18th Annual GCC Theoretical and Computational Neuroscience Annual Conference

Jan 29, 2021

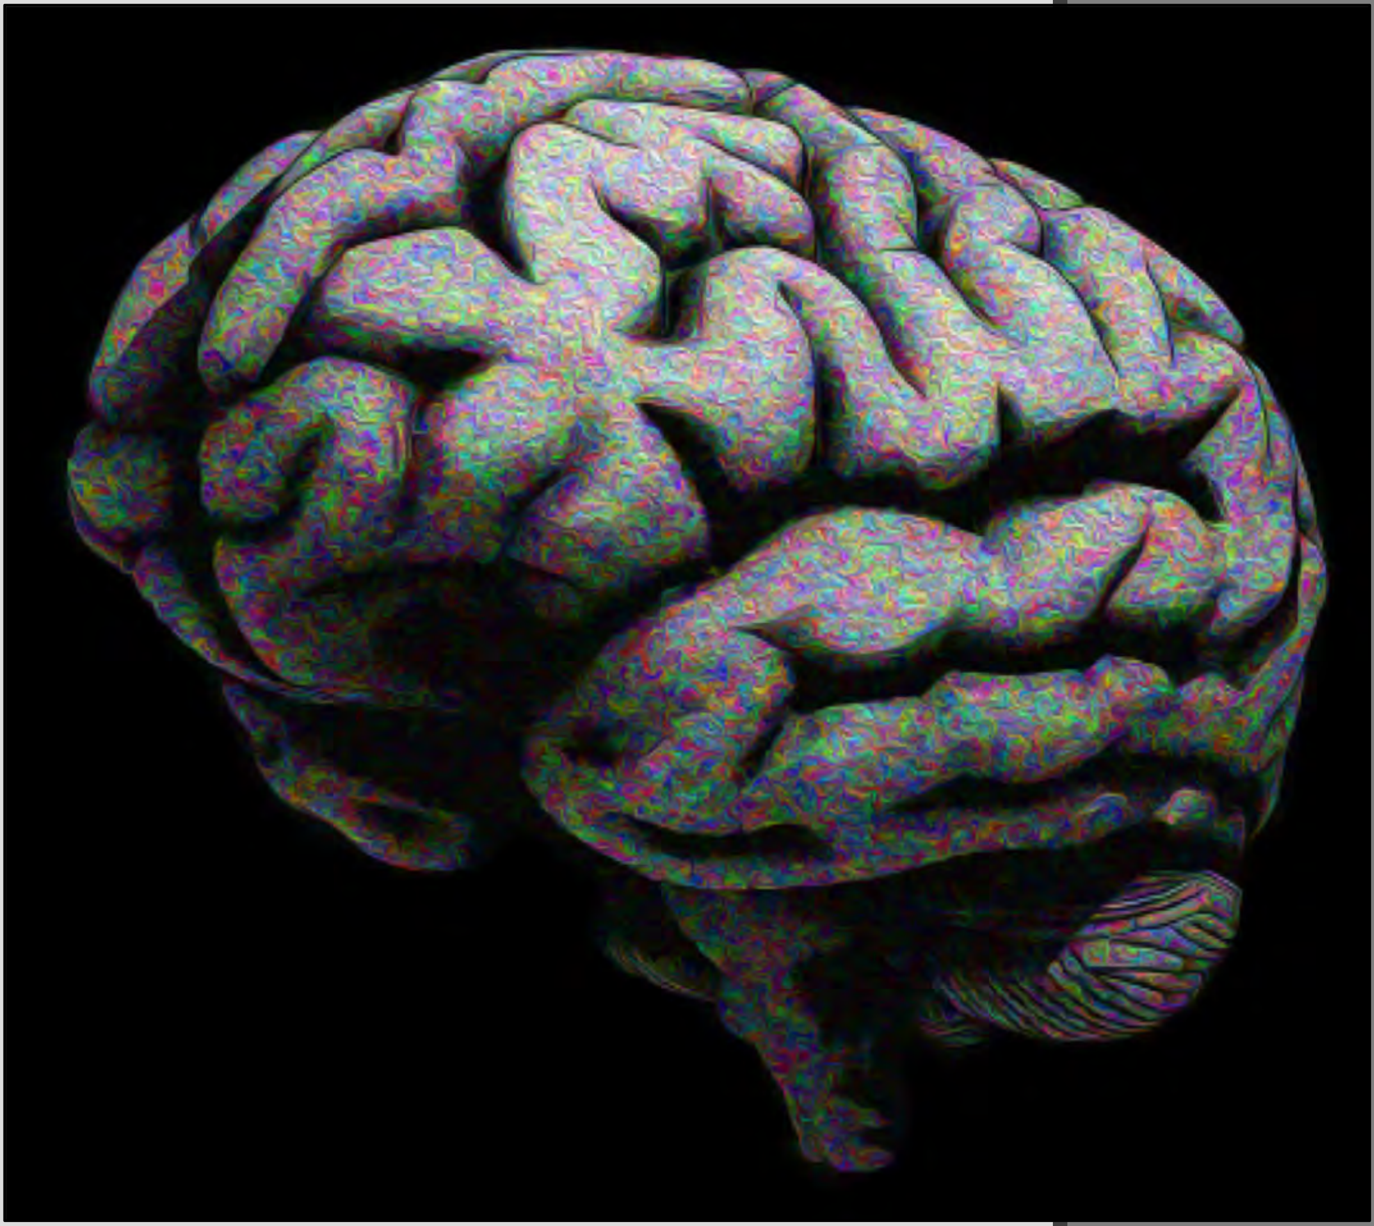

Gulf Coast Consortia

QUANTITATIVE BIOMEDICAL SCIENCES

Gulf Coast Consortia
QUANTITATIVE BIOMEDICAL SCIENCES

The Gulf Coast Consortia (GCC), located in Houston, Texas, is a dynamic, multi- institution collaboration of basic and translational scientists, researchers, clinicians and students in the quantitative biomedical sciences, who benefit from joint training programs, topic-focused research consortia, shared facilities and equipment, and exchange of scientific knowledge. Working together, GCC member institutions provide a cutting-edge collaborative training environment and research infrastructure beyond the capability of any single institution. GCC training programs currently focus on Biomedical Informatics, Computational Cancer Biology, Molecular Biophysics, Pharmacological Sciences, Precision Environmental Health Sciences and Antimicrobial Resistance. GCC research consortia gather interested faculty around research foci within the quantitative biomedical sciences, and currently include AI in Healthcare, Antimicrobial Resistance, Cellular and Molecular Biophysics, Innovative Drug Discovery and Development, Immunology, Mental Health, Regenerative Medicine, Single Cell Omics, Theoretical and Computational Neuroscience, Translational Imaging and Translational Pain Research. Current members include Baylor College of Medicine, Rice University, University of Houston, The University of Texas Health Science Center at Houston, The University of Texas Medical Branch at Galveston, The University of Texas M. D. Anderson Cancer Center, and the Institute of Biosciences and Technology of Texas A&M Health Science Center.

GCCTCN Executive Steering Committee

Krešimir Josić
University of Houston
Chair

Harel Shouval
UT Health Houston
Vice Chair

Fabrizio Gabbiani, Baylor College of Medicine
Fernanda Laezza, UT Medical Branch at Galveston
Xaq Pitkow, Baylor College of Medicine/Rice Univ.
Prahlad Ram, UT MD Anderson Cancer Center

Agenda

January 29, 2021

8:50	Welcome
Moderator:	**Harel Shouval** University of Texas Health Science Center at Houston
9:00	*Theory of Neural Perturbome* **Claudia Clopath**, Imperial College London
10:00	*Pushing the Limits of Intracortical Neural Recording* **Chong Xie**, Rice University
10:30	Break
Moderator:	**Krešo Josić** University of Houston
10:45	*Can Deep Neural Networks Model the Human Faculty of Abstraction?* **Cameron Buckner**, University of Houston
11:15	*Neurally Plausible Mechanisms for Learning Selective and Invariant Representations* **Fabio Anselmi,** Baylor College of Medicine
Moderator:	**Fabrizio Gabbiani** Baylor College of Medicine
11:45	Trainee Short Talks
	Balanced Networks Under Spike-Time Dependent Plasticity **Alan Akil**, Graduate Student, University of Houston
	Normative Adaptive Decision Rules in Static Environments **Nicholas Barendregt**, Graduate Student, University of Colorado Boulder
	Inference as Control **Lokesh Boominathan**, Graduate Student, Rice University
	Spike-Constrained Neural Control **Itzel Olivos-Castillo,** Graduate Student, Rice University
	Behavioral Time Scale Plasticity of Place Fields: Mathematical Analysis **Ian Cone**, Graduate Student, Rice University
	Normative Decision Asymmetries with Symmetric Priors but Asymmetric Evidence **Tahra Eissa**, University of Colorado Boulder, Research Associate
	Learning Accurate Path Integration in a Ring Attractor Model for Heading in Drosophila **Pantelis Vafidis**, Student, California Institute of Technology
12:30	Lunch Break

Agenda

Moderator:	**Fabrizio Gabbiani** Baylor College of Medicine
1:45	*Quantifying Uncertainty in Spikes Estimated from Calcium Imaging Data* **Daniela Witten,** University of Washington, Seattle
Moderator:	**Xaq Pitkow** Baylor College of Medicine/Rice University
2:45	*Large Scale Brain Mapping* **Bobby Kasthuri**, University of Chicago
3:45	Networking Break
Moderator:	**Xaq Pitkow** Baylor College of Medicine/Rice University
4:10	*On the Rational Boundedness of Cognitive Control* **Jonathan Cohen,** Princeton University
5:00	Closing Remarks

Speakers, in alphabetical order

Fabio Anselmi, PhD
Assistant Professor
Center for Neuroscience and Artificial Intelligence

Neurally Plausible Mechanisms for Learning Selective and Invariant Representations

Fabio is an assistant professor at the Center for Neuroscience and Artificial Intelligence at Baylor College of Medicine where he has been faculty member since 2020. He is also affiliated to the Centre for Brains Minds and Machines at MIT.
Fabio completed his PhD in quantum mechanics at Hertforshire University (Uk) and he graduated in Physics at Padova University in Italy.
His work lies at the interface between computational neuroscience and machine learning with focus on development of biologically grounded machine learning algorithms with application to visual cortex.

Coding for visual stimuli in the ventral stream is known to be invariant to object identity preserving nuisance transformations. Indeed, much recent theoretical and experimental work suggests that the main challenge for the visual cortex is to build up such nuisance invariant representations. Recently, artificial convolutional networks have succeeded in both learning such invariant properties and, surprisingly, predicting cortical responses in macaque and mouse visual cortex with unprecedented accuracy. However, some of the key ingredients that enable such success—supervised learning and the backpropagation algorithm—are neurally implausible. This makes it difficult to relate advances in understanding convolutional networks to the brain. In contrast, many of the existing neurally plausible theories of invariant representations in the brain involve unsupervised learning, and have been strongly tied to specific plasticity rules. To close this gap, we study an instantiation of simple-complex cell model and show, for a broad class of unsupervised learning rules (including Hebbian learning), that we can learn object representations that are invariant to nuisance transformations belonging to a finite orthogonal group. These findings may have implications for developing neurally plausible theories and models of how the visual cortex or artificial neural networks build selectivity for discriminating objects and invariance to real-world nuisance transformations.

Cameron Buckner, PhD
Associate Professor
Philosophy

Can Deep Neural Networks Model the Human Faculty of Abstraction?

Cameron Buckner is an Associate Professor in the Department of Philosophy at the University of Houston. He began his academic career in logic-based artificial intelligence. This research inspired an interest into the relationship between classical models of reasoning and the (usually very different) ways that humans and animals actually solve problems, which led him to the discipline of philosophy. He received a PhD in Philosophy at Indiana University in 2011 and an Alexander von Humboldt Postdoctoral Fellowship at Ruhr-University Bochum from 2011 to 2013. His research interests lie at the intersection of philosophy of mind, philosophy of science, animal cognition, and artificial intelligence, and he teaches classes on all these topics. Recent representative publications include “Empiricism without Magic: Transformational Abstraction in Deep Convolutional Neural Networks” (2018, Synthese), and “Rational Inference: The Lowest Bounds” (2017, Philosophy and Phenomenological Research)—the latter of which won the American Philosophical Association's Article Prize for the period of 2016–2018. He is currently writing a book about the philosophy of deep learning (on support from the National Science Foundation).

Abstract: Recently, deep neural networks have accomplished feats that skeptics thought would remain beyond the reach of artificial intelligence for many more years. In evaluating these achievements, one question which has been a topic of stark disagreement—but curiously not yet much explicit debate—is whether these networks can model the human faculty of abstraction. Skeptics have argued that as impressive as these achievements are, these networks’ solutions are fundamentally unlike the strategies that humans use to solve these problems, because deep neural networks are unable to discover and manipulate human abstractions. At the same time, many machine learning researchers take it as obvious that deep neural networks’ distinctive computational strength lies in their ability to perform a hierarchical form of abstraction, which influential neuroscientists have argued is similar to processes of abstraction performed in primate perceptual cortex. In this talk, I argue that empiricist philosophy of mind can help illuminate this debate. Abstraction has long played a key role in empiricist philosophy of mind, and over the millenia philosophers have distinguished

several qualitatively different forms of abstraction. I explain four forms in this talk—abstraction-as-composition, abstraction-as-subtraction, abstraction-as-representation, and abstraction-as-invariance—and discuss ways that deep neural networks may or may not be said to implement them. I conclude by suggesting that deep neural networks may in some ways transcend the limits of human abstraction, which intersects with pressing questions over the desirability of transparency and trustworthiness in scientific applications of deep learning.

Claudia Clopath, PhD
Professor
Bioengineering
Theory of Neural Perturbome

Professor Claudia Clopath is based in the Bioengineering Department at Imperial College London. She is heading the Computational Neuroscience Laboratory.

Her research interests are in the field of neuroscience, especially insofar as it addresses the questions of learning and memory. She uses mathematical and computational tools to model synaptic plasticity, and to study its functional implications in artificial neural networks.

Professor Clopath holds an MSc in Physics from the EPFL and did her PhD in Computer Science under Wulfram Gerstner. Before joining Imperial College, she did postdoctoral fellowships in neuroscience with Nicolas Brunel at Paris Descartes and in the Center for Theoretical Neuroscience at Columbia University. She published highly cited articles in top journals such as Science and Nature, has given dozens of invited talks and keynotes around the world, and received various prizes such as the Google Faculty Award in 2015.

Abstract: To unravel the functional properties of the brain, we need to untangle how neurons interact with each other and coordinate in large-scale recurrent networks. One way to address this question is to measure the functional influence of individual neurons on each other by perturbing them in vivo. Application of such single-neuron perturbations in mouse visual cortex has recently revealed feature-specific suppression between excitatory neurons, despite the presence of highly specific excitatory connectivity, which was deemed to underlie feature-specific amplification. Here, we studied which connectivity profiles are consistent with these seemingly contradictory observations, by modeling the effect of single-neuron perturbations in large-scale neuronal networks. Our numerical simulations and mathematical analysis revealed that, contrary to the prima facie assumption, neither inhibition dominance nor broad inhibition alone were sufficient to explain the experimental findings; instead, strong and functionally specific excitatory–inhibitory connectivity was necessary,

consistent with recent findings in the primary visual cortex of rodents. Such networks had a higher capacity to encode and decode natural images, and this was accompanied by the emergence of response gain nonlinearities at the population level. Our study provides a general computational framework to investigate how single-neuron perturbations are linked to cortical connectivity and sensory coding and paves the road to map the perturbome of neuronal networks in future studies.

Jonathan Cohen, PhD
Professor
On the Rational Boundedness of Cognitive Control

Professor Jonathan Cohen's research focuses on the neural mechanisms underlying cognitive control, and their relationship to the human capacity for general intelligence. Cognitive control is the ability to guide attention, thought and action in accord with goals or intentions. One of the fundamental mysteries of neuroscience is how this capacity for coordinated, purposeful, and flexible behavior arises from the distributed activity of many billions of neurons in the brain. Several decades of cognitive and neuroscientific research have focused on the mechanisms by which control influences processing (e.g., attentional effects in sensory processing, goal directed sequencing of motor output, etc.), and the brain structures upon which these functions depend. However, we still have a poor understanding of how these systems give rise to cognitive control and intelligence. Our work seeks to develop formally rigorous, mechanistically explicit hypotheses about the functioning of these systems, and to test these hypotheses in empirical studies.

Understanding how the human brain gives rise to the remarkable flexibility of the human mind is one of the greatest challenges in science, and work in our laboratory is leveraging the convergence of research in neuroscience, psychology, and computer science that is addressing this challenge. Progress in this area promises both to deepen our understanding of how the human brain gives rise to the mind, and to serve as the foundation for long sought rational approaches to the treatment of neuropsychiatric disorders, as well as the design of machines that can interact more naturally and productively with humans.

Professor Cohen holds a B.A. in Biology and Philosophy from Yale University, an M.D. from University of Pennsylvania, and a Ph.D. in Cognitive Psychology from Carnegie Mellon University. He joined the Princeton faculty in 1998. He has been conferred the highest awards for research in psychology, including the American Psychological Association's Distinguished Scientific Contribution Award and the William James Fellow Award from the Association for Psychological Science.

Abstract: The capacity for cognitive control, while one of the defining characteristics of

human cognition, is also remarkably limited. Typically, people cannot engage in more than a few — and sometimes only a single — control-demanding task at once. Limited capacity was a defining element in the earliest conceptualizations of cognitive control, it remains one of the most widely accepted axioms of cognitive psychology, and is even the basis for some laws (e.g., against the use of mobile devices while driving). Remarkably, however, the source of this limitation remains a mystery. Structural and/or metabolic constraints are commonly, if tacitly, assumed sources. However, these seem unlikely. Cognitive control is known to rely on the function of a part of the brain — the prefrontal cortex — that comprises approximately one third of the neocortex and some 30 billion neurons. It seems unlikely that this poses a structural limitation. Metabolic constraints are equally unlikely. Other functions, such as vision, routinely engage widespread regions of neocortex in an intense and sustained manner. In this talk, I will present an alternative account, that strives to provide a normative explanation for the capacity constraints on cognitive control. This suggests that constraints reflect a fundamental tradeoff in network architectures between the efficacy of learning (generalization) and the efficiency of processing (multitasking). I will describe simulation studies and empirical findings in support of this idea, and discuss its broader implications for both cognitive science and machine intelligence.

Narayanan 'Bobby' Kasthuri, MD
Assistant Professor
Neurobiology
Large Scale Brain Mapping

Dr. Kasthuri is the first Neuroscience Researcher at Argonne National Labs and an Assistant Professor in the Dept. of Neurobiology, University of Chicago. He has an MD from Washington University School of Medicine and a D.Phil. from Oxford University where he studied as a Rhodes scholar. As a post-doctoral fellow, Dr. Kasthuri developed an automated approach to large volume serial electron microscopy (‘connectomics’). Currently, the Kasthuri lab continues to innovate new approaches to brain mapping including the use of high-energy x-rays from synchrotron sources for mapping brains in their entirety. The Kasthuri lab is applying these techniques to in service of answering the question: how do brains grow up, age, and degenerate?

Abstract: The Kasthuri lab at the University of Chicago and Argonne National Laboratory is pioneering new techniques for brain mapping of the fine structure of the nervous system at industrial scale. I will describe these developments including: large volume automated electron microscopy for mapping neuronal connections at the nanoscale, synchrotron source X-ray microscopy for mapping the cellular composition of entire brains, and combining both with genetic cell type specific labeling for multi-scale, multi-modal brain maps. We have applied these tools to brains from octopuses and squids to primates and mice in the service of answering the questions: how do brains grow up and age and how do brains differ across individuals, phylogeny, and disease.

Daniela Witten, PhD

Professor

Statistics

Quantifying Uncertainty in Spikes Estimated from Calcium Imaging Data

Daniela Witten is a professor of Statistics and Biostatistics at University of Washington, and the Dorothy Gilford Endowed Chair in Mathematical Statistics. She develops statistical machine learning methods for high-dimensional data, with a focus on unsupervised learning.

Daniela is the recipient of an NIH Director's Early Independence Award, a Sloan Research Fellowship, an NSF CAREER Award, a Simons Investigator Award in Mathematical Modeling of Living Systems, a David Byar Award, a Gertrude Cox Scholarship, and an NDSEG Research Fellowship. She is also the recipient of the Spiegelman Award from the American Public Health Association for a statistician under age 40 who has made outstanding contributions to statistics for public health, as well as the Leo Breiman Award for contributions to the field of statistical machine learning. She is a Fellow of the American Statistical Association, and an Elected Member of the International Statistical Institute.

Daniela’s work has been featured in the popular media: among other forums, in Forbes Magazine (three times) and Elle Magazine, on KUOW radio (Seattle's local NPR affiliate station), in a NOVA documentary, and as a PopTech Science Fellow.

Daniela is a co-author (with Gareth James, Trevor Hastie, and Rob Tibshirani) of the very popular textbook "Introduction to Statistical Learning". She was a member of the National Academy of Medicine (formerly the Institute of Medicine) committee that released the report "Evolution of Translational Omics".

Daniela completed a BS in Math and Biology with Honors and Distinction at Stanford University in 2005, and a PhD in Statistics at Stanford University in 2010.

Daniela’s abbreviated CV is available here. Her full CV is available upon request.

Abstract: abstract: In recent years, a number of algorithms have been developed to estimate spike times on the basis of calcium imaging data.

In this talk, I will ask a question that arises naturally in applying these algorithms: how can we quantify uncertainty associated with the estimated spikes? I will present a selective inference framework that allows us to compute a p-value associated with each estimated spike. The key idea is that instead of asking the question "what is the probability of seeing such a large increase in fluorescence if there isn't really a spike at this position?", we must instead ask the question "what is the probability of seeing such a large increase in fluorescence if there isn't really a spike at this position, given that we estimated a spike?"

This is joint work with my current PhD student Yiqun Chen and PhD alum Sean Jewell.

Chong Xie, PhD
Associate Professor
Electrical and Computer Engineering and Neuroengineering
Pushing the Limits of Intracortical Neural Recording

Dr. Chong Xie received his BS degree in Physics from the University of Science and Technology of China in 2004, and Ph.D. degree in Materials Science and Engineering from Stanford University in 2011. He did his postdoctoral work at Harvard University in 2011-2014. Before joining Rice ECE, he was an assistant professor of the Department of Bioengineering at University of Texas at Austin in 2014-2019. Dr. Xie's laboratory is primarily interested in applying specially designed functional devices to solve key challenges in fundamental and clinical neuroscience.

Abstract: The brain is a massively-interconnected and constantly-evolving network of specialized circuits, a systematic understanding of which requires an interface that functions at diverse spatial and temporal scales. Implanted electrodes provide a unique approach to decipher brain circuitry by allowing for time-resolved electrical detection of individual neuron activity. However, scalable and stable neural recording that can track and map a large ensemble of neurons across days, weeks and months remains challenging. We recently demonstrated that ultraflexible, cellular-dimensioned neural electrodes afford seamless integration with brain tissue and stable recording of individual neurons for over a year. Building upon this platform, I will also present our recent progress in further decreasing their form factors, and their massive scaling-up of channel count and density in behaving animals. I will finally discuss about our on-going efforts in applying these ultraflexible electrodes in fundamental and translational neurosciences.

Faculty Website

Trainee Short Talk Abstracts

Balanced Networks Under Spike-Time Dependent Plasticity

Akil A[1], Rosenbaum R[2], Josić, K[1]

1. Department of Mathematics, University of Houston
2. Department of Computational and Applied Mathematics and Statistics, University of Notre Dame

Corresponding author: Krešimir Josić, Department of Mathematics, University of Houston, 3551 Cullen Blvd, Houston, TX, josic@math.uh.edu;

The dynamics of local cortical networks are irregular, but correlated. Dynamic excitatory–inhibitory balance is a plausible mechanism that generates such irregular activity, but it remains unclear how balance is achieved and maintained in plastic neural networks. In particular, it is not fully understood how plasticity induced changes in the network affect balance, and in turn, how correlated, balanced activity impacts learning. How does the dynamics of balanced networks change under different plasticity rules? How does correlated spiking activity in recurrent networks change the evolution of weights, their eventual magnitude, and structure across the network? To address these questions, we develop a general theory of plasticity in balanced networks. We show that balance can be attained and maintained under plasticity induced weight changes. We find that correlations in the input mildly, but significantly affect the evolution of synaptic weights. Under certain plasticity rules, we find an emergence of correlations between firing rates and synaptic weights. Under these rules, synaptic weights converge to a stable manifold in weight space with their final configuration dependent on the initial state of the network. Lastly, we show that our framework can also describe the dynamics of plastic balanced networks when subsets of neurons receive targeted optogenetic input.

Normative Adaptive Decision Rules in Static Environments

Barendregt NW[1], Gold JI[2], Josić K[3], Kilpatrick ZP[1]

1. Department of Applied Mathematics, University of Colorado Boulder
2. Department of Neuroscience, University of Pennsylvania
3. Department of Mathematics, University of Houston

Corresponding author: Nicholas W. Barendregt, Department of Applied Mathematics, University of Colorado Boulder, Engineering Center, ECOT 225, 526 UCB, Boulder, Colorado, nicholas.barendregt@colorado.edu

Normative decision-making often requires adaptive forms of evidence accumulation, but less is known about the decision rules needed to achieve optimal performance under these conditions. For example, a foraging animal must account for an uncertain and depleting food yield when deciding whether to leave their current territory. Even in static environments, task constraints and structure (e.g., time) can affect how evidence is interpreted during a task, which can lead to adaptive decision rules. Recent interest in studying adaptive decision rules has resulted in several *phenomenological* models, such as the "urgency-gating model" (UGM), that use collapsing decision thresholds to explain subject behavior in psychophysics experiments. However, we currently lack a general, normative account of adaptive decision rules and their relevance to human decision-making.

Here we show that normative decision-making even in relatively simple, static two-choice tasks involves a broad range of decision-bound dynamics that include non-monotonic forms of urgency. We use dynamic programming to find dynamic bounds used by a normative observer who optimizes their expected reward rate based on knowledge of task parameters. We then validate the model, including its predictions of non-monotonic urgency, using human subject data from a "tokens task", in which the latent parameters of the task are static but evidence quality fluctuates with evidence history. Adaptive, non-monotonic decision bounds underlie the optimal decision policy across broad regions of task parameter space, expanding the definition of decision urgency and suggesting future task paradigms for discriminating subject decision strategies. We thus show that it is just as important to consider the role of potentially complex decision criteria when evaluating subject decision strategies as it is to model evidence accumulation.

Inference as Control

Boominathan L[1], Schrater P[2], Xaq Pitkow X[3,4]

1. Department of Electrical and Computer Engineering, Rice University
2. Department of Psychology and Computer Science, University of Minnesota
3. Department of Neuroscience, Baylor College of Medicine
4. Department of Electrical and Computer Engineering, Rice University

Corresponding author: Lokesh Boominathan, Department of Electrical and Computer Engineering, Rice University, 6100 Main Street, Houston, Texas, E-mail: lb36@rice.edu

A critical computation for the brain is to infer the world's latent variables from ambiguous observations. Computational constraints, including metabolic costs and noisy signals, limit the performance of these inferences. Efficient coding is a prominent theory that describes how limited resources can be used best. In one incarnation, this leads to the theory of predictive coding, which posits that predictions are sent along feedback channels to be subtracted from signals at lower cortical areas; only the difference returns to the higher areas along feedforward channels, reducing the cost of sending redundant signals already known to the higher areas. This theory does not, however, account for the costs or noise associated with the feedback. Depending on the costs for sending predictions and the reliability of signals encoding those predictions, we expect different optimal strategies to perform computationally constrained inferences. For example, if the feedback channel is too unreliable and expensive, we hypothesize that it is not worth sending any predictions at all. Here we offer a more general theory of inference that accounts for the costs and reliabilities of the feedback and feedforward channels, and the relative importance of good inferences about the latent world state. We formulate the inference problem as control via message-passing on a graph, maximizing how well an inference tracks a target state while minimizing the message costs. Messages become actions with their own costs to reduce while improving how well an inference tracks a target state. We solve this problem under Linear-Quadratic-Gaussian (LQG) assumptions: Linear dynamics and transformations, Quadratic state and control message costs, and Gaussian noise for the process, observations, and measurements. Our theory enables us to determine the optimal predictions and how are they are integrated into computationally constrained inference.

Behavioral Time Scale Plasticity of Place Fields: Mathematical Analysis

Cone I,[1,2] Shouval H[1]

1. Department of Neurobiology and Anatomy, University of Texas Medical School at Houston, Houston, TX, USA
2. Applied Physics Program, Rice University, Houston, TX, USA

Traditional synaptic plasticity experiments and models depend on tight temporal correlations between pre- and postsynaptic activity. These tight temporal correlations, on the order of tens of milliseconds, are incompatible with significantly longer behavioral time scales, and as such might not be able to account for plasticity induced by behavior.

Indeed, recent findings in hippocampus suggest that rapid, bidirectional synaptic plasticity which modifies place fields in CA1 operates at behavioral time scales. In these experiments[1-3], place-field plasticity is shown to occur rapidly in response to either naturally occurring or artificially induced dendritic calcium spikes, also known as "plateau potentials". These protocols demonstrate both an increase and a decrease in synaptic efficacies occurring in synapses that were active seconds before or after the plateau potentials. This plasticity, coined "behavioral timescale synaptic plasticity" (BTSP), is therefore unable to be reconciled with forms of synaptic plasticity that depend on tight correlations between pre and postsynaptic activity.

The phenomenon of BTSP may depend on synaptic “eligibility traces”, both for LTP and LTD. Activated by neural activity, eligibility traces act as a marker or “tag” at specific synapses, and can last on the order of seconds. These traces can then be converted to changes in synaptic efficacies by the activation of some sort of reward or instructive signal. A recent paper has shown that for BTSP, these traces likely depend only on presynaptic activity and the magnitude of the existing synaptic efficacy, and that change in synaptic efficacies can depend on the overlap between these traces and an instructive signal that is activated by the plateau potential[3]. The data therefore supports a two-factor model in which the two factors are presynaptic activity and an instructive signal.

The model for BTSP we present and analyze here extends these previous results. We show that the place fields produced by a two-factor eligibility trace model have fixed points, and that these fixed points can be defined and calculated. Our model additionally predicts the convergence rate to these fixed points. In some simple cases these fixed points can be fully solved analytically. Using these solutions, we show how these fixed points depend on the system's parameters such as the shape of the presynaptic place fields and the animal's velocity. We show explicitly that the place fields become broader if the animal has a higher velocity during induction, and predict that LTD far away from the instructive signal has a slow convergence time to the fixed point. These results agree with, and extend upon, existing experiments on BTSP[1-3] and are achieved by a simple and analytically tractable mathematical model.

1. Bittner, K. C. et al. Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. Nature Neuroscience 18, 1133–1142 (2015).
2. Bittner, K. C., Milstein, A. D., Grienberger, C., Romani, S. & Magee, J. C. Behavioral time scale synaptic plasticity underlies CA1 place fields. Science 357, 1033–1036 (2017).
3. Milstein, A. D. et al. Bidirectional synaptic plasticity rapidly modifies hippocampal representations independent of correlated activity. http://biorxiv.org/lookup/doi/10.1101/2020.02.04.934182 (2020) doi:10.1101/2020.02.04.934182.

Normative Decision Asymmetries with Symmetric Priors but Asymmetric Evidence

Eissa TL[1*], Gold JI[2+], Josić K[3,4+], Kilpatrick ZP[1,5+]

1. Department of Applied Mathematics, University of Colorado Boulder, Boulder, CO, USA
2. Department of Neuroscience, University of Pennsylvania, Philadelphia, PA, USA
3. Department of Mathematics, University of Houston, Houston, TX, USA
4. Department of Biology and Biochemistry, University of Houston, Houston, TX, USA
5. Institute of Cognitive Science, University of Colorado Boulder, Boulder, CO, USA

\+ equal contributions
*Corresponding author: Tahra Eissa, Department of Applied Mathematics, University of Colorado Boulder, Engineering Center, ECOT 225, Boulder, CO, tahra.eissa@colorado.edu

Decision asymmetries can reflect many factors but are typically considered normative only when they result from asymmetric priors or values. Previous studies of normative decision asymmetries have focused on restricted sets of conditions that have limited our understanding of potential sources of asymmetry. Most studies assumed that available evidence is symmetric, and evidence is matched in strength for each alternative. Here we examine decisions between two alternatives that are each associated with relatively rare, asymmetric evidence (e.g., deciding which of two slot machines gives higher odds of winning by watching their outputs). We show that when evidence is asymmetric and sparse, decision asymmetries are inevitable. With symmetric priors, these evidence-driven asymmetries can raise a conundrum for normative decision-makers, who must reconcile decision asymmetries with an expectation that no asymmetries should exist. We examined how 200 human participants handled this conundrum and compared their performance to normative and non-normative models.

We used a ball-drawing task in which balls of one color could be rare. Subjects saw a short string of balls drawn with replacement from one of two equally likely jars with known ratios of ball colors and were asked which jar was used. Most participants reported choices supported by more extreme (i.e., rarer) evidence more often, despite both alternatives being equally probable. This discrepancy was consistent with a Bayesian ideal observer that displays similar asymmetries in its decisions.

Notably, many subjects' asymmetries were enhanced relative to the ideal observer, matching Bayesian models that under-weighed the value of rare balls. Model-free results confirmed that deviations from optimality correlated with a decreased use of task-relevant information. These results provide quantitative and theoretically grounded insights into how humans use rare events to make inferences, relevant to predictions in both real-world situations (stock-market crashes) and common laboratory tasks (changes in reward contingencies).

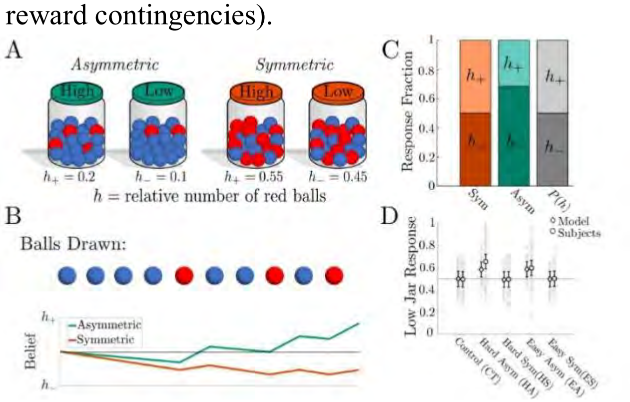

A) Cartoon of task. B) Ideal observer's belief (log-likelihood ratio, LLR) about the jar in use, given a sequence of observations. C) Response fractions across many trials (<10 observations per trial) compared with symmetric prior. D) Subject and model low jar response fractions (mean and bootstrapped confidence interval shown).

This work was funded by CRCNS R01MH115557-01.

Spike-Constrained Neural Control

Olivos-Castillo IC[1], Schrater P[5, 6], Pitkow X[2, 3, 4]

1. Department of Computer Science, Rice University
2. Department of Electrical and Computer Engineering, Rice University
3. Department of Neuroscience, Baylor College of Medicine
4. Center for Neuroscience and Artificial Intelligence, Baylor College of Medicine
5. Department of Psychology, University of Minnesota
6. Department of Computer Science, University of Minnesota

Corresponding author: Itzel Coral Olivos-Castillo, Department of Computer Science, Rice University, 6100 Main St., Houston, Texas, E-mail: itzel.olivos@rice.edu

Abstract:

We develop a version of stochastic control that accounts for computational costs in the brain. Motor control and reinforcement learning both appeal to the conceptual framework of accumulating evidence about the world and selecting actions based on the synthesized information to maximize the total expected reward. However, neither of these approaches consider the costs that the brain pays for performing computations and representing information. Conversely, past studies identified metabolically efficient ways of coding sensory information, but these studies do not consider the consequences for closed-loop control and are restricted to feedforward settings and static environments. To help bridge this gap, here we combine concepts of efficient coding with control theory to analyze Linear Quadratic Gaussian (LQG) control, a well-understood mathematical example of optimal control.

In LQG control, the environment's dynamics are linear, all sources of noise are Gaussian, and a quadratic function dictates task performance. Given these assumptions, an agent that estimates the environment's hidden state using a Kalman filter and computes control policies using Bellman's optimality principle is optimal. Thus, in this work, we: i) Use the activity of a population of neurons with Poisson-like response variability to encode the environment's noisy observations. ii) Use a dynamic Probabilistic Population Code to implement a Kalman filter in which linear projections of spiking neural approximate the natural parameters of a Gaussian posterior over the environment's state. iii) Introduce a representation cost, defined as the total integrated number of spikes that the neural circuit uses to encode observations and inferences. iv) Let the agent select actions that minimize state, action, and representation costs.

By solving this problem, we describe how the optimal spike rate varies with properties of the agent, such as sensory gain, and properties of the system to be controlled, such as stability, process noise, and observation noise. For example, since the task performance increases with precise inferences and the inferences' precision is directly proportional to the number of spikes, an agent with fixed sensory gain can obtain more utility overall by relinquishing some task performance if doing so saves enough spikes. In the case of an agent with active sensing capacity, the strategy to obtain more utility overall consists of allocating more spikes in states where making mistakes is highly punished. The latter strategy is often observed in behavioral experiments and may provide a prediction for arousal signals reflected in neural activity and pupil dilatation. Overall, this work provides a foundation for a new type of bounded rational behavior that could be used to explain suboptimal computations in the brain.

Funding sources: McNair foundation, Fulbright-García Robles Scholarship

Learning Accurate Path Integration in a Ring Attractor Model for Heading in Drosophila

Vafidis P[1], Owald D[2], D'Albis T[3,4], Kempter R[3,4]

1. Division of Biology & Biological Engineering, California Institute of Technology, Pasadena,USA
2. Institute of Neurophysiology, Charite – Universitatsmedizin Berlin, Berlin, Germany
3. Institute for Theoretical Biology, Humboldt-Universität zu Berlin, Berlin, Germany
4. These authors contributed equally

Corresponding author: Pantelis Vafidis, Division of Biology & Biological Engineering, California Institute of Technology, 1200 E. California Blvd., Pasadena, CA 91125, E-mail: pvafeidi@caltech.edu

Head direction (HD) cells can track an animal's HD in darkness by integrating angular velocity signals, a phenomenon called path integration (PI) [1]. The one-dimensional ring attractor has been successful in explaining how HD cells achieve PI [2,3,4], and has received strong experimental support in the fly HD system [5,6,7]. However it remains unclear how a path-integrating circuit of HD cells could self-organize. Previous models required parameter tuning after training [3] or used non-local learning rules [4].

We propose a biologically plausible model for the development of the synaptic connections of a ring attractor network where the visual input acts as a supervising signal during training. The learning rule used is local, and it is motivated from mammalian systems, where pyramidal neurons have been shown to act as coincidence detectors to associate inputs arriving at different compartments [8]. Applied to the architecture of the fly HD system (Fig. 1), the model learns a connectivity that shows striking similarities to the one reported in the fly [6], and achieves gain-1 path integration in darkness, when the visual input is absent, for the full range of angular velocities that the fly displays (Fig. 2). The resulting network is a quasi-continuous attractor, and it reproduces experiments in which optogenetic stimulation artificially controls the internal representation of heading [7], and where the network remaps to integrate with different gains, akin to experiments conducted in virtual reality in rodents [9].

Overall, our model answers how the well-characterized HD system in the fly could self-organize during development, while resolving the age-old question of how to learn continuous attractor networks that achieve gain-1 PI. Although tailored to the fly HD system, our model is general and can be used to learn path integration in architectures that lack the physical topography of a ring [10].

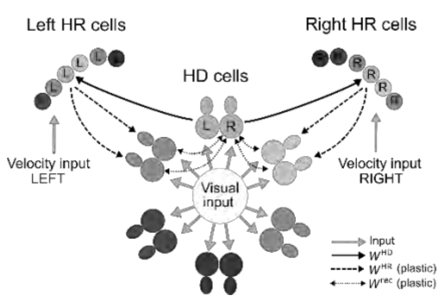

Fig. 1. *Drosophila* central complex. HR = Head Rotation

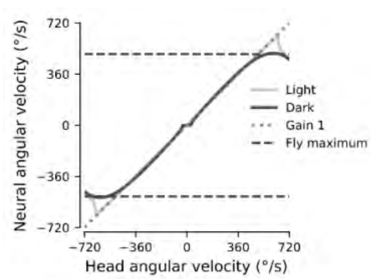

Fig. 2. Network achieves gain-1 PI in darkness

[1] Etienne et al, J. Exp. Biol. 199(1):201209 (1996)
[2] Zhang, J. Neurosci. 16(6):2112–2126 (1996)
[3] Stringer et al, NETCNS 13(2):217 242 (2002)
[4] Hahnloser, Neuroscience 120(3):877891 (2003)
[5] Seelig & Jayaraman, Nature 521(7551):186191 (2015)
[6] Turner-Evans et al, Neuron 108, 1-29 (2020)
[7] Kim et al, Science 356(6340):849–853 (2017)
[8] Larkum, Trends Neurosci. 6(3):141–151 (2013)
[9] Jayakumar et al, Nature 566, 533–537 (2019)
[10] Chaudhuri et al, Nat Neuro 22, 1512-20(2019)

PV funded by Onassis Scholarship and Bing Graduate Fellowship, DO, TD & RK by SFB 1315