THURSDAY, 09/22/2022

Interdisciplinary Science and Engineering Building (ISEB), Room 1200

Session 1 - Neural Representations of Space, Time, and Sequences

09:10 am – 09:45 am

Norbert Fortin
UC Irvine

Linking Sequential Representations and Dynamics in the Hippocampus with Task-Critical Series of Cognitive Operations During Behavior

Technical advances in recent years have given neuroscientists the ability to record massive and multidimensional neural activity datasets during behavior. While this is expected to lead to an unprecedented rate of breakthroughs in our understanding of how the brain forms and organizes our memories, and how we use those memories to guide our behavior, making sense of this data deluge is a considerable statistical challenge. I will begin by outlining the neuroscience objective of the session, which in this case is to identify key neuroscience questions that require new analytical tools to be answered, and the area of expertise of each speaker in the session. I will then discuss our recent work as an example of a successful application of an interdisciplinary approach combining teams of neuroscientists and data scientists, in which statistical machine learning methods were developed to answer fundamental questions about the role of the hippocampus in the temporal organization of memory and behavior.

09:50 am – 10:25 am

Adrien Peyrache
McGill University

Attractor Dynamics in the Head-Direction System and Current Analytical Challenges in Quantifying the Organization of Population Activity

Continuous attractor neural networks, which maintain population activity in a small number of possible states, are believed to support various cognitive functions, from working memory to spatial representations. One example of such networks is the head-direction (HD) circuit, a crucial signal for navigation. During sleep, HD neurons of the thalamus remain coordinated, coding for a “virtual” direction. This observation supports the idea that the HD network is driven by attractor dynamics, yet the origin and function of this coherent population activity during sleep remain unknown. Here, I will present evidence that the thalamic HD cells remain coordinated independent of their inputs from the mammillary bodies during sleep. This observation invalidates the hypothesis of a unique generator in the mammillary bodies and confirms that the thalamus is not a mere relay of sensory signals. However, analyzing this kind of data present many challenges, starting from the number of neurons required to reveal the intrinsic topology of population activity.

10:40 am – 11:15 am

Marc Howard
Boston University

Continuous Neural Representations of Time, Space, and Number

Computational models of human memory tasks have long utilized a distributed representation of temporal context. Over the last decade, a body of work including studies from rodents, monkeys, and humans has shown that so-called ``time cells'' in the hippocampus (and other brain regions) have properties like those required for temporal context. The recent discovery of so-called ``temporal context cells'' in the entorhinal cortex provides dramatic support for a detailed theoretical proposal. Theoretically, time can be seen as a special case of a more general computation. The same equations that generate temporal context cells and time cells (and thus a representation of temporal context) can also generate cells with a broad variety of spatial correlates which together constitute a spatial context. The same equations can also be used to construct decision-making circuits that on the one hand map onto known properties of neurons in evidence accumulation tasks and widely-used cognitive models for response time. Thus, these equations provide a bridge between populations of neurons and models for cognition. They apply broadly throughout the brain. We discuss the need for a general data-analytic tool to find neural representations obeying these equations. To this end, we discuss how to understand the neural circuits in the context of recurrent networks and illustrate properties of the covariance matrices that result from populations of neurons obeying these equations.

11:20 am – 11:55 am

Zoran Tiganj
Indiana University Bloomington

Learning With Mental Maps

Converging evidence from cognitive science and neuroscience suggests that the brain represents physical and abstract variables such as distance, time, and numerosity in a structured form, as mental or cognitive maps. These maps seem to be supported by neurons tuned to a specific magnitude of the represented variable, such as place cells for spatial distance and time cells for elapsed time. Neural recordings also reveal cells with mixed selectivity that are simultaneously tuned to a specific combination of magnitudes of different variables, such as cells tuned to a particular spatial distance and numerosity and cells tuned to a particular amount of elapsed time and stimulus identity. We will discuss neural mechanisms that could give rise to such activity profiles, as well as the potential computational benefits of this rather expensive coding scheme. We will illustrate how mental maps can be combined with deep learning approaches. We will also discuss the need for tools that can provide robust identification and characterization of these cells. Experimental data show strong heterogeneity in firing profiles of individual cells, requiring analysis techniques that can account for different sources of variance.

Session 2 - Neural Mechanisms of Perception, Learning, and Inference

01:00 pm – 01:35 pm

Aaron Bornstein
UC Irvine

Inferring Latent Structure During Latent Structure Inference
As we navigate the world around us, we become aware of regularities linking events, objects, and concepts. Though these associations may be *latent*, in the sense of not immediately affecting our behavior, they can help us make better decisions down the road, especially in new settings where well-worn habits don't apply. Researchers in many areas of cognitive neuroscience are examining how humans learn this sort of structure during tasks ranging from simple choice to spatial navigation to the understanding of social hierarchies. There is a growing understanding that this question is of fundamental importance because it addresses potentially core cognitive computations that appear to generalize widely across tasks within the laboratory, and may thus be more predictive of phenomena outside of it. However, latent structure learning presumes the existence of latent representations of this structure, which are by their nature more difficult to uniquely identify than learning that has immediate behavioral expression. Humans vary in their attention to, learning of, and memory for, the features that inform these representations. Thus, a cohesive approach to discovering latent representations must account for these influences, a problem that requires advances in both methods and theory. I will first introduce this problem, then briefly review a series of studies that examine how humans acquire and represent varieties of associative structure in serial learning and decision-making tasks. Subsequent speakers in this session will discuss their approaches to this problem in distinct domains, with similarly diverse methods and theoretical frameworks. I will close with a discussion of some commonalities underlying these lines of research, and the problems and challenges that we face in unifying them.

01:40 pm – 02:15 pm

Angela Radulescu
Mount Sinai School of Medicine

Towards Naturalistic Reinforcement Learning in Health and Disease
Humans learn more from their experiences than just how to behave in different situations. They also learn to organize experiences into internal representations that facilitate flexible behavior, in domains ranging from simple decision-making to goal-directed action in naturalistic, richly structured environments. In this talk, I will show that such representation learning relies on selective attention to constrain the dimensionality of environments that humans learn from; and that attention is in turn guided by memory-guided inference over what features of the environment are relevant for the task at hand.

In addition, I will talk about some of my new and ongoing trying to model gaze data in VR, framing it as a challenge for generalizing these models to naturalistic domains — i.e., what methods would we need to be able to do that?

02:30 pm – 03:05 pm

Megan Peters
UC Irvine
Challenges and Promise of Current Approaches to the Human Neuroscience of Perceptual Metacognition
Humans have an incredible capacity to extract information from our environments and use it in flexible, adaptive, and goal-driven ways. We also have powerful capacities to evaluate our own ability to use information: We know when we're behaving well and achieving our goals, and we know when we're doing poorly. How do we do this, what neural representations and computations are involved? In this talk I'll discuss ongoing work by my group that works to evaluate the types of information representation and processing that humans use to make decisions about their visual environments and evaluate those decisions in adaptive ways, and what challenges we face in answering these questions with human neuroscience. I'll cover computational models and computational approaches to human neuroimaging processing, as well as what we can learn from how we perform these tasks that might be of interest to the development of smarter and more capable artificial systems.

03:10 pm – 03:45 pm

Jeremy Manning
Dartmouth College

Improving Real-World Learning Using Scalable Automated Teachers
What are the most effective ways to teach, and how quickly can humans learn? To begin to answer these questions, it can be instructive to think back to your personal learning journey. Who has been your most effective teacher? What was your favorite thing to learn about? Which concepts or skills were more difficult for you to learn, and how did you overcome those difficulties? Broadly, the answers to these questions likely depend on what specific content is being learned (e.g., introductory coding in Python versus medieval European history), properties of the individual learner (e.g., a kindergartener versus a graduate student), and myriad other factors. My lab is working to build scalable general purpose solutions to the "teaching problem" -- i.e., teaching any person, any content, as quickly and effectively as possible. Our approaches to optimizing instruction center on characterizing (a) the knowledge and skills the learner has already acquired, (b) the set of concepts the learner still needs to know about or master, and (c) the optimal timings and formats for presenting new material. In addition to drawing on my own experiences as an educator, this line of research incorporates advances in cognitive modeling, natural language processing, natural language understanding, human-computer interaction, and other related fields. I'll talk about some of our prior and ongoing work (and some of our most pressing challenges!) on developing automated teachers that can be scaled to support millions of students and thousands of content areas.