Mind & Machine Virtual Summit Agenda

UC Santa Barbara
February 16 – 17, 2021

February 16, 2022 — Neuroscience to Advance AI

9:00am	**Welcome and Opening Remarks** Miguel Eckstein and William Wang, UC Santa Barbara
9:05am	***Keynote*: Agile Movement & Embodied Intelligence: Computational and Comparative Considerations** Bing Brunton, Biology, University of Washington, Seattle
10:05am	**Introduction to Session 1: Neuroscience to Advance AI** Moderators: Michael Beyeler, Computer Science & Psychological & Brain Sciences, UC Santa Barbara William Wang, Computer Science, UC Santa Barbara
10:10am	**The Neural Computations Underlying Human Social Interaction Perception** Leyla Isik, Cognitive Science, Johns Hopkins University
10:35am	**Inductive Bias of Neural Networks** Cengiz Pehlevan, Applied Mathematics, Harvard University
11:00am	**Knowledge and Data in Neuro-Symbolic Learning** Guy Van den Broeck, Computer Science, UC Los Angeles
11:25am	**Panel Discussion and Q&A**
12:05pm	Meeting adjourns

February 17, 2022 — AI to Advance Brain Understanding

9:00am	**Welcome to Day 2** William Wang and Miguel Eckstein, UC Santa Barbara
9:05am	**Introduction to Session 2: AI to Advance Brain Understanding** Moderators: Scott Grafton, Psychological & Brain Sciences, UC Santa Barbara Michael Goard, Molecular, Cellular, and Developmental Biology & Psychological & Brain Sciences, UC Santa Barbara
9:10am	**Language in Brains and Algorithms** Jean-Rémi King, Facebook AI Research | Ecole Normale Supérieure
9:35am	**Towards Robust Representations of Neural Activity** Eva Dyer, Biomedical Engineering, Georgia Tech
10:00am	**Neural Circuits of Cognition in Artificial and Biological Neural Networks** David Freedman, Neurobiology, The University of Chicago
10:25am	**Panel Discussion and Q&A**

UC SANTA BARBARA Mind & Machine Intelligence
Center for Responsible Machine Learning
ICB
The SAGE Center for the Study of the Mind • UCSB

February 17, 2022 — Learning and Representations in AI and Brains

11:05am **Introduction to Session 3: Learning and Representations in AI and Brains**
Moderators: Miguel Eckstein, Psychological and Brain Sciences, UC Santa Barbara
Ikuko Smith, Molecular, Cellular, and Developmental Biology & Psychological & Brain Sciences, UC Santa Barbara

11:10am **Self-Motivated and Self-Supervised Open-World Continual Learning**
Bing Liu, Computer Science, University of Illinois at Chicago

11:35am **Stimulus-Dependent Representational Drift in Primary Visual Cortex**
Michael Goard, Molecular, Cellular, and Developmental Biology, UC Santa Barbara

12:00am **Deep Recurrent Neural Networks as a Modelling Framework for Understanding Human Vision**
Tim Kietzmann, Donders Institute for Brain, Cognition, and Behavior

12:25am **Panel Discussion and Q&A**

1:00pm Meeting adjourns

February 16, 2022 — Neuroscience to Advance AI

Keynote: Bing Brunton, Biology, University of Washington, Seattle
Agile Movement & Embodied Intelligence: Computational and Comparative Considerations

Abstract: I will tell a story in 3 parts with the broad theme of developing data-intensive approaches to connect brain and behavior. They all highlight my love of natural behaviors, dynamical systems, and open science. These stories feature several fun collaborations, including joint work with experimental neurobiologists, mathematicians, and engineers. In particular, I will describe several research threads asking how insects perform dexterous, coordinated movements in uncertain, complex environments. This ability is enabled by the sensation of mechanical forces to inform rapid corrections in body orientation. Curiously, mechanoreceptor neurons do not faithfully report forces; instead, they are activated by specific time-histories of forcing. In a set of results that combine biomechanics, sparse sensing mathematical theory, and neural encoding, we find that, far from being a bug, neural encoding by biological sensors is a feature that acts as a transformation superbly matched to detect body rotation. Indeed, this encoding further enables surprisingly efficient detection using only a small handful of neurons at key locations. Our ongoing and future work explores biological principles that achieve agile movements, including understanding and designing neural-inspired sensors, actuators, and controllers.

Bio: Bing Wen Brunton is an Associate Professor and H. Stewart Parker Faculty Fellow at the University of Washington in Seattle, Department of Biology, with additional affiliations at the eScience Institute for Data Science, the Paul G. Allen School of Computer Science & Engineering, and the Department of Applied Mathematics. Her lab develops data-driven analytic methods that are applied to, and are inspired by, neuroscience.

[Leyla Isik](), Cognitive Science, Johns Hopkins University

The Neural Computations Underlying Human Social Interaction Perception

Abstract: Humans perceive the world in rich social detail. We effortlessly recognize not only objects and faces in our environment, but also other peoples' social interactions. The ability to perceive others' social interactions is critical for social trait judgement and ultimately guides how humans act in the social world. We recently identified a region that selectively represents others' social interactions in the posterior superior temporal sulcus (pSTS) using controlled experiments with simple stimuli. However, it is unclear how social interactions are processed in the real world where they co-vary with many other sensory and social features. In this talk I will discuss new work using naturalistic video paradigms and novel machine learning analyses to understand how humans process social interactions in natural settings. We find that social interactions guide behavioral judgements and are selectively processed in the brain, even after controlling for the effects of other visual and social information. Finally, I will discuss the computational implications of humans' social interaction selectivity and how we can develop artificial systems that share this core human ability.

Bio: Leyla Isik is the Clare Boothes Luce Assistant Professor of Cognitive Science at Johns Hopkins University. Her research aims to understand humans' vast visual and social abilities using a combination of human neuroimaging and machine learning. Before joining JHU, she was a postdoctoral researcher at MIT and Harvard in the Center for Brains, Minds, and Machines working with Nancy Kanwisher and Gabriel Kreiman. She completed her PhD at MIT with Tomaso Poggio.

[Cengiz Pehlevan](), Applied Mathematics, Harvard University

Inductive Bias of Neural Networks

Abstract: A learner's performance depends crucially on how its internal assumptions, or inductive biases, align with the task at hand. I will present a theory that describes the inductive biases of neural networks in the infinite width limit using kernel methods and statistical mechanics. This theory elucidates an inductive bias to explain data with "simple functions" which are identified by solving a related kernel eigenfunction problem on the data distribution. This notion of simplicity allows us to characterize whether a network is compatible with a learning task, facilitating good generalization performance from a small number of training examples. I will demonstrate that the same theory also applies to linear readout from neuronal population codes, elucidating how population codes enable learning from few examples by shaping inductive bias. I will present applications of the theory to deep networks (at finite width) trained on synthetic and real datasets, and recordings from the mouse primary visual cortex. Finally, I will briefly present an extension of the theory to out-of-distribution generalization.

Bio: Cengiz (pronounced "Jen-ghiz") Pehlevan is an assistant professor of applied mathematics at Harvard SEAS. His research interests are in theoretical neuroscience and theory of neural computation. Cengiz comes to Harvard SEAS from the Flatiron Institute's Center for Computational Biology (CCB), where he was a a research scientist in the neuroscience group. Before CCB, Cengiz was a postdoctoral associate at Janelia Research Campus, and before that a Swartz Fellow at Harvard. Cengiz received a doctorate in physics from Brown University and undergraduate degrees in physics and electrical engineering from Bogazici University, Turkey.

Guy Van den Broeck, Computer Science, UC Los Angeles
Knowledge and Data in Neuro-Symbolic Learning

Abstract: The key challenge in neuro-symbolic machine learning is to bridge between perception from data and reasoning about symbolic knowledge. In this talk, I will present recent work that finds novel ways to unify these two worlds, by using logical and probabilistic reasoning tools (circuits and SAT solvers) to improve the learning capabilities of deep neural networks.

Bio: Guy Van den Broeck is an Associate Professor and Samueli Fellow at UCLA, in the Computer Science Department, where he directs the Statistical and Relational Artificial Intelligence (StarAI) lab. His research interests are in Machine Learning, Knowledge Representation and Reasoning, and Artificial Intelligence in general. His papers have been recognized with awards from key conferences such as AAAI, UAI, KR, OOPSLA, and ILP. He also serves as Associate Editor for the Journal of Artificial Intelligence Research (JAIR). Guy is the recipient of an NSF CAREER award, a Sloan Fellowship, and the IJCAI-19 Computers and Thought Award.

February 17, 2022 — AI to Advance Brain Understanding

Jean-Rémi King, Facebook AI Research | Ecole Normale Supérieure
Language in Brains and Algorithms

Abstract: Deep learning has recently made remarkable progress in natural language processing. Yet, the resulting algorithms fall short of the language abilities of the human brain. To bridge this gap, we here explore the similarities and differences between these two systems using large-scale datasets of magneto/electro-encephalography (M/EEG, n=1,946 subjects), functional Magnetic Resonance Imaging (fMRI, n=589), and intracranial recordings (n=176 patients, 20K electrodes). After investigating where and when deep language algorithms map onto the brain, we show that enhancing these algorithms with long-range forecasts makes them more similar to the brain. Our results further reveal that, unlike current deep language models, the human brain is tuned to generate a hierarchy of long-range predictions, whereby the fronto-parietal cortices forecast more abstract and more distant representations than the temporal cortices. Overall, our studies show how the interface between AI and neuroscience clarifies the computational bases of natural language processing.

Bio: Jean-Rémi King is a CNRS researcher at the Ecole Normale Supérieure detached to Meta AI. Jean-Rémi's group works on understanding the brain and computational bases of human intelligence, with a focus on language. They develop deep learning techniques to decode and model intracranial recordings, magneto-encephalography and functional magnetic resonance imaging.

Eva Dyer, Biomedical Engineering, Georgia Tech
Towards Robust Representations of Neural Activity

Abstract: Even when behavior is stable, neural responses in many areas of the brain are highly variable. Variability in neural activity makes it extremely challenging to build models that generalize to new time points or different sets of neurons. In this talk, I will describe ways in which my lab is tackling this challenge, focusing on our recent efforts to use principles underlying self-supervised learning (SSL) to build invariances into representations of brain states. In SSL, invariances are achieved by encouraging "augmentations" (transformations) of the input to be mapped to similar points in the latent space. We show how this guiding principle can be used to model populations of neurons in diverse brain regions and disentangle different sources of information, all without labels. Our work makes important steps towards the development of new approaches for building models of neural activity that can generalize across time and different sets of neurons, and potentially be used to compare neural representations across different individuals.

Bio: Eva L. Dyer is an Assistant Professor in the Department of Biomedical Engineering at the Georgia Institute of Technology. Dr. Dyer's research combines machine learning and neuroscience to understand the brain, its function, and how neural activity is shaped by disease.

David Freedman, Neurobiology, The University of Chicago
Neural Circuits of Cognition in Artificial and Biological Neural Networks

Abstract: We have a remarkable ability to interpret incoming sensory stimuli and plan task-appropriate behavioral responses. This talk will present parallel experimental and computational approaches aimed at understanding the circuit mechanisms and computations underlying flexible perceptual and categorical decisions. In particular, our work is aimed at understanding how feature encoding in visual cortex is transformed across the cortical hierarchy into flexible task-related encoding in the parietal and prefrontal cortices. The experimental studies utilize multielectrode recording approaches to monitor activity of neuronal populations, as well as reversible cortical inactivation approaches, during performance of visually-based decision making tasks. In parallel, our computational work employs machine learning approaches to train recurrent artificial neural networks to perform the same tasks as in the experimental studies, allowing a deeper investigation of putative neural circuit mechanisms used by both artificial and biological networks to solve cognitively demanding behavioral tasks.

Bio: David Freedman is a Professor at The University of Chicago in the Department of Neurobiology and Neuroscience Institute. His research focuses on the brain mechanisms of higher cognitive functions through large scale electrophysiological recordings of neural activity, as well as translating findings from neuroscience to advance artificial intelligence and machine learning.

February 17, 2022 — Learning and Representations in AI and Brains

Bing Liu, Computer Science, University of Illinois at Chicago
Self-Motivated and Self-Supervised Open-World Continual Learning

Abstract: Humans learn continuously. Novelties or unknowns serve as an intrinsic motivation for our learning. As more and more AI agents are used in practice, it is time to think about how to make these agents learn autonomously by themselves in a self-motivated and self-supervised manner after they are deployed. As the real-world is an open and dynamic environment, the agent needs to detect novelties or unknowns and distribution shifts, characterize and accommodate the novelties, and gather ground-truth training data and incrementally or continually learn them to make it more and more knowledgeable and powerful over time. The key challenge is how to automate the process so that it can be carried out on the agent's own initiative and through its own interactions with humans and the environment rather than on the initiative of human engineers through periodical re-training. In this talk, I will first present this paradigm and then discuss some of our recent work related to the topic.

Bio: Bing Liu is a distinguished professor at the University of Illinois at Chicago. His current research interests include lifelong/continual learning, continual learning dialogue systems, sentiment analysis, machine learning and data mining. Three of his papers have received Test-of-Time awards and another one received Test-of-Time honorable mention. He is the winner of 2018 ACM SIGKDD Innovation Award, and a Fellow of AAAI, ACM, and IEEE.

Michael Goard, Molecular, Cellular, and Developmental Biology, UC Santa Barbara

Stimulus-Dependent Representational Drift in Primary Visual Cortex

Abstract: In order to recognize familiar visual scenes, we assume that our visual cortical neurons must represent visual information faithfully over time. However, recent studies have shown that many brain structures, including sensory regions, exhibit "representational drift"—with neurons changing their response properties to the same stimuli over days to weeks. In this presentation, I will present recent data from the lab in which we used chronic two-photon calcium imaging to track thousands of visual cortical neuron responses to both drifting gratings and natural movies. Surprisingly, we found that while responses to drifting gratings are highly stable, responses to natural movies exhibit progressive representational drift. The observed drift is present across cortical layers and cell types, and cannot be explained by changes in receptive field or behavioral variables. Moreover, we found that the population correlation structure changes progressively across sessions to a greater extent for natural movies than for gratings. As a result, we propose that subnetworks of co-tuned neurons preserve fundamental response properties (like orientation tuning) while allowing more flexible responses to naturalistic stimuli. I will conclude by discussing the ramifications of these findings for the storage of sensory representations in biological and artificial networks.

Bio: Michael Goard received his B.A. in Psychology from Reed College and his Ph.D. in Neuroscience from UC Berkeley. As a postdoctoral fellow at the Massachusetts Institute of Technology, he developed approaches for imaging and manipulating the activity of large populations of neurons in behaving mice. He joined the faculty of UC Santa Barbara in 2016, where his lab investigates the neural circuitry underlying our perceptual, spatial, and cognitive abilities.

Tim Kietzmann, Donders Institute for Brain, Cognition, and Behavior

Deep Recurrent Neural Networks as a Modelling Framework for Understanding Human Vision

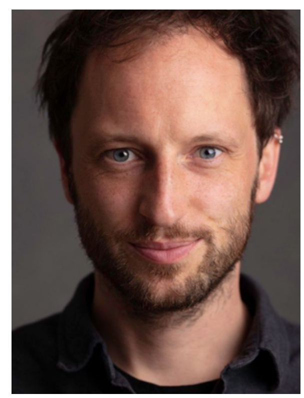

Abstract: This talk will describe our recent advances in understanding information processing in the human brain and artificial vision systems. Operating in the interdisciplinary research area in between deep learning, neuroimaging, and theory, I will here focus on the role and benefits of recurrent, i.e. top-down and lateral, connectivity in deep neural network models. I will demonstrate that neural network architectures with recurrent connectivity (a) provide better models of human cortical dynamics, estimated via representational dynamics across multiple visual areas, (b) are better aligned with human behaviour by allowing for speed-accuracy-tradeoffs, and (c) outperform feedforward networks in terms of categorization accuracy and parametric efficiency. Finally, we will look at how recurrent neural networks treat category-orthogonal object features. I will make the case that, despite being non-diagnostic, such auxiliary variables are not discarded, but rather integrated into recurrent inference to guide and improve performance over time.

Bio: Tim Kietzmann works at the intersection of Machine Learning and Neuroscience, focussing in particular on (visual) information processing in man and machine. After a PhD in Cognitive Science, he moved to the University of Cambridge, and later opened his own lab at the Donders Institute (Radboud University). Recently, he accepted an offer from the Institute of Cognitive Science in Osnabrück, Germany, where he was appointed Full Professor for Machine Learning.