Neuro-cybernetics at scale

Symposium

15 → 17 · 10 · 2025

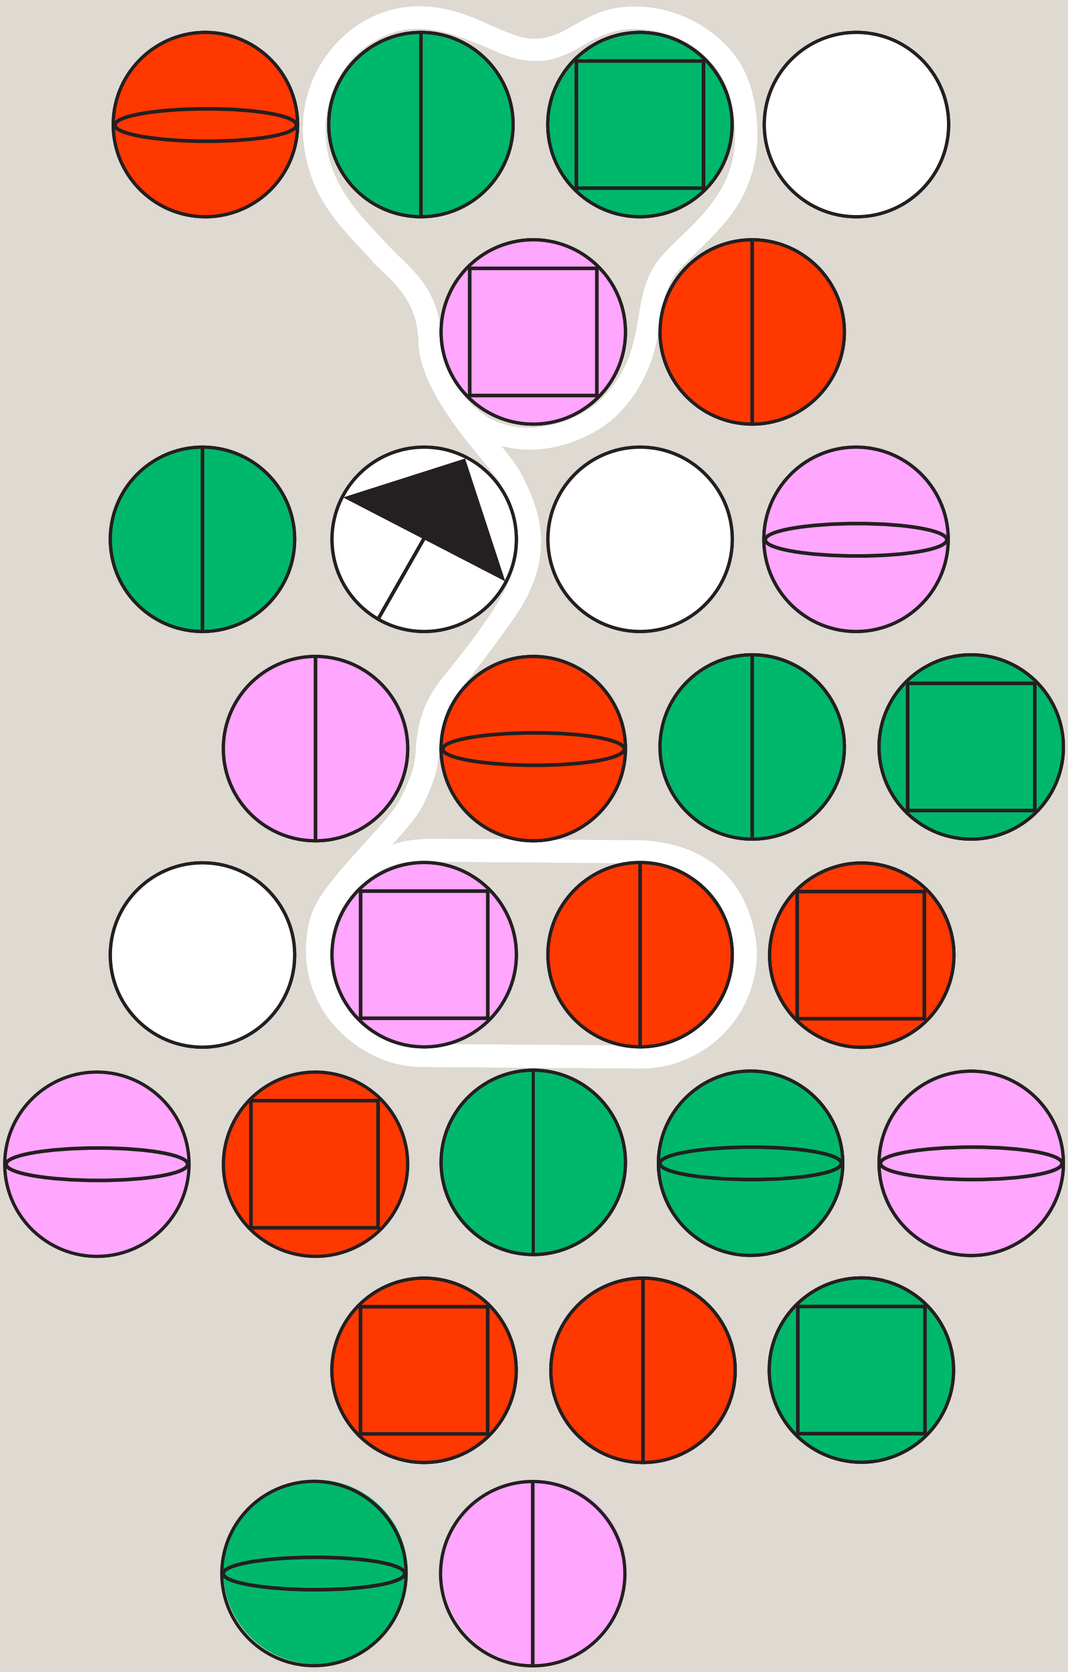

CHAMPALIMAUD RESEARCH SYMPOSIUM

Protocol Labs

partners

initiative

Champalimaud Foundation

support

Protocol Labs

SPN
SOCIEDADE PORTUGUESA DE NEUROCIÊNCIAS

DELTA CAFÉS

welcome

This year's Champalimaud Research Symposium, Neuro-Cybernetics at Scale, explored how behavior emerges through complex feedback loops across neural circuits, bodies, and varied environments. While reductionist approaches have brought deep insights into individual components, understanding behavior as it unfolds in the real world demands a shift toward systems-level thinking—across scales, disciplines, and modalities.

The symposium drew inspiration from the recent advances in AI and machine learning, where scaling has unlocked unprecedented performance. We believe neuroscience may be approaching a similar inflection point. By convening researchers in experimental neuroscience, robotics, machine learning, control theory, and theoretical neuroscience, we aimed to explore how ideas of scaling and feedback can shape the future of systems neuroscience.

This single-track scientific meeting featured four keynote speakers and a lineup of esteemed invited speakers, along with selected talks from submitted abstracts. Attendees had the opportunity to showcase their research during two poster sessions, all while enjoying various networking and social activities.

The Symposium Chairs,

Guillaume Hennequin
Memming Park
Shreya Saxena

numbers

29 speakers

320 participants

72 posters

> 40 institutions

of 11 countries

Bangladesh, Canada, France, Germany, Ireland, Israel, Japan, Portugal, Switzerland, the United Kingdom, the United States of America

team

Events
António J. Monteiro
Ana Rita Mendes
Andreia Pinho
Diana Cadete
João Van Zelst
Thaïs Lindemann
Teresa Fernandes

Communication
Ana Rita Mendes
Andreia Pinho
Catarina Ramos
Diana Cadete
Hedi Young
John Lee
Thaïs Lindemann

Design
Carla Emilie Pereira
Marta Correia

Sound design
João Van Zelst

Chairs
Guillaume Hennequin
Memming Park
Shreya Saxena

Volunteers
Ana Raquel Augusto
Angeline Wang
Hugo Rafael
Isabela Pereira
Karthik Sama
Margarida Seabra Gomes
Sofia Pereira da Silva
Tomás Caldeira

programme

at a glance

wed — 15
scaling up

thu — 16
cybernetic loops

fri — 17
models
and robots

at a glance

programme

	Registration & welcome coffee		
9		Registration & welcome coffee	Registration & welcome coffee
	Opening remarks	Adrienne Fairhall ⊖	Eva Dyer ⊖
10	Misha B. Ahrens ⊖		
		Taro Toyoizumi ▢	Jesse Marshall ▢
11	Coffee break	Coffee break	Coffee break
	Juan Álvaro Gallego ▢	Jonathan W. Pillow ▢	Josue Ortega Caro Christoph Hemmer ⦶
12	Shreya Saxena ▢	Juan R. Castiñeras Jeffrey Walker ⦶	Panel discussion Future of Neurocybernetics
	Ayesha Vermani T. Anderson Keller ⦶	Reece Keller Ivan Voitov ⦶	
13	Lunch	Lunch + Meet the expert	Lunch
14			
	Poster session	Poster session	Srinivas C. Turaga ▢
15			Auke Ijspeert ⊖
16	Coffee break	Coffee break	Coffee break
	Ehud Ahissar ▢	Eugenia Chiappe ▢	Caroline Haimerl Kabir Dabholkar ⦶
17	Bing W. Brunton ▢	Máté Lengyel ▢	Christopher J. Rozell ▢
	Rui Xia Rich Pang ⦶	Sofia Freitas Leonardo Agueci ⦶	Closing remarks
18	Champalimaud tours	Walk & talk sunset drinks	Boat tour
19			
			Farewell dinner sponsored by Protocol Labs
20			

⦶ selected speaker ▢ invited speaker ⊖ keynote speaker

programme
wed — 15

8:30	Registration & welcome coffee
9:30	**Opening Remarks** **Leonor Beleza** President of the Champalimaud Foundation (CF)
10:00	**Misha B. Ahrens** Whole-brain and whole-body computations for behavior and physiology
11:00	Coffee break
11:30	**Juan Alvaro Gallego** Neural manifolds, invariances, and meaning
12:00	**Shreya Saxena** Computational approaches towards cognitive motor control
12:30	**Ayesha Vermani** Charting the space of dynamics with hierarchical state-space models **T. Anderson Keller** Flow equivariant cybernetics
13:00	Lunch
14:30	**Poster session**
16:00	Coffee break
16:30	**Ehud Ahissar** Active embodied perception through brain-world loops
17:00	**Bing W. Brunton** Embodied intelligence: from connectomes to body models
17:30	**Rui Xia** A multi-area RNN model of adaptive motor control explains adaptation-induced reorganization of neural activity **Rich Pang** Learning neural dynamics through instructive signals
18:00	**Champalimaud Tours**

scaling up

wed — 15

keynote lectures

Whole-brain and whole-body computations for behavior and physiology

Misha B. Ahrens

Janelia Research Campus, Howard Hughes Medical Institute
Ashburn, Virginia, USA

Neural computation underlying sensation, information accumulation, memory, decisions, and actions are distributed across many brain areas that exhibit intricate local and long-range connectivity. Understanding the brain's global dynamics may depend on observing and modeling the system as a whole. To study neuronal dynamics simultaneously across all brain regions at cellular resolution, we developed a whole-brain light-sheet imaging system for larval zebrafish as they behave in virtual reality environments. As the animals swim through virtual arenas, they evaluate the outcomes of their actions and, and as they learn from those, adjust their future actions. Whole-brain imaging and neuronal perturbations during such behavior revealed a role for multiple neuromodulatory systems and non-neuronal cells called radial astrocytes in integrating behavioral outcomes to drive motor learning and adaptive changes in behavioral states. To study longer-term learning, we developed a robotic system where fish learn from interactions with virtual agents, in which we discovered that learning occurs through noradrenergic reinforcement signals impinging on forebrain dynamics. To explain such brain dynamics through neuronal connectivity, we developed simultaneous functional imaging and whole-brain connectomics, allowing for novel computational analyses of structure-function relationships. Finally, since bidirectional connections between the brain and organ systems form an integral component of neural computation, we expanded into whole-body imaging to study the relationship between brain dynamics and organ physiology at the scale of the entire organism.

Juan Álvaro Gallego

Shreya Saxena

Ehud Ahissar

Bing W. Brunton

wed — 15

talk sessions

Ayesha Vermani

T. Anderson Keller

Rui Xia

Rich Pang

Neural manifolds, invariances, and meaning

Juan A. Gallego

Champalimaud Foundation
Lisbon, PT

Computational approaches towards cognitive motor control

Shreya Saxena

Department of Biomedical Engineering and Wu Tsai Institute
Yale, United States

Active embodied perception through brain-world loops

Ehud Ahissar

Weizmann Institute of Science
Rehovot, Israel

Mammalian sensation has evolved to be coupled with movement: we actively move our eyes, hands, whiskers, and other sensors to explore the world. Active sensing could, in principle, operate in open- or closed-loop modes, but growing evidence shows that it functions in a closed loop, where sensation and movement continuously shape one another. From this perspective, perception is not a series of static snapshots but an ongoing convergence in which external objects become embedded in brain activity. Research further suggests that perception unfolds in two complementary domains: a physical, analog brain–world (BW) loop, and a mental, digital brain–brain (BB) loop. I will outline this framework and illustrate it with three examples: rodents detecting objects with whiskers, humans recalling objects through language, and artificial networks classifying tiny images.

Embodied intelligence: From connectomes to body models

Bing Wen Brunton

Professor & Richard and Joan Komen University Chair
University of Washington
Seattle, USA

Brains and nervous systems evolved jointly with the bodies they inhabit. Thus, NeuroAI stands to benefit from embracing the embodied nature of neural computations, because I believe such a perspective will be crucial to understanding brains, behavior, and intelligence. In this talk, I will describe some ongoing collaborative work towards putting dynamic models of brains as controllers inside the bodies they evolved to sense and control. Such virtual animals will be capable of interacting with simulations of their environment. These efforts are currently closest to fruition in the fruit fly. Here, we are working with collaborators to develop tools to interface simulations of connectomes — at the resolution of cells and synapses — in feedback with muscle actuators, neural transduction sensors, and whole-animal biomechanics.

As an example of the promise of such an approach to gain new scientific insights, I will present our recent work developing connectome simulations of the fly ventral nerve cord (VNC) to identify a putative central pattern generator (CPG) circuit for walking. We used these simulations as a platform to perform a series of computational screens that are impossible to do in a biological experiment, which allowed us to computationally isolated a three-neuron core CPG circuit that is sufficient to generate leg motor rhythms. In ongoing work, integrating simulations of connectomes with physically realistic body models will allow us to investigate active behavioral sequences and how specific recurrent neural circuits support robust, embodied intelligence in animals.

Charting the space of dynamics with hierarchical state-space models

Ayesha Vermani [1]

Josue Nassar [2]

Justus Kautz [1]

Hyungju Jeon [1]

Matthew Dowling [1]

Memming Park [1]

[1] Champalimaud Foundation, Lisbon, PT
[2] Ryvivyr, NY

Understanding neural computation through dynamics has provided valuable insights into a range of cognitive behaviors. As such, numerous methods have been developed to learn dynamics directly from neural population recordings. However, these methods are typically applied to individual recording sessions and do not readily capture the broad diversity of dynamical repertoires that support flexible and adaptive behavior. In this work, we introduce meta-dynamical state-space models, a hierarchical framework that learns a solution space of dynamics from neural recordings of task-trained animals. Our approach captures variabilities across heterogeneous recordings on a low dimensional manifold which concisely parametrises a family of dynamical systems. We demonstrate the efficacy of our approach in modelling motor cortex recordings spanning multiple sessions, subjects and tasks, and show that it enables sample efficient transfer to novel recordings. Additionally, we illustrate the applicability of this formulation for investigating changes in dynamics over the course of learning, as well as unsupervised discovery of dynamics underlying inter-individual variability. Together, these results highlight the power of learning and reasoning over a structured space of dynamical systems to probe the complexity and functional diversity of biological neural networks.

dynamical systems / state-space models / meta-learning

Flow equivariant cybernetics

T. Anderson Keller

The Kempner Institute For the Study of Natural & Artificial Intelligence at Harvard University, Cambridge, USA

Cybernetic systems experience the world as ‘sensorimotor flows’: continuous streams of sensory input entrained to self-motion. These streams obey smooth, time-parameterized symmetries (e.g. expanding or rotating optic flow), yet most embodied neural network sequence models ignore this structure, and instead laboriously re-learn the same transformations from data. In this work, we introduce ‘Flow Equivariant Cybernetics’, a framework in which a recurrent controller’s hidden state evolves in the co-moving reference frame of the agent’s own motion, driven by both proprioceptive and exteroceptive feedback. Online motion estimates and efference copies drive a coordinate transformation that shifts the latent state of the controller network to flow in step with the environment, thereby yielding provable equivariance to the agents own self-motion [1]. This feedback-modulated reparameterization eliminates the overhead of traditional equivariant neural networks [2], sharply reducing memory usage, and — crucially — cuts the amount of training data required to master long-horizon prediction compared with non-equivariant counterparts. On a 3D maze-navigation benchmark that demands accurate world models [3], flow equivariantagents learn with an order of magnitude fewer samples and consequentially outperform comparable transformer-based world-modeling architectures. The flow equivariant update rule remains stable over hundreds of future rolled-out time-steps, generating a drift-free latent map that tracks the environment regardless of ego-motion —mirroring the efficiency and resilience of biological navigation circuits. By entraining sensory prediction with body-centric feedback, flow equivariance charts a scalable route to data-efficient embodied intelligence and offers systems neuroscience a symmetry-based blueprint for next-generation cybernetic systems.

dynamical systems / state-space models / meta-learning

References

[1] Keller, T. A. (2025). Flow Equivariant Recurrent Neural Networks. Under Review, NeurIPS ‘25;
[2] Cohen, T. S., & Welling, M. (2016). Group equivariant convolutional networks. ICML ‘16;
[3] Beattie, C., et al. (2016). DeepMind Lab. arXiv:1612.03801.

A multi-area RNN model of adaptive motor control explains adaptation-induced reorganization of neural activity

Rui Xia [1]

Marine Schimel [2]

Guillaume Hennequin [1]

[1] University of Cambridge, Cambridge, UK
[2] Stanford University, California, US

Adaptation is essential for biological motor control, enabling rapid modifications of control strategies as the environment changes. Along with the recent development of neuroscience experiments—from simplistic task and single-area neural data to simultaneous large-scale recordings of multiple brain areas during naturalistic behaviour — we now have the opportunities to investigate adaptive motor control which requires system-wide computations. To capture the intricacies of such behaviour, we propose a unified network-level theory through coordinated, multi-area dynamics in the context of motor adaptation.

We developed a multi-area RNN where the major components of motor control – state estimation, planning and feedback control are continuously governed by a higher-level contextual module. Specifically, the network comprises an internal forward model and a planning/control policy network. Critically, the dynamics of both areas, as well as their interactions, are orchestrated by a contextual inference network via gain modulation of trans-thalamic loops, implementing low-rank modifications of effective connectivity.

The network is meta-trained to maintain proficient control of hand reaches in the face of changing environmental dynamics (force fields, FF). Behaviorally, the network exhibits successful few-shot learning of new contexts, with both state estimation and motor preparation/control improving rapidly within few interactions with the environment.

On the level of neural population activities, we observe a systematic reorganization of preparatory activity following FF changes. The ring-structured manifold of preparatory states associated with center-out reaching targets rotates in a way consistent with compensatory re-aiming, an effect most pronounced for reach directions close to the adapted direction. Moreover, adaptation causes this ring manifold to shift along a third dimension that separates the various FF contexts.

These findings closely align with experimental data from monkey primary motor cortex [1].

Together, our model provides a promising framework for integrating neural and behavioral data to advance our understanding of adaptive motor control.

References
[1] Sun, X., O'Shea, D. J., Golub, M. D., Trautmann, E. M., Vyas, S., Ryu, S. I., & Shenoy, K. V. (2022). Cortical preparatory activity indexes learned motor memories. Nature, 602(7896), 274-279.

Learning neural dynamics through instructive signals

Rich Pang

Juncal Arbelaiz

Jonathan Pillow

Princeton University, Princeton, NJ, USA

Computation in the brain depends on coordinated neural activity dynamics and is often studied in flexible, brain-like dynamics models such as recurrent neural networks (RNNs). Neural dynamics also change through learning via synaptic plasticity, sometimes over fast timescales. RNNs, however, are usually trained through either Hebbian rules ("neurons that fire together wire together") or artificial rules like backpropagation through time (BPTT) [1-6]. Hebbian rules are bioplausible and can create specific useful dynamics (e.g. fixed-point attractors that recall memories) but fail to create more flexible dynamics, much less over rapid timescales; BPTT is flexible but not likely bioplausible, limiting its relevance to the brain. How RNNs dynamics emerge from biological plasticity rules to support rapid learning is unknown. Here we present a novel plastic RNN inspired by rapid, non-Hebbian plasticity rules in hippocampus, cerebellum, and mushroom body. These rules, which we term "HIP" (heterosynaptic instructed plasticity), change synapses based on presynaptic activity and an "instructive signal" from another brain area [7-12]. As in the brain, we apply HIP not directly to the RNN, but to a mushroom-body-like feedback circuit coupled to the RNN. We show that via HIP instructive signals can rapidly teach the RNN highly flexible nonlinear dynamics, with precise mathematical interpretation inspired by support vector machines: HIP builds flow fields out of "support states" that can be composed into highly flexible nonlinear dynamics. Via adaptive control theory, we also show how HIP can teach RNNs to learn target dynamics from error signals. Finally, HIP can teach RNNs long-range dependencies that challenge BPTT [13]. Thus, RNNs can rapidly learn flexible dynamics via instruction-mediated plastic feedback circuits, shedding new light on brain function and neurocybernetics.

rapid learning / feedback / recurrent neural network / dynamic / mushroom body

References

1 Barak 2017 Curr Opin Neurobiol.
2 Hopfield 1982 PNAS.
3 Brunel 2016 Nat Neurosci.
4 Werbos 1990.
5 Sussillo & Abbott 2009 Neuron.
6 Driscoll et al Nat Neurosci 2024.
7 Bittner et al 2017 Science.
8 Milstein et al 2021 eLife.
9 Suvrathan et al 2016 Neuron.
10 Silva et al 2024 Nat Neurosci.
11 Handler et al 2019 Cell.
12 Jiang & Litwin-Kumar 2020 PLoS Comput Biol.
13 Pascanu et al 2013 arXiv.

programme

thu — 16

9:00	Registration & welcome coffee
9:30	**Adrienne Fairhall** Brain-wide dynamics of internal model update
10:00	**Taro Toyoizumi** Controlling neural variability
11:00	**Coffee break**
11:30	**Jonathan W. Pillow** Disentangling the roles of distinct cell classes with cell-type dynamical systems
12:00	**Juan Castiñeras** Optimal control of spiking neural networks **Jeffrey Walker** Building a whole-body marmoset musculoskeletal model for simulating ethologically-relevant natural behavior
12:30	**Reece Keller** Autonomous behavior and whole-brain dynamics emerge in embodied zebrafish agents with model-based intrinsic motivation **Ivan Voitov** Intrinsic dynamics underlying flexible sequence generation in the hippocampus
13:00	**Lunch + Meet the expert**
14:30	**Poster session**
16:00	**Coffee break**
16:30	**Eugenia Chiappe** Active gaze control enables competitive pursuit in Drosophila
17:00	**Máté Lengyel** Spatial uncertainty determines how navigation and grid cell tuning are shaped by environmental geometry
17:30	**Sofia Freitas** Temporal control at multiple scales: distinct behavioral and neural signatures of when and how to act **Leonardo Agueci** Distributed learning across fast and
18:00	**Walk & talk + sunset drinks**

cybernetic loops

keynote lectures

Brain-wide dynamics of internal model update

Adrienne Fairhall

University of Washington, Seattle
Washington, USA

Taro Toyoizumi

Jonathan W. Pillow

Eugenia Chiappe

Máté Lengyel

talk sessions

Juan R. Castiñeras

Jeffrey Walker

Reece Keller

Ivan Voitov

Sofia Freitas

Leonardo Agueci

Controlling neural variability

Taro Toyoizumi

RIKEN Center for Brain Science
Wako, Japan

Variability is an intrinsic feature of brain activity. In this talk, I present modeling works on how neural variability can be controlled and exploited. First, I introduce a model that characterizes the role of closed-loop sensory feedback via the body and environment. This sensory feedback—conveying the consequences of actions—regulates neural variability, when sensory input alone is insufficient to constrain neural responses due to ongoing internal dynamics. Second, I show that generating neural variability can facilitate Bayesian computation: our model learns to actively induce chaos-driven variability to perform robust Monte Carlo sampling from the Bayesian posterior distribution.

Disentangling the Roles of Distinct Cell Classes with Cell-Type Dynamical Systems

Jonathan W. Pillow

Princeton University, Princeton
New Jersey, USA

Latent dynamical systems have been widely used to characterize the dynamics of neural population activity in the brain. However, these models typically ignore the fact that the brain contains multiple cell types, which limits their ability to capture the functional roles of distinct cell classes or predict the effects of cell-specific perturbations. To overcome these limitations, we introduce the "cell-type dynamical systems" (CTDS) model, which extends latent linear dynamical systems to contain distinct latent variables for each cell class, with appropriate sign constraints on the interactions between them. In this talk, I will describe the CTDS model and show that fitting in the noiseless case can be reduced to non-negative matrix factorization. I will then show an application of a multi-region model CTDS to simultaneous recordings from rat frontal orienting fields (FOF) and anterior dorsal striatum (ADS) during an auditory decision-making task. Remarkably, the model — fit only to unperturbed neural activity — predicts the time-dependent effects of different optogenetic perturbations on behavior, specifically in FOF, ADS, and FOF-to-ADS axon terminals. I will close by discussing the future directions and other applications for biologically-constrained dynamical models of neural activity and behavior.

Active Gaze Control Enables Competitive Pursuit in Drosophila

Eugenia Chiappe

Sensorimotor Integration Lab,
Champalimaud Foundation
Lisbon, Portugal

The ability to move effortlessly through complex environments, interacting with conspecifics, avoiding predators, or simply locating resources, depends on maintaining a stable gaze. During locomotion, gaze stability relies on the coordinated movements of the eyes, head, and body, orchestrated by distributed neural activity within the central nervous system, an orchestration that remains poorly understood. In this talk, I will describe how we have leveraged Drosophila courtship behavior to study gaze control as males swiftly pursue a female. Using a combination of quantitative analysis of behavior, circuit manipulations, and physiology, we found that males actively move their heads to keep females within a critical region of their visual field essential for maintaining pursuit. This active gaze control depends on brain-body circuit loops that coordinate head rotations with body translations, shaping the male's pursuit strategy and their competitiveness under naturalistic courtship conditions. These findings reveal a critical role for head-body coordination in enabling rapid control of directed walking and provide an opportunity to study how brain-body loops are orchestrated to meet the competing demands of pursuit under biomechanical constraints and unpredictability in the females' behavior.

Spatial uncertainty determines how navigation and grid cell tuning are shaped by environmental geometry

Máté Lengyel

University of Cambridge & Central European University Cambridge, UK & Budapest, Hungary

Variations in environmental geometry, such as the shape and size of an enclosure, have profound effects on navigational behaviour and its neural underpinning in grid cells. Here, we show that these effects arise as a consequence of a single, unifying principle: to navigate efficiently, the brain must maintain and update the uncertainty about one's location. We developed an image-computable Bayesian ideal observer model of navigation, continually combining noisy visual and self-motion inputs, and a neural encoding model representing the spatial uncertainty computed by the ideal observer. Through mathematical analysis and numerical simulations, we show that the ideal observer's spatial uncertainty accounts for a diverse range of sometimes paradoxical distortions of human homing behaviour in anisotropic and deformed environments. Moreover, the neural encoding of this uncertainty accounts for distortions of grid cell responses under identical environmental manipulations. Our results suggest that spatial uncertainty plays a fundamental role in navigation.

Optimal Control of Spiking Neural Networks

Juan R. Castiñeras *

Tiago Costa *

Alfonso Renart

* equal contribution

Champalimaud Foundation, Lisbon, PT

Flexible cognitive tasks, such as covert attention and brain-machine interfaces (BMIs), require steering local neural circuit dynamics using context-dependent inputs. Control theory offers a natural language for describing these phenomena, particularly optimal control, which seeks inputs maximizing long-term value. However, the general optimal control problem for neural networks is often intractable, limiting current approaches to simplified settings (e.g., Linear-Quadratic Regulators[1]). To address this, we present a general mathematical framework for optimal control of recurrent networks of stochastic spiking neurons with low-rank connectivity. A key feature is a control cost that encourages leveraging the network's default dynamics by penalizing deviations from it. We derive a Hamilton-Jacobi-Bellman equation that specifies a Value function over the low-dimensional network state (LDS), and a corresponding optimal control law. This partial differential equation is solved using Deep Galerkin Methods[2].

The resulting law implements a feedback controller, providing external excitatory (inhibitory) synaptic input to the neurons if their spiking activity tends to move the LDS towards regions of higher (lower) Value. We apply this framework to state-steering tasks, mimicking BMIs, and optimal perceptual decision-making problems in which the agent has access to a belief over the correct choice. Our results provide the foundation of a novel approach with broad applicability that unifies bottom-up and top-down perspectives on neural computation. One limitation of this approach is that it is not directly suited to model settings with stochastic input to the network state; we are currently addressing this issue using methods like Evolutionary Strategies, which are more robust to non-Markovian settings.

Optimal Control Theory / Spiking Neural Networks / Deep Learning / Brain Computer Interfaces / Decision Making

References

[1] Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. Transactions of the ASME--Journal of Basic Engineering, 82(1), 35–45. https://doi.org/10.1115/1.3662552;
[2] Sirignano, J., & Spiliopoulos, K. (2018). DGM: A deep learning algorithm for solving partial differential equations. Journal of Computational Physics, 375, 1339–1364. https://doi.org/10.1016/j.jcp.2018.08.029

Building a whole-body marmoset musculoskeletal model for simulating ethologically-relevant natural behavior

Jeffrey Walker [1]

Jason MacLean [2]

Nicholas Hatsopoulos [2]

Shreya Saxena [1]

[1] Yale University, New Haven USA
[2] University of Chicago, Chicago USA

Understanding the neural control of natural primate behavior requires explicit consideration of the biomechanics of the body. Yet models of motor control often neglect the musculoskeletal system, omitting a critical component of behavioral execution. As part of an ongoing effort to address this gap, we develop a whole-body marmoset musculoskeletal model for neuromuscular control simulations to complement experimental studies of sensorimotor control in freely moving marmosets. We construct the model using microCT-derived skeletal geometry in the MuJoCo physics simulator[1] and design tasks in a custom marmoset domain built on the DeepMind Control Suite [2]. Kinematic recordings from marmosets performing prey capture are used to train the model with deep reinforcement learning and imitation learning techniques. We define a reward function that penalizes the distance between recorded and simulated hand positions and apply the Deep Deterministic Policy Gradient (DDPG) algorithm [3] to train multilayer perceptron policy and value networks. Our framework learns to track hand trajectories from both simple and complex episodes of prey capture, including dynamic pursuit with multiple sequential corrective movements. We find that exponential reward scaling improves tracking fidelity over simpler linear rewards. This work provides a foundation for developing biologically inspired control architectures. It enables direct comparison between sensorimotor cortical activity recorded in behaving marmosets and neural activations of artificial agents solving the same

tasks. The simulation framework offers unique access to feedback control variables such as sensorimotor delays, observation noise, and full training history, which are difficult or impossible to measure in biological systems. As such, it offers a powerful complement to experimental approaches for understanding sensorimotor control of natural behavior.

sensorimotor contro / musculoskeletal modeling / natural behavior / marmoset

References

[1] Todorov, E., Erez, T., & Tassa, Y. (2012). MuJoCo: A physics engine for model-based control. IEEE International Conference on Intelligent Robots and Systems, 5026–5033;
[2] Tunyasuvunakool, S., Muldal, A., Doron, Y., Liu, S., Bohez, S., Merel, J., Erez, T., Lillicrap, T., Heess, N., & Tassa, Y. (2020). dm_control: Software and tasks for continuous control. Software Impacts, 6, 100022;
[3] Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., & Wierstra, D. (2016). Continuous control with deep reinforcement learning. arXiv:1509.02971.

Autonomous Behavior and Whole-Brain Dynamics Emerge in Embodied Zebrafish Agents with Model-based Intrinsic Motivation

Reece Keller [1]

Alyn Tornell [1]

Felix Pei [1]

Xaq Pitkow [1]

Leo Kozachkov [2]

Aran Nayebi [1]

[1] Carnegie Mellon University, Pittsburgh, USA
[2] IBM Thomas J. Watson Research Center, Yorktown Heights, USA

Autonomy is a hallmark of animal intelligence, enabling adaptive and intelligent behavior in complex environments without relying on external reward or task structure. Existing reinforcement learning approaches to exploration in sparse reward and reward-free environments, including a class of methods known as intrinsic motivation, exhibit inconsistent exploration patterns and thus fail to produce robust autonomous behaviors observed in animals. Moreover, systems neuroscience has largely overlooked the neural basis of autonomy, focusing instead on experimental paradigms where animals are motivated by external reward rather than engaging in unconstrained, naturalistic and task-independent behavior. To bridge these gaps, we introduce a novel model-based intrinsic drive explicitly designed to capture robust autonomous exploration observed in animals. Our method (3M-Progress) motivates naturalistic behavior by tracking divergence between the agent's current world model and an ethological prior. We demonstrate that artificial embodied agents trained with 3M-Progress capture the explainable variance in behavioral patterns and whole-brain neural-glial dynamics recorded from autonomously-behaving larval zebrafish, introducing the first goal-driven, population-level model of neural-glial computation. Our findings establish a computational framework connecting model-based intrinsic motivation to naturalistic behavior, providing a foundation for building artificial embodied agents with animal-like autonomy.

neuroA / embodied A / reinforcement learning / neural-glial computation / naturalistic behavior

Intrinsic dynamics underlying flexible sequence generation in the hippocampus

Ivan Voitov

Atheer Musad

Charan Santhira-segaran

Abhishek Shah

Attila Losonczy

Columbia University, New York, USA

The hippocampus supports sequential neural activity that spans multiple timescales and flexibly reorganizes across contexts. Recurrent circuitry in area CA3 has been implicated in generating sequential activity, but it remains unclear how distinct patterns of activity arise from fixed network dynamics. Here, we used two-photon targeted optogenetic stimulation to activate CA3 neurons in precisely timed sequences with varied orders and temporal lags, allowing us to map the intrinsic interactions shaping network dynamics. We uncovered dense interactions whose strength and sign – facilitatory or inhibitory – changed gradually with stimulation lag at tens-of-milliseconds resolution. We sampled these empirical interactions to construct a synaptic-kernel RNN, whose dynamics exhibited dense input-dependent phase transitions that enabled it to generate context-specific sequences. To test whether these dynamics support sequences at behavioral timescales, we performed closed-loop experiments by stimulating at millisecond timescales the place field sequences observed during spatial navigation. Intrinsic interactions were stronger for behaviorally relevant sequences and reorganized rapidly upon remapping to novel environments. Furthermore, local network recruitment was highly specific to the active manifold location, a key prediction from our model. Together, our results demonstrate that fast, context-sensitive intrinsic dynamics in CA3 generate reservoirs of flexible spatiotemporal sequences, bridging millisecond-scale network interactions with behavioral timescale computation in the hippocampus.

hippocampus / optogenetics, sequences / two-photon imaging / neural dynamics

Temporal control at multiple scales: distinct behavioral and neural signatures of when and how to act

Sofia Freitas

Mafalda Valente

Alfonso Renart

Joseph J. Paton

Attila Losonczy

Champalimaud Foundation, Lisbon, PT

The ability to act at the right time and with appropriate intensity is central to adaptive behavior. This study investigates the mechanisms underlying timing and vigor in action control, focusing on their independence and neural substrates. Using a Fixed Interval (FI) lever-pressing task, rats were trained to press a lever for water rewards delivered only after a predetermined interval had elapsed.

Behavioral analysis revealed distinct patterns: rats initiated pressing behavior before the interval elapsed (timing), followed by a sustained, constant pressing rate until reward delivery (vigor). These behaviors were quantified as the transition point (a measure of timing) and the pressing rate (a measure of vigor), which align with theoretical distinctions between phasic and tonic dopamine.

Systematic manipulations of the task environment provided evidence for the independence of timing and vigor mechanisms. Increasing the FI duration modulated the transition point without affecting pressing rate, indicating timing-specific adaptations. In contrast, altering reward magnitude and the presence of a click cue marking interval completion selectively affected pressing rate while leaving the transition point unchanged, implicating vigor-related processes.

To explore the neural correlates of these behaviors, we simultaneously recorded dopamine signals using fiber photometry in the dorsomedial striatum and neural ensemble activity using Neuropixels probes in freely moving rats performing the task. This dataset revealed distinct patterns of activity corresponding to timing and vigor, supporting their mechanistic dissociation.

This work provides new insights into the principles of temporal control, demonstrating that when to act and how vigorously to act are governed by independent processes. These findings have broad implications for understanding action regulation and its disruption in neurodegenerative and psychiatric disorders.

Temporal control / fixed interval task / dopamine / timing and vigor / photometry / electrophysiology / striatum

Distributed learning across fast and slow neural systems supports efficient motor adaptation

Leonardo Agueci

N. Alex Cayco Gajic

Laboratoire de Neurosciences Cognitives et Computationnelles, INSERM U960, Département D'Etudes Cognitives, Ecole Normale Supérieure, PSL University, Paris, France

Adaptation is a fundamental component of motor learning, allowing organisms to adjust to environmental changes while preserving stable motor memories. A long-standing hypothesis suggests that this balance is achieved through the interaction of learning systems operating at different timescales [1,2], with the cerebello-cortical loop playing a key role [3]. However, the mechanisms that coordinate fast adaptation with slower, lasting consolidation remain poorly understood.

Here, we introduce a distributed learning model in which adaptation and memory are implemented by two distinct but interacting neural populations. One population rapidly adjusts its activity in response to perturbations, while the other gradually integrates these changes to form stable representations. Fast learning generates predictive signals that guide plasticity in the slower system through biologically plausible local rules. This framework offers a mechanistic perspective on how the cerebellum may support rapid error-based learning and instruct long-term consolidation in cortical circuits during motor adaptation.

Our model generalizes established theories of cerebellar function [4,5], aligns with recent experimental findings on cerebellar characterization [6], and makes testable predictions that further clarify the cerebellum's role in motor learning and memory.

References

[1] Smith, M. A., Ghazizadeh, A., & Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. PLOS Biology, 4(6), e179. https://doi.org/10.1371/journal.pbio.0040179;

[2] Ashby, F. G., Turner, B. O., & Horvitz, J. C. (2010). Cortical and basal ganglia contributions to habit learning and automaticity. Trends in Cognitive Sciences, 14(5), 208–215. https://doi.org/10.1016/j.tics.2010.02.001;

[3] Wagner, M. J., Kim, T. H., Kadmon, J., Nguyen, N. D., Ganguli, S., Schnitzer, M. J., & Luo, L. (2019). Shared cortex–cerebellum dynamics in the execution and learning of a motor task. Cell, 177(3), 669–682.e24. https://doi.org/10.1016/j.cell.2019.02.019

[4] Wolpert, D. M., Diedrichsen, J., & Flanagan, J. R. (2011). Principles of sensorimotor learning. Nature Reviews Neuroscience, 12(12), 739–751. https://doi.org/10.1038/nrn3112

[5] Ito, M. (2006). Cerebellar circuitry as a neuronal machine. Progress in Neurobiology, 78(3–5), 272–303. https://doi.org/10.1016/j.pneurobio.2006.02.006

[6] Sathyamurthy, A., Johnson, K. R., Matson, K. J. E., Dobrott, C. I., Li, L., Ryba, A. R., & Levine, A. J. (2020). Massively parallel single nucleus transcriptional profiling defines spinal cord neurons and their activity during behavior. Cell Reports, 31(1), 107595. https://doi.org/10.1016/j.celrep.2020.107595

programme

fri — 17

models and robots

9:00	Registration & welcome coffee
9:30	**Eva Dyer** Large-scale pretraining on neural data allows for transfer across individuals, tasks and species
10:30	**Jesse Marshall** Building generic neuromotor interfaces for human computer interaction
11:00	**Coffee break**
11:30	**Josue Ortega Caro** Spatiotemporal modeling of cortical cholinergic and calcium signaling dynamics across learning **Christoph Hemmer** DynaMix: zero-shot inference of dynamical systems
12:00	Panel discussion **Future of Cybernetics** **Adrienne Fairhall, PhD Eva Dyer, PhD Jesse Marshall, PhD Joe Paton, PhD, Sean Escola, MD PhD**
13:00	Lunch
14:30	**Srinivas C. Turaga** Simulating the brain and body of the fruit fly
15:00	**Auke Ijspeert** Investigating the neuromechanics of animal locomotion using robots and simulations
16:00	**Coffee break**
16:30	**Caroline Haimerl** World structure emerges within neural networks trained to generate egocentric, goal-directed behavior **Kabir Dabholkar** Finding separatrices of dynamical flows with deep koopman eigenfunctions
17:00	**Christopher J. Rozell** Measuring brain body interactions in humans during complex naturalistic decision making
17:30	**Closing remarks**
18:00	**Boat tour** **Farewell dinner**

fri — 17

keynote lectures

Large-scale pretraining on neural data allows for transfer across individuals, tasks and species

Eva Dyer

University of Pennsylvania
Philadelphia, Pennsylvania, USA

The brain is incredibly complex, with diverse functions that emerge from the coordinated activity of billions of neurons. These functions vary across brain regions and adapt dynamically as we engage in different tasks, process sensory information, or generate behavior. Yet, each neural recording captures only a small glimpse of this immense complexity, offering a limited view of the broader system. This motivates the need for an algorithmic approach to stitch together diverse datasets, integrating neural activity across brain regions, cell types, and individuals. In this talk, I will present our work on building scalable models pretrained on a broad corpus of neural recordings. Our findings demonstrate positive transfer across tasks, cell types, regions, and individuals, effectively bridging gaps between isolated studies. This unified framework opens new possibilities for brain-machine interfaces and cross-species neuroscience, and offers a path toward more generalizable models of brain function.

Investigating the neuromechanics of animal locomotion using robots and simulations

Auke J. Ijspeert

Ecole Polytechnique Fédérale de Lausanne (EPFL)
Lausanne, Switzerland

The ability to efficiently move in complex environments is a fundamental property both for animals and for robots, and the problem of locomotion and movement control is an area in which neuroscience, biomechanics, and robotics can fruitfully interact. In this talk, I will present how biorobots and numerical models can be used to explore the interplay of the four main components underlying animal locomotion, namely central pattern generators (CPGs), reflexes, descending modulation, and the musculoskeletal system. Going from lamprey to human locomotion, I will present a series of models that tend to show that the respective roles of these components might have changed during evolution with a dominant role of CPGs in lamprey and salamander locomotion, and a more important role for sensory feedback and descending modulation in human locomotion. I will also present how deep reinforcement learning can be used to explore questions related to supraspinal learning and planning that takes into account spinal cord dynamics.

Jesse Marshall

Srinivas C. Turaga

Christopher J. Rozell

talk sessions

Josue Ortega Caro

Christoph Hemmer

Caroline Haimerl

Kabir Dabholkar

Building generic neuromotor interfaces for human computer interaction

Jesse Marshall

Meta Platforms Inc
New York, New York, USA

Our team recently described our approach to building the first high-bandwidth neuromotor interface with performant out-of-the-box generalization across people (Kaifosh et al Nature 2025). This research work led to the development of the Meta Neural Wristband, which allows users to use natural hand gestures and handwriting to interact with augmented reality glasses. I will discuss the details of this research work, in particular the scaling laws that demonstrate how — similar to modern language models — there are predictable power-law relationships that describe the improvements in the performance of neuromotor gesture detection models with increasing number of participants in the training data corpus. I will also discuss the potential of these neuromotor interfaces to provide generalized interfaces for higher-dimensional computer inputs such as hand pose estimation (emg2pose, NeurIPS 2024) and typing (emg2qwerty, NeurIPS 2024) and how neuromotor interfaces can help create universal access to computing by enabling control in individuals with spinal cord injury and tremor.

Simulating the brain and body of the fruit fly

Srinivas C. Turaga

Janelia Research Campus, Howard Hughes Medical Institute
Ashburn, Virginia

Through recent advances in microscopy, we now have an unprecedented view of the brain and body of the fruit fly Drosophila melanogaster. We now know the connectivity at single neuron resolution across the whole brain. How do we translate these new measurements into a deeper understanding of how the brain processes sensory information and produces behavior? I will describe two computational efforts to model the brain and the body of the fruit fly. First, I will describe a new modeling method which makes highly accurate predictions of neural activity in the fly visual system as measured in the living brain, using only measurements of its connectivity from a dead brain[1], joint work with Jakob Macke. Second, I will describe a whole body physics simulation of the fruit fly which can accurately reproduce its locomotion behaviors, both flight and walking[2], joint work with Google DeepMind.

References

1 Lappalainen JK, Tschopp FD, Prakhya S, McGill M, Nern A, Shinomiya K, Takemura Sy, Gruntman E, Macke JH, and Turaga SC. Connectome-constrained networks predict neural activity across the fly visual system. Nature, 2024. https://www.nature.com/articles/s41586-024-07939-3

2 Vaxenburg R, Siwanowicz I, Merel J, Robie AA, Morrow C, Novati G, Stefanidi Z, Card GM, Reiser MB, Botvinick MM, Branson KM, Tassa† Y, and Turaga† SC. Whole-body simulation of realistic fruit fly loco- motion with deep reinforcement learning. Nature (in press), 2025. https://www.biorxiv.org/

Measuring brain body interactions in humans during complex naturalistic decision making

Christopher J. Rozell

Sankaraleengam Alagapan

Georgia Institute of Technology

The challenge of understanding how neural activity results in human behavior and cognition in health and disease is a crucial one for neuroscience. Traditional research methods often employ abstract tasks focusing on discrete cognitive processes, which may not fully capture the complexity of real-world behaviors or the underlying multiscale interactions between the brain and the body. For example, deficits in effort based decision-making (EBDM) are pervasive across various disorders such as major depression, but EBDM experiments are often limited in their use of simple effort proxies (not accounting for psychomotor response variability) and their use of single measurement modalities that do not capture the rich dynamics of natural behavior in the real world. We've recently used longitudinal local field potential recordings and explainable AI methods to capture changes in brain dynamics with deep brain stimulation for treatment resistant depression. While this shows that we can measure brain state changes that have behaviorally meaningful outcomes in the real world, the unconstrained nature of daily life leaves a gap in our ability to learn about the complex dynamics that underly specific behaviors. To address this, we are building a new naturalistic EBDM assay in an immersive virtual environment requiring effortful locomotion. The system will measure neural circuit activity synchronized with behavior across decision-making, embodied, affective, and clinical domains, lending itself to new latent variable characterizations of the brain-body interactions underlying EBDM. I will describe our previous work collecting intracranial electrophysiology in the real world, as well as our progress building a new experimental system (requiring integration of many disparate disciplines) that aims to unlock how real world behavior emerges from complex brain body interactions.

human electrophysiology / effort based decision making / multimodal data / computational psychiatry

Spatiotemporal modeling of cortical cholinergic and calcium signaling dynamics across learning

Josue Ortega Caro [1,2]

Sweyta Lohani [1]

Hannah Batchelor [1]

David van Dijk [2,3]

Jessica Cardin [1,2]

[1] Department of Neuroscience, Yale University
[2] Wu Tsai Institute, Yale University
[3] Department of Computer Science, Yale University

Neuromodulatory systems such as acetylcholine (Ach) play a central role in modulating cortical dynamics during behavior and learning. Recent studies have revealed that cholinergic signaling is spatiotemporally heterogeneous and exhibits task- and state-dependent coordination with neural activity [1]. However, how this coordination evolves during learning and how it shapes cortical computation remain unclear. Here, we combined dual-color mesoscopic imaging of Ach and Calcium (Ca) signals using genetically encoded sensors [2,3] with behavioral data in mice learning a visual contrast detection task. We observed large-scale plasticity in both Ach and Ca signals across learning, including increased responses in visual, somatosensory, and motor areas. To model these dynamics, we developed the multimodal transformer (MMT), a self-supervised deep learning architecture capable of combining the Ach, Ca, and behavioral dynamics into a single framework [4]. MMT significantly outperformed single-modality and multimodality baseline models in predicting trial-by-trial behavior and Ca activity across animals. To interpret the model, we analyzed the self-attention of the trained transformer model via attention rollout [5]. This approach revealed discriminative spatiotemporal patterns that differentiate naïve and expert animals, including enhanced attention to frontal cortical ACh signals in expert compared to naïve mice. Furthermore, we found that these changes are contrast-dependent, with Ach playing a critical role in near-threshold trials compared to high-contrast trials. Our results highlight the power of multimodal self-supervised modeling for uncovering interpretable patterns in complex cortical dynamics and neuromodulation during learning.

cortical dynamics / acetylcholine / self-supervised learning / widefield imaging / neural plasticity

References
[1] Lohani, S., Moberly, A. H., Benisty, H., et al. (2022). Spatiotemporally heterogeneous coordination of cholinergic and neocortical activity. Nature Neuroscience, 25(12), 1706–1713.
[2] Jing, M., Zhang, Y., Wang, H., et al. (2020). An optimized acetylcholine sensor for monitoring in vivo cholinergic activity. Nature Methods, 17(11), 1139–1146.
[3] Dana, H., Sun, Y., Mohar, B., et al. (2016). Sensitive red protein calcium indicators for imaging neural activity. eLife, 5, e12727.
[4] He, K., Chen, X., Xie, S., et al. (2022). Masked autoencoders are scalable vision learners. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 16000–16009.
[5] Abnar, S., & Zuidema, W. (2020). Quantifying attention flow in transformers. arXiv preprint arXiv:2005.00928.

DynaMix: Zero-Shot Inference of Dynamical Systems

Christoph Hemmer [1,3]

Daniel Durstewitz [1,2,3]

[1] Central Institute of Mental Health, Mannheim, DE
[2] Interdisciplinary Center for Scientific Computing, Heidelberg, DE
[3] Heidelberg University, Heidelberg, DE

Dynamical systems (DS) govern complex, time dependent phenomena across domains from climate patterns [1] to neural activity [2]. However, current approaches require specialized training for each system, failing to achieve zero-shot generalization that has made recent AI breakthroughs powerful. Current time series foundation models, including Amazon's Chronos [3], face critical limitations in capturing long-range temporal dependencies and fundamental mechanisms driving complex systems. These constraints prevent them from fully representing dynamics found in natural phenomena [4]. We introduce DynaMix, a foundation model bridging dynamical systems theory with modern scaling techniques to generalize across temporal domains [4]. DynaMix represents the first pre-trained foundation model capable of zero-shot forecasting on novel dynamical systems without system-specific training. Our approach leverages a mixture-of-experts architecture built upon Almost-Linear Recurrent Neural Networks [5]. From minimal contextual input, it extrapolates long-horizon system behavior and extracts invariant dynamical properties—capabilities setting it apart from existing foundation models. Remarkably, it achieves exceptional efficiency, using only 0.1% of parameters compared to billion-parameter foundation models while delivering orders of magnitude faster inference times. This efficiency stems from integrating control-theoretically motivated training algorithms [6] and DS principles. Despite being trained exclusively on simulated low-dimensional DS, DynaMix outperforms current time series foundation models in preserving long-term properties and often surpasses them in short-term prediction on e.g. real-world neural data like fMRI or EEG. This capability shows that models built on DS principles may bear a huge potential also for advancing the TS prediction field.

dynamical systems / time series / chaos / recurrent neural networks / foundation model

References

[1] Tziperman, E., Scher, H., Zebiak, S. E., & Cane, M. A. (1997). Controlling spatiotemporal chaos in a realistic El Nino prediction model. Physical Review Letters, 79(6), 1034.;
[2] Buzsáki, G. (2006). Rhythms of the Brain. Oxford university press.;
[3] Ansari, A. F., Stella, L., Turkmen, A. C., Zhang, X., Mercado, P., Shen, H., ... & Wang, B. Chronos: Learning the Language of Time Series. Transactions on Machine Learning Research.;
[4] Hemmer, C. J., & Durstewitz, D. (2025). True zero-shot inference of dynamical systems preserving long-term statistics. arXiv preprint arXiv:2505.13192.;
[5] Brenner, M., Hemmer, C. J., Monfared, Z., & Durstewitz, D. (2024). Almost-Linear RNNs Yield Highly Interpretable Symbolic Codes in Dynamical Systems Reconstruction. arXiv preprint arXiv:2410.14240.;
[6] Mikhaeil, J., Monfared, Z., & Durstewitz, D. (2022). On the difficulty of learning chaotic dynamics with RNNs. Advances in neural information processing systems, 35, 11297-11312.

World structure emerges within neural networks trained to generate egocentric, goal-directed behavior

Caroline Haimerl

Joseph Paton

Daniel McNamee

Champalimaud Foundation, Lisbon, PT

Behavior unfolds across multiple spatiotemporal scales. While fast, direct control is rooted in concrete, egocentric state-action mappings, slower, flexible behavior is suggested to rely on abstract representations such as "world models". How abstract internal reference frames emerge from interactions with the environment remains poorly understood. Here we hypothesize that they are a side product of needing to generate goal-directed behavior from high-dimensional sensory observations.

We train inverse networks to map current and goal sensory states to egocentric action sequences during spatial exploration. To support action decoding we introduce an inductive bias promoting linearized latent space (LLS): we apply a difference operator in the latent space such that actions are decoded from the embedding vector difference between encoded current and goal observations.

We find that models trained to predict egocentric action sequences develop latent spaces that reflect allocentric spatial structure - despite no explicit spatial supervision or predictive loss. Specifically, the latent representations support efficient spatial decoding, exhibit emergent properties such as direction invariance and multiscale periodicity (similar to grid cells), and straighten observation trajectories facilitating generalization. Interestingly, these structured representations are weaker or absent in models trained on allocentric action sequences, and depend crucially on the novel vector difference operator shaping the latent space geometry.

Thus our framework shows that allocentric spatial structure can emerge in deep networks trained to generate behavior in egocentric reference frames. Our findings offer an alternative computational explanation to predictive coding for how allocentric world representations may arise in the brain to support fast task learning and generalization.

action representation learning / world models / spatial navigation / neural network models

Finding separatrices of dynamical flows with deep koopman eigenfunctions

Kabir Dabholkar [1]

Omri Barak [2]

[1] Faculty of Mathematics, Technion - Israel Institute of Technology, Haifa, Israel 3200003
[2] Rappaport Faculty of Medicine and Network Biology Research Laboratory, Technion - Israel Institute of Technology, Haifa, Israel 3200003

Many natural systems, including neural circuits involved in memory [1] and decision making [2], are modeled as high-dimensional dynamical systems with multiple stable states. While existing analytical tools primarily describe behavior near stable equilibria [3], characterizing separatrices—the manifolds that delineate boundaries between different basins of attraction—remains challenging, particularly in high-dimensional settings. This challenge has grown with the increasing use of recurrent neural networks (RNNs) optimized to replicate behavioral or neural data. These models offer richer dynamics and are less structured than human-engineered models, and therefore require additional tools to reverse-engineer [3,4]. Here, we introduce a numerical framework leveraging Koopman Theory [5] combined with deep neural networks to characterize separatrices in such systems. Specifically, we approximate Koopman eigenfunctions (KEFs) associated with real positive eigenvalues, which vanish precisely at the separatrices. These scalar KEFs allow us to efficiently trace decision boundaries via gradient-based optimization, even in complex, black-box systems. We demonstrate the utility of our approach across synthetic benchmarks, ecological network models, and RNNs trained on neuroscience-inspired tasks. Finally, we show how the learned KEFs can be used to design minimal perturbations that push systems across separatrices, enabling targeted interventions—such as those relevant for optogenetic stimulation experiments in neuroscience.

recurrent neural networks / dynamical system

References

[1] Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. Proceedings of the national academy of sciences, 79(8), 2554-2558.
[2] Machens, C. K., Romo, R., & Brody, C. D. (2005). Flexible control of mutual inhibition: a neural model of two-interval discrimination. Science, 307(5712), 1121-1124.
[3] Sussillo, D., & Barak, O. (2013). Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks. Neural computation, 25(3), 626-649.
[4] Barak, O. (2017). Recurrent neural networks as versatile tools of neuroscience research. Current opinion in neurobiology, 46, 1-6.
[5] Mezić, I. (2015, December). On applications of the spectral theory of the Koopman operator in dynamical systems and control theory. In 2015 54th IEEE Conference on Decision and Control (CDC) (pp. 7034-7041). IEEE.

poster session 1

1 Shrivas Chaterji
2 Tatiana Silva
3 Youjing Yu
5 Ryan Lu
6 William F. Podlaski
7 Francesco Trapani
8 Sonica Saraf
9 Corinna Gebehart
10 Claire Rusch
11 Kcénia Bougrova
12 Guillermo Martín-Sánchez
13 Teresa Serradas Duarte, Simon A. Zamora
14 Marta Forcella
15 Simon Altrogge
16 Joana Carmona, Francesca Mastrogiuseppe
17 Lydia Fettweis Neto
18 Cheshta Bhatia
19 Virginia Palieri
20 A. Lucas Martins
21 Saheli Roy
22 João C. Marques
23 Nuno J. Machado
24 Ioana Lazar
25 Mattia Della Vecchia
26 Chenguang Li
27 Alice C. Navalho, Filipe S. Rodrigues
28 Afonso P. Dias, Filipe S. Rodrigues
29 Genji Kawakita
30 Joan Gort
31 Pranshu Malik
32 Inês Laranjeira
33 Renato Sousa
34 Mirjam Heinemans
35 Trinity Chung, Yuchen Shen
36 Juan R. Castiñeiras de Saa
37 Trung Le

The superior colliculus directs goal-oriented forelimb movements

Our ability to perform accurate goal-directed forelimb movements such as reaching and grasping is fundamental for interacting with and manipulating objects in our environments. Despite decades of research into the neural basis of forelimb movements, the underlying neural architecture, particularly the contributions of midbrain regions, remains only partially elucidated. Here we present evidence that the superior colliculus, an evolutionarily-conserved midbrain structure known to control saccadic and orientation movements, is also essential for executing accurate forelimb reaches in mice. Using intersectional circuit monitoring and manipulation techniques, we found that excitatory neurons in the mouse lateral superior colliculus are active during a skilled forelimb-reaching task, and phase-specific silencing of these neurons during reaching movements differentially impairs reach accuracy. Anatomical studies dentified neurons in the deep cerebellar nuclei and the pars reticulata region of the substantia nigra as sources of input to the lateral superior colliculus, and manipulation of these pre-synaptic partners revealed a role for the nigrotectal but not cerebellotectal neurons in facilitating reach accuracy in skilled mice. Notably, inhibition of the nigrotectal pathway resulted in deviation of the paw that was diametrically opposite to the kinematic deviation observed upon collicular silencing. In summary, our findings establish the superior colliculus as a pivotal regulator of skilled forelimb use, highlighting the coordinated action of the substantia nigra and the superior colliculus in controlling reach accuracy and enriching our understanding of how the brain orchestrates complex forelimb movements.

superior colliculus / skilled forelimb control / neural circuit monitoring and manipulation / droplet-retrieval task

Shrivas Chaterji [1,2]

Punarva H. Belliappa [1,3]

Anupama Sathya-murthy [1]

[1] Centre for Neuroscience, Indian Institute of Science, Bengaluru, Karnataka, India
[2] Present address: Champalimaud Centre for the Unknown, Lisbon 1400-038, Portugal - Israel Institute of Technology, Haifa, Israel
[3] Present address: Brown University, USA

Brief Critical Windows for Memory Consolidation are Regulated by Behavioral State

N. Tatiana Silva [1]

M Inês Ribeiro [1,2]

Megan Carey [1]

[1] Champalimaud Research - Neuroscience, Lisbon, PT
[2] OIST, Okinawa, JP

Long-term memories are consolidated over time, progressively becoming more stable and resistant to interference. Memory consolidation occurs offline and often involves transfer of memories from one brain site to another. For many motor memories, consolidation is thought to involve early learning in the cerebellar cortex that is subsequently transferred to the cerebellar nuclei. Here we report that in mice, engaging in locomotor activity during training in a classical conditioning task shifts the critical time window for memory consolidation, from just after training sessions, to between trials, within sessions. This temporal shift requires natural patterns of cerebellar granule cell activity during intertrial intervals and is accompanied by earlier involvement of the downstream cerebellar nucleus. These results reveal that the critical time window for cerebellar memory consolidation can be surprisingly brief, on a timescale from seconds to minutes, and that it is dynamically regulated by behavioral state. In ongoing experiments we are imaging granule cell activity to investigate how their activity patterns change during different memory consolidation windows and behavioral states.

associative learning / memory consolidation / behavioral state / locomotor activity

Second-order forward-mode optimization of recurrent neural networks for neuroscience

Youjing Yu

Rui Xia

Qingxi Ma

Máté Lengyel

Guillaume Hennequin

University of Cambridge, Cambridge, UK

Training recurrent neural networks (RNNs) to perform neuroscience tasks can be challenging. Unlike in machine learning where any architectural modification of an RNN (e.g. GRU or LSTM) is acceptable if it facilitates training, the RNN models trained as models of brain dynamics are subject to plausibility constraints that fundamentally exclude the usual machine learning hacks. The "vanilla" RNNs commonly used in computational neuroscience find themselves plagued by ill-conditioned loss surfaces that complicate training and significantly hinder our capacity to investigate the brain dynamics underlying complex tasks. Moreover, some tasks may require very long time horizons which backpropagation cannot handle given typical GPU memory limits. In earlier work [1], we developed SOFO, a second-order random subspace optimizer that efficiently explores loss surfaces without requiring backpropagation. Instead, SOFO relies on parallelizable batched forward-mode differentiation, yielding constant memory cost over time. By exploiting second-order curvature, SOFO significantly outperforms Adam on various RNN tasks. In more recent work, we improve SOFO by using curvature-guided navigation of the parameter space, replacing random subspace sampling in each iteration. We show that the improved SOFO converges faster than 1 across tasks including a delayed addition task [2], a challenging double-reach motor task, and a meta continual-learning task using a Hebbian learning rule. By accelerating and scaling the training of biologically grounded network models, SOFO greatly facilitates research into neural networks modeling behavioral tasks.

biologically plausible recurrent neural networks / behaviour modelling / optimization / memory-efficient training / second-order method

References

[1] Yu, Y., Xia, R., Ma, Q., Lengyel, M., & Hennequin, G. (2024). Second-order forward-mode optimization of recurrent neural networks for neuroscience. Advances in Neural Information Processing Systems, 37, 92101-92129;

[2] Schmidt, D., Koppe, G., Monfared, Z., Beutelspacher, M., & Durstewitz, D. (2019). Identifying nonlinear dynamical systems with multiple time scales and long-range dependencies. arXiv preprint arXiv:1910.03471.

Towards a neuromechanical model of swimming in the jellyfish Clytia hemisphaerica

Long-term memories are consolidated over time, progressively becoming more stable and resistant to interference. Memory consolidation occurs offline and often involves transfer of memories from one brain site to another. For many motor memories, consolidation is thought to involve early learning in the cerebellar cortex that is subsequently transferred to the cerebellar nuclei. Here we report that in mice, engaging in locomotor activity during training in a classical conditioning task shifts the critical time window for memory consolidation, from just after training sessions, to between trials, within sessions. This temporal shift requires natural patterns of cerebellar granule cell activity during intertrial intervals and is accompanied by earlier involvement of the downstream cerebellar nucleus. These results reveal that the critical time window for cerebellar memory consolidation can be surprisingly brief, on a timescale from seconds to minutes, and that it is dynamically regulated by behavioral state. In ongoing experiments we are imaging granule cell activity to investigate how their activity patterns change during different memory consolidation windows and behavioral states.

associative learning / memory consolidation / behavioral state / locomotor activity

Ryan Lu [1]

Karen Cunningham [2]

James Fitzgerald [1]

Brady Weissbourd [2]

Ann Kennedy [3]

[1] Northwestern University, Chicago, US
[2] Massachusetts Institute of Technology, Cambridge, US
[3] Scripps Research, San Diego, US

Exploring the inner workings and internal representations of predictive coding networks in comparison to usual feedforward neural networks

William F. Podlaski

Christian K. Machens

Champalimaud Foundation, Lisbon, PT

Deep feedforward and recurrent networks have become popular choices for models of neural computation, but they do not provide a clear link to the details of the underlying biophysical processes. Here we argue that the reason this link has been difficult to establish is because biological circuits operate in a fundamentally different way. To illustrate this, we present a new theory of neural computation in spiking networks, which captures the precise computational role of each neuron and every spike. By assuming low-rank recurrent connectivity, we show that spiking population activity is confined to a well-defined nonlinear manifold in a low-dimensional latent space, with the spikes of individual neurons pushing the latent dynamics along this manifold. We then show that the network's recurrent connectivity can be factorized into a part that determines the manifold geometry and another part that determines the manifold dynamics. The stability of the on-manifold dynamics can be enforced through sign constraints — either through an all-inhibitory network with a constant background input, or through an (inhibition-stabilized) excitatory inhibitory network — thereby suggesting a functional role for Dale's law. We show that such networks can approximate arbitrary continuous dynamical systems, and demonstrate several examples including a limit cycle, ring attractor, and a set of Hopfield-like fixed-point attractors. Overall, our work proposes a new way of understanding the dynamics and computations of spiking neural networks, and suggests the possibility of an intriguing distinction between biological and artificial computation.

theoretical neuroscience / spiking neural networks / computation through dynamics / low-dimensional manifolds

A normative dynamical theory of intuitive human reasoning

Francesco Trapani

Daniel McNamee

Champalimaud Foundation, Lisbon, PT

Humans routinely solve complex, high-dimensional problems with efficiency, often relying on intuitive reasoning that yields near-optimal solutions with minimal deliberation. This suggests an inference mechanism that is fast, directed, yet capable of implicitly evaluating long-range, multi-step action sequences. Early cognitive theories, notably those of H. Simon [1], framed human reasoning as heuristic-guided search over symbolic representations. These accounts emphasized meta-cognitive strategies for navigating large problem spaces. While foundational, such models were limited to static symbolic operations and overlooked the dynamical nature of cognitive processes. We introduce a normative framework in which key meta-cognitive principles emerge from a sub-symbolic, dynamical inference process, formalized within fully specified Markov Decision Processes. Our model generalizes Simon's means–ends analysis as the core unifying perspective by replacing explicit forward models with an implicit representation of global problem structure, guiding inference through gradient-based policy optimization that converges to globally optimal solutions. This architecture gives rise to hallmark traits of human reasoning as emergent properties of optimization dynamics, including hierarchical abstraction, forward-backward search, and least commitment—e.g., selectively refining policies at sub-goals while avoiding premature full-plan specification. The inference process is inherently parallel and supports flexible adaptation through second-order mechanisms such as situational focusing or preconditioning. We show that model predictions align with empirical observations in tasks involving spatial navigation, physical reasoning, and puzzle-solving, capturing both behavioral outcomes and the process-level dynamics underlying intuitive human reasoning through a single unifying normative principle.

intuitive reasoning / dynamical systems

References

[1] Newell, A., & Simon, H. A. (1972). Human problem solving (Vol. 104, No. 9). Englewood Cliffs, NJ: Prentice-hall.

Variations in neuronal selectivity create efficient representational geometries for perception

Sonica Saraf [1]

J. Anthony Movshon [1]

SueYeon Chung [2]

[1] New York University, New York, USA

[2] Flatiron Institute, New York, USA

Our visual capabilities depend on neural response properties in visual areas of our brains. Populations of neurons exhibit a wide variety of selective response properties, but the reasons for this diversity are unknown. Here, we related the distribution of neuronal tuning properties to the information capacity of the population. In particular, we study how tuning heterogeneity affects geometry and coding efficiency in populations of neurons. Our results from theory, simulations, and analysis of recordings from macaque primary visual cortex (V1) reveal that diversity of amplitude and bandwidth drive complementary changes to the representational geometry of a population. Amplitude diversity pushes the centers of the representations further apart, whereas bandwidth heterogeneity decorrelates the center locations. These geometric changes separate out representations for distinct stimuli, creating more efficient encoding. We study how both types of diversity affect the population code in the context of two different perceptual tasks: discrimination and identification. While both types of diversity improve encoding for both tasks, their distinct impacts on geometry make each more beneficial for one of the two tasks. Amplitude diversity impacts coding efficiency more for discrimination than it does for identification, while bandwidth diversity has a stronger impact on identification. These complementary effects indicate the importance of both types of diversity for perception. Because tuning diversity exists across species and brain areas, our results suggest a fundamental population coding strategy that may be applicable to a wide range of behavior. Finally, this work establishes connections between multiple levels of neural processing - from the distributions of single neuron properties, to population-level representations, and to perception.

vision / psychophysical discrimination / psychophysical identification / neural geometry / population coding

Angular Self-Motion Estimation for Action Selection in Walking Drosophila

Corinna Gebehart [1]

Tomás Cruz [2]

M. Eugenia Chiappe [1]

[1] Champalimaud Foundation, Lisbon, PT
[2] Friedrich Miescher Institute for Biomedical Research, Basel, CH

Circuits involved in self-motion estimation provide a unique opportunity to understand how various aspects of walking control - from rapid course adjustments to slower navigational decisions - are implemented by distributed yet interconnected networks. Recently, we found that in darkness, exploratory walking flies (Drosophila melanogaster) select the direction of their turns based on the drift direction of a preceding forward run, a phenomenon we termed drift-to-saccade (DTS). DTS shows that flies use angular estimation to plan a future action. Previous work in the lab identified a visuomotor network in flies that monitors rotations during walking. One branch of this network projects to the ventral nerve cord (VNC) to mediate real-time steering adjustments. The other branch projects to central regions intimately connected with the central complex, the insect's navigational center. In addition to processed visual inputs, this central network receives ascending inputs from the VNC and feedback from the central complex. This suggests it serves as a key integrative node for angular state estimation. Indeed, individually silencing identified neurons within this network disrupts DTS without impairing forward runs or rapid body turns.

Using whole-cell patch clamp recordings in walking flies, we are currently characterizing a central population of neurons within this network. We found that these neurons are sensitive to both instantaneous angular velocity and cumulative orientation within a walking bout. Moreover, their activity resets following rapid, direction selective turns, suggesting the network monitors left-right

asymmetries during forward runs that are reset after rapid changes in course direction.

Together, our findings reveal a multilayer network likely responsible for encoding angular state estimation by integrating diverse sensory and motor-related signals from visual, VNC, and central complex inputs. We are currently analysing how different network components shape these signals and contribute to turning decisions. Our goal is to provide a detailed functional map of a circuit involved in state estimation for continuous movement control and discrete action selection.

sensorimotor integration / electrophysiology / self-motion estimation / action selection / movement control

References

[1] M. E. Chiappe, 'Circuits for self-motion estimation and walking control in Drosophila', Curr. Opin. Neurobiol., vol. 81, p. 102748, Aug. 2023, doi: 10.1016/j.conb.2023.102748.
[2] A. J. Kim, L. M. Fenk, C. Lyu, and G. Maimon, 'Quantitative Predictions Orchestrate Visual Signaling in Drosophila', Cell, vol. 168, no. 1–2, pp. 280-294.e12, Jan. 2017, doi: 10.1016/j.cell.2016.12.005.
[3] T. Fujiwara, T. L. Cruz, J. P. Bohnslav, and M. E. Chiappe, 'A faithful internal representation of walking movements in the Drosophila visual system', Nat. Neurosci., vol. 20, no. 1, pp. 72–81, Jan. 2017, doi: 10.1038/nn.4435.
[4] C. Busch, A. Borst, and A. S. Mauss, 'Bi-directional Control of Walking Behavior by Horizontal Optic Flow Sensors', Curr. Biol., vol. 28, no. 24, pp. 4037-4045.e5, Dec. 2018, doi: 10.1016/j.cub.2018.11.010.
[5] T. Fujiwara, M. Brotas, and M. E. Chiappe, 'Walking strides direct rapid and flexible recruitment of visual circuits for course control in Drosophila', Neuron, vol. 110, no. 13, pp. 2124-2138.e8, Jul. 2022, doi: 10.1016/j.neuron.2022.04.008.
[6] B. D. DeAngelis, J. A. Zavatone-Veth, and D. A. Clark, 'The manifold structure of limb coordination in walking Drosophila', eLife, vol. 8, p. e46409, Jun. 2019, doi: 10.7554/eLife.46409.
[7] T. L. Cruz, S. M. Pérez, and M. E. Chiappe, 'Fast tuning of posture control by visual feedback underlies gaze stabilization in walking Drosophila', Curr. Biol., vol. 31, no. 20, pp. 4596-4607.e5, Oct. 2021, doi: 10.1016/j.cub.2021.08.041.

Sensorimotor calibration in optic flow processing circuits

Claire Rusch

Eugenia Chiappe

Champalimaud Foundation, Lisbon, PT

Exploratory animals structure their behavior to maximize gaze stability, thereby facilitating the acquisition of visual and spatial information while minimizing retinal slip. Such gaze control depends on multisensory integration, but the circuit mechanisms underlying precise multimodal calibration during locomotion remain unclear. Partially, this is because of the highly distributed nature of sensorimotor circuits, making it challenging to identify circuits involved in goal-directed gaze control. Here we take advantage of Drosophila melanogaster's compact Central Nervous System to examine multimodal calibration during exploratory walking. We tested whether flies maintain gaze stability under visual perturbations by immersing them in a virtual world and subjected them to constant rotations of this world. Flies adjusted their velocity to preserve gaze stability, underscoring visuomotor recalibration. To begin to understand the neural underpinnings of such recalibration, we adapted this parading to head-fixed flies walking on a spherical treadmill to record neural activity simultaneously from a population of genetically identified neurons involved in gaze control. The GABAergic bIPS cells receive multimodal information from integrative brain regions and the VNC (the insect analogue of the spinal cord), providing an anatomical substrate for calibration. Recordings from bIPS in walking flies showed that they congruently combine retinal and extra-retinal signals. Moreover, this congruent multimodal combination sharpens the neuron's sensitivity to the body's translation and rotation. Ongoing experiments are testing the activity of bIPS under visual perturbations to examine recalibration at the

level of bIPS activity. Together, our data underscores the properties of an integrative inhibitory hub involved in steering during locomotion. Future work leveraging the EM connectomics datasets will test the mechanisms by which bIPS combine and calibrate multimodal information for gaze control in the context of exploration.

sensorimotor integration / sensorimotor adaptation / visual processing / calcium imaging

Comparing the activity of four neuromodulators during decision-making

Kcénia Bougrova [1]

Matteo Carandini [2]

Davide Crombie [1]

Eric deWitt [1]

Carolina Quadrado [2]

The four principal neuromodulators (NMs) - dopamine (DA), serotonin (5-HT), norepinephrine (NE), and acetylcholine (ACh) - are implicated in decision-making processes such as reward prediction errors, learning, and uncertainty, yet the unique contributions and mechanisms by which they interactively shape decision-making remain unclear since most studies target individual NMs in varying paradigms, preventing direct comparisons.
To address this, we leveraged the International Brain Laboratory (IBL) framework, using a highly standardized, reproducible decision-making task. We employed fiber photometry to record bulk calcium activity of genetically targeted NM populations expressing GCaMP6f across ~750,000 trials in ~40 double-transgenic mice during learning and proficient stages of a visually guided decision-making task.

Consistent with prior research, DA activity increased after rewards and decreased after errors, with graded amplitude changes depending on the trial difficulty, reflecting reward prediction errors. In contrast, 5-HT exhibited strong responses to incorrect choices - suggesting an unsigned prediction error signal, and also showed suppression before stimulus onset, potentially linked to maintaining task focus.
NE signals were notably variable across and within mice, pointing to possible specialized roles in different subpopulations. ACh activity was higher in early training stages, especially during correct trials, indicating a role in learning processes.

Together, these findings highlight distinct and complementary roles for each NM in encoding

[1] Champalimaud Foundation, Lisbon, PT
[2] University College London, London, UK
[3] University of Geneva, Geneva, CH

reward-related contingencies that may guide future behavior. Integrating fiber photometry data with IBL's behavioral and electrophysiological datasets offers a powerful platform to dissect how these neuromodulatory systems interact during learning and decision-making.

neuromodulators / decision making / behavior / dopamine / serotonin

Georg Raiser [1]

Laura Silva [1]

Olivier Winter [1]

Florian Rau [1]

Gaelle Chapius [3]

Zachary Mainen [1]

[1] Champalimaud Foundation, Lisbon, PT
[2] University College London, London, UK
[3] University of Geneva, Geneva, CH

Solving Constrained Minimax Problems with Spiking Neural Networks

Spiking neural networks (SNNs) serve as key models of brain function in neuroscience [1] and as energy efficient algorithms in engineering [2]. However, they remain difficult to build and interpret. In this work, we establish a theoretical connection between SNNs and minimax optimization — a broad class of optimization problems with links to decision making under uncertainty, zero-sum games, and optimal control. Building on existing results that link SNNs to convex optimization [4,5,6,7], we show that the dynamics of certain low-rank SNNs can solve minimax problems with quadratic objectives and linear constraints. We provide geometrical intuitions for how the minimax problem maps onto a network's latent space, and how the latent dynamics of the spiking networks reaches the optimal solution. With this work, we hope to make a step forward in interpretability and usability of spike-based computation — not only serving as a viable framework to understand biological networks, but also opening avenues for neuromorphic implementations of energy-efficient optimization solvers.

Spiking neural networks / constrained optimization / minimax optimization / population geometry / low-dimensional dynamics

Guillermo Martín-Sánchez

William F. Podlaski

Christian K. Machens

Champalimaud Foundation, Lisbon, PT

References

[1] Abbott, L., DePasquale, B., & Memmesheimer, R. M. (2016). Building functional networks of spiking model neurons. Nature Neuroscience, 19, 350–355. https://doi.org/10.1038/nn.4241;
[2] Mead, C. (1990). Neuromorphic electronic systems. Proceedings of the IEEE , 78 (10), 1629-1636. doi: 10.1109/5.58356 ;
[3] Boyd, S., & Vandenberghe, L. (2004). Convex optimization. Cambridge university press.;
[4] Tang, P., Lin, T.-H., & Davies, M. (2017, 05). Sparse coding by spiking neural networks: Convergence theory and computational results. doi: 10.48550/arXiv.1705.05475;
[5] Chou, C.-N., Chung, K.-M., & Lu, C.-J. (2018, 03). On the algorithmic power of spiking neural networks. doi: 10.48550/arXiv.1803.10375;

Contributions of Basal Ganglia and Cerebellum to Discrete and Continuous Behavioral Control

Teresa Serradas Duarte

Simon A. Zamora

Coralie Hérent

Megan R. Carey

Joseph J. Paton

Champalimaud Foundation, Lisbon, PT

Adaptive behavior requires selecting appropriate actions and executing them effectively, a process that depends on coordinated activity across distributed brain systems, including the basal ganglia (BG) and cerebellum (CB). Lesions to either structure produce distinct impairments, suggesting they function as dissociable modules within a hierarchically organized yet parallel network [1,2,3,4]. To assess their respective roles in motor control, we examined the effects of BG and CB ablations in mice performing multiple variants of a water reaching task. We found a striking double dissociation: CB-ablated mice succeeded in the head-fixed task but failed in the freely moving version, which requires full-body coordination; BG-ablated mice showed the opposite pattern, failing in the head-fixed task where demands on action selection are greater.

Additionally, CB-ablated mice could recruit piecewise primitives for reaching, but their movements were less smooth and reach trajectories were more stereotyped, within and across animals. In contrast, BG-ablated mice preserved movement execution but exhibited biases towards richer sensory-feedback movements, abnormal action transitions, and failure to suppress competing actions. To build on these results, we developed the select-and-collect task, a paradigm designed to temporally parse demands on action selection and execution.

Simultaneous recordings from BG and CB revealed distinct dynamics: BG activity spanned the trial and peaked at reward collection, while CB activity was locked to movement and modulated by ongoing kinematics. Notably, CB activity consistently led index

reach kinematics, whereas BG activity lagged. These findings support a functional dissociation, where BG supports discrete control–intermittently selecting actions and modulating vigor; while CB supports continuous control–refining movement in high-dimensional spaces.

motor control / basal ganglia / cerebellum / behavior /

References

[1] Middleton, F. A., & Strick, P. L. (2000). Basal ganglia and cerebellar loops: Motor and cognitive circuits. Brain Research Reviews, 31(2–3), 236–250. doi: 10.1016/S0165-0173(99)00040-5
[2] Shadmehr, R., & Krakauer, J. W. (2008). A computational neuroanatomy for motor control. Experimental Brain Research, 185(3), 359–381. doi: 10.1007/s00221-008-1280-5
[3] Doya, K. (2000). Complementary roles of basal ganglia and cerebellum in learning and motor control. Current Opinion in Neurobiology, 10(6), 732–739. doi: 10.1016/S0959-4388(00)00153-7
[4] Haimerl, C., Rodrigues, F. S., & Paton, J. J. (2025). Time, Control, and the Nervous System. Annual Review of Neuroscience. doi: 10.1146/

Modulation of Granule Cell Activity by Locomotion During Associative Learning

Marta Forcella Sánchez

Merit Kruse

Romain Sala

Tatiana Silva

Megan R. Carey

Champalimaud Foundation, Lisbon, PT

Associative learning–the process by which organisms form predictive associations between events–enables behavioral adaptation to dynamic conditions [1]. This requires the representation of events (e.g. stimuli) and mapping them to appropriate responses. Such associations are often formed while animals are engaged in active behaviors like walking or running, which may influence the neural representations relevant for learning.

Previous work from our lab has shown that locomotion facilitates the acquisition, consolidation and expression of delay eyeblink conditioning, a form of cerebellum-dependent associative learning [2,3]. Specifically, increased locomotor speed accelerates the rate and amplitude of learning in mice. Locomotor state and task-related sensory inputs are conveyed via mossy fibers (MFs) to granule cells (GCs), the input layer of the cerebellum. Notably, optogenetic stimulation at the MF-GC synapses mimicking the enhanced synaptic transmission during locomotion [2,4] has been shown to amplify the expression of conditioned responses [2]. In this project, we are investigating the mechanism by which the locomotor state influences GC activity, thus modulating learning.

As an initial step, we collected two-photon calcium imaging data from cerebellar GCs before and after eyeblink conditioning, while the animals ran either on a motorized or self-paced treadmill. Early analyses suggest that stimulus processing at the GC population level differs depending on locomotor signals. These preliminary findings point to a potential modulatory mechanism by which behavioral state influences associative learning and lay the groundwork for more detailed future investigations.

cerebellar associative learning / granule cells / locomotion / two-photon calcium imaging

References

[1] A. Mason, E. A. Ludvig, and C. R. Madan, “12 Conditioning and Associative Learning”.
[2] C. Albergaria, N. T. Silva, D. L. Pritchett, and M. R. Carey, “Locomotor activity modulates associative learning in mouse cerebellum,” Nat. Neurosci., vol. 21, no. 5, pp. 725–735, May 2018, doi: 10.1038/s41593-018-0129-x.
[3] N. T. Silva, M. I. Ribeiro, and M. R. Carey, “Behavioral state regulates the dynamics of memory consolidation,” Jul. 01, 2024, Cold Spring Harbor Laboratory. doi: 10.1101/2024.06.28.601231.
[4] K. Powell, A. Mathy, I. Duguid, and M. Häusser, “Synaptic representation of locomotion in single cerebellar granule cells,” eLife, vol. 4, p. e07290, Jun. 2015, doi: 10.7554/eLife.07290.

A Fully 3D Model of Jellyfish Swimming, Integrating Nerve Nets, Muscles, Body Biomechanics and Fluid-Structure Interaction

Jellyfish are among the earliest freely swimming animals and their behaviors—including forward propulsion and evasive as well as approaching turning maneuvers—are generated by a comparatively simple, decentralized nervous system composed of neurons homologous to those of bilaterians [1]. It remains unclear, however, how the spatial structure and neural activity of jellyfish nerve nets give rise to both symmetric and asymmetric, tightly coordinated muscle contractions and, as a result, highly energy-efficient locomotion in water [2]. To address this, we develop a computational model of the scyphozoan jellyfish Aurelia aurita that integrates nerve nets, muscles, body biomechanics and fluid-structure interaction in—for the first time—three dimensions, to bridge the gap between neural activity and behavior. To reduce computational cost, we introduce an abstract model of nerve net activity as signal transmission on a graph whose nodes correspond to synapses. This model captures the spatiotemporal propagation of neural signals and closely approximates the dynamics of a nerve net model of interconnected Hodgkin-Huxley-type neurons [3]. To simulate the interaction between the deformable jellyfish body and the surrounding fluid, we employ smoothed particle hydrodynamics, a particle-based method that models both the elastic body and the water in a unified framework and naturally handles large deformations. Our model spans multiple domains, from nerve net architecture and dynamics to muscle activation and fluid-body interaction, providing a platform to study how behavior emerges from the interplay of nervous system, muscles, body and environment. In addition, it offers insights into the function of early nervous systems and the origins of autonomous movement.

computational neuroscience / jellyfish / locomotion / embodied behavio / fluid-structure interaction

Simon Altrogge

Raoul-Martin Memme-sheimer

University of Bonn, Bonn, DE

References

[1] Cunningham, K., Anderson, D. J., & Weissbourd, B. (2024). Jellyfish for the study of nervous system evolution and function. Current Opinion in Neurobiology, 88, 102903. https://doi.org/10.1016/j.conb.2024.102903

[2] Gemmell, B. J., Costello, J. H., Colin, S. P., Stewart, C. J., Dabiri, J. O., Tafti, D., & Priya, S. (2013). Passive energy recapture in jellyfish contributes to propulsive advantage over other metazoans. Proceedings of the National Academy of Sciences, 110(44), 17904–17909. https://doi.org/10.1073/pnas.1306983110

[3] Pallasdies, F., Goedeke, S., Braun, W., & Memmesheimer, R.-M. (2019). From single neurons to behavior in the jellyfish Aurelia aurita. eLife, 8, e50084. https://doi.org/10.7554/eLife.50084

Input-dependent Directionality of Interactions Between Cortical Areas

A long-standing question in systems neuroscience is how different brain areas communicate. One approach to studying information flow is to analyze the covariation of activity across areas, with the temporal structure of this covariation offering clues about the directionality of signaling [1–3]. Recent studies applying this method to neural populations across areas have revealed that the directionality of interactions can shift rapidly, depending on stimuli and task demands [4,5]. These findings suggest that experimentally measured covariation metrics reflect not only the underlying synaptic connectivity but also dynamic influences such as inputs from other regions. Understanding how these metrics relate to circuit-level mechanisms remains a challenging task that warrants theoretical approaches [6,7].

Here, we develop a theoretical framework for the emergence of directional interactions in recurrent circuits, based on recurrent neural network models driven by stochastic inputs [8]. Our theory leverages the analysis of network activity along particular directions, given by the eigenvectors of the synaptic connectivity matrix. We apply our framework to study how external inputs can flexibly shape the directionality of inter-area communication in mesoscopic cortical circuits. Our analysis reveals that inputs targeting excitatory (E) and inhibitory (I) populations play different roles in inter-areal interactions. Specifically, inputs to E significantly influence directionality by inducing a systematic bias from the area receiving the strongest input, while inputs to I regulate the amplitude and timescale of activity with minimal impact on directionality. In circuits with feature-specific connectivity and

Joana Carmona [1]

Francesca Mastrogiu-seppe [1]

Byron M. Yu [2]

Adam Kohn [3]

Christian K. Machens [1]

[1] Champalimaud Foundation, Lisbon, PT
[2] Carnegie Mellon University, Pittsburgh, US
[3] Albert Einstein College of Medicine, New York, US * these authors contributed equally

inputs, the effect of inputs to E on directionality is most apparent at the level of latent variables reflecting co-fluctuations between stimulus-selective and unselective units. Overall, our work provides a theoretical foundation for the interpretation of experimentally measured covariation metrics, and advances our understanding on how to link the functional and anatomical substrates of neural interactions.

cross-covariances / directionality / inter-areal interactions / recurrent neural networks

references

[1] Bair, W., Zohary, E., & Newsome, W. T. (2001). Correlated firing in macaque visual area MT: time scales and relationship to behavior. The Journal of neuroscience : the official journal of the Society for Neuroscience, 21(5), 1676–1697. https://doi.org/10.1523/JNEUROSCI.21-05-01676.2001;
[2] Gochin, P. M., Miller, E. K., Gross, C. G., & Gerstein, G. L. (1991). Functional interactions among neurons in inferior temporal cortex of the awake macaque. Experimental brain research, 84(3), 505–516. https://doi.org/10.1007/BF00230962;
[3] Roe, A. W., & Ts'o, D. Y. (1999). Specificity of color connectivity between primate V1 and V2. Journal of neurophysiology, 82(5), 2719–2730. https://doi.org/10.1152/jn.1999.82.5.2719;
[4] van Kempen, J., Gieselmann, M. A., Boyd, M., Steinmetz, N. A., Moore, T., Engel, T. A., & Thiele, A. (2021). Top-down coordination of local cortical state during selective attention. Neuron, 109(5), 894–904.e8. https://doi.org/10.1016/j.neuron.2020.12.013;
[5] Semedo, J. D., Jasper, A. I., Zandvakili, A., Krishna, A., Aschner, A., Machens, C. K., Kohn, A., & Yu, B. M. (2022). Feedforward and feedback interactions between visual cortical areas use different population activity patterns. Nature communications, 13(1), 1099. https://doi.org/10.1038/s41467-022-28552-w;
[6] Ostojic, S., Brunel, N., & Hakim, V. (2009). How connectivity, background activity, and synaptic properties shape the cross-correlation between spike trains. The Journal of neuroscience : the official journal of the Society for Neuroscience, 29(33), 10234–10253. https://doi.org/10.1523/JNEUROSCI.1275-09.2009;
[7] Trousdale, J., Hu, Y., Shea-Brown, E., & Josić, K. (2012). Impact of network structure and cellular response on spike time correlations. PLoS computational biology, 8(3), e1002408. https://doi.org/10.1371/journal.pcbi.1002408;
[8] Godrèche, C. & Luck, J.-M. (2018). Characterising the nonequilibrium stationary states of Ornstein–Uhlenbeck processes. Journal of Physics A: Mathematical and Theoretical, 52(3), 035002. https://dx.doi.org/10.1088/1751-8121/aaf190.

Hallucination of visual instability is promoted by psychedelics and suppressed by gaze shifts

Lydia Fettweis Neto

Scott Renie

Zachary Mainen

Champalimaud Foundation, Lisbon, PT

Substances such as LSD and psilocybin, known as classical psychedelics, cause profound alterations in consciousness, including various alterations in visual experience [1]. Psychophysical analysis of these visual effects could help clarify how psychedelics alter brain function. A key aspect of vision is the interplay between foveal and peripheral processing, orchestrated by saccadic movements and fixations through which visual scenes are constructed [2,3].

In this study, we investigated a visual effect anecdotally reported under psychedelic: the perception of stationary surfaces as unstable, dynamically warping, or drifting [4]. Participants viewed projected still images of natural textures (e.g., grass, gravel) and reported perceived motion. Gaze trajectory was recorded using a wearable eye-tracker [5,6]. Importantly, on half the trials, a fixation cross instructed subjects to hold their gaze steady. Subjects (n = 45) were recruited and tested in a natural environment in which psychedelic use is common, under a protocol approved by our institutional ethics committee. Texture instability was reported markedly more often in the fixation condition, and this effect was stronger in subjects reporting recent use of a classical psychedelic.

The dependence of visual instability on gaze behaviour lends support to enactive sensorimotor theories of visual perception that hypothesize "hallucinations" arise when active exploration of a scene is reduced [7]. Notably, reports of movement without psychedelics [8,9] suggest that psychedelic visual effects do not necessarily require abnormal cortical processing but may arise from attending to normally unattended features of visual processing in embodied sensorimotor loops. These results lend support to the use of visual psychophysics to reveal general principles of psychedelic effects on perception in relation to normal brain function.

simple visual hallucinations / gaze behavior / eye-tracking / psychedelic / attention

References

Kwan AC, Olson DE, Preller KH, Roth BL. The neural basis of psychedelic action. Nat Neurosci. 2022 Nov;25(11):1407–19.
Skaramagkas V, Giannakakis G, Ktistakis E, Manousos D, Karatzanis I, Tachos N, et al. Review of eye tracking metrics involved in emotional and cognitive processes. IEEE Rev Biomed Eng. 2023 Jan 5;16:260–77.
Groner R, Groner MT. Attention and eye movement control: An overview. Eur Arch Psychiatry Neurol Sci. 1989 Jan;239(1):9–16.
Kohn B, Bryden MP. The effect of lysergic acid diethylamide (LSD-25) on perception with stabilized images. Psychopharmacologia. 1965 Apr 7;7(5):311–20.
Baumann C, Dierkes K. Neon Accuracy Test Report. Pupil Labs. 2023.
Pfeffer T, Dierkes K. Neon Pupillometry Test Report. Pupil Labs. 2024.
Gibson JJ. On the Relation between Hallucination and Perception. Leonardo. 1970 Oct;3(4):425.
Linszen MMJ, de Boer JN, Schutte MJL, Begemann MJH, de Vries J, Koops S, et al. Occurrence and phenomenology of hallucinations in the general population: A large online survey. Schizophrenia (Heidelb). 2022 Apr 23;8(1):41.
Linscott RJ, van Os J. An updated and conservative systematic review and meta-analysis of epidemiological evidence on psychotic experiences in children and adults: on the pathway from proneness to persistence to dimensional expression across mental disorders. Psychol Med. 2013 Jun;43(6):1133–49.

How Cortical and Thalamic inputs to the Sensorimotor striatum guide motor sequence execution

Cheshta Bhatia

Bence Olveczky

Department of Organismic and Evolutionary Biology, Harvard University

The brain's ability to sequence movements enables rich and adaptive behavior, such as playing piano sonatas from sheet music. However, such flexibility can result in slow and error-prone performances [1,2]. Practicing the same motor sequence repeatedly can render its execution fast, effortless, and 'automatic' [3,4,5]. The sensorimotor striatum (DLS) within the Basal Ganglia underlies both flexible and automatic sequence execution, yet the roles of cortical (motor cortex, MC) and thalamic inputs to DLS in executing motor sequences and whether and how these depend on task demands is unclear. One possibility is that extensive practice of the same sequence transfers control to the thalamostriatal pathway, such that the sequence can be unambiguously defined in terms of how past (thalamus) begets future behavior (DLS) [6]. However, if there are competing demands for flexibility, such that the elements used in overtrained behavior are also used in other sequences, this could prevent subcortical consolidation of overtrained behavior, rendering it dependent on MC inputs to DLS [7]. To test our hypothesis, we probed neural circuits in expert rats trained to perform motor sequences of three lever presses. We trained two cohorts of rats - one, in which rats were overtrained to perform a single sequence (automatic-only), and second, in which rats performed the same motor sequence also in a visually-guided context (flexible+automatic). We chronically silenced thalamostriatal projections in both cohorts using an intersectional viral approach and observed that performance and kinematics were disrupted in the automatic-only cohort. Interestingly, automatic sequence execution in the flexible+automatic cohort was unaffected. One downside of this technique is that it also silences collateral projections. To specifically target Thalamic inputs to DLS, we use transient optogenetic inhibition using eOPN3 and observed the same result Optogenetic silencing of corticostriatal projections in rats performing the flexible+automatic task impaired kinematics in both modes but only disrupted automatic sequencing. Together, this work provides insights into the hierarchical control of motor sequences, showing that the thalamostriatal circuit controls automatic sequence execution only when there are no demands for flexibility. However, if there is a need to flexibly re-use motor elements, the execution of automatic sequences is controlled by corticostriatal projections.

References

[1] Schneider, W. & Shiffrin, R. M. Controlled and automatic human information processing: I. Detection, search, and attention. Psychol. Rev.84, 1–66 (1977).

[2] Pashler, H. Dual-task interference in simple tasks: data and theory. Psychol. Bull. 116, 220–244 (1994).

[3] Wiestler, T. & Diedrichsen, J. Skill learning strengthens cortical representations of motor sequences. eLife 2, e00801 (2013).

[4] Wymbs, N. F. & Grafton, S. T. The human motor system supports sequence-specific r Exploring the Inner epresentations over multiple training-dependent timescales. Cereb. Cortex 25, 4213–4225 (2015).

[5] Ashby, F. G., Turner, B. O. & Horvitz, J. C. Cortical and basal ganglia contributions to habit learning and automaticity. Trends Cogn. Sci. 14, 208–215 (2010).

[6] Steffen B. E. Wolff et al. Distinct roles for motor cortical and thalamic inputs to striatum during motor skill learning and execution.Sci. Adv.8,eabk0231(2022).

[7] Mizes, K.G.C., Lindsey, J., Escola, G.S. et al. The role of motor cortex in motor sequence execution depends on demands for flexibility. Nat Neurosci 27, 2466–2475 (2024).

Angular Self-Motion Estimation Guides Evasive Turning Decisions in Walking Drosophila

During navigation, animals continuously monitor their movements relative to internal goals and past actions to detect and correct deviations from their intended trajectory. Estimation of self-motion from visual cues as well as vestibular and proprioceptive signals enables movement correction, planning, and execution of new actions. While the role of self-motion in continuous movement adjustments is well studied, its influence on discrete, punctuated actions remains poorly understood. To address this, we leveraged the structure of Drosophila exploratory behavior, where walking trajectories consist of relatively straight runs interrupted by rapid turns, or body saccades. We developed a paradigm in which a brief, non-directional infrared laser pulse—delivered in darkness—transiently increases body temperature and triggers evasive, rapid turns. This approach enables precise temporal control over the transitions between running and turning. In absence of directional cues, we found that angular self-motion, estimated during the preceding run, predicts the direction of the evasive turn (e.g., leftward-drifting runs lead to rightward turns, and vice versa). This "drift-to-contraversive evasive turn" rule mirrors patterns observed in spontaneous body saccades, suggesting that information accumulated about angular drift, during runs, modulates both spontaneous and stimulus-evoked rapid turns. We are currently investigating the circuits underlying evasive turns. Our goal is to understand how angular drift estimates, emerging through feedback loops involving the ventral nerve cord (the insect analogue of the spinal cord), the central complex (the fly's navigational center), and the lateral accessory lobe (a higher premotor brain region) influence the direction of triggered evasive saccades.

self-motion estimation / decision making / sensorimotor transformation / behaviour / drosophila

Virginia Palieri

Eugenia Chiappe

Champalimaud Foundation, Lisbon, PT

Brain-wide Functional Imaging Reveals Neuronal Dynamics of Working Memory in Larval Zebrafish

Working memory is a central component of cognition, where sensory information is temporally retained in a form that can be processed and used to assist future decision-making. However, the dynamic nature of working memory, with activity patterns that evolve across multiple areas, presents a challenge for studying its neural circuit mechanisms. We developed a working memory-like assay in larval zebrafish, in which larvae use past visual cues to decide where to escape when an acoustic stimulus is presented after a time delay. We show that this memory consists not only of an overt component, observable during spontaneous swimming, but also a longer-lasting covert component revealed only when the fish performs a long-latency C-start, a delayed escape that facilitates information integration. Furthermore, we found that this memory is used selectively and that it resets after being enacted, sharing characteristics of mammalian working memory. Leveraging the optical transparency of the zebrafish larva, we conducted brain-wide calcium imaging using a volume-scanning light-sheet microscope during this memory paradigm. We found brain areas involved in the creation, maintenance, and usage of this memory, sharing key features with the behavioral observations. In particular, we found regions in the anterior hindbrain showing persistent activity during the delay period with a timescale matching that of working memory, and a region exhibiting reset dynamics, which might be part of a memory gating mechanism. We propose and constrain models of the circuits underlying this memory and predict escape behavior from high-dimensional neuronal data. Our results

A. Lucas Martins

Catarina Matos

Margarida Pereira

João C. Marques

Michael Orger

Champalimaud Foundation, Lisbon, PT

show that even simple behaviors in teleost larvae can combine to produce an unexpected level of cognitive function, presenting the larval zebrafish as an excellent model to study subcortical brain regions involved in working memory.

working memory / larval zebrafish / brain-wide imaging / subcortical cognition / decision-making

Survival of Self vs Species: a Fly's Perspective

Animals in the wild are often exposed to different stimuli at the same time, leading to behavioral prioritizations. One such context is how reproductive decisions change in presence of approaching threats. Our study aims to address this question of behavioral hierarchy between defensive responses (ie., choosing one's own safety) versus egg laying (ie, the final step in species propagation) using Drosophila melanogaster. We created this naturalistic context in our lab by building a recording setup with sufficient resolution to automatically characterize both egg laying and defensive responses, while recording multiple freely-moving flies at a time. The subsequent assays revealed an interesting dynamic between both the behaviors – they can co-occur within a short temporal window, but it leads to modifications in certain characteristics of each behavior, to variable extents. Here, we will be presenting this behavioral data along with some preliminary analysis to uncover the mechanistic details.

behavioral hierarchy / innate behavior / naturalistic setting / drosophila / reproduction

Saheli Roy

Matheus Farias

Hugo Marques

Nuno Machado

Sara Santos

Marta Moita

Maria Luisa Vasconcelos

Champalimaud Foundation, Lisbon, PT

The heading change of a preceding movement predicts the escape bias in larval zebrafish

To navigate the environment, an animal must know where it is heading. Recent studies have identified a heading direction system in larval zebrafish, whose circuit architecture resembles the fly's central complex. This suggests that zebrafish, like other animals, might use heading information to guide their movements and decisions. However, no behavior has yet been found in larval zebrafish that depends on heading information—supporting the idea that, although they possess a functional heading system, they don't actually use it. In this study we challenge this idea by using high speed behavioral tracking and automated movement classification in one week old zebrafish doing the acoustic startle behavior. At this age larvae cannot know the sound location and respond to it by doing escapes to the left or right that seem to be unbiased. These escapes fall into two types: long latency C-starts (LLCs) and short latency C-starts (SLCs). We then asked if the heading change of the previous movement modulates the direction of the escape. Our results show that the direction of the immediately preceding swim defines the direction of both LLC and SLC escapes. This effect varied by escape type. LLCs exhibited a strong tendency to repeat the direction of the previous bout, regardless of its angular magnitude. Surprisingly, SLCs showed a graded modulation dependent on the heading angle of the prior movement which wrapped around for angles larger than 180 degrees. Relying on the prior movement heading to guide the escape direction may be a simple mechanism to allow the animal to rapidly escape in a direction consistent with an ongoing goal. Our findings provide the first behavioral evidence that larval zebrafish use head-direction information to guide behavior and open the possibility of studying the neural mechanisms that support this ability.

heading direction / navigation / larval zebrafish / acoustic startle / C-starts

João C. Marques

Catarina Matos

A. Lucas Martins

Michael B. Orger

Champalimaud Foundation, Lisbon, PT

Sml sensory neurons regulate reproductive motor output in Drosophila

Nuno J. Machado

Sara Santos

Luisa Vasconcelos

Champalimaud Foundation, Lisbon, PT

Complex behaviors emerge through sophisticated control systems that integrate sensory information with motor output across distributed neural circuits. Using Drosophila egg-laying as a model to study the neuronal circuits responsible for flexible innate behaviors, we identified and characterized a minimal sensory input of 6 neurons in the reproductive system that, when manipulated, leads to marked behavioral output changes. This system exemplifies how sparse neural populations can regulate behavioral control through multi-level circuit interactions. Using split-Gal4 intersectional genetics, we identified a genetically accessible population of sensory neurons ("Sml" that form a direct sensory-to-central pathway from the uterus to the abdominal ganglion (Abg) of the ventral nerve cord (VNC) of the fly, where abdominal motor control is processed. Trans-synaptic labeling revealed that these 6 uterine sensory neurons interface with a high number of neurons in the Abg forming a complex network: local Abg interneurons, ascending projections to brain centers, and descending motor control pathways - forming a distributed circuit architecture. Optogenetic activation of this minimal sensory input triggered coordinated motor outputs, while circuit silencing reduced reproductive output by ~66%, demonstrating the functional significance of this feedback loop. The circuit architecture suggests a control system where uterine sensory neurons likely monitor internal reproductive state (potentially egg position/movement), relay this information to central processing networks, which then modulate motor programs that influence egg progression. The dramatic behavioral complexity arising from activation of just 6 sensory neurons illustrates how distributed neural networks can amplify minimal sensory signals into coordinated motor responses through circuit-level computation.

drosophila / innate behavior / reproductive behavior / neuronal circuits / ventral nerve cord

Covert strategies determine the preservation of neural population dynamics across individuals during a cognitive task

The coordinated activity of neural populations—their latent dynamics— within motor regions of the brain is preserved across individuals from the same species performing a similar behaviour (Safaie, Chang et al. [1]). Such preservation is linked to behavioural similarity and presumably reflects the shared connectivity patterns of their neural circuits. Here we asked whether a cognitive task that can be solved using different covert strategies would reveal latent dynamics whose preservation is linked to the similarity in the strategies themselves. We analysed prefrontal cortex recordings from an associative memory task in which monkeys had to select the target associated with an initial visual cue, following a working memory period in which no information was presented (Tremblay et al. [2]). We computed session-specific latent dynamics using Principal Component Analysis and defined the covert "strategy" based on the evolving relationship of the latent dynamics with the past cue and the future selected target during the memory period. We tested the preservation of latent dynamics across individuals using Canonical Correlation Analysis. Animals solved this task using different "strategies" which clustered into distinct groups. Using this classification, we found that both the preservation of latent dynamics and the ability to decode future target across animals were associated to the similarity of their covert strategies. Moreover, monkeys exhibited idiosyncratic fidgets during the working memory period that were predictive of the future selected target. These fidgets had a stable relationship with the prefrontal cortex latent dynamics and largely drove their across animals. Neural population dynamics thus capture fundamental differences and similarities in neural processes across individuals during both sensorimotor and cognitive processes.

neural dynamics / neural population / working memory / prefrontal cortex /m embodiment

Ioana Lazar

Mostafa Safaie

Juan Alvaro Gallego

Department of Bioengineering, Imperial College London, London, UK

References

[1] Safaie, M., Chang, J., Park, J., Miller, L. E., Dudman, J. T., Perich, M. G., & Gallego, J. A. (2023). Preserved neural dynamics across animals performing similar behaviour. Nature, 623, 765–771. https://doi.org/10.1038/s41586-023-06714-0

[2] Tremblay, S., Testard, C., DiTullio, R. W., Inchauspé, J., & Petrides, M. (2022). Neural cognitive signals during spontaneous movements in the macaque. Nature Neuroscience, 26, 295–305. https://doi.org/10.1038/s41593-022-01220-4

Deep Reinforcement Learning Sheds Light on Neural Manifold Dynamics in Motor Adaptation

Animals can remarkably adapt their behavior to new conditions or in response to external perturbations. This capacity for motor adaptation is a fundamental aspect of motor control and has been the focus of extensive research. However, the mechanisms by which the brain achieves this ability and the neural computations underlying it remain a topic of debate. To explore these mechanisms, we simulate a virtual arm performing target-directed reaching movements, controlled by an artificial neural network trained via deep reinforcement learning. Following baseline training, in which the network generates hand trajectories closely matching empirical observations, we introduce systematic external perturbations that induce motor errors. The model subsequently adjusts its control strategies, recovering near-baseline performance. The resulting adaptation dynamics, including error correction and learning curves, mirror those observed in non-human primate experiments [1]. To examine the underlying neural representations, we apply dimensionality reduction techniques [2] to the network's activity during baseline, adaptation, and washout phases. We find that preparatory neural states organize into a ring-shaped manifold reflecting the geometry of the target space, which rotates during adaptation, consistent with observations in cortical preparatory activity [3]. Moreover, the post-washout manifold differs from its baseline counterpart, despite identical task demands, indicating a persistent neural signature of the perturbation, suggestive of a memory trace similar to what has been observed in primary motor cortex recordings [4]. These results highlight how reinforcement learning models can reproduce both behavioral and neural features of motor adaptation, offering a powerful framework for investigating the computational principles of sensorimotor learning [5].

computational neuroscience / motor adaptation / deep reinforcement learning / neural manifolds

Mattia Della Vecchia [1]

Mehdi Khamassi [2]

N. Alex Cayco Gajic [1]

[1] Ecole Normale Superieure, Paris, FR
[2] Sorbonne Universite, Paris, FR

Economic Behaviors Emerge from Local Learning Rules under Fast Environmental Feedback

Chenguang Li [1]

Hanming Ye [2]

Jeffrey Erlich [1]

[1] University College London, London, UK
[2] Harvard University, Cambridge, MA, USA

Repeatedly choosing well between options is vital for animal survival. Researchers have categorized such decisions in three broad paradigms: (1) cued decision-making, guided by sensory cues about rewards; (2) foraging, based on experience with a resource's value; and (3) strategic games, where choices depend on competitor behavior. Prior work has produced an abundance of models for this behavior, e.g. drift diffusion models [1,2], evidence accumulation [2], marginal value theorem [3,4], bistable attractors [5,6]. But these efforts typically focus on a single behavior at a time, which limits cross-task generalization. To fill this gap, we present a biologically plausible model that combines prior work from Li et al. [7], Pereira-Obilinovic et al. [8], and Loewenstein & Seung [9] into one unified algorithm: a noisy neural network that uses only local synaptic learning rules and acts in continuous, closed-loop interaction with a task environment. The model can explain behaviors across all three behavioral paradigms in two-alternative forced choice (2AFC) tasks: cued choice 2AFC; two-armed bandit tasks (foraging); and matching pennies, a zero-sum game where players win if their choices match a computer opponent's [10]. Through tight environmental feedback loops, models replicate key behavioral features that have previously required distinct models to achieve: (1) random walk dynamics in cued-choice tasks [11], (2) probability matching in two-armed bandit tasks [12-14], (3) learning to be stochastic in matching pennies [10], and (4) autonomous transitions between distinct behavioral states in all tasks. Overall, this work contributes toward a more integrated view of behavior, showing that a diverse set of economic decision-making behaviors can be performed not by a patchwork of specialized algorithms, but by one flexible mechanistic model shaped by environmental demands.

modelling / economic decisions / synaptic learnings rules / theory / behavior

References

[1] R. Bogacz, E. Brown, J. Moehlis, P. Holmes, and J. D. Cohen, “The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks,” Psychological Review, vol. 113, no. 4, pp. 700–765, 2006, doi: 10.1037/0033-295X.113.4.700.

[2] B. B. Scott, C. M. Constantinople, J. C. Erlich, D. W. Tank, and C. D. Brody, “Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats,” eLife, vol.4, p. e11308, Dec. 2015, doi: 10.7554/eLife.11308.

[3] M. H. Cassini, G. Lichtenstein, J. P. Ongay, and A. Kacelnik, “Foraging behaviour in guinea pigs: further tests of the marginal value theorem,” Behavioural Processes, vol. 29, no. 1, pp. 99–112, Apr. 1993, doi: 10.1016/0376-6357(93)90030-U.

[4] E. L. Charnov, “Optimal foraging, the marginal value theorem,” Theor. Popul. Biol., vol. 9, pp. 129–136, 1976.

[5] J. C. Erlich et al., “Mice dynamically adapt to opponents in competitive multi-player games,” May 01, 2025, bioRxiv. doi: 10.1101/2025.02.14.638359.

[6] A. T. Piet, J. C. Erlich, C. D. Kopec, and C. D. Brody, “Rat Prefrontal Cortex Inactivations during Decision Making Are Explained by Bistable Attractor Dynamics,” Neural Computation, vol. 29, no. 11, pp. 2861–2886, Nov. 2017, doi: 10.1162/neco_a_01005.

[7] C. Li, J. Brenner, A. Boesky, S. Ramanathan, and G. Kreiman, “Neuron-level Prediction and Noise can Implement Flexible Reward-Seeking Behavior,” bioRxiv, p. 2024.05.22.595306, May 2024, doi: 10.1101/2024.05.22.595306.

[8] U. Pereira-Obilinovic, H. Hou, K. Svoboda, and X.-J. Wang, “Brain mechanism of foraging: Reward-dependent synaptic plasticity versus neural integration of values,” Proceedings of the National Academy of Sciences, vol. 121, no. 14, p. e2318521121, Apr. 2024, doi: 10.1073/pnas.2318521121.

[9] Y. Loewenstein and H. S. Seung, “Operant matching is a generic outcome of synaptic plasticity based on the covariance between reward and neural activity,” Proceedings of the National Academy of Sciences, vol. 103, no. 41, pp. 15224–15229, Oct. 2006, doi: 10.1073/pnas.0505220103.

[10] H. Wang and A. C. Kwan, “Competitive and cooperative games for probing the neural basis of social decision-making in animals,” Neuroscience & Biobehavioral Reviews, vol. 149, p. 105158, Jun. 2023, doi: 10.1016/j.neubiorev.2023.105158.

[11] K. S. Chaisanguanthum, H. H. Shen, and P. N. Sabes, “Motor Variability Arises from a Slow Random Walk in Neural State,” J. Neurosci., vol. 34, no. 36, pp. 12071–12080, Sep. 2014, doi: 10.1523/JNEUROSCI.3001-13.2014.

[12] R. J. Herrnstein, “Relative and Absolute Strength of Response as a Function of Frequency of Reinforcement,” Journal of the Experimental Analysis of Behavior, vol. 4, no. 3, pp. 267–272, 1961, doi: 10.1901/jeab.1961.4-267.

[13] J. Kubanek and L. H. Snyder, “Matching Behavior as a Tradeoff Between Reward Maximization and Demands on Neural Computation,” F1000Res, vol.4, p. 147, Oct. 2015, doi: 10.12688/f1000research.6574.2.

[14] L. P. Sugrue, G. S. Corrado, and W. T. Newsome, “Matching Behavior and the Representation of Value in the Parietal Cortex,” Science, vol. 304, no. 5678, pp. 1782–1787, Jun. 2004, doi: 10.1126/science.1094765.

(V)Repetition without (V)Repetition: using virtual reality tennis to replicate and expand on a foundational result in motor control

Alice C. Navalho [1]

Filipe S. Rodrigues [2]

Joseph J. Paton [2]

[1] Instituto Superior Técnico, Lisbon, PT
[2] Champalimaud Foundation, Lisbon, PT

It has been nearly 100 years since Bernstein's seminal experiments on expert blacksmiths established repetition without repetion—the idea that reproducible task performance variables arise from variable elemental movements—as a core aspect of dexterous motor control [1]. Here we combined virtual reality (VR), tennis and motion tracking in an attempt to replicate and expand on this foundational study on motor control. A tennis-inspired VR task was developed in Unity with realistic physics simulation to enable the characterization of forehand movement stereotypy within and across tracked features, and conditional on subject expertise and task demands. The task consisted of four phases, starting with only a proximal goal of simply hitting a virtual tennis ball and then progressing to nested goals with a distal component—hitting the ball such that it landed on specific court locations—with increasing precision requirements. In a pilot experiment, expert participants showed more stereotypical trajectories in the center of the racket stringbed than in any other tracked point, echoing Bernstein's original observations. Moreover, ball trajectories became more stereotypical as target locations became narrower, whereas racket movement stereotypy seemed to decrease. These preliminary findings comport with the minimum intervention principle [2], which posits that the nervous system stabilizes task-relevant performance variables rather than intermediate elemental variables. More generally, this work demonstrates VR's potential for studying the human motor system under complex and dynamic conditions without sacrificing experimental control [3].

motor control / DoF Problem/ virtual reality / movement stereotypy / tennis / minimum intervention / principle

[1] Bernstein, Nicholai A. 2015. Dexterity and Its Development. Edited by Mark L. Latash and Michael T. Turvey. London, England: Psychology Press.
[2] Mangalam, Madhur, Mathew Yarossi, Mariusz P. Furmanek, John W. Krakauer, and Eugene Tunik. 2023. "Investigating and Acquiring Motor Expertise Using Virtual Reality." Journal of Neurophysiology 129 (6): 1482–91.
[3] Todorov, E., and Michael I. Jordan. 2002. "A Minimal Intervention Principle for Coordinated Movement." Neural Information Processing Systems, 27–34.

Neuranimat: neurobiologically inspired sensorimotor control based on inverse action objectives

The NEURANIMAT project explores a neurobiologically inspired approach to controlling soft robotic systems, drawing from the architecture and function of the vertebrate spinal cord. As a first step towards reaching that goal, here we leverage recurrent neural networks (RNNs) trained via distillation learning to mimic spinal circuits that transform high-level goals into low-level motor commands. A two-joint four-muscle arm was implemented using the MuJoCo physics engine [1] and controlled by an RNN trained in a combination of reinforcement and supervised learning to reach random targets in horizontal space. After training, the in silico system reproduced a core feature of spinal motor control—convergent force fields (CFFs)—when tested in a simulated replicate of a seminal experiment in the field [2]. However, these patterns did not directly map onto neurons with the initially expected characteristics. Nonetheless, these preliminary results suggest promising directions for identifying general principles in neural motor control, and support the viability of biologically plausible controllers for soft robotic systems.

reinforcement learning / convergent force field / motor synergy / recurrent neural network / soft robotics / motor control / distillation learning / mujoco

Afonso P. Dias [1]

Filipe S. Rodrigues [2]

Beatriz Lourenço [2]

Artur Silva [2]

Caroline Haimerl [2]

Tiago Costa [2]

Joseph J. Paton [2]

[1] Instituto Superior Técnico, Lisbon, PT
[2] Champalimaud Foundation, Lisbon, PT

[1] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 5026–5033. IEEE, 2012.

[2] Emilio Bizzi, Simon F. Giszter, and Ferdinando A. Mussa-Ivaldi. Computations underlying the execution of movement: a biological perspective. Science, 253(5017):287–291, 1991.

Modeling neural adaptation to tendon transfer

Genji Kawakita

Juan Alvaro Gallego

Imperial College London, London, UK

Motor adaptation has been mostly studied in the context of learning to counteract changes in the environment or external forces applied to the body. Yet, adaptation to changes in the neuromusculoskeletal system remains poorly understood. Previous work examined this question through tendon transfer surgeries. A human study [1] investigated within-session adaptation to a virtual tendon transfer under two conditions: compatible surgeries that preserved existing muscle synergies, and incompatible surgeries that disrupted them, thus requiring new patterns of muscle coordination. Incompatible surgeries led to dramatically slower adaptation. A study in monkeys [2] showed that multiple processes may underpin the multi-week process of regaining hand function after incompatible tendon transfer. Thus, different learning processes with distinct timescales may underpin adaptation to changes in the neuromusculoskeletal system. Yet, their details remain unclear. We used a musculoskeletal model (MyoSuite) and reinforcement learning to investigate adaptation to different types of tendon transfer surgeries. We focused on how agents learned to counteract tendon transfers in a hand model performing pose control and object manipulation tasks. We simulated both compatible and incompatible tendon transfers affecting one to four fingers, and analyzed adaptation dynamics and control policy changes following each surgery. Our first results confirm that incompatible transfers require longer adaptation periods with larger initial performance drops compared to compatible transfers, with adaptation difficulty scaling to the number of affected fingers. Even without additional constraints, compatible transfers preserved the original muscle synergies, but incompatible transfers led to new synergies. We aim to gain insights into neural adaptation following persistent neuromusculoskeletal changes.

tendon transfer / reinforcement learning / musculoskeletal model / adaptation / motor control

[1] Berger, D. J., Gentner, R., Edmunds, T., Pai, D. K., & d'Avella, A. (2013). Differences in adaptation rates after virtual surgeries provide direct evidence for modularity. The Journal of neuroscience : the official journal of the Society for Neuroscience, 33(30), 12384–12394. https://doi.org/10.1523/JNEUROSCI.0122-13.2013;
[2] Philipp, R., Hara, Y., Ohta, N., Uchida, N., Oya, T., Funato, T., & Seki, K. (2025). Multi-timescale neural adaptation underlying long-term musculoskeletal reorganization [Preprint]. bioRxiv. https://doi.org/10.1101/2024.10.18.618983;
[3] Caggiano, V., Wang, H., Durandau, G., Sartori, M., & Kumar, V. (2022). MyoSuite – A contact-rich simulation suite for musculoskeletal motor control [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2205.13600

A Unified Perspective on Computation and Stochastic Spiking Neural Activity

Joan Gort

Alfonso Renart

Champalimaud Foundation, Lisbon, PT

Current models of neuronal computation cannot account for the nature of the irregular spike trains observed in the cortex. While being good at implementing computationally useful low dimensional dynamical systems [1,2], the link of current RNN models with real neural dynamics is not well understood [3].

Meanwhile, classical models of irregular desynchronized neocortical activity based on balanced excitation and inhibition, are not designed to implement the general and flexible dynamics neural systems should produce to drive behavior and cognition, despite being able to explain the irregular nature of spike trains and predict most cortical activity statistics [4-6]. Here we present a mathematical model of a spiking neural network that is capable of codifying distributed latent variables while maintaining E-I balance, bridging the advantages of both of the previous paradigms. A mean-field analysis reveals that the combination of these two architectures entails novel dynamical features that are not present in either of the previous models. We show that, when activity in the latent subspace evolves in time, maintaining balance requires coherent fluctuations in average membrane potential across the neurons, and also nontrivial dynamics in the distribution of firing rates, which coexist with constant and decorrelated activity at the level of the global activity of the network. These counterintuitive results provide novel predictions which we are currently trying to test.

In terms of neural tuning curves, in contrast to current low-rank models – whose units generate monotonous and symmetric mixed responses – our network shows more complex kinds of nonlinear mixed selectivity, closer to those observed [7,8].

Our results provide a long-sought bridge between bottom up and top down views of computation as instantiated through the dynamics of recurrent neural circuits.

References

[1] Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. nature, 503(7474), 78-84.

[2] Mastrogiuseppe, F., & Ostojic, S. (2018). Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. Neuron, 99(3), 609-623.

[3] DePasquale, B., Sussillo, D., Abbott, L. F., & Churchland, M. M. (2023). The centrality of population-level factors to network computation is demonstrated by a versatile approach for training spiking networks. Neuron, 111(5), 631-649.

[4] Van Vreeswijk, C., & Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. Science, 274(5293), 1724-1726.

[5] Renart, A., De La Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., & Harris, K. D. (2010). The asynchronous state in cortical circuits. science, 327(5965), 587-590.

[6] Roxin, A., Brunel, N., Hansel, D., Mongillo, G., & Van Vreeswijk, C. (2011). On the distribution of firing rates in networks of cortical neurons. Journal of Neuroscience, 31(45), 16217-16226.

[7] Rigotti, M., Barak, O., Warden, M. R., Wang, X. J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. Nature, 497(7451), 585-590.

[8] Nogueira, R., Rodgers, C. C., Bruno, R. M., & Fusi, S. (2023). The geometry of cortical representations of touch in rodents. Nature Neuroscience, 26(2), 239-250.

Integration Timescales of Visuospatial Task Constraints and Composition Strategies in Sensory Feedback Responses

Pranshu Malik [1]

Paul Gribble [2]

Andrew Pruszynski [2]

[1] Okinawa Institute of Science and Technology Graduate University, Okinawa, JA
[2] University of Western Ontario, London, CA

Changing contexts and intentions continually shape our movement goals, requiring replanning and real-time adjustments. While it is established that stretch reflexes are tuned to our movement goals, the precise timescales at which they update during ongoing movements and the principles linking external perturbations to goal-directed reflex responses still remain unclear. To address these issues, we introduced new spatial targets during a reaching task and randomly applied various mechanical perturbations to the arm after one of four preview delays (0-200 ms relative to target onset), serving as cues to rapidly adjust the reach toward the new targets. We found that stretch reflexes were first modulated in the R3 epoch ~120 ms after target onset, closely following visuomotor reaction times and occurring faster than previously reported in a comparable postural task. Interestingly, we also observed smaller yet reliable reflex modulation within the R1 epoch after longer preview delays, suggesting that even the motor periphery can be independently and quickly tuned toward the updated goal. Overall, these results highlight a tight coupling between reflexive and voluntary control during dynamic motor tasks. Furthermore, nuanced analyses suggest a resource- and time-efficient strategy in the composition of reflex responses, in which local sensory information — such as changes in intramuscular pressure — can evoke task-aligned reflex responses without needing sensory integration from other muscles or sites, particularly during the early phase of long-latency reflexes. This distributed and nested control strategy provides a coherent framework reconciling past and recent findings while addressing overlooked factors in previous literature. We propose novel hypotheses and experimental predictions that extend these insights into movement control, planning, and adaptation.

motor control / stretch reflex / upper limb / online planning and control / electromyography

Decoding individuality in the micro-structure of mouse behavior

Individuality is an intrinsic and essential aspect of animal behavior that emerges even in genetically identical organisms experiencing the same environmental conditions [1]. In the International Brain Laboratory (IBL), over one hundred mice were trained on a visual decision-making task with the explicit goal of establishing a rigorously standardized experimental protocol. This effort led to an automated pipeline that produced proficient mice whose behavior was indistinguishable across seven different labs, when considering trial-level descriptors of behavior [2]. Nevertheless, substantial inter-individual variability was evident in both training time and proficient behavior [2,3], but its nature remains poorly characterized. To address this, we developed a semi-supervised modular segmentation approach to characterize the finegrained temporal structure across multiple behavioral variables (e.g. wheel movement, whisking and licking). This yielded a discrete latent space of syllables which we further analyzed at different timescales.

Variability in the expression of behavioral syllables was highly structured revealing systematic differences across mice, which allowed individual identity to be decoded. Importantly, behavioral signatures were stable across sessions for most mice, indicative of different strategy types or even mouse personality traits. Furthermore, the micro-behavioral structure identified during proficient behavior after training predicted the learning speed of individual mice across training. This ability to retrospectively relate micro-behavioral structure from the proficient phase to the learning phase is

Inês Laranjeira [1]

Matthew Whiteway [2]

Zachary Mainen [1]

Daniel McNamee [1]

[1] Champalimaud Foundation, Lisbon, PT
[2] Center for Theoretical Neuroscience, Columbia University, New York, NY, USA

indicative of long-term stability and suggestive of meaningful differences in behavioral idiosyncrasies that impact learning trajectories or systematic variation in adaptive cognitive strategies.

mouse behavior / behavioral variability / individuality / hidden markov models

References

[1] Bierbach, D., Laskowski, K. L., & Wolf, M. (2017). Behavioural individuality in clonal fish arises despite near-identical rearing conditions. Nature Communications, 8(1). https://doi.org/10.1038/ncomms15361

[2] The International Brain Laboratory, Aguillon-Rodriguez, V., Angelaki, D., Bayer, H., Bonacchi, N., Carandini, M., Cazettes, F., Chapuis, G., Churchland, A. K., Dan, Y., Dewitt, E., Faulkner, M., Forrest, H., Haetzel, L., Häusser, M., Hofer, S. B., Hu, F., Khanal, A., Krasniak, C., & Laranjeira, I. (2021). Standardized and reproducible measurement of decision-making in mice. ELife, 10, e63711. https://doi.org/10.7554/eLife.63711

[3] A Pan-Vazquez, Araujo, Y. S., B McMannon, M Louka, Bandi, A., L Haetzel, Pillow, J., Daw, N., & Witten, I. (2024). Pre-existing visual responses in a projection-defined dopamine population explain individual learning trajectories. BioRxiv (Cold Spring Harbor Laboratory). https://doi.org/10.1101/2024.02.26.582199

Layered, Hierarchical Behavioral Control Underlies Dopamine Signals Across the Striatum During Decision-Making

Midbrain dopamine (DA) neurons send teaching signals, reward prediction errors (RPEs), to all regions of the striatum [1]. However, these signals converge with regionally segregated inputs from all levels of the cortical hierarchy [2] that vary in degree of abstraction and temporal extendedness [3]. Recent work highlights how regionally diverse DA RPEs might update parallel estimates of future reward as informed by region-specific inputs [4,5], yet the behavioral implications of such parallel processing are poorly understood. We used photometry to measure DA release across the striatum in mice performing a value-based decision task. Importantly, two choice options were baited with reward on each trial with a fixed probability in trial blocks, meaning rewards remained baited until harvested. This introduced qualitative differences between the value of choices computed by linearly weighting past choices and rewards, and their true reward probability, which grew the longer a choice was not selected, both of which influenced behavior. DA exhibited quantitative increases in timescale, but also qualitative differences in sensitivity to task structure, from dorsolateral (DLS) to ventral striatum (VS). DLS DA reflected RPEs in relation to a weighted average of past choices and rewards, whereas VS DA reflected knowledge about the growing value of unchosen options. These data indicate that dopaminergic RPEs in different striatal regions derive from, and likely act on, a hierarchy of representations for behavioral control. We hypothesize that this reflects a scenario in which hierarchically-organized behavioral policies are learned in parallel by basal ganglia circuits, which jointly influence action selection through a computationally layered architecture for the overall control of behavior.

dopamine / striatum / hierarchical control / decision-making

Renato Sousa

Rodrigo Martins

Joseph J. Paton

Champalimaud Foundation, Lisbon, PT

References

[1] Schultz, Wolfram, et al. Science (1997);
[2] Hunnicutt, Barbara J., et al. Elife (2016);
[3] Murray, John D., et al. Nature Neuroscience (2014);
[4] Mohebi, Ali, et al. Nature Neuroscience (2024);
[5] Lee,Rachel, et al. Nature Neuroscience (2024).

Flies Rely on Numerosity to Assess the Relative Safety of the Group

Numerosity, the ability to estimate the number or proportion of items in a set, has been reported throughout the animal kingdom. Drosophila melanogaster is no exemption and has previously been shown to be able to discriminate between groups with fewer or more members, and to prefer bigger group sizes [1,2]. Interestingly, we have found that flies exhibit numerosity sensitivity in a predation context, namely the safety in numbers effect. Previously we had found that moving conspecifics provide social safety cues that reduce freezing in response to looming threats [3]. Here we investigated how flies of both sexes integrate social motion cues across different group sizes. We uncovered that flies decrease their sensitivity to the total social motion in bigger groups and that flies respond to the fraction of moving flies independently of group size. Specifically, the probability of exiting freezing is similar when the same proportion of the group is moving, regardless of group size. Our results suggest that flies rely on relative, not absolute, measures of group activity. These results also indicate that numerosity processing enables scale-invariant assessment of social movement, allowing flies to flexibly gauge safety cues in a group. This work reveals a novel role for numerosity in social modulation of defensive responses, showing how simple nervous systems combine social sensory information to guide survival decisions.

numerosity / freezing / social buffering / behavior / safety in numbers

Mirjam Heinemans [1]

Marta Moita [1]

Clara Ferreira [2]

[1] Champalimaud Foundation, Lisbon, PT
[2] Northumbria University, Newcastle, UK

References

[1] Bengochea, M., Sitt, J. D., Izard, V., Preat, T., Cohen, L., & Hassan, B. A. (2023). Numerical discrimination in Drosophila melanogaster. Cell Reports;
[2] Rooke, R., Rasool, A., Schneider, J., & Levine, J. D. (2020). Drosophila melanogaster behaviour changes in different social environments based on group size and density. Communications Biology;
[3] Ferreira, C. H., & Moita, M. A. (2020). Behavioral and neuronal underpinnings of safety in numbers in fruit flies. Nature Communications

Task-Optimized Convolutional Recurrent Networks Align with Tactile Processing in the Rodent Brain

Tactile sensing remains far less understood in neuroscience and less effective in artificial systems compared to more mature modalities such as vision and language. We bridge these gaps by introducing a novel Encoder-Attender-Decoder (EAD) framework to systematically explore the space of task-optimized temporal neural networks trained on realistic tactile input sequences from a customized rodent whisker-array simulator. We identify convolutional recurrent neural networks (ConvRNNs) as superior encoders to purely feedforward and state-space architectures for tactile categorization. Crucially, these ConvRNN-encoder-based EAD models achieve neural representations closely matching rodent somatosensory cortex, saturating the explainable neural variability and revealing a clear linear relationship between supervised categorization performance and neural alignment. Furthermore, contrastive self-supervised ConvRNN-encoder-based EADs, trained with tactile-specific augmentations, match supervised neural fits, serving as an ethologically-relevant, label-free proxy. For neuroscience, our findings highlight nonlinear recurrent processing as important for general-purpose tactile representations in somatosensory cortex, providing the first quantitative characterization of the underlying inductive biases in this system. For embodied AI, our results emphasize the importance of recurrent EAD architectures to handle realistic tactile inputs, along with tailored self-supervised learning methods for achieving robust tactile perception with the same type of sensors animals use to sense in unstructured environments.

embodied neuroai / tactile/ somatosensory cortex / convrnn / rodent whisking

Trinity Chung [1]

Yuchen Shen [2]

Nathan C. L. Kong [4]

Aran Nayebi [1,2,3]

[1] Robotics Institute, Carnegie Mellon University; Pittsburgh, PA 15213
[2] Machine Learning Department, Carnegie Mellon University; Pittsburgh, PA 15213
[3] Neuroscience Institute, Carnegie Mellon University; Pittsburgh, PA 15213
[4] Department of Psychology, University of Pennsylvania; Philadelphia, PA 19104

Identifying Termination Rules in Perceptual Decision Making

Juan R. Castiñeiras de Saa

Mafalda Valente

Ulysses Tsai

Alfonso Renart

Champalimaud Foundation, Lisbon, PT

Both descriptive and normative theories of decision making over the last decades [1] have consolidated a model whereby decision commitment takes place when the relative evidence favoring one alternative reaches a certain decision bound. However, sensory evidence is often terminated before a decision bound has been reached.

How is commitment achieved under incomplete evidence? Almost universally, it has been assumed that decision-makers will, in these conditions, choose the option to which the decision variable is closest when evidence is terminated [2]. However, the empirical validity of this rule has not been scrutinized against alternatives. Here, we studied this problem in a rat sound lateralization task which we have previously shown is very well described by a model based on bounded accumulation of evidence [3].

To study early termination, we imposed varied maximum sound durations, sometimes shorter than the typical reaction time (RT). The baseline RT behavior was well described by a variant of the drift diffusion model (DDM) including both stimulus dependent and anticipatory processes [4]. However, the standard early termination rule failed to describe choice accuracy for short duration stimuli, and had to be modified in two ways: first, by considering that evidence integration continues after offset (approximated by exponential decay of the neural activity providing the sensory input). Second, by assuming that if the decision bound is not reached during this period, the subsequent choice is random, implying that rats appear to fail to translate accumulated evidence into choices unless the decision bound is hit within the extended integration period. Overall, this work quantifies decision-making under temporal constraints, indicating that bound-crossing is a fundamental requirement to reliably convert integrated sensory evidence into a specific behavioral choice.

embodied neuroai / tactile/ somatosensory cortex / convrnn / rodent whisking

References

[1] Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. Neu-ral computation, 20(4), 873–922.
[2] Green, D. M., Swets, J. A., et al. (1966). Signal detection theory and psychophysics (Vol. 1). Wiley New York.
[3] Pardo-Vazquez, J. L., Castiñeiras-de Saa, J. R., Valente, M.,Damião, I., Costa, T., Vicente, M. I., . . . Renart, A. (2019). The mechanistic foundation of weber's law. Nature neuroscience, 22(9), 1493–1502.
[4] Hernández-Navarro, L., Hermoso-Mendizabal, A., Duque, D., De la Rocha, J., & Hyafil, A. (2021). Proactive and reactive accumulation-to-bound processes compete during perceptual decisions. Nature communications, 12(1), 7148.

Addressing non-stationarities in neural population dynamics for long-term behavioral decoding

Intracortical Brain-Computer Interfaces (iBCI) decode behavior from neural population dynamics to restore motor functions and communication abilities in individuals with motor impairments.

A central challenge for long-term iBCI deployment is the nonstationarity of the recorded population activity, where electrode instability changes the composition and tuning of the recorded population across sessions. Existing approaches address this issue with explicit alignment techniques; however, they rely on fixed neural identities and require test-time labels and parameter updates, limiting their generalization across sessions and imposing extra computational burden during deployment. We introduce SPINT, a Spatial Permutation-Invariant Neural Transformer framework for behavioral decoding that operates directly on unordered sets of neural units. Our approach is centered on a novel context-dependent positional embedding that dynamically infers unit-specific identities, enabling flexible generalization across sessions. SPINT supports inference on variable-size populations and allows few-shot, gradient-free adaptation using only a small amount of unlabeled data from the test session. We evaluate our approach on three motor decoding tasks from the FALCON Benchmark, demonstrating robust cross-session generalization that outperforms existing zero-shot and few-shot unsupervised baselines while eliminating the need for test-time alignment and fine-tuning. Our work contributes an initial step toward a flexible and practical framework for robust, scalable neural decoding in long-term iBCI applications.

intracortical brain-computer interface

Trung Le [1]

Hao Fang [1]

Jingyuan Li [1]

Tung Nguyen [2]

Lu Mi [3]

Amy Orsborn [1]

Uygar Sümbül [4]

Eli Shlizerman [1]

[1] University of Washington, Seattle, US
[2] University of California, Los Angeles, Los Angeles, US
[3] Tsinghua University, Beijing, CN
[4] Allen Institute, Seattle, US

poster session

2

Measuring the effects of serotonergic psychedelics on the structure of neural population activity

Davide Crombie [1]
Jaime Arlandis [1]
Joana Catarino [2]
Daniel McNamee [2]
Guido Meijer [3]
Martin Zuniga [4]
Zachary Mainen [1]

[1] Champalimaud Foundation, Lisbon, PT
[2] Karolinska Institute, Solna, SE
[3] Donders Institute, Nijmegen, NL
[4] Princeton University, Princeton, US

Classical psychedelics primarily act via serotonin 2A receptors to cause changes in perception and cognition, and have shown promise in the treatment of various mental disorders [1]. Studies in humans often report that psychedelics increase the brain-wide spatiotemporal complexity of aggregate measures of neural activity [2], but such effects have not been examined at the level of populations of individual neurons [3].

We performed high-yield electrophysiological recordings from awake mice (n=16) to record the spiking activity of >17,000 neurons across forebrain structures including the prefrontal cortex, striatum, thalamus, amygdala, and hippocampus, while administering either LSD or saline via intragastric catheter. We found that the brain-wide spatiotemporal complexity of neural activity increased after LSD administration, but, surprisingly, this was also true after saline. While LSD doesn't drive broad increases in the complexity of neural population activity, it may be that specific structuring motifs that normally act to reduce the effective complexity of population activity – such as movement and inter-regional connectivity [4,5] – are specifically disrupted, and our ongoing analyses assess these relationships.

Overall, we found that ~80% of neurons significantly changed their firing rate under LSD, but observed a similar proportion of modulation in the control condition. However, the two conditions could be distinguished by the distinct anatomical profile of firing rate modulation, implicating the prefrontal cortex and hippocampus as important mediators

for the effects of LSD. Our results raise critical questions about psychedelic-induced complexity increases measured in humans, and highlight that the effects of LSD may be driven by more subtle and anatomically specific changes in neural population activity than often suggested by contemporary results.

psychedelics / LSD / neuropixels / neural population activity / entropy

References

[1] Vollenweider, F. X., & Preller, K. H. (2020). Psychedelic drugs: neurobiology and potential for treatment of psychiatric disorders. Nature Reviews Neuroscience, 21(11), 611-624.
[2] Carhart-Harris, R. L., & Friston, K. J. (2019). REBUS and the anarchic brain: toward a unified model of the brain action of psychedelics. Pharmacological reviews, 71(3), 316-344.
[3] Kwan, A. C., Olson, D. E., Preller, K. H., & Roth, B. L. (2022). The neural basis of psychedelic action. Nature Neuroscience, 25(11), 1407-1419.
[4] Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., & Harris, K. D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. Science, 364(6437), eaav7893.
[5] MacDowell, C. J., Libby, A., Jahn, C. I., Tafazoli, S., Ardalan, A., & Buschman, T. J. (2025). Multiplexed subspaces route neural activity across brain-wide networks. Nature communications, 16(1), 3359.

Cortico-subcortical dynamics during complex locomotion

Locomotion is a fundamental behavior across the animal kingdom. Much research has focused on the role of central pattern generators in locomotion. Yet, the role of supraspinal structures in modulating this behavior, especially during perturbations, remains elusive. To address this question, we designed a task based on delivering rapid mechanical perturbations to head-fixed mice running on a large spherical treadmill. We applied random, unpredictable perturbations from 12 different positions spanning six locations and two different elevations. This allowed us to elicit a broad range of behavioral responses determined by the direction and duration of the perturbation as well as its timing relative to the mouse' ongoing gait cycle. To study sensorimotor integration along the cortico-basal pathways, we recorded from limb-specific subregions of the sensorimotor cortices, downstream basal ganglia projections, and relay centers in the motor thalamus projecting back to cortex. Simultaneously, we tracked whole-body 3D kinematics to define how neural dynamics within these regions contribute to sensorimotor corrections during locomotion. Our analysis is focused on understanding how the cortico-basal ganglia pathways translate sensory responses into motor corrections. Our first findings show the presence of rapid perturbation-specific responses in all the examined cortico-thalamo-striatal regions, with similar latencies. However, the perturbations reveal a hierarchy of functional interactions between these areas, which remains obscured during unperturbed running. Our data also indicate a switch in overall neural "state" between and initial unperturbed period and the inter-perturbation running periods despite the similarity of the kinematics between these two epochs. This change, which was most prominent in the motor cortex, is potentially suggestive of a switch in control policy.

electrophysiology / population dynamics / locomotion

Martín Esparza-Iaizzo [1,2]

Ioana Lazar [1]

Catia Fortunato [1]

Mostafa Safaie [1]

Juan Álvaro Gallego [1]

[1] Department of Bioengineering, Imperial College London, London, UK
[2] INBRAIN Neuroelectronics S.L, Barcelona, Spain

Eigenmode analysis of spherical brain activity via neural field theory

K. N. Mukta [1,2]
J. N. MacLaurin [1,2]
P. A. Robinson [1,2]

[1] School of Physics, University of Sydney, New South Wales 2006, Australia
[2] Center for Integrative Brain Function, University of Sydney, New South Wales 2006, Australia

Corticothalamic neural field theory (NFT) has successfully explained a wide variety of phenomena, ranging from EEG spectra and evoked potentials to nonlinear phenomena such as seizures and Parkinsonian oscillations. NFT has also been used to understand brain connectivities and its eigenmodes have been employed to solve the inverse problem of determining brain structure from functional connectivity. Most recently, its unihemispheric eigenmodes have been shown to be remarkably similar to spherical harmonics in structure. They are also the building blocks for bihemispheric modes, whose structure and symmetry properties explain many features of resting state and task-related activity. This eigenmode expansion is of use because it helps us understand the dynamics of the brain's activity in terms of its natural modes.

Here, corticothalamic NFT is analyzed on a sphere and used to derive the transfer function, the power spectrum, the correlation function, and the cross spectrum in terms of spherical harmonics. The results are analyzed and compared with planar NFT in both finite and infinite geometries. The results of spherical and finite-planar geometries converge to the infinite-planar geometry in the limit of large brain size. The main effects of the spherical modal structure are explored, particularly to understand the number of modes that contribute significantly to these observable quantities and the effects of the finite spatial extent of the cortex. Key results are that when we truncate the modal series it is found that, for physiology plausible parameters, only the lowest few spatial eigenmodes are needed for an accurate representation of macroscopic brain activity. Cortical

modal effects can lead to a double alpha peak structure in the power spectrum, although the main determinant of the alpha peak is corticothalamic feedback. In the spherical geometry, the coherence function between points decays monotonically as their separation increases at a fixed frequency, but persists further at resonant frequencies. The correlation between two points is found to be positive, regardless of the time lag and spatial separation, but decays monotonically as the separation increases at fixed time lag. At fixed distance the correlation has peaks at multiples of the period of the dominant frequency of system activity.

This analysis of physiologically-based corticothalamic NFT in a spherical geometry will enable more realistic modeling and analysis of experimental brain signals in future.

corticothalamic neural field theory / eigenmodes / EEG power spectrum

Freezing Duration in Drosophila Reflects Bounded Accumulation of Evidence about Safety

Marco Colnaghi [1]
Juan Castiñeiras [1]
Mirjam Heinemans [1]
Clara Ferreira [2]
Marta Moita [1]
Alfonso Renart [1]

[1] Champalimaud Foundation, Lisbon, PT
[2] Northumbria University, Newcastle, UK

Freezing is a conserved defensive strategy observed in many species in response to threat. In Drosophila, both the probability and duration of freezing depend on several factors, including the animal's social context and internal state [1,2]. To gain a mechanistic understanding of how these factors quantitatively influence freezing duration, we analyzed data from a paradigm in which flies were exposed to looming stimuli of different speeds in the presence of conspecifics whose movement was under experimental control [3]. We developed a theory in which the dynamics of freezing are governed by bounded accumulation of evidence (BAE; [4]) about safety, where a single decision bound represents the amount of safety deemed appropriate by the animal to resume movement.

Our analysis revealed two coexisting processes even within individuals - a short and a long freezing mode - which are selected probabilistically by flies at the time of loom. How do the drift and decision bound of each freezing mode, as well as the probability of entering either, depend on experimental conditions such as loom speed? Freezing duration is very accurately described by a model where freezing mode probability is a logistic function of all experimental conditions, but where the evidence that is integrated towards the bound depends only on the moment-to-moment fluctuations in the amount of social motion experienced by the fly.

Our model accurately describes the whole spectrum of variation of the full freezing duration distribution across all conditions in our experiments with great accuracy and parsimony. The mathematical

precision of the behavior together with the powerful toolkit for circuit mapping in Drosophila, uniquely positions this paradigm as an ideal model for unraveling a neural implementation of BAE, a canonical computational motif for the temporal sequencing of behavior.

freezing / social / behavioral modelling / evidence accumulation / drosophila

References

[1] Ferreira, C. H., & Moita, M. A. (2020). Behavioral and neuronal underpinnings of safety in numbers in fruit flies. Nature Communications, 11(1), 4182.
[2] Zacarias, R., Namiki, S., Card, G. M., Vasconcelos, M. L., & Moita, M. A. (2018). Speed dependent descending control of freezing behavior in Drosophila melanogaster. Nature Communications, 9(1), 1.
[3] Ferreira, C. H., Heinemans, M., Farias, M., Gonçalves, R., & Moita, M. A. (2022). Social Cues of Safety Can Override Differences in Threat Level. Frontiers in Ecology and Evolution, 10.
[4] Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. Neural Computation, 20(4), 873–922.

Contributions of the ACx to Fine Discriminations of Sound Lateralization

Mafalda Valente

Juan Castiñeiras

Raghavendra Kaushik Archak

Alfonso Renart

Champalimaud Foundation, Lisbon, PT

While the Auditory Cortex (ACx) is clearly involved in processing auditory stimuli, its necessity in simple auditory decisions—such as sound localization, based on interaural level differences (ILDs)—remains controversial. Lesion studies in rodents suggest minimal impairment in simple localization tasks after ACx damage [1–3]; however, deficits are often observed in finer perceptual decision-making tasks [4,5] Therefore, ACx may play a modulatory role. We used ibotenic acid lesions and transient optogenetic silencing of ACx neurons to directly examine their contribution to an ILD discrimination task.

While lesioning the cortex yielded no substantial effects, optogenetic silencing significantly impaired decision accuracy and increased anticipatory behaviors, with higher rates of fixation aborts. Somewhat unexpectedly, reaction times (RTs) were shorter, contrary to our initial predictions that sensory disruption would slow decisions. We used an adapted drift-diffusion model (DDM), incorporating two parallel processes: a proactive (anticipatory) process, and a reactive (stimulus-driven) process of evidence accumulation [6,7]. While simply increasing the anticipatory (proactive) drive can mimic the increase in anticipatory responses, it did not account for the reduced accuracy nor the decreased RTs.

Reducing the firing rate of sensory neurons, and increasing noise within the evidence accumulation process, were required in the model to accurately reproduce the full pattern of behavioral changes. Thus, ACx silencing appears to simultaneously degrade the sensory evidence and increase anticipatory behaviors. Our work demonstrates the strength of combining targeted neural manipulations with detailed computational modelling, providing deeper insight into how cortical areas modulate perceptual decisions, and these findings support a nuanced role for ACx in auditory decision-making.

auditory cortex / drift-diffusion model / optogenetics / sound localization

References

[1] Kelly, Jack B., and Bruce Masterton. "Auditory sensitivity of the albino rat." Journal of comparative and physiological psychology 91.4 (1977): 930.;
[2] Kelly, JACK B. "Effects of auditory cortical lesions on sound localization by the rat." Journal of neurophysiology 44.6 (1980): 1161-1174.;
[3] Pai, Shraddha, et al. "Minimal impairment in a rat model of duration discrimination following excitotoxic lesions of primary auditory and prefrontal cortices." Frontiers in systems neuroscience 5 (2011): 74.;
[4] Znamenskiy, Petr, and Anthony M. Zador. "Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination." Nature 497.7450 (2013): 482-485.;
[5] Guo, Lan, Nicholas D. Ponvert, and Santiago Jaramillo. "The role of sensory cortex in behavioral flexibility." Neuroscience 345 (2017): 3-11.;
[6] Pardo-Vazquez, Jose L., et al. "The mechanistic foundation of Weber's law." Nature neuroscience 22.9 (2019): 1493-1502.;
[7] Hernández-Navarro, Lluís, et al. "Proactive and reactive accumulation-to-bound processes compete during perceptual decisions." Nature communications 12.1 (2021): 7148.

Motor Cortex Encodes Speed Through Temporal Scaling of Latent Population Dynamics

Skilled movements require precise control of their speed, yet how the motor cortex encodes speed remains unresolved. One proposal—the axis hypothesis—suggests that different speeds of the same movement are generated by distinct trajectories of neural activity along a speed axis. An alternative, the traversal hypothesis, posits that different speeds are generated by how quickly a fixed neural trajectory unfolds over time. To evaluate these competing accounts, we analysed intracortical recordings from PMd and M1 in monkeys performing a random-target arm reaching task [1]. Since movement speed and length co-vary in reaching tasks, we isolated the effects of speed by analysing only reaches of a fixed length. We found that neural trajectories diverged before movement onset in proportion to the angular difference between upcoming movement directions. RNN modelling of the traversal hypothesis predicted that these trajectories' maximum divergence should occur later for slower movements and that their maximum traversal rate should correlate with upcoming movement speed; both predictions were confirmed in PMd and M1 data. To directly contrast the hypotheses, we asked if neural trajectories for the same direction of upcoming arm movement were ordered by its speed, but found they were not; indeed they were no more distinct than were trajectories of the same upcoming movement speed before the start of movement preparation, contradicting the axis hypothesis. These results support the traversal hypothesis, suggesting that movement speed is encoded not by switching between distinct neural trajectories, but by modulating the timing of transitions within a shared neural state sequence. Understanding this encoding mechanism could improve the design of brain-machine interfaces by enabling more precise decoding of intended movement speed from neural activity.

motor cortex / speed encoding / population dynamics / motor control

Kumar Neelabh [1]

Andrea Colins Rodriguez [2]

Mark D Humphries [3]

[1] Champalimaud Foundation, Lisbon, PT
[2] Universidad Adolfo Ibáñez, Santiago, CL
[3] University of Nottingham, Nottingham, UK

[1] Perich, Matthew G.; Miller, Lee E.; Azabou, Mehdi; Dyer, Eva L. (2025) Long-term recordings of motor and premotor cortical spiking activity during reaching in monkeys (Version draft) [Data set]. DANDI Archive. https://dandiarchive.org/dandiset/000688/draft

Mouse lockbox: a sequential mechanical decision-making task for freely moving mice

Advances in automated tracking tools have sparked a growing interest in studying naturalistic behavior in animals. Yet, traditional decision-making tasks remain the norm for assessing learning behavior in neuroscience. Here, we present an alternative sequential decision-making task to study mouse behavior. We developed a 3D-printed mechanical puzzle, a so-called lockbox, that requires a sequence of four steps to be solved in a specific order. Each mechanism is easily movable by the mice, but requires distinct manipulations, such that the mechanisms cannot be solved by accident. During the task, the mice move around freely, enabling the emergence of complex behavioral patterns.

We observed that mice learned this relatively complex, four-step task surprisingly quickly compared to conventional operant tasks, demonstrating its potential utility of more ethologically relevant challenges. To delineate different contributions to the learning process, we recorded their behavior in a multi-camera setup and developed a data analysis pipeline to automatically detect the interactions of the mice with the different lockbox mechanisms for a large corpus of video footage (> 110 h, 12 mice). Our analysis pipeline allowed detailed quantification of behavior, revealing that learning involved both improvement in object manipulation skills and the emergence of a task-specific strategy. While the rapidly increasing task performance seems primarily due to the improvement in object manipulation, clear signs of a cognitive strategy for the task appear during later trials. The lockbox

Marcus Nicolaas Boon [1,2]

Niek Andresen [1,2,4]

Sole Traverso [1,2]

Sophia Meier [1,2]

[1] Cluster of Excellence "Science of Intelligence", Berlin, Germany, DE
[2] Technische Universität Berlin, Germany, DE
[3] German Centre for the Protection of Laboratory Animals, German Federal Institute for Risk Assessment, Berlin, Germany, DE
[4] Freie Universität Berlin, Berlin, Germany, DE

task therefore could offer a promising paradigm for studying how low-level motor learning interacts with high-level decision-making strategies within a single, ethologically relevant task.

behavior / decision-making / mice / mechanical puzzle

↓

Friedrich Schuessler [1,2]

Olaf Hellwich [1,3,4]

Lars Lewejohann [1,3,4]

Christa Thöne-Reineke [1,4]

Dynamical Archetype Analysis: Autonomous Computation

Abel Sagodi

Memming Park

Champalimaud Foundation, Lisbon, PT

The study of neural computation aims to understand the function of a neural system as an information processing machine. Neural systems are undoubtedly complex, necessitating principled and automated tools to abstract away details to organize and incrementally build intuition. We argue that systems with the same effective behavior should be abstracted by their ideal representative, i.e., archetype, defined by its asymptotic dynamical structure. We propose a library of archetypical computations and a new measure of dissimilarity that allows us to group systems based on their effective behavior by explicitly considering both deformations that break topological conjugacy as well as diffeomorphisms that preserve it. The proposed dissimilarity can be estimated from observed trajectories. Numerical experiments demonstrate our method's ability to overcome previously reported fragility of existing (dis)similarity measures for approximate continuous attractors and high-dimensional recurrent neural networks. For example, our model can reconstruct deformed and perturbed ring attractors, while the methods Dynamical Similarity Analysis [1] and Smooth Prototype Equivalences [2] struggle to identify expected similarities and differences. Our experiments focus on working memory systems – the backbone of most temporal computation, but our theoretical approach naturally extends to general mechanistic interpretation of recurrent dynamics in both biological and artificial neural systems. We argue that abstract dynamical archetypes, rather than detailed dynamical systems, offer a more useful vocabulary for describing neural computation.

interpretability / dynamical dissimilarity / dynamical archetypes / neural computation / working memory / computation through dynamics

References

[1] Ostrow, M., Eisen, A., Kozachkov, L., & Fiete, I. (2023). Beyond geometry: Comparing the temporal structure of computation in neural circuits with dynamical similarity analysis. Advances in Neural Information Processing Systems, 36, 33824-33837.

[2] Friedman, R., Moriel, N., Ricci, M., Pelc, G., Weiss, Y., & Nitzan, M. (2025). Characterizing Nonlinear Dynamics via Smooth Prototype Equivalences. arXiv preprint arXiv:2503.10336.

Facilitating cognitive synergy: AI’s ability to support collaborative cognitive search strategies

Eric Lacosse [1]

Mariana Duarte [1]

Graham Todd [2]

Peter M. Todd [3]

Daniel C. McNamee [1]

[1] Champalimaud Foundation, Lisbon PT
[2] New York University, NYC
[3] Indiana University, Bloomington, IN

Recent advances in large language models (LLMs) have garnered attention for their unprecedented text generation capabilities. Indeed, in many language-based tasks, the need for human input has been greatly reduced or even eliminated. However, rather than replacing humans, LLMs may serve as cognitive aids to humans particularly when they are engaged in a task requiring open-ended conceptual exploration and creativity. Our perspective is that we are yet to understand how these models may enhance human cognitive abilities in such a paradigm of human-AI interaction. In this study, we explore and evaluate distinct interaction mechanisms designed to enhance the synergy between the cognitive processes of humans and their collaborative LLMs in order to optimize collective performance. By focusing on semantic memory search, a classic cognitive experimental task used to interrogate convergent and divergent thinking in humans, we will be able to tractably analyze large-scale datasets of human-LLM interactions using well-established theoretical tools. In particular, we will test whether such cognitive synergy can enable the paired human-LLM collaboration partners to surpass individual or paired human search performance, and thus overcome previously observed collaborative inhibition effects. More generally, our core hypothesis to be tested in this study is that human-AI collaboration can be designed to synergize cognitive strategies, whereby human contextual understanding and AI’s broad associative capabilities may fruitfully combine to explore conceptual spaces more effectively than either can achieve alone. We believe that this would provide solid empirical evidence and guidance regarding an alternative framework for collaborative human-AI engagement.

LLMs / cognitive synergies / cognitive alignment

Larval Zebrafish's Sophisticated Short-term Memory Shares Key Features of Working Memory with Mammals

Catarina Matos

A. Lucas Martins

João C. Marques

Michael B. Orger

Champalimaud Foundation, Lisbon PT

Working memory is a central component of cognition, where sensory information is kept in mind and processed to make better-informed decisions. Working memory is considered to exist only in a few select species, namely humans and monkeys, and to a lesser degree rodents, adult zebrafish and honeybees. In this project, we are challenging this idea by showing that one-week-old larval zebrafish also have complex memories, sharing many of the central characteristics of working memory with mammals. These characteristics include being short term, lasting for around 20 to 30 seconds, its use being dependent on context, being covert, resetting after use and the ability to manipulate information kept in memory. The detailed characterization of larval zebrafish's' natural behaviours allows us to use them as building blocks to create behavioural assays. As such, we developed a working memory-like assay in larval zebrafish, which includes the use of a past visual cue as information to guide decision-making when the fish escapes an acoustic stimulus, presented after a time delay. Leveraging this assay, we found evidence that zebrafish larvae can hold the memory of the cue for around 20 seconds, showing a timescale comparable to mammalian working memory. After a first decision is made however, a second decision is unbiased, showing the memory is reset. We also found that this memory is not visible on the animal's behaviour, being covert and is only used in contexts where there is lack of information. Additionally, we modified the original assay to create a memory manipulation assay. This allowed us to show that zebrafish larvae can manipulate information obtained from at least two distinct cues to use them together to make a more accurate decision. In this project, we were able to show that zebrafish larvae form memories that share main characteristics of mammalian working memory.

working memory / larval zebrafish / behaviour / cognition

Head-body coordination for high-performance courtship pursuit in Drosophila

Nuno Rito [1]
Martina Canova [1,2]
Miguel Paço [1]
Claire Rusch [1]
Sofia Pinto [1]
M.Eugenia Chiappe [1]

[1] Champalimaud Foundation, Lisbon PT
[2] University of Bonn, Bonn, Germany

Pursuit behaviors require the coordination of movements across the body, often balancing posture control for the body's mechanical stability with gaze control for image stabilization of the moving target, enabling accurate sensorimotor calculations critical for predatory and mating decisions. However, how visual information is integrated with ongoing movement across the body to direct the pursuer's subsequent actions remains poorly understood. We address this question by evaluating the role of head-body coordination on pursuit success, leveraging the elaborate courtship behavior of Drosophila melanogaster.

Through a combination of quantitative analysis of behavior, optogenetics, circuit manipulations, and physiology, we found that males move their heads in exquisite coordination with their bodies to keep a female within a specialized functional region in their retina. Masking this retinal region compromises pursuit, while fixing the head of males increases the steering demands of the body and the overall stability of pursuit. Moreover, when head-fixed males are challenged by the addition of competitors, their chances of copulation are reduced below the level expected by competition. These observations underscore the critical role of gaze control on pursuit performance in naturalistic conditions. Using a VR system in which aroused, tethered males track fictive females, we found that by fixing the flies' gaze at different angles, males integrate head position—rather than head velocity—with retinal information to orient toward female avatars. Analysis of large-scale electron microscopy datasets of the adult's fly brain, in combination with circuit manipulation, identified

a circuit involved in the coordination of head and body movements, from the visual sensory pathway to the motor neurons of the neck that control head movements, and back to the brain.

In summary, our findings suggest that courting Drosophila males integrate retinal input with gaze direction, transforming these signals into a coordinate system suitable for guiding pursuit, dynamically tuning head movements for steering stability while responding quickly to the female's uncertain movements.

head-body coordination / gaze stabilization / visuomotor integration

Planning to Learn: Active Inference of Latent Dynamics with Fisher Information

Understanding the latent dynamical structure that governs behavior and neural activity is a central challenge in computational neuroscience and many complex system domains. However, performing joint state and system identification for high-dimensional nonlinear dynamics from partial, noisy observations remains data-intensive and often sample-inefficient under naïve data acquisition strategies.

We propose a computationally efficient active learning framework for inferring latent dynamics models that plans optimized input sequences to maximize future information gain. Our method combines local Fisher Information metrics, which quantify parameter sensitivity in probabilistic latent dynamics, with predictive uncertainty reduction strategies derived from ensemble disagreement or posterior entropy. To address the compounding effect of early model inaccuracies, we introduce discount factors that regularize information gathered under high epistemic uncertainty in the latent posterior, ensuring robustness in belief updates over time. We validate our approach through computational experiments in various nonlinear dynamics environments from classical control, and simulated neural population dynamics, showing improved sample efficiency over myopic and uncertainty-driven baselines. Our method offers a flexible and general framework for structured system identification and active experimental design in partially observed dynamical systems.

neural dynamics / active learning / fisher information / state-space models

Hyungju Jeon [1]

Isabel Rosário [1,2]

Il Memming Park [1]

[1] Champalimaud Foundation, Lisbon PT
[2] Instituto Superior Técnico, Lisbon, PT

Exploring Human Information Processing in Virtual Reality during Physical Inverse Design

We develop an immersive Virtual Reality (VR) paradigm to study the cognitive and behavioral mechanisms underlying complex human planning in high-dimensional design problems. In particular, participants must shape a deformable virtual clay to optimize fluid flow—an open-ended inverse design task inspired by Differentiable Learned Simulators [1]. By analyzing patterns of surface deformation and fine-grained hand and eye movement data, we investigate how internal planning strategies unfold over time and how they adapt across dynamic and static contexts. VR enables not only realistic interactions but also the introduction of non-physical affordances—such as symmetry activations—that support the deformation process and reveal deeper layers of planning.

motor planning / inverse design / virtual reality / hand-tracking / deformable objects

Daniel McNamee

Sara Monteiro

Champalimaud Foundation, Lisbon PT

References

[1] Allen, K. R., Lopez-Guevara, T., & Stachenfeld, K. L. (2022). Physical design using differentiable learned simulators.

Forecasting motor cortical population activity during delayed reaching using XFADS

We applied XFADS (eXponential FAmily Dynamical Systems (XFADS): Large-scale nonlinear Gaussian state-space modeling) [1], a nonlinear variational state-space model, to neural population recordings from motor and premotor cortex during a delayed reach task with barriers.

XFADS learns a latent dynamical system that enables causal inference and long-range forecasting—beyond the capabilities of standard latent trajectory methods. Given only the early preparatory activity, the GO cue timing, and the reaction time of the animal—prior to movement onset—XFADS was able to forecast neural activity up to ~800 ms into the future. It revealed that target and movement information was reliably encoded well before the movement initiation, even though the reaction time remained undecodable without explicit go cue input. Accuracy of the forecasted behavior was better than baselines, particularly on trials with large deviation from the condition-average. XFADS also generalized to unseen task conditions, including novel reach directions and error trials, despite being trained only on stereotypical trials. It decoded error behaviors more accurately than condition-averaged templates.

These results demonstrate that XFADS captures causal spatiotemporal structure in motor cortical activity and highlight the potential of generative modeling for understanding motor planning at the single-trial level.

atent dynamical systems / variational inference / single-trial forecasting / motor preparationn / causal inference

Mahmoud Elmakki [1]

Matthew Dowling [1,2]

Yuan Zhao [3]

Il Memming Park [1]

[1] Champalimaud Foundation, Lisbon PT
[2] New York University, New York, USA
[3] National Institute of Mental Health, Maryland, USA

References

[1] Dowling, M., Zhao, Y., & Park, I. M. (2024). eXponential FAmily Dynamical Systems (XFADS): Large-scale nonlinear Gaussian state-space modeling. March 3, 2024.

Contextual reconfiguration of mouse V1 representations driven by top-down input

Dimitra Maoutsa [1,2]

[1] School of Life Sciences, Technical University of Munich, Munich, DE
[2] Independent, DE

Neural circuits adapt rapidly to sensory stimuli through experience and ongoing learning in a non-trivial manner. Extensive experimental evidence suggests that they operate in distinct processing modes based on contextual information characterised by differential responses to the same stimuli. In early sensory cortices, different neuronal subtypes contribute differentially to gating sensory information. In early sensory cortices, different neuronal subtypes are thought to have specialized functional roles in controlling sensory information flow through the cortical hierarchy. Specifically, in mouse primary visual cortex (V1), populations of pyramidal and vasoactive-intestinal-peptide-expressing (VIP) cells show pronounced responses to novel and unexpected visual stimuli, while they rapidly adapt to elicit moderate responses to familiar or expected ones. However, the exact mechanisms that establish these distinct responses are not yet fully characterised. Local recurrent plasticity within V1 has been proposed as the origin of these context-dependent representations [1,2], but the full range of mechanisms remains unclear.

Here, we explore an alternative explanation for these differential representations based on feedback-driven interactions from adjacent visual areas to V1. We analysed electrophysiological recordings from primary visual cortex and adjacent visual areas of mice performing a sequential change detection task that establishes contextual stimulus expectations. Our analyses identified distinct inter-area communication patterns, based on stimulus condition, suggesting selective, context-driven information routing to V1. To test whether such feedback is sufficient to explain

References
[1] Ito, S., Piet, A., Bennett, C., Durand, S., Belski, H., Garrett, M., ... & Arkhipov, A. (2024). Coordinated changes in a cortical circuit sculpt effects of novelty on neural dynamics. Cell reports, 43(9).;
[2] Maoutsa, D., Gjorgjieva, J. (2024). Identifying plasticity mechanisms underlying experience-driven adaptation in cortical circuits. Bernstein Conference for Computational Neuroscience.

the experimentally observed response profiles, we built a multi-population firing-rate model receiving temporally resolved top-down input according to our empirical findings. We explore under which conditions feedback interactions suffice for explaining the experimentally observed context-dependent stimulus representations, and compare with a previous model of recurrently-driven changes that explains the context-dependent response profiles. Our work provides an alternative mechanistic explanation of how cortical circuits adjust their activity in response to contextual cues and potential stimulus expectations, and provides a mechanistic link between contextual feedback and adaptive sensory processing.

computational neuroscience / modelling / top-down feedback / inter-area communication

Altered States and Basic Taste Perception: Evaluating Detection and Recognition Thresholds Under Psychedelics

Zuzanna Dedyk [1,2,3]

Razvan Sandru [1]

Zachary Mainen [1]

[1] Champalimaud Foundation, Lisbon, PT
[2] École Pratique des Hautes Études, Paris, FR
[3] Institut Lyfe Research & Innovation Center, Ecully, FR

Altered states of consciousness, such as those induced by psychedelics, are increasingly studied for their impact on perception, cognition, and sensory integration. However, despite anecdotal reports [1] there is a significant gap in modern scientific literature regarding how these states influence basic taste thresholds, with prior research primarily focusing on visual and auditory enhancements [2,3] To address this, we designed a basic taste threshold study building on earlier work [4] and pilot findings. We aimed to evaluate basic sweet (sucrose) and sour (citric acid) taste thresholds in healthy individuals under the effects of psychedelics in a naturalistic setting, using a paper strips test [5]. We hypothesised that individuals experiencing psychedelic-induced altered states would exhibit heightened taste sensitivity, leading to lower detection and recognition thresholds for sucrose and citric acid.

Participants (n=120), including ones in altered states and sober controls, underwent an adapted taste strips test [5,6,7,] containing various concentrations of sucrose and citric acid [8]. Data collected for each strip included detection (yes/no), identification (correct/incorrect), and intensity ratings (0-10). Additionally, participants provided self-reported drug use information and completed a sensory questionnaire. Statistical analysis will employ mixed-effects models to estimate group differences in detection and recognition thresholds.

Data collection has been successfully completed, demonstrating the feasibility and suitability of the taste strip method in naturalistic environments.

Data analysis is currently in progress. The results have the potential to enhance our understanding of how altered states affect basic sensory processing, with possible broader implications for multisensory integration and flavor perception research.

psychedelics / Taste perception / Gustatory thresholds / Psychophysics / Sensory Neuroscience

References

[1] Kovacevich, A., Carhart-Harris, R. L., & Preller, K. H. (2023). Subjective reports of altered olfaction and taste after psilocybin administration. Frontiers in Pharmacology, 14, 1172895.;
[2] Abramson, H. A., Rolo, A., & Gorman, F. J. (1955). Hallucinogens and visual perception. Journal of Nervous & Mental Disease, 122(6), 560-562.;
[3] Gouzoulis-Mayfrank, E., Thelen, B., Maier, S., Heekeren, K., Kovar, K. A., Sass, H., & Spitzer, M. (2002). Effects of the hallucinogen psilocybin on covert orienting of visual attention in humans. Neuropsychobiology, 45(4), 205–212.;
[4] Fischner, G., & Kaebling, R. (1966). Taste acuity under psychedelics. Psychonomic Science, 4(2), 79-80.;
[5] Landis, B. N., Konrad, M., & Braun, H. (2009). The “taste strips” test. Laryngoscope, 119(9), 1836-1840.;
[6] Wolf, A., Landis, B. N., & Luther, C. A. (2016). A comparative study of taste strips and whole mouth taste testing. European Archives of Oto-Rhino-Laryngology, 273(7), 1779-1784.;
[7] Ribeiro, A., Ferreira, A., & Carvalho, R. (2022). Adaptation and validation of the “taste strips” test for a Portuguese population. Journal of Clinical and Experimental Neuropsychology, 44(4), 312-320.;
[8] Yamauchi, Y., Watanabe, T., & Minami, T. (2002). Development of taste strips for clinical use. Acta Oto-Laryngologica, 122(5), 527-531.;

The Rest of the Fish: Light-Sheet Microscopy Imputation with Transformers

Whole-brain light-sheet microscopy in larval zebrafish enables volumetric imaging of neural activity across the entire brain at relatively low frame rates, or of partial brain regions at higher temporal resolutions. However, the fast dynamics of modern calcium indicators, which operate on millisecond timescales, are often undersampled—leading to significant loss of temporal detail. Slower indicators, meanwhile, obscure brief but important neural events. To address this challenge, we propose a transformer-based model that leverages the intrinsic correlations in neural activity to impute missing data—both temporally and spatially. Our model interpolates missing timepoints between sparsely acquired full-brain volumes and in-paints omitted planes in high-speed partial volumes, effectively reconstructing a high-resolution recording across space and time.
We employ a BERT-style masked-token transformer, trained to recover masked frames and planes. This architecture offers several key advantages:
- Temporal flexibility via learned positional encodings, enabling irregular sampling and the omission of nuisance frames (e.g., flyback planes);
- Insights can be derived through attention maps, which highlight influential regions in the imputation process;
- Bidirectional inference, allowing both forward and backward context;
- Scalable inference, as the model accommodates variable token lengths to adapt to hardware constraints.
The imputed data can then be used to recover, for example, fluorescence of identified ROIs. We hope to extend the model to include a deep Gaussian

Dean Rance

Aaron Ostrovsky

Lucas Martins

João Marques

Pedro Tomás Silva

Carolina Gonçalves

↓

Champalimaud Foundation, Lisbon, PT

process layer, allowing the user to have full posterior estimates of activity alongside the ability to regularise the smoothness of the fit [2]. Ultimately, this framework could enable experimenters to achieve high-framerate, whole-brain imaging with only modest adjustments to current sampling protocols.

zebrafish / light-sheet microscopy / imputation / transformer

Tiago Costa

Michael Orger

Gonzalo G. de Polavieja

References

[1] Zhang, Y., Rózsa, M., Liang, Y., Bushey, D., Wei, Z., Zheng, J., Reep, D., Broussard, G. J., Tsang, A., Tsegaye, G., Narayan, S., Obara, C. J., Lim, J.-X., Patel, R., Zhang, R., Ahrens, M. B., Turner, G. C., Wang, S. S.-H., Korff, W. L., … Looger, L. L. (2023). Fast and sensitive GCaMP calcium indicators for imaging neural populations. Nature, 615(7954), Article 7954.;
[2] Kanagawa, M., Hennig, P., Sejdinovic, D., & Sriperumbudur, B. K. (2018). Gaussian Processes and Kernel Methods: A Review on Connections and Equivalences (No. arXiv:1807.02582). arXiv.

Exploring trade-offs between sequential and parallel computation in recurrent neural networks

Parallel and serial processing modes have been studied widely in humans and other animals [1]. While serial processing is thought to enable few-shot generalization and rule-based reasoning [1,2], extensive practice of specific skills yields more parallel, habit-like computation, increasing speed by reducing the number of reasoning steps [3], but likely at the cost of increased representational resources [4]. Here we explore the trade-off between processing time and representational resources in a recurrent neural network model. Starting from a simple feedforward auto-encoder, we add a recurrent decoder layer and a feedback connection from this layer to the encoder. When the bottleneck is narrower than the intrinsic dimensionality of the data, this recurrence can be used to iteratively "gate" extra information, improving reconstruction relative to the feedforward architecture at the cost of additional time steps. We train the model end-to-end to reconstruct static noise vectors with a mean-squared error loss and backpropagation through time. The bottleneck is either continuous or discrete; in the latter case we use the Gumbel-softmax reparametrization trick. In both cases the network finds near-optimal latent codes, e.g. compressing the inputs into a sequential binary code to traverse a binary bottleneck. More generally, systematically varying the bottleneck width forces either sequential, mixed or parallel processing. Furthermore, preliminary continual-learning experiments suggest that the code reorganizes when the input distribution shifts: the network continually adopts the encoding that maximizes information transmission rate, mirroring adaptive habit formation [3]. Future work will scale the framework to richer datasets, test whether the learned dynamics support compositional generalisation, and derive quantitative predictions for neural data.

Marcel Graetz

Caroline Haimerl

Daniel McNamee

Alfonso Renart

Champalimaud Foundation, Lisbon, PT

References

[1] Evans, J. S. B. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. Annu. Rev. Psychol., 59(1), 255-278.
[2] Kahneman, D. (2011). Thinking, fast and slow. macmillan.
[3] Du, Y., Krakauer, J. W., & Haith, A. M. (2022). The relationship between habits and motor skills in humans. Trends in cognitive sciences, 26(5), 371-387.
[4] Musslick, S., & Cohen, J. D. (2021). Rationalizing constraints on the capacity for cognitive control. Trends in cognitive sciences, 25(9), 757-775.

Startle Responses and Behavioural State Switching

Parallel and serial processing modes have been studied widely in humans and other animals [1]. While serial processing is thought to enable few-shot generalization and rule-based reasoning [1,2], extensive practice of specific skills yields more parallel, habit-like computation, increasing speed by reducing the number of reasoning steps [3], but likely at the cost of increased representational resources [4]. Here we explore the trade-off between processing time and representational resources in a recurrent neural network model. Starting from a simple feedforward auto-encoder, we add a recurrent decoder layer and a feedback connection from this layer to the encoder. When the bottleneck is narrower than the intrinsic dimensionality of the data, this recurrence can be used to iteratively "gate" extra information, improving reconstruction relative to the feedforward architecture at the cost of additional time steps. We train the model end-to-end to reconstruct static noise vectors with a mean-squared error loss and backpropagation through time. The bottleneck is either continuous or discrete; in the latter case we use the Gumbel-softmax reparametrization trick. In both cases the network finds near-optimal latent codes, e.g. compressing the inputs into a sequential binary code to traverse a binary bottleneck. More generally, systematically varying the bottleneck width forces either sequential, mixed or parallel processing. Furthermore, preliminary continual-learning experiments suggest that the code reorganizes when the input distribution shifts: the network continually adopts the encoding that maximizes information transmission rate, mirroring adaptive habit formation [3]. Future work will scale the framework to richer datasets, test whether the learned dynamics support compositional generalisation, and derive quantitative predictions for neural data.

Charlie Rosher [1]

Roian Egnor [1]

Alice Robie [2]

Violetta LaFranca [1]

Matheus Farias [1]

Kristin Branson [1]

Marta Moita [1]

[1] Champalimaud Foundation, Lisbon, PT
[2] Janelia Research Campus, Virginia, USA

Architecture as Neuro-Cybernetic Feedback: Designing Spaces That Regulate the Human Nervous System

Arianna Bavuso

AB+AC Architects, Lisbon, PT

As environments become increasingly dense, artificial, and overstimulating, the built environment emerges as a key player in shaping cognitive and physiological health. This submission proposes a new design framework—architecture as a neuro-cybernetic feedback system—where space is intentionally designed to interact with the human nervous system. My work at AB+AC Architects focuses on immersive environments that act as therapeutic interfaces: regulating stress, supporting attention, and enhancing emotional states through sensory-driven design strategies. while this approach is not yet supported by empirical neuroscience data, it is grounded in interdisciplinary theory and practice, informed by principles from neuroarchitecture, psychology of space, ecopsychology, and integrative medicine. Projects like "The Luxury of Less" (dubai design week 2024) and "Back to the Future" (milan design week 2023) offer real-life case studies of multi-sensory spatial systems. These pavilions combine materiality, light, acoustics, and scale to influence neurophysiological response and invite calm, introspection, and affective presence. Although current evidence is anecdotal, the aim of this presentation is to frame future research collaborations with neuroscience labs to develop data-informed spatial prototypes using tools such as EEG, biosensing, and environmental biometrics. This proposal invites discussion on the future of built environments not only as structures but as dynamic regulatory systems—interfaces that harmonize with the nervous system and influence behavior, healthspan, and well-being. It is a call for interdisciplinary partnership across architecture, neuroscience, and systems thinking to prototype responsive spaces for the next era of human-centered design.

neuroarchitecture / psychology of spaces / environmental homeostasis / holistic well-being / spatial cognition

Meta-learning three-factor learning rules for structured credit assignment in recurrent neural networks

Dimitra Maoutsa [1,2]

[1] Technical University of Berlin, Berlin, DE
[2] Independent, DE

Biological neural networks learn to execute complex tasks using feedback that is both sparse and delayed, yet synaptic plasticity operates locally both in space and time. How can delayed, global reinforcement signals drive the precise synaptic updates required for successful task execution? Previous work tackling this problem has considered either continuous error signals that guide learning throughout the trial period, or hand-crafted local updates. Yet, the space of plasticity rules capable of supporting learning from delayed reinforcement, and how these rules relate to task structure and resulting representations remains largely unexplored. Here, we introduce a meta-learning framework for discovering biologically plausible three-factor learning rules for structured credit assignment in recurrent neural networks.
Each rule couples a reward prediction error signal to synapse-specific eligibility traces that accumulate nonlinear interactions of pre- and post-synaptic activity over task execution. We optimise the trace coefficients in a meta-optimisation loop that employs a REINFORCE-like gradient approximation for the plasticity parameter updates, avoiding memory intensive backpropagation-through-timeover the entire learning period, while maintaining accurate gradient estimates. We systematically explore how resulting representations depend on the employed form of discovered plasticity rules, and how these rules relate to task parameters. Our work offers insights into biologically grounded mechanisms for learning in recurrent circuits, showing how local plasticity interacts with reinforcement signals to give rise to behaviourally relevant representations.

neuroAI / biologically plausible learning / plasticity rule inference / RNNs

Multitasking Recurrent Networks Utilize Compositional Strategies for Control of Movement

The brain and body comprise a complex control system that can flexibly perform a diverse range of movements. Despite the high-dimensionality of the musculoskeletal system, both humans and other species are able to quickly adapt their existing repertoire to novel settings. A strategy likely employed by the brain to accomplish such a feat is known as compositionality, or the ability to combine learned computations to perform novel, yet similar tasks. Previous works have demonstrated that recurrent neural networks (RNNs) utilize a compositional structure to perform diverse cognitive tasks. However, the attractor-based computations required for cognition are largely distinct from those required for the generation of movement, and it is unclear whether compositional structure extends to RNNs producing complex movements. To address this question, we train a multitasking RNN in feedback with a musculoskeletal model to perform ten distinct movements at various speeds and directions. We find a shared manifold implemented across all tasks in the network space, while the manifolds across different epochs of the tasks lie in orthogonal subspaces. The network utilizes a compositional representation (new tasks can be built from learned sub-parts of others) for movements with similar kinematic and rotational properties, while also acquiring the ability to stitch together distinct kinematic motifs in a sequential fashion. Additionally, we show dynamical similarities across tasks residing in the same subspaces in the form of similar fixed point structures and trajectory disparities. Lastly, our model is able to reuse learned motifs in novel settings by simply training the input weight for a scalar rule input. Our framework sheds light on how the brain might flexibly perform a diverse range of movements through the reuse of pre-existing computations on a shared manifold.

RNN / motor / compositionality / multi-tasking / feedback

John Lazzari

Shreya Saxena

Yale University, New Haven, USA

Multi-Agent Reinforcement Learning to understand Cooperation in Animals

Social interactions are an important yet complex behavior for animals to learn. When learning behavioral policies for social settings, they must take into account the dynamics of the environment and the actions of other agents. The effects of other agents' actions can be difficult to separate from environmental dynamics and other agents' policies might be non-stationary. In cooperative social experiments, these challenges are compounded as animals receive reward only when they coordinate their actions with others. In past cooperative experiments, conditions such as obscured visual access and familiarity were shown to modulate cooperative success [1]. To better understand the strategies animals use to solve cooperative tasks and how experimental conditions affect these strategies, we used multi-agent deep reinforcement learning to model animal behavior during cooperative social tasks [2]. Models were trained in a goal-driven manner in a virtual environment that mimicked the experimental paradigm of a rat cooperative social behavior task. In this task, both rats had to press a lever within a certain time window to receive a reward.

As visual occlusion experimentally modulated cooperative success, the restriction of partial observability was imposed on the virtual environment. In the partially observable setting, agents would receive information about other agents only when a gaze action was taken. To model familiarity in our multi-agent setting, we included a 'social model': an internal predictive model that learns the policy of the other agent. When training agents independently [3] and in

Amelia Johnson

S William Li

Henry W Kietzman

David Backer Peral

R. Allie Cauchon

Nina Pelias

Jane R. Taylor

Shreya Saxena

Yale University, New Haven, USA

partially observable conditions, an internal predictive model of the other agent was necessary to achieve high reward on the cooperative task. Modeling cooperative behavior with deep reinforcement learning also provides a hypothesis for neural computations underlying social behavior.

reinforcement learning / social behavior/ deep learning / social cooperation

References

[1] Meisner Olivia C, Shi Weikang, Fagan Nicholas A, Greenwood Joel, Jadi Monika P, Nandy Anirvan S, Chang Steve WC (2024) Development of a Marmoset Apparatus for Automated Pulling (MarmoAAP) to Study Cooperative Behaviors eLife 13:RP97088

[2] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. "Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments". In: Advances in Neural Information Processing Systems. Vol. 30. 2017, pp. 6382–6393.

[3] Richard S. Sutton and Andrew G. Barto. "Reinforcement Learning: An Introduction". In: IEEE Transactions on Neural Networks. Adaptive Computation and Machine Learning 9.5 (1998).

Towards a Disease Model of Cerebellar Ataxia Using Supervised Learning

The cerebellum is believed to use a forward supervised learning approach to guide motor control and learning [1]. While the circuitry of the cerebellum has been well researched, there remains questions of how the cerebellum performs a wide range of computations, and how changes in these computations map to cognitive and motor dysfunction. We model the cells of the cerebellum using a recurrent neural network (RNN) [2], to then drive a cortical RNN [3] in control of a simple 2-joint arm model. We start with a fully functioning model that is able to achieve task specific movement while receiving time dependent input regarding the state of the arm. We aim to build a disease model of ataxia, a motor disorder marked by a lack of control over voluntary behavior, such as unwanted trembling of the limbs or uncoordinated movement. It is associated with unusual cerebellar activity, specifically of the purkinje cells (PCs), the sole output cells of the cerebellar cortex. Proposed sources of ataxia are changes in the intrinsic firing rate of PCs [4], and damage at the PC and climbing fiber (CF) synapse [5] CF inputs signal error and induce complex spiking of the PCs, altering cerebellar output and influencing the resulting motor behavior. Towards modeling ataxia, we implement changes in intrinsic firing of the PCs in the cerebellar RNN by altering the bias term of the PC units while the model is attempting to perform previously learned tasks. We represent damage at the CF-PC synapse by modifying the strength of connections between PC and CF units, interrupting the flow of hypothesized error signals to the PCs, as the cerebellar network works to learn a task. This implementation is a step towards developing a model of cerebellar activity capable of producing ataxic patterns of movement in models of the limbs, guiding therapeutic efforts for cerebellar ataxia.

ataxia / cerebellum / motor disorder / supervised learning

Latavia Thompson

Heather Snell

Yale University, New Haven, USA

References

1 Raymond, JL., and Medina JF (2018). Computational Principles of Supervised Learning in the Cerebellum. Annu Rev Neurosci., 41:233-253. doi: 10.1146/annurev-neuro-080317-061948.
2 Lazzari J., mRNNTorch, https://github.com/JohnLazzari/mRNNTorch
3 Pemberton, J., Chadderton, P. & Costa, R.P. (2024). Cerebellar-driven cortical dynamics can enable task acquisition, switching and consolidation. Nat Commun, 15, 10913. https://doi.org/10.1038/s41467-024-55315-6
4 Cook, A., Fields, E., & Watt, A. (2020). Losing the Beat: Contribution of Purkinje Cell Firing Dysfunction to Disease, and Its Reversal. Neuroscience. 462. 10.1016/j.neuroscience.2020.06.008.
5 Smeets, C.J.L.M., Verbeek, D.S. (2016). Climbing fibers in spinocerebellar ataxia: A mechanism for the loss of motor control. Neurobiology of Disease, 88, 96-106. https://doi.org/10.1016/j.nbd.2016.01.009.

Adaptive behavioral strategies in Drosophila courtship pursuit

The cerebellum is believed to use a forward supervised learning approach to guide motor control and learning [1]. While the circuitry of the cerebellum has been well researched, there remains questions of how the cerebellum performs a wide range of computations, and how changes in these computations map to cognitive and motor dysfunction. We model the cells of the cerebellum using a recurrent neural network (RNN) [2], to then drive a cortical RNN [3] in control of a simple 2-joint arm model. We start with a fully functioning model that is able to achieve task specific movement while receiving time dependent input regarding the state of the arm. We aim to build a disease model of ataxia, a motor disorder marked by a lack of control over voluntary behavior, such as unwanted trembling of the limbs or uncoordinated movement. It is associated with unusual cerebellar activity, specifically of the purkinje cells (PCs), the sole output cells of the cerebellar cortex. Proposed sources of ataxia are changes in the intrinsic firing rate of PCs [4], and damage at the PC and climbing fiber (CF) synapse [5] CF inputs signal error and induce complex spiking of the PCs, altering cerebellar output and influencing the resulting motor behavior. Towards modeling ataxia, we implement changes in intrinsic firing of the PCs in the cerebellar RNN by altering the bias term of the PC units while the model is attempting to perform previously learned tasks. We represent damage at the CF-PC synapse by modifying the strength of connections between PC and CF units, interrupting the flow of hypothesized error signals to the PCs, as the cerebellar network works to learn a task. This implementation is a step towards developing a model of cerebellar activity capable of producing ataxic patterns of movement in models of the limbs, guiding therapeutic efforts for cerebellar ataxia.

ataxia / cerebellum / motor disorder / supervised learning

Miguel Paço [1]

Nuno Rito [1]

Roman Goulard [2]

Barbara Webb [2]

Eugenia Chiappe [1]

[1] **Champalimaud Foundation, Lisbon, PT**
[2] **University of Edinburgh, UK**

References

[1] Cook, R. (1979). The courtship tracking of Drosophila melanogaster. Biological Cybernetics, 34(2), 91–106. https://doi.org/10.1007/BF00365473

[2] Sten, T., Li, R., Otopalik, A., & Ruta, V. (2021). Sexual arousal gates visual processing during Drosophila courtship. Nature, 595(7868), 549–553. https://doi.org/10.1038/s41586-021-03714-w

Supervised Mechanisms for Locomotor Learning Driven by Climbing Fibers

Motor learning is essential to move in a continuously changing environment, but the underlying neural mechanisms are still poorly understood. This is critical for locomotion, a fundamental but complex behavior, which requires simultaneous control of multiple limbs across the body. The cerebellum plays a key role in motor learning [1], generating anticipatory corrective movements through supervised error-based mechanisms, driven by instructive signals in response to perturbations - in accordance with the canonical theory [2,3]. For simple tasks, climbing fibers originating in the inferior olive and projecting to the cerebellum drive supervised learning [4-6].

However, for complex whole-body behaviors like locomotion, whether and how climbing fibers encode instructive signals is still unknown. We recently demonstrated [7] that mice running on a split-belt treadmill - where the speeds under each side of the body are controlled independently - exhibit a form of motor learning called locomotor adaptation, which is cerebellum-dependent and reflects learned changes in interlimb symmetry. Here, we applied closed-loop optogenetics controlled by real-time tracking of limb kinematics to test the hypothesis that the timing of climbing fiber signals encodes instructive signals for supervised learning during locomotion.

We find that optogenetic perturbations of climbing fibers, precisely locked to specific phases of the locomotor cycle, are sufficient to drive bidirectional, learned changes in interlimb symmetry, depending on the phase of the locomotor cycle in which they occur.

Alice Geminiani

Ana Gonçalves

Hugo G. Marques

Tatiana Silva

Francesco Costantino

Megan R. Carey

Champalimaud Foundation, Lisbon, PT

We are now combining these behavioral results with electrophysiological measurements, temporal basis functions and supervised learning rules to model how this remarkable precision of climbing fiber timing is used by the cerebellum to effectively and robustly drive learning during locomotion.

locomotor adaptation / supervised learning / cerebellum / optogenetics / electrophysiology

References

[1] Ito, M. (2011). The Cerebellum: Brain for an implicit self.
[2] Marr, D. (1969). A theory of cerebellar cortex. J Physiol, 202, 437–470.
[3] Albus, J.S. (1971). A theory of cerebellar function. Math Biosci, 10, 25–61.
[4] Silva, N.T., Ramírez-Buriticá, J., Pritchett., D.L., Carey, M.R. (2024). Climbing fibers provide essential instructive signals for associative learning. Nat Neurosci, 27(5), 940-951. doi: 10.1038/s41593-024-01594-7.
[5] Ohmae, S., Medina, J. (2015). Climbing fibers encode a temporal-difference prediction error during cerebellar learning in mice. Nat Neurosci, 18, 1798–1803. doi: https://doi.org/10.1038/nn.4167.
[6] Medina, J. F. & Lisberger, S. G. (2008). Links from complex spikes to local plasticity and motor learning in the cerebellum of awake-behaving monkeys. Nat Neurosci, 11, 1185–1192.
[7] Darmohray, D.M., Jacobs, J.R., Marques, H.G., Carey, M.R. (2019). Spatial and temporal locomotor learning in mouse cerebellum. Neuron, 102(1). doi: 10.1016/j.neuron.2019.01.038.

A unifying theory for hippocampal remapping through the lens of contextual inference

Hippocampal remapping describes the phenomenon of substantial reshuffling of the spatial preferences of hippocampal place cells in response to environmental changes. Disparate neural representations in different environments could reduce interference between competing memories, hence functionally supporting robust context-dependent spatial cognition. Experimental data on existing remapping studies exhibit a diverse range of physiological properties, sometimes with seemingly conflicting results. In particular, subtle changes in experimental conditions lead to qualitatively different results. Despite previous attempts, no existing models provide a coherent normative explanation for these mixed remapping responses reported in the literature. Here, we interpret remapping through the lens of probabilistic inference over ongoing task context, instead of being based solely on environmental changes. This subtle but important difference allows us to probe how remapping profiles change as a function of a multitude of task-dependent features. Through performing Bayesian inference over the latent contextual variable, we derive precise mathematical identities showing that remapping exhibits explicit dependence on the degree of environmental difference during pre-training and the amount of pre-training experience. Moreover, we numerically show that remapping is implicitly dependent on the order of presentation during probing trials. We model three influential experimental studies that lead to different remapping profiles under similar experimental conditions [1,2], and that previous theoretical accounts were not able to reconcile. We show that our theory captures the qualitative features of remapping in all three studies. Our work suggests that the computations underlying spatial remapping generalise beyond spatial context, and support a more general role of the hippocampal formation in computations centering around contextual inference.

Changmin Yu [1]

Máté Lengyel [2]

[1] University of Cambridge, Cambridge, UK
[2] Central European University, Budapest, Hungary

References

[1] Wills, T.J., Lever, C., Cacucci, F., Burgess, N. and O'Keefe, J., 2005. Attractor dynamics in the hippocampal representation of the local environment. Science, 308(5723), pp.873-876.
[2] Leutgeb, J.K., Leutgeb, S., Treves, A., Meyer, R., Barnes, C.A., McNaughton, B.L., Moser, M.B. and Moser, E.I., 2005. Progressive transformation of hippocampal neuronal representations in "morphed" environments. Neuron, 48(2), pp.345-358.

Deep reinforcement learning to mimic neuromechanical control: realistic locomotion in an embodied fly

Elliott T. T. Abe [1]

Charles Zhang [2]

Raveena Chhibber [1]

Grant M. Chou [1]

Dang P. Q. Truong [1]

Bence P. Ölveczky [2]

[1] University of Washington, Seattle, USA
[2] Harvard University, Boston, USA
[3] Salk Institute for Biological Studies, San Diego, USA

Most locomoting animals can maintain motor control after limb injury or amputation, where coordinating movements of a drastically altered body requires dynamic interactions between the nervous system, the biomechanics of the body, and the physical environment. Until recently, it has been challenging to link neural and biomechanical models to investigate adaptive motor recovery because it requires coupling neural control in closed loop with the environment. Further, while previous approaches have been able to simulate walking in biomechanically realistic bodies, the joint kinematics and ground reaction forces have been unrealistic or unvalidated. In this project, we developed and trained an agent with deep reinforcement learning (DRL) to imitate real Drosophila walking and flying using a biomechanically realistic fly body model in the physics simulator MuJoCo [1]. We implemented a variational encoder-decoder architecture [2], to build an interpretable latent space for motor control. For training data, we used high-fidelity 3D tracking data to imitate behaviors. We show that our model closely resembles real fly locomotion while reproducing accurate movement dynamics (i.e. ground reaction forces in walking). Historically, measuring forces produced by such small animals has been impossible. We validate our walking MuJoCo model with the first-of-their-kind measurements of ground reaction forces in freely walking fruit flies, demonstrating that simulated ground contact forces during walking closely match experimental measurements. Additionally, we implemented a biologically realistic aerodynamics model [3] to simulate the forces experienced by flying insects. To train our agent on higher-level

tasks while maintaining realistic movements, we built a high-level controller by freezing the decoder, which translates the latent space into realistic movement, and trained a new encoder with a task driven computational goal. Using this pipeline, we can simulate in real time the full neuromechanical control in closed loop with the environment during locomotion to test and generate hypotheses.More broadly, this work is a key step in using embodied agents to understand the neural mechanisms controlling robust movement with a dynamically changing body and environment.

embodied agents / neuromechanical control / reinforcement learning

Nathan Sniadecki [1]

John C. Tuthill [1]

Talmo D. Pereira [3]

Bing W. Brunton [1]

References

[1] Todorov, E., Erez, T., & Tassa, Y. (2012). MuJoCo: A physics engine for model-based control. 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, 5026–5033. https://doi.org/10.1109/IROS.2012.6386109
[2] Aldarondo, D., Merel, J., Marshall, J. D., Hasenclever, L., Klibaite, U., Gellis, A., Tassa, Y., Wayne, G., Botvinick, M., & Ölveczky, B. P. (2024). A virtual rodent predicts the structure of neural activity across behaviors. Nature. https://doi.org/10.1038/s41586-024-07633-4
[3] Sane, S. P., & Dickinson, M. H. (2001). The control of flight force by a flapping wing: Lift and drag production. Journal of Experimental Biology, 204(15), 2607–2626. https://doi.org/10.1242/jeb.204.15.2607

Evolution Imposes an Inductive Bias that Alters and Accelerates Learning Dynamics

Benjamin Midler

Alejandro Pan Vazquez

Princeton Neuroscience Institute, Princeton, USA

The learning dynamics of biological brains and artificial neural networks are of interest to both neuroscience and machine learning. A key difference between them is that neural networks are often trained from a randomly initialized state whereas each brain is the product of generations of evolutionary optimization, yielding innate structures that enable few-shot learning and inbuilt reflexes. Artificial neural networks, by contrast, require non-ethological quantities of training data to attain comparable performance. To investigate the effect of evolutionary optimization on the learning dynamics of neural networks, we combined algorithms simulating natural selection and online learning to form a new method for evolutionarily conditioning artificial neural networks, and applied it to both supervised and reinforcement learning problems. Evolutionarily conditioned networks learned both task types from a few dozen training examples—an order of magnitude less than randomly initialized baselines. Furthermore, the learning dynamics of evolutionarily conditioned networks differed significantly from those of conventionally trained networks, and did not respond to the same mathematical regularities in training data. These results indicate that evolutionary optimization of the population imposes an inductive bias on the learning dynamics of the individual, effectively tuning networks' initial configuration to the learning environment it will occupy.

evolution / learning / neural networks / inductive bias

A Transient Neural Code for Feedback-Driven Motor Corrections During Reaching

33

The learning dynamics of biological brains and artificial neural networks are of interest to both neuroscience and machine learning. A key difference between them is that neural networks are often trained from a randomly initialized state whereas each brain is the product of generations of evolutionary optimization, yielding innate structures that enable few-shot learning and inbuilt reflexes. Artificial neural networks, by contrast, require non-ethological quantities of training data to attain comparable performance. To investigate the effect of evolutionary optimization on the learning dynamics of neural networks, we combined algorithms simulating natural selection and online learning to form a new method for evolutionarily conditioning artificial neural networks, and applied it to both supervised and reinforcement learning problems. Evolutionarily conditioned networks learned both task types from a few dozen training examples—an order of magnitude less than randomly initialized baselines. Furthermore, the learning dynamics of evolutionarily conditioned networks differed significantly from those of conventionally trained networks, and did not respond to the same mathematical regularities in training data. These results indicate that evolutionary optimization of the population imposes an inductive bias on the learning dynamics of the individual, effectively tuning networks' initial configuration to the learning environment it will occupy.

evolution / learning / neural networks / inductive bias

Nina Kudryashova [1]

Cole Hurwitz [2]

Matthew G. Perich [3,4]

Matthias H. Hennig [1]

[1] School of Informatics, University of Edinburgh; Informatics Forum, 10 Crichton St, Newington, Edinburgh EH8 9AB, United Kingdom
[2] Zuckerman Institute, Columbia University; 3227 Broadway, New York, NY 10027, United States
[3] Département de neurosciences, Faculté de médecine, Université de Montréal; Pavillon Roger-Gaudry, 2900 Edouard Montpetit Blvd, Montreal, Quebec H3T 1J4, Canada
[4] Mila, Quebec Artificial Intelligence Institute; 6666 Rue Saint-Urbain, Montréal, QC H2S 3H1, Canada

A Bio-Inspired Algorithm Enables Scalable Training of Spiking Neural Networks Using Feedback Control Across Layers

Jonathan Haag [1]

Christian Metzner [1]

Chiara De Luca [1,2]

Dmitrii Zendrikov [1]

Giacomo Indiveri [1]

Benjamin Grewe [1,3]

Matteo Saponati [1]

[1] Institute of Neuroinformatics, University of Zurich and ETH Zurich, Zürich, CH
[2] Digital Society Initiative, University of Zurich, Zürich, CH
[3] ETH AI Center, ETH Zurich, Zürich, CH

Unlike artificial networks, biological neural networks communicate via spikes and learn with local plasticity, allowing scalable credit assignment through feedforward and feedback signals [1,2].

However, existing bio-plausible learning rules can only train shallow Spiking Neural Networks (SNNs), which reduces their ability to accurately model how biological systems learn. This also limits their effectiveness for on-device training using neuromorphic hardware, which mimics key brain features like spike-based communication and limited energy availability [3,4]. To address this, we introduce a learning algorithm for SNNs that leverages feedback control to compute weight updates locally in space and time [5,6].

Our framework uses spiking control neurons to guide network activity towards a desired target by sending top-down feedback signals into the apical compartment of every neuron and updating weights to minimize this feedback signal over training. This enables supervised learning across multiple layers, while being fully compatible with the constraints of biological and neuromorphic systems. We evaluate our algorithm by training SNNs on a mixed-signal neuromorphic device, the DYNAP-SE [7], across various classification tasks. We achieve the expected test-time accuracy, with results that are consistent with simulations both quantitatively and qualitatively.

Our results thus demonstrate bio-inspired, real-time training of multi-layer SNNs on specialized neuromorphic hardware, paving the way for

next generation AI. Strikingly, the competitive performance of our algorithm shows that learning by minimizing feedback control can scale effectively, even under constraints similar to those found in the brain, such as substrate variability, spike-based communication, and imprecise synaptic weights.

feedback control / neuromorphic computing / bio-inspired learning / spiking neural networks / on-device training

References

[1] Richards, B. A., Lillicrap, T. P., Dendritic solutions to the credit assignment problem. Current Opinion in Neurobiology, 54, 28–36 (2019);
[2] Poirazi, P., Brannon, T., & Mel, B. W., Pyramidal neuron as two-layer neural network. Neuron, 37(6), 989-999. (2003);
[3] Rubino, A., Cartiglia, M., Payvand, M., & Indiveri, G., Neuromorphic analog circuits for robust on-chip always-on learning in spiking networks. 2023 IEEE 5th International Conference on Artificial Intelligence Circuits and Systems (AICAS), 1–5 (2023);
[4] Chicca, E., Stefanini, F., Bartolozzi, C., & Indiveri, G., Neuromorphic electronic circuits for building autonomous cognitive systems. Proceedings of the IEEE, 102(9), 1367–1388 (2014);
[5] Vilimelis Aceituno, P., Tristany Farinha, M., Loidl, R., & Grewe, B. F., Learning cortical hierarchies with temporal Hebbian updates. Frontiers in Computational Neuroscience, 17, 1136010 (2023);
[6] Saponati, M., De Luca, C., Indiveri, G. & Grewe, B. F., A feedback control optimizer for online and hardware-aware training of Spiking Neural Networks. 2025 Neuro-Inspired Computational Elements Conference (NICE), (2025);
[7] Moradi, S., Qiao, N., Stefanini, F., & Indiveri, G., A Scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic processors. IEEE Transactions on Biomedical Circuits and Systems, 12(1), 106–122 (2018)

From Principles to Production: How Visual Design Shapes Understanding

Ágota Végső [1]

Paulo Nuno Vicente [1]

Ana Figueiras [2]

[1] ICNOVA, NOVA FCSH University of Lisbon, Lisbon, PT
[2] CICANT, Lusófona University, Lisbon, PT

Animation is one of the most popular audio-visual educational tools, with a growing audience, thanks to accessible technological innovations and the changing digital media consumption habits of young audiences. In response to the continued rapid development of the digital industry and the increasing reliance of young generations on online information and artificial intelligence, this article focuses on enhancing learning from Science Animation through academic and industry collaboration and bridging the culture of art and science by translating the principles for first-time science animation collaborators. To address the communication gap between scientists and visual storytellers, a literature search was conducted on science animation, encompassing overviews, reviews, and meta-analyses, which established guiding steps for creating accurate and compelling science visualizations. Drawing on Communication Models, Communication Design principles, Cognitive Theory of Multimedia Learning definitions and Social Semiotics, four initial key stages were identified: (1) Project Identification, the extraction and framing of the scientific core to be visualized, (2) Project Description, guiding the visual articulation of scientific content, (3) Animation Design, wherein animation history and animation-related multimedia learning principles inform the fundamental design rules of the concept art of animation, and (4) Animation Analysis, where systematic characterization complimented with social semiotic approach is compared to industry based animation analysis. Stages represent different complexities of talking, discussing, understanding, analyzing science animation to support effective interdisciplinary collaboration.

animation / communication design / interdisciplinary collaboration / multimedia learning / science visualisation

Dynamic Scale Equivariance in Retinal Neural Codes Supports Object Tracking

S. Azeglio [1,2]
R. Baroux [1]
U. Ferrari [1]
O. Marre [1]

[1] Institut de la Vision, Sorbonne Université, INSERM, CNRS, 17 Rue Moreau, 75012, Paris, FR
[2] Laboratoire des Systèmes Perceptifs, Ecole Normale Supérieure, 24 Rue Lhomond, 75005, Paris, FR

Various animals, including mice, can approach and catch prey in different environments. Their visual system must track moving targets across visual space, and be robust to changes in position and scale. Retinal ganglion cell (RGC) mosaics provide translation equivariance: if an object is represented by a feature map of ganglion cells, translating this object will trigger a similar translation in the feature map, making the code robust to changes in positions. However, how encoding is robust to changes of scale is unclear, and we don't know if there is a similar equivariance to scale in the retina. Here, we demonstrate that specific RGC types achieve functional scale equivariance, enabling a robust representation of a target across scales during ethological tasks like prey capture. We developed a novel paradigm combining multi-scale visual stimuli simulating the visual experience of a mouse during prey hunting with high-density multielectrode array recordings from ex vivo mouse retina responding to these stimuli. We adapted Scale-Equivariant Steerable Networks (SESN) from geometric deep learning [1], together with a Scale-Factorized Readout layer to model and predict RGC responses. Our SESN model significantly outperformed standard CNNs in predicting RGC responses with greatest gains for transient OFF alpha cells, which are implicated in looming detection and object tracking [2,3]. Systematic ablation experiments revealed these cells implement a key trade-off: strong center-surround antagonism impairs scale equivariance, so these cells have a weak surround to ensure scale equivariance, balancing contrast sensitivity with scale-tracking. These findings highlight equivariance as a valuable framework for understanding neural computation, linking the symmetries of natural stimuli to sensory system responses and complementing recent developments in geometric deep learning.

equivariance / retina, neural coding / object tracking / geometric deep learning

References

[1] Sosnovik I, et al. Scale-equivariant steerable networks. ICLR. 2020.
[2] Wang F, et al. OFF-transient alpha RGCs mediate looming-triggered innate defensive response. Curr Biol. 2021;31(13):2786-2797.
[3] Kerschensteiner D. Feature detection by retinal ganglion cells. Annu Rev Vis Sci. 2022;8:135-169.

Computational and Representational contributions to intensity invariance in Weber's Law

Behavioral invariances are typically assumed to depend on invariant neural representations, i.e., to have a representational origin. A case in point concerns the mechanisms underlying Weber's law (WL;[1]), which states that the discriminability between two stimuli depends on the ratio of their intensities, and is thus invariant to overall multiplicative changes in amplitude. We recently showed that WL in sound lateralization coexists with reaction times (RTs) that are lawful but not invariant to overall level [2]. Both RT and accuracy are very accurately described by a model based on bounded accumulation of evidence where the decision variable (DV) does not possess invariance to overall intensity. Instead, WL holds due to the dependence of the commitment probability on the mean and variance of the evidence [3]. We have now analyzed a massive dataset of 30 rats performing this task in over 1.6 million trials using a wider range of sound lateralization levels. The data displays tight adherence to WL and its associated RT regularity, but our previous model is not accurate for the largest lateralization levels. We show that accounting for these differences requires a partial level of explicit divisive normalization (DN)[4] in the sensory neurons conveying the evidence. A single model parameter interpolates between varying levels of DN. At one extreme (purely representational invariance) the DV is completely level invariant and so are accuracy and RT (unlike the experiments). At the other extreme lies our previous model (purely computational invariance), which fails for large level differences. A model with an intermediate level of DN accurately describes accuracy and RT across all conditions, suggesting that both evidence accumulation and divisive normalization are necessary. We conclude that both computational and representational mechanisms contribute to Weber's Law.
weber's law / perceptual decision making / drift diffusion model / reaction time / divisiven normalization.

Raghavendra Kaushik Archak

Mafalda Valente

Juan R. Castiñeiras de Saa

Alfonso Renart

Champalimaud Foundation, Lisbon, PT

Bridging scales between chemical space and behavioral phenotype

A core challenge in neuroscience is to model organism-level behavior without abstracting away relevant low-level biological processes. Foundational models in neuroscience learn to relate neuronal activity with behavior [1], while parallel models in pharmacology aim to map chemical structure to cellular effects [2]. Connecting these fields, we developed a multiscale model of how animal behavior can be manipulated at the chemical level.

We mapped molecular structure to behavioral phenotypes using a dataset of over 18,000 compounds screened in zebrafish larvae across three experimental paradigms: visuomotor behavior [3], sensory habituation [4], and circadian rhythm activity [5]. Using a pre-trained molecular embedding network and contrastive learning, we created a joint embedding space that aligns chemical structures with their corresponding behavioral fingerprints.

The resulting model generalizes, predicting learned features of the behavior for unseen compounds. The learned joint embedding space is structured, clustering compounds by mechanism of action and revealing relationships within chemical space. We also study the model's scaling properties as a function of both the chemical library and the behavioral readouts.

animal behavior / drug discovery / foundation model

Adrien Jouaryk

Dean Rance

Gonzalo G. de Polavieja

Christian K. Machens

Michael B. Orger

Champalimaud Foundation, Lisbon, PT

References

[1] Azabou, M., Arora, V., Ganesh, V., Mao, X., Nachimuthu, S., Mendelson, M., Richards, B., Perich, M., Lajoie, G., & Dyer, E. (2023). A unified, scalable framework for neural population decoding. Advances in Neural Information Processing Systems.

[2] Blucher, A. S., Celik, S., Jensen, J. D., Taylor, J., Cuccarese, M. F., Cooper, J. C., Rinaldi, J. M., et al. (2021). Mapping biology with a unified representation space for genomic and chemical perturbations to enable accelerated drug discovery. In Learning Meaningful Representation of Life Workshop at NeurIPS.

[3] Gendelev, L., Taylor, J., Myers-Turnbull, D., Chen, S., McCarroll, M. N., Arkin, M. R., Kokel, D., & Keiser, M. J. (2024). Deep phenotypic profiling of neuroactive drugs in larval zebrafish. Nature Communications.

[4] Lamiré, L. A., Haesemeyer, M., Engert, F., Granato, M., & Randlett, O. (2023). Functional and pharmacological analyses of visual habituation learning in larval zebrafish. Elife.

[5] Kroll, F., Donnelly, J., Gürel Özcan, G., Mackay, E., & Rihel, J. (2025). Behavioural pharmacology predicts disrupted signalling pathways and candidate therapeutics from zebrafish mutants of Alzheimer’s disease risk genes. Elife.

Learning What Matters: An Episodic Memory Perspective on Video-Language Models

Inês Baptista
Daniel MacNamee

Champalimaud Foundation, Lisbon, PT

Human memory is remarkably efficient, actively compressing less important details while prioritizing and retaining what is novel, emotionally significant, or functionally useful. Episodic memory, in particular, orchestrates our continuous experience by dynamically segmenting it into discrete, meaningful events. This event segmentation supports crucial cognitive functions like abstraction, generalization from sparse data, and the understanding of narratives. Neurocognitive research suggests that this segmentation process is tightly linked to surprise — moments of high prediction error — triggered when expectations about the unfolding environment fail. At these points, the hippocampus initiates memory updates, marking event boundaries and selectively encoding contextually significant moments.

On the other hand, current AI models struggle to process extended, continuous and high-dimensional data like long videos and movies. Despite new developments in the field, Video-language models (VLMs) struggle in organizing asynchronously visual, auditory, and textual inputs; interpret nuanced emotional tone; and understand connections between temporally distant events in a narrative — for example, understanding how a character's early decision shapes the outcome of a later scene. Crucially, unlike humans who naturally structure their past experiences into coherent "mental videos" with a clear sense of how one event leads to another, current VLMs lack mechanisms to automatically break down a continuous stream of data into meaningful individual events or to encode the temporal coherence vital for narrative understanding.

[1] Santos, S., Farinhas, A., McNamee, D. C., & Martins, A. F. T. (2025). ∞-VIDEO: A training-free approach to long video understanding via continuous-time memory consolidation. Retrieved May 8, 2025, from https://github.com/deep-spin/Infinite-Video

In this work, we investigate whether the neurocognitive framework of surprise-driven event segmentation provides a viable model for enhancing memory in VLMs. We also study how latent internal representations in VLMs evolve over long videos and whether cues like prediction failure (surprise), contextual shifts, or affective salience correlate with changes in the model's latent space. Ultimately, our goal is to guide the development of more brain-aligned VLMs — enhancing their ability to select and encode the most memorable events through mechanisms inspired by episodic recall, and improving performance on video question answering and long-range temporal reasoning tasks.

episodic memory / video-language models / event segmentation/
computer vision / artificial intelligence / interpretability

wed 15 — 18:00 Champalimaud Tours
The Communication, Events & Outreach Team guided groups through the corridors of the foundation, offering an engaging overview of the cutting-edge research facilities while enjoying the stunning surroundings.

thu 16 — 18:00 Walk and Talk + Sunset Drinks
Situated in Terraços do Carmo, just behind Largo do Carmo and near the Elevador de Santa Justa, this rooftop offers stunning views of downtown Lisbon and provides a perfect setting to mingle with fellow attendees. Champalimaud researchers guides the participants from the Champalimaud Foundation to the venue, giving an informal tour of Lisbon's downtown along the way.

social activities

fri 17 — 18:00 Boat Tour – Hippotrip
After 3 days of engaging scientific discussions, it was time for some fun with a Hippotrip adventure! The Hippotrip, an amphibious vehicle, lets you explore Lisbon by both land and river. The tour began at the Champalimaud Foundation, traveling along the Tagus River to downtown Lisbon, and ending at LIV Beato, the venue for the farewell dinner.

fri 17 — 19:30 Farewell Dinner + Party
On the final day, it was time for a farewell dinner at LIV Beato, with delicious Portuguese cuisine in a setting that captures the vibrant spirit of Portugal. Following the dinner, the evening continued with a lively party at the same venue. It was perfect to enjoy great food, connect with fellow researchers, and immerse in the relaxed atmosphere of one of Lisbon's historic neighborhoods.

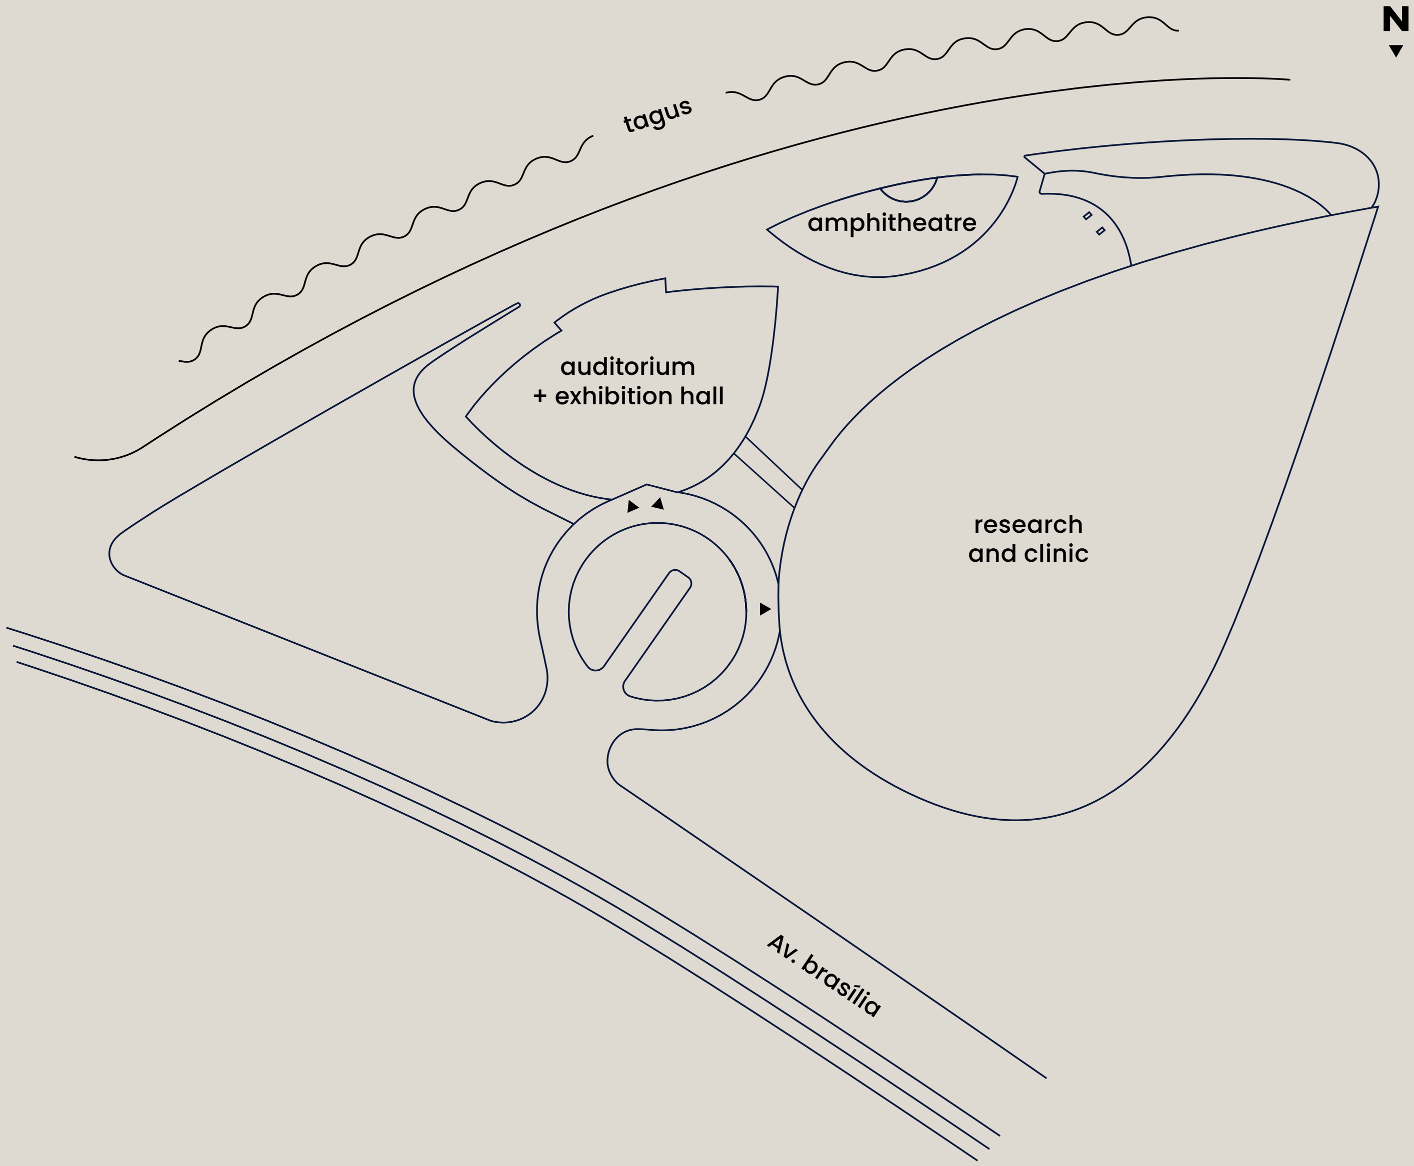

+ info

contacts

Champalimaud Centre for the Unknown

(+351) 210 480 200

Avenida Brasília 1400-038, Lisboa, PT

https://symposium.fchampalimaud.science

follow us

bluesky **instagram** **facebook** **youtube**

Champalimaud
Foundation