojima, Mimi H Kao, Allison J Doupe, and Michael S Brainard. The avian basal ganglia are a source of rapid behavioral variation that enables vocal motor exploration. *Journal of Neuroscience*, 38(45):9635–9647, 2018.
- [51] R Becket Ebitz, Jiaxin Cindy Tu, and Benjamin Y Hayden. Rule adherence warps decision-making. *BioRxiv*, 2019.
- [52] Siddhartha Joshi and Joshua I Gold. Pupil size as a window on neural substrates of cognition. *Trends in Cognitive Sciences*, 2020.

- [53] R Becket Ebitz and Tirin Moore. Both a gauge and a filter: Cognitive modulations of pupil size. *Frontiers in neurology*, 9:1190, 2019.
- [54] Marieke Jepma and Sander Nieuwenhuis. Pupil diameter predicts changes in the exploration–exploitation trade-off: Evidence for the adaptive gain theory. *Journal of cognitive neuroscience*, 23(7):1587–1596, 2011.
- [55] D. G. Tervo, M. Proskurin, M. Manakov, M. Kabra, A. Vollmer, K. Branson, and A. Y. Karpova. Behavioral variability through stochastic choice and its gating by anterior cingulate cortex. *Cell*, 159(1):21–32, 2014.
- [56] Hideki Ohira, Naho Ichikawa, Kenta Kimura, Seisuke Fukuyama, Jun Shinoda, and Jitsuihiro Yamada. Neural and sympathetic activity associated with exploration in decision-making: further evidence for involvement of insula. *Frontiers in behavioral neuroscience*, 8:381, 2014.
- [57] R Becket Ebitz, John M Pearson, and Michael L Platt. Pupil size and social vigilance in rhesus macaques. *Frontiers in neuroscience*, 8:100, 2014.
- [58] Ana Raquel O Martins and Robert C Froemke. Coordinated forms of noradrenergic plasticity in the locus coeruleus and primary auditory cortex. *Nature neuroscience*, 18(10):1483–1492, 2015.
- [59] Christopher M Warren, Robert C Wilson, Nic J Van Der Wee, Eric J Giltay, Martijn S Van Noorden, Jonathan D Cohen, and Sander Nieuwenhuis. The effect of atomoxetine on random and directed exploration in humans. *PloS one*, 12(4):e0176034, 2017.
- [60] Marieke Jepma, Erik T Te Beek, Eric-Jan Wagenmakers, Joop Van Gerven, and Sander Nieuwenhuis. The role of the noradrenergic system in the exploration–exploitation trade-off: a pharmacological study. *Frontiers in human neuroscience*, 4:170, 2010.
- [61] François Cinotti, Virginie Fresno, Nassim Aklil, Etienne Coutureau, Benoît Girard, Alain R Marchand, and Mehdi Khamassi. Dopamine blockade impairs the exploration–exploitation trade-off in rats. *Scientific reports*, 9(1):1–14, 2019.
- [62] Agata Budzillo, Alison Duffy, Kimberly E Miller, Adrienne L Fairhall, and David J Perkel. Dopaminergic modulation of basal ganglia output through coupled excitation–inhibition. *Proceedings of the National Academy of Sciences*, 114(22):5713–5718, 2017.
- [63] Laura E Schulz and Elizabeth Baraff Bonawitz. Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental psychology*, 43(4):1045, 2007.

- [64] Elizabeth Baraff Bonawitz, Tessa JP van Schijndel, Daniel Friel, and Laura Schulz. Children balance theories and evidence in exploration, explanation, and learning. *Cognitive psychology*, 64(4):215–234, 2012.
- [65] Aimee E Stahl and Lisa Feigenson. Observing the unexpected enhances infants’ learning and exploration. *Science*, 348(6230):91–94, 2015.
- [66] J. Wang, Y. Yang, C. Macias, and E. Bonawitz. Children with more immature intuitive theories seek domain-relevant information. in revision.
- [67] Eric Schulz, Charley M Wu, Azzurra Ruggeri, and Björn Meder. Searching for rewards like a child means less generalization and more directed exploration. *Psychological science*, 30(11):1561–1572, 2019.
- [68] Björn Meder, Charley M Wu, Eric Schulz, and Azzurra Ruggeri. Development of directed and random exploration in children. 2020.
- [69] Elizabeth Bonawitz, Ilona Bass, and Elizabeth Lapidow. Choosing to learn: Evidence evaluation for active learning and teaching in early childhood. In *Active Learning from Infancy to Childhood*, pages 213–231. Springer, 2018.
- [70] E. Lapidow and E. Bonawitz. Explore-exploit: Ambiguity, expectation, and information gain influence preschooler’s choices in exploration. in review.
- [71] Leah H Somerville, Stephanie F Sasse, Megan C Garrad, Andrew T Drysdale, Nadine Abi Akar, Catherine Insel, and Robert C Wilson. Charting the expansion of strategic exploratory behavior during adolescence. *Journal of experimental psychology: general*, 146(2):155, 2017.
- [72] Jack-Morgan Mizell, Siyu Wang, Alec Frisvold, Lily Alvarado, Alex Farrell-Skupny, Waitsang Keung, Mark H. Sundman, Mary-Kathryn Franchetti, Ying-hui Chou, Gene E. Alexander, and Robert C. Wilson. Differential impacts of healthy cognitive aging on directed and random exploration.
- [73] Rista C Plate, Jacqueline M Fulvio, Kristin Shutts, C Shawn Green, and Seth D Polak. Probability learning: Changes in behavior across time and development. *Child development*, 89(1):205–218, 2018.
- [74] Mark Steyvers, Michael D Lee, and Eric-Jan Wagenmakers. A bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, 53(3):168–179, 2009.
- [75] Michael D Lee, Shunan Zhang, Miles Munro, and Mark Steyvers. Psychological models of human and optimal performance in bandit problems. *Cognitive Systems Research*, 12(2):164–174, 2011.

- [76] R Becket Ebitz, Brianna J Sleezer, Hank P Jedema, Charles W Bradberry, and Benjamin Y Hayden. Tonic exploration governs both flexibility and lapses. *PLoS computational biology*, 15(11):e1007475, 2019.
- [77] Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. Deep exploration via bootstrapped dqn. In *Advances in neural information processing systems*, pages 4026–4034, 2016.
- [78] Robert Wilson, Siyu Wang, Hashem Sadeghiyeh, and Jonathan D Cohen. Deep exploration as a unifying account of explore-exploit behavior. 2020.
- [79] Irene Cogliati Dezza, Axel Cleeremans, and William Alexander. Should we control? the interplay between cognitive control and information integration in the resolution of the exploration-exploitation dilemma. *Journal of Experimental Psychology: General*, 148(6):977, 2019.
- [80] Bruno B Averbeck, Atbin Djamshidian, Sean S O’Sullivan, Charlotte R Housden, Jonathan P Roiser, and Andrew J Lees. Uncertainty about mapping future actions into rewards may underlie performance on multiple measures of impulsivity in behavioral addiction: Evidence from parkinson’s disease. *Behavioral neuroscience*, 127(2):245, 2013.
- [81] Gregory P Strauss, Michael J Frank, James A Waltz, Zuzana Kasanova, Ellen S Herbener, and James M Gold. Deficits in positive reinforcement learning and uncertainty-driven exploration are associated with distinct aspects of negative symptoms in schizophrenia. *Biological psychiatry*, 69(5):424–431, 2011.
- [82] Antonius Wiehler, Karima Chakroun, and Jan Peters. Attenuated directed exploration during reinforcement learning in gambling disorder. *BioRxiv*, page 823583, 2019.
- [83] Irene Cogliati Dezza, Xavier Noel, Axel Cleeremans, and Angela Yu. Novelty-seeking impairment in addiction. *BioRxiv*, 2020.
- [84] J. A. Waltz, R. C. Wilson, M. A. Albrecht, M. J. Frank, and J. M. Gold. Differential effects of psychotic illness on directed and random exploration. *Computational Psychiatry*, 2020.
- [85] Flurin Cathomas, Federica Klaus, Karoline Guetter, Hui-Kuan Chung, Anjali Raja Beharelle, Tobias Spiller, Rebecca Schlegel, Erich Seifritz, Matthias Hartmann-Riemer, Philippe N Tobler, et al. Increased random exploration in schizophrenia is associated with inflammation. *bioRxiv*, 2020.
- [86] Anjali Raja Beharelle, Rafael Polanía, Todd A Hare, and Christian C Ruff. Transcranial stimulation over frontopolar cortex elucidates the choice attributes and neural mechanisms used to resolve exploration–exploitation trade-offs. *Journal of Neuroscience*, 35(43):14544–14556, 2015.

- [87] Bettina von Helversen, Rui Mata, Gregory R Samanez-Larkin, and Andreas Wilke. Foraging, exploration, or search? on the (lack of) convergent validity between three behavioral paradigms. *Evolutionary Behavioral Sciences*, 12(3):152, 2018.
- [88] Elizabeth Bonawitz, Stephanie Denison, Thomas L Griffiths, and Alison Gopnik. Probabilistic models, learning algorithms, and response variability: sampling in cognitive development. *Trends in cognitive sciences*, 18(10):497–500, 2014.
- [89] Elizabeth Bonawitz, Stephanie Denison, Alison Gopnik, and Thomas L Griffiths. Win-stay, lose-sample: A simple sequential algorithm for approximating bayesian inference. *Cognitive psychology*, 74:35–65, 2014.
- [90] Elizabeth Bonawitz, Tomer D Ullman, Sophie Bridgers, Alison Gopnik, and Joshua B Tenenbaum. Sticking to the evidence? a behavioral and computational case study of micro-theory change in the domain of magnetism. *Cognitive science*, 43(8):e12765, 2019.
- [91] Tomer D Ullman, Noah D Goodman, and Joshua B Tenenbaum. Theory learning as stochastic search in the language of thought. *Cognitive Development*, 27(4):455–480, 2012.
- [92] Daniel E Berlyne. Curiosity and exploration. *Science*, 153(3731):25–33, 1966.
- [93] Celeste Kidd and Benjamin Y Hayden. The psychology and neuroscience of curiosity. *Neuron*, 88(3):449–460, 2015.
- [94] Jacqueline Gottlieb and Pierre-Yves Oudeyer. Towards a neuroscience of active sampling and curiosity. *Nature Reviews Neuroscience*, 19(12):758–770, 2018.
- [95] Andra Geana, Robert Wilson, Nathaniel D Daw, and Jonathan D Cohen. Boredom, information-seeking and exploration. In *CogSci*, 2016.
- [96] Suzanne Hidi and K Ann Renninger. The four-phase model of interest development. *Educational psychologist*, 41(2):111–127, 2006.
- [97] Joseph Colantonio and Elizabeth Bonawitz. Awesome play: Awe increases preschooler’s exploration and discovery. 2018.
- [98] R Blythe Towal and Mitra JZ Hartmann. Variability in velocity profiles during free-air whisking behavior of unrestrained rats. *Journal of neurophysiology*, 100(2):740–752, 2008.
- [99] Peter Redgrave and Kevin Gurney. The short-latency dopamine signal: a role in discovering novel actions? *Nature reviews neuroscience*, 7(12):967–975, 2006.
- [100] C. M. Wu, E. Schulz, M. Speekenbrink, J. D. Nelson, and B. Meder. Generalization guides human exploration in vast decision spaces. *Nature human behaviour*, 2(12):915–924, 2018.