Neural Computations Supporting Cognition: Rumelhart Prize Symposium in Honor of Peter Dayan

Participants

Kenji Doya (doya@oist.jap)
Neural Computation Unit, Okinawa Institute of Science and Technology, 1919-1 Tancha, Onna
Okinawa 904-0495 Japan

John O'Doherty (jodoherty@caltech.edu)
Division of the Humanities and Social Sciences, California Institute of Technology, MC 228-77
Pasadena, CA 91125 USA

Alexandre Pouget (alex@cvs.rochester.edu)
Département de neuroscience fondementale, Université de Genéve, 1 rue Michel-Servet
CH-1211 Geneva 4, Switzerland

Peter Bossaerts (pbs@hss.caltech.edu)
Division of the Humanities and Social Sciences, California Institute of Technology, MC 228-77
Pasadena, CA 91125 USA

Organizors

Nathaniel Daw (daw@cns.nyu.edu)
Center for Neural Science
New York University
New York, NY, 10003

Yael Niv (yael@princeton.edu)
Princeton Neuroscience Institute and Psychology Department
Princeton University,
Princeton, NJ, 08544

Keywords: neural computation; reinforcement learning; inference; uncertainty

Motivation

Principles of sound statistical inference underpin prominent accounts for a variety of cognitive phenomena, including perception, learning, and decision-making. Linking these building blocks of cognition to the biological substrate that supports them, recent work has investigated how the brain implements probabilistic inference and learning under uncertainty. The interplay between the psychological and biological levels of analysis has shed light on the structure of cognition and computation at both levels.

This symposium builds on Peter Dayan's seminal contributions to linking psychological, neural and computational phenomena. In particular, speakers will discuss recent work growing out of two areas where Dayan made early and fundamental contributions: the brain's mechanisms for reinforcement learning, and neural representations supporting probabilistic inference under uncertainty.

Reinforcement learning and the basal ganglia

Authors: Kenji Doya and Makoto Ito

Abstract: The discovery of the parallel between the firing of dopamine neurons and the temporal difference error signal of the reinforcement theory in the 1990s brought a breakthrough in understanding the function of the basal ganglia. Previously the most enigmatic part of the brain is now considered as the center for linking perception, action, and reward. After more than a decade from the discovery, however, there still remain questions to be answered, such as what striatal neuron firing represents, how and where an action is selected, and how negative reinforcement is realized. Here we review Peter Dayan's seminal contributions and recent developments.

Fractionating model-based reinforcement-learning its component neural processes

Author: John P. O'Doherty

Abstract: It has recently been proposed that action-selection in the mammalian brain depends on at least two distinct mechanisms: a model-free reinforcement learning (RL) mechanism in which actions are selected on the basis of cached values acquired through trial and error, and a model-based RL system in which actions are chosen using values computed on-line by means of a rich cognitive model of the decision problem and knowledge of the current incentive value of goals. While much is now known about the putative neural substrates of the model-free RL system and its concomitant temporal difference prediction error, much less is known about how model-based RL is implemented at the neural level. In this talk I will review recent evidence from a series of functional neuroimaging studies in humans supporting the presence of neural signals within a wide expanse of cortex that are relevant to model-based RL. These include, a state-action based prediction error signal within a fronto-parietal network that could mediate learning of the cognitive model, a goal-value signal encoding the value of putative goal-outcomes within the ventromedial prefrontal cortex, computations corresponding to action-contingency within inferior parietal cortex, and the representation of the effort costs of an action within the dorsomedial frontal cortex. These different computations then need to be integrated in order to construct an overall model-based action-value. Taken together, this evidence suggests that model-based reinforcement-learning theory provides a scaffold upon which a deeper understanding of the functions of a large extent of cortical territory within the mammalian brain can be built.

Probabilistic inferences in neural circuits using probabilistic population codes

Author: Alexandre Pouget

Abstract: A wide range of seemingly unrelated behaviors can be formalized as instances of probabilistic inferences. This includes odor recognition, sensorimotor transformations, decision making, simple arithmetics and visual search, to name just a few. We will present a neural theory of probabilistic inference in which neurons encode probability distributions using a basis function decomposition of the log probability or log likelihood. This approach makes very specific predictions about the form of the variability in neural responses, as well as about the neural implementation of various probabilistic inferences such as product of distributions, marginalization and sampling. We will discuss several experimental tests of these predictions in the context of arithmetics, visual search and bistable perception.

The neural process of subjective belief formation in humans

Author: Peter Bossaerts

Abstract: We present a general experimental paradigm with which to study human belief formation from experience. We first establish that human learning proceeds along Bayesian principles, but from subjective albeit robust priors rather than the true prior. Second, properly dissociating neural encoding of values and beliefs, we identify the default mode network as the locus of beliefs learned from frequencies. Third, we study the neural basis for combining objective frequentist information with prior beliefs and discover that Bayesian posterior beliefs are encoded bilaterally in the lateral prefrontal cortex.