Dopamine as a neural substrate of reward prediction and psychopathology

Nobel Mini-symposium

8th and 9th of September 2022

Nobel Forum, Karolinska Institutet, Stockholm, Sweden

Nobelförsamlingen
The Nobel Assembly at Karolinska Institutet

Karolinska Institutet

Schedule

Thursday the 8th of September

8:30-9:00	Registration
9:00-9:10	**Welcoming words** **Thomas Perlmann**, Professor at Karolinska Institutet and Secretary for the Nobel committee for Physiology and Medicine
9:10-9:20	**Introduction** **Marc Guitart-Masip**, Associate Professor at Karolinska Institutet
9:20-10:20	**Characteristics of the Dopamine Reward Prediction Error Signal: Reinforcement Learning and Economic Choice** **Wolfram Schultz**, Professor at the University of Cambridge, UK
10:20-11:00	**Coffee break**
11:00-12:00	**Toward and Untoward Influences on Dopamine** **Peter Dayan**, Professor at the University of Tübingen, Germany
12:00-13:00	**Lunch** *Registration to the mini-symposium does not include lunch*
13:00-14:00	**Decoding human neuromodulatory signaling and its connection to reinforcement learning** **Read Montague**, Professor at Virginia Tech, USA
14:00-15:00	**Dopamine prediction errors and hidden task structure** **Yael Niv**, Professor at Princeton University, USA
15:00-15:30	**Coffee break**
15:30-16:30	**Tuning striatal dopamine signals to optimize reinforcement learning across tasks** **Michael Frank**, Professor at Brown University, USA

Friday the 9th of September

9:00-10:00	**Role of dopamine in human learning, motivation and control**	**Roshan Cools**, Professor at Radboud University Medical Center, The Netherlands
10:00-10:40	**Coffee Break**	
10:40-11:40	**Prediction error encoding in mental disorders**	**Andreas Heinz**, Professor at Charité - University Medical Center, Germany
11:40-12:00	**Concluding remarks**	**Predrag Petrovic**, Associate Professor at Karolinska Institutet

Registration

Due to limited space, only registered participants can attend the symposium.

In order to register, mail **Britt-Marie Höglund (email: hoglundbm@gmail.com)** no later than 26th of August 2022. You can register for either one or both days.

We would appreciate that you only register if you are sure you will be attending and unregister if you cannot make it (to the same email).

The Nobel Mini-symposium will be posted online and accessible for the general public after the event.

Wolfram Schultz

Professor at the University of Cambridge, UK

Wolfram Schultz is a graduate in medicine from the University of Heidelberg. After post-doctoral stays in Germany, USA and Sweden, and a faculty position in Switzerland, he works currently at the University of Cambridge. He combines behavioural, neurophysiological and neuroimaging techniques to investigate the neural mechanisms of learning, goal-directed behaviour and economic decision making. He uses behavioural concepts from animal learning theory and economic decision theories to study the neurophysiology and neuro-imaging of reward and risk in individual neurons and in specific brain regions, including the dopamine system, striatum, orbitofrontal cortex and amygdala.

Characteristics of the Dopamine Reward Prediction Error Signal: Reinforcement Learning and Economic Choice

The phasic dopamine reward prediction error (RPE) signal is suitable for updating reward predictions that can be used for making informed economic choices. As the basics of the dopamine RPE signal are fairly well known, we now addressed the issue whether this signal reflects reward value in a meaningful way. Reward value is specific for the individual biological decision maker and thus is subjective, and so should be meaningful neuronal reward signals (although subjectivity is immaterial for computer algorithms of reinforcement learning). We therefore tested dopamine signals in tasks informed by economic concepts. Monkeys made meaningful choices that complied with first, second and third order stochastic dominance, suggesting well-ordered processing of reward value according to goodness and risk. Based on this necessary background, we estimated utility functions from choice under risk (Von Neumann-Morgenstern utility). We found that the dopamine RPE signal was scaled in the metric of utility, thus constituting a utility prediction error (UPE) signal. Consistent with this characteristic, the dopamine response complied with second-order stochastic dominance that tests the incoporation of risk into subjective value. These data demonstrate that the phasic dopamine RPE signal incorporates the basic requirements for a biologically meaningful reinforcement signal. The data unite concepts from animal learning theory and economic decision theory at the level of single reward neurons.

Peter Dayan

Professor at the University of Tübingen, Germany

Peter Dayan was an assistant professor in the Department of Brain and Cognitive Sciences at MIT, and was a founding faculty member of the Gatsby Computational Neuroscience Unit at UCL, which he then ran for 15 years. He is currently a Director at the Max Planck Institute for Biological Cybernetics and a Professor at the University of Tübingen. His interests include affective decision making and neural reinforcement learning.

Toward and Untoward Influences on Dopamine

There is substantial evidence that the activity of many dopamine neurons and the release of dopamine in target structures code temporal difference prediction errors for reward, with due effects over Pavlovian predictions and instrumental behaviour, However, there is also important heterogeneity in this activity, and a variety of other factors that apparently influence dopamine activity and release, including counterfactual reward, punishment, reward-free information and even forms of fictive rewards in response to outcomes that are not allostatically beneficial. I will discuss various computational sources of these seemingly anomalous signals, showing that many of them are prediction errors in disguise.

Read Montague

Professor at Virginia Tech, USA

Read Montague earned a doctoral degree in biophysics from the University of Alabama at Birmingham School of Medicine. He completed a fellowship in Theoretical Neurobiology at the Neurosciences Institute at Rockefeller University, mentored by Prof. Gerald Edelman and subsequently was a HHMI fellow in the Computational Neurobiology Lab led by Prof. Terrence Sejnowski. Montague's research focuses on the connection between physical mechanisms present in real neural tissue and the computational functions that these mechanisms embody. From 2005-2006 he was a member of the Institute for Advanced Study at Princeton and he was a Principal Research Fellow at Wellcome Centre for Human Neuroimaging from 2011-2018. He was the founding director of the Human Neuroimaging Lab and Computational Psychiatry Unit at Baylor College of Medicine in Houston, before moving to his current post as founding director of the Center for Human Neuroscience Research at the Fralin Biomedical Research Institute at Virginia Tech Carilion.

Decoding human neuromodulatory signaling and its connection to reinforcement learning

In recent years, a kind of renaissance of reinforcement learning applications has occurred separately in two domains. The first is the deployment of reinforcement learning systems to learning problems that seem on their face to be extremely complicated, and thus amaze our sensibilities about the kinds of problems solvable by such 'simple' learning systems. The second is the use of the 'reinforcement learning platform' to inform and interpret neurobiological experiments that seek to connect reinforcement learning models to various features of neuromodulatory signaling. In rodents, these latter applications abound and now make use of a collection of superb molecular tools for fast, selective detection of important neuromodulatory transmitters or selective stimulation of their parent neurons. This rodent work includes fast measurements of dopamine, noradrenaline, acetylcholine and more. There has been no comparable work in humans, which opens up a gap for understanding the neuromodulatory substrates of human decision-making. In this talk, I will present new methodology for sub-second detection of monoamines in conscious humans and demonstrate its deployment in a range of simple valence and arousal processing tasks.

Yael Niv

Professor at Princeton University, USA

Yael Niv is a professor of neuroscience and psychology at Princeton University. Her lab studies the computational processes underlying reinforcement learning, focusing on how attention, memory and learning interact to construct task representations that allow efficient learning through optimal generalization. Niv is co-founder and co-director of the Rutgers-Princeton Center for Computational Cognitive Neuropsychiatry, where she is applying ideas from reinforcement learning to understanding and treating mental illness. Her proudest career accomplishment is winning a graduate mentoring award. In her nonexistent spare time, she is a mom to two awesome boys, and an activist within and outside academia.

Dopamine prediction errors and hidden task structure

Phasic dopamine responses are thought to encode a prediction-error signal – the difference between expected and received rewards. These same errors are then used to update the predictions. However, predictions are tied to context, to the current ”state” of the world. These states, critical in reinforcement learning theory and in models of dopamine function, have to be learned themselves, and are not trivial as they are often partly observable and must be inferred. In this talk, I will focus on prediction errors in the context of temporally-extended states, how we infer that one state has ended and another has begun, and when do we infer brand new states in light of new observations. I will connect these ideas to mental health implications when we infer too few, or too many environmental states.

Michael J Frank

Professor at Brown University, USA

Michael J. Frank directs the Center for Computational Brain Science within the Carney Institute for Brain Science. He received his PhD in Neuroscience and Psychology in 2004 at the University of Colorado, following undergraduate and master's degrees in electrical engineering. Frank's work focuses primarily on theoretical models of frontostriatal circuits and their modulation by dopamine, especially their cognitive functions and implications for neurological and psychiatric disorders. The models are tested and refined with experiments across species, neural recording methods, and neuromodulation. Honors include the Troland Research Award from the National Academy of Sciences (2021), Kavli Fellow (2016), the Cognitive Neuroscience Society Young Investigator Award (2011), and the Janet T Spence Award for early career transformative contributions (Association for Psychological Science, 2010). Dr Frank is a senior editor for eLife.

Tuning striatal dopamine signals to optimize reinforcement learning across tasks

The basal ganglia and dopaminergic (DA) systems are well studied for their roles in reinforcement learning. A wealth of evidence supports the dopamine-RPE hypothesis, but a closer look at the biology suggests an architecture that differs from that typically assumed in artificial agents, whereby RPEs are scalar and globally influence downstream targets. Computational and empirical considerations suggest that DA signals may be enriched beyond scalar signals to support structured learning in striatal circuits. First, the nonlinear effect of RPEs within opponent striatal D1 vs. D2 pathways, together with dynamic DA levels that evolve with reward history, enables preferential credit to the pathway best specialized for reward statistics of the given task. Models that incorporate such mechanisms provide robust advantages over traditional monolithic RL models over a range of environments, and suggest that empirical observations of altered learning and decision-making inpatient populations reflect a byproduct of an otherwise normative mechanism. Second, dopamine RPE signals are not global and synchronous across the striatum. Rather, spatiotemporal dynamics (in the form of traveling waves of dopaminergic activity) provide a mechanism to support credit assignment to appropriate striatal subregions. I will present experimental data and computational models suggesting that such dynamics enable agents to reinforce striatal subregions that are most well tuned to the generative task statistics, facilitating adaptive structured behavior.

Roshan Cools

Professor at Radboud University Medical Center, The Netherlands

Roshan Cools is is an expert in the chemical neuromodulation of human cognition and motivation, focusing on the functions of dopamine and serotonin. She completed her undergraduate degree in Experimental Psychology at the University of Groningen, Netherlands, in 1998, then moved to the University of Cambridge, UK, for an MPhil degree (1999), a PhD degree (2002) and Royal Society Fellowship funded postdoctoral work until 2007. She also spent two post-doc years at UC Berkeley working with Mark D'Esposito from 2003. In November 2007 she returned to The Netherlands, where she now runs her Motivational and Cognitive Control group. Cools is a member of the Royal Netherlands Academy of Arts and Sciences (KNAW), the Academia Europea and the (Dutch Government) Advisory Council for Science, Technology, and Innovation (AWTI, since 2014). She holds an ERC Advanced grant, a Vici award from NWO and an EBRAINS Lead Scientist Voucher from the Human Brain Project.

Role of dopamine in human learning, motivation and control

The human brain faces a variety of computational dilemmas, including the flexibility/stability, the speed/accuracy and the labor/leisure tradeoff. Given its role in reinforcement learning and motivation, striatal dopamine is particularly well suited to dynamically regulate these computational tradeoffs depending on constantly changing task demands. To illustrate this, I will discuss evidence from studies on learning, motivation and cognitive control in human volunteers, using chemical PET imaging, psychopharmacology, and/or fMRI.

These studies also begin to elucidate the mechanisms underlying the large variability in dopaminergic drug effects across different individuals and across different task contexts.

Andreas Heinz

Professor at Charité - University Medical Center, Germany

Andreas Heinz, MD PhD, is professor of psychiatry and the director and chair of the department of Psychiatry and Neurosciences, Charité - University Medical Center, CCM Berlin. He studied medicine, philosophy and anthropology at the Ruhr Universität Bochum, Freie Universität Berlin and Howard University, Washington DC. His main areas of research include addiction and schizophrenia as well as intercultural psychiatry and psychotherapy. He is a member of the National Academy of Sciences, Leopoldina, and the Mainz Academy of Literature and Science.

Prediction error encoding in mental disorders

Dopamine dysfunction is implicated in several mental disorders including addiction and psychosis. All drugs of abuse release dopamine. In psychosis, stress-related or chaotic phasic release of dopamine has been implicated in delusion formation. These observations have been explained by phasic dopamine encoding stimulus novelty and errors of reward prediction. In drug addiction, drug intake can thus elicit dopamine release and attribute salience to cues and contexts associated with drug consumption. In psychosis, imprecise cortical information processing may be compensated by elevated phasic dopamine release, which increases the signal-to-noise ratio, albeit at the price of attributing salience to otherwise irrelevant cues. Functional MRI studies confirmed alterations in phasic activation of the ventral striatum elicited by reward prediction errors in accordance with the hypotheses outlined above. Time resolution of in vivo dopamine imaging is limited in humans, so associations of errors of reward prediction with dopamine neurotransmission are largely correlational. Nevertheless, multimodal imaging studies confirm loss of associations between dopamine and reward prediction encoding in patients with psychosis and addiction compared to controls. Clinically, these studies have implications for the pathogenesis of addictive behavior and psychotic experiences, and they suggest to dose neuroleptic medication as low as possible in order to avoid impairments of reward prediction encoding and motivation.

This Mini-symposium is organized by Predrag Petrovic and Marc Guitart-Masip, Karolinska Institutet and sponsored by the Nobel Assembly at the Karolinska Institutet.

ki.se

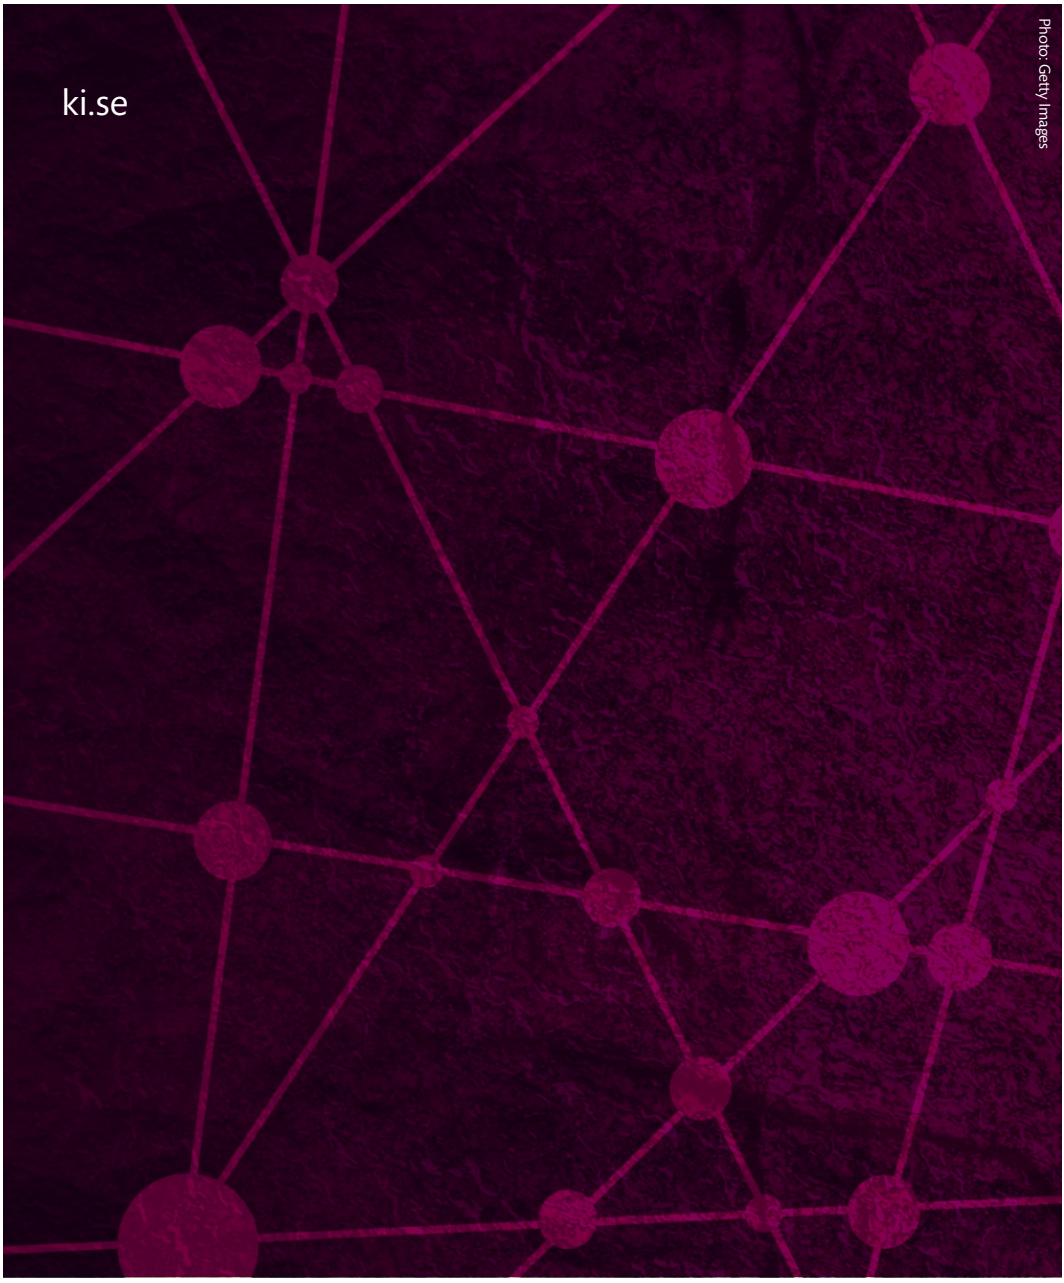

Photo: Getty Images

Nobelförsamlingen
The Nobel Assembly at Karolinska Institutet

Karolinska Institutet