NEXTEN: Envisaging Theoretical and Computational Neuroscience for the Next 10 Years

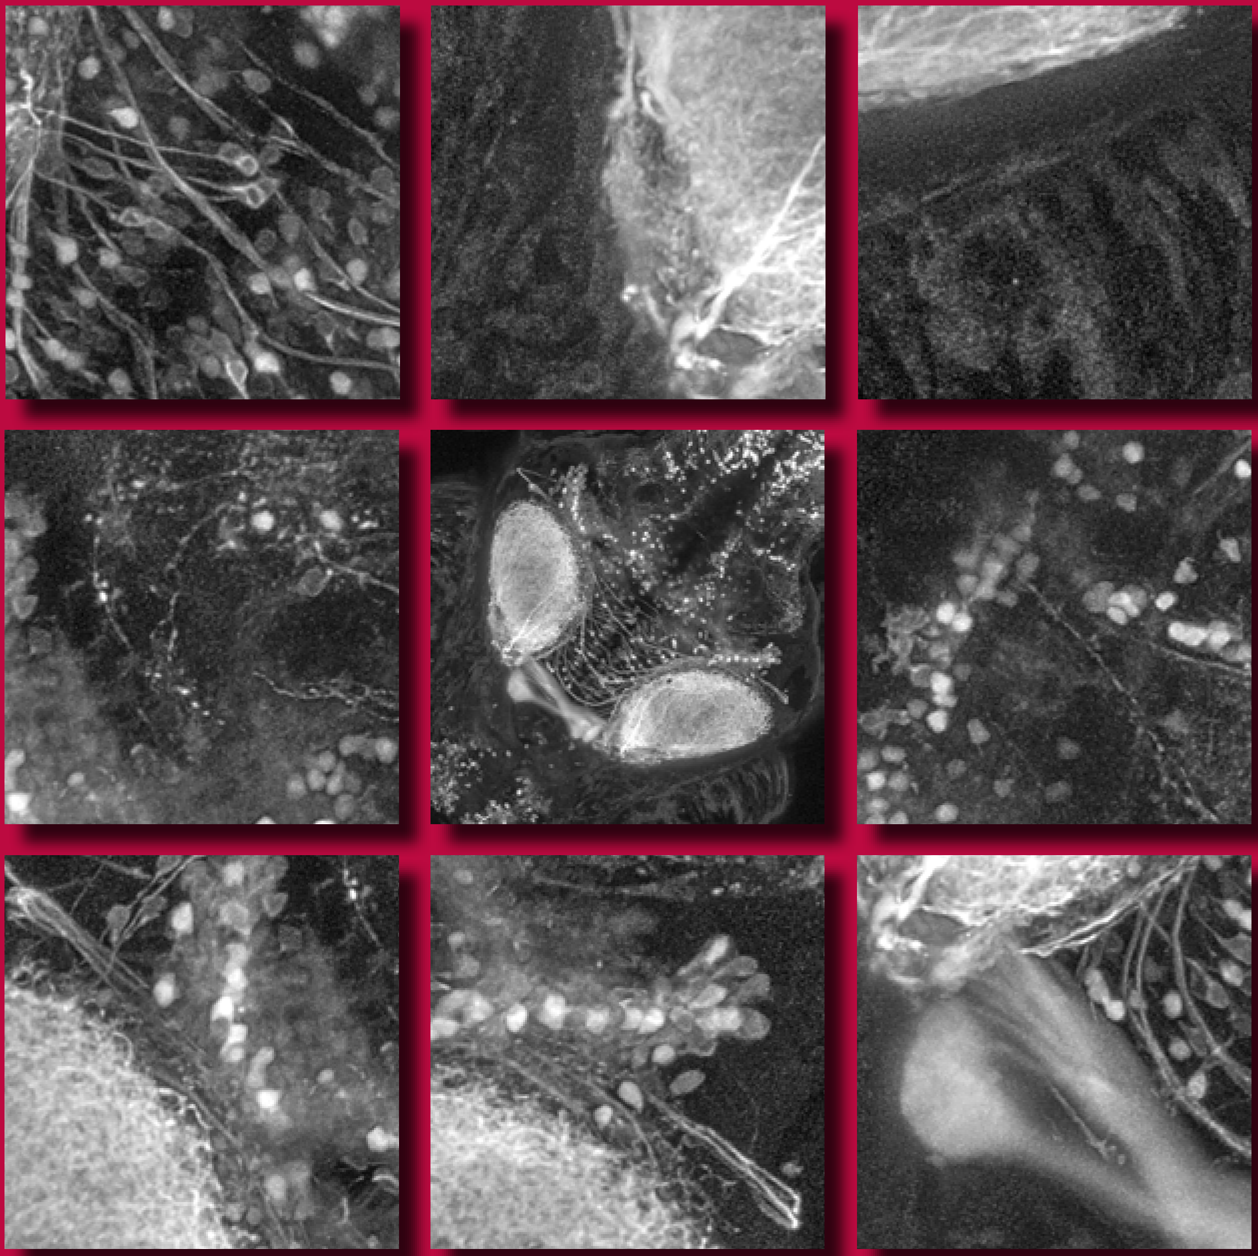

May 16-17, 2024
Washington University in St. Louis

Washington
University in St. Louis

Contents

Welcome

Dear Conference Attendees,

On behalf of the Center for Theoretical and Computational Neuroscience at Washington University in St Louis, welcome to NEXTEN! We're delighted you have chosen to join us for what will be a very exciting two days of science.

The past several years have witnessed expansive growth in both neural/behavioral recording technologies and Artificial Intelligence, opening up new vistas in terms of what is now possible and what might be possible for neuroscience and brain-derived technologies in the future. The goals of NEXTEN are three-fold: (i) to review some of the most exciting recent developments in the field of computational and theoretical neuroscience, (ii) to make projections about what progress the next decade will bring and (iii) to promote lively discussions about the state of the field. We are very much looking forward to your participation in this endeavor. Please note that NEXTEN will finish with a panel discussion that is sure to be a spirited debate!

And finally, please do stay if you can after the conference for the first CTCN Annual Public Lecture on Minds and Machines, given by Terry Sejnowski (Friday 5pm, register here).
https://ctcn.wustl.edu/items/the-ctcn-annual-public-lecture-on-minds-and-machines

Best wishes,

CTCN Executive Committee:
Geoff Goodhill
ShiNung Ching
Gaia Tavoni
Jeff Zacks

Location

Jeffrey T. Fort Neuroscience Research Building (NRB)
Washington University School of Medicine
4370 Duncan Avenue
St. Louis, MO 63110

Talks are held in the NRB Auditorium
Poster session, breaks, and reception are held in the lobby

Postdoctoral Opportunities

The **Center for Theoretical and Computational Neuroscience** is now hiring an exceptional cohort of postdocs to work at the interface between theoretical and experimental labs and help forge new collaborations.

CTCN Fellows receive a competitive salary, generous benefits and $3,000/year for conference travel.

WashU ranks in the top 10 universities in the world for the study of Neuroscience and Behavior and continues to expand in this area, including the new 600,000 sq ft building dedicated to neuroscience research.

St. Louis is a city of close to 3 million people rich in culture, green spaces, free museums, and a thriving music and arts scene.

ctcn.wustl.edu

Center for Theoretical & Computational Neuroscience

Washington University in St. Louis
School of Medicine

Agenda

THURSDAY, MAY 16, 2024

Jeffrey T. Fort Neuroscience Research Building
Washington University School of Medicine

8:30 a.m.	BREAKFAST	
9:00 a.m.	Opening remarks	Geoff Goodhill, PhD CTCN Director Professor of Neuroscience and Developmental Biology Washington University in St. Louis
9:05 a.m.	**Keynote** State-space dynamics in cortex and transformers	Terrence Sejnowski, PhD Professor and Laboratory Head Computational Neurobiology Laboratory Francis Crick Chair Salk Institute for Biological Studies
10:05 a.m.	What an engineer learned studying the insect olfactory system	Barani Raman, PhD Professor of Biomedical Engineering Washington University in St. Louis
10:35 a.m.	BREAK	
11:00 a.m.	Balancing memory and learning: an analytical model of knowledge transfer and forgetting in neural networks	Naoki Hiratani, PhD Assistant Professor of Neuroscience Washington University in St. Louis
11:30 a.m.	Low-dimensional criticality embedded in high-dimensional awake brain dynamics	Woodrow Shew, PhD Associate Professor of Physics University of Arkansas
11:45 a.m.	Desegregation of neuronal predictive processing	Bin Wang, PhD Postoctoral Researcher Columbia University
12:00 p.m.	LUNCH	
1:15 p.m.	**Keynote** Learnable latent dynamics	Mackenzie Mathis, PhD Assistant Professor, EPFL Bartarelli Foundation Chair of Integrative Neuroscience European Laboratory for Learning and Intelligent Systems (ELLIS) Scholar

2:15 p.m.	Unsupervised discovery of a decision making algorithm from neural activity	Adam Kepecs, PhD Robert J. Terry Professor of Neuroscience Professor of Psychiatry BJC Investigator Washington University in St. Louis
2:45 p.m.	Geometry of naturalistic object representations in models of working memory	Xiaoxuan Lei Graduate student McGill University
3:00 p.m.	Addressing temporal credit assignment in recurrent networks using dynamical systems theory	Rainer Engelken, PhD Postdoctoral researcher Columbia University
3:15 p.m.	BREAK	
3:30–5:30 p.m.	**Poster session**	
5:30–7:00 p.m.	Opening reception	

Friday, May 17, 2024

Jeffrey T. Fort Neuroscience Research Building
Washington University School of Medicine in St. Louis

8:30 a.m.	BREAKFAST	
9:00 a.m.	**Keynote** Integrating past and present in continual learning	Richard Zemel, PhD Trianthe Dakolias Professor of Engineering and Applied Science Professor, Department of Computer Science Columbia University
10:00 a.m.	Connecting principles of optimal coding in feedback driven sensory networks	Gaia Tavoni, PhD Assistant Professor of Neuroscience Washington University in St. Louis
10:30 a.m.	BREAK	
11:00 a.m.	Integrating astrocytes as contextual multiplexers of neural dynamics and computation	ShiNung Ching, PhD Associate Professor of Electrical & Systems Engineering Washington University in St. Louis

11:30 a.m.	Optimal learning of long-memory tasks using recurrent network interactions instead of single neurons	Roxana Zeraati Graduate student University of Tübingen
11:45 a.m.	Suppression of high-dimensional chaos by clustered synaptic connectivity	Lukasz Kusmierz, PhD Scientist Allen Institute
12:00 p.m.	Emergent behavior and neural dynamics in artificial agents tracking odor plumes	Satpreet Singh, PhD Postdoctoral researcher Harvard University
12:15 p.m.	LUNCH	
1:30 p.m.	**Keynote** A less Artificial Intelligence	Andreas Tolias, PhD Professor Department of Neuroscience Baylor College of Medicine
2:30 p.m.	**Panel discussion**	
3:30 p.m.	Afternoon tea & conference conclusion	

Abstracts

1 Using interpretable generative models to probe the multiregional interactions underlying movement initiation

Mark Agrios[1], Amy Kristl[1], Natalie Koh[1], Sajishnu Savya[1], Sarah Hsu[1], Sara Solla[2], Andrew Miri[1]

[1] Northwestern University Department of Neurobiology,
[2] Northwestern University Department of Neuroscience

2 Physics-based machine learning in high-content screening

Lina Ali, Graham Bachman, Mallory Wright, Jason Waligorski, Colin Kremitzki, Serena Elia, Diana Grigore, Manny Gerbi, Uma Kaushik, Saul Weiss, Nico Zaharia, Jimin Lee, Purva Patel, Josh Langmade, Waleed Minzal, Aldrin Yim, Josh Milbrandt, Paul Hime, Rob Mitra, Jeff Milbrandt, William Buchser

FIVE@MGI, McDonnell Genome Institute, Department of Genetics, Washington University in St. Louis

3 Dopamine signal auditory Sensory Prediction Error in the auditory striatum

Eleonora Bano, Amelia Christensen, Fengrui Zhang, Heejae Choi, Adam Kepecs

Department of Neuroscience, Washington University in St. Louis

4 Mediodorsal thalamus regulates sensory and mapping uncertainties in flexible decision making

Xiaohan Zhang[1], Michael M. Halassa[2], Zhe Sage Chen[1,3,4]

[1] Department of Psychiatry, New York University Grossman School of Medicine, New York, NY, USA
[2] Department of Neuroscience, Tufts University School of Medicine, Boston, MA, USA
[3] Neuroscience Institute, New York University Grossman School of Medicine, New York, NY, USA
[4] Department of Biomedical Engineering, New York University Tandon School of Engineering, Brooklyn, NY, USA

5 Planar, spiral, and concentric traveling waves distinguish cognitive states in human memory

Anup Das[1], Erfan Zabeh[1], Bard Ermentrout[2], Joshua Jacobs[1,3]

[1] Department of Biomedical Engineering, Columbia University, New York, NY 10027
[2] Department of Mathematics, University of Pittsburgh, Pittsburgh, PA 15260
[3] Department of Neurological Surgery, Columbia University, New York, NY 10027

6 Generalized Contrastive PCA (gcPCA): a generalized framework for finding subspaces that differ between experimental conditions

Eliezyer Fermino de Oliveira[1], Pranjal Garg[3], Lucas Sjulson[1,2]

[1] Dominick P. Purpura Department of Neuroscience, Albert Einstein College of Medicine, Bronx, NY, USA
[2] Department of Psychiatry and Behavioral Sciences, Albert Einstein College of Medicine, Bronx, NY, USA
[3] All India Institute of Medical Sciences, Rishikesh, India

7 Comparing functional MRI data representations based on distributions of repeatable topological features

Ty Easley, Kevin Freese, Elizabeth Much, Janine Bijsterbosch

Washington University in St. Louis

8 Addressing temporal credit assignment in recurrent networks using dynamical systems theory Talk and Poster

Rainer Engelken, Larry Abbott

Columbia University

9 Laying the groundwork for semantic decoding of object categories in naturalistic movies via high-density diffuse optical tomography

Wiete Fehner[1], Zachary Markow[1], Morgan Fogarty[1], Aahana Bajracharya[1], Dana Wilhelm[1], Alexander G. Huth[2], Joseph P Culver[1]

[1] Washington University in St. Louis, St. Louis, Missouri 63130, USA
[2] The University of Texas at Austin, Austin, 78712, TX, USA

10 Optimal dynamics in the manifold for spontaneous behavior

Antonio J. Fontenele[1], J. Samuel Sooter[1], Nivaldo A. P. de Vasconcelos[2] and Woodrow L. Shew[1]

[1] Department of Physics, University of Arkansas, Fayetteville, 72701, AR, USA
[2] Department of Biomedical Engineering, Federal University of Pernambuco (UFPE), Recife, PE 50670-901, Brazil

11 Low-dimensional criticality embedded in high-dimensional awake brain dynamics

Talk and Poster

Antonio J. Fontenele, J. Samuel Sooter, V. Kindler Norman, Shree Hari Gautam, and Woodrow L. Shew

University of Arkansas, Department of Physics, UA Integrative Systems Neuroscience Group, Fayetteville, 72701, AR, USA

12 Quasicriticality in light and dark conditions

Leandro J Fosque, Yifan Xu, Aidan Schneider, and Keith Hengen

Washington University in St. Louis

13 Parvalbumin expression does not account for discrete electrophysiological profiles of glutamatergic ventral pallidal subpopulations

Robert D Graham[1], Lisa Z Fang[1], Jessica R Tooley[1,2], Vani Kalyanaraman[1], Mary Christine Stander[1], Darshan Sapkota[1], Joseph Dougherty[2,3], Bryan Copits[1,2], Meaghan C Creed[1,2]

[1] Department of Anesthesiology, Washington University in St. Louis
[2] Graduate program in neuroscience, Division of Biological and Biomedical Sciences
[3] Department of Psychiatry, Washington University in St. Louis
[4] Department of Biological Sciences, University of Texas at Dallas

14 Cue-paced training protocols: An assessment of their impact on modulating lower limb excitability

Rachel Hawthorn[1,2], Rodolfo Keesey[1,2], Carolyn Atkinson[1,2], Haolin Nie[1], Ismael Seáñez[1,2,3]

[1] Department of Biomedical Engineering, Washington University in St. Louis
[2] Division of Neurotechnology, Washington University School of Medicine in St. Louis
[3] Department of Neurosurgery, Washington University School of Medicine in St. Louis

Suppression of high-dimensional chaos by clustered synaptic connectivity Talk only

Łukasz Kuśmierz[1], Ulises Pereira-Obilinovic[1], Zhixin Lu[1], Dana Mastrovito[1], and Stefan Mihalas[1]

[1] Allen Institute, Seattle, WA

15 Geometry of naturalistic object representations in models of working memory

Talk and Poster

Xiaoxuan Lei[1], Takuya Ito[2], Pouya Bashivan[1]

[1] Department of Physiology, McGill University, Canada
[2] IBM Research, New York, United States

16 Mechanisms of neuronal coordination in multi-layer networks

Tianfeng Lu[1], Shuyu I. Zhu[1], Robert Wong[1,2], Geoffrey J. Goodhill[1]

[1] Departments of Developmental Biology and Neuroscience, Washington University in St. Louis
[2] Departments of Electrical & Systems Engineering, Washington University in St. Louis

17 Neurogenesis enhances olfactory coding efficiency, especially in variable environments

Ryan McGee[1], Gaia Tavoni[1]

[1] Department of Neuroscience, Washington University in St Louis

18 Neuro Translate: Using surface image transformers to translate shared latent information across functional MRI brain representations

Samuel Naranjo Rincon, Fyzeen Ahmad, Ty Easly, Ulugbek Kamilov, Janine Bijsterbosch

Washington University in St. Louis

19 Generative modeling of the functional connectome in first-episode psychosis

Jacob Pine, Deanna Barch, and Matthew Singh

Washington University in St. Louis, Department of Psychological and Brain Sciences

20 Noise gives brains super-Turing computation computers are missing out

Emmett Redd

Physics, Astronomy and Materials Science, Missouri State University, Springfield, Missouri USA

21 Decision confidence and belief updating: quantitative behavioral phenotyping of perceptual decision-making in a rat *CHD8+/-* model of autism

Steven C. Ryu, Jack McGuire, Adam Kepecs

Washington University in St. Louis, School of Medicine, MO, USA

22 Structural localization is embedded in the spike trains of single neurons

Schneider, G. Tolossa, K. Hengen

Department of Biology, Washington University in St. Louis

23 Relative neural population size modulates learnability of cyclic features of neural codes

Niko Schonsheck, Chad Giusti

University of Delaware, Oregon State University

Emergent behavior and neural dynamics in artificial agents tracking odor plumes

Talk only

Satpreet H. Singh[1], Floris van Breugel[2], Rajesh P. N. Rao[1], and Bingni W. Brunton[1]

[1] University of Washington, Seattle, WA 98195,
[2] University of Nevada, Reno, NV 89557

24 Visual cortex is nearest to criticality at intermediate level of arousal: a temporal renormalization group approach

J. Samuel Sooter[1], Andrea K. Barreiro[2],Cheng Ly[3], Antonio J. Fontenele[1],
Nivaldo A.P. de Vasconcelos[4], Woodrow L. Shew[1]

[1] Department of Physics, UA Integrative Systems Neuroscience, University of Arkansas
[2] Department of Mathematics, Southern Methodist University
[3] Department of Statistical Sciences and Operations Research, Virginia Commonwealth University
[4] Department of Biomedical Engineering, Federal University of Pernambuco

25 Circumstantial evidence and explanatory models for synapses in large-scale spike recordings

Ian H. Stevenson

University of Connecticut

26 Desegregation of neuronal predictive processing **Talk and Poster**

Bin Wang[1], Nicholas J Audette[2], David M Schneider[2] and Johnatan Aljadeff[3]

[1] Center for Theoretical Neuroscience, Columbia University
[2] Center for Neural Science, New York University
[3] Department of Neurobiology, University of California San Diego

27 Dimensionality and developmental trajectory of signal and noise in the brain

Robert Wong[1,3,4], Alisha Tromp[2], Zac Pujic[2], Biao Sun[2], Naoki Hiratani[3], Geoffrey Goodhill[3,4]

[1] Department of Electrical & Systems Engineering, Washington University in St. Louis, MO, USA
[2] Queensland Brain Institute, Australia
[3] Department of Neuroscience, Washington University School of Medicine, MO, USA
[4] Department of Developmental Biology, Washington University School of Medicine, MO, USA

28 Border ownership signals emerge in an artificial neural network trained to predict future visual input

Zeyuan Ye, Ralf Wessel, Tom P. Franken

Washington University in St. Louis

29 A recurrent neural network model of excitatory-inhibitory imbalance in autism spectrum disorder

Daniel Zavitz[1], ShiNung Ching[2], Geoffrey Goodhill[1]

[1] Departments of Developmental Biology and Neuroscience
[2] Department of Electrical and Systems Engineering
Washington University in St. Louis

Optimal learning of long-memory tasks using recurrent network interactions instead of single neurons Talk only

Roxana Zeraati[1,2], Sina Khajehabdollahi[1,3], Emmanouil Giannakakis[1,2], Tim Jakob Schäfer[1,2], Georg Martius[1,3], Anna Levina[1,2,4]

[1] University of Tübingen, Germany
[2] Max Planck Institute for Biological Cybernetics, Tübingen, Germany
[3] Max Planck Institute for Intelligent Systems, Tübingen, Germany
[4] Bernstein Network Computational Neuroscience Tübingen, Germany

30 Capacity of networks with arbitrary topologies and neuron activation probabilities

Kaining Zhang, Gaia Tavoni

Department of Neuroscience, Washington University School of Medicine

Suppression of high-dimensional chaos by clustered synaptic connectivity

Łukasz Kuśmierz,[1] Ulises Pereira-Obilinovic,[1] Zhixin Lu,[1]
Dana Mastrovito,[1] and Stefan Mihalas[1]

[1] Allen Institute, Seattle, WA

Convergent evidence suggests that neural dynamics is constrained to low-dimensional manifolds which has been linked to low-rank connectivity. Additionally, it is known that synaptic weights are clustered and that this can lead to doubly stochastic dynamics with unpredictable spike times and fluctuating firing rates. However, strictly lowrank connectivity is unlikely to be implemented by biological neural networks due to the difficulties of developing and maintaining the required correlated connectivity through mostly local synaptic plasticity driven by noisy neural dynamics. Moreover, the functional relevance of clustered connectivity is largely unknown. Thus, we set out to understand how low-rank clustered synaptic structures embedded in high-rank random connectivity affect neural dynamics and information processing.

To this end, we developed a novel model of random hierarchical connectivity and studied the resulting neural dynamics utilizing mean-field theory and computer simulations. In our model the synaptic weight matrix is generated in blocks corresponding to populations, i.e. clusters of neurons with similar connectivity patterns. Such clusters give rise to approximately low-rank connectivity and can be interpreted as emerging from shared spatial locations, morphologically or genetically defined cell types, or neuronal ensembles.

Our analysis revealed that clustered connectivity can effectively suppress high-dimensional chaos and give rise to low-dimensional attractors. More generally, we found asymmetric interactions that can be most easily understood in terms of an effective competition between intra- and inter-cluster dynamics. In the regime where these two forms of dynamics coexist without one dominating the other, the overall network dynamics features relatively low values of the maximal Lyapunov exponent and participation ratio dimension, with a significantly higher value of the Lyapunov dimension.

Our model offers a normative perspective on the experimentally observed multiscale stochastic neural dynamics. Indeed, our results indicate that clustered, approximately low-rank synaptic connectivity supports flexible dynamics on low-dimensional manifolds and close to the edge of chaos and suggest an interesting computational regime that maintains the balance between intra- and inter-cluster dynamics.

Talk only

Emergent behavior and neural dynamics in artificial agents tracking odor plumes

Satpreet H. Singh[1], Floris van Breugel[2], Rajesh P. N. Rao[1], and Bingni W. Brunton[1]

[1] University of Washington, Seattle, WA 98195,
[2] University of Nevada, Reno, NV 89557

Tracking a turbulent plume to locate its source is a complex control problem requiring robust multi-sensory integration in the face of intermittent odors, changing wind direction, and variable plume shape. This task is routinely performed by flying insects, often over long distances, in pursuit of food or mates. Several aspects of this remarkable behavior have been studied in detail in many experimental studies. Here, we take a complementary *in silico* approach, using artificial agents trained with reinforcement learning to develop an integrated understanding of the behaviors and neural computations that support plume tracking. Specifically, we use Deep Reinforcement Learning (DRL) to train Recurrent Neural Network (RNN) based agents to locate the source of simulated turbulent plumes. Interestingly, the agents' emergent behaviors resemble those of flying insects, and the RNNs learn to represent task-relevant variables, such as head direction and time since last odor encounter. Our analyses suggest an intriguing experimentally testable hypothesis for tracking plumes in changing wind direction—that agents follow local plume shape rather than the current wind direction. While reflexive short-memory behaviors are sufficient for tracking plumes in constant wind, longer timescales of memory are essential for tracking plumes that switch direction. At the level of neural dynamics, the RNNs' population activity is lowdimensional and organized into distinct dynamical structures, with some correspondence to the uncovered behavioral modules. Our *in silico* approach provides key intuitions for turbulent plume tracking strategies and motivates future targeted experimental and theoretical developments.

Talk only

Optimal learning of long-memory tasks using recurrent network interactions instead of single neurons

Roxana Zeraati[1,2,*], Sina Khajehabdollahi[1,3,*], Emmanouil Giannakakis[1,2],
Tim Jakob Schäfer[1,2], Georg Martius[1,3], Anna Levina[1,2,4]

[1] University of Tübingen, Germany
[2] Max Planck Institute for Biological Cybernetics, Tübingen, Germany
[3] Max Planck Institute for Intelligent Systems, Tübingen, Germany
[4] Bernstein Network Computational Neuroscience Tübingen, Germany

*These authors contributed equally to this work.

The brain can solve tasks with long-term credit assignments by maintaining the memory of previous inputs over extended periods. Long timescales required for solving such tasks may arise from the biophysical properties of individual neurons (single-neuron timescale, e.g., membrane or synaptic time constants) or recurrent network interactions. While both mechanisms operate in the brain, the contribution of each mechanism for optimally solving memory-dependent tasks remains poorly understood. To uncover the role of each mechanism and their interactions, we train recurrent neural networks (RNNs) to solve N-parity and Ndelayed match-to-sample tasks with increasing memory requirements controlled by N, by simultaneously optimizing recurrent weights and individual neurons' time constants. To improve learning long-memory tasks, RNNs are trained with two curricula, gradually increasing difficulty N: (i) in single-N curriculum, networks learn a new N at each curriculum step; (ii) in multi-N curriculum, they learn a new N while maintaining the solutions for previous Ns, similar to biological learning. We estimate the network-mediated timescales from the autocorrelation decay of population neural activity. In both curricula, RNNs develop longer timescales with increasing N, but via distinct mechanisms. Single-N RNNs operate in a strong inhibitory regime and mainly rely on increasing their single-neuron timescales with N. However, multi-N RNNs operate closer to a balanced regime and use only recurrent connectivity to develop long timescales, while keeping their single-neuron timescales constant. The behavior of multi-N RNNs is in line with our previous findings that neural timescales in the primate visual cortex adapt to the task state via network interactions rather than biophysical time constants. Furthermore, we show that using network-mediated mechanisms to develop long timescales increases training speed and stability to perturbations, and allows generalization to tasks beyond the training set. Our results suggest that adapting timescales to task requirements via recurrent connectivity enables learning more complex objectives (holding multiple concurrent memories) and improves computational robustness, which can be an optimal strategy for implementing brain computations.

Talk only

8

Addressing Temporal Credit Assignment in Recurrent Networks using Dynamical Systems Theory

Rainer Engelken, Larry Abbott

Columbia University

Discovering associations between temporally distant cues is crucial for an animal's survival. Solving this temporal credit assignment problem allows organisms to bridge the gap between when a stimulus cue arrives and when its effect unfolds. Gradient-based training of recurrent neural circuit models for temporal tasks with long time horizons presents challenges, potentially leading to vanishing or exploding gradients due issues with gradients across many time steps. Recent research has connected this issue to the Lyapunov exponents of the forward dynamics, describing how nearby perturbations grow or shrink in forward passes. Here, we propose a novel approach, termed gradient flossing, to address the gradient instability in recurrent spiking and firing rate networks by controlling the Lyapunov exponents of the forward dynamics throughout learning. We achieve this by backpropagating through the Lyapunov exponents, which enables us to ’floss’ the gradients and improve network training. Our method enhances the success rate of RNNs on typical neuroscience tasks involving the bridging many time steps, and we demonstrate that applying gradient flossing during training further improves trainability for challenging temporal credit assignment tasks. We establish a connection between the Lyapunov exponents and the dimensionality of the gradient signal in backpropagation. Additionally, we demonstrate the effectiveness of our approach both on spiking and firing rate networks. Our results suggest that flossing gradients through dynamic control of Lyapunov exponents can significantly improve the stability and effectiveness of RNN training. We speculate on the optimization of neural dynamics in animals over evolutionary timescales to bridge long time horizons.

11

Low-dimensional criticality embedded in high-dimensional awake brain dynamics

Antonio J. Fontenele, J. Samuel Sooter, V. Kindler Norman, Shree Hari Gautam, and Woodrow L. Shew

University of Arkansas, Department of Physics, UA Integrative Systems Neuroscience Group,
Fayetteville, 72701, AR, USA

Whether cortical neurons operate in a strongly or weakly correlated dynamical regime determines fundamental information processing capabilities and has fueled decades of debate. Here we offer a resolution of this debate; we show that two important dynamical regimes, typically considered incompatible, can coexist in the same local cortical circuit by separating them into two different subspaces. In awake mouse motor cortex, we find a low-dimensional subspace with large fluctuations consistent with criticality – a dynamical regime with moderate correlations and multi-scale information capacity and transmission. Orthogonal to this critical subspace, we find a high-dimensional subspace containing a desynchronized dynamical regime, which may optimize input discrimination. The critical subspace is apparent only at long timescales, which reconciles discrepancies among some previous studies and dispels several controversies about criticality. Using a computational model, we show that the emergence of a low-dimensional critical subspace at large timescale agrees with established theory of critical dynamics. Our results suggest that cortex leverages its high dimensionality to multiplex dynamical regimes across different subspaces.

Talk and Poster

15

Geometry of naturalistic object representations in models of working memory

Xiaoxuan Lei[1], Takuya Ito[2], Pouya Bashivan[1]

[1] Department of Physiology, McGill University, Canada
[2] IBM Research, New York, United States

Working memory (WM) is a central cognitive ability crucial for intelligent decision-making. Prior computational models of WM have focused on how neural networks maintain categorical stimuli rather than multidimensional naturalistic inputs and only considered single or a limited number of cognitive tasks. Here, we investigate how Recurrent Neural Networks (RNN) manipulate properties of multidimensional representations of naturalistic visual inputs during different stages of WM according to task demands. Specifically, we examined: 1) How do task-optimized RNNs select task-relevant properties of naturalistic objects during WM? 2) What are the computational strategies that RNNs employ to dynamically maintain object properties according to task demands?

To address these questions, we trained sensory-cognitive models consisting of a convolutional neural network (CNN) and a RNN (i.e., perceptual and WM processors) across 9 interrelated N-back tasks. By examining the RNN's latent space, we found that 1) compared to the perceptual space, the RNN orthogonalizes the axes along which distinct object properties are represented, enabling enhanced separation of object properties. Moreover, we found that task-relevant object properties are represented along more orthogonalized axes relative to task-irrelevant ones. Furthermore, the embedding spaces are more aligned between tasks with similar WM demands; 2) Across time, the model preserves the geometric structure of the latent subspace representing various object properties. Interestingly, the rotational dynamics governing the transformation of WM encodings (i.e., embedding of visual inputs in the RNN latent space) into memories was shared across stimuli, yet the transformations governing the retention of an encoding in the face of incoming stimuli were distinct across time. Our findings reveal the ways in which goal-driven RNNs adaptively modify their latent representations in response to task requirements, optimizing their capacity to maintain memories of naturalistic objects within high-dimensional spaces.

Talk and Poster

26

Desegregation of neuronal predictive processing

Bin Wang[1], Nicholas J Audette[2], David M Schneider[2] and Johnatan Aljadeff[3]

[1] Center for Theoretical Neuroscience, Columbia University
[2] Center for Neural Science, New York University
[3] Department of Neurobiology, University of California San Diego

Neural circuits across brain regions are thought to construct internal 'world-models' to guide behavior. The predictive processing framework posits that neural correlates of expected stimulus values, and violations of those expectations ('prediction-errors'), are signatures of those internal models. Most existing experimental and theoretical studies focus on simple unitary sensory predictions. However, in the real world, the brain forms predictions for multimodal, high-dimensional sensorimotor signals, which may require different circuit mechanisms from these simple scenarios. Thus, it remains a major challenge to decipher the brain's algorithms and circuitry to perform predictive processing in support of naturalistic behaviors. Here we address this challenge by investigating how multimodal and high-dimensional predictive representations emerge in a class of recurrent neural networks during learning. By analytically solving the underlying learning dynamics, we find that in natural conditions, the most accurate and robust predictive processing arises in an optimal regime of loose excitatory/inhibitory balance. In this regime, the network exhibits a functional desegregation of stimulus and prediction-error representations at the cellular-level. We tested and confirmed these model predictions in recent behavioral and electrophysiological experiments that probe predictive coding circuits using a rich set of stimulus associations and expectation violations. Overall, our results reveal that in natural conditions, neural representations of internal models are highly distributed yet structured to support effective readout of behaviorally relevant information. Our work will thus help identify the neural mechanisms of predictive processing in a wide range of brain regions and species, through integration of modern neuroscience experimental data into the theoretical and modelling framework developed.

Talk and Poster

Posters

1

Using Interpretable Generative Models to Probe the Multiregional Interactions Underlying Movement Initiation

Mark Agrios[1], Amy Kristl[1], Natalie Koh[1], Sajishnu Savya[1], Sarah Hsu[1], Sara Solla[2], Andrew Miri[1]

[1] Northwestern University Department of Neurobiology,
[2] Northwestern University Department of Neuroscience

While many previous studies have identified numerous brain circuits involved in initiating specific movements in response to cues, the precise neural mechanisms associated with a spontaneous decision to act remain unknown. Traditionally, experiments studying this question have been hampered by a limited ability to capture the simultaneous activity of a large number of neurons. We addressed this challenge by performing simultaneous recordings using four neuropixels in a head-fixed mouse; this allowed us to record the activity of hundreds of neurons in premotor cortex, primary motor cortex, primary somatosensory cortex, striatum, and thalamus, while mice perform a self-initiated naturalistic climbing task. We have implemented generalized linear models (GLMs) and demonstrated that they can be a powerful tool for building generative models that can capture the observed neural dynamics across large populations, while maintaining fidelity to first and second order statics of the recorded population activity. The spiking activity of each neuron is used as input to a GLM that learns pairwise statistical relationships between individual neurons and describes how the activity of one neuron excites or inhibits the activity of the other across varying time lags. Once trained, these models can be used to simulate each individual neuron’s activity adhering to the learned relationships between neurons. We sought to validate these models by comparing the first and second order spiking statistics of the simulated activity against the experimentally measured activity and found that our models are able to produce spike trains with striking similarity to what we observe in the real data. Interestingly, spontaneous ramps in the simulated neural activity are produced with similar frequency and duration as those observed in the neural recordings during climbing. This observation implies that the complex neural dynamics that drive the behavioral switching between non-movement and movement states are captured by our GLM. Additionally, we leverage the interpretability of GLMs by performing analyses directly on the set of pairwise neuron-to-neuron relationships learned by the model during training to test hypothesized interactions across brain regions underlying movement initiation. Our results shed light on the interactions that underlie volition and motor planning, from the level of single neurons to large multiregional populations. Our results indicate that large scale neural recordings combined with powerful data analysis techniques provide a unique tool for investigating the mechanisms that underlie neural system function in cognition, motor behaviors, and beyond.

2

Physics-based machine learning in high-content screening

Lina Ali, Graham Bachman, Mallory Wright, Jason Waligorski, Colin Kremitzki, Serena Elia, Diana Grigore, Manny Gerbi, Uma Kaushik, Saul Weiss, Nico Zaharia, Jimin Lee, Purva Patel, Josh Langmade, Waleed Minzal, Aldrin Yim, Josh Milbrandt, Paul Hime, Rob Mitra, Jeff Milbrandt, William Buchser

FIVE@MGI, McDonnell Genome Institute, Department of Genetics, Washington University in St. Louis, St. Louis, MO

FIVE@MGI conducts high-throughput phenotyping and cell capture screening experiments that connect high content images with genomics, proteomics, and transcriptomics at single cell resolution. Deep learning models are often used to draw inferences quickly and efficiently on images. Such models excel with extensive and representative annotated data, yet manual annotation methods face challenges such as time constraints, subjectivity, and limited resolution, resulting in under-sampling of real-life variations. These challenges are particularly pronounced for small subcellular structures like mitochondria, exacerbated by microscopes' blurring point spread function, and limited pixel and optical resolution – render subcellular annotation a ground-truth deficient problem.

To address these issues, we propose a simulation-supervised deep learning paradigm that leverages 3D computer graphics software to generate ground truth simulation images closely resembling real microscopy images. Our physics-based approach offers two key advantages: it enables supra-resolution feature representation and communicates the biophysical underpinnings of features of interest (e.g., perturbed mitochondria). This methodology allows for controlled sampling of image features, ensuring representativeness and unbiased annotations.

We demonstrate the effectiveness of our simulation-supervised encoder-decoders in various tasks. We train encoder-decoder artificial neural networks on both the synthetic images and the underlying physics-based measurements and deploy the trained models on real image datasets. This enables the reconstruction of previously inaccessible phenotypes from real images while simultaneously communicating their biophysical underpinnings and instance segmentation in more complex images such as human sciatic nerve tissues, axon untangling in neural networks, or identifying lipid droplets in adipocytes.

Future work aims to extend this methodology to model artificial neural networks on Raft-Seq, enabling the creation and refinement of neural connections using chemical or optic stimuli in iPSC-derived cortical neurons. These advancements hold promise for advancing our understanding of cellular dynamics and neural network formation.

3

Dopamine signal auditory Sensory Prediction Error in the auditory striatum.

Eleonora Bano[1], Amelia Christensen[1], Fengrui Zhang[1], Heejae Choi[1], Adam Kepecs[1]

[1] Department of Neuroscience, Washington University in St. Louis, St. Louis, MO, USA.

To behave optimally in a dynamic environment, perceptual priors should adapt when the statistics of the environment change. To study this process, we trained mice on an auditory detection task, with a latent block structure that governed the probability of cue presence on each trial. Mice were able to dynamically update their decision thresholds to track cue probability. Performance of similar auditory decision-making tasks relies on monosynaptic input from the auditory cortex to the auditory striatum (TS); thus, to study how mice learned and used this cue probability prior, we focused on the TS. We reasoned that dopaminergic input to TS might be ideally positioned to update auditory priors. When we measured dopamine release in the TS during the auditory detection task, we observed large dopamine transients upon cue delivery, precisely time-locked only to perceived cues, but not to report initiation. These cue evoked *dopamine transients reflected the difference between prior expectation and sound volume*, reminiscent of a *prediction error signal* comparing the perceived sound to a prior expectation of its presence. Importantly, if this prediction error serves as a teaching signal for *updating of auditory priors*, this signal should adjust expectations for future events. Indeed, we found that after adjusting for tone volume, *cue evoked dopamine release was more predictive of cue response probability on the next trial than it was of cue detection on the current trial*. To test whether mouse cue probability tracking behavior in our task was consistent with online reinforcement learning of statistical priors, we built and fit a hierarchical belief-state reinforcement learning model. This model, with just two parameters, effectively captures the trial-to-trial variations in detection threshold as a function of past perceptual experience. Our findings add support to an emerging view that the cortical-striatal-basal ganglia circuit plays a key role in performing statistical inference.

4

Mediodorsal thalamus regulates sensory and mapping uncertainties in flexible decision making

Xiaohan Zhang[1], Michael M. Halassa[2] and Zhe Sage Chen[1,3,4] *

[1] Department of Psychiatry, New York University Grossman School of Medicine, New York, NY, USA
[2] Department of Neuroscience, Tufts University School of Medicine, Boston, MA, USA
[3] Neuroscience Institute, New York University Grossman School of Medicine, New York, NY, USA
[4] Department of Biomedical Engineering, New York University Tandon School of Engineering, Brooklyn, NY, USA.

*Correspondence: zhe.chen@nyulangone.org (Z.S.C.)

Successful execution of complex decision-making tasks requires identification and processing of multiple sources of uncertainty, such as sensory and mapping uncertainties. The mediodorsal (MD) thalamus is a critical partner for the prefrontal cortex (PFC) in cognitive flexibility and resolving uncertainty in decision making. Animal experiments have shown that the MD enhances prefrontal signal-to-noise ratio in decision making under uncertainty. However, the computational mechanisms by which the MD enhances prefrontal activity in decision making under uncertainty remain unclear. In addition, how the newly discovered cellular diversity in MD contributes to such computations is unexplored. Here we used a reverse-engineering approach and trained biologicallyconstrained computational models to delineate these mechanisms. We found that the inclusion of an MD-like feedforward module increased robustness to sensory noise, enhanced working memory and enabled rapid context switching in the recurrent excitatory-inhibitory PFC network performing two versions of context-dependent decisionmaking tasks with sensory and mapping uncertainties. The sensory uncertainty was represented by the coherence or congruence of sensory cues, whereas cue-to-rule mapping uncertainty was reversed during context switching. Incorporating genetically identified thalamocortical pathways and interneuron cell types replicated neurophysiological findings of neuronal tuning and uncovered attractor-like population dynamics. PFC excitatory units showed context-invariant rule tunings, whereas MD units showed modulation with respect to context and cue uncertainty. The MD also enhanced working memory maintenance during the task delay. Our model revealed key computational mechanisms of context-invariant MD in regulating cueing uncertainty and context switching. It also made experimentally testable predictions linking cognitive deficits with disrupted thalamocortical connectivity, prefrontal excitation-inhibition (E/I) imbalance and dysfunctional inhibitory cell types. In summary, our biologically realistic PFC-MD models replicate key behavioral and neurophysiological findings in mice and provide insight into parsing cognitive deficits for mental disorders such as schizophrenia and autism spectral disorder.

5

Planar, Spiral, and Concentric Traveling Waves Distinguish Cognitive States in Human Memory

Anup Das[1], Erfan Zabeh[1], Bard Ermentrout[2], Joshua Jacobs[1, 3]

[1] Department of Biomedical Engineering, Columbia University, New York, NY 10027
[2] Department of Mathematics, University of Pittsburgh, Pittsburgh, PA 15260
[3] Department of Neurological Surgery, Columbia University, New York, NY 10027

A fundamental challenge in neuroscience is explaining how widespread brain regions flexibly interact to support behaviors. We hypothesize that a mechanism of neural coordination is brain oscillations that propagate as traveling waves across the cortex in distinctive patterns that control when and where different regions interact. To test this hypothesis, we used direct recordings from humans performing multiple memory experiments and a novel analytical framework that measures the directional propagation of oscillations. We found that traveling waves propagated along the cortex in not only plane waves as seen previously, but also in spirals, sources and sinks, and more complex patterns. These traveling wave patterns correlated with various aspects of behavior, with specific propagation patterns reflecting broad cognitive processes, and the identities of specific remembered items. Together, these findings show that complex spatiotemporal patterns of traveling waves underlie human cognition and are relevant practically for neural decoding and brain–computer interfacing.

6

Generalized Contrastive PCA (gcPCA): a generalized framework for finding subspaces that differ between experimental conditions

Eliezyer Fermino de Oliveira[1], Pranjal Garg[3], Lucas Sjulson[1,2]

[1] Dominick P. Purpura Department of Neuroscience, Albert Einstein College of Medicine, Bronx, NY, USA
[2] Department of Psychiatry and Behavioral Sciences, Albert Einstein College of Medicine, Bronx, NY, USA
[3] All India Institute of Medical Sciences, Rishikesh, India

Dimensionality reduction methods are used widely across neuroscience. Their applications range from data visualization to interpretation of high-dimensional data. However, these methods typically explain the low-dimensional structure within a single dataset. In many cases, the experimenter is interested in finding the differences between two conditions, such as stimulus vs. control or drug vs. saline. Contrastive PCA (cPCA) was developed as a tool for this purpose, but it depends critically on computationally-expensive optimization of an opaque free parameter, of which it is not always possible to determine the "correct" value. Here we intuitively explain the purpose of this parameter and propose Generalized Contrastive PCA (gcPCA), a framework that achieves optimal performance without it. gcPCA comprises three distinct algorithms that are appropriate for different use cases, and we demonstrate their utility in different neuroscientific datasets.

7

Comparing functional MRI data representations based on distributions of repeatable topological features

Ty Easley, Kevin Freese, Elizabeth Much, Janine Bijsterbosch

Washington University in St. Louis

Despite much recent attention, comparing reduced-rank representations of resting-state functional MRI (rfMRI) data remains challenging. We approach this problem by examining the effect of varying choices of representation of a single dataset on varying measures of inter-subject variability within the dataset. We offer a framework for evaluating the topological similarity of reduced-rank representations of resting-state fMRI data. Because the dimensions of both the original data and its reduced-rank representations are typically large compared to sample size, topological structure in this problem regime is more likely captured in the distribution of topological “noise” than by a few prominent generators. Treating each representational map as a metric embedding, we compute the Vietoris-Rips persistence of its image. We then use the topological bootstrap to analyze the re-sampling stability of each representation, assigning a “prevalence score” for each nontrivial basis element of its persistence module. Finally, we compare the persistent homology of representations using a prevalence-weighted variant of the Wasserstein distance. Notably, our method is able to compare representations derived from different samples of the same distribution and, in particular, is not restricted to comparisons of graphs on the same vertex set. In addition, representations need not lie in the same metric space. We apply this analysis to a cross-sectional sample of representations of rfMRI data in a large cohort and hierarchically cluster under the prevalence-weighted Wasserstein. In our analysis, resampling stability of topological features carried more distinguishability than persistence. We find that the type of rfMRI representation is a stronger predictor of the number and stability of topological features than its decomposition rank. Specifically, our findings suggest that rfMRI brain representations clusters distinguish between amplitude-based representations that reflect network strength and connectivity-based representations that reflect network spatial organization and/or temporal connectivity.

9

Laying the Groundwork for Semantic Decoding of Object Categories in Naturalistic Movies via High-Density Diffuse Optical Tomography

Wiete Fehner[1], Zachary Markow[1], Morgan Fogarty[1], Aahana Bajracharya[1], Dana Wilhelm[1], Alexander G. Huth[2], Joseph P Culver[1]

[1] Washington University in St. Louis, St. Louis, Missouri 63130, USA
[2] The University of Texas at Austin, Austin, 78712, TX, USA

Author e-mail address: f.wiete@wustl.edu

Functional magnetic resonance imaging (fMRI) studies have validated advanced decoding approaches for visual scenes, semantic categories, and natural language using naturalistic stimuli like movies and podcasts (Huth et al., 2012; Huth et al., 2012; Tang et al., 2023). These advancements are essential for understanding and mapping language recovery in individuals experiencing language impairments, like aphasia, following a stroke (Baldassarre et al., 2016). However, fMRI's usefulness is limited in everyday settings due to its size, cost, and physical requirements. On the other hand, high-density diffuse optical tomography (HD-DOT) provides a portable, cost-effective alternative for brain mapping in natural settings. This study aims to develop an experimental approach to encode and decode semantic categories from natural movie scenes using very high-density DOT (VHD-DOT).

One participant completed three 89-minute VHD-DOT imaging sessions. To ensure session repeatability, cap placement followed a precision protocol developed by (Bajracharya et al. 2023). Each session included four 10-minute training, and three 10-minute testing movie runs. Movie clips were validated in prior fMRI studies and labeled using semantic categories from WordNet (Huth et al., 2012; Huth et al., 2012).

Data processing utilized a standard DOT pipeline (Eggebrecht et al., 2014), encompassing pre-processing, reconstruction, and spectroscopy. Brain activation maps for individual runs were created using a general linear model for auditory and visual localizer tasks. Correlation analysis and explainable variance (EV) showed high repeatability across imaging sessions for localizer tasks and test movies. EV represents the variance of the signal explainable by the stimulus. A higher EV suggests that the model can predict a greater portion of the observed signal in the test set, indicating a better model fit.

10

Optimal dynamics in the manifold for spontaneous behavior

Antonio J. Fontenele[1], J. Samuel Sooter[1], Nivaldo A. P. de Vasconcelos[2] and Woodrow L. Shew[1]

[1] Department of Physics, University of Arkansas, Fayetteville, 72701, AR, USA
[2] Department of Biomedical Engineering, Federal University of Pernambuco (UFPE), Recife, PE 50670-901, Brazil.

Sensory, cognitive, and motor functions emerge from the interplay of many neurons in brain circuits. Simultaneous recording of behavior and activity from multiple neurons together with dimensionality reduction techniques suggest that high-dimensional neuronal space can be decomposed into different subspaces for encoding different behaviors. Well-defined, repetitive tasks have been associated with low dimensional manifolds, few neuronal modes explain the task-related data variance. On the other hand, recent experiments have shown that spontaneous behavior is linked to a high-dimensional neuronal state space when no specific task constrains the animal behavior. This scenario raises several questions, including how a variety of behaviors shapes neuronal manifolds and whether these manifolds share neural modes across specific behaviors. Another theory, heavily inspired by principles of critical phenomena, hypothesizes that the cerebral cortex operates close to a critical phase transition. This hypothesis offers an explanation for the observed complexity of brain dynamics and is biologically significant due to potential computational advantages near criticality. Recently, it was shown that the dominant neuronal modes (first few PCs), identified from the high-dimensional space during spontaneous behavior, exhibit fingerprints of optimal dynamics (i.e. criticality). Considering the computational advantages associated with criticality, it may be that the critical manifold drives specific behaviors. However, previous work did not determine how this manifold for optimal dynamics is related to behavior. Thus, we investigated how the neuronal manifold encoding spontaneous behavior relates to the neuronal modes spanning the critical manifold. We have found that the critical manifold is linked to the neuronal subspace governing many spontaneous behaviors. Ultimately, we postulate that criticality underlies non-task-constrained behaviors, such as foraging, a mechanism that has been honed through the evolutionary process. This concept suggests that criticality provides benefits by maximizing the likelihood of survival.

12

Quasicriticality in Light and Dark Conditions

Leandro J Fosque, Yifan Xu, Aidan Schneider, and Keith Hengen

Washington University in St. Louis

The primary visual cortex (V1) must adapt to varying light conditions to maintain an optimal computational regime. Previous research has shown that firing rates in V1 are actively stabilized during transitions between light and dark environments (Pacheco et al., 2019). However, the impact of these changes on network dynamics and computational efficiency has not been extensively explored. Furthermore, it has been demonstrated that networks operating near criticality can achieve an optimal computational regime (Beggs, 2008). Yet, the influence of external input on the scaling exponents of neuronal avalanches remains poorly understood. To address this gap, our study focuses on the primary visual system in young rodents, a period when the network is highly plastic. We employ the Quasicriticality Hypothesis (qC) as a framework to understand brain dynamics, accounting for variations in external input (Fosque et al., 2021; Williams-Garcia, 2014). In this study, we present experimental findings that corroborate the predictions of qC, shedding light on the effects of light and dark conditions on the network dynamics of V1, in freely behaving animals. The implications of our findings further the search for a unifying framework of brain dynamics, offering a deeper comprehension of the dynamic balance between neural stability and flexibility, crucial for effective information processing in the brain.

13

Parvalbumin expression does not account for discrete electrophysiological profiles of glutamatergic ventral pallidal subpopulations

Robert D Graham[1,*], Lisa Z Fang[1], Jessica R Tooley[1,2], Vani Kalyanaraman[1], Mary Christine Stander[1], Darshan Sapkota[3,4], Joseph Dougherty[2,3], Bryan Copits[1,2], Meaghan C Creed[1,2]

[1] Department of Anesthesiology, Washington University in St. Louis
[2] Graduate program in neuroscience, Division of Biological and Biomedical Sciences
[3] Department of Psychiatry, Washington University in St. Louis
[4] Department of Biological Sciences, University of Texas at Dallas

*presenting author (robertg@wustl.edu)

The ventral pallidum (VP) is a critical node in the mesolimbic reward system. Modulating the VP can impact reward valuation and seeking, making it an attractive target for neuromodulation therapies that manage substance use disorders (SUDs). To understand how to rationally modulate the VP, we need a better understanding of the electrophysiological properties of VP neurons and the biophysical determinants of these properties. Here, we used patch-clamp electrophysiology to characterize the intrinsic properties of glutamatergic VP (VP$_{Glu}$) neurons and observed two distinct electrophysiological profiles: VP$_{Glu}$ neurons that undergo depolarization block in response to progressively increasing current injection amplitudes and those that were resistant to depolarization block. To explore the mechanisms that could contribute to these distinct profiles, we used targeted ribosome affinity purification to isolate actively transcribed mRNA selectively from VP$_{Glu}$ neurons to identify ion channel subunits present in these cells. We used this transcriptomic information to implement a Markov Chain Monte Carlo method to parameterize a large population of biophysically distinct multicompartment cable models of VP$_{Glu}$ neurons conforming to each subpopulation. Prior literature suggests parvalbumin (PV) is expressed in a subset of VP$_{Glu}$ neurons and that PV expression governs the firing properties of those neurons. We tested the hypothesis that PV expression accounted for differences in firing properties in these subgroups. Our models determined that PV expression at physiological levels does not determine firing phenotype. However, supraphysiological expression levels of PV induced depolarization block in previously depolarization block-resistant neuron models, suggesting that other calcium-buffering proteins could play a role in determining the firing phenotype of VP$_{Glu}$ neurons. We corroborated this result with singlecell patch-clamp RT-PCR, which confirmed that PV expression did not distinguish these subpopulations. Together, these findings establish that VP$_{Glu}$ neurons are composed of biophysically distinct subpopulations that have not been appreciated in prior studies. With the advent of novel tools for cell-type specific pharmacology and targeted neurostimulation, this understanding will be critical for developing strategies to rationally modulate VP$_{Glu}$ cells to treat SUDs.

14

Cue-Paced Training Protocols: An Assessment of their Impact on Modulating Lower Limb Excitability

Rachel Hawthorn[1,2], Rodolfo Keesey[1,2], Carolyn Atkinson[1,2], Haolin Nie[1], Ismael Seáñez[1,2,3]

[1] Department of Biomedical Engineering, Washington University in St. Louis
[2] Division of Neurotechnology, Washington University School of Medicine in St. Louis
[3] Department of Neurosurgery, Washington University School of Medicine in St. Louis

Spinal cord injury (SCI) leads to long-lasting motor impairments affecting hundreds of thousands of individuals in the United States, and currently, there is no cure for paralysis. When paired with activity-based training, spinal cord stimulation (SCS) has demonstrated functional enhancements in people with SCI. However, the neural mechanisms potentially responsible for this neurorecovery remain poorly understood. Moreover, the specific characteristics that make certain rehabilitation exercises more likely to enhance recovery remain unknown. In this work, we evaluate short-term changes in neural excitability induced by 30 minutes of cue-paced motor-skill and isometric resistance training (IRT) protocols. Transcranial magnetic stimulation (TMS) was used to record motor-evoked potentials (MEPs) for examining the contributions of the corticospinal tract. Peripheral nerve stimulation generated motor, H-reflex, and F-wave responses in the tibialis anterior and medial gastrocnemius, allowing examination of spinal motoneuron tract excitability. Cue-paced motor-skill training involved a non-invasive body-machine interface (BoMI), where participants utilized their right foot to control a 2D cursor, completing five sets of reaching tasks with 1Hz auditory cues. Cue-paced IRT utilized a biodex isokinetic dynamometer, with participants executing three sets of ankle flexion and extension exercises with two minutes of rest between each set, again with 1Hz auditory cues pacing the sets. Excitability assessments were conducted before and after 30-minute training sessions, during rest, and at 15% active dorsiflexion. We hypothesize that both training protocols will increase corticospinal excitability, while spinal excitability will remain unchanged. This study refines a 30-minute training protocol for lower limbs in control participants, utilizing cue-paced motor-skill training and isometric resistance training. The outcomes of this research will facilitate future translational investigations into neural plasticity in individuals with SCI undergoing SCS-assisted rehabilitation.

16

Mechanisms of neuronal coordination in multi-layer networks

Tianfeng Lu[1], Shuyu I. Zhu[1], Robert Wong[1, 2], Geoffrey J. Goodhill[1]

[1] Departments of Developmental Biology and Neuroscience, Washington University in St. Louis
[2] Departments of Electrical & Systems Engineering, Washington University in St. Louis

---

Anatomical hierarchies in neural systems support detection and response towards external input, thus the downstream region should reliably read from noisy upstream population[1]. The readout strategy in either biological of artificial neural networks is still not clear. Here we combine large scale neuronal imaging, recurrent neural network modelling and mathematical analysis to demonstrate a naturally emerged network implementation that enhances readout against correlated noise in both biological and artificial regimes.

Building on recent work in mice showing that inter-area communication supports reliable sensory processing despite variability of neuronal representation[2], we take advantage of the transparency of the larval zebrafish to observe brain-wide neural population activity during a visuomotor transformation. We show that, while moving visual spots drive localized neural coordination in the tectum, the neural representation of the resulting tail responses is coordinated across the whole brain and dominates the majority of trial-to-trial fluctuations. Furthermore, the population covariance structure is essential for decoding movement but not the visual information. A multi-layer *in silico* network simulation also exhibits similar properties. By linking optimal decoding theory and the training objective function, we demonstrate that the differentiated readout strategy across layers is a consequence of parameter optimization. Our findings reveal an approach for both the brain and artificial neural networks to extract information from noisy population activity, and indicate a computational principle underlying brain-wide neural dynamics.

[1] Rowland, J. M. *et al*. Propagation of activity through the cortical hierarchy and perception are determined by neural variability. *Nature Neuroscience 26*, 1584-1594, doi:10.1038/s41593-023-01413-5 (2023).

[2] Ebrahimi, S. *et al*. Emergent reliability in sensory cortical coding and inter-area communication. *Nature 605*, 713-721, doi:10.1038/s41586-022-04724-y (2022).

17

Neurogenesis enhances olfactory coding efficiency, especially in variable environments

Ryan McGee[1], Gaia Tavoni[1]

[1] Department of Neuroscience, Washington University in St Louis

The olfactory bulb (OB) is one of a few brain regions where new neurons continue to be integrated throughout the lifespan, with many of the OB's local inhibitory granule cells (GCs) being generated in adulthood. The effects of adult neurogenesis are believed to contribute to the odor discrimination function of the OB, but the impact of neuronal replacement has not been assessed (a) in terms of optimal coding benchmarks, or (b) in relation to the concurrent role of synaptic plasticity. Through a combination of theory and computational modeling, we show that synergy between local Hebbian learning and GC survival rules enables a network to learn the global statistical structure of the inputs, which enhances the coding efficiency of olfactory circuits beyond what can be achieved with Hebbian plasticity alone. That is, plausible mechanisms of neuronal replacement and integration facilitate population codes that better maximize the decodability of neural representations. Together, the effects of neurogenesis significantly improve the OB's ability to maintain high coding efficiency when stimuli statistics vary over time. These results connect normative predictions derived from optimal coding theory with biologically plausible processes, providing a foundation for further study of the role of neurogenesis in sensory computation.

18

Neuro Translate: Using Surface Image Transformers to translate shared latent information across functional MRI brain representations

Samuel Naranjo Rincon, Fyzeen Ahmad, Ty Easly, Ulugbek Kamilov, Janine Bijsterbosch

Washington University in St. Louis

Resting state functional MRI has become a main source of data for identifying biomarkers across a range of disorders for over a decade, yet there has been little translation to clinical practice. A key factor in this gap is the lack of standardized brain representations to analyze the fMRI data. A brain representation consists of a low-dimensional spatial description of the brain into regions or networks and a metric of interest such as rfMRI amplitude or connectivity. As a result, two labs can have the same preprocessed fMRI data, research question, and hypothesis, but adopt different brain representations that make it impossible to compare the final results. The goal of this study is to develop a brain representation translator to translate results obtained from one brain representation into results using a different brain representation, without costly and often infeasible re-analysis. As a starting point, we developed a translation between two common brain representations: Independent Component Analyses (ICA) and Profumo using the Human Connectome Project Young Adult sample (N = 1002). We built upon a previous deep learning model known as a Surface Image Transformer (SiT) that embeds meshes of MRI brain topology for each brain representation. We exploit the utility of attention from vanilla transformers to identify the crucial topological segments that best inform translation from input brain representation to output brain representation. Our results showed that the SiT performs with high accuracy to decode Profumo results based on ICA input at the group level (mean absolute error rate of 0.05). In future work, we will expand to a larger set of brain representations. We hope that by comparing many pairwise translations between brain representations and identifying shared layers between deep learning models, we may be able to identify a principal latent fMRI subspace.

19

Generative Modeling of the Functional Connectome in First-Episode Psychosis

Jacob Pine, Deanna Barch, and Matthew Singh

Washington University in St. Louis, Department of Psychological and Brain Sciences

Psychotic disorders are thought to manifest in excitation-inhibition (E:I) balance abnormalities. This study used Mesoscale Individualized Neurodynamic (MINDy) modeling to investigate brain-wide E:I-associated BOLD dynamics among individuals with First-Episode Psychosis (FEP). MINDy models the fMRI BOLD timeseries using sets of abstracted, biophysically-inspired, neural populations. A set of equations describes how the activity of each neural population evolves over time, and how this activity manifests in the fMRI BOLD signal. Existing work suggests that systematic changes in E:I properties across the cortical hierarchy contribute to heterogeneous timescales of activity, which are evident in fMRI BOLD dynamics. MINDy's decay rate parameter captures systematic changes in fMRI BOLD dynamics across the cortical hierarchy, and may be useful for investigating E:I abnormalities. This study analyzed differences in decay rates between those with FEP and otherwise healthy individuals as a method for investigating E:I abnormalities among those with FEP. This study utilized data from the Human Connectome Project – Early Psychosis study. The final sample included 139 individuals with FEP and 61 without. Statistical significance was determined using bootstrapped estimates of parameter differences, adjusted for multiple comparisons. Individuals with FEP had significantly greater decay rates in all functional networks, except for the Somatomotor network and subcortical areas. When normalizing decay rates across parcels within an individual's brain, which reflects relative differences in decay rates across individuals, those with FEP had significantly greater decay rate for Dorsal Attention (p=0.001) and Limbic (p=0.002) networks, which normally have slow dynamics, and had a significantly lower decay rate in the Somatomotor network (P=0.001), which normally has fast dynamics. This indicates that there may be reduced hierarchical differentiation in FEP, which may be related to E:I abnormalities.

20

Noise Gives Brains Super-Turing Computation Computers Are Missing Out

Emmett Redd

Physics, Astronomy and Materials Science, Missouri State University, Springfield, Missouri USA

Although computers implement float-point numbers, their memory and operations only involve integers. Models of neural operations (either artificial or natural) have inputs and outputs that are at most countably unbounded. Newell[1] lamented that there was nothing beyond a Universal Turing Machine (UTM) to account for "…a human[‘s] immense variety of functions…" He thought that brain operation needed to work with the uncountably infinite which amounts to a real number-sized set. Our efforts to develop more brain-like artificial intelligence need to access real numbers. This submission will show a relief to Newell's lament and a way to implement it in brain and artificial intelligence models. It simply involves adding the proper size of noise in the internal computer operations. The quantized universe is like digital computers in not having real-numbered signals but also provides noise that computers are designed to eliminate. The brain works in the quantized and noisy universe. Siegelman[2] proved this is sufficient for higher complexity computation than the integer bound UTM. Adding the proper size random numbers to the internal signals of a recurrent neural network allows computers to provide outputs that gain indirect access to the real number space.[3] If pseudorandom numbers have high enough quality, they are sufficient for computers to perform higher complexity computations.[4] These are simple calculations; they rely on the properties of a quantized universe, avoiding any explicit quantum mechanical calculations. The noise implements the physical phenomenon of stochastic resonance. Smart[5] questions if the next breakthrough in either artificial intelligence or neuroscience would come from it. Russell, et.al.[6] demonstrate its use in paddlefish sensing. Many practitioners already add noise to training data for artificial neural networks to better generalize. A careful, slight change in the location of noise application will gain greater computational complexity that approaches brain activity. Neurons take advantage of discrete, noisy signals in their operation. The next step is to follow that lead of using discrete, noisy signals so computations are more intelligent.

1. Newell, A. *Unified Theories of Cognition*, Cambridge, MA: Harvard Univ. Press, 1990.
2. Siegelmann, H. T. Stochastic Analog Networks and Computational Complexity, Journal of Complexity, 15(4) 451-475, 1999.
3. Redd, E., Senger, S., Obafemi-Ajayi, T.: Noise optimizes super-turing computation in recurrent neural networks. Phys. Rev. Res. 3(1), 013120 (2021)
4. Redd, E., & Obafemi-Ajayi, T., Noise Quality and Super-Turing Computation in Recurrent Neural Networks. In *Artificial Neural Networks and Machine Learning–ICANN 2021: 30th International Conference on Artificial Neural Networks, Bratislava, Slovakia, September 14–17, 2021, Proceedings, Part IV 30* (pp. 469-478). Springer International Publishing (2021).
5. Smart, A., Is noise the key to artificial general intelligence? converging evidence indicates noise plays a fundamental role in the brain, https://www.psychologytoday.com/intl/blog/machine-psychology/201606/is-noise-the-key-artificial-general-intelligence, 2016.
6. Russell, D., Wilkens, L. & Moss, F. Use of behavioural stochastic resonance by paddle fish for feeding. *Nature 402*, 291–294 (1999). https://doi.org/10.1038/46279

21

Decision confidence and belief updating: quantitative behavioral phenotyping of perceptual decision-making in a rat *CHD8*+/- model of autism

Steven C. Ryu, Jack McGuire, Adam Kepecs

Washington University in St. Louis, School of Medicine, MO, USA

Metacognition, or the ability to monitor and evaluate one's own mental states and processes, may be a crucial cognitive trait altered in autism spectrum disorder. Questions such as "How confident am I in what I just decided?" and "Should I adjust my thinking based on past experiences?" are not mere introspection — they are vital tools for navigating the complexities of sensory environments and social interactions, both of which present significant challenges to individuals with autism. Given vast genetic heterogeneity in autism's causes and variation in behavioral symptoms, identifying core cognitive features of autism is a considerable challenge that requires quantitative behavioral phenotyping of individuals. Our laboratory has developed methodologies for the crossspecies identification of computational signatures of psychiatric diseases using perceptual decision-making tasks. Here, we aimed to determine whether metacognition is impaired in a transgenic rat model of autism with a heterozygous knockout of *CHD8*, a high-risk autism gene. Rats are trained on an auditory decision task in which animals invest time in their choices by waiting for delayed rewards. The amount of time they invest reflects their statistical confidence in each decision. Additionally, subsequent decisions are biased by the confidence in and the outcome of previous decisions, allowing us to measure how much subjects' beliefs are updated by past experience. Initial analyses suggest that relative to littermate controls, *CHD8*+/- rats exhibit equal perceptual performance, moderate but heterogenous reduction in discriminability between confidence reports in correct and error trials, and reduced belief updating. Our results position us to make analogous investigations in other transgenic models of autism and genotype-matched human subjects to establish whether altered metacognition is a core cognitive feature of autism.

22

Structural localization is embedded in the spike trains of single neurons

A. Schneider[1], G. Tolossa[1], K. Hengen[1]

[1] Department of Biology, Washington University in St. Louis

Neurons within a brain region have long been observed to exhibit correlated dynamics while neurons from distinct regions receive diverse inputs and display different connectivity patterns. As a result, the electrophysiological features of neurons may vary across functional brain regions, potentially possessing information about their spatial location. However, the extent to which structural location can reliably be extracted from the neural code remains uncertain. Further, the degree to which the neural activity signatures emerge to be apparent at the circuit level or manifest within individual neurons is unknown. In this study, we investigated the differences in spike trains among neurons across various functional brain regions using machine learning. Using the visual coding neuropixels data from the Allen Brain Observatory, we extracted spike train features and trained machine learning models to classify neuronal units into their functional brain regions. Our findings demonstrate that discernible patterns exist within individual neurons, which support their localization at arbitrary levels of structural hierarchy such as functional visual cortex subareas, hippocampal subregions, and thalamic nuclei. This suggests the presence of spatial fingerprints within neuronal spike trains, which may provide an additional level of computational advantage in the brain. Beyond an improved understanding of neuronal computation, these findings have practical implications for neuronlevel brain-region localization directly from electrophysiological recording.

23

Relative neural population size modulates learnability of cyclic features of neural codes

Niko Schonsheck, Chad Giusti

University of Delaware, Oregon State University

Cyclic structures are a class of mesoscale features ubiquitous in both experimental stimuli and the activity of neural populations encoding them. Important examples include encoding of head direction, grid cells in spatial navigation, and orientation tuning in visual cortex. The central question of our present work is: how does the brain faithfully transmit cyclic structures between regions? Is this a generic feature of neural circuits, or must this be learned? If so, how?

While cyclic structures are difficult to detect and analyze with classical methods, tools from algebraic topology have proven to be particularly effective in understanding cyclic structures. Recently, work of Yoon et al. develops a topological framework to match cyclic coding patterns in distinct populations that encode the same information. We leverage this framework to study the efficacy of Hebbian learning rules in propagating cyclic structures through neural systems.

Our primary results are 1) feedforward networks with connections drawn from inhibitory-biased random distributions do not reliably propagate cyclic features of neural coding 2) updating network connections with a biologically realistic Hebbian learning rule modeling spike timing dependent plasticity robustly constructs networks that do propagate cyclic features and 3) under biologically plausible parameter choices, the inhibition and propagation of such features can be modulated by the size of the output neuron population. These results suggest that \unpacking” geometry through high-dimensional en-codings is necessary when applying biologically inspired learning rules in the context of naive neural connectivity. However, when the ratio of dimensions is large enough, such rules can generically produce faithful representations of complex stimulus spaces.

24

Visual cortex is nearest to criticality at intermediate level of arousal: a temporal renormalization group approach

J. Samuel Sooter[1], Andrea K. Barreiro[2], Cheng Ly[3], Antonio J. Fontenele[1], Nivaldo A.P. de Vasconcelos[4], Woodrow L. Shew[1]

[1] Department of Physics, UA Integrative Systems Neuroscience, University of Arkansas
[2] Department of Mathematics, Southern Methodist University
[3] Department of Statistical Sciences and Operations Research, Virginia Commonwealth University
[4] Department of Biomedical Engineering, Federal University of Pernambuco

---

Cerebral cortex has long been hypothesized to operate close to a critical phase transition. The main evidence for this hypothesis is that bursts of cortical activity, called avalanches, are often approximately scale-free, with power-law distributions of size and duration. Avalanche-based criticality signatures require long experiments (many avalanches are needed to estimate the avalanche distributions) and can be challenging to interpret. Here we develop a fundamental theory of multi-timescale brain dynamics–temporal renormalization group theory (tRG)–and use it to obtain a more direct measure of nearness to criticality. We performed tRG analysis on autoregressive (AR) models of arbitrary order, identifying multiple types of criticality (i.e. tRG fixed points) and the conditions on the AR model's history kernel necessary to flow into them. We also developed practical tRG-driven tools to measure proximity to criticality for experimental data. More precisely, we calculate the Euclidean distances from a best-fit AR model to the basins of attraction of each of the tRG fixed points. As a first application, we analyzed extracellular electrophysiological recordings of unit activity from primary visual cortex of freely behaving mice. We found that the best-fit AR model closely reproduces the avalanche distributions and spectral properties of the population activity during both NREM sleep and wakefulness. The distances from the best-fit AR model to the basins of attraction of each of the fixed points are then a reasonable estimate of the true distance to criticality. We found that the best-fit AR models in NREM are much further from criticality than those in the awake state. Comparing avalanche distributions, power spectra, and DFA fluctuation functions between the two states corroborates this conclusion. To better understand the relationship between the animal's behavioral state and nearness to criticality, we next segmented the experiments into minute-long windows, performing an AR model fit and calculating the distance to criticality for each window separately. We found a prominent minimum in distance to criticality at an intermediate level of alertness (measured by EMG), suggesting that the cortex is closer to criticality in the quiescent awake state than in either the hyper-alert awake state or sleep.

25

Circumstantial evidence and explanatory models for synapses in large-scale spike recordings

Ian H. Stevenson

University of Connecticut

Whether, when, and how causal interactions between neurons can be meaningfully studied from observations of neural activity alone are vital questions in neural data analysis. Here we aim to better outline the concept of functional connectivity for the specific situation where systems neuroscientists aim to study synapses using spike train recordings. In some cases, cross-correlations between the spikes of two neurons are such that, although we may not be able to say that a relationship is causal without experimental manipulations, models based on synaptic connections provide precise explanations of the data. Additionally, there is often strong circumstantial evidence that pairs of neurons are monosynaptically connected based on anatomy, spike waveforms, patterns of plasticity, and patterns across neurons. Here we illustrate how circumstantial evidence for or against synapses can be systematically assessed and show how models of synaptic effects can provide testable predictions for pair-wise spike statistics. We use case studies from large-scale multi-electrode spike recordings to illustrate key points and to demonstrate how modeling synaptic effects using large-scale spike recordings opens a wide range of data analytic questions.

27

Dimensionality and developmental trajectory of signal and noise in the brain

Robert Wong[1,3,4], Alisha Tromp[2], Zac Pujic[2], Biao Sun[2], Naoki Hiratani[3], Geoffrey Goodhill[3,4]

[1] Department of Electrical & Systems Engineering, Washington University in St. Louis, MO, USA
[2] Queensland Brain Institute, Australia
[3] Department of Neuroscience, Washington University School of Medicine, MO, USA
[4] Department of Developmental Biology, Washington University School of Medicine, MO, USA

A major question in systems neuroscience is how the brain encodes and processes visual information using stochastic neural activity. Previous work has improved our understanding of coding mechanisms and noise structure by recording from subsets of neurons in a range of visual areas, and recent work has focused on extending these analyses to larger-scale recordings. Despite the large number of neurons, the intrinsic dimensionality of this neural code is often much lower. However the representation of the signal and the relevant noise correlation structure across multiple brain regions, and how this changes over development, are unknown. Here we employ recently developed statistical and computational tools to probe the neural code within the entire zebrafish brain in response to smoothly parameterized low-dimensional visual stimuli, and examine the developmental trajectory of these responses, focusing on the change in both the signal and the noise correlation structure. We find that the decoding accuracy of visual information from the optic tectum improves over development due to both better separation of the signals and changes to the noise correlation structure which enhance signal discriminability. While decoding performance for visual stimuli is still quite accurate in brain regions downstream of the tectum, its representation displays more complex structure, and shows greater inter-subject variability than for the optic tectum. This study advances our understanding of the developmental trajectory of neural encoding, offering novel insights into the signal and noise structures of visual processing.

28

Border ownership signals emerge in an artificial neural network trained to predict future visual input

Zeyuan Ye, Ralf Wessel, Tom P. Franken

Washington University in St. Louis

To identify the objects that surround us, the brain needs to segment visual scenes into organized collections of objects. A dominant segmentation signal in the visual cortex of the primate brain is border ownership: border ownership neurons signal which side of a border is owned by a foreground surface. Neural border ownership is known to display hysteresis with scene changes or object motion, suggesting that these neurons aid in processing dynamic visual scenes. Here we explore whether border ownership signals emerge in a neural network trained to predict future visual input.

We evaluated whether units in PredNet, an artificial neural network trained to predict the next frame in natural videos, are selective for border ownership. We measured the response of units in PredNet to static scenes of an isoluminant square on an isoluminant background. We focused our analysis on units for which the classical receptive field only contained a luminance contrast border which could – in different trials – be owned by a square on one or the other side (border ownership). Scenes also varied in border orientation and luminance contrast polarity.

We found that R and E units in PredNet are often selective for border ownership, irrespective of the luminance contrast polarity. The preferred side of ownership was remarkably tolerant to border orientation, similar as in our data from the non-human primate brain. The proportion of border ownership units was higher in deeper layers (L1,2,3>L0). The magnitude of border ownership signals increased with depth (L3>L2>L1>L0).

Our data show that units selective for border ownership emerge in PredNet even though this network was not explicitly trained to segment visual scenes, but instead to predict the next frame in natural videos. This suggests that border ownership units in neural networks aid in efficiently predicting future input in natural videos.

29

A recurrent neural network model of excitatory-inhibitory imbalance in autism spectrum disorder

Daniel Zavitz[1], ShiNung Ching[2], Geoffrey Goodhill[1]

[1] Departments of Developmental Biology and Neuroscience
[2] Department of Electrical and Systems Engineering
Washington University in St. Louis

The proportion of excitation (E) and inhibition (I) that neurons receive has a fundamental influence on neural network activity and is critical for effective computation. An E/I imbalance is a leading hypothesis for the mechanism underlying autism spectrum disorders (ASDs), amongst other disorders. Numerous studies of ASD mouse models have found genetic, neuronal, and circuit changes that disrupt E or I balance. However, it remains unclear the precise sense in which E and I are imbalanced in ASDs, what computational changes result from such an imbalance, and how this selectively affects different circuit functions. To address these questions, we constructed a recurrent neural network (RNN) model of cortical circuits constrained to have separate excitatory and inhibitory populations, then trained this RNN to perform a variety of cognitive tasks. Once trained, we performed systematic sensitivity analysis on network parameters, inducing hyperexcitable or hypoexcitable regimes and tracking the resulting shifts in task performance. We exploit symmetries in RNN dynamics to perform formal model reduction into a 1D E/I imbalance space, and thus reveal structure in the task-specific degradation in computational performance. By analyzing fixed-points we demonstrate that the changes induced in fixed-point geometry by E/I imbalance mirror changes in task performance. To close the loop between circuit structure, task performance, and dynamics, we apply contraction theory to derive fixed-point geometry from simple circuit features. Together, our study illuminates the computational significance of excitatoryinhibitory balance and its implications for ASDs and other psychiatric disorders.

30

Capacity of Networks with Arbitrary Topologies and Neuron Activation Probabilities

Kaining Zhang and Gaia Tavoni

Department of Neuroscience, Washington University School of Medicine

An important issue in systems neuroscience is estimating the memory capacity of neural networks. The Hopfield model, a type of neural network storing memory patterns as stable attractors of binary units, offers a framework for this inquiry. Gardner's analysis [1] calculated the capacity of fully-connected networks with homogeneous activation probabilities. However, in the brain, memories are stored across heterogeneous neural populations. Here, we present an analytical solution for networks with arbitrary topologies and activation probabilities, leveraging a perceptron model with arbitrary input statistics.

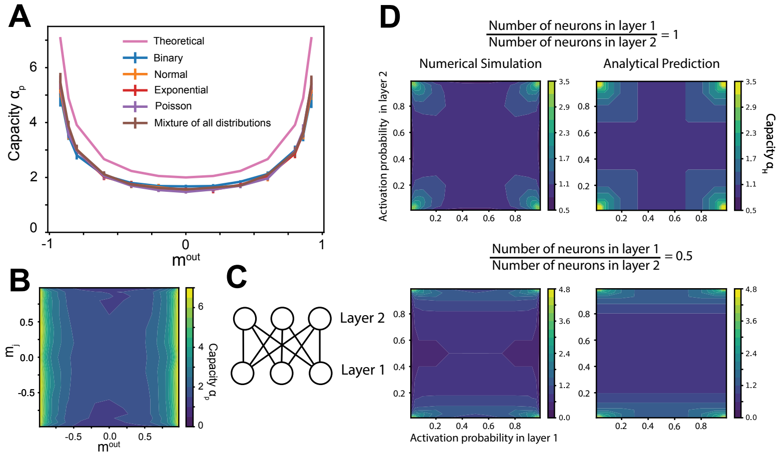

Figure 1: **A.** The analytical perceptron capacity (magenta curve) provides an upper bound (and a close approximation) to numerical estimates obtained for different input state distributions. **B.** The capacity is mainly just a function of the output state statistics and not the input state statistics. m_j: Expectation of the input binary distribution for unit j, here all m_j have the same value. **C.** Two-layer network representing the hippocampal-cortical memory system. **D.** Capacity of the network in (C) for different activation probabilities and relative number of neurons in the two layers.

A perceptron, a simple one-layer feedforward network, learns to associate input patterns with binary labels. We generalize Gardner's method to compute the capacity of perceptrons with arbitrary output binary distributions and input distributions. The resulting capacity expression (Equation 1) solely depends on the expectation of the output binary distribution with states $\{-1, 1\}$, m^{out}, or (equivalently) the output activation probability $p^{out} = (m^{out} + 1)/2$; it does not depend on the input distributions (Figure 1A,B). $\alpha_p \triangleq C_p/N$ is the maximal number of patterns per neuron that the perceptron can store.

$$\alpha_p\left(m^{out}\right) = 2\left\{1 - m^{out}\int_{-M}^{M} Dy\right\}^{-1}, \quad \frac{1+m^{out}}{2}\int_{M}^{\infty} Dy\,(y-M) = \frac{1-m^{out}}{2}\int_{-M}^{\infty} Dy\,(y+M) \quad (1)$$

Using the analytical solution for the perceptron capacity, we can derive the capacity of Hopfield models with arbitrary network topologies and heterogeneous activation probabilities. Notice that a Hopfield model with N units can be seen as a collection of N perceptrons with each unit as the output, and all units projecting to it as the input. The capacity of the Hopfield model is dictated by the smallest of the capacities of the N perceptrons: $\alpha_H = \frac{1}{N}\min_i\left(C_p(m_i, N_i)\right) = \min_i\left(\frac{N_i}{N} * \alpha_p(2p_i - 1)\right)$

Utilizing this, we calculate the capacity of a two-layer neural network resembling the hippocampal-cortical memory system (Figure 1C,D). The network maximizes capacity when activation probabilities in both layers are either low or high.

Furthermore, based on the network capacity equation, we propose a conjecture regarding the number of input connections of each neuron and its activation probabilities, which is testable in the biological memory system. In networks with optimized capacity, neurons with low activation probabilities ($p_i \ll 1/2$, i.e., high $\alpha_p(2p_i - 1)$) should be more sparsely connected (lower N_i) than neurons with higher activation probabilities ($p_i \lesssim 1/2$).

[1] Gardner, E. (1988). The space of interactions in neural network models. *Journal of Physics A: Mathematical and general*, *21*(1), 257.

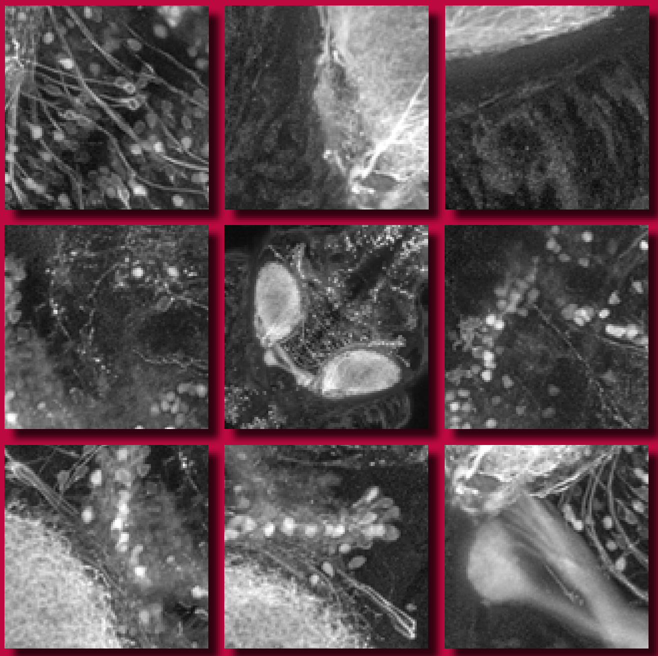

Washington
University in St. Louis