AREADNE 2024

Research in Encoding and Decoding of Neural Ensembles

Eliopoulos Conference Center, Milos, Greece

25–29 June 2024

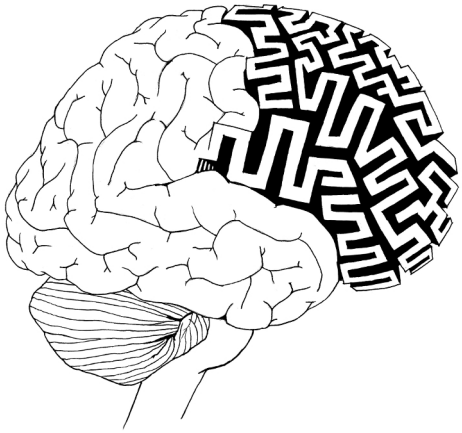

Conference Information

Schedule and Program

Invited Speaker Abstracts

Poster Abstracts

Attendee and Author Index

AREADNE 2024
Research in Encoding and Decoding of Neural Ensembles
Eliopoulos Conference Center, Milos, Greece, 25–29 June 2024
John S. Pezaris, Nicholas G. Hatsopoulos, editors

Published by The AREADNE Foundation, Inc., Cambridge, Massachusetts, USA, https://areadne.org, info@areadne.org
Single copy price USD 50
ISSN 2154-6819 (on-line)
ISSN 2155-3203 (print)
ISSN 2155-319X (CD-ROM)

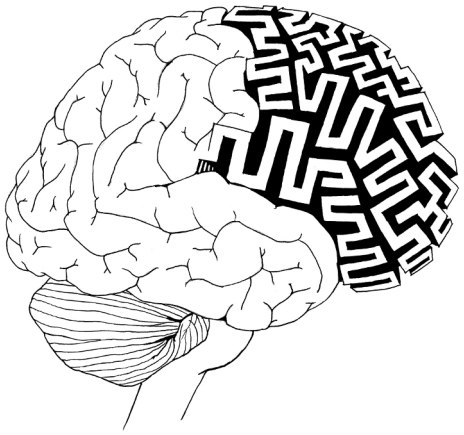

FOREWORD

Foreword

This year's meeting marks a new chapter for AREADNE as we move the conference from the volcanic drama of Santorini to the gentle serenity of Milos. As K. P. Kavafy's poem "Ithaka" [1] reminds us, the adventures and knowledge gained along a path to a far-off destination can outshine the eventual arrival. And so, our move to Milos promises fresh perspectives and renewed inspiration that we relish as a vital part of scientific inquiry. Milos, often described as a hidden gem of the Aegean, offers tranquil and less well-trodden surroundings. From its unusual landscapes to its vernacular architecture, from the other-worldly Sarakiniko beach to the *sirma* garages used to store traditional fishing boats, Milos offers to us an atmosphere of deep beauty, peace and reflection.

This transition of location is not just a change of scenery; it is a metaphor for the journey of scientific exploration. In Milos, we find an ideal setting to pause, reflect, and delve into deeper contemplation of our research on the functioning of neuronal ensembles. Following Kavafy's advice to bar mythical threats from our minds to prevent them from haunting us, Milos encourages us to elevate our thoughts to discourse and discovery.

As we gather in the Eliopoulos Conference Center, overlooking the elegant bay, let us think on the importance of our efforts. Our work in understanding how networks of neurons give rise to higher-order brain functions such as perception, learning, memory, cognition, and action continues to advance with breathtaking results. The development of brain-machine interfaces, which enhance the capabilities of individuals with sensory and motor deficits, exemplifies the profound impact of our research. The natural beauty of Milos will undoubtedly foster collaboration and innovation, allowing us to investigate the mysteries of the brain with fresh eyes and renewed vigor.

In the spirit of "Ithaka," we embrace this change, the journey to a new location, with the excitement that new horizons bring. As we pursue groundbreaking research, we will find joy in our discoveries, forging a path to knowledge that is full of adventure and surprise.

John S. Pezaris, Ph.D.

Nicholas G. Hatsopoulos, Ph.D.

1. K. P. Kavafy, "Ithaka" ($I\theta\acute{\alpha}\kappa\eta$) 1911

WELCOME

Welcome

Welcome to AREADNE 2024, the tenth AREADNE Conference on Research in Encoding and Decoding of Neural Ensembles.

A central challenge in neuroscience is to understand how the activity within networks of neurons gives rise to the higher order functions of the brain including learning, memory, cognition, perception, action and ultimately conscious awareness. Over fifty years of electrophysiological recordings in behaving animals have produced significant insights into what the firing patterns of single neurons encode in isolation, but many of the mysteries of how collections of neurons interact to perform these functions remain.

Technological advances have allowed us to glimpse into the global functioning of the brain. Tools such as multielectrode electrophysiology, multi-photon microscopy and connectomics have expanded our understanding beyond single neurons and into ensembles. We routinely observe the activity of dozens and even thousands of individual neurons simultaneously, and deduce the connectivity between them.

At the same time, our understanding of how neuronal ensembles carry information has allowed the development of brain-machine interfaces (BMIs) to enhance the capabilities of patients with sensory and motor deficits. Knowledge of how neuronal ensembles encode sensory stimuli has made it possible to develop perceptual BMIs for the hearing and visually impaired. Likewise, research in how neuronal ensembles decode motor intentions has resulted in motor BMIs by which people with severe motor disabilities can control external devices.

Conference Mission Statement

There are three major goals of this conference. First and foremost, this conference is intended to bring scientific leaders from around the world to present their most recent findings on the functioning of neuronal ensembles. Second, the meeting will provide an informal and beautiful setting on Milos in which attendees can discuss and share ideas outside of the presentations at the conference center to develop professional relationships and collaborations. Third, this conference continues our long term goals to promote systems neuroscience within Greece by providing a forum for scientists from around the world to interact with Greek researchers and students.

Organizing Committee

The Organizing Committee for AREADNE 2024 has been co-chaired by John Pezaris and Nicholas Hatsopoulos, with highly-valued contributions from members Dora Angelaki, Kenny Blum, Yiota Poirazi, Thanos Siapas, and Andreas Tolias.

Local organization effort has been provided by Nike Makres with assistance from Olympia Tziampiri and Ariadne Pangalos.

Sponsors and Support

Our conference is being sponsored with generous gifts from Mrs. Daphne Hatsopoulos through the NIMA Foundation, and The Gatsby Charitable Foundation to the University of Chicago, along

with a gift from the Simons Foundation to the Massachusetts General Hospital. We have received generous in-kind support from Foley & Lardner, LLC, and both the University of Chicago and Massachusetts General Hospital, where the conference is co-administered.

Sarah J. Caddick, Ph.D.
Advisor to Gatsby Charitable Foundation
CEO and Founder, THALAMIC Ltd
sjc@sciencecuratedby.com

Alyssa Picchini Schaffer, Ph.D.
Senior Scientist
Simons Foundation
apschaffer@simonsfoundation.org

Lawrence W. Vernaglia, J.D., M.P.H.
Chair, Health Care Industry Team
Foley & Lardner, LLP
lvernaglia@foley.com

John Maunsell, Ph.D.
Director, Grossman Institute for Neuroscience
Unversity of Chicago
maunsell@uchicago.edu

Bob S. Carter, M.D., Ph.D.
Chief, Neurosurgical Service
Massachusetts General Hospital
bcarter@mgh.harvard.edu

The Myth of Ariadne

The conference name AREADNE is a combination of the conference title, Research in Encoding and Decoding of Neural Ensembles, and the name of the mythological figure Ariadne. Our brain-to-maze logo was inspired by the central role Ariadne played in the myth of Theseus and the Labyrinth.

In Greek Mythology, Ariadne was the daughter of Minos, king of Crete. King Minos built a large, intricate maze called the Labyrinth to house the Minotaur, a fearsome creature that was half bull, half human. Any who attempted to face the Minotaur perished, either by becoming lost in the maze or from the Minotaur's vicious attack. When the hero Theseus came from Athens to slay the Minotaur, Ariadne gave him a sword and a ball of silk thread. Theseus tied one end of the thread at the Labyrinth entrance and unwound it as he went along, so that after he had found and slain the Minotaur, he could escape from the maze by following the thread back out.

LOCAL INFORMATION

We have assembled a small selection of local information on Adámas and the island of Milos. For more information, select among the many guidebooks written for travel in Milos.

Restaurant Information

Greeks normally eat their evening meal quite late, with restaurants being busiest 10 PM to midnight. The largest meal of the day is often lunch, leading naturally to the habitual afternoon siesta. Tipping at restaurants is not expected, as the cost of service is normally included in the price of the meal. Each euro symbol in the list below is about € 10.

Restaurants in Adámas

O! Hamos!	+30-22870-21672	€ €	traditional Greek
Yankos Restaurant	+30-22870-23615	€ € €	Greek
Nostos	+30-697-196-4981	€ € €	seafood
Alevromilos	+30-22870-23117	€ € €	Greek
O Zygos	+30-22870-23120	€	Greek grilled meats
Mikros Apoplous	+30-22870-24207	€ € €	seafood
Flisvos	+30-22870-22275	€ € €	seafood, traditional Greek; by the port
Volta Restaurant	+30-22870-22858	€ € €	seafood, traditional Greek; by the port

Restaurants and Bars in Plaka

Glaronisia	+30-22870-23480	€ €	Greek food
Methysemi Politia	+30-22870-23100	€ € €	Greek grilled meats
Utopia Cafe-Bar	+30-22870-23678	€ €	great sunset view

Recommended Activities

Milos offers many diversions, including excellent beaches, interesting ancient as well as modern history, high-quality museums, enticing restaurants, and, above all, natural beauty. Some ideas are given below.

Mining Museum
open 10:00–14:00 and 18:00–21:00 (closed Mondays), +30-22870-22481, located on the main road in Adámas

Archeological Museum
open 08:30–15:30 Monday, Wednesday–Saturday; 09:00–22:00 Sunday; closed Tuesday, +30-22870-28026, located in Plaka

Catacombs of Milos
open 09:00-18:45 Monday–Sunday; closed Tuesday, +30-22870-21625, located near Trypiti
Among the most important early Christian monuments in Greece.

Klima Village
A charming fishing village with traditional colorful houses built right on the water’s edge.

Trypiti Village
A quaint village known for its windmills and the ancient theater of Milos.

Sulfur Mines
Abandoned sulfur mines give a glimpse into the island’s industrial past.

Sunset at Plaka
Watch a breathtaking sunset from the Utopia bar in Plaka, offering one of the best views on the island.

Main Beaches

The beaches on Milos are beautiful and varied. They appear on both on the inner bay and around the outside of the island. Arrays of lounges and umbrellas that appear on many beaches are available for rent. A short while after sitting down, someone will approach you and ask for payment, in cash. Typical rates are EUR 5 to 10 for the day.

Tsigrado Beach, not suitable for children
Firiplaka Beach, narrow, pebbly, volcanic cliffs
Provatas Beach, can get crowded, family friendly
Sarakiniko Beach, small, the most photographed on Milos
Papafragas Caves, stunning, tiny, not suitable for small children
Achivadolimni Beach, longest sandy beach on Milos, shallow, family friendly
Firopotamos Beach, small-pebbled beach, family-friendly, no restaurants
Paleochori Beach, most popular, water sports available
Agia Kyriaki Beach, sand and fine pebbles, no sun beds

Conference Center Map

Oral presentations will be held in the main auditorium of Building A (map below) at the Eliopoulos Conference Center. Posters will be hung on stands placed on the terrace, or near the adjacent Building B (not shown), depending on the wind. Coffee breaks and lunches will be on the terrace or near Building B, again based on current conditions. Restrooms are in Building B and the lower level of Building A, while a first aid station is available in Building B, along with a room for extended internet use.

Please do not use the Terrace Entrance when we are in session to avoid disturbing the talks. Instead, please come in through either the Main or Side Entrances.

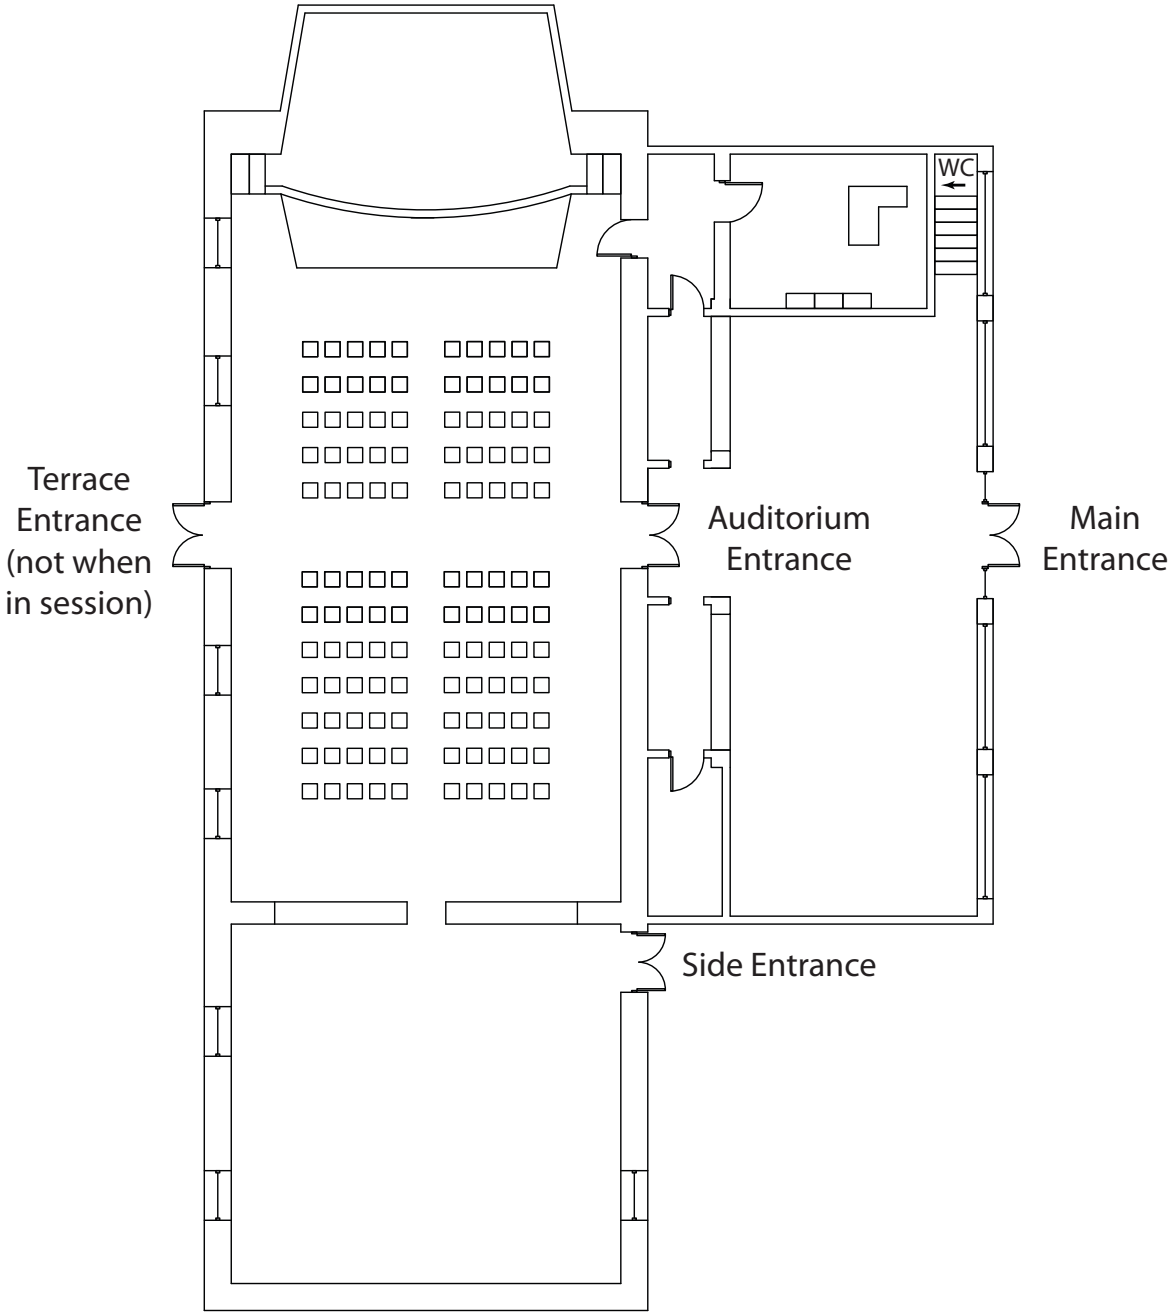

DAILY SCHEDULE AND PROGRAM

Overall Schedule

The schedule for the four-day conference follows the Greek lifestyle of having a long lunch, with the afternoon free for siestas or swimming, and a late dinner.

Tuesday
19:30–22:00 Welcome Reception and Registration

Wednesday
08:00–08:45 Registration
08:45–09:00 Opening Remarks
09:00–12:30 Lectures and Coffee Break
12:30–14:00 Lunch
17:00–21:30 Lectures and Coffee Break, Posters

Thursday
09:00–12:30 Lectures and Coffee Break
12:30–14:00 Lunch
17:00–21:30 Lectures and Coffee Break, Posters

Friday
09:00–13:00 Optional Excursions (no lunch provided)
17:00–21:30 Lectures and Coffee Break, Posters

Saturday
09:00–12:30 Lectures and Coffee Break
12:30–14:00 Lunch
17:00–19:00 Lectures and Coffee Break
19:00–19:15 Closing Remarks
20:30–24:00 Banquet Dinner at Kipos Cafe in Kipos Beach

TUESDAY, 25 JUNE 2024

19:30–22:00 welcome reception at Eliopoulos Conference Center

WEDNESDAY, 26 JUNE 2024

08:00–08:45 registration

08:45–09:00 opening remarks

MORNING SESSION Barry Richmond, moderator

09:00–09:45 **Elizabeth Buffalo** (University of Washington)
Neural dynamics of memory formation in the primate hippocampus, 29

09:45–10:30 **James Fitzgerald** (Northwestern University)
Learning produces a hippocampal cognitive map in the form of an orthogonalized state machine, 35

10:30–11:00 coffee break

11:00–11:45 **Terence Sanger** (University of California, Irvine)
Neural hashcodes and encrypted representations in human basal ganglia, 46

11:45–12:30 **Anne-Marie Oswald** (University of Chicago)
Synaptically defined cortical assemblies correspond to rewarded stimuli following olfactory discrimination behavior, 41

12:30–14:00 lunch

AFTERNOON SESSION Leslie Osborne, moderator

17:00–17:45 **Alexander Ecker** (University of Goettingen)
Most discriminative stimuli for functional cell type identification, 32

17:45–18:15 coffee and light snacks

18:15–19:00 **Melissa Warden** (Cornell University)
Neuromodulation and the balance between goal-directed and reactive behavior, 50

19:00–19:20 **Desmond Patterson** (University of Texas at Austin)
The geology of Milos: An introduction, 42

19:20–19:40 **Andronike Makres** (University of the Peloponnese)
Might makes right: The Melian Dialogue (416 BCE), 38

20:00–21:30 posters (*i*), presenting author

Christina Brozi (University of Crete)
Identifying neuronal computational and communication modules of the functional connectivity architecture in area V1, 59

Alexander Ecker (University of Goettingen)
Most discriminative stimuli for functional cell type identification

Jamal Esmaily (Ludwig Maximilian University of Munich)
Dissociation of decoding of subjective confidence and stimulus uncertainty in EEG and pupil signals in human, 66

Paul Fahey (Stanford University School of Medicine)
Functional connectomics reveals general wiring rule in mouse visual cortex, 67

Valeria Fascianelli (Columbia University)
Neural signatures of stress susceptibility and resilience in the amygdala-hippocampal network, 68

Maureen Hagan (Monash University)
Visual attention increases spike-LFP coherence across cortical layers in the posterior parietal cortex, 75

Dean Halperin (Weizmann Institute)
Modulation of functional connectivity and representations of value in primate single-neurons using a closed-loop BCI framework, 76

Samantha Johnson (University of Chicago)
Differences in neural ensembles between imagined and native limb movement raise new challenges for brain-computer interfaces, 78

Sze Chai Kwok (Duke Kunshan University)
Cortical replay of temporal memory triggered by inter- and intra- hippocampal ripples, 82

Bram Nuttin (Neuro-Electronics Research Flanders)
Dendritic properties shape the encoding of behaviorally relevant stimuli in collicular wide-field neurons, 80

David O'Reilly (University of Leeds)
Quantifying the diverse contributions of hierarchical muscle interactions to motor function, 89

Ganna Palagina (Brigham and Women's Hospital)
Layer 1: The gate of the bistable visual perception, 92

THURSDAY, 27 JUNE 2024

MORNING SESSION Andreas Tolias, moderator

09:00–09:45 **Haim Sompolinsky** (Harvard University)
Neural representation of concepts in vision and language, 47

09:45–10:30 **Sophia Sanborn** (Science Corporation)
Symmetry and universality, 45

10:30–11:00 coffee break

11:00–11:45 **Tatiana Engel** (Princeton University)
Functional cell types and the linear dimension of neural manifolds, 33

11:45–12:30 **Cengiz Pehlevan** (Harvard University)
Partial observation can induce mechanistic mismatches in data-constrained models of neural dynamics, 43

12:30–14:00 lunch

AFTERNOON SESSION Kenneth Blum, moderator

17:00–17:45 **Jiannis Taxidis** (The Hospital for Sick Children / Univ. of Toronto)
The role of inhibition in shaping memory-encoding hippocampal spiking sequences, 48

17:45–18:15 coffee and light snacks

18:15–19:00 **Lee Miller** (Northwestern University)
In pursuit of a universal, biomimetic iBCI decoder: Exploring the manifold representations of action in the motor cortex, 39

19:00–19:45 **Scott Imbrie** (University of Chicago)
My life as a human cyborg, 36

20:00–21:30 posters (*ii*), presenting author

Amirmasoud Ahmadi (MPI biological Intelligence)
Neural decoding of temporal features of zebra finches song, 55

Dani Bassett (University of Pennsylvania)
Event graph structure determines fidelity of neural representations, 58

R. James Cotton (Shirley Ryan AbilityLab / Northwestern University)
Development of a control system for a next generation high-density, wireless, bidirectional brain-computer interface, 61

FRIDAY, 28 JUNE 2024

09:00–13:00 optional excursion (no lunch provided)

AFTERNOON SESSION Nicholas Hatsopoulos, moderator

17:00–17:45 **Edward Chang** (University of California, San Francisco)
Human precentral gyrus neurons encode and transform speech from perception to production, 30

17:45–18:15 coffee and light snacks

18:15–19:00 **Ziv Williams** (Harvard Medical School)
Studying human speech at cellular scale through ultrahigh resolution recordings, 51

19:00–19:45 **Michale Fee** (MIT)
Neural clock underlying temporal structure of an auditory memory, 34

20:00–21:30 posters (*iii*), presenting author

Edward Bader (Albert Einstein College of Medicine)
The rostral zona incerta: An integrative hub in the circuitry of reward processing?, 57

Federico Benitez (University of Bern)
Order from chaos: Interplay of development and learning in recurrent networks of structured neurons, 79

Ryan Canfield (University of Washington)
Data-driven sampling of motor cortical networks during reaching, 60

Rares Dorcioman (University of Amsterdam)
Local field potential simulation across a V1 cortical model, 64

Becket Ebitz (Université de Montréal)
Pupil size predicts critical transitions in prefrontal neuronal population activity at the onset of exploration, 103

Emmanouil Giannakakis (MPI Biological Cybernetics)
Overlapping E/I neuronal assemblies generate rich network dynamics and enable complex computations, 73

Dominic Gonschorek (University of Tuebingen)
Nitric oxide mediates differential effects in mouse retinal ganglion cells, 74

Benjamin Grannan (University of Washington)
One shot learning in the human brain, 72

SATURDAY, 29 JUNE 2024

MORNING SESSION John Pezaris, moderator

09:00–09:45 **Martin Usrey** (University of California, Davis)
Feedforward and feedback interactions between thalamus and cortex for vision, 49

09:45–10:30 **Robert Desimone** (MIT)
Prefrontal feedback to visual cortex for feature attention, 31

10:30–11:00 coffee break

11:00–11:45 **Carlos Ponce** (Harvard Medical School)
Linking population activity to single neurons via generative networks, 44

11:45–12:30 **Tirin Moore** (Stanford University/HHMI)
Intermittent rate coding and cue-specific neuronal ensembles support working memory, 40

12:30–14:00 lunch

AFTERNOON SESSION Nicholas Hatsopoulos, moderator

17:00–17:45 **Gilles Laurent** (Max Planck Institute for Brain Research)
Of sleep and CPGs — or sleep as nested dynamics, 37

17:45–18:15 coffee and light snacks

18:15–19:00 **Anthony Zador** (Cold Spring Harbor Laboratory)
Brain wiring through the genomic bottleneck, 52

19:00–19:15 closing remarks

20:30–24:00 banquet dinner at Kipos Cafe in Kipos Beach

INVITED SPEAKER ABSTRACTS
(in alphabetical order by speaker)

NEURAL DYNAMICS OF MEMORY FORMATION IN THE PRIMATE HIPPOCAMPUS

Elizabeth A. Buffalo

Department of Physiology and Biophysics, University of Washington Seattle, Wash., USA
Washington National Primate Research Center, Seattle, Wash., USA
ebuffalo@uw.edu

Our understanding of the hippocampus has been framed by two landmark discoveries: the discovery by Scoville and Milner that hippocampal damage causes profound and persistent amnesia and the discovery by O'Keefe and Dostrovsky of hippocampal place cells in rodents. However, it has been unclear to what extent spatial representations are present in the primate brain and how to reconcile spatial representations with the known mnemonic function of this region. A concept bridging spatial representation and memory that may be a powerful principle for the primate hippocampus is that of the cognitive map, a form of internal model which allows one to organize knowledge gained from experience, assimilate new information and plan. I will discuss a series of new experiments that have examined neural activity in the hippocampus in monkeys performing behavioral tasks including foraging and spatial memory tasks in a virtual environment. Data from these experiments demonstrate that behavioral task structure has a significant influence on hippocampal activity, potentially providing a neural instantiation of a cognitive map that extends to non-spatial domains and serves as an important scaffold for memory formation.

Acknowledgments

The studies were performed by John Rueckemann, Yoni Browning, and Sofia Landi in the Buffalo lab at UW in collaboration with our NIH U19 team of PIs and trainees. The work was supported by the NIH BRAIN Initiative and the Simons Collaboration on the Global Brain.

HUMAN PRECENTRAL GYRUS NEURONS ENCODE AND TRANSFORM SPEECH FROM PERCEPTION TO PRODUCTION

Duo Xu, Alex Silva, Jason Chung, Matthew Leonard, Quinn Grecious, Edward Chang *

Weill Institute for Neurosciences, University of California, San Francisco, Calif., USA
* edward.chang@ucsf.edu

Perception and production are deeply interconnected aspects of human speech. They often form cycles of processing where speech information is extracted from auditory input, maintained in working memory, transformed to motor output, then the self-voice is perceived as feedback. The mechanisms that enable us to link a sequence of speech elements from perception to production are particularly mysterious and computationally challenging.

In this study, we used Neuropixels recordings from the precentral gyrus (preCG) in participants who performed a delayed sentence repetition task. We found single preCG neurons that were auditory (responding selectively to speech sounds), preparation, and speech production (speaking). Single cortical columns in preCG contained neuronal responses that tile the cycle of perception, working memory, speech initiation, and production.

We discovered *mirror* neurons in preCG that are transiently activated after perceiving specific speech content during listening and before producing the same content during speaking. More interestingly, another group of *bridge* neurons maintained elevated firing between the same speech content in perception and production, bridging over the working memory and/or speech initiation phases.

Clinically, the preCG is considered the primary motor cortex, however our results are more consistent with premotor functions. We found little evidence in the direct encoding of articulatory movements. Instead, neurons are tuned to and modulated by specific speech sequences.

PREFRONTAL FEEDBACK TO VISUAL CORTEX FOR FEATURE ATTENTION

Robert Desimone

McGovern Institute, MIT, Cambridge, Mass., USA
desimone@mit.edu

To find an object in a complex scene, we use feature-based attention to guide our search, typically in conjunction with spatial attention and targeting eye movements. When searching for our keys on a table, for example, the features of the keys are used as an attentional template that guides the eyes to the various objects sharing features with keys until the keys are found. Work from our own and other labs has found that objects with attended features or attended locations are processed more efficiently in the visual cortex, while the processing of unattended, distracting objects is suppressed.

The interconnected structures important for the control of attention have many common features. At the surface level, these common features suggest there may be little difference in their functions. However, our results show surprising specificity instead. We find that the ventral pre-arcuate area (VPA) and the frontal eye fields (FEF) in the prefrontal cortex (PFC) have different functions in visual search. Specifically, VPA appears to mediate the selection of likely targets based on their features, and FEF directs spatial attention and gaze to those possible targets until the object of the search, the target, is found [1]. We hypothesize that VPA and FEF work together as an interconnected system for guiding gaze to objects we are searching for, and that they also provide attentional feedback to the occipital and temporal cortex. This specificity of PFC feedback is also supported by our studies using es-fMRI, which reveal topographic connections between PFC and posterior cortex [2].

Although there is strong evidence that VPA provides feedback to visual cortex important for feature attention, it has been less clear if this feedback is important for the working memory of the target features. In a recent study, we optogenetically inactivated the VPA region while monkeys performed a feature attention task [3]. While bilateral inactivation of VPA during attentional selection impaired performance, inactivation during a working memory delay had no behavioral effect. Likewise, inactivation during attentional selection impaired attentional selection by cells in MST and LIP but inactivation during the working memory delay had a much smaller effect. Thus, VPA plays an important role in feature attention but other structures, possibly in other parts of PFC or posterior parietal cortex, appear to play a more important role in working memory for features.

Acknowledgments

This work has been supported by NIH EY029666.

References

1. Bichot, *et al.*, 2019, *Nature Comm*, 10(1):5727
2. Xu, *et al.*, 2022, *Neuron*, 110(2):312–327 e7
3. Mendoza-Halliday, *et al.*, 2024, *Neuron*, 112(5):850–863 e6

MOST DISCRIMINATIVE STIMULI FOR FUNCTIONAL CELL TYPE IDENTIFICATION

Max F. Burg[1–3], Thomas Zenkel[4,5], Michaela Vystrèilová[2], Jonathan Oesterle[4,5], Larissa Höfling[4,5], Konstantin F. Willeke[1,2,6], Jan Lause[3,7], Sarah Müller[1,3,7], Paul G. Fahey[8,9], Zhiwei Ding[8,9], Kelli Restivo[8,9], Shashwat Sridhar[10,11], Tim Gollisch[10–12], Philipp Berens[3,7], Andreas S. Tolias[8,9,13], Thomas Euler[4,5], Matthias Bethge[3,5], Alexander S. Ecker[2,14,]*

[1] Intl. Max Planck Research School for Intelligent Systems, Tübingen, Germany
[2] Institute of Comp. Sci. and Campus Inst. Data Science, Univ. of Göttingen, Germany
[3] Tübingen AI Center, University of Tübingen, Germany
[4] Institute of Ophthalmic Research, University of Tübingen, Germany
[5] Centre for Integrative Neuroscience, University of Tübingen, Germany
[6] Institute for Bioinformatics and Medical Informatics, Tübingen University, Germany
[7] Hertie Institute for AI in Brain Health, University of Tübingen, Germany
[8] Department of Neuroscience, Baylor College of Medicine, Houston, Texas, USA
[9] Center for Neuroscience and AI, Baylor College of Medicine, Houston, Texas, USA
[10] Department of Ophthalmology, University Medical Center Göttingen, Germany
[11] Bernstein Center for Computational Neuroscience, Göttingen, Germany
[12] Cluster of Excellence “Multiscale Bioimaging”, University of Göttingen, Germany
[13] Department of Electrical and Computer Engineering, Rice Univ., Houston, Texas, USA
[14] Max Planck Institute for Dynamics and Self-Organization, Göttingen, Germany
[*] ecker@cs.uni-goettingen.de

Identifying cell types and understanding their functional properties is crucial for unraveling the mechanisms underlying perception and cognition. For example in the retina, functional types can be identified by a carefully selected and manually curated battery of stimuli. However, this requires expert domain knowledge and biases the procedure towards previously known cell types. In the visual cortex, it is still unknown what functional types exist and how to identify them. Thus, for unbiased identification of the functional cell types in retina and visual cortex, new approaches are needed. Here we propose an optimization-based clustering approach using deep predictive models to obtain functional clusters of neurons using Maximally Discriminative Stimuli (MDS). Our approach alternates between stimulus optimization with cluster reassignment akin to an expectation-maximization algorithm. The algorithm recovers functional clusters in mouse retina, marmoset retina and macaque visual area V4. This demonstrates that our approach can successfully find discriminative stimuli across species, stages of the visual system and recording techniques. Presenting maximally discriminative stimuli during data acquisition allows for on-the-fly assignment to functional cell types, and paves the way for experiments that were previously limited by experimental time. Crucially, MDS are interpretable: they visualize the distinctive stimulus patterns that most unambiguously identify a specific type of neuron.

Acknowledgments

Funding by ERC Grant numbers 101041669 & 101039115, DFG project ids 432680300 (SFB 1456), 276693517 (SFB 1233), 390727645 (EXC 2064/1), IARPA via DoI/IBC contract number D16PC00003NSF NeuroNex grant 1707400, NIH award U19MH114830, NEI award numbers R01 EY026927 & T32-EY-002520–37.

FUNCTIONAL CELL TYPES AND THE LINEAR DIMENSION OF NEURAL MANIFOLDS

*Christopher Langdon, Tatiana A. Engel**

Princeton Neuroscience Institute, Princeton University, Princeton, New Jersey, USA
*tatiana.engel@princeton.edu

Single neurons show complex responses during cognitive tasks, and dimensionality reduction has been a successful approach for finding structure in heterogeneous neural activity. When applied to neural response patterns in a population state space, in which each axis represents activity of one neuron, dimensionality reduction methods often reveal low-dimensional manifolds encoding behaviorally relevant variables. Alternatively, we can view the same responses in a selectivity space, in which each axis represents a response feature (*e.g.*, task condition) and each point is a neuron, to reveal the distribution of tuning properties across the population. Traditionally, these two views on neural activity were linked through the dynamics of neural circuits, in which recurrently connected subpopulations generate low-dimensional response manifolds. The subpopulations are groups of neurons with similar response properties, which therefore form clusters in the selectivity space, called functional cell types. However, recordings form prefrontal cortex indicate that although population activity lies on a low-dimensional manifold, individual neurons respond to seemingly random mixtures of task variables so that boundaries defining functional types are not apparent [1]. Yet, data from orbitofrontal cortex consistently show functional clustering of neurons and indicate that revealing functional types requires the correct choice of features [2]. Theoretical studies, on the other hand, proposed that functional types emerge only in tasks requiring cognitive flexibility [3], but this idea conflicts with reports of random mixed selectivity of cortical neurons in these tasks [1]. Thus, it is unclear when functional types emerge in neural circuits and whether random mixed selectivity is compatible with low-dimensional manifolds.

Using a firing-rate recurrent network model, we mathematically prove that functional types emerge in neural circuits exactly when neural manifolds are constrained to lie in low-dimensional linear subspaces. Moreover, the number of functional types is approximately equal to the linear dimension of the underlying neural manifold, with equality arising in the limit of strict low dimensionality. We confirm these predictions in both recurrent neural networks trained to perform various cognitive tasks and in brain-wide neural recordings from mice during a decision-making behavior. Our findings explain the emergence of functional types in neural circuits and show that low-dimensional neural dynamics can be understood in terms of single cell responses.

Acknowledgments

Supported by the NIH grant RF1DA055666.

References

1. Mante, *et al.*, 2013, *Nature*, 503:78–84
2. Hirokawa, *et al.*, 2019, *Nature*, 576:446–451
3. Dubreuil, *et al.*, 2022, *Nat Neurosci*, 25:783–794

NEURAL CLOCK UNDERLYING TEMPORAL STRUCTURE OF AN AUDITORY MEMORY

Michale S. Fee

Department of Brain and Cognitive Sciences, MIT, Cambridge, Mass., USA
fee@mit.edu

Songbirds, such as zebra finches, store an auditory memory of their tutor's song and, over the course of several months during learning, gradually refine their own vocalizations to produce a mature song closely matching the tutor song. A young zebra finch can learn to imitate a song it hears only a few times, even though their songs are temporally complex and are extremely precisely reproduced. How is such a stable auditory memory of the tutor song formed? Song is controlled by a sparse neuronal sequence in the pre-motor nucleus HVC. Observations that local cooling of HVC slows the song, while cooling the downstream nucleus RA has no effect, suggest that timing is controlled by local dynamics within the HVC. However, recent work also shows that disruption of HVC or its auditory inputs and outputs during tutoring impairs vocal learning, suggesting a role for HVC in tutor memory formation and recall. Here we propose a specific model by which neural dynamics in HVC form during tutoring and act as a precise neural clock, not just for singing, but also for laying down the auditory tutor memory and synchronizing the recall of the tutor memory during singing, allowing the temporally precise readout of vocal errors. Motivated by this hypothesis, we tested a critical prediction of our theory: that cooling of HVC during tutoring leads to a sped-up tutor song memory, analogous to slowing the motor of a tape recorder during recording and then playing the recording back at normal speed. We designed a modular thermoelectric cooling device and found that transient cooling during tutoring caused birds to produce faster imitations, consistent with our theory. This work gives insight into formation of tutor song memories, and how they can be stably recalled months later to guide precise sensory-motor learning.

LEARNING PRODUCES A HIPPOCAMPAL COGNITIVE MAP IN THE FORM OF AN ORTHOGONALIZED STATE MACHINE

Weinan Sun[1,†], *Johan Winnubst*[1,†], *Maanasa Natrajan*[1–3], *Chongxi Lai*[1], *Koichiro Kajikawa*[1], *Michalis Michaelos*[1], *Rachel Gattoni*[1], *Carsen Stringer*[1], *Daniel Flickinger*[1], *James E. Fitzgerald*[1,2,*], *Nelson Spruston*[1]

[1] Janelia Research Campus, Howard Hughes Medical Institute, Ashburn, Virginia, USA
[2] Department of Neurobiology, Northwestern University, Evanston, Illinois, USA
[3] Department of Neuroscience, Johns Hopkins University, Baltimore, Maryland, USA
[†] equal contribution
[*] james.fitzgerald@northwestern.edu

Cognitive maps confer animals with flexible intelligence by representing spatial, temporal, and abstract relationships that can be used to shape thought, planning, and behavior. Cognitive maps have been observed in the hippocampus, but their algorithmic form and the processes by which they are learned remain obscure.

Here, we employed large-scale, longitudinal two-photon calcium imaging to record activity from thousands of neurons in the CA1 region of the hippocampus while mice learned to efficiently collect rewards from two subtly different versions of linear tracks in virtual reality. The results provide a detailed view of the formation of a cognitive map in the hippocampus. Throughout learning, both the animal behavior and hippocampal neural activity progressed through multiple intermediate stages, gradually revealing improved task representation that mirrored improved behavioral efficiency. The learning process led to progressive decorrelations in initially similar hippocampal neural activity within and across tracks, ultimately resulting in orthogonalized representations resembling a state machine capturing the inherent structure of the task. We further demonstrate that mice exhibited adaptive behavior in novel task settings, with neural activity reflecting flexible deployment of the state machine.

We then showed that a Hidden Markov Model (HMM) and a biologically plausible recurrent neural network trained using Hebbian learning can both capture the orthogonalized representational structure in neural activity and certain aspects of the learning dynamics. However, only the HMM trained through the Baum-Welch (BW) algorithm could capture the core sequence of representational changes that occur during *de novo* cognitive map formation. These computational properties are non-generic, as we found that other machine-learning-inspired sequence models, such as Long Short-Term Memory networks (LSTMs) and Transformers, do not naturally produce such orthogonalized representations.

These findings shed light on the mathematical form of cognitive maps, the learning rules that sculpt them, and the algorithms that promote adaptive behavior in animals. The work thus charts a course toward a deeper understanding of biological intelligence and offers insights toward developing more robust learning algorithms in artificial intelligence.

Acknowledgments

Supported by HHMI.

MY LIFE AS A HUMAN CYBORG

Scott Imbrie[1], *Lee Miller*[2], *Nicholas G. Hatsopoulos*[1,*]

[1] University of Chicago, Chicago, Illinois, USA
[2] Northwestern University, Chicago, Illinois, USA
[*] nicho@uchicago.edu

Over three years ago, Scott Imbrie became the first subject in the Chicago area to be implanted with electrodes in the motor and somatosensory cortices as part of our Cortical Bionics Research Group (CBRG), a consortium which includes teams from the University of Chicago, the University of Pittsburgh, and Northwestern University. The goal of the CBRG research program is to further develop a bi-directional brain computer interface for dexterous object manipulation, using control signals from motor cortex to produce movement, in combination with electrical stimulation of somatosensory cortex to provide artificial tactile feedback. In this presentation, Scott will discuss what led him to participate in our study, what he likes and dislikes about it, and what his prosthetic sense of touch feels like. Scott will discuss these topics in an interview format with Lee Miller and Nicho Hatsopoulos, prior to taking questions from the audience.

OF SLEEP AND CPGs — OR SLEEP AS NESTED DYNAMICS

Gilles Laurent[*], *Lorenz Fenk, Luis Riquelme, Emma Morris*

Max Planck Institute for Brain Research, Frankfurt, Germany
[*] gilles.laurent@brain.mpg.de

Behavioral sleep seems to be ubiquitous among animals; yet its electrophysiological attributes vary a lot between species. Mammalian sleep is typically split into two states, called slow wave (SW) and Rapid Eye Movement (REM). REM is also known as active or paradoxical sleep (because of the similarity of its EEG signature to that of awake states). These two states (SW and REM) typically alternate during sleep with a frequency, regularity and duty cycle that vary from species to species.

While the existence of SW and REM states was accepted early for most placental mammals, sleep in non-mammalian vertebrates (fish, amphibians, reptiles including birds) was considered to be analogous to mammalian NREM sleep, and to lack REM. Some hints of the existence of REM in birds and reptiles existed in the 1960s already but REM became truly accepted in avians (but not in non-avian reptiles) only in the early 2000s. Because birds branched off the non-avian reptiles (via the dinosaurs), because REM was not clearly established in non-avian reptiles and because birds and mammals are both homeotherms (while non-avian reptiles are not), it was proposed that REM arose by convergence in both birds and mammals as a result of high cognitive demands linked to thermoregulation.

In 2016, we provided evidence that a non-avian reptile, the Australian dragon, produces two distinctive sleep states; one of them was associated with rapid-eye movements and awake-like electrophysiological activity [1]. This and associated evidence suggested that REM may in fact be an ancestral feature of sleep that existed already in the common ancestor of all amniotes (*i.e.*, mammals and reptiles, including birds), a vertebrate that colonized land some 320 million years ago. If so, a model system that expresses the primitive features of REM and SW sleep could be extremely useful to understand both the evolution of sleep but also its most basic mechanisms and functions.

The dragon model proved to be very useful to study the physiology and mechanisms of sleep because its features are extremely robust, and its rhythm is both fast and regular. In this talk, I will present new results that suggest that the alternation of the two sleep states (SW and REM) is under the control of a pair of central pattern generators, themselves under circadian chemical control. These results, added to our earlier ones on bilateral competition during REM and the roles of the hindbrain's isthmus and the forebrain's claustrum for sleep control, reveal a complex but remarkably structured picture of sleep dynamics over 6 orders of magnitude of time.

Acknowledgments

Supported by the Max Planck Society and the European Research Council

Reference

1. Shein-Idelson, et al., 2016, *Proc AREADNE 2016*

MIGHT MAKES RIGHT: THE MELIAN DIALOGUE (416 BCE)

Nike Makres

Hellenic Education Research Center, Athens, Greece
info@herc.org.gr

One of the most breathtaking moments in the history of Ancient Athens is when, in the 5th century BCE, Athens became a superpower in the Aegean and had to manage her relations with her allies. This proved to be a very difficult and controversial task as Athens was democratic in the running of her internal affairs but autocratic in her relations with other cities. The best example of Athenian autocratic attitude when dealing with others is the famous Melian Dialogue, the party in that instance confronting Athens being the island of Melos. This dialogue is preserved fully in the work of Thucydides. In my talk I shall present and explain this very dramatic historical moment that still poses a crucial question: Can a Democracy rule an Empire? It also teaches a big lesson for today: Might makes Right.

IN PURSUIT OF A UNIVERSAL, BIOMIMETIC iBCI DECODER: EXPLORING THE MANIFOLD REPRESENTATIONS OF ACTION IN THE MOTOR CORTEX

Xuan Ma[1], *Fabio Rizzoglio*[1], *Ege Altan*[2], *Lee E. Miller*[1–4,*]

[1] Dept. Neuroscience, Northwestern University, Chicago, Illinois, USA
[2] Dept. Biomedical Engineering, Northwestern University, Evanston, Illinois, USA
[3] Dept. Physical Medicine and Rehabilitation, Northwestern University, Chicago, Illinois, USA
[4] Shirley Ryan AbilityLab, Chicago, Illinois, USA
[*] lm@northwestern.edu

My group pioneered the development of a novel intracortical brain computer interface (iBCI) that decodes muscle activity (EMG) from signals recorded in the motor cortex (M1) of monkeys [1]. We have used these predicted EMG signals to control Functional Electrical Stimulation (FES), which causes the muscles to contract and thereby restore rudimentary voluntary control of the monkey's temporarily paralyzed hand during simple grasp tasks. Our goal is to extend this approach to more complex motor actions, and ultimately to translate the methods to humans.

The ability to record simultaneously from many neurons which made iBCIs possible, also led to the development of the concept of a neural state space and a low-dimensional manifold [2], and has led to new insights about the neural control of movement. However, the nonlinearities and context dependence of decoding within this manifold present significant challenges. We built a large plastic cage that allows us to record wirelessly from both M1 and muscles of the arm and hand of completely unconstrained monkeys. We have examined the representation in M1, of a broad range of behaviors, including both well-learned, stereotyped movements in the lab, and the more natural grasping movements related to food retrieval and locomotion in the monkey's home cage. We can also study more freeform activities like grooming and foraging.

These different motor actions occupy different regions of neural state space. The cage behaviors tend to have dimensionality about 50% larger than do the lab behaviors. They also tend to be somewhat more nonlinear. As a consequence, linear decoders of EMG trained on one behavior type generalize very poorly to other behaviors. A linear decoder trained on all behaviors is more accurate, but significantly less so than an RNN trained on the same data. However, in an effort to better understand the representation of these diverse behaviors within the neural manifold, we also computed an unsupervised, piecewise-linear decoder. This decoding approach initially finds clusters within either a linear or nonlinear M1 manifold, then computes a linear, Wiener filter decoder for each cluster. Subsequent data points are classified, then decoded using the appropriate decoder. Surprisingly, this piecewise approach outperformed even the RNN decoder. Our most recent efforts have been devoted to understanding the shape of the overall manifold, its stability over time, and the relation between neighboring clusters.

References

1. Ethier, Oby, Bauman, Miller, 2010, *Nature*, 485(7398):368–371, 10.1038/nature10987
2. Gallego, Perich, Miller, Solla, 2017, *Neuron*, 94(5):978–984, 10.1016/j.neuron.2017.05.025

INTERMITTENT RATE CODING AND CUE-SPECIFIC NEURONAL ENSEMBLES SUPPORT WORKING MEMORY

Matthew F. Panichello, Donatas Jonikaitis, Yu Jin Oh, Shude Zhu, Ethan B. Trepka, Tirin Moore *

Dept. Neurobiol. and Howard Hughes Medical Inst., Stanford University, Stanford, Calif., USA
*tirin@stanford.edu

Persistent, memorandum-specific neuronal spiking activity has long been hypothesized to underlie working memory. However, emerging evidence suggests a possible role for *activity-silent* synaptic mechanisms. This issue remains controversial because evidence for either view has largely relied on datasets that fail to capture single-trial population dynamics or on indirect measures of neuronal spiking.

We recently addressed this by examining the dynamics of mnemonic information on single trials obtained from large, local populations of prefrontal neurons recorded simultaneously in monkeys performing a working memory task. We found evidence that mnemonic information does not persist in the spiking activity of neuronal populations during memory delays, but instead alternates between coordinated *On* and *Off* states. At the level of single neurons, Off states are driven both by a loss of selectivity for memoranda and a return of firing rates to baseline.

Further exploiting the large-scale recordings, we show that mnemonic information is available in the patterns of functional connections among neuronal ensembles throughout the memory delay, even during Off states. Our results suggest that intermittent periods of memoranda-specific spiking coexist with synaptic mechanisms to support working memory.

Acknowledgments

Supported by NIH Grants EY014924 and NS116623.

SYNAPTICALLY DEFINED CORTICAL ASSEMBLIES CORRESPOND TO REWARDED STIMULI FOLLOWING OLFACTORY DISCRIMINATION BEHAVIOR

Anne-Marie M. Oswald

Department of Neurobiology, University of Chicago, Chicago, Illinois, USA
amoswald@uchicago.edu

Throughout the brain, sensory stimuli, decisions, and motor plans can be decoded from the joint activity of populations of neurons. In olfactory cortex, odor stimuli are similarly represented by the activity of unique ensembles neurons. However, recent studies have suggested that additional information related to reward and context may also be represented in cortical activity. The circuitry of the olfactory cortex is akin to the hippocampus and, as such, has the potential to support both sensory and associative roles in olfactory processing. We investigated the stabilization of olfactory sensory ensemble representations through synaptic plasticity to form neural assemblies. Consistent with an associative role, we find that assembly formation depends on both reward and the features of the olfactory task.

We trained mice to perform a foraging task that involved digging in odorized pots for food rewards. Mice learned to discriminate two overlapping odor mixtures that differed by a single component. Within 1–2 days of training, mice were able to reliably dig in the pot scented with rewarded mixture and refrain from digging in the unrewarded odorized pot. Once mice were proficient at the task, we used targeted recombination in active populations (TRAP) [1] to conditionally express channelrhodopsin and/or tdTomato in ensembles that are activated by either the rewarded or unrewarded odor. We then recorded synaptic strength between neurons within each ensemble *in vitro*. We found that assemblies of strongly interconnected neurons formed in response to the rewarded odor mixture but not the unrewarded mixture. Moreover, ensemble size for the rewarded mixture is greater than for the unrewarded mixture. In further behavioral studies, we found that mice can use the differential information corresponding to a single component to perform the task. Altogether, these findings suggest that assembly formation corresponds to task-relevant information while redundant information is suppressed. Our current experimental and computational studies address the underlying circuit mechanisms that mediate selective assembly formation during learning as well as long-term stabilization following learning.

Acknowledgments

Supported by NIH BRAIN Initiative RF1NSI133598.

Reference

1. Guenthner, Miyamichi, Yang, *et al.*, 2013, *Neuron*, 78:773–784, 10.1016/j.neuron.2013.03.025

THE GEOLOGY OF MILOS: AN INTRODUCTION

Des Patterson

University of Austin at Texas, Austin, Texas, USA
des@helium-labs.com

The Milos Archipelago is a complex of small predominantly volcanic islands within the Southern Aegean Volcanic Arc which includes (west to east) Crommyonia, Methana, Aegina, Milos, Santorini, Nisyros, and Kos. Volcanism is related to the ongoing subduction of the African tectonic plate beneath the Aegean. This presentation will focus on Milos, although the other three islands in the local cluster, Antimilos, Kimolos, and Polyegos, share a geologic history that is broadly the same.

Milos covers an area of 151 km^2 and is predominantly, but not exclusively, of volcanic origin (composed of andesite to dacite rock types). Pre-volcanic rocks include restricted outcrops of Mesozoic (252–66 Ma) schists along the southeast coast, which are overlain by Neogene (23–2.6 Ma) marine sediments which provide the sub-volcanic platform for the island and are found to a limited extent at the surface in the south-central area. Milos is tectonically active with numerous faults forming horst and graben structures (areas of high and low relief) with vertical offsets that can exceed several hundred meters. The most obvious of these is the central Gulf of Milos which is not a volcanic crater nor a caldera, but is a large bay formed by vertical movement of these faults.

Volcanic activity began approximately 3.5 Ma ago and has continued until historical times. The eruptive history can be divided into three phases reflecting a progression from purely submarine (3–2 Ma), through transitional, to dominantly subaerial (in open air, or *under the air*) (1.5 Ma to Recent) phases. Individual eruptions appear to have been comparatively small, with a maximum total thickness of 700 meters. There is no evidence for any large scale cataclysmic events. Eruptions have occurred from over 20 individual centers scattered across all parts of the island. Volcanic deposits are highly variable depending on interaction with seawater and proximity to the vent, and deposits from different vent locations are commonly interleaved. When combined with extensive post eruption hydrothermal alteration and significant erosion, the field geology of Milos is challenging and requires careful detailed mapping to untangle the history. As a generalization the oldest volcanics are in the western areas and along the eastern coast (3–2 Ma), with younger deposits in the north (Trachilas, 500–300 ka), and south-central regions (Fyriplaka, 70–60 ka). The youngest events are dated 200 BCE–200 CE by ^{1}4C ages of carbonized wood samples near Agia Kiriaki in the southeast. This area also hosts a large geothermal reservoir which feeds active surface fumaroles and warm springs.

Milos has a long history of mining. High quality obsidian from two locations (one top of hill behind Adamas Port) was traded throughout the neolithic Aegean (from circa 7000 BCE). Trachyte lava used for millstones was also mined and exported from Milos since about 400 BCE. Extensive hydrothermal alteration of the volcanic deposits has resulted in economically important deposits of various clays (kaolin, bentonite, perlite) as well as sulphur, barite and manganese (Cape Vani). Today Milos is a significant producer of industrial perlite and bentonite from open pit mines in the northwest near Pollonia.

The Milos Mining Museum has an excellent display which includes an extensive collection of obsidian stone tools. The museum also runs a short documentary film loop of historical interviews with local Mileans who worked the mines in the early 20th Century — an evocative glimpse of a lifestyle long past and is highly recommended.

PARTIAL OBSERVATION CAN INDUCE MECHANISTIC MISMATCHES IN DATA-CONSTRAINED MODELS OF NEURAL DYNAMICS

Cengiz Pehlevan

Center for Brain Science, Harvard University, Cambridge, Mass., USA
Kempner Inst. of Natural and Artificial Intelligence, Harvard University, Cambridge, Mass., USA
cpehlevan@seas.harvard.edu

One of the central goals of computational neuroscience is to understand how the dynamics of neural circuits give rise to their observed function. A popular approach towards this end is to train recurrent neural networks (RNNs) to reproduce experimental recordings of neural activity. These trained RNNs are then treated as surrogate models of biological neural circuits, whose properties can be dissected via dynamical systems analysis. While recent advances in population-level recording technologies have allowed simultaneous recording of up to tens of thousands of neurons, this represents only a tiny fraction of most cortical circuits. Here we show that partial observation can create mechanistic mismatches between a simulated teacher network and a data-constrained student, even when the two networks have otherwise matching architectures. In particular, we show that partial observation of models of low-dimensional cortical dynamics based on functionally feedforward or low-rank connectivity can lead to surrogate models with spurious attractor structure. In total, our results illustrate the challenges inherent in accurately uncovering neural mechanisms from single-trial data, and suggest the need for new methods of validating data-constrained models for neural dynamics.

LINKING POPULATION ACTIVITY TO SINGLE NEURONS VIA GENERATIVE NETWORKS

Alireza A. Dehaqani[1], *Binxu Wang*[2], *Carlos R. Ponce*[1,*]

[1] Department of Neurobiology, Harvard Medical School, Boston, Mass., USA
[2] Kempner Inst. for Natural and Artificial Intelligence, Harvard Univ., Cambridge, Mass., USA
[*] carlos@hms.harvard.edu

Every behavior, percept, and decision arises from the collective activity of neuronal populations. However, the fundamental level of information transfer comprises single spikes emitted by single neurons. How is population-level information distributed across individual neurons? Here, we explored the hypothesis that the population code is emergent (reflecting a complicated interaction of individual neurons) against the hypothesis that the population code is reducible to the parallel activity of its constituent neurons and microclusters.

We tested these hypotheses by measuring the function of single neurons and local microclusters (*multiunits*) in three visual cortical areas of the macaque monkey: primary visual cortex (V1), extrastriate cortex (V4), and inferotemporal cortex (n = 64–96 chronically implanted electrodes per monkey, two monkeys). We defined the function of each given neuron or microcluster by optimizing synthetic images, each image containing visual features that were highly activating. We achieved this optimization by using deep generative networks in conjunction with evolutionary search algorithms. Then, we localized the most activating and inhibiting regions of this synthetic image by (1) scrambling local image regions, and (2) grafting strongly activating regions onto natural images. We show that these interventions reduced and then increased the excitatory potential of random natural images. Further, this approach revealed when and why neurons responded to random images sampled across a 50,000 image set (ImageNet validation set).

After defining the local features that best activated each site in our population, we measured the activity of all units and multiunits across all chronic microelectrode arrays, estimating a population activity vector in response to target images randomly selected from the reference set of natural images. We adapted the generative-image approach to work at the population level, in order to approximate each target image population vector. We found that the reconstructed images reflected the individual similarity between local target-image patches and local optimized-image patches. Our results suggest that the broader visual population code was reducible to the discrete contributions of individual neurons and microclusters. Overall, we conclude that to understand perception and decision-making, one must understand neuronal population codes; yet to understand population codes, we must study the single neuron.

SYMMETRY AND UNIVERSALITY

Giovanni Luca Di Marchetti[1], *Christopher Hillar*[2], *Danica Kragic*[1], *Sophia Sanborn*[3,*]

[1] KTH Royal Institute of Technology, Stockholm, Sweden
[2] Redwood Center for Theoretical Neuroscience, University of California, Berkeley, Calif., USA
[3] Science Corporation, Alameda, California, USA
[*] sophia.sanborn@gmail.com

Artificial neural networks trained on natural data exhibit a striking phenomenon: regardless of exact initialization, dataset, or training objective, models trained on the same data domain frequently converge to similar learned representations. For example, the early layer weights of diverse image models tend to converge to Gabor filters and color-contrast detectors. Remarkably, many of these same features are observed in the visual cortex, suggesting a form of representational universality that transcends biological and artificial substrates. While such findings are empirically well-established, the field lacks theoretical explanations.

Spatially localized versions of 2D Fourier basis functions, such as Gabor filters or wavelets, are perhaps the most frequently observed universal features in image models. They commonly arise in the early layers of vision models as well as in the primary visual cortices of diverse mammals. Non-localized Fourier features have been observed in networks trained to solve tasks that permit cyclic wraparound. In the domain of spatial navigation, the grid cells of the entorhinal cortex display periodic firing patterns at different spatial frequencies as they build a map of space. Their response properties are naturally modeled with the harmonics of the twisted torus. Similar features also emerge in artificial neural networks trained to solve spatial navigation tasks. The ubiquity of these features across diverse learning systems is both striking and unexplained.

In this work, we provide a mathematical explanation for the emergence of Fourier features in learning systems such as neural networks. We argue that the mechanism responsible for this emergence is the downstream invariance of the learner to the action of a group of symmetries (*e.g.*, planar translation or rotation). Since natural data typically possess symmetries, invariance is a fundamental bias that is injected both implicitly and sometimes explicitly into learning systems. Motivated by this, we derive theoretical guarantees for the presence of Fourier features in invariant learners that apply to a broad class of machine learning models. Here, we provide a non-technical overview of the theoretical results presented in this work. Our main result can be summarized as follows.

Informal Theorem 1. If $\varphi(W, x)$ is a parametric function of a certain kind that is invariant in the input variable x to the action of a finite group G, then each component of its weights W coincides with a harmonic of G up to a linear transformation. In particular, when the weights are orthonormal, W coincides with the Fourier transform of G up to linear transformations.

Informal Theorem 2. If $\varphi(W, x)$ is *almost invariant* to G according to certain functional bounds and the weights are *almost orthonormal*, then the multiplicative table of G can be recovered from W.

Our results prove that Fourier features are guaranteed to emerge in a wide array of neural systems solving common learning problems and serve as a theoretical explanation for the ubiquity of Fourier features in biological and artificial neural networks, grounded in the mathematics of symmetry.

NEURAL HASHCODES AND ENCRYPTED REPRESENTATIONS IN HUMAN BASAL GANGLIA

Terence D. Sanger

Electrical Engineering and Computer Science, University of California, Irvine, Calif., USA
Movement Disorders and Neurology, Childrens Hospital Orange County, Orange, Calif., USA
terry@sangerlab.net

Linear readout methods including population codes and principal component methods have been remarkably successful for understanding representations in cortex, and analysis of linear population dynamics uncovers latent variables that change smoothly with time. We have applied similar methods to attempt to understand representations of movement in basal ganglia, motor thalamus, and movement-related brainstem regions, based on data obtained during evaluation of targets for deep brain stimulation in children with severe movement disorders. This evaluation includes the implantation of temporary electrodes capable of stimulating and recording at up to 160 sites in awake unrestrained children. Isolation of up to 250 distinct spike shapes provides the ability to look at the pattern of activity at high temporal resolution over multiple deep brain regions.

Unlike the results in cortex, we find that movement variables such as direction or speed of movement are not linearly encoded, the linear population vector is not predictive of direction of movement, and principal components and independent components do not find meaningful latent variables. Furthermore, the temporal dynamics of firing patterns do not appear to be smooth in time, so that the Euclidean distance between firing patterns close in time is not less than the distance between firing patterns far in time.

We suggest that the multi-electrode recording data are consistent with the hypothesis of *neural hashcodes*, in which distinct random patterns are assigned to discrete states of the system. This representation facilitates one-shot or few-shot learning and allows precise representation of multidimensional states, at the expense of poor generalization and lack of smoothness in time. Decryption recognizes that such codes function as a substitution cipher, in which fixed random patterns are substituted for the underlying smooth latent variables. While full decryption is not possible due to significant under-sampling (we identify 250 distinct spikes in a region with 300 million cells), it is nevertheless possible to identify latent variables by identifying patterns that occur closely spaced in time and mapping these patterns nonlinearly to smooth variables in a latent space. We show that such mappings permit a reduced-dimensionality latent space that is a low order predictor of the basal ganglia, thalamic, and brainstem representation.

The neural hashcode representation is expected to be unique to the individual subject, and thus the same decryption method finds differing representations across subjects. Nevertheless, our methods show that meaningful temporally smooth latent variables can be extracted from a hashcode representation and may be helpful for understanding the encoding of motor variables in deep brain regions.

Acknowledgments

Supported by the Crowley Carter Foundation and the UC Irvine department of Electrical Engineering.

NEURAL REPRESENTATION OF CONCEPTS IN VISION AND LANGUAGE

Haim Sompolinsky

Department of Molecular and Cellular Biology, Harvard University, Cambridge, Mass., USA
hsompolinsky@mcb.harvard.edu

The atomic units of cognitive functions in humans and primates are concepts such as objects, categories, rules, and, for humans, words. A key challenge for cognitive and brain sciences is understanding how the brain extracts these discrete entities from continuous analog streams of stimuli. Similar questions apply to deep networks in current artificial intelligence (AI) systems, including DNNs for vision and large language models (LLMs), which can serve as useful testing beds for theories of neural encoding of concepts in biological brains.

In this talk, I propose that concepts are encoded as neural manifolds with geometries that allow downstream networks to perform computations based on concept identities. I will present evidence for the existence of concept manifolds and discuss their properties in the contexts of visual object recognition and language processing in both biological neural and AI neural networks. By exploring the parallels between these systems, we can gain insights into the fundamental mechanisms underlying concept representation and manipulation in cognitive systems.

THE ROLE OF INHIBITION IN SHAPING MEMORY-ENCODING HIPPOCAMPAL SPIKING SEQUENCES

Jiannis Taxidis

SickKids Research Inst., Neurosci. and Mental Health Program, Univ. Toronto, Toronto, Canada
Dept. of Physiology, University of Toronto, Toronto, Canada
jiannis.taxidis@sickkids.ca

How does the brain keep track of events we need to remember as well as the intervals between them? This process involves the hippocampus. When a series of sensory cues is experienced by mice, hippocampal spiking sequences encode these cues and link them in memory by tiling the time gaps between them. At each timepoint, these sequences retain information on the identity of the most recent cue and the time elapsed since its presentation. They, therefore, form activity trajectories in *memory space*. But the role of inhibitory circuits in shaping these memory-encoding sequences remains unclear. I will present pioneering, longitudinal voltage imaging of cell-type-specific CA1 interneurons while mice perform a memory task. Combined with two-photon calcium imaging and electrophysiological data, these recordings demonstrate that CA1 interneurons increase the signal-to-noise ratio of hippocampal sequences during cue presentation but not during time intervals between cues. Therefore, inhibition is crucial for efficient memory encoding but less so for memory linking across time.

FEEDFORWARD AND FEEDBACK INTERACTIONS BETWEEN THALAMUS AND CORTEX FOR VISION

W. Martin Usrey

Center for Neuroscience, University of California, Davis, California, USA
wmusrey@ucdavis.edu

The thalamus and cerebral cortex are interconnected by a dense network of feedforward and feedback circuits. In the visual system, the response properties of neurons in the lateral geniculate nucleus (LGN) of the thalamus and primary visual cortex (V1) are governed by the anatomical organization of these connections and the temporal patterns of impulse arrival. Results will be presented from experiments using multielectrode recordings and optogenetic manipulation to examine the specificity of neuronal circuits and the role of spike timing and behavioral modulation in the reciprocal exchange of information between the LGN and V1 in the alert macaque monkey. These results reveal a striking relationship between the parallel feedforward and feedback processing streams and the biophysical properties that govern spike transfer and the encoding of visual information in neuronal spike trains.

NEUROMODULATION AND THE BALANCE BETWEEN GOAL-DIRECTED AND REACTIVE BEHAVIOR

Melissa R. Warden

Department of Neurobiology and Behavior, Cornell University, Ithaca, New York, USA
mrwarden@gmail.com

Striking an adaptive balance between persistently pursuing goals and reacting quickly to important environmental events is essential for survival. Here, I will discuss the role of neuromodulation in regulating this balance, and will describe our recent work using optical methods to monitor and control the activity of serotonin and dopamine neurons in freely behaving mice. First, I will present evidence that phasic activity in dorsal raphe serotonin neurons promotes fast, state-dependent behavioral/emotional reactions, and will discuss the implications of these findings for the therapeutic efficacy of drugs that target the serotonin system. Then, I will discuss ventral tegmental area dopamine neural activity during approach to goals and the use of an internal model of progress. Finally, I will discuss the role of inputs to the ventral tegmental area in continued engagement in goal-directed behavior.

STUDYING HUMAN SPEECH AT CELLULAR SCALE THROUGH ULTRAHIGH RESOLUTION RECORDINGS

Arjun Khanna[1,2], *William Munoz*[1,2], *Joon Kim*[3], *Ziv Williams*[1,2,4]

[1] Department of Neurosurgery, Massachusetts General Hospital, Boston, Massachusetts, USA
[2] Department of Neurosurgery, Harvard Medical School, Boston, Massachusetts, USA
[3] Harvard Medical School, Boston, Massachusetts, USA
[4] Program in Neuroscience, Harvard Medical School, Boston, Massachusetts, USA
[*] zwilliams@mgh.harvard.edu

Humans can produce a remarkably wide array of word sounds to convey specific meanings. To produce fluent speech, linguistic analyses suggest a structured succession of processes involved in planning the arrangement and structure of phonemes in individual words. These processes are thought to occur rapidly during natural speech and to recruit prefrontal regions in parts of the broader language network known to be involved in word planning and sentence construction. Understanding the basic cellular mechanisms underlying language production in humans, however, has remained a significant challenge. Here, we describe the use of recently developed ultrahigh-density Neuropixels arrays for acute intraoperative neuronal recordings [1, 2] in combination with dynamical systems modeling and decoding techniques for studying human language [3]. By tracking their collective activities, we find prefrontal neurons that closely mirrored the way in which the word sounds were produced during speech, meaning that they reflected how individual planned phonemes were generated through specific articulators. Moreover, rather than representing phonemes independently of their order or structure, many of the neurons coded for their composition in the upcoming words. They also reliably predicted the arrangement and segmentation of phonemes into distinct syllables and morphemes, suggesting a process that could allow the structure and order of articulatory events to be encoded at a cellular level. We describe how these cells are organized along the cortex and how they track the composition of consonant and vowel sounds during perception. We also describe the potential use of these recording techniques of studying other aspects of human cognition and communication such as theory of mind. Taken collectively, these initial studies begin to reveal some of the basic cellular computations by which humans produce speech and offer a prospective platform for studying human-unique cognitive processes at cellular scale.

Acknowledgments

Supported by the NIH and BRAIN initiative.

References

1. Paulk, Kfir, Khanna, *et al.*, 2022, *Nat Neurosci*, 25:252–63, 10.1038/s41593-021-00997-0
2. Coughlin, Munoz, Kfir, *et al.*, 2023, *Nat Protoc*, 18:2927–53, 10.1038/s41596-023-00871-2
3. Khanna, Munoz, Kim, *et al.*, 2024, *Nature*, 626:603–10, 10.1038/s41586-023-06982-w

BRAIN WIRING THROUGH THE GENOMIC BOTTLENECK

Anthony Zador

Cold Spring Harbor Laboratory, Cold Spring Harbor, New York, USA
prof.zador@gmail.com

Animals are born with highly structured brain connectivity, which equips them with the ability to function well soon after birth with little or no experience. Because the wiring diagram is far too complex to be specified explicitly in the genome, it must be compressed through a genomic bottleneck. I will discuss theoretical and experimental results regarding the nature of the rules governing the genomic bottleneck and how high-throughput molecular connectomics based on new sequencing technologies can help uncover these rules.

POSTER ABSTRACTS

(in alphabetical order by first author)

NEURAL DECODING OF TEMPORAL FEATURES OF ZEBRA FINCHES SONG

Amirmasoud Ahmadi[1,*], *Hermina Robotka*[1], *Frederic Theunissen*[1,2], *Manfred Gahr*[1]

[1] Max Planck Institute for Biological Intelligence, Seewiesen, Germany
[2] University of California, Berkeley, California, USA
[*] amirmasoud.ahmadi@bi.mpg.de

Auditory decoding of temporal features is a complex process in which vocalizations such as birdsongs are segmented into distinct auditory units, a crucial step in transforming sound into meaning. The songbird auditory system offers a valuable model for studying the neural representation of complex sounds. The hierarchical structure of the Zebra Finch song makes it ideal for in-depth analysis of auditory decoding.

We use stacked Bidirectional Long Short-Term Memory (BiLSTM) deep neural networks to decode the amplitude envelope and the time-locked envelope features of zebra finch songs. To assess the neural activity's efficacy at segmenting continuous songs into units and decoding amplitude, the network was trained with local field potential (LFP) and multi-unit activity envelope (MUAe).

In ensemble responses, both the amplitude envelope and time-locked features could be accurately decoded using LFP and MUA. The performance of LFP and MUAe was very similar, but MUAe gave slightly better results for envelope decoding. It was observed that temporal information might not have been present everywhere in the brain/auditory pallium, and this segmentation function could be modulated by other factors (such as attention). Notably, the envelopes of the introductory notes and the first motif were significantly better decoded than the second motif. This result suggested that these specific parts of songs were receiving more attention. Additionally, network accuracy and inter-trial phase coherence exhibited a positive linear relationship in LFP and MUA signals, indicating the importance of neural synchrony.

High-performance decoding of temporal features has shown how neural representations of these features facilitate or reflect the segmentation of songs. It provides valuable insights for future research into the intricate processes involved in vocal communication.

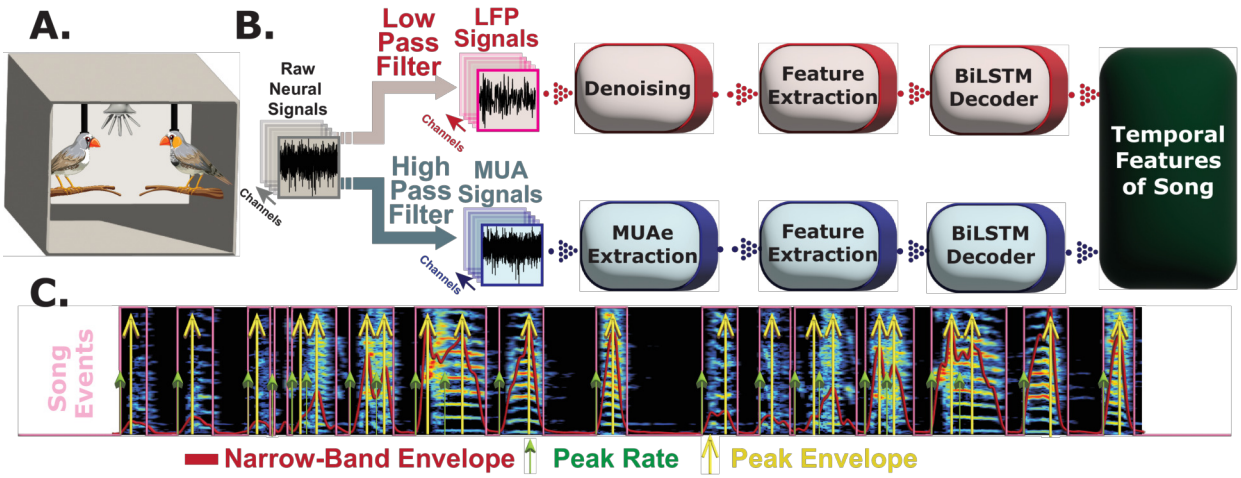

Figure 1. (**A**) Recording setup. (**B**) LFP and MUAe-based decoding system. (**C**) Decoded temporal features.

FAST DYNAMICS OF SENSORY CORTICAL NETWORKS DURING ACTIVE PERCEPTUAL DECISION MAKING INFERRED FROM MASSIVE SPIKING DATASETS

Alex Armstrong, Alex Schwing, Yurii Vlasov *

Electrical and Computer Engineering, University of Illinois, Urbana-Champaign, Illinois, USA
* yvlasov@illinois.edu

We combine experimental and computational methods specifically developed to determine how neuronal populations in mice primary (S1) and secondary (S2) somatosensory cortices are organized in dynamic functional networks that drive perceptually driven behavioral choices.

First, to load cortical circuits with difficult yet manageable cognitive tasks and, at the same time, elicit an ethologically relevant behavior, we designed naturalistic tactile virtual reality in which head-fixed mice while running on a 3D treadmill, are navigating left and right turns provided by closed-loop whisker stimulation [1].

Second, we perform massive recordings of neural activity with Neuropixel probes implanted in principle barrel S1 and S2 cortices during the choice period: the time interval when the animal first senses the approaching virtual wall with whiskers until it makes a choice to change the direction of his run to avoid the collision.

Third, to analyze these massive datasets of simultaneously recorded neural spiking activity and animal behavioral parameters (run angle, run speed and acceleration, paws positions, gait switching from walk to trot, etc.) we developed a data pipeline that incorporates an unsupervised probabilistic model. This DyNetCP model infers dynamic connectivity across the recorded neuronal population from a synchrony of their spiking activity. This approach is fundamentally different from most latent variable models that infer the dynamics of the entire recorded population in a low-dimensional latent feature space. Instead, to enable interpretation as directed time-dependent network connectivity across distinct cellular-level neuronal ensembles, DyNetCP extracts latent variables that are assigned to individual neurons.

The extracted latent temporal dynamics of cellular-level interactions is further correlated with behavioral parameters to uncover how the network connectivity within each functional area (S1 and S2) and between areas is changing during the choice period. The results indicate that extracted patterns of dynamic connectivity change rapidly at specific times corresponding to stimulus presentation, motor action preparation, and choice execution when the running direction is changed. Observed correlation of whisker-guided behavior with dynamics of intra- and inter-area network connectivity reveal therefore specific dynamic organization of cortical networks along the major pathway of tactile information flow that drives perceptual decision making.

References

1. Sofroniew, Vlasov, Hires, *et al.*, 2015, *eLife*, 4:e12559, 10.7554/eLife.12559

THE ROSTRAL ZONA INCERTA: AN INTEGRATIVE HUB IN THE CIRCUITRY OF REWARD PROCESSING?

Edward R. Bader[*], *Emad N. Eskandar, Nathaniel J. Killian*

Department of Neurological Surgery, Albert Einstein College of Medicine, Bronx, New York, USA
[*] edward.bader@einsteinmed.edu

The zona incerta is a nucleus in the subthalamus with extensive interconnectivity with both cortical and subcortical structures and has been implicated in a variety of behaviors including sensorimotor action and novelty. However, its involvement in reward processing itself remains incompletely characterized. Given its connectivity, the zona incerta is ideally placed to serve as a subcortical hub for regulating top-down and bottom-up control of behavior [1]. In particular, its rostral subcomponent, the Rostral Zona Incerta (ZIr), is well positioned to modulate reward circuitry given its monosynaptic connectivity with the ventral tegmental area and lateral habenula [1, 2].

To characterize the neurophysiology of the ZIr during reward processing, we trained C57BL/6 mice ($n = 2$) to perform a head-fixed, classical conditioning task under water restriction. Over repeat presentations, they learned to associate discrete audiovisual cues with either 100%, 90%, or 0% water reward probability. Neurons were analyzed for their cue and reward selectivity in Matlab. Granger causality between pairs of spike trains was used to investigate the internal functional connectivity of ZIr neurons [3].

We recorded a total of 194 ZIr neurons, across 200 trials. Of recorded neurons, 72 (37%) demonstrated significant modulation of firing in response to reward-predicting cues, relative to cues predicting no reward. Additionally, 25 (13%) demonstrated a significant modulation of firing in response to reward delivery itself, relative to absent rewards. GC analyses revealed that during presentation of reward-predicting stimuli, there is enhanced intrinsic functional connectivity relative to during reward delivery (χ^2, $p < 0.00001$). This suggests that ZIr neurons may coordinate circuitry that assigns value to reward-predicting cues, consistent with previous evidence that zona incerta neurons readily form functional ensembles both externally and internally [4].

These results provide preliminary evidence that the ZIr may be a functional hub coordinating diverse brain regions during reward processing. Furthermore, these data suggest that the ZIr may be a promising DBS target for the treatment of diseases with disordered reward processing, including obsessive compulsive disorder and obesity.

Acknowledgments

ERB was supported by the Einstein-Montefiore Clinical and Translational Science Award Hub (1UM1TR004400; MPI: Keller, Kim) and the Predoctoral T32 at Albert Einstein College of Medicine (T32TR004537; MPI: Marantz, Hosgood).

References

1. Haber, Lehman, *et al.*, 2023, *Biol Psych*, 93(11):1010–22, 10.1016/j.biopsych.2023.01.006
2. Yang, Jiang, *et al.*, 2022, *Neurosci Bull*, 38(11):1315–29, 10.1007/s12264-022-00930-w
3. Kim, Putrino, *et al.*, 2011, *PLoS Comput Biol*, 7(3):e1001110, 10.1371/journal.pcbi.1001110
4. Londei, Arena, Ferrucci, *et al.*, 2024, *iScience*, 27(1):108761, 10.1016/j.isci.2023.108761

EVENT GRAPH STRUCTURE DETERMINES FIDELITY OF NEURAL REPRESENTATIONS

Dani S. Bassett[1,*], *Sophie E. Loman*[1], *Ari E. Kahn*[2]

[1] Department of Bioengineering, University of Pennsylvania, Philadelphia, Pennsylvania, USA
[2] Princeton Neuroscience Institute, Princeton University, Princeton, New Jersey, USA
[*] dsb@seas.upenn.edu

Humans receive information from the world around them in sequences of discrete items—from words in language or notes in music to abstract concepts in books and websites on the Internet [1]. To model their environment, humans engage in graph learning, a form of structure learning in which their mental model encodes the graph of event-to-event transition probabilities, typically in medial temporal cortex (MTC).

Recent evidence suggests that some network structures are easier to learn than others [1, 2], and that offers principles for the design of structures that are easiest to learn [3]. But precisely how this differential ease might be explained by differential neural encoding remains unknown. Here we use human fMRI data acquired over a two-day learning experiment to show that the network structure of a temporal sequence of stimuli influences the fidelity with which those stimuli are represented in the brain.

Participants learned a set of stimulus-motor associations following one of two graph structures: modular and lattice. Whereas the motor response could be decoded from neural representations in postcentral gyrus (PCG), the stimulus shape could be decoded from lateral occipital cortex (LOC). The structure of the graph impacted the nature of neural representations: when the graph was modular as opposed to lattice-like, BOLD representations in visual areas better predicted trial identity in a held-out run and displayed higher intrinsic dimensionality.

Notably, classification accuracy was predicted by individual differences in functional connectivity (FC) on the previous day of task performance. Specifically, high LOC-PCG FC and low LOC-MTC FC predicts high PCG classification accuracy, whereas the opposite pattern predicts hippocampal classification accuracy. These competing functional and/or representational dynamics exist atop a broader pattern of changing functional connectivity patterns that track with individual differences in response time, closing the loop between behavioral and neural markers of learning.

Our work demonstrates that graph structure determines the fidelity and dimensionality of event representations, in part through inter-regional functional coupling. Future work could build upon these findings to design, optimize, and adapt network contexts for distinct types of learning over different timescales.

Acknowledgments

Supported by the National Institutes of Mental Health (1-R21-MH-124121–01) and Army Research Office (Grafton-W911NF-16–1-0474 & DCIST-W911NF-17–2-0181).

References

1. Kahn, Karuza, *et al.*, 2018, *Nat Hum Behav*, 2(12):936–947, 10.1038/s41562-018-0463-8
2. Lynn, Papadopoulos, Kahn, *et al.*, 2020, *Nat Phys*, 16:965–973, 10.1038/s41567-020-0924-7
3. Qian, Lynn, Klishin, *et al.*, 2022, *PNAS*, 119(35):e2121338119, 10.1073/pnas.2121338119

IDENTIFYING NEURONAL COMPUTATIONAL AND COMMUNICATION MODULES OF THE FUNCTIONAL CONNECTIVITY ARCHITECTURE IN AREA V1

Christina Brozi[1,2,†], *Eleftheria Psilou*[1,2,†], *Mario Alexios Savaglio*[1,2,†], *Ioannis Smyrnakis*[2,3], *Manos Koniotakis*[2], *Ganna Palagina*[4], *Andreas Tolias*[5], *Maria Papadopouli*[1,2,‡,*], *Stelios M. Smirnakis*[4,‡,⁑]

[1] Dept. of Computer Science, University of Crete, Heraklion, Greece
[2] ICS, Foundation for Research & Technology Hellas, Heraklion, Greece
[3] Hellenic Mediterranean University, Heraklion, Greece
[4] Dept. of Neurology, BW's & JP VA Hospitals, Harvard Medical School, Boston, Mass., USA
[5] Dept. of Neuroscience, Baylor College of Medicine, Houston, Texas, USA
[†] equal contribution, [‡] equal contribution
[*] maria@csd.uoc.gr, [⁑] smsmirnakis@bwh.harvard.edu

Ensembles of neurons firing in synchrony are likely to be more efficient at relaying information to downstream targets [1] and to belong to networks subserving similar functions [2, 3]. Functional connectivity analysis performed during epochs of spontaneous activity (absence of visual stimulation) allows us to identify groups of neurons that participate in synchronous activity patterns, and thus, have the potential to coordinate in processing information. Identifying functional-connectivity patterns under spontaneous conditions and studying potential implications they may have for visual stimulus encoding is important: for example, it has been suggested that spontaneous activity patterns may span a *vocabulary space* shared with population activity patterns elicited during sensory responses. Nevertheless, the relationship between stimulus-driven and spontaneous activity patterns remains obscure.

Here, we used the spike time tiling coefficient metric of correlation strength [4] to measure pairwise functional connectivity between pyramidal neurons from a large population in granular (L4) and supra-granular layers (L2/3) of mouse area V1 recorded via large field-of-view 2-photon microscopy. Approximately 15–25% of neuronal pairs are functionally connected at high statistical significance (z score above 4) both within and across L2/3 and L4. Compared to L2/3, L4 exhibits a higher percentage of strong pairwise correlations, higher clustering coefficients, and higher and more uniformly distributed degrees of connectivity, suggesting a less hierarchical, more robust hub structure. Of particular interest are L4-pyramidal neurons that are functionally connected to a single L2/3 neuron, forming its putative input group. The probability of an L2/3 neuron firing depends strongly on its L4 putative input group and is modulated strongly by brain state as reflected by the pupil size and overall V1-population activity. Interestingly, L2/3 neurons with small L4 putative input groups behave differently than those with large ones both in the slope of the prediction and their modulation by behavioral state.

In sum, resting-state functional connectivity analysis in area V1 allows us to classify pyramidal neurons into different classes depending on their functional connectivity, with interesting implications for stimulus encoding and rules of communication across the V1 circuit.

Acknowledgments

HFRI neuronXnet 2285, neuron-AD 4058; EU MSCA neuronsXnets 101007926; NINDS NS088457.

References

1. Smyrnakis, Smirnakis, 2013, *Neural Comput*, 25(9):2265–2302, 10.1162/neco_a_00476
2. Panzeri, *et al.*, 2022, *Nat Rev Neurosci*, 23(9):551–567, 10.1038/s41583-022-00606-4
3. Kenet, Bibitchkov, *et al.*, 2003, *Nature*, 425(6961):954–6, 10.1038/nature02078
4. Cutts, Eglen, 2014, *J Neurosci*, 34(43):14288-303, 10.1523/jneurosci.2767-14.2014

DATA-DRIVEN SAMPLING OF MOTOR CORTICAL NETWORKS DURING REACHING

Ryan Canfield[1,*], *Tomohiro Ouchi*[2], *Leo Scholl*[2], *Pavithra Rajeswaran*[1], *Lydia Smith*[2], *Amy Orsborn*[1–3]

[1] Bioengineering, University of Washington, Seattle, Washington, USA
[2] Electrical & Computer Engineering, University of Washington, Seattle, Washington, USA
[3] Washington National Primate Research Center, Seattle, Washington, USA
[*] rcanfiel@uw.edu

The coordinated activity of neuronal populations, or networks, within motor cortices facilitates movements. Motor cortex is often studied by broadly sampling neurons and then inferring population-level computations. Neurons, however, comprise well-structured networks defined by anatomy. Coarse sampling across large areas may therefore obscure our ability to accurately infer how neural populations implement movement-related computations. We present a new experimental paradigm for doing data-driven sampling of primate motor cortices, which has yielded preliminary data suggesting complex spatial organization of task information.

We developed a method to target networks in frontal motor cortical areas. Our two-part approach first leverages high-density micro-electrocorticography (µECoG) to map functional responses across multiple motor cortical areas during movement tasks (Figure 1a). We then used these maps to target high-density, laminar, microelectrode arrays (Neuropixels) to study networks across large spatial areas at the single neuron level within the same animal (Figure 1b).

Our µECoG measurements revealed localized differences in movement-related information across motor cortex in one rhesus macaque (male, 10 years old). These results were supported at the neuron level in dorsal pre-motor cortex (PMd), where we found that task information was correlated with µECoG measurements. However, this correlation was not seen in single neuron recordings in primary motor cortex (M1) (Figure 1c). Our results demonstrate a method for data-driven sampling and has revealed that the spatial distribution of task information is complex. The discrepancy of results between motor cortical regions highlights challenges in investigating functionally defined networks across multiple areas. Our results also show the potential importance of combining data-driven sampling with population analysis to better understand motor encoding.

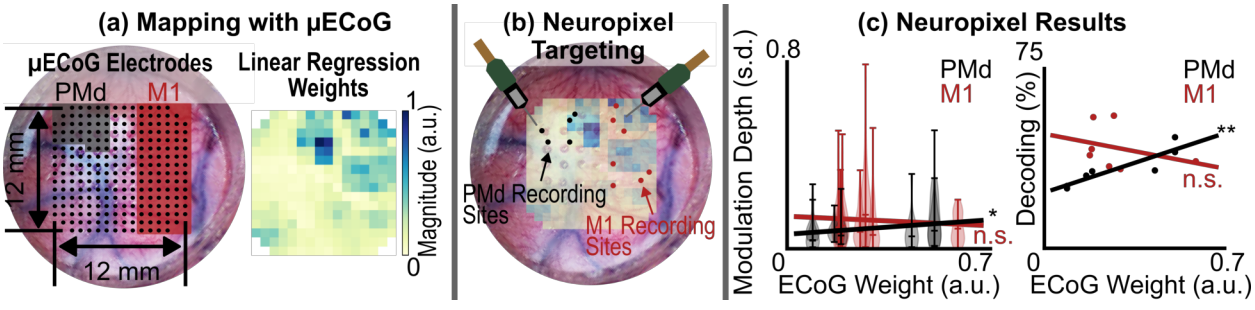

Figure 1. (**a**) µECoG array coverage and map of movement-related information. (**b**) Neuropixel targeting based on the µECoG maps. (**c**) Neuropixel movement-related information compared with that of µECoG using single-neuron tuning (left) and neuron population decoding accuracy (right).

DEVELOPMENT OF A CONTROL SYSTEM FOR A NEXT GENERATION HIGH-DENSITY, WIRELESS, BIDIRECTIONAL BRAIN-COMPUTER INTERFACE

R. James Cotton[1,2,*], *Taesung Jung*[3], *Nanyu Zeng*[3], *Zhe Li*[4–6], *Konstantin Willeke*[4–6], *Jason Fabbri*[3], *Guy Eichler*[3], *Saumil Patel*[4–6], *Andreas S. Tolias*[4–6] *Ken Shepard*[3]

[1] Physical Medicine and Rehabilitation, Northwestern University, Chicago, Illinois, USA
[2] Shirley Ryan AbilityLab, Chicago, IL, USA
[3] Electrical Engineering, Columbia University, New York City, New York, USA
[4] Ctr. Neuroscience and Artificial Intelligence, Baylor College of Medicine, Houston, Texas, USA
[5] Neuroscience, Baylor College of Medicine, Houston, Texas, USA
[6] Electrical and Computer Engineering, Rice University, Houston, Texas, USA
[*] rcotton@sralab.org

A substantial barrier to the wider clinical adoption of brain-computer interfaces is the requirement for percutaneous connections. The Bioelectric Interfacing to Sensory Cortex (BISC) is a next-generation, wireless, implantable, bidirectional, 65,536 channel implant that overcomes this barrier. The fully integrated complementary metal-oxide semiconductor design enables a flexible, 50 µm thin device with 6.8 mm by 7.4 mm surface area that can be slipped into the subdural space and offers unparalleled channel density per unit volume. We present the results of the initial testing of this novel device after chronic implantation in the visual system of a non-human primate (NHP).

BISC is powered by and communicated with a wireless antenna placed over the closed skull and scalp. This antenna is controlled by a relay station that uses custom FPGA logic to manage the 100 Mbps data stream and device communication protocol. Serialized data is sent via ethernet to our existing NHP stimulation acquisition framework, for real-time visualization and subsequent analysis. Daily alignment of this antenna was straightforward and reliable. The wireless link was stable with negligible packet loss, even in the presence of animal movements.

We presented a variety of stimuli as the NHP fixated on the screen. These include receptive field mapping stimuli, orientation mapping stimuli, and natural images. We performed these experiments under a variety of recording motifs supported by BISC with varying electrode sampling densities. These experiments replicated classical findings of the NHP visual system. For example, high-density recordings over V1 showed spatially compact receptive fields with the expected retinotopic gradients and a gradient reversal at the border of V1. These receptive fields were stable over sessions spanning 2 months. We also trained artificial neural networks to predict the responses across the array (V1 to V4) and computed the most exciting images (MEI) for these models. The MEIs captured the localized receptive fields in V1. There was a sharp qualitative transition in MEI style over V4, where textures and colors became much more prevalent. These predictive models can also be used to reconstruct the stimulus shown, which shows promising results when applied to BISC responses.

The results from this initial test demonstrate the power and promise of the BISC design. The flexible, wireless, and highly manufacturable subdural implant produced a stable, high-density interface to the visual cortex of an NHP. The wireless power and data link allows continuous streaming of 100 Mbps of neural recordings. This ultra-high density microelectrocorticography regime is novel, and an important future direction is to optimize the signal processing and recording configurations to maximize the effective bandwidth. For this purpose, the extensively studied and high-acuity NHP visual system, therefore, serves as a strong benchmark for this purpose.

USING DEEP LEARNING SYNTHESIS OF OPTIMAL STIMULI TO STUDY OBJECT RECOGNITION ACROSS MOUSE LATERAL VISUAL HIERARCHY

Maria Diamantaki[1,4,*], *Konstantin Willeke*[2,3], *Saumil Patel*[4,5], *Katrin Franke*[4,5], *Fabian Sinz*[2,3], *Andreas Tolias*[4,5], *Emmanouil Froudarakis*[1,6]

[1] Inst. Molecular Biol. and Biotechnol., Found. for Res. and Technol. Hellas, Heraklion, Greece
[2] Institute for Bioinformatics and Medical Informatics, Tübingen University, Tübingen, Germany
[3] Inst. for Computer Sci. and Campus Insti. Data Sci., Univ. Göttingen, Göttingen, Germany
[4] Department of Neuroscience, Baylor College of Medicine, Houston, Texas, USA
[5] Center for Neurosci. and Artificial Intelligence, Baylor College of Medicine, Houston, Texas, USA
[6] School of Medicine, University of Crete, Heraklion, Greece
[*] maria_diamantaki@imbb.forth.gr

Invariant object recognition is the ability of animals to rapidly recognize objects irrespective of variations in their appearance. In primates, this remarkable ability is mediated by the ventral visual stream, a set of hierarchically organized interconnected visual areas. Due to the hierarchical organization of deep neural networks, they have been routinely used to both model and predict the stimuli that optimally activate the neurons along ventral visual stream. In mice, anatomical and physiological studies have revealed a network of lateral higher-order cortical visual areas (LVAs) which are believed to form the mouse ventral visual stream. While deep learning approaches have explored optimal stimuli in mouse visual areas, similar to primate studies, they have mostly been restricted to functional imaging data from the primary visual cortex (V1). Here we aim for a detailed characterization of the optimal excitable stimuli of LVAs and their role for invariant object representation.

Therefore, we used large-scale electrophysiological data from multiple visual cortical areas to generate a digital twin of the mouse visual cortex. The digital twin model allowed us to synthesize optimal visual stimuli that maximally drive the neurons (most exciting inputs, MEIs). Specifically, by using the Neuropixels probes we simultaneously recorded the activity of hundreds of neurons *in vivo* in mouse V1 and LVAs in response to natural images. We then trained a convolutional neural network to predict the responses of each neuron recorded across the different areas and generated a set of MEIs that would optimally excite the recorded neurons. Subsequently, we showed the optimized stimuli back to the mice and recorded the activity of the same neurons.

This closed loop approach enabled us to verify *in vivo* the results of the *in silico* model predictions. Furthermore, we showed differences in the size and similarity of V1 and LVAs MEIs, and that MEIs of LVAs carry and extract more information about the object identity. Additional investigations are ongoing to uncover neural invariances across the visual hierarchy and how are these affected by the low-level properties we have characterized in this study. Identifying the neural invariances across LVAs will allow us to dissect the role of hierarchical processing in complex cortical computations such as invariant object recognition.

Acknowledgments

Supported by the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement No. 101025482.

HOW DOES WORKING MEMORY AFFECT NEURAL POPULATION ACTIVITY ACROSS THE CORTEX

E. Mika Diamanti[1,*], *Lucas Pinto*[2,3], *Manuel Schottdorf*[1], *Stephan Y. Thiberge*[1], *Carlos D. Brody*[1], *David W. Tank*[1]

[1] Princeton Neuroscience Institute, Princeton University, Princeton, New Jersey, USA
[2] Department of Neuroscience, Northwestern University, Chicago, Illinois, USA
[3] Department of Neurobiology, Northwestern University, Evanston, Illinois, USA
[*] emdiamanti@princeton.edu

Complex behaviors requiring working memory (WM) recruit multiple brain regions across the cortex, suggesting a distributed neural mechanism. Yet it is not clear whether the recruitment of distributed neural populations directly relates to the working memory component per se.

We trained 9 mice expressing the Calcium indicator GCaMP6s to perform a WM task, *towers*, in virtual reality, which combined navigation with sensory-evidence accumulation. As mice ran down the stem of a T-maze, they had to sample the cues of towers that flashed on the side walls, retain this information in memory during a delay period and finally turn towards the side with most cues. Crucially, the towers task was alternated in blocks with a simpler *visually guided* task featuring the same virtual environment but without requiring WM, since a prominent visual landmark at the end of the maze indicated the correct turn throughout the trial. Mice performed the two tasks with the same motor behavior. Thus, given the similarity in sensory and motor variables between the two tasks, any difference in neural activity should be attributed to the WM component.

We used a random-access two-photon microscope to image hundreds of excitatory neurons in three distinct cortical areas simultaneously, visual AM, retrosplenial and premotor M2. We observed only subtle differences in the neural population activity between the two tasks, suggesting that it mainly captured common task variables. For instance, trial-averaged activity across each area's population formed sequences that were indistinguishable between tasks. However, when we decoded evidence from each area, we found that only M2 exhibited higher decoding accuracy during the delay period of the towers task, compared to the visually guided one. Next, we asked whether differences in evidence encoding between the two tasks could be observed at the level of collective dynamics across areas. We decoded sensory evidence directly from trial-by-trial pairwise correlations, separated into within- and across-areas, the latter acting as a proxy for inter-areal interactions. Although pairwise correlations contained information about task variables, only M2's across-area correlations encoded evidence more strongly during the delay of the towers task. Finally, we asked if the subtle difference in neural activity between the two tasks was also reflected in the intrinsic dimensionality of each area. Only M2's intrinsic dimensionality was higher in the towers than the visually guided task by a few dimensions.

Therefore, WM related neural activity is not readily observable at the trial-averaged level. Rather, our task comparison approach shows that trial-by-trial fluctuations in the activity of M2 neurons, and their interactions with other cortical areas, have a unique role in supporting memory maintenance of accumulated evidence.

LOCAL FIELD POTENTIAL SIMULATION ACROSS A V1 CORTICAL MODEL

Rares A. Dorcioman, Giulia Moreni, Cyriel Pennartz, Jorge F. Mejias[*]

Cognitive and Systems Neuroscience Group, SILS, UvA, Amsterdam, Netherlands
[*]j.f.mejias@uva.nl

The local field potential (LFP) has become an increasingly important brain signal to study due to its role in neural mechanisms such as visual perception and decision making, as well as its potential to be utilized in developing brain-computer interfaces. Its exact origin is still in question, with sources such as synaptic currents, intrinsic potential oscillations, Ca^{++} and Na+ spikes contributing to it. Given its complexity, computational models have been employed to understand the individual contributions to the LFP and to interpret the extracellular potential as it relates to the underlying neural activity. Most commonly, these models incorporate detailed biophysical information of neurons to generate LFPs. However, such detailed models require a high number of GPUs and CPUs, as well as large running times to compute, making them less suitable to simulate large networks. Moreover, the complexity of the biophysical details involved causes them to be less accessible to the wider community.

To address these limitations, we implemented a computational model of a cortical column representing the mouse V1 area, composed of neurons with simplified but realistic morphological structure and dynamics, suitable for large-network simulations. The model consists of 17 sub-populations spanning cortical layers 1 to 6, including pyramidal neurons, PV, SST and VIP interneurons. Using dendritic structural information from the Allen Brain Atlas, the model successfully reproduced experimental firing rates across the column under spontaneous conditions.

A novel feature of our model is a tool that allows users to design custom-made probes (*e.g.*, Neuropixels, NeuroNexus) and record the LFP at any point across the column, with flexible number of channels, inclination angle and depth. We simulated a visual stimulation experiment while recording the LFP with a virtual probe emulating a 64-channels NeuroNexus Technologies Polytrode, and compared the resulting Current Source Density (CSD) profile with *in vivo* mouse data. We found that the modeled CSD exhibited the expected combination of sink and source rising upon visual stimulation. To further illustrate the model's capacity to simulate complex conditions, we conducted two proof-of-concept visual stimulation experiments: the first with optogenetic inhibition of all PV interneurons within the column and the second with blocking part of the AMPA receptors' activity across the network. Thus, depending on the research question a user is investigating, the model allows a wide variety of settings, such as inhibition of specific sub-populations or cortical layers, as well as the targeted or complete blocking of certain receptors.

These preliminary results suggest that our simplified model can predict relevant extracellular potentials across the cortical region. Due to its versatility and significantly reduced computational requirements compared to its biophysically detailed counterparts, this tool could assist scientists in planning experimental work and in testing and forming hypotheses through *in silico* experiments.

A RETINO-CORTICAL MODEL OF ANTICIPATION

Jerôme Emonet[1,*], *Alain Destexhe*[2], *Frédéric Chavane*[3], *Bruno Cessac*[1]

[1] Biovision, Inria Sophia-Antipolis, Valbonne, France
[2] Neurosciences computationnelles, NeuroPsi, Gif-sur-Yvette, France
[3] NeOpTo, Institut de Neurosciences de la Timone, Marseille, France
[*] jerome.emonet@inria.fr

When processing a moving object, the visual system uses several anticipation mechanisms to compensate the delay induced by the photo-transduction [1]. On one hand, retinal gain control advances the peak in sensory cells response. On the other hand, due to the lateral connectivity, the activity begins to raise before the moving stimulus enters in the receptive field of the sensory cells. This last effect, a latency shift, is observed both in the cortex and in the retina [2].

We present a model assembling a realistic retinal input fed by a video stimulus to a cortical model of V1. We address the respective role of retinal and cortical anticipation mechanisms [3]. Our study unravels the role of retinal amacrine cells in the latency shift.

The figure to the right presents a simulation of the V1 cortical activity, in response to a moving bar, obtained (i) with amacrine cells and gain control inactivated (control conditions, upper plot) and (ii) with amacrine cells activated (lower plot). In both conditions the activity starts well before the time of the stimulus arrival in the center of the receptive field of the cortical column. This activity is the anticipation induced by latency shift. With amacrine cells engaged, the peak shifts quite significantly, here about 47 ms before the control conditions.

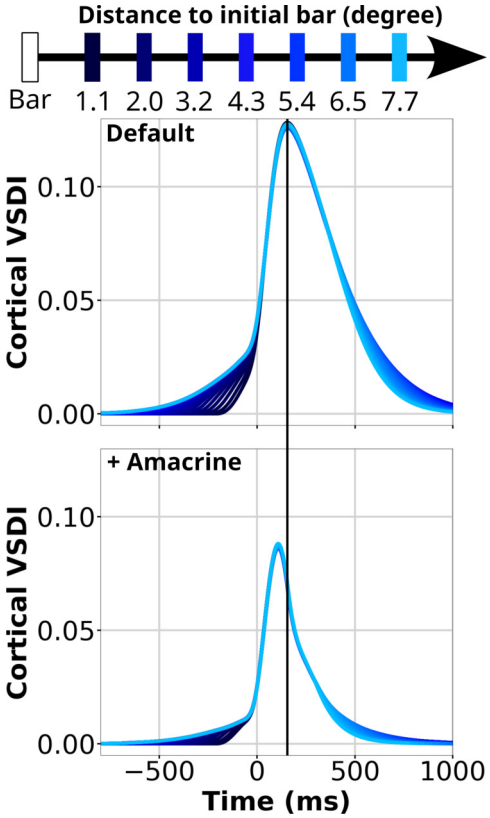

Figure 1: Effect of amacrine activity on cortical response and anticipation. The simulated voltage sensitive dye imaging (VSDI) activity of the middle grid row have been superposed and centered on the arrival of the center of the bar in the middle of their receptive field. The blue color gradient correspond to the distance (in degrees) between the cortical column and the initial position of the bar. The black vertical line correspond to the VSDI peak with default parameters.

Acknowledgments

Supported provided by the ANR Shooting Star (ANR-20-CE37-0018).

References

1. Berry, Brivalou, Jordan, Meister, 1999, *Nature*, 398(6725):334–8, 10.1038/18678
2. Johnston, Lagnado, 2015, *eLife*, 4:e06250, 10.7554/eLife.06250
3. Souihel, Cessac, 2019, ICMNS, inria.hal.science/hal-02167737

DISSOCIATION OF DECODING OF SUBJECTIVE CONFIDENCE AND STIMULUS UNCERTAINTY IN EEG AND PUPIL SIGNALS IN HUMAN

Jamal Esmaily[1,2,*], *Mustafa Yavuz*[2], *Ophelia Deroy*[3,4], *Bahador Bahrami*[1]

[1] Depat. General Psych. and Education, Ludwig Maximilians University Munich, Munich, Germany
[2] Grad. School Systemic Neurosci., Ludwig Maximilians University Munich, Munich, Germany
[3] Faculty of Philosophy, Ludwig Maximilian University, Munich, Germany
[4] Munich Centre for Neuroscience, Ludwig Maximilian University, Munich, Germany
[*] jamal.esmaily@campus.lmu.de

Numerous neuroscientific works show that in perceptual decision-making, subjective confidence in choices depends on external factors such as stimulus noise as well as internal factors such as state of arousal and neuronal noise. Aiming to help disentangle the contributions of these internal and external factors to subjective confidence, we present data from 3 separate experiments (n = 12, n = 15, and n = 29) where we had participants report choice and confidence in a classical random dot motion task while we measured both EEG and pupil size signals.

Our results replicate previous works, with subjective confidence and motion coherence covarying with choice accuracy. Yet they reveal a novel separation: while the trial-by-trial variations in external stimulus uncertainty could be decoded from (reading through) EEG, pupil data captured the variations in subjective confidence. Importantly, these dissociations remained, albeit to a lesser yet still significant degree, even after removing the covariation between coherence and confidence.

These results help clarify the cognitive processes underlying decision making by identifying pupil and EEG signals as separate windows into the processing of external and internal variables that contribute to our subjective sense of confidence.

FUNCTIONAL CONNECTOMICS REVEALS GENERAL WIRING RULE IN MOUSE VISUAL CORTEX

Paul G. Fahey[1,*], *Zhuokun Ding*[1], *Stelios Papadopoulos*[1], *Eric Wang*[1], *Christos Papadopoulos*[1], *Anthony Ramos*[1], *MICrONS Consortium*[1,3,4], *Kresimir Josic*[7], *Robert Rosenbaum*[8], *H. Sebastian Seung*[4], *Forrest Collman*[3], *Nuno Macarico da Costa*[3], *R. Clay Reid*[3], *Xaq Pitkow*[1,2,9,10], *Edgar Y. Walker*[5,6], *Jacob Reimer*[1], *Andreas S. Tolias*[1,2]

[1] Department of Neuroscience, Baylor College of Medicine, Houston, Texas, USA
[2] Department of Electrical and Computer Engineering, Rice University, Houston, Texas, USA
[3] Allen Institute for Brain Science, Seattle, Washington, USA
[4] Princeton Neuroscience Institute, Princeton University, Princeton, New Jersey, USA
[5] Department of Physiology and Biophysics, University of Washington, Seattle, Washington, USA
[6] UW Computational Neuroscience Center, University of Washington, Seattle, Washington, USA
[7] Department of Mathematics, University of Houston, Houston, Texas, USA
[8] Department of ACMS, University of Notre Dame, Notre Dame, Indiana, USA
[9] Neuroscience Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA
[10] Department of Machine Learning, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA
[*] pgfahey@stanford.edu

To understand how the brain processes information, it is informative to study how neural circuit connectivity relates to neural function. Previous research has shown that excitatory neurons in layer 2/3 of mouse primary visual cortex with similar response properties are more likely to form connections [1, 2], a relationship referred to as *like-to-like* connectivity. However, technical challenges of combining synaptic connectivity and functional measurements have limited these studies to few, highly local connections. With the millimeter scale and nanometer resolution of the MICrONS dataset, we extended the study of the connectivity-function relationship to excitatory neurons of the mouse visual cortex across inter-laminar and inter-area projections, assessing connection selectivity at the coarse axon trajectory and fine synaptic formation levels. A digital twin model of this mouse, that accurately predicted responses to arbitrary video stimuli, enabled a comprehensive characterization of the function of neurons. We found that neurons with highly correlated responses to natural videos tended to be connected with each other, not only within the same cortical area but also across multiple layers and visual areas, including feedforward and feedback connections. The digital twin model separated each neuron's tuning into a feature component (what the neuron responds to) and a spatial component (where in the visual field the neuron responds to). We show that the feature, but not the spatial component, predicted which neurons were connected at fine synaptic resolution. Our results demonstrate that the like-to-like connectivity rule generalizes to multiple connection types, and the rich MICrONS dataset is suitable to further refine a mechanistic understanding of circuit structure and function.

Acknowledgments

Supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/Interior Business Center (DoI/IBC) contract number D16PC00003.

References

1. Ko, Hofer, Pichler, *et al.*, 2011, *Nature*, 473(7345):87–91, 10.1038/nature09880
2. Lee, Bonin, Reed, *et al.*, 2016, *Nature*, 532(7599):370–374, 10.1038/nature17192

NEURAL SIGNATURES OF STRESS SUSCEPTIBILITY AND RESILIENCE IN THE AMYGDALA-HIPPOCAMPAL NETWORK

Valeria Fascianelli[1,†,*], *Frances Xia*[2,†], *Nina Vishwakarma*[2], *Frances Grace Ghinger*[2], *Stefano Fusi*[1,3], *Mazen A Kheirbek*[2–4]

[1] Center for Theoretical Neuroscience, Columbia University, New York, USA
[2] Department of Psychiatry, University of California, San Francisco, California, USA
[3] Kavli Institute for Brain Science, Columbia University Irving Medical Center, New York, USA
[4] Weill Institute for Neurosciences, University of California, San Francisco, California, USA
[†] equal contribution
[*] vf2266@columbia.edu

Mood disorders are characterized by behavioral changes arising from network-level alterations, with the basolateral amygdala (BLA)-ventral CA1 (vCA1) pathway being a central component. Anhedonia, the reduced ability to experience pleasure, is a core feature of mood disorders. However, while BLA and vCA1 are known to encode outcome-specific reward representations to guide decision-making, how these functions are impacted by changes in emotional state remain unclear. Furthermore, how neural dynamics in the BLA or vCA1 may differ in mice susceptible or resilient to chronic stress, and how circuit-specific interventions may reduce susceptible phenotypes remain unknown.

Here, we used Neuropixels probes to record BLA and vCA1 single units in mice following chronic social defeat stress (CSDS) to assess how stress modulates reward choice representations and reward-seeking behavior. First, in a reward-choice task, we found that basolateral amygdala (BLA) activity in resilient mice showed enhanced discrimination of upcoming reward choices. Next, we explored the origins of anhedonic behavior in susceptible mice and found that they tended to switch to, and stay on, low-value reward choices. In line with this, a linear decoder could classify the intention to switch or stay on a previously chosen reward in BLA of susceptible mice, reminiscent of a rumination-like state observed in individuals with depression.

Furthermore, in the absence of overt task stimuli, we found that the geometry of the spatial structure of hidden neural states during spontaneous activity in the BLA of susceptible mice was higher dimensional than controls, reflecting the exploration of a larger number of distinct neural states. This may reflect the emergence of intrusive activity patterns in the BLA, like intrusive thought patterns observed in depressed patients. Notably, we found that features of neural activity in the BLA during a task-free period were more effective than behavioral readouts in distinguishing between control mice and those with a history of CSDS. This suggests the intriguing possibility that resting-state activity patterns in the BLA hold significant potential as a powerful biomarker for predicting individuals who have experienced a stressful event. Finally, manipulation of vCA1 inputs to the BLA reversed dysfunctional BLA neural dynamics in susceptible mice and rescued anhedonic behavior.

In summary, by looking at the geometry of neural representation, we identified resilience and susceptibility signatures in the BLA-vCA1 network and highlighted the BLA-vCA1 circuit as a promising target for neuromodulation in mood disorder treatments.

AN INTERACTIVE FLOOR FOR STRUCTURED FREELY MOVING FORAGING BEHAVIOR

Tucker G. Fisher[*], *Marielena Sosa, Lisa M. Giocomo*

Department of Neurobiology, Stanford University School of Medicine, Stanford, California, USA
[*]tgfisher@stanford.edu

Although entorhinal cortex (EC) is most often associated with grid cells, much of the activity of this area is heterogeneous and not well characterized [1]. For example, medial EC and other parahippocampal regions remap in line with task performance during more demanding freely moving tasks [2]. Even grid cells are more exhibit complex coding properties outside of the common free foraging tasks. We have set out to probe these regions with a structured foraging-like task aimed at driving spatially variable, but overt, trajectories in a 2-D environment. In addition our task is designed to decouple reward from one given location and hinges on the reward free goal acquisition an intermediate target.

Up to now, two adult rats with chronic Neuropixels 1.0 implants ran instructed trajectories to targets in an open arena. Targets were circular pools of light projected up through the floor (Fig. 1 left, two example trajectories visiting the same target. One with an error in phase II.). Trajectories had three phases: (I) Cued trial initiation well to target collision, (II) from target to remembered reward well, and (III) from reward to an LED indicated initiation well. Animals waited in a rest box between blocks. In order to promote trajectories with complete spatial coverage, we selected target centers at random from a 4 by 6 grid with about 25cm spacing. This supported target acquisitions with decent spatial uniformity and good spatial coverage (Fig. 1 right, colored dots and smoke lines).

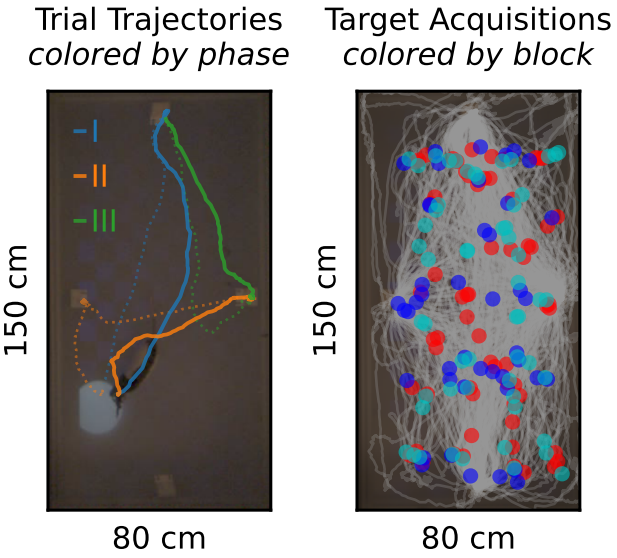

Figure 1. Instructed trajectories to targets with uniform spatial coverage.

We plan to focus our investigations on the movement information that is invariant to position, such as changing coding properties around key transitional phases of the task. An initial look into these trial-based patterns suggests position agnostic trajectory coding. After aligning trials to target collision and referencing spikes based on the fraction of the trajectory completed (*i.e.* stretching or compressing trials around the target collision) we observe cells with some consistent banding or ramping, and look forward to addressing the population activity more formally soon.

Acknowledgments

NIH Grants 1R01MH126904–01A1 (LMG), RO1MH130452 (LMG), BRAIN Initiative U19NS118284 (LMG), The James S. McDonnell Foundation (LMG), The Simons Foundation 542987SPI (LMG)

References

1. Nayebi, Attinger, Campbell, *et al.*, 2021, *Proc NeurIPS 34*
2. Gonzalez, Giocomo, 2024, *eLife* Reviewed Preprint, 10.7554/eLife.85646.2

CHARACTERIZATION OF THE DYNAMICS OF NEURONAL ENGAGEMENT TO THE MICRO-PROGRESSION OF ACUTE SEIZURES IN MOUSE CORTICAL CIRCUITS

Ioannis Fotis[1,2], *George Touloumes*[4], *Joseph Lombardo*[4], *Maria Markaki*[2], *Maria Papadopouli*[1–3,*], *Stelios M. Smirnakis*[4,*]

[1] Dept. of Computer Science, University of Crete, Heraklion, Greece
[2] ICS, Foundation for Research and Technology Hellas, Heraklion, Greece
[3] Archimedes Research Unit, Athena Research Center, Athens, Greece
[4] Dept. of Neurology, BW's and JP VA Hospitals, Harvard Medical School, Boston, Mass., USA
[*] maria@csd.uoc.gr, smsmirnakis@bwh.harvard.edu

Epilepsy is a common neurological disorder, whose circuit mechanisms are poorly understood [1]. Focal injury in acquired forms of epilepsy is thought to induce an abnormal balance between excitation and inhibition that drives neuronal networks into self-perpetuating oscillatory activity states [2] (seizure). A major unsolved question is how neurons get recruited, *in vivo*, during the evolution of seizure events. Specifically, it is not known whether neurons fire in a stereotyped pattern or sequence per seizure event, whether this happens reliably, and how it depends on neuronal type.

Here, we injected the chemoconvulsant 4-aminopyridine (4-AP) in the mature neocortex (layer V), performed *in vivo* 2P calcium imaging in the primary visual cortex (V1), and simultaneously recorded the ECoG signal of awake mice. For each neuron, we defined its *local plateaus* as periods of abnormally increased and prolonged calcium activity based on thresholds derived from recordings before 4-AP injection, and *bursts* as periods when at least 10% of neurons showed local plateaus simultaneously. Neurons were then divided into temporal quartiles based on the time intercept of the tangent at the half maximum of the sigmoidal fit of their local plateaus. Acute seizure micro-progression was characterized by a cyclical alternation of spatial expansion and contraction in the engagement of neuronal ensembles starting from the injection site. This spatio-temporally organized pattern is repeated until the organization fades away and only a few sparse neurons are engaged. The use of ECoG, in addition to Ca signal, allows us to distinguish three types of bursts, namely the spreading depression bursts, the seizures, and the epileptiform spiking ones. Spreading depression bursts have a stereotypical profile, characterized by a suppressed ECoG signal, contrary to the seizure and epileptiform spiking ones. Additionally, we identified individual epileptiform spikes in the ECoG recordings and calculated the average Ca signal modulation around the spikes. Epileptiform spikes are consistently associated with sharp increases in the Ca signal. In sum, this work develops a methodology that uses EEG and Ca signals to analyze the spatiotemporal dynamics of recruitment of individual neurons into epileptic events, *in vivo*, and studies the patterns of engagement. This methodology could potentially form a basis for the development of new circuit-based therapeutic strategies targeting specific cell classes.

Acknowledgments

Funded by NINDS R21 NS088457, BLR&D VA I01BX004729-01A1, EU Horizon2020 MSCA neuronsXnets 101007926, Fondation Sante, and HFRI neuronXnet 2285 and neuron-AD 4058.

References

1. Hirtz, *et al.*, 2007, *Neurol*, 68(5):326–337, 10.1212/01.wnl.0000252807.38124.a3
2. Sisodiya, Mefford, 2011, *Curr Opin Neurol*, 24(2):140–145, 10.1097/WCO.0b013e328344062f

IDENTIFYING THE NEURONAL SELECTIVITY LANDSCAPE IN MACAQUE AREA V4: A DEEP LEARNING AND INCEPTION LOOP APPROACH

Katrin Franke[1,2,*], *Nikos Karantzas*[1,2,‡], *Kandan Ramakrishnan*[1,2], *Konstantin Willeke*[3], *Pavithra Elumalai*[3], *Paul Fahey*[1,2], *Kelli Restivo*[1,2], *Tori Shinn*[1,2], *Cate Nealley*[1,2], *Gabrielle Rodriguez*[1,2], *Saumil Patel*[1,2], *Alexander Ecker*[3], *Fabian Sinz*[3], *Andreas Tolias*[1,2]

[1] Department of Neuroscience, Baylor College of Medicine (BCM), Houston, Texas, USA
[2] Center for Neuroscience and Artificial Intelligence, BCM, Houston, Texas, USA
[3] Institute of Computer Science, University of Göttingen, Göttingen, Germany
[*] kafranke@stanford.edu [‡] nikoskar@stanford.edu

Area V4 in the primate visual system is primarily involved in processing visual information related to object features such as shape, color, and texture, playing a crucial role in object detection and discrimination. Most research to date has characterized the selectivity of V4 neurons to individual stimulus features (*i.e.* color or texture) and, therefore, a systematic understanding of the selectivity landscape of V4 neurons across features is missing. Here, we addressed this important problem in the macaque as a model system.

We performed large-scale laminar probe recordings in fixating macaques (over 300 neurons from $n = 1$ animal) while showing natural images exhibiting complex shape, color and texture compositions (Fig. 1a) to the monkey. Based on this data, we trained convolutional neural network (CNN) models, which accurately predicted the firing rate of single V4 neurons to hold-out test images (Fig. 1b), allowing us to perform a thorough *in silico* analysis of stimulus selectivity, jointly for multiple object features at the same time (examples in Fig. 1c).

For the *in silico* analysis, we used natural-scene inspired rendered scenes that provide tight control over distinct stimulus features, while still matching natural scene statistics. Then, we systematically quantified each neuron's selectivity to different features and their combinations by predicting the firing rate to manipulated scenes with the trained convolutional neural network (CNN) model. We are currently recording data from a second monkey and are confirming model predictions *in vivo* by performing closed-loop experiments. Our results will reveal how V4 neurons tile the landscape of object feature selectivity, thereby significantly contributing to a deeper understanding of the neuronal correlates of object recognition.

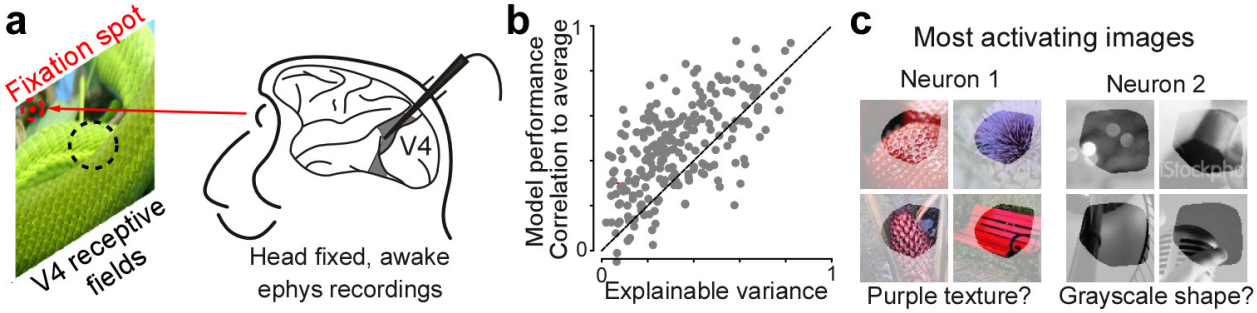

Figure 1. Determining most activating images for V4 neurons and building CNNs to model single unit responses.

ONE SHOT LEARNING IN THE HUMAN BRAIN

Megha Ghosh, Sophia Lowe-Hines, Adam Crandall, Andrew Ko, Jeffrey Ojemann, Benjamin L. Grannan *

Department of Neurosurgery, University of Washington, Seattle, Washington, USA
* bgrannan@uw.edu

Humans have a remarkable ability to learn the meaning of a novel word based on preceding context [1, 2]. How does the human brain achieve such one-shot learning after a single exposure to a new item? Here, we study neural dynamics during a one-shot language learning task to answer this question.

Human participants (n = 9) undergoing intracranial-EEG monitoring participated in the task where they were presented with sentences ending with either a real (non-learning) or pseudo word (learning) (Fig. 1A). Behnke-Fried micro-macro electrodes were used for isolating activity of single neurons from the hippocampus (HPC).

We found significantly higher theta power during the last word of the sentence for learning compared to non-learning trials in the language dominant left HPC, middle temporal gyrus (MTG) and anterior nuclei of the thalamus (ANT) (Fig. 1C, D). To assess network activity, we performed phase-locking analysis of HPC units to theta rhythms in the above regions. Single HPC units showed higher phase locking to theta rhythms during learning trials compared to non-learning (Fig. 1E, F). Together, these results support a network level response spanning cortical and neocortical structures, that may mediate rapid linguistic learning. Further understanding of this learning mechanism can inform therapeutics for individuals with verbal learning deficits and provide keen insights into development of artificial language models being engineered to achieve one-shot word learning.

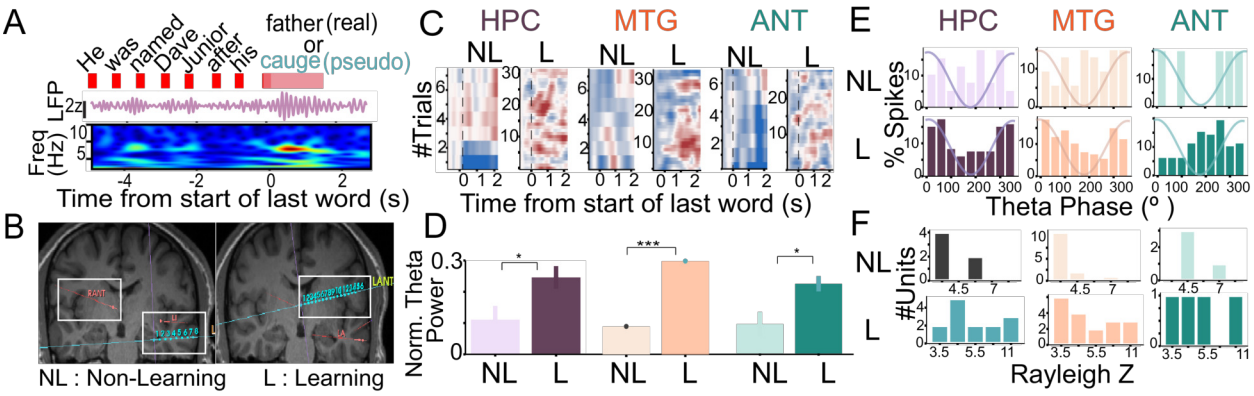

Figure 1. (**A**) Task structure with simultaneous intracranial recordings. (**B**) Example location of electrodes in the HPC, MTG and ANT. (**C**) Theta power changes during the last word across trials for a single subject. Warmer colors refer to higher power. (**D**) Population comparisons of theta power (HPC:28 channels, MTG:80, ANT:6, paired t-test) (*: $p < 0.05$, ***: $p < 0.001$) during learning (L) and non-learning (NL) trials. (**E**) Significant and no phase locking of example hippocampal units to theta in the HPC, MTG and ANT during L and NL trials respectively. (**F**) Number of units with significant phase locking (Rayleigh $Z > 3$) for L and NL trials for the above 3 brain regions.

References

1. Borovsky, Kutas, Elman, 2010, *Cognition*, 116(2):289–296, 10.1016/j.cognition.2010.05.004
2. Lee, O'Doherty, Shimojo, 2015, *PLoS Biol*, 13(4):e1002137, 10.1371/journal.pbio.1002137

OVERLAPPING E/I NEURONAL ASSEMBLIES GENERATE RICH NETWORK DYNAMICS AND ENABLE COMPLEX COMPUTATIONS

Emmanouil Giannakakis[1,2,†,*], *Victor Buendia*[1,2,†], *Oleg Vinogradov*[1,2], *Sina Khajehabdollahi*[1], *Anna Levina*[1,2]

[1] Department of Computer Science, University of Tübingen, Germany
[2] Max Planck Institute for Biological Cybernetics, Tübingen, Germany
[†] equal contribution
[*] giannakakismanos@gmail.com

The connectivity of cortical networks often includes functional clusters of strongly interconnected, similarly tuned neurons. While these assemblies are known to include both excitatory and inhibitory neurons, there is a high variability of specificity and connectivity patterns of different neuron types. Although there are theoretical studies linking the presence of excitatory assemblies with the emergence of rich dynamics in spiking networks [1], the impact of neuron type-specific connectivity patterns on the dynamics and computational capabilities of recurrent networks remains poorly understood.

Here, we use mean-field theory to assess the impact of E and I assemblies of varying strengths on the dynamics of a balanced recurrent network. We demonstrate that in networks with sufficiently strong coupling, different levels of clustering among E and I neurons can control the distance from a chaotic transition. In particular, we show that relatively weak I assemblies combined with stronger E assemblies can maintain the network's dynamics at the edge of chaos.

We then evaluate how the topology-induced dynamics impact the computational capabilities of the recurrent network using a reservoir computing model. Specifically, we train our network on the complex task of simultaneously predicting multiple chaotic timeseries. We find that the performance can be optimized for any given average coupling strength (within a reasonable range) by adjusting the relative strength of E and I assemblies.

Finally, we use simulation-based inference [2] to evaluate the interaction between the specificity of E and I input, the relative strength of E and I assemblies, and the network performance. We find that in near-chaotic dynamical regimes, broader E than I input can stabilize activity and boost performance.

Our findings provide a description of how diverse connectivity between different neuron types can generate complex dynamics associated with superior performance in computational tasks. The connectivity patterns we identify largely coincide with experimentally observed structures in the mammalian cortex, suggesting a potential link between network topology, population dynamics, and the ability to perform challenging computations.

Acknowledgments

Supported by the Alexander von Humboldt Foundation, Joachim Herz Foundation, BMBF through the Tübingen AI Center and the Machine Learning Cluster of Excellence.

References

1. Litwin-Kumar, Doiron, 2012, *Nat Neurosci*, 15:1498–1505, 10.1038/nn.3220
2. Gonçalves, Lueckmann, *et al.*, 2020, *eLife*, 9:e56261, 10.7554/eLife.56261

NITRIC OXIDE MEDIATES DIFFERENTIAL EFFECTS IN MOUSE RETINAL GANGLION CELLS

Dominic Gonschorek[1,2,*], *Matías A. Goldin*[3], *Jonathan Oesterle*[1,2], *Tom Schwerd-Kleine*[1,2], *Ryan Arlinghaus*[1,2], *Zhijian Zhao*[2], *Timm Schubert*[1,2], *Olivier Marre*[3], *Thomas Euler*[1,2]

[1] CIN, University of Tübingen, Tübingen, Germany
[2] Institute for Ophthalmic Research, University of Tübingen, Tübingen, Germany
[3] Institut de la Vision, Sorbonne Université, INSERM, CNRS, Paris, France
[*] dominic.gonschorek@cin.uni-tuebingen.de

Nitric oxide (NO) is a prominent neuromodulator within various neuronal circuits that has been extensively studied in the retina across species. However, while NO has been associated with several aspects of retinal function such as light adaptation, gain control, and gap junction coupling, its effect on the retinal output remains poorly understood. In this study, we aim to discern which of the approximately 40 mouse retinal ganglion cell (RGC) types [1] are affected by NO to understand its functional role for vision.

We recorded two paired datasets of light-induced RGC responses using two-photon calcium imaging: (i) an NO dataset, where an NO-donor was added after the first recording ($n = 1838$), and (ii) a control dataset, where RGC responses were recorded twice without pharmacological manipulation ($n = 1590$). To identify type-specific effects, we employed an RGC type classifier for both datasets. Our findings revealed that RGCs exhibited heterogeneous changes in their responses not only in the NO, but also in the control dataset. This implies that any changes in the NO dataset may be obscured by non-drug induced adaptational, time-dependent effects. By comparing the effects in both datasets, we discovered that the majority of adaptational effects were type-specific. However, a specific RGC type, known for being suppressed-by-contrast (SbC), demonstrated significant differences in the NO dataset. Specifically, its activity increased during NO application, indicating that NO reduces the cell's inhibition by contrast. As this type was proposed to consist of more types [1], a functional clustering unveiled three distinct clusters with marginal adaptation but varying modulations in response to NO. As NO is associated with gap junction coupling, we expected NO-induced effects on receptive field structures. Interestingly, NO did not influence the spatial features of any tested RGC type, highlighting a highly type-specific modulatory effect. To validate our findings and further understand those neuromodulatory effects, we conducted multi-electrode array recordings on RGCs ($n = 391$). Preliminary results suggest a subdivision of the same SbC RGC type into three clusters, each displaying response modulations consistent with the calcium data.

In this study, we found that one particular neuromodulator exhibits a specific effect in a distinct group of SbC RGC types and that adaptational response changes should be considered when analyzing pharmacological effects. This highlights the importance of studying neuromodulators to understand their functional role in vision. Given that these effects already occur in the retina, they hold broader implications for downstream areas targeted by RGCs.

Reference

1. Baden, Berens, Franke, *et al.*, 2016, *Nature*, 529(7586):345–50, 10.1038/nature16468

VISUAL ATTENTION INCREASES SPIKE-LFP COHERENCE ACROSS CORTICAL LAYERS IN THE POSTERIOR PARIETAL CORTEX

Maureen A. Hagan[1,*], *Joanita F. D'Souza*[1], *Jessima M. Rich*[1], *Shaun L. Cloherty*[2], *Nicholas S. C. Price*[1]

[1] Department of Physiology, Monash University, Clayton, Australia
[2] School of Engineering, RMIT, Melbourne, Australia
[*] maureen.hagan@monash.edu

Brain areas communicate with temporal precision to control complex, cognitive behaviors like visual attention. One strategy the brain uses to keep track of inputs across areas is the laminar organization of the cortex. Inputs from different brain areas arrive at different cortical layers, yet we have little understanding of how individual cells keep track of different streams of information. One hypothesis is that the frequency components of the local field potential (LFP) may synchronize inputs from different streams of information [1].

To test whether coherence between spiking and the phase of the LFP changes across cortical layers, we recorded from the posterior parietal cortex of adult marmoset monkeys using semi-chronic 32-channel linear arrays (100ðİlJĞm spacing, NeuroNexus) that allowed us to recorded spiking and LFP activity across cortical layers. Marmosets were trained to perform a visually guided saccade task. The phase of spike-LFP coherence varied depending on frequency band (theta 4 Hz, alpha 12 Hz, beta 20 Hz, and gamma 40 Hz). Shifts in the phase of spike-LFP coherence have been shown to correspond to layer IV, the transition between superficial and deep layers [2]. Consistent with this, we found a significant shift in spike-LFP phase between superficial and deep layers ($p < 0.05$), random permutation test.

To test whether spike-LFP phase coherence is related to cognitive behavior, we trained the marmosets to do a cued-saccade task. On a subset of trials, a visual cue indicated the upcoming direction of the saccade with 85% accuracy. For cells with that were tuned to saccade directions near the cued direction (± 20 deg), there was a significant increase in spike-LFP coherence in the beta band (20 Hz) on cued trials compared to uncued trials ($p < 0.05$, random permutation test), and superficial cortical layers had higher spike-LFP coherence compared to deep cortical layers ($p < 0.05$). When the preferred saccade direction was far from the cued direction (±140 deg) there was a trend for spike-LFP coherence to be greater in the deeper cortical layers ($p = 0.05$). Therefore, spike-LFP coherence may be a mechanism for cells to support cognitive behavior. Furthermore, marmosets provide a unique animal model for dissecting the functional roles of cortical lamina.

Acknowledgments

Generously supported by the Australian Research Council and National Health and Medical Research Council.

References

1. Fries, 2015, *Neuron*, 88(1):220–235, 10.1016/j.neuron.2015.09.034
2. Davis, Dotson, *et al.*, 2023, *eLife*, 12:e84512, 10.7554/eLife.84512

MODULATION OF FUNCTIONAL CONNECTIVITY AND REPRESENTATIONS OF VALUE IN PRIMATE SINGLE-NEURONS USING A CLOSED-LOOP BCI FRAMEWORK

Dean Halperin[*], *Liliya Iskhakova, Rita Perets, Evgeniya Kornilova, Eyal Hizmi, Michal Israelashvili, Tamar Reitich-Stolero, Rony Paz*

Department of Brain Sciences, Weizmann Institute of Science, Rehovot, Israel
[*] dean-adam.halperin@weizmann.ac.il

Adaptive behavior and emotional modulation depend on the flexible assignment of value to stimuli, mediated by modulatory inputs from subcortical and cortical brain regions into the primate amygdala. Yet, the neural mechanisms that govern this plasticity remain poorly understood. The Substantia Innominata (SI) mediates plasticity via its cholinergic projections. However, much less is known about its role in the primate brain and how it facilitates value representations in the amygdala. We developed a novel framework utilizing a brain-computer interface (BCI) to guide functional connectivity between the SI and the amygdala of primates. Experimental sessions consisted of auditory classical conditioning containing aversive and appetitive stimuli, followed by a BCI block and an additional classical conditioning block (retention). We recorded single-neuron activity (more than 1000 neurons, and more than 100 sessions) simultaneously from the SI, amygdala, and another control region (dorsal anterior cingulate cortex, dACC). During BCI blocks, Auditory feedback was provided based on the strength and directionality of the cross-correlation between small groups of neurons, calculated in real-time and in a continuous manner. Performance in a session was measured by the increase in cross-correlation strength and its directionality, measured by the normalized proportion of amygdala spikes following either SI or ACC spikes within a pre-determined time window.

A significant proportion of SI→amygdala sessions were successful, whereas ACC→amygdala sessions were only marginally so. During successful sessions, there was an increase in the strength and the center of mass (CoM) of the cross-correlation, indicating that amygdala spikes followed SI spikes in a locked manner. This change persisted after the BCI block was over and into the retention block, suggesting that the change resulted from network plasticity and reorganization during the previous BCI block. Behaviorally, there was an increase in the conditioned response (CR, blinking) that follows an aversive CS, whereas a decrease occurred in the unconditioned response (UR) following the air puff (US). This behavioral change was specific to successful BCI sessions, and moreover, the behavioral response was correlated at a session level with the directionality of the cross-correlation during the BCI and during the following retention block.

In conclusion, our findings indicate that synchronization between two brain regions can be modulated using BCI and that we can modulate the strength and directionality of SI-to-amygdala networks. The network re-organization lasts beyond the BCI session itself and results in better adaptive behavioral responses to learned stimuli. We discuss the use of this tool to examine the constraints imposed by the network, the representations of value in the amygdala and SI, and the potential for restoration of adaptive behavior.

BEHAVIORAL BIASES FROM CELLULAR-RESOLUTION OPTOGENETICS IN A COMPLEX (ACCUMULATING EVIDENCE) TASK

Mark Ioffe[*], *Stephan Thiberge, Carlos Brody, David Tank*

Princeton Neuroscience Institute, Princeton University, Princeton, New Jersey, USA
[*]mioffe@princeton.edu

How many neurons does it take to change a mind? Optogenetic perturbations can affect animal behavior in many tasks, ranging from freezing to biases in decision-making, but the numbers of perturbed neurons in those experiments are typically large. Recently, a handful of cellular-resolution perturbation studies have shown small, statistically-significant behavioral effects in simple tasks. Does the small scale of the behavioral effects reflect a fundamental limit of what is possible with a small number of activated neurons, or could larger effects be obtained in a more complex task, where neural dynamics are more sensitive to perturbations?

Here, we simultaneously activated 5–120 choice-coding neurons in mice performing an accumulation-of-evidence task [1], with a (custom) cellular-resolution all-optical imaging and stimulation microscope. We found large behavioral effects from small perturbations in a complex task: animal choice was biased about 10%, dependent on the cellular-resolution details of the designed perturbations. Ongoing experiments are probing how these effects depend on functional connectivity and hemisphere-specificity in populations of excitatory neurons.

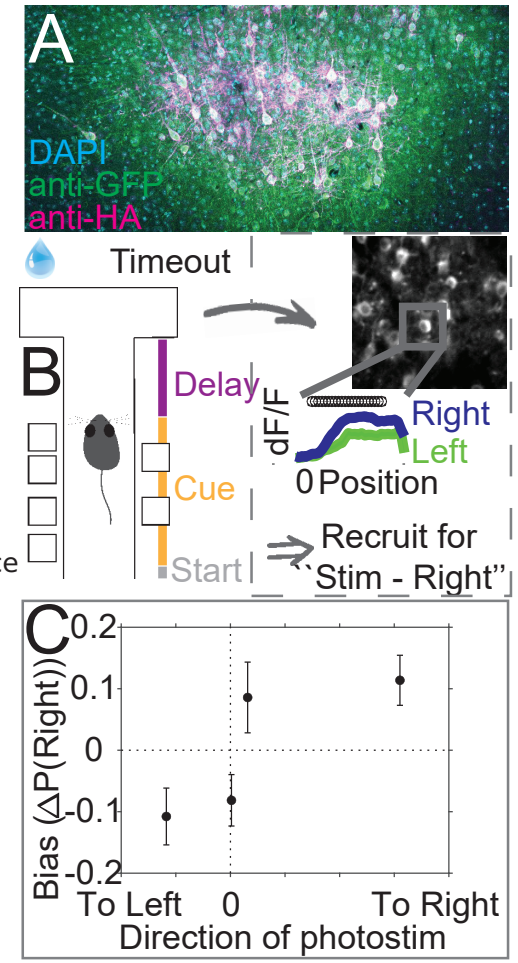

Figure 1. (**A**). Co-expression of GCaMP (green) and opsin (magenta) in retrosplenial cortex. (**B**). Neurons are recruited into population photostimulation patterns based on choice coding in the task [1]. (**C**) Behavioral biases (7 mice, 14 sessions) are elicited when neural activity is perturbed along the choice coding direction. Calculated similarly to Daie and colleagues [2].

Acknowledgments

Supported by NIH grant 5U19NS132720.

References

1. Pinto, Koay, *et al.*, 2018, *Front Behav Neurosci*, 12(36):1–19, 10.3389/fnbeh.2018.00036
2. Daie, *et al.*, 2021, *Nat Neuro*, 24:259–265, 10.1038/s41593-020-00776-3

DIFFERENCES IN NEURAL ENSEMBLES BETWEEN IMAGINED AND NATIVE LIMB MOVEMENT RAISE NEW CHALLENGES FOR BRAIN-COMPUTER INTERFACES

Samantha N. Johnson[1], *Milan Rybar*[2], *Viraj Bodiwala*[3], *Linnea Endsley*[2], *John E. Downey*[2], *Nicholas G. Hatsopoulos*[2,*]

[1] Committee on Computational Neuroscience, University of Chicago, Chicago, Illinois, USA
[2] Department of Organismal Biology and Anatomy, University of Chicago, Chicago, Illinois, USA
[3] The College, University of Chicago, Chicago, Illinois, USA
[*] nicho@uchicago.edu

The involvement of the motor cortex in controlling both native limb and imagined movements is well-established. However, the extent of overlap between the neural representations and dynamics of motor control in these two contexts is unknown.

This question holds particular significance in the context of brain-computer interfaces (BCIs), which decode neural signals to control output devices like dexterous bionic hands for individuals with limited motor control. For BCI users lacking motor control, the standard practice involves training a BCI decoder using neural activity while the subjects observe and imagine performing movements presented to them. Therefore, significant differences in neural ensembles between observation-based imagined and native limb movements could have important implications for the BCI community.

In this study, we attempt to explore the similarity between these two contexts through a decoding approach. We took advantage of a fortuitous opportunity to work with two subjects implanted with chronic electrode arrays in the precentral gyrus, each of whom had residual movement in the upper-limb following incomplete spinal cord injury. Center-out reaching movements were performed under three conditions: imagined reaches while observing a virtual arm completing the task, native limb reaches while observing the virtual arm make the same movements, and native limb movements without observing the virtual arm.

Kalman filter decoders were trained on two-dimensional velocity vectors in each of these three conditions and tested across different conditions. While the decoder trained on the imagined condition effectively decoded the velocity of the virtual arm, it could not generalize to decoding native limb movement in either of the two native limb conditions. Additionally, decoders trained on native limb conditions performed well within those conditions, but neither could generalize to decode imagined movements with precision.

To explore the population dynamics underlying these results, we conducted a dynamical systems analysis and discovered a fundamental difference in the rotational dynamics of the imagined and native limb movement conditions.

Consequently, we assert that the difficulty of generalizable decoders and the contrasting population dynamics reveal important differences in neural representation between native and imagined movements. These results prompt careful consideration in the application of BCI technologies.

Acknowledgments

This work has been supported by NIH Grant UH3NS107714.

ORDER FROM CHAOS: INTERPLAY OF DEVELOPMENT AND LEARNING IN RECURRENT NETWORKS OF STRUCTURED NEURONS

Laura Kriener[*], *Kristin Völk*, *Federico Benitez*, *Ben von Hünerbein*, *Walter Senn*, *Mihai A. Petrovici*

Department of Physiology, University of Bern, Bern, Switzerland
[*] laura.kriener@unibe.ch

Behavior can be described as a temporal sequence of actions driven by neural activity. To learn complex sequential patterns in neural networks, memories of past activities need to persist on significantly longer timescales than relaxation times of single-neuron activity. While recurrent networks can produce such long transients, training these networks is a challenge.

One approach has been reservoir computing, where only weights from a recurrent network to a readout are learned [1]. Other models achieve learning of recurrent synaptic weights using propagated errors. However, their biological plausibility suffers from issues with locality (BPTT [2]), resource allocation (RTRL [3]), or parameter scales (E-PROP [4]).

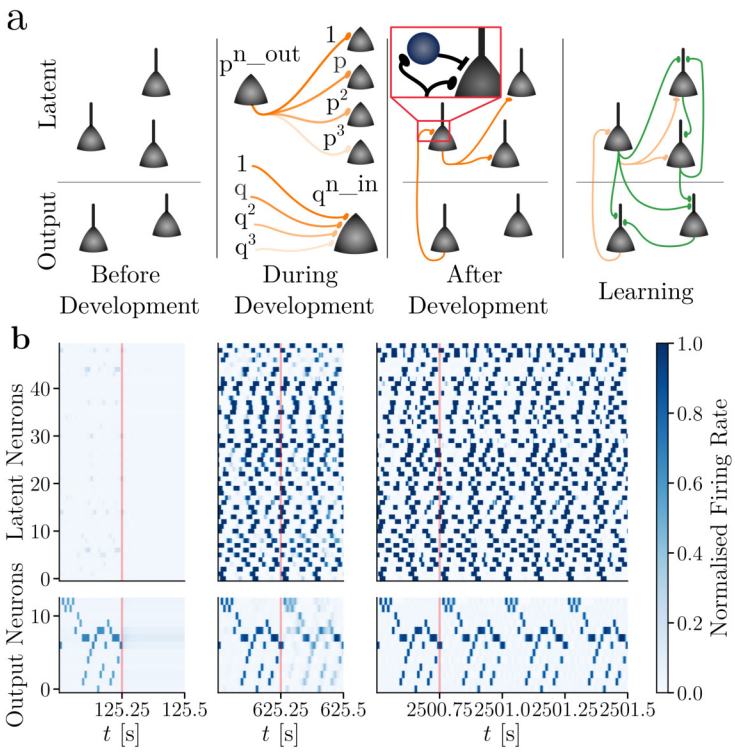

Figure 1. (**a**) Development and learning in a two-population model of latent and output neurons. (**b**) Evolution of network activity during learning.

We suggest that many of these issues can be alleviated by considering dendritic information storage and computation. By applying a fully local, always-on plasticity rule we are able to learn complex sequences in a recurrent network comprised of two populations (latent & output, Figure 1a). Importantly, our model is resource-efficient, enabling learning of complex sequences with only a small number of neurons.

We demonstrate these features in a mock-up of birdsong learning, in which our networks first learn a long, non-Markovian sequence (a sample of Beethoven's "Für Elise", Figure 1b) that they can then reproduce robustly despite external disturbances.

References

1. Jäger, Lukosevicius, 2009, *Comp Sci Rev*, 3(3):127–149, 10.1016/j.cosrev.2009.03.005
2. Werbos, 1990, *Proc IEEE*, 78(10):1550–1560, 10.1109/5.58337
3. Williams, Zipser, 1989, *Connection Sci*, 1(1):87–111, 10.1162/neco.1989.1.2.270
4. Bellec, Scherr, *et al.*, 2020, *Nat Comm*, 11(1):3625, 10.1038/s41467-020-17236-y

DENDRITIC PROPERTIES SHAPE THE ENCODING OF BEHAVIORALLY RELEVANT STIMULI IN COLLICULAR WIDE-FIELD NEURONS

Norma K. Kühn[1,*], *Chen Li*[2], *Bram Nuttin*[1], *Natalia Baimacheva*[1], *Katja Reinhard*[3], *Vincent Bonin*[1], *Karl Farrow*[1]

[1] Neuro-Electronics Research Flanders (NERF), Leuven, Belgium
[2] Yale School of Medicine, New Haven, Connecticut, USA
[3] SISSA, Trieste, Italy
[*] norma.kuehn@gmail.com

Dendrites are important for filtering relevant information along sensory-motor pathways. To fully comprehend their contribution, it is crucial to identify their inputs and outputs and the distinctive features of the circuit. Wide-field neurons of the mouse superior colliculus are genetically targetable, receive direct input from the retina, and trigger innate defensive and appetitive behaviors (Fig. 1A). We used a three-step approach to identify the rules by which these neurons filter their inputs to guide behavior. First, we measured their responses to relevant visual cues in dendrites and cell bodies. Second, we examined inputs from the retina and collicular inhibitory interneurons using a combination of transsynaptic virus tracing and two-photon calcium imaging. Third, to replicate wide-field neuron responses and identify important physiological parameters, we tested computational models of increasing dendritic complexity (Fig. 1B).

We made three important findings: First, twelve retinal ganglion cell types provide layered input to wide-field neuron dendrites. Second, wide-field neuron cell bodies, known to respond to expanding discs, turn out to also respond strongly to shrinking discs and cluster into three types based on contrast preference: *on*, *off*, and *on-off*. While an expanding disc is commonly associated with looming overhead predators, a shrinking disc might emulate receding prey. Third, a nonlinear point neuron model is not sufficient to capture wide-field neuron responses to the receding prey stimulus. Instead, dendritic filtering properties and the location of inputs along wide-field neuron dendrites need to be explicitly modeled to accurately capture wide-field cell body responses (Fig. 1B). This demonstrates that we need to include dendrites into our models of the brain, otherwise responses to important stimuli will be missed and a neuron's function within a circuit misinterpreted.

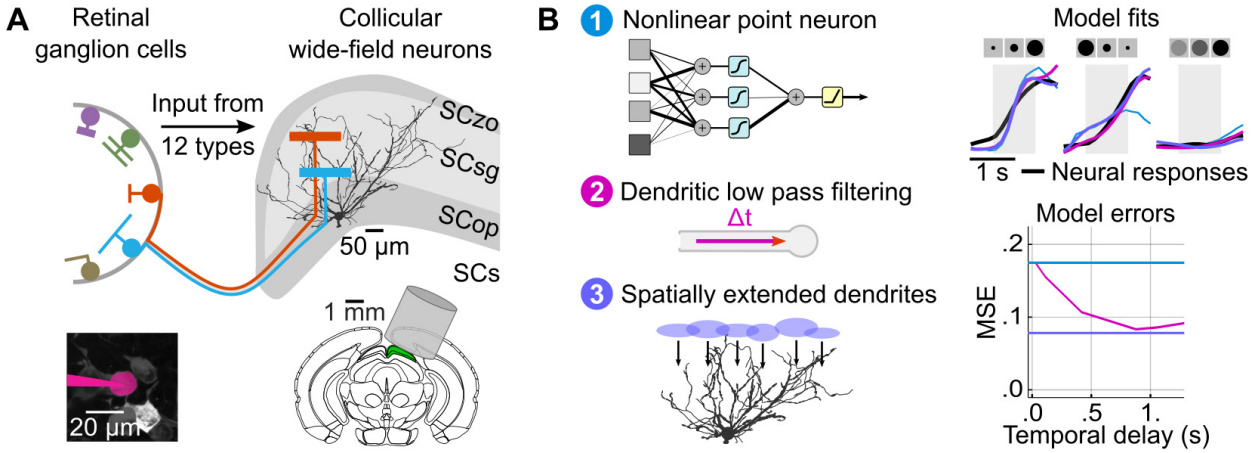

Figure 1. (**A**) Measuring input and output of collicular wide-field neurons. (**B**) Applied models with increasing dendritic complexity (left), average fitted responses and errors (right).

AN INFORMATION-THEORETICAL APPROACH TO OPTIMIZE TASK DESIGN FOR DISTINGUISHING PROBABILISTIC CODES IN NEURAL POPULATIONS

Po-Chen Kuo[*], *Edgar Y. Walker*

Department of Physiology and Biophysics, University of Washington, Seattle, USA
[*]pckuo@uw.edu

Bayesian inference provides a framework to understand perceptual decision-making. Two competing theories of probabilistic representation, probabilistic population code and neural sampling code, posit that sensory population responses r encode likelihood functions $p(r|\Theta)$ or posterior distributions $p(\Theta|r)$ over the stimulus Θ (*e.g.*, orientation), respectively. The key distinction lies in if priors over the stimulus modulate population responses. However, distinguishing the two codes experimentally has been challenging as it is unclear what stimulus distributions would maximally differentiate populations that follow likelihood and posterior codes.

To address this challenge, we derived an information-theoretical measure to assess the expected difference in decodable information when applying decoders to decode either the likelihood (likelihood decoder) or the posterior (posterior decoder) (Fig. 1A). The decoders are implemented using deep neural networks following our recent work [1]. We show the expected difference in decoder performance can be evaluated as the information gap when decoding mismatched probabilistic quantity, *i.e.*, the KL divergence between the true posterior and a posterior marginalized over different stimulus distributions $\Delta^* = D_{KL}(p(\Theta|r) \,||\, \hat{p}(\Theta|r))$.

This measure helps assess how effective a task design can differentiate between likelihood and posterior codes. We show on simulated likelihood-coding and posterior-coding populations, the difference in likelihood and posterior decoders matched the expected information gap (Fig. 1B, D top). Critically, maximizing the expected information gap renders stimulus distributions that maximally differentiate between the two codes (Fig. 1C, D bottom). This information-theoretical approach drives targeted task designs for identifying the probabilistic neural code.

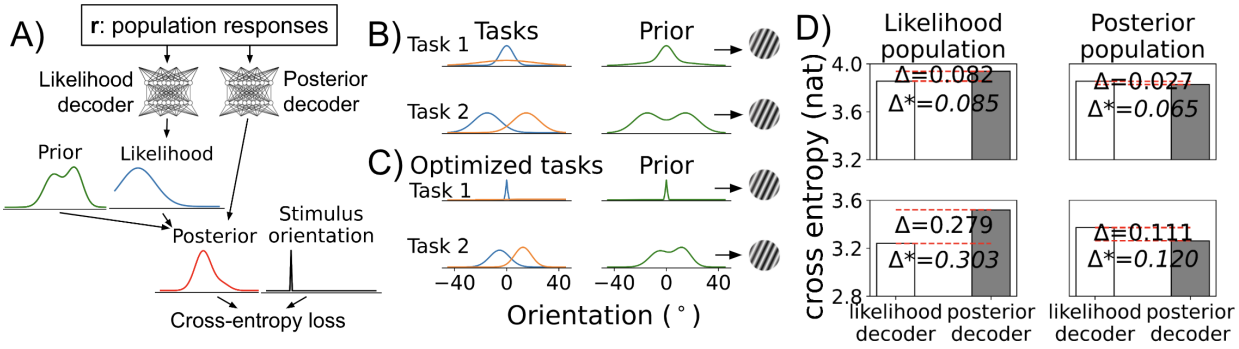

Figure 1. (**A**) Likelihood decoders and posterior decoders. (**B**) Two probabilistic visual categorization tasks with different priors. (**C**) Optimized task design by maximizing the expected information gap. (**D**) Difference in decoding performance of likelihood and posterior decoders (Δ) on simulated likelihood and posterior coding populations can distinguish the underlying population code. Δ^* denotes the expected information gap. Top for tasks in B, bottom for tasks in D.

Reference

1. Walker, Cotton, Ma, *et al.*, 2020, *Nat Neurosci*, 23(1):122-129, 10.1038/s41593-019-0554-5

CORTICAL REPLAY OF TEMPORAL MEMORY TRIGGERED BY INTER- AND INTRA-HIPPOCAMPAL RIPPLES

Sze Chai Kwok[1,*], *Xuanlong Zhu*[1], *Diogo Santos-Pata*[1], *Hongjie Jiang*[2], *Shaomin Zhang*[2], *Zhaoxin Wang*[3]

[1] Duke Kunshan University, Kunshan, China
[2] Zhejiang University, Hangzhou, China
[3] East China Normal University, Shanghai, China
[*] sk695@duke.edu

Replay is instrumental for episodic memory retrieval and consolidation. In replay, patterns of neural activity recapitulate past or possibly predict future trajectories. While evidence of neural replay for the *where* (location) and *what* (object) sequence are known in rodents and humans, the replay of the *when* remain understudied. We have 18 epileptic human volunteers (Fig. 1a) participate in a temporal order judgment task in which they recalled the order of events in videos they had watched.

We acquired task-based intracranial electroencephalography data and combined a replay detection framework and non-linear classifiers to differentiate neural representations of content-invariant temporal order within whole-brain space. We found evidence in the magnitude of the temporal dependence among reactivation levels for rapid replay reenacting the order of previously experienced video segments during memory retrieval (sequenceness: 0.014 vs. permutation threshold: 0.013) and subsequent offline resting period (sequenceness: 0.012 vs. chance level: 0.011).

We then asked whether the intrinsic temporal structure of the replay is manifested in relation to hippocampal sharp-wave ripples (SWRs). By taking the reactivation at the peak of each SWRs, we found replay evidence across multiple SWRs (*i.e.*, inter-SWR replay score: 0.21, $p < 0.001$) as the replayed segments are chained across time (Fig. 1b). We confirmed replay happen within each of these SWRs (*i.e.*, intra-SWR replay score: 0.09, $p < 0.001$) (Fig. 1c). Our findings suggest that temporal information in episodic memory is retrieved via a hippocampal SWRs triggered whole cortex replay mechanism.

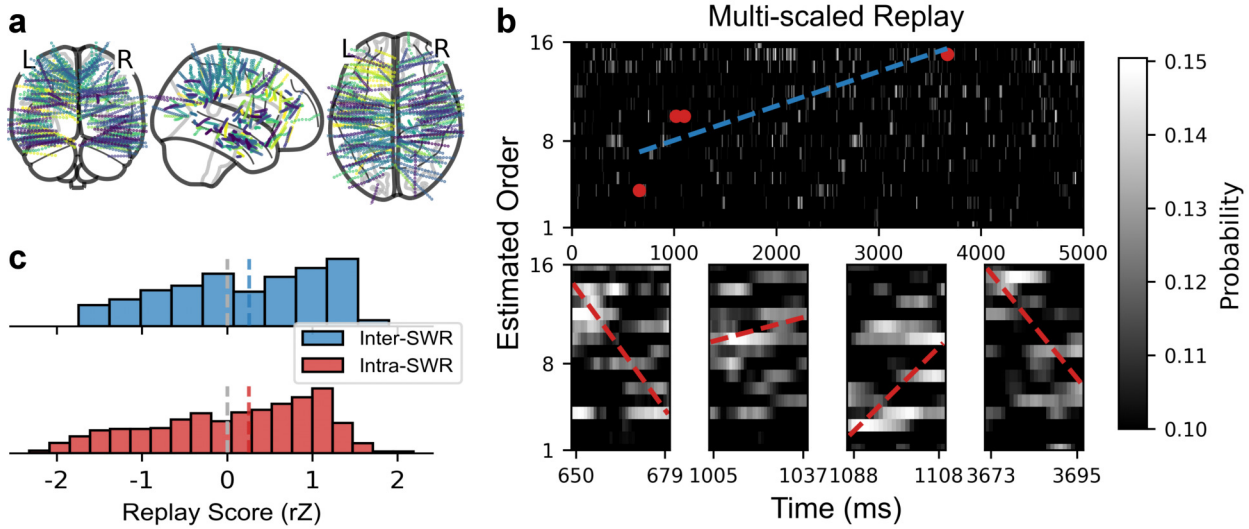

Figure 1. (**a**) 2492 recording sites from 18 patients. (**b**) Replay in a single trial (upper) with four hippocampal SWRs (bottom). (**c**) Significant statistics for intra- and inter-SWR replay.

NEURAL REPRESENTATIONS OF TEXTURE IN THE MACAQUE VENTRAL VISUAL STREAM AND THEIR RELATION TO PERCEPTION

Gerick M. Lee[*], *C. L. Rodriguez-Deliz, Brittany N. Bushnell, Najib J. Majaj, J. Anthony Movshon, Lynne Kiorpes*

Center for Neural Science, New York University, New York, USA
[*]gerick@cns.nyu.edu

Selective neuronal activity in mid-level ventral visual cortical areas, like V2 and V4, represents information about visual texture. This selectivity is thought to contribute to the representation of visual objects, an ability linked with inferotemporal cortex (IT). How the representation of textures in V2 and V4 relates to responses in IT remains unclear.

To better understand the relationship between these areas, we recorded multiunit neural activity from areas V1, V2, V4, and IT in response to visual texture images matched in local spectral content, but differing in statistical similarity to natural images. These naturalistic statistics are known to modulate activity in V2 and V4 [1, 2], but not in V1. We used textures from 35 families generated from 35 different original natural images that have been used in previous psychophysical experiments [1].

As described previously [1, 2], naturalistic texture modulation was weak in V1, stronger in V2, and stronger still in V4. These textures have not previously been used to study IT — selectivity was weaker there than in V4, and similar to V2. We also correlated neural sensitivities with human perceptual sensitivities for the 35 different texture families. These correlations were weak in V1, stronger in V2, stronger still in V4 and IT. Therefore, despite the lower overall sensitivity of IT to naturalistic texture, its correlation with human perceptual sensitivity was as high as in V4.

These similarities between V4 and IT in how texture information is represented preserves their mutual correlation with perception. This might reflect a process which first measures texture-like statistics from a scene, before using that signal as a platform from which features like object identity can be extracted. The shift in how texture information is represented in V4 and IT indicates that the signals in those areas are more plausibly directly involved in perceptual judgements than are the signals in earlier areas.

References

1. Freeman, Ziemba, *et al.*, *Nat Neurosci*, 16(7):974–981, 10.1038/nn.3402
2. Okazawa, *et al.*, *Cereb Cortex*, 27(10):4867–4880, 10.1093/cercor/bhw282

RESPONSES OF NEURAL POPULATIONS IN MACAQUE V4 TO OBJECT AND TEXTURE IMAGES

Justin D. Lieber[1,*], *Timothy D. Oleskiw*[1,2], *Laura Palmieri*[1], *Eero P. Simoncelli*[1,2], *J. Anthony Movshon*[1]

[1] Center for Neural Science, New York University, New York, New York, USA
[2] Flatiron Institute, New York, New York, USA
[*] justinlieber@nyu.edu

Humans and monkeys can effortlessly recognize objects in natural scenes. This ability relies on neural computations in the ventral stream of visual cortex. The intermediate computations that lead to object selectivity are not well understood, but previous studies implicate V4 as an early site of selectivity for object shape. To explore the mechanisms of this selectivity, we transformed two sets of images from photographs to create a continuum of images that span the space between the original images and scrambled textures. This was achieved by using the Portilla and Simoncelli texture synthesis procedure, which preserves the local statistics of the original image while discarding information about scene and shape. To create a continuum of images (an *image family*) that smoothly varies between fully scrambled textures and natural images, we varied the size of scrambling regions from small localized regions to the whole image. For all sizes, these scrambling regions seamlessly covered the whole image, with modest overlap.

Using single electrodes, linear multielectrode arrays, and chronically implanted multielectrode arrays, we measured the responses of both well-isolated single units and multi-unit channels in awake macaque V4 to these scrambled images. On average, V4 neurons were slightly more active in response to the original photographs than to their scrambled counterparts. However, responses in V4 varied widely both across different cells and different sets of images. An important determinant of this variation was the effectiveness of image families at driving strong neural responses. Across the full V4 population, a cell's average evoked firing rate for a family reliably predicts that family's preference for natural over scrambled images. Accordingly, the cells that respond most strongly to each image family showed a much stronger difference between natural and scrambled images and a graded level of modulation for images of intermediate pooling sizes. This preference for natural images was delayed until about 50 ms after the onset of neuronal activity and did not peak in strength until 130 ms after activity onset.

Finally, V4 neural responses strongly separated natural images from all partial and full scrambling conditions, despite the fact that the least scrambled images in our set appear similar to the original natural images. We hypothesized that this separation might be better explained by the exquisite sensitivity of observers to minor degradations in the structure of natural images. To test this, we analyzed our image set with the Deep Image Structure and Texture Similarity metric (DISTS), an image-computable similarity measurement that predicts human judgements of image degradation. Distances measured with DISTS also showed a categorical separation of natural images from all scrambling conditions, and predicted distances measured from V4 neural responses better than simpler metrics like image pixel distance or the Structural Similarity Index Measure (SSIM). This suggests that V4 responses are highly sensitive to small deviations from natural image structure.

ALPHA OSCILLATIONS DRIVE ALTERNATION OF INHIBITION AND EXCITATION IN THE AWAKE RESTING STATE

Fabrizio Lombardi[1,*], *Hans J. Herrmann*[2], *Liborio Parrino*[3], *Dietmar Plenz*[4], *Silvia Scarpetta*[5], *Anna E. Vaudano*[3], *Lucilla de Arcangelis*[6], *Oren Shriki*[7]

[1] Department of Biomedical Sciences, University of Padova, Padova, Italy
[2] PMMH, ESPCI, Paris, France
[3] Department of Neuroscience, University of Parma, Parma, Italy
[4] NIH, Bethesda, Maryland, USA
[5] Department of Physics, University of Salerno, Salerno, Italy
[6] Department of Physics, University of Campania, Caserta, Italy
[7] Ben-Gurion University of the Negev, Beer-sheva, Israel
[*] fabrizio.lombardi@unipd.it

Alpha oscillations are considered to reflect cortical inhibition. However, the nature of this inhibition, as well as its effects on collective neural dynamics and information processing, is still poorly understood. The pulsed inhibition hypothesis states that alpha oscillations mediate inhibition in rhythmic pulses, a mechanism that might allow for selective information processing during the periods of relative excitation [1]. Verifying directly this hypothesis requires simultaneous multiscale recordings of neural activity. However, we argue that it should be possible to identify the hallmarks of alpha-mediated, pulsed inhibition through the analysis of neural activity that is readily accessible. To this end, we note that, within the pulsed inhibition hypothesis, alternating states of inhibition and excitation should result in periods of attenuation and enhancement of collective neural activity modulated by alpha oscillations.

To verify this assumption and clarify the functional role of alpha, we study collective neural activity during resting wake (MEG/EEG) and NREM sleep (EEG), a state with marginal presence of alpha. We show that, during resting wake, alpha drives an alternation of attenuation and amplification bouts of neural activity. Our analysis indicates that inhibition is activated in pulses that last a single alpha cycle and gradually suppress neural activity, while excitation is successively enhanced over a few alpha cycles to amplify neural activity. Furthermore, we show that long-term alpha amplitude fluctuations—known as the *waxing* and *waning* phenomenon—are associated with an attenuation-amplification mechanism acting over the timescales of several seconds and described by a power law decay of the activity rate in the waning phase. Importantly, we do not observe such dynamics during NREM.

The results suggest that alpha oscillations control the alternation of inhibition and excitation bouts across multiple timescales, the waxing and waning being a long-term control mechanism of cortical excitability—the high-amplitude alpha bursts being main inhibitory events of this mechanism. The amplification regime observed beyond the timescales of the individual alpha cycle indicates in turn that alpha oscillations might modulate the intensity of neural activity not only through pulses of inhibition, as proposed in the pulsed inhibition hypothesis, but also by timely enhancing excitation (or disinhibition).

Reference

1. Klimesch, 2012, *Trends Cogn Sci*, 16:606–617, 10.1016/j.tics.2012.10.007

LONG RANGE CORTICAL INTERACTIONS DURING COMPARISON OF SENSORY AND COGNITIVE INFORMATION

Camila Losada[*], *Alexis Monnet-Aimard, Guilhem Ibos*[**]

Institut de Neurosciences de la Timone, CNRS, Aix-Marseille Univ., Marseille, France
[*] camila.losada-gomez@univ-amu.fr
[**] guilhem.ibos@univ-amu.fr

Decision-making processes are complex and usually require taking into account large amount of information. For example, looking for a friend in a crowd involves comparing visual information from the environment to cognitive representation of what we are looking for. This project aims at understanding how visual cortex (V4-encoding visual information) and fronto-parietal networks (LIP and PFC-encoding working memory related signals) interact during comparing sensory working-memory information. We specifically test the hypothesis that LIP integrates, compares and transform sensory and cognitive-related information into decision-related signals [1].

We recorded the activity of approximately 2000 V4, 2000 PFC and 500 LIP neurons of one macaque monkey performing a delayed conjunction-matching task. In this task, one sample stimulus (composed of one of two colors, one of two orientations, either inside or outside the RF of the recorded neurons) is followed by a variable delay and a succession of test 1 to 5 stimuli (composed of one of eight color and one of eight orientations each). Monkey's task is to match the location, color and orientation of sample and test stimuli. In a subset of trials, neutral sample stimuli cue monkeys to wait passively (no engagement in active vision).

Preliminary data show a clear dissociation between sensory and cognitive/memory-related signals in two overlapping networks. Selectivity to sensory-related information emerge primarily and preferentially in V4 and LIP networks. Cognitive signals emerge preferentially within PFC networks and seem to be leaking within LIP. Applying SVM linear classifiers to PFC data allowed us to decode the identity of sample stimuli all along the trials (sample presentation, delay and test presentation) from the activity of PFC neurons. However, this encoding seems extremely dynamic and dependent of the task epoch. Using similar methods, we could decode sample's identity from V4 neurons only during stimulus presentation (sample and test presentation), suggesting sensory encoding modulated by selective attention. Interestingly, LIP neurons showed strong encoding of sensory information during sample presentation as well as limited but significant encoding of sample's identity during the delay.

These preliminary results confirm LIP's central position in integrating and comparing bottom-up (from V4) and top-down (from PFC) information. However, additional analyses are required. We will use encoding models (*e.g.*, GLMs) to investigate how tasks events modulate neuronal processing in each area and quantify the information transfer withing and between populations.

Reference

1. Freedman, Ibos, 2018, *Neuron*, 97:1219–1234, 10.1016/j.neuron.2018.01.044

ESTIMATING LOCAL FIELD POTENTIALS FROM PRESYNAPTIC FIRING RATES

Nicolò Meneghetti[*], *Alberto Mazzoni*

The Biorobotics Institute, Sant'Anna School of Advanced Studies, Pisa, Italy
[*] nicolo.meneghetti@santannapisa.it

The intricate dynamics of neural networks in the mammalian visual cortex are pivotal for decoding external stimuli. While understanding interactions among individual neuronal families poses a challenge, computational models serve as valuable tools in simulating and decoding sensory information processing. Computational models, however, can only be properly tested against experimental evidence when they can quantitatively predict empirical measures of local activity, such as local field potentials (LFPs). LFPs estimation require however multi-compartmental neuronal networks, whose computational demands, especially in large-scale networks, can limit their widespread use for predicting extracellular signals.

In this study, we employed a previously developed linearized framework [1] to estimate LFPs of a detailed multi-compartmental network of mouse primary visual cortex (V1) [2, 3]. Briefly, LFPs estimates are obtained by convolution of firing rates with appropriate spatiotemporal kernels, representing postsynaptic spike-signal impulse responses averaged over pairs of pre- and postsynaptic populations. The linearized framework effectively captured LFPs generation in the layer 2/3 of the network, provided non-stationarity of membrane potentials and firing dynamics were considered.

Our findings highlighted that upper layers LFPs are primarily influenced by external sources, with over 85% of variance explained by synaptic activity from thalamic afferents, feedback from higher visual cortical areas, and background activity (simulating inputs from surrounding cortical areas). Intracortical sources, including one excitatory and three different interneuron families, had a notably lower impact, with parvalbumin-positive interneurons emerging as the most significant intracortical contributors.

Extending our exploration, we examined computing LFPs using a point-neuron leaky-integrate-and-fire network alongside the linear approach. Our results showed that a linear combination of synaptic currents could effectively capture LFPs dynamics. In conclusion, our study provides insights into how multiple neuronal families interact to shape extracellular potentials in mice V1 upper layers. This not only advances our understanding of neural processing but also facilitates broader accessibility to computationally efficient tools for addressing extracellular signal predictions within the scientific community.

Acknowledgments

This work was supported by the Italian Ministry of University and Research (MIUR) PRIN2017, PROTECTION, Project 20178L7WRS.

References

1. Hagen, *et al.*, 2022, *PLoS Comput Biol*, 18(8):e1010353, 10.1371/journal.pcbi.1010353
2. Rimehaug, Stasik, Hagen, *et al.*, 2023, *eLife*, 12:e87169, 10.7554/eLife.87169
3. Billeh, Cai, Gratiy, *et al.*, 2020, *Neuron*, 106(3):388–403, 10.1016/j.neuron.2020.01.040

DYNAMICAL ANALYSIS OF NEURONAL POPULATION SPIKING ACTIVITY VIA DIFFUSION APPROXIMATION

Nosratullah Mohammadi[1,*], *Wilson Truccolo*[2,3], *Anne-Lise Giraud*[1,4], *Timothée Proix*[1]

[1] Department of Basic Neurosciences, University of Geneva, Geneva, Switzerland
[2] Department of Neuroscience, Brown University, Providence, Rhode Island, USA
[3] Carney Institute for Brain Sciences, Brown University, Providence, Rhode Island, USA
[4] Institut Pasteur, Université Paris Cité, Paris, France
[*] nosratullah.mohammadi@unige.ch

Recent advancements in brain recording techniques have enabled to simultaneously record spiking activity from thousands of individual neurons. Modeling and interpreting such data is key to understanding the neural dynamics governing brain functions. A powerful model to relate spiking activity to sensory and motor functions is the nonlinear Hawkes process [1, 2]. However, the stochastic and non-Markovian nature of Hawkes processes makes their direct analyses challenging.

Here, we employed a diffusion approximation [1, 3] to transform Hawkes processes with exponential nonlinearity into a system of coupled stochastic differential equations. We then studied the dynamics and bifurcations of the deterministic part of the approximated system using tools from dynamical systems theory. For a single neuron, fixed points are determined by a closed-form solution called Lambert-w function, offering an intuitive understanding of their long-term behavior and bifurcations. The full stochastic system can be approximated with moment closure methods [3], which crucially reconcile discrepancies between previous analysis by explaining unstable behavior in bistable regimes [4]. Extending this approach to multiple neurons, we analytically determined the fixed points and their stabilities from the deterministic part of the neuronal network equations. We found that specifying auto- and cross-history coupling kernels allows setting the network in flexible regimes with rich dynamics, as indicated by the number of fixed points. We further showed that time scale separation allows studying the low-dimensional coordinations of neural dynamics.

We applied this method to spiking neural activity recorded from the human superior temporal gyrus during speech processing. Goodness-of-fit was validated with Kolmogorov-Smirnov statistical test. We demonstrated that the dynamical properties of each fixed point could be monitored throughout the behavioral task. This revealed transient changes in the qualitative nature of these fixed points through saddle-node bifurcation as speech envelope increases. This work introduces a novel and promising approach for analyzing the low-dimensional dynamics emerging from population spiking activity and relating it to behavioral factors.

Acknowledgments

This work was funded by Swiss National Science Foundation career grant 193542.

References

1. Truccolo, 2016, *J Physiology (Paris)*, 10.1016/j.jphysparis.2017.02.004
2. Pillow, *et al.*, 2008, *Nature*, 10.1038/nature07140
3. Rule, Sanguinetti, 2018, *Neural Comput*, 10.1162/neco_a_0112
4. Gerhard, *et al.*, 2017, *PLoS Comput Biol*, 10.1371/journal.pcbi.1005390

QUANTIFYING THE DIVERSE CONTRIBUTIONS OF HIERARCHICAL MUSCLE INTERACTIONS TO MOTOR FUNCTION

David Ó'Reilly[1,*], *William Shaw*[1], *Pauline Hilt*[2], *Rafael de Castro Aguiar*[1], *Sarah L. Astill*[1], *Ioannis Delis*[1]

[1] School of Biomedical sciences, University of Leeds, Leeds, United Kingdom
[2] Université Bourgogne Franche-Comté, Dijon, France
[*] bsdor@leeds.ac.uk

The synergy concept is a pervasive phenomenon describing the interactions underpinning the self-organisation of living systems. Within the motor neurosciences, the muscle synergy approach applies this concept to the problem of coordinating the many degrees-of-freedom of the human body. It suggests that the human motor system is organised into functional modules comprised of muscles working together towards common task goals. However, recent innovative work has added further nuance to this idea [1], showing how muscles may also work together towards functionally different and independent task-goals, representing crucial attributes of flexible motor behaviour. Here we probed this proposed functional neural architecture by building upon an established theoretical framework to characterise distinct types of muscle interactions, *i.e.*, functionally similar (redundant), complementary (synergistic), and independent (unique), across scales.

Building on traditional approaches and our recent work [2, 3], we introduced the Partial Information Decomposition (PID) approach to the motor control field. We integrated this approach into our computational framework which leverages network- and information-theoretic statistical tools along with machine-learning to extract low-dimensional motor components from human movement data.

Through a novel application of the PID approach to large-scale muscle activations, we unveiled complex networks of inter- and intra-muscular interactions with distinct functional roles as well as independent muscle contributions to task performance. We showcased the effectiveness of this approach by extracting hierarchical and functionally diverse motor components that were (i) generalisable across participants and tasks and (ii) predictive of balance performance across trials and of differences in motor variability between young and older adults.

In aligning muscle synergy analysis with the forefront of understanding on human movement modularity [1], our findings suggest the proposed methodology can offer novel biological insights into the neural control of movement and research opportunities towards health and engineering applications.

References

1. Hug, Avrillon, Ibáñez, Farina, 2023, *J Physiol*, 601(1):11–20, 10.1113/JP283698
2. Ó'Reilly, Delis, 2022, *J Neur Eng*, 19(1):016031, 10.1088/1741-2552/ac5150
3. Ó'Reilly, Delis, 2023, *eLife*, 12:RP87651, 10.7554/eLife.87651.4

HUBEL AND WIESEL AFTER DARK: EXPLORING THE ORIGIN OF CORTICAL ACTIVITY ACROSS ADAPTATION STATES

Ronan T. O'Shea[1,*], *Xue-Xin Wei*[2], *Nicholas J. Priebe*[2]

[1] Department of Psychology, University of Texas at Austin, Austin, Texas, USA
[2] Institute for Neuroscience, University of Texas at Austin, Austin, Texas, USA
[*] ronan_oshea@utexas.edu

Each day, we traverse disparate sensory environments- from hot midday sun to cool starlit night, loud barrooms and silent hallways, all the while maintaining sensitivity to the world around us. This feat is accomplished by rapid adaptation of neural responses to changes in environmental statistics. At the sensory periphery, adaptation alters neural receptive fields and correlated activity between neurons. Less is known about how adaptation at the periphery affects encoding in sensory cortex. Here, we focus on luminance adaptation in the visual system, which has well-defined effects on the output of the retina in terms of shifts in both spatiotemporal tuning and correlated variability. We describe how the shift from rod (scotopic) to cone (photopic) vision affects encoding by neural populations in both the lateral geniculate nucleus (LGN) and primary visual cortex (V1) of the mouse.

Using the neuropixels probe to densely sample LGN activity across the retinotopic map, we find that receptive fields shift with luminance in a manner consistent with the known changes in the retina. Namely, LGN neurons become more high-pass in the spatial and temporal frequency tuning under photopic conditions. In addition, we find that the noise correlations between nearby LGN cell pairs are larger in the scotopic than photopic condition, in line with prior results from the retina [1]. Taken together, these results suggest that shifts in both the signal and noise of the LGN population code could substantially impact the encoding of visual stimuli across light levels.

Remarkably, we find that the functional signatures of light adaptation state in retinal and geniculate populations do not persist in V1. We measure more stable receptive field properties across light levels than in LGN. In addition, using multi-electrode arrays and 2-photon calcium imaging to sample V1 activity across the retinotopic map, we find that noise correlations are invariant to luminance. In line with these results showing invariance in both the signal and noise of the V1 population code, we find that V1 encodes fine orientation differences equally well across light levels.

We show that our data are consistent with a model of the visual hierarchy in which retinal parallel processing streams remain segregated in the LGN before converging at the thalamocortical synapse. Our findings add a new dimension to Hubel and Wiesel's classic model of the V1 receptive field, by demonstrating how the convergence of functionally diverse thalamic afferents acts to maintain a constant representation of the sensory environment despite the changing structure of peripheral inputs.

Acknowledgments

This work has been supported by NIH R01EY028657.

Reference

1. Ruda, Zylberberg, Field, 2020, *Nat Commun*, 11:4605, 10.1038/s41467-020-18436-2

NEURAL MANIFOLDS CARRY REACTIVATION OF PHONETIC REPRESENTATIONS DURING SEMANTIC PROCESSING

Pavo Orepić[1,*], *Wilson Truccolo*[2,3], *Eric Halgren*[4], *Sydney S. Cash*[5], *Anne-Lise Giraud*[1,6], *Timothée Proix*[1]

[1] Department of Basic Neurosciences, University of Geneva, Geneva, Switzerland
[2] Department of Neuroscience, Brown University, Providence, United States
[3] Carney Institute of Brain Science, Brown University, Providence, Rhode Island, USA
[4] Department of Neurosciences, University of California, San Diego, California, USA
[5] Department of Neurology, Harvard Medical School, Boston, Massachusetts, USA
[6] Institut Pasteur, Université Paris Cité, Paris, France
[*] pavo.orepic@gmail.com

Traditional models of speech perception posit that neural activity encodes speech through a hierarchy of cognitive processes, from low-level representations of acoustic and phonetic features to high-level semantic encoding. Yet it remains unknown how neural representations are transformed across the levels of speech hierarchy. To carefully answer this question, it may be necessary to look beyond standard human neuroimaging and delve into human single-neuron recordings, which remain very rare.

Here, we investigated phonetic-to-semantic transformations in unique Utah array recordings from the human left anterior superior temporal gyrus, a brain region at the interface of phonetic and semantic speech processing, during an auditory semantic categorization task and natural speech perception. We hypothesized that the relevant functional read-out of phonetic and semantic processing would emerge at the level of low-dimensional neuronal population dynamics, known as neural manifolds. For this, we first regressed phonetic and semantic features to the activity of 176 neurons, and then performed PCA on the resulting regression weights to explore the latent dynamics on the corresponding neural manifolds.

We identified distinct neural manifolds for phonetic and semantic features, with a functional separation of their corresponding low-dimensional trajectories. The same low-dimensional dynamics generalized from a controlled experiment to natural speech perception, indicating invariant representations across different contexts. During natural speech, early phonetic features were re-encoded 400 ms after word onset, concurrently with the encoding of semantic features. Simultaneously, concurrent bottom-up and top-down processes were reflected in low-gamma and beta local field potentials across the Utah array. These effects were specific to neuronal population dynamics and could not be detected even on the most adjacent ECoG contact.

Our findings identify a neuronal substrate for phonetic-to-semantic transformations, substantiating the analysis-by-synthesis framework for speech, where early-processed phonemes are combined into a first semantic guess that is compared to phonological inputs at about 400 ms. We show, for the first time in humans, that neural manifolds provide a suitable lens to observe the fine-grained speech encoding mechanisms.

Acknowledgments

This work was funded by Swiss National Science Foundation career grant 193542.

LAYER 1: THE GATE OF THE BISTABLE VISUAL PERCEPTION

Ganna Palagina[1–3,*], *Sepehr Mani*[1], *Andreas Sarantopoulos*[4], *Daria Bogatova*[1], *Stelios Smirnakis*[1–3]

[1] Brigham and Women's Hospital, Boston, Massachusetts, USA
[2] Harvard Medical School, Boston, Massachusetts, USA
[3] Jamaica Plain VA Hospital, Boston, Massachusetts, USA
[4] European University Cyprus, Nicosia, Cyprus
[*] gpalagina@bwh.harvard.edu

Bi-stable visual stimuli cause subjective perception to reverse between two mutually exclusive interpretations of the same input, offering a well-controlled approach to the study of percept maintenance. In prior work [1] we combined bi-stable visual perception task with mesoscale 2-photon imaging of the cortical hemisphere in mice and identified a cortex-wide subnetwork of perceptual reversal-sensitive cells across V1, visuomotor area VRL, visual association cortex VA and secondary motor cortex M2. In this study, we continue this research line and investigate if the cell populations in layer 1 of the visual areas and M2 also contribute to bistable visual perception.

In layer 1 of the V1, VA and visual portion of M2 we found two distinct subnetworks of reversal-locked neurons: (i) sparse subnetwork of choline acetyltransferase positive (ChAT+) cells located at the border of layers 1 and 2, and (ii) denser subnetwork of ChAT-negative layer 1 cells. Remarkably, nearly all ChAT+ cells (97%) were reversal-locked. Out of the remaining layer 1 cells, only about 36% showed reversal-locked activity. Both populations showed co-activation with pyramidal cell ensembles located deeper in the layer 2/3 of the cortical column. Further, ChAT+ neurons were engaged by changes in behavioral state, walking velocity and pupil dynamics.

We next used chemogenetic control to see if the ChAT+ cells were necessary for the perceptual reversals to occur. We simultaneously expressed excitatory Gq-DREADD and inhibitory KORD in the ChAT+ cells in V1, VA and M2. We found that increasing the firing of ChAT+ cells via DREADD moderately increased the reversal rate. On the contrary, the application of KORD ligand Salvinorin B suppressed ChAT+ neuron firing and dramatically decreased the reversal rate.

We propose that ChAT+ neurons of layer 1 serve as a perceptual gate, linking local sensory processing units to behavioral variables and attention. They likely drive perceptual reversals by selectively activating pyramidal and SOM+ cell ensembles belonging to the reversal-locked subnetwork.

Acknowledgments

Supported by NIH R21 EY031537 to GP.

Reference

1. Palagina *et al.*, 2022, *Proc AREADNE 2022*, https://areadne.org/2022.

ORIGINS OF VARIABLE CORTICAL VARIABILITY

Jagruti J. Pattadkal[1,*], *Ronan O'Shea*[1], *David Hansel*[2], *Thibaud Taillefumier*[1], *Darrin Brager*[3], *Nicholas J. Priebe*[1]

[1] Department of Neuroscience, The University of Texas at Austin, Austin, Texas, USA
[2] CNRS, Université de Paris, Paris, France
[3] Department of Life Sciences, University of Nevada, Las Vegas, Nevada, USA
[*] jagruti@austin.utexas.edu

Across the neocortex, responses are variable when encoding sensory stimuli or generating motor actions. Understanding the source of variability is essential to studying its effects on encoding and decoding of neuronal signals. We present a thorough exploration into the origins of cortical Poisson and super-Poisson variability using *in vivo* conductance measurements, *in vitro* dynamic clamp electrophysiology, large-scale *in vivo* population variability measurements and simulations.

Variability can emerge from the synaptic inputs or intrinsic neuronal noise. Neurons respond with high precision when repeatedly injected with large fluctuating current [1]. Are neuronal responses to physiological inputs similarly deterministic? We measured stimulus-evoked conductances *in vivo* in mouse visual cortex and injected them in neurons in cortical slices, using dynamic clamp. Neurons *in vitro* showed minimal variability when a single trial conductance was injected repeatedly; and showed *in vivo* like Poisson variability when conductances from different trials of same stimulus were injected. Under physiological conditions, intrinsic noise plays little role and input variability is sufficient to generate Poisson spiking variability.

Uncorrelated inputs cannot account for observed *in vivo* variability, but it can be matched by synchronous inputs [2]. We measured synchrony in cortical populations using neuropixels in awake marmosets and mice. Cortical populations across species exhibit small but non-zero synchrony. We show that weak cortical synchrony is sufficient to cause Poisson-like variability by generating synthetic conductances from weakly synchronous neuronal input population and injecting them in neurons *in vitro*.

Cortical variability varies from super-Poisson (during spontaneous responses) to Poisson. How does this change in variability statistics emerge? Changing synchrony was not the cause as cortical synchrony does not change between stimulus-evoked versus spontaneous states. Instead, we find that the time-scale of synchrony varies between states *in vivo*; fluctuations are slow during spontaneous epochs and fast during driven responses. Injecting neurons *in vitro* with conductances from populations with fast versus slow fluctuations is sufficient to change responses from Poisson to super-Poisson.

Acknowledgments

Supported by NIH R01 EY025102.

References

1. Mainen, Sejnowski, 1995, *Science*, 268:1503–1506, 10.1126/science.7770778
2. Stevens, Zador, 1998, *Nat Neurosci*, 1:210–217, 10.1038/659

THE NEURAL GEOMETRY OF EMOTIONAL STATES IN THE AMYGDALA

Lorenzo Posani[†,*], *Pia-Kelsey O'Neill*[†], *Jozsef Meszaros, Phebe Warren, Carl E. Schoonover, Andrew J. P. Fink, Stefano Fusi, C. Daniel Salzman*

Mortimer B. Zuckerman Mind Brain Behavior Institute, Columbia University, New York, USA
[†] equal contribution
[*] lp2878@columbia.edu

Emotional responses to salient events can elicit multiple behaviors, such as freezing or fleeing in response to threats. Whether the brain uses distinct neural mechanisms to represent these reactions is an open question. Here, we analyzed the coding properties of the basolateral amygdala (BLA), a region that connects emotionally relevant stimuli to behavioral responses [1], in mice in a virtual burrow assay (Fig. 1a) [2]. Subjects were presented with aversively conditioned or neutral stimuli, to which they could respond with different defensive behaviors, including trembling (freezing) and ingressing into the burrow (fleeing to safety).

We used a geometric decoding approach [3] to analyze the population code of BLA neurons in response to stimulus identity, stimulus valence, trembling behavior, and ingressed state. We found that these variables were all encoded in the activity with mixed selectivity (Fig. 1b, c), challenging the notion that the amygdala encodes different variables with specialized sub-populations. We found that the representational geometries of valence during tremble or after ingressing were markedly different. While valence and tremble were represented in a low-dimensional, disentangled geometry (Fig. 1d), valence was not decodable in the ingressed state (Fig. 1e). In fact, ingressing in the burrow modulated the conjunctive coding of valence and state, forming a geometry that indicates an abstract representation of safety (Fig. 1e).

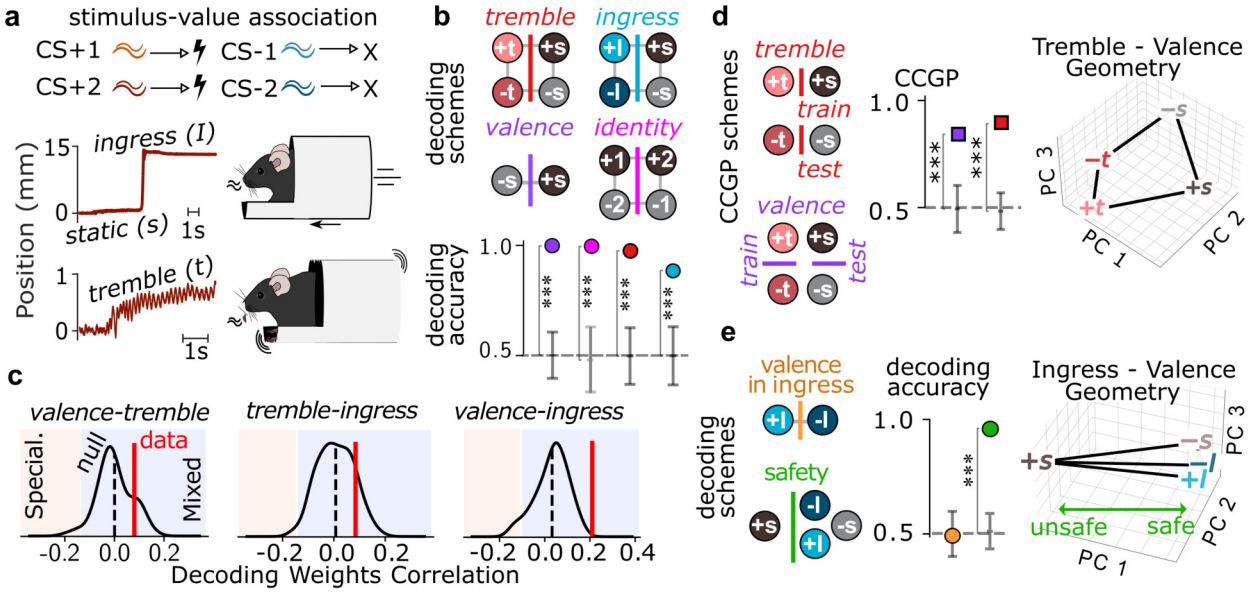

Figure 1. Decoding basolateral amygdala responses with factors of stimulus identity, stimulus valence, trembling, and ingressed state.

References

1. Janak, Tye, 2015, *Nature*, 517, 284–292, 10.1038/nature14188
2. Fink, Axel, Schoonover, 2019, *eLife*, 8:e45658, 10.7554/eLife.45658
3. Bernardi, Benna, Rigotti, *et al.*, 2020, *Cell*, 183:954–967, 10.1016/j.cell.2020.09.031

A MULTIVARIATE POINT PROCESS POPULATION CODE FOR SIMULTANEOUSLY RECORDED SPIKE TRAINS

Reza Ramezan[1,*], *Meixi Chen*[1], *Martin Lysy*[1], *David Moorman*[2]

[1] Department of Statistics and Actuarial Science, University of Waterloo, Ontario, Canada
[2] Department. of Psychological and Brain Sciences, Univ. of Massachusetts, Amherst, Mass., USA
[*] rramezan@uwaterloo.ca

Statistical analysis of simultaneously recorded neural spike trains from an ensemble of neurons is a challenging problem from both statistical and computational points of view. The Skellam Process with Resetting (SPR) is a continuous-time discrete-state model tailored for neural spike trains, but it is not scalable in its current form [1, 2]. We introduce a generalization of our previous SPR framework under a continuous-time latent factor model (LFM). To obtain scalability, we use computationally efficient approximate Bayesian inference, and employ a latent factor model to reduce the dimensionality (Fig. 1). Our point-process model can handle larger neuronal ensembles compared to alternative approaches. To the best of our knowledge, this model is the first continuous-time multivariate LFM for studying neuronal interactions and functional connectivity.

In a classical conditioning study, OFC recordings were collected from 6 rats where each subject underwent 6 different experimental conditions to understand whether OFC neurons predict the outcome (sucrose, quinine, ethanol) based on two delivery paradigms: active (at a well) and passive (via a catheter in their mouths). Our population coding results based on the LFM model show that in passive experiments, OFC neurons have a robust predictive power for ethanol delivery.

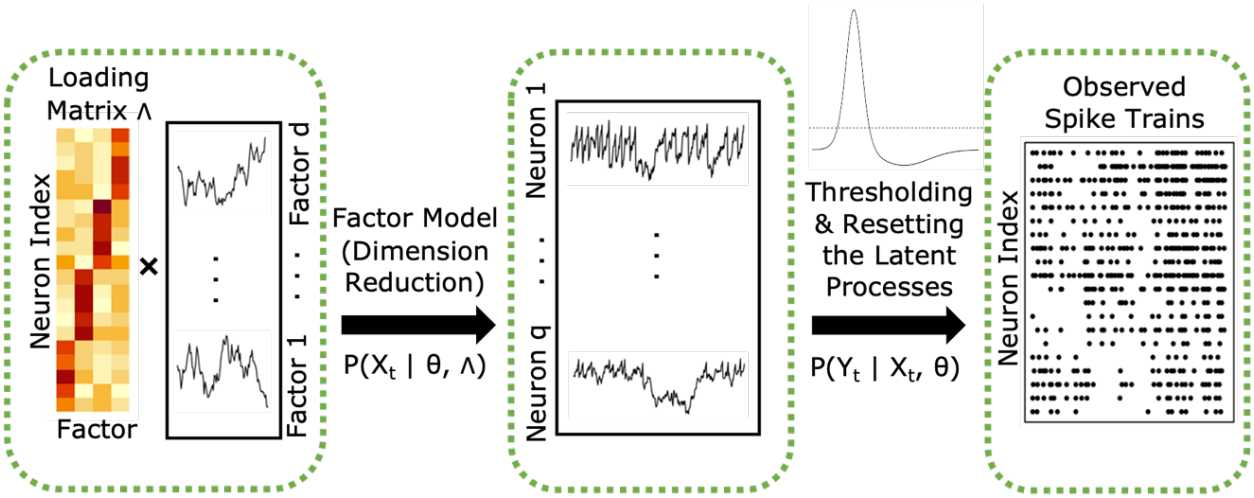

Figure 1. The schematic representation of the LFM model.

References

1. Ramezan, Marriott, Chenouri, 2016, *Stat Med*, 35(30):5717–5729, 10.1002/sim.7127
2. Ramezan, Chen, Lysy, *et al.*, 2022, *Proc Conf Cog Comp Neurosci*, 189–191

MIRROR NEURON POPULATIONS DO MIRROR

Vassilis Raos[1–3,*], *Konstantinos Chatzimichail*[1,2], *Eleftheria Tzamali*[1], *Vassilis Papadourakis*[1]

[1] Lab of Movement Physiology, Medical School, University of Crete, Heraklion, Greece
[2] Graduate Program in the Brain and Mind Sciences, UoC & FORTH Heraklion, Greece
[3] Institute of Applied & Computational Mathematics, FORTH, Heraklion, Greece
[*] raos@uoc.gr

Mirror neurons (MirNs) respond not only when a monkey performs an action but also when it observes a similar action executed by another. The prevailing idea is that this response pattern contributes to action comprehension through a mechanism known as the *direct-matching hypothesis*. This hypothesis posits that MirNs, as implied by the term *mirror*, exhibit congruence between the actions they encode during observation and execution. Contrary to this assumption, available electrophysiological studies consistently reveal that MirNs categorized as *strictly congruent*, closely mirroring observed and executed actions, constitute a minority. Our recent finding that strictly congruent neurons are encountered at chance-level frequencies within the recorded MirNs population, underscores the significant limitations of the MirNs theory. We hypothesized that the matching process might occur in a distributed manner across specific elements of the neuronal ensemble. Consequently, while the mirror property may not be a characteristic of each MirN individually, it could define the MirN population as a collective entity.

To investigate this hypothesis, we utilized our MirNs database, which includes spiking neural activity data recorded from the forelimb representations of dorsal and ventral premotor cortical areas, while macaque monkeys were either observing or performing reaching-to-grasp actions. Applying pattern classifiers, we found that grip-related information can be reliably decoded from the population activity of MirNs in both conditions (execution and observation) over extended time periods. Throughout movement, whether observed or executed, we revealed a distributed dynamic code. Notably, in the observation condition, the classification accuracy saw improvement when the grip most frequently misidentified, as indicated by the confusion matrices, was excluded. It is worth mentioning that the excluded grip shared significant kinematic similarities with others. Cross-conditional decoding revealed an extended period spanning from the middle of the movement to the middle of the holding period, with significant readout. This readout improved when (1) the grip most frequently misidentified in the observation condition decoding was excluded, (2) only the congruent units of the population were used (as expected) and (3) when a shared space between the two conditions was identified using neural manifold optimization. These findings offer compelling evidence that neural population codes for observed and executed actions align, endorsing our hypothesis.

Acknowledgments

Supported by a grant co-financed by Greece and the European Union (European Social Fund-ESF) through the Operational Programme “Human Resources Development, Education and Lifelong Learning 2014–2020” in the context of project MIS 5048512.

PROPAGATING SPATIO-TEMPORAL PATTERNS ACROSS THE PRIMARY MOTOR CORTEX DECODE KINEMATICS

Milan Rybar[1,*], *Wei Liang*[2], *Karthikeyan Balasubramanian*[1], *Nicholas G. Hatsopoulos*[1]

[1] Department of Organismal Biology and Anatomy, University of Chicago, Chicago, Illinois, USA
[2] Committee on Computational Neuroscience, University of Chicago, Chicago, Illinois, USA
[*] milanrybar@uchicago.edu

The primary motor cortex (M1) is well-known to be somatotopically organized. However, this classical view is static and does not consider temporal and spatial dynamics. Previous studies have shown that spatially organized patterns of excitability along a rostro-caudal direction play a role in the initiation of movement but do not specify the details of a particular movement [1]. These propagating patterns were observed in the attenuation times of low-frequency beta oscillation (15–35 Hz) amplitude of the local field potential (LFP). Moreover, at movement onset, amplification times of high-frequency gamma band (200–400 Hz) of the LFP spatially propagate across M1 during a center-out reaching task, and the direction of this propagating pattern carries kinematic information about the movement direction [2]. Since high-gamma LFP amplitude serves as a close proxy to multi-unit activity, these findings suggest that an orderly recruitment sequence of activity across the motor cortical sheet provides behaviorally relevant information.

Here, we investigate instantaneous spatial patterns of high-gamma at every moment in time and not just at movement onset. We analyze neural data recorded from two rhesus macaques with implanted neural arrays in M1 while they performed a planar center-out task in eight different directions [2]. A plane is fitted to normalized high-gamma amplitudes across the array for each time point. From a spatial gradient of each plane, we can estimate the gradient direction and its magnitude. The gradient direction across the motor cortex dynamically varies throughout a reaching movement and differentiates between different reach directions. Using the time-varying gradient direction, it is possible to decode hand velocity continuously over time using a linear decoder. The decoder performs particularly well at the beginning of movement but then the performance degrades over time suggesting that these dynamic spatial patterns encode movement parameters intermittently. To partially examine this question, we use a subset of trials from one of the two monkeys which was overtrained to make movements to one target. This behavior led to movement paths that were initially directed to the overtrained target and then switched direction to reach the neighboring instructed target resulting in bent paths. By decoding the hand velocity during these bent path trials using dynamic spatial patterns, we demonstrate that these patterns carry intermittent information about hand velocity at the beginning of the movement and after the transition point where the hand changes direction. Overall, these results offer a novel perspective on how information about movement initiation and execution is encoded in M1 in the form of spatio-temporal patterns across the cortical sheet.

Acknowledgments

This work has been supported by R01NS111982 from the NIH.

References

1. Balasubramanian, *et al.*, 2023, *iScience*, 26(4):106518, 10.1016/j.isci.2023.106518
2. Liang, Balasubramanian, *et al.*, 2023, *PNAS*, 120(4), 10.1073/pnas.2212227120

MOTION DIRECTION DECODING IN MOUSE V1: PREDICTIVE POWER RELATES TO FUNCTIONAL CONNECTIVITY ORGANIZATION

Mario Alexios Savaglio[1,2,†], *Cristina Brozi*[1,2,†], *Eleftheria Psilou*[2,†], *Chryssalenia Koumpouzi*[2], *Christos Vasilakos*[2], *Spyros Nessis*[2], *Emmanouel Smyrnakis*[2], *Vasileios Akoutas*[2], *Emmanouel Koniotakis*[2], *Georgios Keliris*[2], *Ganna Palagina*[5], *Stelios Smirnakis*[5,‡,*], *Maria Papadopouli*[1–3,‡,⁑]

[1] Dept. of Computer Science, University of Crete, Heraklion, Greece
[2] ICS, Foundation for Research and Technology Hellas, Heraklion, Greece
[3] Archimedes Research Unit, Athena Research Center, Athens, Greece
[4] Hellenic Mediterranean University, Heraklion, Greece
[5] Dept. of Neurology, BW's and JP VA Hospitals, Harvard Medical School, Boston, Mass., USA
[†] equal contribution
[‡] equal contribution
[*] smsmirnakis@bwh.harvard.edu, [⁑] maria@csd.uoc.gr

Variability in single neuron responses presents a challenge in establishing reliable representations of visual stimuli essential for driving behavior. To enhance accuracy, it is necessary to integrate responses from multiple neurons. Studies use decoding to analyze how accurately visual features are represented in the brain [1, 2]. Here we recorded from a large population of neurons in mouse area V1 using mesoscopic 2-photon-calcium-imaging under spontaneous activity conditions and visual stimulation along 16 randomly shuffled directions of motion. We systematically evaluated various decoding strategies and settled on using Linear Support Vector Classification in conjunction with recursive feature elimination to identify neurons with high power in predicting the stimulus direction of motion. We found that neurons displaying high predictive power (HPP) for direction-of-stimulus decoding exhibit a more reliable orientation preference, higher firing rates relative to the rest of the neurons, and are preferentially located in L2/3. As expected, under visual stimulation, functional connectivity, computed using the STTC metric [3], among HPP neurons is stronger and denser compared to other neurons. Intriguingly, in the absence of stimulus, this pattern reverses, with HPP neurons exhibiting, a smaller degree of connectivity, shorter range, lower functional connectivity strength, and lower firing rates than control groups. This suggests that HPP neurons are less susceptible to being driven via stimulus-unrelated internal brain state modulations. Furthermore, in contrast to the frequently reported salt and pepper organization of direction selectivity in the mouse visual cortex, stimulus-driven functional connections among HPP neurons occur more frequently at specific distances, suggesting a clustered organization. Our findings shed light on the organization of neuronal ensembles important for decoding visual motion direction in mouse area V1, contributing to the understanding of information processing in the mouse visual cortex.

Acknowledgments

Funded by HFRI neuronXnet 2285 and neuron-AD 4058, EU MSCA GA101007926, NINDS R21 NS088457, and MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 under the NextGenerationEU. We thank Taliah Muhammad for help with the experiments as well as Marianna Papadostefanaki and Nikolaos Tzanakis for their analysis.

References

1. Stringer, Michaelos, *et al.*, 2021, *Cell*, 184:2767–2778, 10.1016/j.cell.2021.03.042
2. Graf, Kohn, et al., 2011, *Nat Neurosci*, 14(2):239–45, 10.1038/nn.2733
3. Cutts, Eglen, 2014, *J Neurosci*, 34(43):14288–303, 10.1523/jneurosci.2767-14.2014

FUNCTIONAL ORGANIZATION OF THE VISUAL INPUT IN MOUSE SUPERIOR COLLICULUS

Lisa Schmors[1,*], *Na Zhou*[2], *Yongrong Qiu*[3], *Andreas Tolias*[2], *Fabian Sinz*[3], *Philipp Berens*[1], *Jacob Reimer*[2], *Katrin Franke*[2]

[1] Hertie AI, University of Tübingen, Tübingen, Germany
[2] Department of Neuroscience, Baylor College of Medicine, Houston, Texas, USA
[3] Institute for Computer Science & CIDAS, Uni Göttingen, Göttingen, Germany
[*] lisa.schmors@uni-tuebingen.de

Visual processing in mammals starts in the retina and proceeds through two major targets, the dorsal lateral geniculate nucleus and the superior colliculus (SC), before reaching downstream areas like the visual cortex. Studies from recent years have identified about 40 distinct types of retinal ganglion cells (RGCs) through a multi-modal classification [1]. However, a thorough understanding as to how this diverse retinal signal is organized and processed by downstream targets is still missing, especially for the SC.

Here, we aimed to investigate the general organizational principles of the SC, focusing on the visual input from RGCs. We employed intraocular injections of a viral vector to express GCaMP8m, a fast genetically encoded calcium indicator, in all RGCs types. For chronic functional imaging, we then implanted a custom-built and 3D printed window over the SC, maintaining the integrity of the cortex overlaying the SC. This approach enabled us to record the stimulus-evoked excitatory drive from over 200,000 individual RGC axon terminals using two-photon calcium imaging in awake, head-fixed mice. We employed a Gaussian mixture model (GMM) to cluster responses to well-defined artificial stimuli into distinct functional retinal input modules, based on feature weights extracted using sparse PCA. The stimuli included moving bars to test direction selectivity and a full-field chirp stimulus to assess light polarity and response transience.

We demonstrate that the functional diversity of RGC axon terminals in the SC exceeds the diversity observed at the retina. Currently, we are training CNN models using responses to natural movies for an in-depth *in silico* characterization and to test the tuning to behaviorally relevant stimuli.

Finally, we are currently performing additional experiments to record from SC somas to investigate the rules of integration. This study provides significant insights into how sensory input is functionally organized in the SC. It is an important step for identifying key computational principles of visual processing, enhancing our understanding of the complex mechanisms of visual perception of behaviorally relevant visual features in mammals.

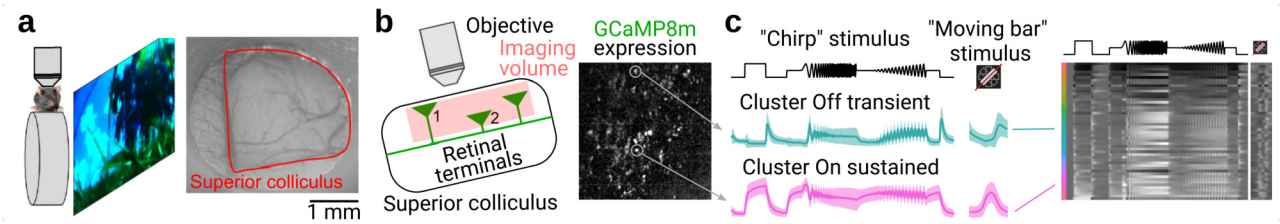

Figure 1. Recording and clustering of visual activity in the SC of awake, head-fixed mice. (**a**) Experimental setup. (**b**) Recording of retinal axons. (**c**) Clustering using GMM.

Reference

1. Baden, Berens, Franke, *et al.*, 2016, *Nature*, 529(7586):345–50, 10.1038/nature16468

ON THE EMERGENCE OF INTERPRETABLE RECEPTIVE FIELDS IN CNNS

Fabio Seel[1,*], *Sacha Sokoloski*[1], *Philipp Berens*[1,2]

[1] Hertie AI, Tübingen, Germany
[2] Tübingen AI Center, Tübingen, Germany
[*] fabio.seel@uni-tuebingen.de

Convolutional Neural Networks (CNNs) are often applied to modelling the neural circuits of animal vision. Similarities have been reported between the receptive fields (RFs) of individual channels in a CNN when compared to the RFs of neurons in the visual pathway [1]. Here, we demonstrate that the emergence of interpretable RFs in CNNs is highly dependent on the choice of architecture and hyperparameters, and we show how to construct CNN models that exhibit RFs that are similar to what is recorded in visual circuits.

Downsampling, whether through stride or pooling, is an architectural feature that allows CNNs to efficiently process high-resolution inputs. To assess the impact of downsampling on RFs, we train various CNNs to classify natural images, and augment the images by rescaling and translating to increase overall image resolution and account for multiscale effects. We find that the quality of RFs depends both on the downsampling factor and the kernel size (Fig. 1). In particular, if the ratio of the downsampling factor over the kernel size is greater than 0.5, RFs in the network exhibit high frequency artifacts, and eventually become uninterpretable. Moreover, uninterpretable RFs can even emerge in CNNs with relatively high classification performance. We have also found that the emergence of RFs can depend on various other choices from activation function, to data augmentation, to regularization (not shown).

In conclusion, careful attention needs to be paid to hyperparameters in order to capture qualitative features of information processing in neural circuits with CNNs. In particular, CNNs can exhibit high-task performance, while failing to exhibit the canonical RFs of visual neural circuits.

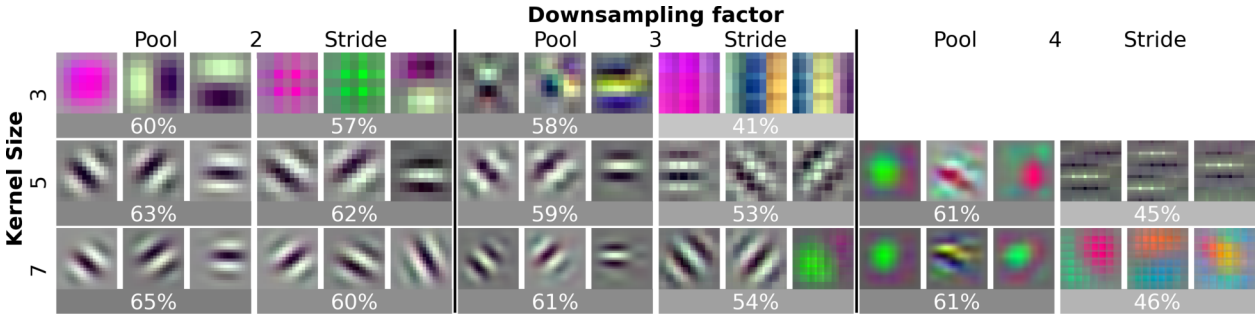

Figure 1. Comparison of the 2nd layer RFs (three strongest each) for different kernel sizes and downsampling. Below the validation performance of the network is shown.

Reference

1. Lindsey, Ocko, Ganguli, *et al.*, 2019, *ICLR*, 10.48550/arXiv.1901.00945

SPATIOTEMPORAL EEG CHARACTERISATION OF MULTISENSORY PROCESSING IN AUTISM AND SCHZOPHRENIA

Amirreza Nadimi Shahraki[1,*], *Maida Toumaian*[2], *Nikolaos Smyrnis*[3], *Ioannis Delis*[4]

[1] Amirreza Nadimi Shahraki, University of Leeds, Leeds, United Kingdom
[2] Lab Cog Neurosci and Sensorimotor Control, Univ Mental Health, Athens, Greece
[3] University General Hospital "ATTIKON", Athens, Greece
[4] School of Biomedical Sciences, University of Leeds, Leeds, United Kingdom
[*] scans@leeds.ac.uk

Multisensory integration (MSI) is the capacity to combine information from multiple senses, such as vision and audition, leading to improvements in behavioral performance, such as faster responses [1]. This process can be affected in individuals with developmental or mental disorders [2, 3].

Here, we introduce tensor decomposition as a data-driven approach to elucidate neural representations of multisensory processing that may transcend diagnostic boundaries among individuals with Autism Spectrum Disorder (ASD), Schizophrenia (SZ), and controls (CN). A total of 32 CN, 23 ASD, and 35 SZ individuals were instructed to respond as quickly as possible to Auditory (A), Visual (V), Audio-visual (AV), or no (catch — C) cues while their electroencephalograms (EEG) were recorded.

Significant reaction time (RT) effects were observed for population ($p < 0.01$) and sensory condition ($p \ll 0.01$), with CN displaying the fastest responses, followed by ASD and SZ. The AV conditions yielded faster responses, indicating multisensory behavioral benefits for all three populations. Based on the presence or absence of multisensory benefit, participants were divided into integrators (I)/non-integrators (N), with a higher N prevalence in ASD (74%) followed by SZ (51%) and CN (22%).

Non-negative Canonical Polyadic Decomposition (NCP) [4] was applied to the EEG recordings to reveal spatiotemporal components of audiovisual processing across populations. We identified two components distinctly activated by the I and N groups: (i) a centro-parietal EEG component peaking at 300–400 ms post-stimulus presentation, showing AV enhancement over unisensory trials, and (ii) a parieto-occipital EEG component with a double peak at about 300 ms and 500 ms that was equally activated in AV and V conditions.

Our study provides evidence for distinct neural signatures of AV integration (or lack thereof), which are shared between ASD and SZ populations and underlie the presence or absence of multisensory behavioral benefits.

References

1. Frassinetti, Bolognini, Làdavas, 2002, *Exp Brain Res*, 147(3):332–343
2. Ainsworth, Ostrolenk, Irion, Bertone, 2021, *Cortex*, 134:195–206
3. Williams, Light, Braff, Ramachandran, 2010, *Neuropsychologia*, 48(10):3128–3136
4. Sidiropoulos, De Lathauwer, Fu, *et al.*, 2017, *IEEE Trans Sig Proc*, 65:3551–3582

NEURAL REPRESENTATION GEOMETRY IN VISUAL RELATIONAL REASONING

Jiaqi Shang[1], *Gabriel Kreiman*[2–4], *Haim Sompolinsky*[4,5,*]

[1] Program in Neuroscience, Harvard Medical School, Boston, Massachusetts, USA
[2] Boston Children's Hospital, Harvard Medical School, Boston, Massachusetts, USA
[3] Center for Brains, Minds, and Machines, Cambridge, Massachusetts, USA
[4] Center for Brain Science, Harvard University, Cambridge, Massachusetts, USA
[5] Edmond and Lily Safra Center for Brain Sciences, Hebrew University, Jerusalem, Israel
[*] hsompolinsky@mcb.harvard.edu

When perceiving visual scenes, we recognize individual objects and their relations. Previous studies have identified neurons in areas such as the prefrontal cortex (PFC) that encode relational information [1]. However, deciphering the mechanisms employed by these neural populations remains challenging. Our study uses a simplified Raven's Progressive Matrices (RPM) task [2] to investigate relational reasoning. We posit that the brain encodes visual relational information in rule-specific neural manifolds and apply a geometrical theory [3] to study how neural representations facilitate relational reasoning.

We introduce a geometrical theory that links RPM performance to key relational representation geometries: signal (manifold centroid distances) and dimensionality. To explore potential mechanisms, we trained a convolutional model on an RPM dataset comprising images with various relational rules. The theory accurately predicted the model's task performance and proposed two potential mechanisms: effective relational reasoning can result from a large signal or high dimensionality. The model revealed a signal-dimensionality trade-off, where optimizing for large signals reduces dimensionality and vice versa. Which mechanism dominates in the brain can be determined by analyzing the geometry of neural population activities along the visual stream from the inferior temporal (IT) to higher-order PFC areas.

We conducted psychological experiments with human subjects using the same questions tested on the model. Human performance closely matched that of the model, displaying similar error patterns across different relation types. This behavioral alignment suggests that the brain may employ similar geometrical strategies for relational reasoning as observed in the model. To further explore this, we presented macaques with RPM images and recorded neural activities in IT. We analyzed the geometry of neural population activity and decoded relational information as the animals passively observed the images. Next, we will train the animals to perform the RPM task and record neural activities in IT and PFC. We will examine the geometry of neural population activity before and after training to understand how these representations support relational reasoning. Overall, our research provides new insights into the neural mechanisms of relational reasoning through the lens of representation geometry.

Acknowledgments

This research is supported by NIH grant No. R01EY026025, the Swartz Foundation, ONR grant No. N0014–23-1–2051, and a generous gift from Amazon.

References

1. Davis, Goldwater, Giron, 2017, *Cerebral Cortex*, 27(4):2652–2670, 10.1093/cercor/bhw099
2. Raven, 1936, MSC thesis Univ London
3. Sorscher, Ganguli, Sompolinsky, 2022, *PNAS*, 19(43):e2200800119, 10.1073/pnas.2200800119

PUPIL SIZE PREDICTS CRITICAL TRANSITIONS IN PREFRONTAL NEURONAL POPULATION ACTIVITY AT THE ONSET OF EXPLORATION

Akram Shourkeshti[1], *Gabriel Marrocco*[1], *Katarzyna Jurewicz*[1], *Tirin Moore*[2,3], *R. Becket Ebitz*[1,*]

[1] Department of Neurosciences, Université de Montréal, Montréal, Québec, Canada
[2] Department of Neurobiology, Stanford University, Stanford, California, USA
[3] Howard Hughes Medical Institute, Chevy Chase, Maryland, USA
[*] becket@ebitzlab.com

In uncertain environments, intelligent decision-makers exploit actions that have been rewarding in the past, but also occasionally explore alternatives that could be even better [1]. Neural recordings suggests that exploratory decision-making involves increased noise in prefrontal neural populations [2, 3]. Understanding the cause of noise could be the key to understanding exploratory discovery. However, in part because past studies conflated the decision to explore with evidence that exploration is warranted, the mechanisms responsible for generating this noise remain unknown.

Here, we examined the behavior and neural activity of 2 rhesus macaques making spontaneous decisions to explore or exploit in a dynamic reward-based decision-making task. This allowed us to dissociate the relative contribution of critical variables like reward, uncertainty, and arousal to (i) the likelihood of exploration, and (ii) exploratory noise in prefrontal neuronal population activity.

We found two main results. First, one measure of arousal (pupil size under constant luminance) predicted the onset of exploration, beyond what could be explained by rewards. In fact, the onset of exploration was entrained by spontaneous oscillations in pupil size that were completely unrelated to reward. Second, pupil size independently predicted noisy patterns of prefrontal neural activity in single neurons and in the population, even within periods of exploitation. This noise was entrained by spontaneous pupil oscillations, but also had certain features that characteristic of critical transitions in complex systems, like networks of neurons.

Together, these results suggest a model in which pupil-linked mechanisms promote exploratory sampling via driving the prefrontal cortex through a critical tipping point in which prefrontal population dynamics become disorganized and then must re-form.

Acknowledgments

This work was supported by the Natural Sciences and Engineering Research Council of Canada (RGPIN-2020-05577), the Fonds de Recherche du Québec-Santé (Junior 1 Chercheur-Boursier 284309), the Jacobs Foundation and CIFAR Azrieli Global Scholars.

References

1. Wilson, *et al.*, 2021, *Curr Opin Beh Sci*, 38:49–56, 10.1016/j.cobeha.2020.10.001
2. Ebitz, Albarran, Moore, 2018, *Neuron*, 97(2):450–61, 10.1016/j.neuron.2017.12.007
3. Ebitz, Tu, Hayden, 2020, *PLoS Biol*, 18(11):e3000951, 10.1371/journal.pbio.3000951

HIGH GAMMA ACTIVITY RELATES TO SYNCHRONY IN THE SPIKING ENSEMBLE MANIFOLD IN MOTOR CORTEX

Marc W. Slutzky[1,2,*], *Tianhao Lei*[1], *Michael Scheid*[1], *Joshua Glaser*[1]

[1] Department of Neurology, Northwestern University, Chicago, Illinois, USA
[2] Departments of Neurosci., PM&R, and Biomed. Eng., Northwestern University, Chicago, Ill., USA
[*] mslutzky@northwestern.edu

Local field potentials (LFPs) and electrocorticography (ECoG) in the cerebral cortex are derived from a variety of brain sources, which are reflected in different frequency bands. The most information in these signals typically is in the high-gamma band, which is variably defined as anywhere from 70–300 Hz. This band is often suggested to be a proxy of local neuronal spiking activity. Here, we investigated whether high-gamma activity (HGA) could be separated from multi-unit spiking activity on the same intracortical electrode on a Utah array in primary motor cortex of 3 monkeys.

We designed an orthogonal neural feedback (ONF) brain machine interface paradigm that mapped spiking activity and HGA (200–300 Hz) from the same control electrode to orthogonal directions of cursor movement. Thus, the monkeys needed to independently modulate spiking and HGA on that electrode to reach targets in the cardinal directions. The monkeys reliably dissociated spiking activity from HGA, reducing the correlation between them from 0.5±0.2 (during normal hand control) to 0.1±0.1 (during ONF). This indicated that HGA is not simply a sum of local spiking activity.

We sought to further characterize the relationship between HGA and the spiking activity in the entire recorded population. We first used factor analysis to quantify how the variance of the individual control channel was shared with the neural population during ONF control. We found that the shared variance of the spike rates were consistently and significantly higher when the monkeys were solely modulating HGA (HGonly) than when modulating spiking (SPonly) on the control channel. Further, the first factor (co-firing in the ensemble) was highly correlated with HGA on the control channel, but not correlated with the spiking activity on that channel. This suggests that HGA is produced not simply by the sum of local neuronal activity, but by the synchronous firing (co-firing) of multiple neurons across a population.

Finally, we examined (using principal angles) the alignment of the low-dimensional subspaces representing ensemble activity during HGonly versus SPonly conditions during ONF control and during hand control. The principal angles between these two conditions were much smaller than chance during both ONF control and hand control, but slightly greater during ONF control. This suggests that modulating HGA involves a slightly different subspace of the main ensemble manifold than does modulating spiking only.

Acknowledgments

Supported by NIH grants K08NS060223, R01NS112942, and T32NS047987.

A DORSAL MEDIAL PREFRONTAL MOTOR CIRCUITS ENCODES INITIATION OF PERSISTANT MOVEMENT

Qian-Quan Sun[*], *Yihan Wang*

Wyoming Sensory Biol. Ctr. of Biomed. Res. Excellence, Univ. of Wyoming, Laramie, Wy., USA
[*]neuron@uwyo.edu

Medial prefrontal cortex (mPFC) regulates decision-making by amplifying certain information while suppressing others. Recently we reported a group of mPFC motor projecting (MP) neurons [1], which majorly projects to the primary motor neurons in almost all motor cortices and the striatum, but less to other deep brain regions and local non-MP neurons. Therefore, MP neurons may be involved in the most downstream mPFC circuit, which collects all filtered information to affect subsequent behavior. As such, we asked (Q1) if the MP neurons play a role on instructing subsequent movement in decision-making. MP neurons receive unidirectional inputs from the insular cortex (IC), which encodes valence, and from the basal lateral amygdala (BLA), which is responsible for valence assignment (internal belief), here we asked (Q2) what type of information (contextual or valence) the MP neurons encode during a persistent movement. Using single-unit extracellular recordings and opto-tagging in awake mice, we demonstrated that MP neurons in the dMPFC selectively encode contextual information, rather than natural valence, triggering a persistent movement. Inactivation of dmPFC MP neurons impairs the initiation of persistent movement and reduces neuronal activity in the insular and motor cortex. Finally, a computational model suggests that a successive sensory stimulus acts as an input signal for the dmPFC MP neurons to initiate a persistent movement. These results reveal a neural initiation mechanism on the persistent movement.

Acknowledgments

Supported by NIH grants from National Institute of Mental Health (1R21MH131363) and from National Institute of General Medical Sciences (2P20GM121310).

Reference

1. Wang, Sun, 2021, *Cell Rep*, 36(12):109733, 10.1016/j.celrep.2021.109733.

ATYPICAL EXTRACELLULAR WAVEFORMS AND FUNCTIONAL RESPONSES IN AWAKE MONKEY LGN

Shi Hai Sun[1,2,*], *Nathan J. Killian*[3], *John S. Pezaris*[1,2]

[1] Department of Neurosurgery, Massachusetts General Hospital, Boston, Massachusetts, USA
[2] Department of Neurosurgery, Harvard Medical School, Boston, Massachusetts, USA
[3] Department of Neurological Surgery, Albert Einstein College of Medicine, Bronx, New York, USA
[*] shsun@mgh.harvard.edu

In primates, the LGN is anatomically separated into layers of magnocellular (M), parvocellular (P), and koniocellular (K). Electrophysiological studies have shown that M, P, and K neurons differ in their stimulus and response dynamics: M neurons have fast temporal dynamics, well suited for motion detection; P neurons are sensitive to chromatic dynamics, well suited for color and form processing; and K neurons receive input from short-wavelength photoreceptors.

Although these divisions have been extensively characterized, much of our knowledge about the function of the LGN has been obtained from electrophysiological recordings using single-channel electrodes in the form of extracellular spikes. With the recent advancement of dense multi-electrode arrays and sophisticated spike-sorting algorithms, there has been a substantial increase in the variation of spike waveform shapes reported in multiple species and brain regions that are atypical and poorly understood. Accordingly, key populations may have been overlooked in earlier studies; identifying the gamut of extracellular signals from the LGN may help understand the visual pathway's processing.

Thus, in this study, we agnostically survey the extracellular space in the LGN of rhesus macaques (n = 281), employing a variety of stimulus and electrode properties. For every LGN unit recorded, we recovered its receptive field (RF) class (M, 48%; P, 23%; or K, 6%) and extracellular spike waveform class (Narrow, 33%; Broad, 22%; Triphasic, 13%; or Positive, 23%), along with several response metrics, and relationships between spike shape and possible neuronal type were identified. We also recovered a set of units without an estimated RF (non-RF, 24%) that are not regularly reported, with most of these units consistently responsive to the visually presented stimulus at a lower and more sustained response rate than units with an RF. The presence of N cells hints that the LGN may require understanding beyond the classification into M, P, and K responses, and further insight into these nuances (*e.g.*, spatial transcriptomics) could reveal computations in the visual system that have yet to be discovered. These new findings have interesting implications for our understanding of LGN function that can be applied broadly and extend their relevance to the broader aspect of neuroscience.

Acknowledgments

Supported by NIH EY027888 and the William M. Wood Foundation, Bank of America, Trustee.

ONE NETWORK TO RULE THEM ALL: RESERVOIR COMPUTING WITH NETWORK MODELS OF ASSOCIATIVE MEMORIES

Gianni Valerio Vinci[*], *Andrea Galluzzi, Maurizio Mattia*

Natl. Ctr. for Radiation Protection and Compu. Phys., Istituto Superiore di Sanità, Rome, Italy
[*]gianni.vinci@iss.it

The brain is a complex dynamical system that in order to function properly needs to encode and decode external information with an high degree of accuracy and resilience. To do so it is essential to memorize familiar stimuli and recall them if necessary. Hebbian learning allows to store associative memories as fixed point, or attractors, of recurrent neural networks with symmetric synaptic couplings [1, 2]. However, realistic memories are not static pictures but rather complex dynamical sequences. On the other hand, reservoir computing [3] is a reinforcement learning technique that allows to learn complex sequences harnessing the nonlinearity of recurrent networks but it is mostly assumed to work only with highly asymmetric networks. Our main result is to show that it is indeed possible to implement the reservoir computing paradigm with associative networks bridging the gap between two pillars of theoretical neuroscience. Leveraging the analytical tractability of this models we were able to formulate a complete theory that explains the mechanism by witch these networks learn to mimic arbitrary dynamical systems and even predict the success in reproducing them given the structural parameters of the network. For example we found that the best performances are found with networks that have a number of stored pattern beyond the perfect recall phase.

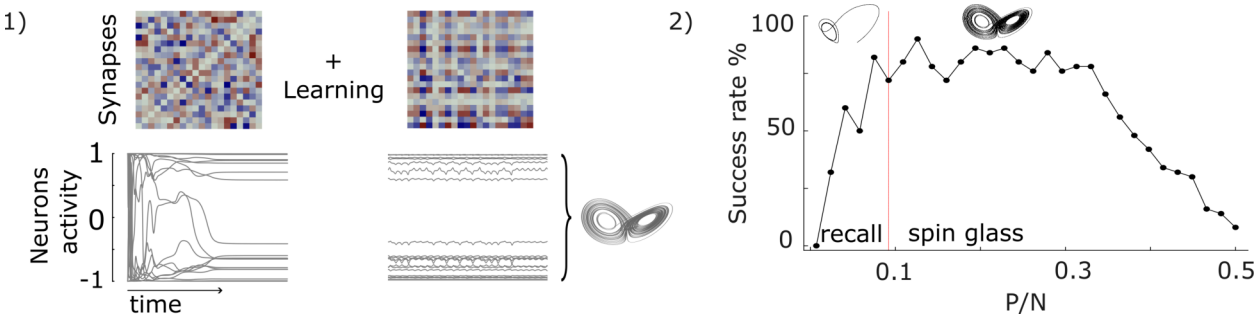

Figure 1. The Hopfield network has symmetric synapses and the neural activity relaxes to fixed point. (**1**) Learning introduces asymmetries that allows the network to autonomously reproduce the target dynamical system. (**2**) The success rate strongly depends on the number of patterns (P) stored in the synaptic matrix before learning (The number of neurons is fixed at $n = 10^3$).

Acknowledgments

Work partially funded by the Italian National Recovery and Resilience Plan (PNRR), M4C2, funded by the European Union — NextGenerationEU (Project IR0000011, CUP B51E22000150006, EBRAINS-Italy) to M.M.

References

1. Amari, 1977, *Biol Cyber*, 27:77–87, 10.1007/BF00337259
2. Hopfield, 1982, *PNAS*, 79(8):2554–2558, 10.1073/pnas.79.8.2554
3. Jaeger, 2001, *German Nat Res Ctr Inf Tech*, GMD Technical Report 148

LEVERAGING CONVOLUTIONAL NEURAL NETWORKS TO STUDY WIDE-FIELD INHIBITION IN THE RETINA

Michaela Vystrčilová[1,*], *Shashwat Sridhar*[2], *Max F. Burg*[1,3], *Tim Gollisch*[2], *Alexander S. Ecker*[1,4]

[1] Inst. Comp. Sci. and Campus Inst. Data Science, University of Göttingen, Göttingen, Germany
[2] Department of Ophthalmology, University Medical Center Göttingen, Göttingen, Germany
[3] Tübingen AI Center, University of Tübingen, Germany
[4] Max Planck Institute for Dynamics and Self-Organization, Göttingen, Germany
[*] michaela.vystrcilova@uni-goettingen.de

Understanding the retina's complex non-linear encoding mechanisms is fundamental to understand vision. A key aspect of these mechanisms is wide-field inhibition that is believed to be mediated by wide-field amacrine cells (ACs) [1]. It is, however, experimentally challenging to record from ACs directly. Hence, we propose training a digital twin of the retina that predicts neural activity with high performance, thus allowing us to study activity inhibition indirectly through the twin. Specifically, we trained a space-time separated convolutional neural network on the spiking activity of a large population of marmoset retinal ganglion cells (RGCs) recorded using multi-electrode arrays in response to naturalistic movies. We extracted stimulus features from the digital twin that drive wide-field inhibition by synthesizing input stimuli in a two step process. First, we optimized most exciting inputs (MEIs) [2] which maximize the predicted response of RGCs and reveal visual features which they are most sensitive to. Then, we fixed the center of the MEIs and optimized only their surrounds to minimize RGCs' predicted responses (Figure 1). We found surrounds suppressing up to 90% of the MEI-elicited activation for most cells, suggesting a strong surround inhibition mechanism. Furthermore, we found structure within the surrounds that revealed possible receptive fields of ACs.

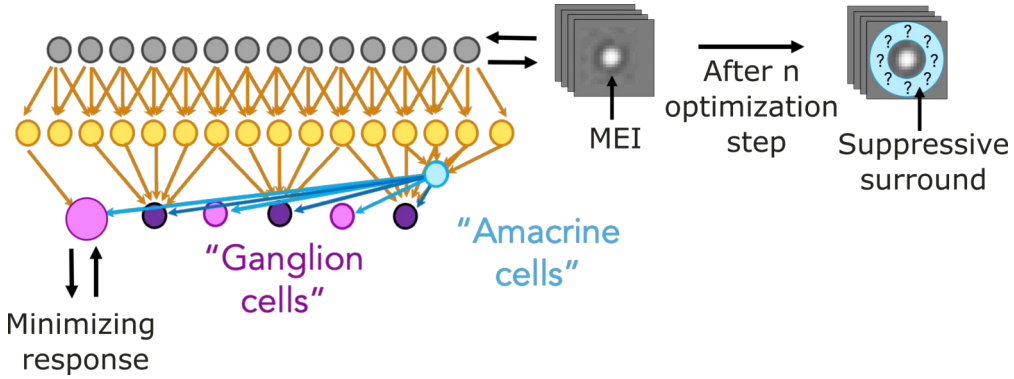

Figure 1. Suppressive surround optimization scheme.

Acknowledgments

Supported by the German Research Foundation (DFG) — project IDs 154113120 (SFB 889, project C01); 432680300 (SFB 1456, project B05); by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement numbers 724822 and 101041669). MFB thanks the IMPRS-IS.

References

1. Franke, Baden, 2017, *J Physiol*, 595(16):5507–15, 10.1113/JP273648
2. Walker, Sinz, Cobos, *et al.*, 2019, *Nat Neurosci*, 22:2060–5, 10.1038/s41593-019-0517-x

CONTEXT-DEPENDENT VALUES ENCODED BY HIPPOCAMPAL-PREFRONTAL CIRCUITS

Joni D. Wallis[*], *Thomas W. Elston*

Department of Neuroscience, University of California, Berkeley, California, USA
[*] wallis@berkeley.edu

The value of a choice option can vary depending on context. For example, a glass of ouzo before dinner may have high value, but not if one has an early flight the next morning. Thus, optimal choice depends on flexibly assigning value to choice options in a contextually appropriate manner. Despite the ubiquity of such contextual reasoning outside of the laboratory, how the brain flexibly switches between value-representations in different contexts remains unknown.

We addressed this question by performing simultaneous high-channel count recordings from the hippocampus (HPC) and orbitofrontal cortex (OFC) from two rhesus monkeys using Neuropixels and V-probes. Our task required animals to flexibly update the values of 8 choice options based on a context-cue that varied from trial-to-trial. Critically, the context-cue preceded and was not present during choice, requiring the animals to simultaneously update the 8 option values prior to seeing them.

We report evidence of a functional double-dissociation across HPC and OFC: HPC neurons (but not OFC) abstractly encoded context during the cue phase whereas OFC strongly represented context during the choice phase. In addition, whereas HPC encoded an abstract value independent of context, OFC encoded a state-dependent value. Analysis of communication through oscillatory coherence revealed a ramping HPC-to-OFC signal in the theta band (4–8 Hz) that started as HPC neurons began encoding context and peaked as the contextually appropriate OFC value subcircuit came online.

The large number of neurons that we were able to record simultaneously with the Neuropixels also greatly improved our ability to decode decision-making signals compared to our previous studies. We could decode value at around 90% accuracy (chance is 8.3%), enabling us to detect previously unobserved dynamics. For example, we observed a transient representation of the irrelevant state the strength of which predicted whether the monkey would make multiple saccades to the choice options.

Our results suggest that context information, a compressed representation of the relevant valuation schema, is initially encoded in HPC and then broadcast to OFC via theta synchronization to select a contextually appropriate value subcircuit, thus allowing for contextual reasoning in value-based choice.

Acknowledgments

Supported by NIMH R01-MH117763 and NIMH R01-MH121448.

UNCOVERING COMPLEX FEATURE TUNING DIMENSIONS IN PRIMATE AREA V4 USING DIGITAL TWIN MODELLING AND CONTRASTIVE LEARNING

Konstantin F. Willeke[1,*], *Katrin Franke*[2], *Kelli Restivo*[2], *Nikos Karantzas*[2], *Kandan Ramakrishnan*[2], *Tori Shinn*[2], *Cate Nealley*[2], *Gabrielle Rodriguez*[2], *Saumil Patel*[2], *Alexander Ecker*[1], *Fabian Sinz*[1,†], *Andreas Tolias*[2,†]

[1] Georg-August University Goettingen, Goettingen, Germany
[2] Department of Neuroscience, Baylor College of Medicine, Houston, Texas, USA
[†] equal contribution
[*] konstantin.willeke@gmail.com

Characterizing the tuning properties of visual neurons, particularly in mid- to higher level visual cortex which are selective for complex features like objects and textures, remains a challenging task. In this work, we address this challenge by focusing on a large population of 1,244 neurons, recorded with laminar probes in primate visual area V4. We trained a deep neural network model on the entire population given the responses of neurons to a large set of natural images to obtain our digital twin. We subsequently tested this model in closed loop experiments, demonstrating its ability to predict neuronal responses to previously unseen as well as model-synthesized stimuli. Our primary objective was then to employ this digital twin model to investigate the visual tuning characteristics of individual neurons. Thus, we trained a self-supervised model on an extensive dataset of over 4,000 natural textures, resulting in an interpretable two-dimensional embedding space that categorizes these textures, where small distances in the embedding between textures were also perceptually similar (Figure 1). Interestingly, neuronal population activity predicted by our V4 digital twin varied smoothly in the identified texture space, suggesting the existence of an interpretable population texture encoding. Currently, we are identifying general tuning axes in this space that faithfully capture the tuning of individual V4 neurons. Ultimately, our goal is to expand this approach to uncover novel tuning axes for the representation of complex visual features in the visual stream.

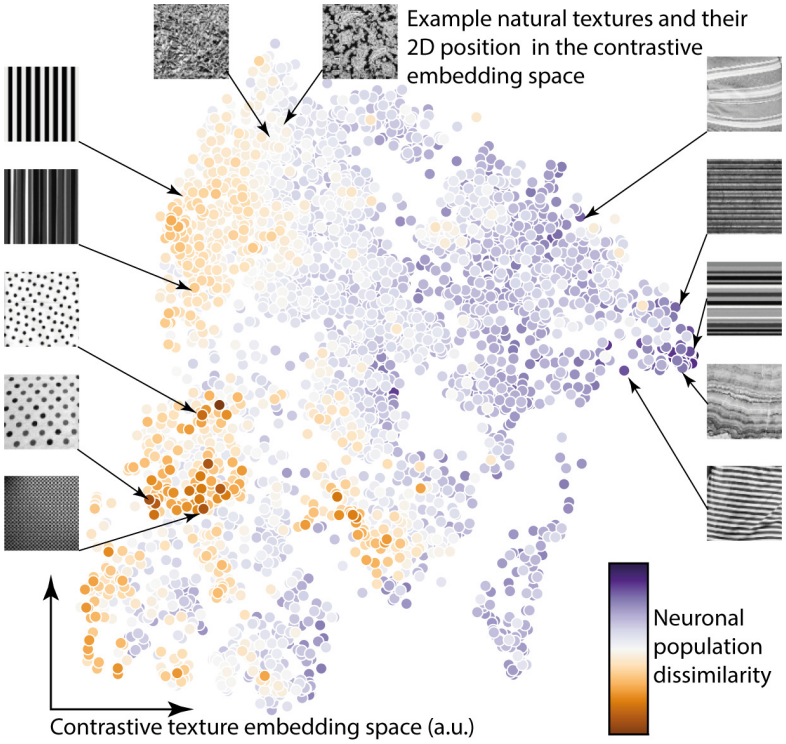

Figure 1. Overview of the approach. We first obtained a low dimensional (2D) embedding space for a large number of natural textures using contrastive learning. We then showed our V4 digital twin model all of the natural textures, to obtain the representational similarity in neuronal response space. Here we show an example for representational similarity of one target texture (darkest blue dot) compared to all other textures. Transferring this approach to single neurons, we are now able to uncover tuning axes within low dimensional representations of complex visual features.

SEGREGATED NEURONAL POPULATIONS IN PREFRONTAL CORTEX ENCODE TASK VARIABLES DURING WORKING MEMORY

Klaus Wimmer[1,*], *Bijan Pesaran*[2], *Nicolás Pollán Hauer*[1]

[1] Centre de Recerca Matemàtica, Barcelona, Spain
[2] University of Pennsylvania, Philadelphia, Pennsylvania, USA
[*] kwimmer@crm.cat

Non-linear mixed selectivity, with neurons responding to diverse combinations of task-relevant variables, has been proposed as a key mechanism to enable flexible behavior and cognition [1]. However, it is debated whether the structure of neural population responses in fronto-parietal cortices is better described as random mixed-selective [1–3] or as non-random, that is, in terms of multiple subpopulations with stereotypical response profiles [4, 5].

Here, we show that neural activity in macaque prefrontal cortex during a working memory and a visual response task is organized into subpopulations that provide a comprehensive description of the low-dimensional population dynamics. First, analysis of the demixed Principal Components shows that the neural code faithfully represents stimulus identity, task condition and elapsed time during the trial. Second, a model-free analysis of the population structure reveals a significant degree of clustering, implying a non-random distribution of feature selectivity that is incompatible with random mixed selectivity. Closer inspection shows stimulus-selective neurons also tend to be task-selective. Third, examining the contribution of stimulus-selective neurons to task condition-related variance reveals two contrasting activity profiles that correspond to functionally different populations. One population responds during visual stimulation while the other activates during memory maintenance. Finally, the observed neural geometry explains how stable task and stimulus information can be read out from the population response using a linear decoder.

Our results highlight that despite the heterogeneity of prefrontal responses during working memory, neurons do not represent random mixtures of task features but are structured according to neural subpopulations.

Acknowledgments

Supported by MCIN/AEI/10.13039/501100011033.

References

1. Rigotti, Barak, Warden, *et al.*, 2013, *Nature*, 497(7451):585–90, 10.1038/nature12160
2. Raposo, Kaufman, Churchland, 2014, *Nat Neurosci*, 17(12):1784–92, 10.1038/nn.3865
3. Dang, Jaffe, Qi, *et al.*, 2021, *J Neurosci*, 41(35):7420–34, 10.1523/jneurosci.2814-20.2021
4. Hirokawa, Vaughan, *et al.*, 2019, *Nature*, 576(7787):446–51, 10.1038/s41586-019-1816-9
5. Yang, Tipparaju, *et al.*, 2022, *Nat Neurosci*, 25(10):1339–52, 10.1038/s41593-022-01171-w

PUBLIC VS PRIVATE PERCEPTION IN EEG AND PUPILLOMETRY

Mustafa Yavuz[1,*], *Jamal Esmaily*[1,2], *Bahador Bahrami*[2], *Ophelia Deroy*[3,4]

[1] Grad. School Systemic Neurosci., Ludwig Maximilians University Munich, Munich, Germany
[2] Dept. General Psych. and Education, Ludwig Maximilians University Munich, Munich, Germany
[3] Faculty of Philosophy, Ludwig Maximilian University, Munich, Germany
[4] Munich Centre for Neuroscience, Ludwig Maximilian University, Munich, Germany
[*] yavuzmustafa11235@gmail.com

The social context influences information processing in the human brain. Here, we investigate whether the mere awareness that an object is visible to others, independently of other interactions, is sufficient to affect individual perception and if so what the neural and physiological correlates of these phenomena are.

To answer these questions, we employed a Random Dot Motions paradigm while recording EEG activity and pupil size changes. In total, 33 right-handed, healthy participants completed 640 trials with 5 levels of difficulty. To distinguish joint-perception from mere social presence, the social presence of a confederate sitting next to the participant was kept constant all through the experiment. The screen was separated in two halves by a divider. On a trial-by-trial basis, the visual stimulus appeared either on the side of the divider that was visible both to the participant and the confederate (joint perception) or on the side that was only visible to the participant (individual perception). Participants and confederates did not make eye contact or interact, meaning that the only difference from the participant's perspective was that the stimulus could be seen in common or privately.

We compared behavioral measures such as accuracy and reaction times between joint and individual perception. Similarly, metacognitive accuracy was also compared between conditions. To see whether individual and joint perception conditions represented differently at the neural level, we used Representational Similarity Analysis and compared our model dissimilarity matrices with neural dissimilarity matrices. Finally, we used regression models to predict pupil size changes in response to motion stimulus by using accuracy, reaction times, confidence judgements.

Accuracy, reaction times and metacognitive accuracy did not differ between conditions. However, RSA analysis showed significant correlation between our model (which assumes joint and individual perception represented with maximum dissimilarity in the brain) and actual brain signals. Importantly, this correlation was observed after the trial type was known and before the start of the motion stimulus. Changes in pupil size in response to the motion stimulus could be predicted by joint or individual experimental conditions, as well reaction times, confidence judgements and task difficulty.

We show that the public or private character of perception is represented at the neural level, and it is encoded in pupil behavior despite no observable behavioral difference.

PREDICTING SINGLE-TRIAL EYE SPEED AND DIRECTION FROM NEURAL POPULATION RESPONSES IN MARMOSET AREA MT

Hoi Ming Ken Yip[1,*], *Elizabeth Zavitz*[3], *Shaun L. Cloherty*[2], *Maureen A. Hagan*[1], *Nicholas S. C. Price*[1]

[1] Department of Physiology, Monash University, Clayton, Victoria, Australia
[2] School of Engineering, RMIT University, Melbourne, Victoria, Australia
[3] Department of ECSE, Monash University, Clayton, Victoria, Australia
[*] hoi.yip@monash.edu

How sensory information in the neuronal population is transformed in the brain to generate motor responses is largely unknown. Ocular following is an ultra-fast, reflexive eye movement triggered by motion of a large visual stimulus [1] which requires rapid sensory-motor transformations. Previous single-unit recordings in macaques suggest that the middle temporal area (MT) plays a critical role in generating ocular following responses [2]. However, single-neuron studies were not able to assess how neuronal variability correlates with behavioural variability on the level of single trials. In 11 recording sessions, we used Neuropixels probes to simultaneously record neurons (n = 29–139) in area MT of a marmoset. The monkey viewed broadband motion stimuli drifting in one of the eight directions (four cardinal and four oblique directions) that evoked robust ocular following.

While MT activity is known to represent visual motion properties, we wondered how well the population activity correlated with the timing and velocity of eye movements. We used partial-least-square regression (PLSR) with leave-one-out cross-validation to predict horizontal and vertical eye velocities from the population spiking rates on a trial-by-trial basis. For each trial, we extracted the mean firing rates of each neuron and mean eye velocities in eight non-overlapping 20 ms time windows (*e.g.*, 0–19 ms, 20–39 ms) after stimulus onset. We first conducted PLSR with the neuronal activities and eye velocities in the same time window. The model started to predict eye velocities well at 60–79 ms after stimulus motion onset (r = 0.58), when the eyes just started moving. Model predictions remained excellent (r = 0.72–0.74) in the subsequent time windows (80–159 ms). To examine the delay between MT neural activity and behavioural response, we used early spiking rates to predict later eye velocities. We found that eye velocities at 60–159 ms can be well predicted by neural activity at 20–39 ms (r = 0.52–0.63), which is a sensible delay. In conclusion, eye velocities can be predicted with MT population activity on a single trial level.

Acknowledgments

This work was supported by the Australian Research Council and the National Health and Medical Research Council.

References

1. Miles, Kawano, Optican, 1986, *J Neurophysiol*, 56(5):1321–54, 10.1152/jn.1986.56.5.1321
2. Kawano, Shidara, *et al.*, 1994, *J Neurophysiol*, 71(6):2305–24, 10.1152/jn.1994.71.6.2305

ATTENDEE AND AUTHOR INDEX

Entries for attendees are listed with last name, first name, affiliation, and email. Page numbers in bold (**00**) indicate presenting authorship, in italics (*00*) indicate session moderation, underlined (<u>00</u>) indicate organizing committee reference, and in normal typeface (00) indicate non-presenting authorship or other reference in the program text.