

# **Cognitive Science Society Workshop: Understanding Exploration-Exploitation Trade-offs**

**Elizabeth Bonawitz (lbaraff@gmail.com)**

Department of Psychology, Rutgers University – Newark  
101 Warren Street; Newark, NJ 07102

**Alison Gopnik (gopnik@berkeley.edu)**

Department of Psychology, University of California, Berkeley  
3317 Tolman Hall; Berkeley, CA 94720

**Celeste Kidd (celestekidd@gmail.com)**

Department of Brain & Cognitive Sciences, University of Rochester  
Meliora Hall, 500 Joseph C. Wilson Blvd., Rochester, NY 14627

**Keywords:** explore-exploit, information seeking, reward, curiosity, reinforcement learning

## **Workshop summary**

All cognitive systems, almost by definition, both gather information and use that information to make decisions –they explore and exploit. But there is an intrinsic tension between these two aspects of cognition. Exploration is costly – the resources that go towards information-seeking are not available for action. Moreover, actions have to be taken before all the relevant information is available. When should a cognitive system seek new information and when should it stop seeking information and simply decide how to act? How do people and other organisms resolve this tension between exploration and exploitation? Is there an optimal way of doing so? And how can systems be motivated to seek information when that pursuit is costly and the immediate utility of the information may not be apparent?

Recently, a wide range of researchers from very different disciplines and traditions have begun to address these questions, and in particular, have used computational methods to characterize explore/exploit tensions. This workshop brings together an interdisciplinary group of researchers, with expertise in neuroscience, computation, robotics, evolutionary biology, and developmental psychology to attempt to connect these separate literatures, draw some general conclusions, and forge a research path for the future.

In reinforcement learning, researchers have used “multi-armed bandit” tasks to determine when learners switch from a more rewarding action to a less rewarding one in search of more information, and which switching techniques are optimal, and have also explored neural realizations of these procedures. In neuroscience and robotics, researchers have explored the way that reward

systems might be adapted to incentivize new information as well as more standard types of reward, and have tried to characterize the neural bases of information-seeking and curiosity. In evolutionary biology researchers have used explore-exploit contrasts to characterize and explain differences among species, such as differences in life histories or differences in reliance on learning versus innate computations. And in cognitive development researchers have increasingly focused on children’s “active learning” the ways that they actively seek information in the environment. In particular, computational approaches quantifying information gain turn out to provide an interesting characterization of children’s exploratory behavior. Children are also characteristically exploratory – as evidenced in their play, for example, and childhood itself may be seen as a solution to explore-exploit trade-offs.

The classic examples of information-seeking involve extracting information from the external world, as in curiosity-driven exploration. However, exploration can take place internally as well as externally. Learners have to trade-off internal forms of exploration, such as those involved in mental simulations, or in exploring the range of possible hypotheses to explain a pattern of data. Activities like pretend play, hypothetical or counterfactual thinking, simulation, and even fiction can be understood as forms of this kind of internal exploration. These more internal forms of exploration also come with substantial computational costs, which must be traded off against the benefits of more accurate actions.

Questions about exploration and exploitation are also interestingly related to questions about variability and randomness. Exploration often involves the production of apparently random variability in behaviors or hypotheses. But different types of exploration may be more or less strategic or more or less random. At one extreme, a learner may carefully calculate exactly which actions are most

likely to lead to an increase in information and evaluate whether that information will be useful or relevant for a particular decision. At the other extreme, a learner may simply become bored or habituated and randomly pursue new policies or test new hypotheses and observe the consequences. Again there may be interesting trade-offs between these strategies, and between the computational work that is required for strategic exploration and its benefits. It also appears plausible that children become more strategic over development. But the relation between these two types of exploration and their costs and benefits have yet to be explored.

By bringing together researchers who would normally belong to different communities we hope to begin to answer this foundational set of questions about the nature of cognition.

### Workshop structure

The workshop will feature well-known experts from different fields. The workshop will also invite poster submissions from the broader cognitive science community, with “poster teasers” flash talks and an opportunity to view new studies investigating exploration-exploitation. Additionally, the schedule has built in ample time for questions for mini-panels of each sub-area of exploration-exploitation, ensuring maximum opportunity for audience engagement.

### Proposed schedule

#### 9:00-9:15: Opening Remarks

#### 9:15-10:45: Curiosity and Information-Seeking in Development

**Celeste Kidd:** *“The role of information gaps and learning rate on exploration”*

**Elizabeth Bonawitz:** *“Drivers of exploration: information, expectation, and affect”*

**Rista Plate:** *“How social information influences exploration and exploitation: Implications for learning”*

#### 10:45-11:00: Coffee break

#### 11:00-12:00: Evolutionary Perspectives on Explore-Exploit

**Emilie Snell-Rood:** *“Exploration in development: common principles from gene expression to neurons”*

**Willem Frankenhuis:** *“Balancing sampling and specialization: Adaptationist models of incremental cognitive development”*

#### 12:00-12:20: Poster teasers

#### 12:20-1:30: Lunch

#### 1:30-2:00: Poster viewing

#### 2:00-3:30 Explore-exploit in Reinforcement Learning and the Brain

**Cate Hartley:** *“The value of variability: Dynamic changes in reinforcement learning from childhood to adulthood”*

**Robert Wilson:** *“The role of information and randomization in exploration and exploitation”*

**Ben Hayden:** *“Explore-exploit tradeoffs as a driver of risky choice and learning in rhesus macaques”*

#### 3:30-3:45: Coffee break

#### 3:45-5:15: Explore-exploit in the Mind

**Thomas Griffiths:** *“Exploring on the inside: Rational models of the uses of cognitive resources”*

**Marcelo Mattar:** *“Prioritized exploration of mental simulations”*

**Alison Gopnik:** *“Childhood as a solution to explore-exploit tensions; changes in exploration and hypothesis search across life history”*

#### 5:15-5:30: Closing remarks

### Acknowledgments

The organizers would like to thank the Cognitive Science Society, the Jacobs Foundation (EB & CK), as well as NSF-SES 1627971 (EB).