UK Neural Computation 2025

Programme

July 10-11th 2025

IMPERIAL

Advanced Research + Invention Agency

ARIA

THE FRANCIS CRICK INSTITUTE

www.neuralcomputation.uk

Contents

Huxley Building Concourse Level 3 (Main Lecture Theatre 340 & Posters Rooms 341, 342),
180 Queen's Gate, South Kensington Campus, Imperial College London, SW7 2AZ

Thursday 10th July 2025

Time	Speaker	Titles
11:30 – 12:30	**Registration & Lunch**	
12:30 – 12:45	**UKNC Organisers**	Welcome
12:45 – 13:00	**Dr Jacques Carolan** The Advanced Research + Invention Agency (ARIA)	Precision Neurotechnologies for Human Therapeutics
13:00 – 13:30	**Dr Jonathan Cornford** University of Leeds	Normative brain-like learning algorithms
13:30 – 14:00	**Dr Ian Hawes** University of Edinburgh **Selected from Abstracts**	Context-specific speed integration enables memory generalisation
14:00 – 14:30	**Prof Jennifer Bizley** University College London	How does the brain construct auditory space?
14:30 – 15:00	**Coffee**	
15:00 – 15:30	**Dr Flavia Mancini** University of Cambridge	The neural and computational blueprint for surviving pain
15:30 – 16:00	**Dr Andrea Colins Rodriguez** Universidad Adolfo Ibañez **Selected from Abstracts**	Rhythmic and discrete arm movements arise from the same control strategy in Primary Motor Cortex but not in the Supplementary Motor Area
16:00 – 16:30	**Dr Petr Znamenskiy** Francis Crick Institute	A depth map of visual space in the primary visual cortex
16:30 – 18:30	**Posters 1 & Networking Reception** Sponsored by The Francis Crick Institute (Partnership Networking Fund)	

IMPERIAL

Advanced Research + Invention Agency ARIA

THE FRANCIS CRICK INSTITUTE

Friday 11th July

Time	Speaker	Titles
9:30 – 10:00	**Prof Simon Schultz** Imperial College London	Neural manifold inference by maximising information
10:00 – 10:30	**Jack Cook** University of Oxford **Selected from Abstracts**	Brain-like pathways form in models with heterogeneous experts
10:30 – 11:00	**Prof Claudia Clopath** Imperial College London	Estimating the uncertainty of feedforward and feedback inputs with prediction-error circuits
11:00 – 11:30	**Coffee**	
11:30 – 13:00	**Breakout Session** Sponsored by the Advanced Research + Invention Agency (ARIA)	
13:00	**Lunch** Overlaps with start of poster session	
13:30 – 15:00	**Posters 2**	
15:00 – 15:30	**Prof Petra Vértes** **University of Cambridge**	Reservoir computing as a window into structure-function relationships in neural systems
15:30 – 16:00	**Sebastian Castedo** Laboratoire de Physique de l'ENS **Selected from Abstracts**	Energy-efficient neural coding under food restriction: structure, noise, and resilience
16:00 – 16:30	**Prof Paul Graham** University of Sussex	Spatial computations in shallow, insect-inspired networks
16:30 – 17:00	Coffee	
17:00 – 17:45	**Keynote: Prof Rafal Bogacz** University of Oxford	Predictive coding: effective learning with local plasticity
17:45	**UKNC Organisers**	Closing

Chair's note

Welcome to UK Neural Computation 2025!

Back in 2019 the time seemed ripe to launch a national meeting for the UK community working on the computational side of neuroscience, be it experimental or theoretical. Our community was thriving, with major investments and recruitment at institutions across the UK; the AI explosion was driving new approaches to understanding the brain; and the ever-present climate crisis was driving a greater awareness of our carbon footprint, for which a strong local meeting is one small way to reduce it.

The inaugural UK Neural Computation meeting in 2019 at the University of Nottingham was a wonderful, buzzing event, drawing attendees from across the UK. Post-pandemic gave us the belated sequel in Sheffield, and now we're delighted to bring you the 3rd edition, at Imperial College.

Our programme showcases the breadth and depth of UK computational work, from the single synapse and single neuron, through ensembles, networks, and circuits, up to the large-scale networks in the human brain. And, borrowing a fine COSYNE tradition, the first evening's poster session is accompanied by a drinks reception to encourage relaxed, convivial scientific discussion.

Hope you enjoy it!

Mark Humphries (Nottingham) on behalf of the program committee:

Danyal Akarca (Imperial College)

Dan Goodman (Imperial College)

Thanks

The conference organisers would like to thank Monika Ozelyte for helping with conference organisation and both Marcus Ghosh and Jascha Achterberg for their help in anonymously reviewing abstracts.

Sponsors

We thank Jacques Carolan and team at ARIA's Scalable Neural Interfaces programme for sponsoring the meeting and the Conference Fellowships. We thank the Francis Crick Institute for sponsoring the meeting and the networking event.

Posters at-a-glance

Poster session 1 (Thursday 10th)

1. Physical Network Constraints Define the Multiplicative Architecture of the Brain's Connectome
Piazza, B., Barabási, D.L., Ferreira Castro, A., Menichetti, G., Barabási, A.-L.

2. Predictive modelling of Alzheimer's Disease progression using Machine Learning Algorithms
Alex G.S, Ibeachu P.C

3. Ion channel distributions underlying sex differences in corticotroph cell excitability
Byttner W, Wedgwood K, Tabak J

4. Topological turning points across the human lifespan
Mousley A, Bethlehem RAI, Yeh FC, Astle DE

5. Common cross-scale cortical mechanisms for working memory deficits and auditory-verbal hallucinations in schizophrenia
Ivanov, TG, Oh, Y, Slifstein, M, Jiang, Y, Chang, X, Snellenberg, JV, Jones M, Hong, SJ, Feng, J, Abi-Dargham, A, Froudist-Walsh, S

6. Transmission Delay Determines In-Phase or Anti-Phase Functional Connectivity: Evidence from EEG data, Simulation and Stability Analysis
Omurtag A, Dragomir A, Lytton WW

7. Biophysical details matter: immature neuronal dynamics and dominance of NMDAR synaptic currents protect spatial information in the developing visual system
Tikidji-Hamburyan R.A., Colonnese M.T.

8. A Canonical Microcircuit of Predictive Coding Under Efficient Coding Principles
Nemati E, Davey CE, Meffin H, Burkitt AN.

9. Predictive learning of cognitive maps with apical dendrites in the hippocampus
Gillon CJ, Clopath C

10. Reinforcement Learning Drives the Emergence of Prefrontal Dynamics in a Recurrent Neural Network
Voce A, Cone I, Lak A, Costa RP

11. Modelling Demonstrates Phase-Reset is an Important Feature of Neural Speech Processing
Shannon AJ, Homer M, Barton DAW, Houghton CJ.

12. OpenWorm Project updates - towards a biologically constrained computational model of C. elegans locomotion and the development of a worm specific Large Language Model
Gleeson P, Vickneswaran Y, Ponzi A, Uher J, Kusmanov D, Sinha A, Larson SD

13. Using machine learning to compare categorical vs dimensional models of Autism phenotype data
Quigley H, McDaid L, Gardiner B, O'Donnell C.

14. Communication constraints promote efficient low entropy solutions in recurrent neural networks
Sharma J, Akarca D, Goodman D

15. Modelling the Impact of Neural Variability in Autism
Shepherdly L, McDaid L, Gardiner B, O'Donnell C

16. A biologically constrained spiking network model of prefrontal cortex using AdEx neurons
Pinheiro PR, Roque AC

17. Computational mechanisms of nonlinear multisensory integration in the superior colliculus for temporal disparities
Yuan A, Bianchini G, Iacaruso MF, Sadeh S

18. Uncovering brain-wide planning strategies with deep RL: Lessons from the Tower of Hanoi
Andrews A, Garibbo M, Achterberg J, Costa R

19. Spatial Navigation Engages a Distributed Neural Representation
van Beest EH, Terry B, Harris KD, Carandini M

20. A unifying neuronal network model framework addressing perceptual multistability and autism
Hikino K, Nakayama M, Wu Y, Barranca V

21. Controlling PFC dynamics for rapid learning and generalisation
Wójcik MJ, Achterberg J, Pemberton J, Costa RP

22. FlashPC: Fast Predictive Coding With Layer-Wise Parallelism
Cook J, Oliviers G, Bogacz R

23. μPC: Scaling Predictive Coding to 100+ Layer Networks
Innocenti F, Achour EM, Buckley CL

24. AI-Inspired Algorithms for Decoding Neural Codes: A Computational Neuroscience Approach
Phyllis Lee

25. Spikes can transmit neurons' subthreshold membrane potentials
Schmutz V

26. Spontaneous emergence of action decisions and slow ramping in a deep, brain-constrained model of frontotemporal areas
Griffin N, Schurger A, Garagnani M

27. Structure of spontaneous activity in mouse visual cortex
Haydaroglu A, Schmutz V, Krumin M, Reddy CB, Xu L, Shinn M, Skriabine S, Harris K, Carandini M

28. SST interneurons encode decision speed via causal regulation despite similar functional tuning across inhibitory neuron types in mouse barrel cortex
Asadpour A, Fekos C, Eley M, Maravall M

29. Computational modelling of memory coding with a dynamic neural network
Boscaglia M., Gastaldi C., Gerstner W., Quian Quiroga R.

30. Hierarchical neural encoding of context and frequency during the oddball paradigm
Hockley A, Bohórquez LH, Malmierca MS

31. Hebbian construction of inhibition-stabilized networks with multiple inhibitory neuron subtypes supports contextual modulation
Ianov-Vitanov R, Eckmann S

32. Degeneracy in Embodied Choice
Baker SA, Lepora NF

33. Learning dynamics in linear recurrent neural networks
Proca AM, Dominé CCJ, Shanahan M, Mediano PAM

34. Multiencoder VAE for cross-subject alignment of brain responses
Papathanasiou A, Achterberg J, Cone I, Nichols TE, Ponte Costa R

35. Cortical feedback circuits for online hierarchical credit assignment
Greedy W, Zhu HW, Duriez A, McCarthy P, Nejad K, Pemberton J, Ponte Costa R

36. Parallelising Neuronal Dynamics for Scalable Sequence Modelling
Stan M.I., Rhodes O.

37. Rescorla-Wagner Update Rule Predicts Metacognitive Bias
Lindersson C, Sax AL, Baddaley R, Ponte Costa R

Poster session 2 (Friday 11th)

1. Synaptic depression outperforms potentiation in learned stimulus discrimination under relative integration of opposing outputs
Abdelrahman NY*, Greenin-Whitehead K*, Jung JS*, Tan MW*, Jiang J, Yamada D, Hige T, van Rossum MCW, Lin AC *equal contribution, alphabetical order

2. Mechanisms of Memory Retrieval in a CA1 Hippocampal Microcircuit
Andreakos N, Shigang Y, Cutsuridis V.

3. A half-centre oscillator encodes sleep need in the Drosophila brain
Hasenhuetl PS, Sarnataro R, Vrontou E, Rorsman HO, Talbot CB, Brain R, Miesenböck G (PSH and RS contributed equally)

4. Computational mechanisms underlying how humans adapt their choices to the average effort of the environment
Scholey EV, Mehta NM, Apps MAJ

5. Autapses improve sensitivity to dynamic inputs in biological neurons.
McSweeney L, Vasilaki E, Toutounji H.

6. Theoretical Frameworks for Credit Assignment via Behavioral Timescale Synaptic Plasticity
Cone I, Clopath C, Costa RP

7. Dopamine D1 receptor expression in prefrontal parvalbumin neurons increases distractibility in marmosets compared with macaques
Ivanov TG, Joyce MKP, Krienen FM, Mitchell JF, Ma S, Inoue W, Nandy AS, Datta D, Duque A, Arellano J, Gupta R, Gonzalez-Burgos G, Lewis DA, Sestan N, McCarroll SA, Martinez-Trujillo J, Froudist-Walsh S, Arnsten AFT

8. Communication versus computation: The hidden costs, shaping the brain's architecture
Fakhar K, Akarca D, Luppi A, Oldham S, Hadaeghi F, Vertes P, Hilgetag C, Astle D.

9. Tracking the Brain's Path from Perception to Categorization of Mathematical Objects
Karami A, Debray S, Valerio D, Caute M, Dehaene S

10. Synaptic strength fluctuations from a model of stochastic gene expression in neurons
O'Donnell C, Veltz R, Senkevich O

11. Long delays reduce the need for precise weights in spiking neural networks
Pengfei Sun, Jascha Achterberg, Dan F. M. Goodman, Danyal Akarca

12. Computational Roles of Higher Order Thalamocortical Feedback in Context-Dependent Sensory Processing
McCarthy PT, Bruno RM, Costa RP

13. Category-biased patches encircle core domain-general regions in the human lateral prefrontal cortex
Assem M, Shashidhara S, Glasser M, Duncan J

14. Biological deep learning for simulating the role of SST Interneurons in cognitive resilience
Duisberg BFE, Marmouset-De La Taille I, Greedy W, Underwood S, Carlyle B & Costa RC

15. Overtraining Enhances Behavioural Flexibility on a Serial Reversal Learning Task: A Reinforcement Learning Perspective
Maggi S, Renstrom J, Grasmeder Allen R, Juty J and Bast T

16. Semantic Enrichment of Episodic Memories through Bidirectional Replay
Albesa-González A, Clopath C

17. Emergence of Auditory Receptive Fields Based on Surprise at Multiple Timescales
Yashaswini, Dash Sneha, Bandyopadhyay Sharba

18. Modeling Neural Adaptation to Tendon Transfer Surgery
Kawakita G, Gallego J

19. Adaptive learning in cortical networks: a computational model of Cholinergic neuromodulation
Filipovica M, Kermani Nejad K, Greedy W, Zhu HW, Mellor J, Ponte Costa R

20. Balancing external-internal learning: a theory of serotonin
Iris Marmouset-de la Taille, Simon J. B. Butt, Florencia Iacaruso, Rui Ponte Costa

21. Balancing stability and plasticity through gain neuromodulation and inhibitory gating
Rodriguez-Garcia A, Ramaswamy S

22. Selective inhibitory neuronal contributions to local field potentials in flexible perceptual decisions: A mean-field model prediction
Azimi A, Lenfesty B, Wong-Lin KF

23. Data-driven biophysically detailed computational modeling of neuronal circuits with the NeuroML standard and software ecosystem
Sinha A, Gleeson P, Ponzi A, Ray S, Panagiotou S, Marin B, Silver RA

24. Structure-function relationships in connectome-based echo state networks
McAllister J, Houghton C, Wade J, O'Donnell C

25. Dimensional Reconfiguration of Functional Network Manifolds Underlies Working Memory Performance in Neurofibromatosis Type 1
Khanbeigi AA, Litwińczuk MC, Garg S, Parkes LM, Madadi-Asl M, Lea-Carnall CA

26. Layer-specific input-output functions can shape representational geometry
Willard KJ, Ponte Costa R, Bruno RM

27. Mice wiggle a wheel to boost the salience of low visual contrast stimuli
Ghani, N. and The International Brain Laboratory

28. Cortically-Embedded RNNs for integration of cortex-wide neuroscience data into recurrent neural network models
Sevenster E, Thrivikraman A, Davies G, Klatzmann U, Pedamonti D, Froudist-Walsh S,

29. Cerebellum drives hippocampal networks for stable goal-driven navigation
Pedamonti D, Pemberton J, Costa RP

30. Distributional dopaminergic codes for economic decision-making
Salmasi M, Dolan R

31. Connectome-constrained recurrent neural networks
Rovný M, Akarca D, Achterberg J, Duncan J, Astle D

32. Spike timing in vivo for identified transcriptomic cell types
Shinn M, Zhou B, Prankerd I, Maat C, Bourdenx M, Nicoloutsopoulos D, Tilbury R, Shuker P, Bugeon S, Harris KD

33. Higher-order interactions reveal circuit motifs across brain regions and behavioural states
Rashid Shomali S, Rasuli SN, Shimazaki H, Sadeh S

34. Theta sweeps along the Papez's circuit
Ji Zilong, Burgess Neil

35. The impact of transcranial ultrasound induced heating on thalamic neurons and networks
Kour R, Jameel A, Smith J, Bain P, Nandi D, Jones B, Quest R, Gedroyc W, Borisyuk R, Yousif N

36. Cortico-basal ganglia dynamics underlying skilled locomotion
Esparza-Iaizzo M, Lazar I, Fortunato C, Safaie M, Gallego JA

Poster Abstracts

Poster session 1 (Thursday 10th)

1. Physical Network Constraints Define the Multiplicative Architecture of the Brain's Connectome

Piazza, B., Barabási, D.L., Ferreira Castro, A., Menichetti, G., Barabási, A.-L.

The brain has long been conceptualized as a network of neurons connected by synapses. However, attempts to describe the connectome using established network science models have yielded conflicting outcomes, leaving the architecture of neural networks unresolved. Here, by performing a comparative analysis of eight experimentally mapped connectomes, we find that their degree distributions cannot be captured by the well-established random or scale-free models. Instead, the node degrees and strengths are well approximated by lognormal distributions, although these lack a mechanistic explanation in the context of the brain. By acknowledging the physical network nature of the brain, we show that neuron size is governed by a multiplicative process, which allows us to analytically derive the lognormal nature of the neuron length distribution. Our framework not only predicts the degree and strength distributions across each of the eight connectomes, but also yields a series of novel and empirically falsifiable relationships between different neuron characteristics. The resulting multiplicative network represents a novel architecture for network science, whose distinctive quantitative features bridge critical gaps between neural structure and function, with implications for brain dynamics, robustness, and synchronization.

2. Predictive modelling of Alzheimer's Disease progression using Machine Learning Algorithms

Alex G.S, Ibeachu P.C

Background: Alzheimer's disease is the most common form of dementia, a significant contributor to the global burden of disability and death globally. Early diagnosis and accurate prediction of disease progression are crucial for effective management and treatment plan.

Objectives: To utilize machine learning models to predict the progression of Alzheimer's disease and enhance early detection.

Methodology: We utilized publicly available neuroimaging dataset comprising 2D axial MRI scans from three classes: Alzheimer's disease, Congnitive Impaired and Cognitive normal. Preprocessing of data was performed using OpenCV (Open source computer vision library). The models employed were ResNet50 and DenseNet121. Model performance was evaluated using accuracy, loss, AUC and confusion matrices

Results: Our results showed that both ResNet50 and DenseNet121 achieved exceptional accuracy of 99.8% and 100% respectivly. Confusion matrices of both models showed excellent performance in classifying images with AUC value exceeding 0.99.

Discussion: The study demonstrates the potential of machine learning in predicting the progression of Alzheimer's disease. Our finding suggest that machine learning models can help clinicians and healthcare professionals

Conclusion: Machine learning can enhance the accurate diagnosis of Alzheimer's disease. Future studies can build robust framework and explore the application of these models in clinical settings.

3. Ion channel distributions underlying sex differences in corticotroph cell excitability

Byttner W, Wedgwood K, Tabak J

Pituitary corticotroph cells are important players in our response to stress. They regulate the release of cortisol from the adrenals by releasing Adrenocorticotropic Hormone (ACTH) in response to hypothalamic signals. These cells are electrically active, like neurones, and the pattern of electrical activity determines ACTH release. Duncan et al. (2023) observed that female mice had more “A-type” corticotroph cells, which exhibit sharp action potentials. Male mice, instead, had more “B-type” corticotrophs, with wider spikes.

A-type corticotrophs are less easily stimulated to produce bursts of action potentials than B-type corticotrophs. Since bursts release more ACTH than single action potentials, the differences between A-type and B-type excitability may be linked to gender differences in the regulation of the stress response. Here, we aim to explain the differences between A-type and B-type excitability in corticotrophs. We adapt a corticotroph cell model from Fletcher et al. (2017) to reproduce A-type and B-type excitability, and ask what differentiates the two types. The parameters of this models are the ionic conductances that are responsible for excitability.

We first generate a database a models, from which we extract the parameter sets that produce clear A-type and B-type excitability. We then use machine learning tools to understand the combinations of parameters that are important for determining excitability type. This produces prediction about how to change ion channel conductances to switch a cell from one type to another. Preliminary results suggest that increasing the conductance of voltage-dependent Ca2+ channels and K+ channels may switch a cell from type B to type A excitability. This prediction is testable in the lab, using the dynamic clamp technique to add virtual ionic conductance models to real corticotroph cells.

4. Topological turning points across the human lifespan

Mousley A, Bethlehem RAI, Yeh FC, Astle DE

Structural brain topology, the complex motifs within which neural connections are organised, develops across the lifespan and is associated with key cognitive, behavioural, and mental health outcomes. While topology-outcome relationships have been established within relatively

narrow age ranges, the fundamental principles of topological change across the lifespan remain unclear. Are there key points in our lifespans wherein the brain transitions into a different phase of topological development? To investigate this, we analysed diffusion imaging from nine datasets with a collective age range of zero to 90 years old (N = 4,216). We analysed how 12 graph theory metrics of organisation change across the lifespan using Generalised Additive Models. We also employed two data-driven methods — Principal Components Analysis (PCA) and Uniform Manifold Projection and Approximation (UMAP) — to characterise changes in topological development. Within the manifold spaces generated UMAPs, we identified four major topological turning points across the lifespan – at eight, 32, 62, and 85 years old. These ages defined five major epochs of topological development, each with distinctive age-related changes in topology. PCA further revealed shifts in the primary sources of variance across these epochs. Additionally, all turning points mark shifts in organisational properties driving the age-topology relationship, however, only ages eight and 32 showed changes in the direction of topological measures. Beyond topological changes, these four turning points align with important cognitive, behavioural and mental health milestones, including the onset of puberty around eight years old and increased dementia risk in the 60s. Our study underscores the complex, non-linear nature of human development, with district phases of topological maturation, which can only be illumined with a multivariate, lifespan, population-level perspective.

5. Common cross-scale cortical mechanisms for working memory deficits and auditory-verbal hallucinations in schizophrenia

Ivanov, TG, Oh, Y, Slifstein, M, Jiang, Y, Chang, X, Snellenberg, JV, Jones M, Hong, SJ, Feng, J, Abi-Dargham, A, Froudist-Walsh, S

Cognitive deficits and hallucinations in schizophrenia are typically studied separately. We propose these symptoms share molecular, cellular, microcircuit, and large-scale mechanisms. Hallucinations may result from failures to suppress noise-evoked activations, allowing them to reach the threshold for conscious access, which critically relies on the superior cortical layers of the frontoparietal control network (FPCN), like working memory maintenance. We study the cross-level mechanisms in a biologically faithful dynamical system model of the neocortex, aiming to replicate cortical activity patterns linked with cognitive deficits and hallucinations in schizophrenia. As constraints, we use cortex-wide structural, chemical, and functional imaging datasets. Notably, schizophrenia is defined by supralaminar thinning, pronounced in frontotemporal cortices, which relates to reduced basal dendrite complexity in layer III pyramidal neurons. Also, uniformly lower dopamine release is seen in the cortex in schizophrenia, which is compensated by D1 receptor (D1R) overexpression but only in prefrontal areas. D1Rs modulate pyramidal and PV neuron excitability. Thus, using regional supralaminar thinning measures and D1R stimulation estimates, we modulate recurrent pyramidal excitation and dendritic SST and somatic PV inhibition in 360 area-specific microcircuits, embedded in the Glasser parcellation. Connections between microcircuits were weighted and directed based on the brain state (i.e. cognitive task or rest) closest to the respective symptoms. Overall, the differential dysregulation of cortical microcircuits may cause reduced persistent activity in FPCN, impairing working memory maintenance, and increased noise-evoked signals from the language network propagating to FPCN, where their gating on

the basal dendrites of pyramidal neurons is diminished due to the SST dysfunction, contributing to hallucinations. Crucially, our modelling allows testing new treatments in silico.

6. Transmission Delay Determines In-Phase or Anti-Phase Functional Connectivity: Evidence from EEG data, Simulation and Stability Analysis

Omurtag A, Dragomir A, Lytton WW

Measurement of ubiquitous long-distance in- and anti-phase synchronisation in mammalian cortex suggests a key role for these oscillations for information processing. Although brain synchronisation has been studied extensively using electroencephalography (EEG), we have lacked tools for identifying detailed phase relationships among these signals. Therefore, we developed a new metric, Phase-Relationship Index (PRI), to demonstrate the widespread patterns of inter-area in-phase and anti-phase synchronization in scalp EEG.

Study of 31 healthy participants (mean age 22 years, 17 female) with 19 channel EEG showed that phase clustering was dominated by in-phase relationships when the cortical distance between the areas was under ~80 mm, and switched to anti-phase as the cortical distance grew beyond ~120 mm. In/anti-phase clustered connections had markedly different topography and PRI showed significant potential as a biomarker as task performance conditions affected in/anti-phase clustered connections differently.

In simulations, the dynamics of delay-coupled populations of 400 purely excitatory Leaky Integrate-and-Fire neurons were sufficient to explain the underlying mechanism, suggesting that the in/anti-phase clustering was driven by the associated axonal conduction delays. The distance corresponding to the in/anti-phase switch (from PRI data) and the relevant delay (from simulation) were consistent with known biophysical propagation speeds. Further, Linear Stability Analysis of interacting populations, extending prior work (Sirovich, Omurtag, & Lubliner, 2006), demonstrated that a Hopf bifurcation underlies the emergence of in- and anti-phase clustering dynamics.

7. Biophysical details matter: immature neuronal dynamics and dominance of NMDAR synaptic currents protect spatial information in the developing visual system

Tikidji-Hamburyan R.A., Colonnese M.T.

Visual perception requires establishing precise connections between neurons in different brain areas during development. Many of the main computational elements of the visual cortex, such as visual topography and orientation-selectivity, are present at the eye-opening. Formation of these elements requires positional (spatial) information (SI) provided by retinal spontaneous wave-like activity that correlates spiking among neighboring ganglion cells (rGC) and encodes spatial information in inter-neuron spike intervals. The initial formation and refinement of thalamocortical and corticocortical synapses happen in parallel with the thalamus, and therefore, the SI must pass the thalamic network even though thalamic circuits are not fully refined. Foundational work by Butts et al. showed that the SI is encoded by the correlations of rGC at specific slow (>100ms) timescales. However, at this age, thalamocortical relay neurons (TCs) receive 10 to 20 inputs from rGC, which should cause correlation in input currents and, as

a result, strong correlations between TCs spikes at the millisecond timescale. To understand how SI is transferred in the developing thalamus, we established a detailed biophysical model calibrated to reproduce neuronal and synaptic dynamics at postnatal day 7 of mouse development. We showed that fast correlations are suppressed by the slow "immature" dynamics of developing neurons in combination with the synaptic currents dominated by slow NMDA receptors. Such fast correlations would be detrimental to cortical development because they reduce the SI in TCs firing. The decorrelation mechanism is a peculiar phenomenon because it suppresses correlations in a specific timescale, leaving informative correlations intact. Moreover, adult neuronal dynamics or synaptic composition removes this suppression, and therefore, accurate replication of synaptic and neuronal dynamics is essential for modeling information processing in the developing visual system.

8. A Canonical Microcircuit of Predictive Coding Under Efficient Coding Principles

Nemati E, Davey CE, Meffin H, Burkitt AN.

Predictive coding describes how the brain integrates sensory inputs with expectations by minimizing expectation errors. Increased neural activity in cortical layer 2/3 during sensory mismatches is vital for understanding cognitive phenomena and neurological disorders. Existing canonical models, however, often ignore detailed spiking dynamics, inhibitory mechanisms, adherence to Dale's law, and distinctions between cortical layers.

This study advances predictive coding by assigning biologically inspired Gabor receptive fields to cortical layer 4 neurons, enabling realistic V1 feature extraction. Layer 2/3 neurons employ a two-compartment structure to signal prediction errors robustly within a balanced excitatory-inhibitory network. The inhibitory network utilizes simplified populations inspired by Parvalbumin (PV) and Somatostatin (SOM) neurons, ensuring physiological accuracy. Layer 5/6 neurons generate predictions modulated by these inhibitory populations, enhancing biological plausibility.

The model employs Leaky Integrate-and-Fire neuronal dynamics to process images via ON/OFF channels, replicating LGN responses. Results show layer 4 neurons attain orientation and phase selectivity consistent with biological data, confirming their role as effective feature extractors. Layer 2/3 neurons reliably signal prediction errors across matched, mismatched, and feedforward- or feedback-only conditions, closely aligning with empirical observations. Importantly, layer 5/6 neurons effectively integrate these prediction errors from layer 2/3, significantly reducing sensory reconstruction errors and confirming their central predictive coding function.

This research highlights predictive coding's explanatory power for cortical functions, emphasizing the critical role of detailed neuronal compartments and inhibitory interneuron dynamics, forming a robust basis for future perceptual neuroscience and neuromorphic system research.

9. Predictive learning of cognitive maps with apical dendrites in the hippocampus

Gillon CJ, Clopath C

When an animal enters a new environment, neurons in the hippocampus begin to map out the new space. They become selectively responsive to features like the animal's location, its head-direction, the location of rewards, and the presence of stimuli relevant to navigating or performing a given task. With experience, hippocampal neurons also develop behaviour-related biases. It is thus believed that the hippocampus encodes a cognitive map of an environment that is not purely spatial, but rather multi-dimensional, and behaviourally-relevant.

Predictive learning provides a possible explanation of how such complex maps are developed. In this framework, as an animal interacts with an environment, the hippocampus learns to predict salient features of its experience, and, in so doing, maps out the environment in a behaviourally-relevant way. In sensory cortex, predictive learning provides a compelling explanation of anticipatory and error-like sensory responses. It is thought to rely on pyramidal neurons which receive top-down and bottom-up inputs, respectively, to their basal and apical dendritic compartments, allowing them to act as comparison units, signalling discrepancies between predictive and sensory inputs.

In the hippocampus, the potential link between pyramidal neurons and predictive learning is still under explored. Here, we investigate the possibility that pyramidal neurons perform a similar comparator function in the hippocampus. In our model, two-compartment pyramidal neurons receive spatial inputs basally which are compared to apically received sensory information about salient features of the environment. We demonstrate how predictive learning in this circuit using behavioural timescale plasticity can explain the emergence of the wide variety of spatial and behaviourally-relevant features encoded in the hippocampus. As such, our work helps bridge this important gap in the literature between predictive learning and pyramidal neurons in the hippocampus.

10. Reinforcement Learning Drives the Emergence of Prefrontal Dynamics in a Recurrent Neural Network

Voce A, Cone I, Lak A, Costa RP

The ability to learn associations between environmental cues, actions, and outcomes is central to flexible goal-directed behaviour. The prefrontal cortex (PFC) plays a key role in this process, exhibiting activity related to predicted rewards, action outcomes, and task representations. However, its contribution to associative learning remains difficult to elucidate due to the heterogeneous dynamics of prefrontal neurons and the PFC's recurrent, highly interconnected architecture. Computational modelling using recurrent neural networks (RNNs) offers a powerful tool for probing such mechanisms, granting full access to network structure and activity. Yet, many existing models rely on biologically implausible architectures, limiting their translational interpretability. The present study investigated whether a minimalistic RNN could replicate naturalistic, rodent-like behaviour in a cue-reward learning task, and whether this conferred the emergence of PFC-like neural signatures observed in animals during task performance. A single-layer Elman RNN was embedded within an Actor-Critic reinforcement learning framework and trained to perform a 'lick' action to obtain reward following a reward-predictive cue, but not a non-rewarded cue. The actor and critic outputs were fixed

linear readouts from the RNN, requiring the network to support both valuation and action selection. The model developed increasingly efficient and selective behavioural policies, including anticipatory licking to reward-predictive cues. Hidden unit dynamics showed increasing task-related activity, state value encoding, and cue-specific selectivity. Population-level activity across trials became progressively distinct and reliably decodable by cue-type. These findings demonstrate the potential of a simplistic RNN to capture key signatures of PFC-like neural dynamics during reinforcement learning. This project is currently being extended through direct comparison with mouse neural data.

11. Modelling Demonstrates Phase-Reset is an Important Feature of Neural Speech Processing

Shannon AJ, Homer M, Barton DAW, Houghton CJ.

Syllable segregation and source separation are foundational components of neural speech processing, yet consensus on their underlying mechanism remains elusive. Several hypotheses have been proposed, suggesting that the brain may align its activity to incoming linguistic stimuli via evoked responses, entrainment of endogenous oscillations, or some combination of the two. We investigate the origin of oscillatory behaviour in syllable segregation by modelling the dynamical response to periodic linguistic stimuli. We compare the response of a biophysically accurate neural mass model and a phase-resetting oscillator with prior experimental EEG data. We find that a correlation between neural activity entrainment strength and the sharpness of incoming phonemes, identified in the EEG experiment, is readily reproduced by both the neural mass model and the oscillator. However, when the phase-resetting dynamics in the oscillator are removed, the oscillator fails to reproduce the correlation. This demonstrates that phase-resetting is required for sharpness specific tuning of neural entrainment to speech. Identifying the neural correlates of this phenomenon may be possible through interrogation of the biophysical features of the neural mass model.

12. OpenWorm Project updates - towards a biologically constrained computational model of C. elegans locomotion and the development of a worm specific Large Language Model

Gleeson P, Vickneswaran Y, Ponzi A, Uher J, Kusmanov D, Sinha A, Larson SD

The OpenWorm project (http://openworm.org) is a global, online collaboration of computational and experimental neuroscientists, software developers and interested volunteers with an ambitious long-term goal: creating a cell-by-cell computer model of the worm C. elegans which reproduces the behaviour of the real animal in as much detail as possible. The project takes a unique Open Science approach to development, and provides a community resource which consolidates our anatomical and physiological knowledge of the worm, allowing investigators to examine the mechanistic underpinnings of how behaviour is generated by a complete nervous system.

OpenWorm contributors have already developed a 3D worm body model (Sibernetic) which incorporates a fluid mechanics simulator for modelling the interactions between the worm body, driven by its musculature, and the external environment. While this can be used for detailed simulations, it is computationally intensive. We have adapted published 2D worm

body simulators for use in the project as more efficient alternatives for testing the generation of behaviour in the body by the nervous system. We will also describe biophysically detailed neuronal cell models which have been translated to standardised NeuroML format to ease incorporation into our simulations. Additionally, we have consolidated multimodal connectivity information from the worm into the C. elegans Connectome Toolbox (https://openworm.org/ConnectomeToolbox).

LLMs hold great promise for facilitating access to huge amounts of scientific literature across multiple domains. We will outline our work to create a corpus of scientific literature and data related to C. elegans which can be used to fine-tune LLMs to allow extraction of scientific knowledge related to the worm. An online interface to this is available at https://openworm.ai. We demonstrate how this custom LLM can be used to help validate computational models of worm anatomy, physiology and behaviour.

13. Using machine learning to compare categorical vs dimensional models of Autism phenotype data

Quigley H, McDaid L, Gardiner B, O'Donnell C.

Autism Spectrum Disorder (ASD) has broad diagnostic criteria with a large range of symptoms, making accurate diagnosis and effective support challenging. To address this, some researchers attempt to categorise ASD into distinct subtypes, while others focused on describing individuals using multiple traits. The raises the open question of which approach more accurately characterises ASD. The objective of this project is to use Variational Auto-Encoders (VAEs) to analyse a large dataset of over 300,000 individuals containing phenotypic data aiming to develop a continuous model that defines a person based on their position along a scale, determined by multidimensional traits identified by the VAE's bottleneck layer. We compare the results to traditional statistical models (Factor Analysis and Gaussian Mixture Models) applied to the same dataset. By running all three models, we aim to determine if ASD is better described as distinct subtypes or as a spectrum of multidimensional traits. These findings will help improve how ASD is diagnosed and understood, paving the way for more targeted research and support strategies.

14. Communication constraints promote efficient low entropy solutions in recurrent neural networks

Sharma J, Akarca D, Goodman D

Neural networks trained to balance task performance with communication constraints develop structural properties in their weight matrices reminiscent of empirical brain networks. We investigate this relationship by introducing one recently proposed promising communication model as a constraint into the training of recurrent neural networks (RNNs), allowing us to systematically manipulate the extent to which short or longer communication pathways in the network are prioritised. Our task requires networks to recover the ground truth mean of features sampled from a multivariate normal distribution, thus measuring the network's ability to reconstruct the mean from the underlying covariance structure. Our findings reveal that communication constraints make RNNs quicker to train and this is maintained under strong

regularisation. Importantly, these networks exhibit a distinctive form of sparsity differing fundamentally from generic sparsity induced by conventional regularisation—characterised by concentrated weights through small numbers of neurons (indicated by lower Shannon entropy) and increased modularity when biased toward shorter communication pathways. These modular networks achieved the best performance in recovering the ground truth distributions, suggesting that the brain's characteristic network organisation—featuring modularity, small-world properties, and these specific brain-like sparsity patterns, may emerge naturally from competing pressures to perform statistical inference while minimising communication costs.

15. Modelling the Impact of Neural Variability in Autism

Shepherdly L, McDaid L, Gardiner B, O'Donnell C

Autism is a highly heterogeneous neurodevelopmental condition, characterised by considerable variability in behavioural traits, cognitive profiles, and neural properties across individual neuron populations. This heterogeneity complicates the identification of underlying biological mechanisms. Moreover, many autism studies rely on traditional case–control paradigms that focus on mean group differences, potentially masking the contribution of increased variability itself in shaping autistic phenotypes.

In this project we investigate the hypothesis that reduced developmental robustness leads to greater variance in neuronal parameters, contributing to the broad range of cognitive and behavioural outcomes observed in autism.

As a preliminary step, we trained convolutional neural networks (CNNs) on image classification tasks under varying levels of synaptic weight noise. Increased parameter variance led to broader and less predictable performance outcomes, and higher noise levels impaired training stability and generalisation—demonstrating that excess variance alone can disrupt computational function.

Building on this proof of concept, we use detailed biophysical neuron models to assess how variability at the single-neuron level shapes computational output. Using a biologically realistic Layer 5 pyramidal cell (L5PC) model, the project aims to simulate populations of neurons with systematically varied biophysical parameters. By analysing how such variability impacts input-output mappings and neuronal performance, we aim to establish a computational framework linking excess neural variance to functional consequences.

This integrated approach provides a novel pathway for understanding the mechanistic role of increased neural variability in autism, laying the groundwork for future experimental validation and therapeutic exploration.

16. A biologically constrained spiking network model of prefrontal cortex using AdEx neurons

Pinheiro PR, Roque AC

We present a data-driven spiking neuronal network model of the prefrontal cortex (PFC), developed using the adaptive exponential integrate-and-fire (AdEx) neuron model and implemented in the NEST simulator. Building on the architecture proposed by Hass et al. (2016), our model replaces the simplified AdEx (simpAdEx) neurons with full AdEx units, thereby enhancing the dynamical richness of individual neurons while preserving the original network topology. The network consists of 1,000 neurons distributed across two excitatory and eight inhibitory populations, spanning cortical layers 2/3 and 5. Synaptic interactions include AMPA, NMDA, and GABAergic receptors, all modelled with double-exponential kinetics and incorporating a 30% synaptic failure rate. Synaptic weights and delays are sampled from log-normal and normal distributions, respectively, and connectivity is defined by population-specific connection probabilities. Short-term synaptic plasticity follows the Tsodyks–Markram model, supporting facilitating, depressing, and combined dynamics. Our aim is to provide the computational neuroscience community with an openly accessible and biologically grounded PFC model, readily reproducible through its implementation in a widely adopted simulation framework. Preliminary simulations demonstrate that the network generates realistic membrane potential traces and supports asynchronous-irregular firing patterns consistent with in vivo activity in the awake state. This model offers a flexible platform for probing prefrontal cortical dynamics and serves as a foundation for large-scale investigations into both normal and pathological brain states.

Reference

Hass J, Hertäg L, Durstewitz D. (2016) A detailed data-driven network model of prefrontal cortex reproduces key features of in vivo activity. PLoS Comput. Biol. 12:e1004930.

17. Computational mechanisms of nonlinear multisensory integration in the superior colliculus for temporal disparities

Yuan A, Bianchini G, Iacaruso MF, Sadeh S

Multisensory integration is a fundamental neural process where the brain combines information from different sensory modalities to enhance perception and behavioural responses. The superior colliculus (SC) is a key brain region that integrates multisensory inputs, particularly audiovisual (AV) signals. In this study, we investigated how neurons in the SC encode temporal disparities between AV stimuli. By presenting AV stimuli with varying temporal delays, our electrophysiological recordings revealed that SC neurons integrate AV inputs nonlinearly, producing responses that differ from the linear sum of unisensory responses, thereby enhancing the encoding for temporal disparities. To understand the mechanisms of this nonlinear integration, we developed a computational network model of the SC based on our experimental findings. The model incorporates the estimated neuronal connectivity between subpopulations of neurons combined with various nonlinear transfer functions for individual neurons. By systematically varying stimulus timing in our simulations, we demonstrated that supralinear integration, where the multisensory response exceeds the linear sum of unisensory responses, significantly improves the encoding of AV delays. This improvement is positively correlated with the degree of nonlinear integration. Further analysis revealed that this nonlinear integration originates from neuronal nonlinear transfer function, which is subsequently

amplified by the network's recurrent and feedforward connectivity patterns. Among these connections, the visual-to-multisensory projections emerged as the primary contributors to nonlinear integration and encoding of AV delays, suggesting a dominant role of visual inputs in multisensory temporal processing in the SC. Our experimental and computational results shed light on how local circuit architecture in the SC supports nonlinear multisensory integration to enhance the precision of temporal processing in multisensory contexts.

18. Uncovering brain-wide planning strategies with deep RL: Lessons from the Tower of Hanoi

Andrews A, Garibbo M, Achterberg J, Costa R

Human planning involves generating and executing action sequences under environmental constraints (Mattar & Lengyel, 2022). Experimental studies have identified that areas such as the prefrontal cortex (PFC) , hippocampus and cerebellum play important roles during planning (Grafman et al., 1992; Goel & Grafman, 1995). We propose that the architectures of deep reinforcement learning agents capable of solving human-level planning tasks can offer a normative framework for understanding the involvement of different brain regions in planning. To demonstrate this, we use MuZero (Schrittwieser et al.,2020) and a widely used task to study goal-directed planning and behavior, Tower-of-Hanoi (ToH). We evaluate the performance of MuZero on the ToH under targeted network ablations to simulate brain region-specific lesion studies. Ablating the value network reproduces the behavior observed in patients with PFC damage, while ablating the policy network mimics cerebellar damage. Our preliminary results suggest that deep RL architectures may provide a brain-wide account of human planning.

19. Spatial Navigation Engages a Distributed Neural Representation

van Beest EH, Terry B, Harris KD, Carandini M

Introduction:

Spatial navigation is a complex behaviour involving spatial, sensory, reward, and motor processes. While research has primarily focused on the hippocampal formation and its well-known place and grid cells, spatially tuned activity also appears in regions like retrosplenial, visual, and olfactory cortex. Building on this, we investigate how broadly spatial navigation recruits distributed representations across the brain.

Methods:

To dissociate spatial from sensory, motor, and reward processes, we used a virtual reality corridor and Neuropixels recordings in visual, somatosensory, retrosplenial, motor cortex, hippocampus, dorsal thalamus, striatum, and midbrain. We modelled neuronal activity using reduced-rank ridge regression, combining spatial predictors (place fields) with temporal kernels for sensory stimuli, reward, and running speed. Significance of each predictor was assessed by circularly shifting it 500 times and comparing the resulting variance explained.

Results:

Our model explained over 1% of the explainable variance in 62.5 ± 12.0% of neurons across areas. Non-spatial predictors, particularly running speed, significantly improved model fits in 50.5 ± 14.6% of neurons. Spatial predictors also significantly contributed in 12.8 ± 6.7% of neurons, especially in anterior, primary, and posteromedial visual cortex, limb-related somatosensory cortex, retrosplenial and motor cortex, several thalamic regions, and hippocampus.

Conclusion:

Spatial navigation engages distributed coding of spatial and non-spatial variables across the brain. Neurons show mixed selectivity, indicating that spatial behaviour is supported by widespread, overlapping neural representations.

20. A unifying neuronal network model framework addressing perceptual multistability and autism

Hikino K, Nakayama M, Wu Y, Barranca V

While perceptual multistability arises from many types of stimuli across different sensory systems, there are common dynamical features that may be rooted in universal organizing principles underlying perception. We probe the mechanisms responsible for visual multistability using a neuronal network model framework in which a set of realistic images directly drives competing pools of neurons with nonlinear dynamics. Incorporating balanced network architecture, long-range connections from excitatory neurons in one pool to inhibitory neurons in the other pools, and a dynamic spiking threshold, the model produces irregular percept switching and replicates key experimental observations. Using a sequence of short-time observations of neuronal dynamics, we derive a methodology for reconstructing the dynamic percept that generalizes to an arbitrary number of percepts. The model dynamics illustrate that perceptual alternations are potentially rooted in the breakdown of balance between excitation and inhibition, with more balanced dynamics generally facilitating longer dominance durations. Finally, we show increasing the ratio of excitatory to inhibitory inputs in the network, either by increasing excitatory connection strengths or decreasing inhibitory connection strengths, systematically yields longer dominance durations as observed for individuals with autism, and we thus demonstrate support for the excitation/inhibition imbalance hypothesis for autism.

21. Controlling PFC dynamics for rapid learning and generalisation

Wójcik MJ, Achterberg J, Pemberton J, Costa RP

The prefrontal cortex (PFC) enables rapid adaptation and flexible behaviour in response to novel cognitive demands, yet the computational mechanisms underlying this flexibility remain poorly understood. Here we introduce a biologically inspired recurrent neural network (RNN) model in which task performance is dynamically guided by an external controller operating on two distinct timescales. On a fast timescale, the controller modulates RNN activity to induce high-dimensional task-relevant representations without modifying synaptic weights. On a

slower timescale, plasticity within the RNN consolidates these representations into low-dimensional, abstract formats.

Geometric and dimensionality analyses reveal that the controller-based model reproduces key features of PFC population activity observed in non-human primates learning a cognitive XOR task. During the early stages of training, high-dimensional, control-driven representations support flexible encoding, whereas later stages involve a reduction in dimensionality consistent with representational refinement through recurrent plasticity. This temporal dissociation between control-driven dynamics and synaptic learning mirrors experimental observations, and supports the hypothesis that, during early learning, the PFC functions as a reservoir whose activity is shaped by modulatory control signals.

Furthermore, we show that once trained on a task, the model generalises rapidly to new task instances that preserve abstract structure but differ in input features. During this generalisation phase, plasticity in the controller alone is sufficient to produce correct output behaviour, with fixed recurrent weights—recapitulating behavioural and neural signatures of cross-context generalisation observed in our primate dataset.

Together, these findings offer a mechanistic account of how the PFC may jointly leverage dynamic control and synaptic plasticity to support flexible cognition and generalisation across tasks.

22. FlashPC: Fast Predictive Coding With Layer-Wise Parallelism

Cook J, Oliviers G, Bogacz R

Predictive coding has made gains in recent years as a biologically plausible alternative to backpropagation. It was recently found to achieve comparable performance to backpropagation on simple image recognition tasks, such as MNIST and CIFAR10. However, a lack of high-performance implementations of predictive coding has made it difficult to train larger models on more difficult tasks, preventing its widespread adoption. In this work, we introduce FlashPC, a GPU kernel that implements a fast, highly parallel implementation of predictive coding. After performing a standard forward pass at the start of each training step, FlashPC performs the subsequent relaxation steps in parallel across the layers of a model. As a result, FlashPC yields a 6.8x speedup over a standard implementation of predictive coding when training VGG-9, and a 10.7x speedup when training ResNet-18. We then use this GPU kernel to train larger predictive coding models with improved performance, outperforming previous results on several image recognition tasks. We release our code and results, with the hope of speeding up future work on predictive coding.

23. μPC: Scaling Predictive Coding to 100+ Layer Networks

Innocenti F, Achour EM, Buckley CL

The biological implausibility of backpropagation (BP) has motivated many alternative, brain-inspired algorithms that attempt to rely only on local information, such as predictive coding

(PC) and equilibrium propagation. However, these algorithms have notoriously struggled to train very deep networks, preventing them from competing with BP in large-scale settings. Indeed, scaling PC networks (PCNs) has recently been posed as a challenge for the community. Here, we show that 100+ layer PCNs can be trained reliably using a Depth-μP parameterisation which we call ``μPC". Through an extensive analysis of the scaling behaviour of PCNs, we reveal several pathologies that make standard PCNs difficult to train at large depths. We then show that, despite addressing only some of these instabilities, μPC can train very deep (up to 128-layer) residual networks on simple classification tasks with competitive performance and little tuning compared to current benchmarks. Moreover, μPC enables zero-shot transfer of both weight and activity learning rates across widths and depths. Our results have implications for other local algorithms and could be extended to convolutional and transformer-based architectures. Code for μPC is made available as part of a JAX library for PCNs at https://github.com/thebuckleylab/jpc.

24. AI-Inspired Algorithms for Decoding Neural Codes: A Computational Neuroscience Approach

Phyllis Lee

Decoding neural codes represents a pivotal challenge in computational neuroscience, bearing substantial implications for elucidating brain function and advancing brain-computer interfaces. Conventional methods for neural decoding frequently depend on statistical techniques and linear models. Nevertheless, the intricate and non-linear nature of neural systems implies that more advanced algorithms are required to precisely decode neural activity. Recent progress in AI-inspired algorithms, notably deep learning models, has demonstrated considerable promise in various decoding tasks, frequently surpassing traditional methodologies. This article examines the utilization of AI-inspired algorithms for decoding neural codes, emphasizing their capacity to enhance our comprehension of the brain and refine neural interfaces. Furthermore, we address the challenges and opportunities associated with employing these algorithms, including issues related to data availability, interpretability, and computational resources. By addressing these concerns, we can harness the full potential of AI-inspired algorithms to unlock the secrets of the neural code and pave the way for new breakthroughs in neuroscience and neuroengineering. The convergence of neuroscience and artificial intelligence is revolutionizing our ability to decipher the intricate language of the brain. By drawing inspiration from the structure and function of biological neural networks, AI algorithms are providing unprecedented insights into neural coding and information processing.

Deep learning, a subfield of AI, has emerged as a particularly powerful tool for decoding neural activity. Deep learning models, inspired by the architecture of the brain, consist of multiple layers of interconnected nodes that learn to extract complex features from raw data.

25. Spikes can transmit neurons' subthreshold membrane potentials

Schmutz V

Neurons in the mammalian brain mainly communicate through the emission of large, pulse-like depolarizations of their membrane potential called spikes. To fire a spike, a neuron's membrane potential needs to cross a threshold. Between spikes, the subthreshold membrane potential fluctuations of a neuron are, by definition, not transmitted to other neurons, but they carry precious information as they reflect the total synaptic input received by the neuron. The richness of sensory and behavioral information contained in the subthreshold membrane potential dynamics, compared to that contained in the timing of spikes, has put forward the idea that spikes are just "the 'tip of the iceberg' in terms of neuronal activity." But if a single presynaptic neuron does not transmit any information about its subthreshold membrane potential fluctuations, does this imply that subthreshold information is lost, in the sense that it is not accessible to a postsynaptic neuron? At the single-neuron level, the loss is real; optimal estimation of the membrane potential from spikes only allows for partial recovery of the membrane potential. Adopting a population-level perspective, we prove that the membrane potential fluctuations of a presynaptic population of neurons emitting sparse spikes can be fully and perfectly transmitted to a postsynaptic population of neurons. Our proof combines ideas from high-dimensional probability with recent results on concentration of measure in networks of spiking neurons. Incidentally, this mathematical result provides a possible explanation for why the processing of sensory signals is improved in the 'desynchronized state' of cortical activity: weak correlations in the desynchronized state may be the reflection of a high-dimensional regime enabling population-level transmission of subthreshold information.

26. Spontaneous emergence of action decisions and slow ramping in a deep, brain-constrained model of frontotemporal areas

Griffin N, Schurger A, Garagnani M

When placed in front of a button and instructed to press it whenever they "feel the urge" to do so, humans typically make a "spontaneous" decision to press at seemingly random times, usually between 5 and 30sec from trial start. In such a situation, how does our brain decide when to act?

We used a brain-constrained neural network model of six relevant areas of the cortex to investigate this issue. Building on previous results, the network was trained to induce the formation of distributed perception-action circuits linking visual input patterns to corresponding motor actions. We then analysed the network's behaviour in absence of any input and observed spontaneous ignitions of such circuits (induced by noise, simulating baseline neuronal firing), which we took as model correlates of endogenous decisions to act. By repeatedly measuring simulation time between trial start and first ignition, we built the network's "waiting time" (WT) distribution, which we compared with experimental data.

We found that the network autonomously "decided" to act (or wait) in a way that accurately reflected participant behaviour - the simulated and experimental WT distributions were statistically indistinguishable. We also found that the network activity (total neuronal firing) prior to a spontaneous ignition event replicated the pre-movement slow ramping neural signal (known as the readiness potential, RP) experimentally recorded in humans using EEG.

Replicating both behavioural and brain indexes of spontaneous voluntary movements, the present brain-constrained model offers a neuro-mechanistic explanation for the emergence of endogenous decisions to act in the human brain. Importantly, our results speak to the interpretation of the RP: rather than reflecting a decision outcome, the present results suggest that the emerging slow ramping signal is pre-decisional, with commitment to movement happening when an action-perception circuit's internal activity overcomes an inherent threshold.

27. Structure of spontaneous activity in mouse visual cortex

Haydaroglu A, Schmutz V, Krumin M, Reddy CB, Xu L, Shinn M, Skriabine S, Harris K, Carandini M

Spontaneous activity in visual cortex is high-dimensional and state-dependent, yet its organizing principles are unclear. We combine large-scale imaging, modelling, and theory to ask:

1. Does spontaneous activity share the structure of evoked activity?

2. Is spontaneous activity organized spatially?

We constructed a modified Light Beads Microscope to image calcium in excitatory neurons across large volumes. We developed an open-source, GPU-accelerated volumetric cell extraction pipeline. These methods enabled simultaneous recording of 30,000+ cells from awake mice during spontaneous activity and visual stimulation.

Cells with similar visual preferences did not have higher spontaneous correlations. The covariance structure of spontaneous activity was not predicted by that of evoked activity, showing that evoked and spontaneous activity do not share the same structure. Spontaneous activity exhibited a weakly spatial structure, with pairwise correlations decaying over ~1mm. We developed Spatial Shared Variance Component Analysis to quantify shared activity between spatially separate populations of cells. The magnitude and dimensionality of shared activity decayed with distance between populations. Cortical correlations could thus be characterized by a few, strong, global dimensions and many weaker local dimensions. We modelled spontaneous activity using a linear RNN with distance-dependent connectivity driven by low-dimensional inputs, which was sufficient to produce high-dimensional activity with spatiotemporal structure similar to that of neural data. We showed analytically that this model can produce high dimensional activity from low-dimensional inputs.

Spontaneous activity does not share its structure with evoked activity. It follows a spatial organization with a few global and many spatially localized modes. The global modes may correspond to a low-dimensional arousal state, while remaining dimensions may be "reverberations" predicted by our model.

28. SST interneurons encode decision speed via causal regulation despite similar functional tuning across inhibitory neuron types in mouse barrel cortex

Asadpour A, Fekos C, Eley M, Maravall M

Decision-making relies on the coordinated activity of excitatory and inhibitory neurons. Previous studies have explored the functional contributions of the principal distinct inhibitory interneuron subtypes—somatostatin (SST), parvalbumin (PV), and vasoactive intestinal peptide (VIP) neurons—often linking them to stimulus-, action-, or reward-related processing in isolation, e.g., Ramamurthy et al. (2023). However, establishing separate functional and causal roles across types would require experiments enabling direct comparison.

To fill this gap, we conducted two-photon calcium imaging of genetically identified interneurons in head-fixed mice performing an established somatosensory go/no-go discrimination task (Bale et al., 2021). Mice were trained to discriminate whisker stimuli and rewarded with water for hit trials, while miss and correct rejection (CR) trials received no reward and false alarm (FA) trials incurred a time-out.

Using a recurrent neural network (RNN) classifier, we first determined whether each neuron was responsive in hit, FA, CR, or miss trials, then used a truth-table framework to categorise their functional responses. This revealed that SST and PV interneurons are distributed across multiple behavioural roles rather than adhering to a single specialised function. Preliminary analysis suggests VIP neurons follow a similar pattern.

Building on this, we investigated whether these interneuron subtypes differ in their causal influence during decision-making. Using Granger Causality analysis, we examined pairwise directional interactions across neuron types during hit and FA trials, subdivided by decision speed (fast, medium, slow). SST neurons consistently exerted stronger causal influence, particularly in hit and slow FA trials ($p < 0.001$). Assuming decision speed as an indirect measure for confidence (Kepecs et al., 2008), these findings suggest SST neurons may play a specialised role in encoding decision confidence.

29. Computational modelling of memory coding with a dynamic neural network

Boscaglia M., Gastaldi C., Gerstner W., Quian Quiroga R.

Memory retrieval improves with repetition and, in line with this observation, experimental evidence shows that familiar items recruit larger neural populations than novel ones. To capture the dynamic nature of such neural representations, we built an attractor network with dynamic synapses whose fundamental elements are: online Hebbian plasticity, background firing rate, neuronal adaptation, and heterosynaptic plasticity. Ongoing Hebbian learning in the presence of spontaneous activity led to the size increase of memory assemblies with stimulus frequency. Notably, assemblies representing uncorrelated memories changed their sizes without interfering with each other (i.e. their neural representations remained orthogonal), in line with results from human single-cell recordings suggesting that partial overlaps between neural assemblies represent meaningful associations between the corresponding memories. We also observed that neural assemblies that were not further stimulated were eventually forgotten, and their neurons became available to create or reinforce other representations. We then asked how independent memories interact when experience links them. Two initially

orthogonal assemblies were repeatedly activated either separately or in pair. The resulting overlap between the initially orthogonal assemblies depended systematically on the frequency of paired stimulation. Specifically: high frequencies drove the assemblies to merge; intermediate frequencies yielded a stable partial overlap, allowing the associated memories to be activated jointly or separately; in case of low frequencies (i.e. paired activations were relatively rare) assemblies preserved their orthogonality. Together, these simulations show that the competitive balance between ongoing Hebbian learning, spontaneous activity, and heterosynaptic mechanisms can explain how memory strength, forgetting, and concepts' association emerge within the same system.

30. Hierarchical neural encoding of context and frequency during the oddball paradigm

Hockley A, Bohórquez LH, Malmierca MS

Neural processing of sensory stimuli is modulated by context. Predicted stimuli in regular patterns are modelled internally, suppressing responses. Responses to unexpected stimuli are increased in the form of prediction errors, which enhance the salience of, and guide attention toward, unpredictable stimuli. Predictive neural mechanisms are still debated, however studies have revealed hierarchically organized predictive circuits, both within cortical layers and between brain regions. Here we study hierarchical encoding of frequency and context information during the auditory oddball paradigm. We decode between neural responses to different sets of auditory stimuli to determine if neurons encode frequency or context, or a combination of both. We focused on the central nucleus of the inferior colliculus (CNIC), primary auditory cortex (A1) and medial prefrontal cortex regions (mPFC) as representative areas along the hierarchy, with the hypothesis that lower sensory processing regions would encode frequency, whereas higher-order areas would encode context. We found that CNIC neurons do not encode context and weakly encode stimulus frequency during the oddball paradigm. A1 neurons also weakly encode frequency but strongly encode predictability context as prediction errors. Prediction errors in A1 encode frequency change direction, most strongly in infragranular layers. Neurons of the mPFC do not encode frequency information, but encode contextual information such as predictability, with a preference for ascending frequency change. These data reveal hierarchical organization and interplay of context and frequency encoding during the auditory oddball paradigm.

31. Hebbian construction of inhibition-stabilized networks with multiple inhibitory neuron subtypes supports contextual modulation

Ianov-Vitanov R, Eckmann S

Surround suppression in visual cortex emerges from recurrent interactions among excitatory pyramidal (PYR) neurons and multiple subtypes of inhibitory neurons. In particular, somatostatin-expressing (SOM) neurons show increased firing when their receptive-field surround is stimulated, thus inhibiting parvalbumin-positive (PV) and PYR neurons in their receptive-field centre. Classical models relied on supralinear inhibition-stabilized networks (SSNs) of only two neuron types (PYR and PV) to explain these effects. Later models also considered SOM neurons, but their synaptic weights were either hand-tuned or trained with

computationally expensive machine-learning algorithms, limiting their scalability to larger network sizes.

Here, we present a principled way to efficiently construct large-scale inhibition-stabilized network models of multiple interneuron subtypes. Similar to previous work, we employed a synapse-type-specific competitive Hebbian plasticity rule and analytically solved for the steady-state weight matrix, assuming that during learning, the network is dominated by external feedforward input. To relate our model to classic SSNs, we construct networks with and without SOM neurons that retain similar firing rate dynamics by analytically deriving an approximate equivalence condition that becomes exact when SOM neurons have linear activation functions.

Our model reproduced key experimental findings—SOM firing monotonically increased with stimulus diameter, whereas PYR and PV responses peaked at intermediate sizes and were suppressed by large stimuli. Silencing SOM cells elevated excitatory rates and weakened surround suppression, reflecting perturbation experiments in biological circuits. In summary, we present a scalable, robust framework that offers a compact mechanistic account of contextual modulation and a principled scaffold for future multi-cell-type network models.

32. Degeneracy in Embodied Choice

Baker SA, Lepora NF

Perceptual decision-making provides a framework for understanding how organisms translate sensory evidence into actions, but traditional models face challenges in explaining choice phenomena and motor integration. Despite evidence of both covert and overt motor processes during deliberation, most frameworks treat movement as merely implementing a completed decision.

We explore the relationship between action and decision making by extending a proposed framework for embodied choice and independently varying the influence of motor feedback on internal choice variables and the contribution of evidence to action. This new model, Degenerate Embodied Choice (DEC), arbitrates between parallel and embodied theories of choice. We demonstrate that DEC replicates the speed-accuracy trade-off (SAT) degenerately, with embodiment proving both necessary and unique for trading speed and accuracy across urgent and accuracy-emphasised tasks. DEC emulates empirical data both qualitatively and quantitatively, with model-fitted parameters falling exclusively within the embodied set and producing congruent predictive SAT values within a narrow band. We then introduce the Optimality Framework for Embodied Choice (OFEC) as a lens for examining embodied choice through optimality principles.

Our findings suggest that complex decision behaviours can emerge from simple underlying principles, whether through geometric properties of decision boundaries or motor-cognitive integration.

33. Learning dynamics in linear recurrent neural networks

Proca AM, Dominé CCJ, Shanahan M, Mediano PAM

Recurrent neural networks (RNNs) are powerful models used widely in both machine learning and neuroscience to learn tasks with temporal dependencies and to model neural dynamics. However, despite significant advancements in the theory of RNNs, there is still limited understanding of their learning process and the impact of the temporal structure of data. Here, we aim to bridge this gap by analyzing the learning dynamics in linear RNNs (LRNNs) analytically. We derive an analytical solution to the learning dynamics of LRNNs under certain conditions, enabled by a novel framework that accounts for task dynamics. Our mathematical result further reveals three key properties of LRNNs: (1) Task dynamics impact solution stability and extrapolation. (2) The tradeoff between recurrent and feedforward computation is governed by a phase transition that leads to low-rank solutions. (3) Recurrence facilitates rich learning, as shown through a novel derivation of the neural tangent kernel for finite-width LRNNs. As a final proof-of-concept, we apply our theoretical framework to explain the behavior of LRNNs performing sensory integration tasks. Our work provides a first analytical treatment of the relationship between the temporal dependencies in tasks and learning dynamics in LRNNs, building a foundation for understanding how complex dynamic behavior emerges in cognitive models.

34. Multiencoder VAE for cross-subject alignment of brain responses

Papathanasiou A, Achterberg J, Cone I, Nichols TE, Ponte Costa R

Neural responses to identical stimuli vary considerably across individuals despite similar behavioural outcomes. Recent research demonstrates preserved latent neural dynamics in motor cortical populations across monkeys performing identical motor tasks. Inspired by these observations we introduce a multiencoder variational autoencoder (VAE) to model visual cortex responses. Our approach transforms subject-specific fMRI responses from natural scene viewing into a common latent space while predicting artificial neural network (ANN) activations elicited by identical stimuli. Using the Natural Scenes Dataset (NSD), our method outperforms traditional alignment techniques by capturing cross-subject representational similarities. The VAE architecture implements subject-specific encoders which project occipitotemporal cortex responses into a shared latent manifold that preserves semantic organisation while accommodating neuroanatomical variability. Simultaneously, the decoder establishes a computational correspondence between this latent representation and ResNet-50 activations. This approach creates a framework for investigating shared neural representations across individuals while quantifying systematic relationships between biological and artificial NNs.

35. Cortical feedback circuits for online hierarchical credit assignment

Greedy W, Zhu HW, Duriez A, McCarthy P, Nejad K, Pemberton J, Ponte Costa R

In recent years, deep learning models have experienced tremendous growth in their learning capabilities. These successes have predominantly been driven by new network architectures, larger datasets, and more powerful hardware, while the fundamental approach to credit

assignment, the error backpropagation (backprop) algorithm, has remained unchanged. In the brain, however, such a ubiquitous learning mechanism has yet to be identified, and many have long believed that backprop is biologically implausible or insufficient to explain the brain's wide array of learning capabilities. Recent models have challenged this view by demonstrating efficient and biologically plausible backprop-like learning.

We build upon our previously proposed Bursting Cortico-Cortical Network (BurstCCN), which provides a biologically plausible mechanism of backprop in which error signals are represented in the apical dendrites of pyramidal neurons that control the rate of bursting activity. These bursts induce local burst-dependent plasticity and give rise to a burst multiplexing code that facilitates the backward propagation of errors through connection-type-specific STP feedback connections. We show that the BurstCCN can use a single-phase learning process to effectively backpropagate error signals that both empirically and analytically approximate backprop-derived gradients. We demonstrate the BurstCCN's ability to effectively learn complex image classification tasks (MNIST, CIFAR-10, and ImageNet) and from scalar reward signals. Our model also offers a unifying explanation for experimentally observed cortex-wide interneuron gradients and plasticity dynamics. Finally, we extend the BurstCCN model to enforce Dale's principle, which requires a strict separation of excitatory and inhibitory neuron populations. We show this model still learns successfully and that ablating specific interneuron populations reproduces experimentally observed changes in synaptic plasticity.

36. Parallelising Neuronal Dynamics for Scalable Sequence Modelling

Stan M.I., Rhodes O.

The success of modern deep learning systems, such as Large Language Models, has been driven in part by the scalability of GPU-parallelisable deep architectures, including Transformers and State Space Models (SSMs). However, recent work has also highlighted how fully parallel models cannot match the range of dynamics of Recurrent Neural Networks (RNNs), with implications for topical tasks such as reasoning. Thus, there has been renewed interest in reintroducing recurrent dynamics to large-scale models while avoiding whole-sequence iterative simulation. State-of-the-art neuromorphic architectures have also seen a shift towards parallelisation, at the expense of expressivity. Nevertheless, it is established that diverse recurrent dynamics support brain computations. If these dynamics could be parallelised and, thus, scaled, they could be beneficial in machine learning applications as well. For example, neuronal bistability can support long-range temporal modelling, while intricate spiking patterns can provide effective and sparse (efficient) communication between network blocks. Accordingly, this study aims to investigate the parallelisation of the bistable regime of the FitzHugh-Nagumo model.

The proposed approach to parallelising these proof-of-concept dynamics is based on the linear interpolation of SSM and Dilated RNN states during input-driven transient dynamics. While SSMs can be efficiently implemented using parallel scans, Dilated RNNs, which take the generic form $h_t = \varphi(h_{t-k} + x_t)$, reduce the number of sequential steps by a factor of k. Preliminary results show that, individually, SSMs cannot model bistability, and Dilated RNNs fail to train for relatively small dilation factors. However, linear interpolation appears to

increase the dilation range for which models can learn and even speed up learning over baseline RNNs. While this accelerates neural simulations, it also paves the way to using more complex neural dynamics in machine learning tasks.

37. Rescorla-Wagner Update Rule Predicts Metacognitive Bias

Lindersson C, Sax AL, Baddaley R, Ponte Costa R

Metacognition is the ability to reflect on and evaluate one's own cognitive processes and is negatively biased in depression. Prior work has shown that depressive symptoms are linked to lower confidence despite intact performance, and to greater sensitivity to new evidence when updating confidence. These findings suggest that metacognitive biases in depression may arise from a distorted updating mechanism. To test this, we examine whether trial-by-trial confidence ratings are better explained by autoregressive learning models, which explicitly capture updating dynamics, or by static linear models. We also ask whether the model parameters predict individuals' metacognitive bias. In each trial, participants reported how many dots they saw, how confident they were about their estimate, and were then given performance feedback. The participants completed three sessions, where each session used either accurate, falsely positive, or falsely negative performance feedback. Confidence ratings varied significantly across feedback conditions: falsely positive feedback elicited the highest confidence, followed by neutral feedback, with falsely negative feedback leading to the lowest confidence. While both task performance and previous feedback predicted confidence, performance remained consistent across the feedback conditions. Cross-validation of model fits across a battery of linear models and autoregressive models revealed that the best fitting model used a Rescorla-Wagner update rule which aligned feedback with a weighted sum of previous feedback and the current performance. The learning rate, feedback weight, and performance weight predicted participant's metacognitive bias which in turn predicted depression scores. These preliminary results support that confidence is better captured by an autoregressive learning model than by static linear models, and that the metacognitive bias in depression may stem from altered learning mechanisms.

Poster session 2 (Friday 11th)

1. Synaptic depression outperforms potentiation in learned stimulus discrimination under relative integration of opposing outputs

Abdelrahman NY*, Greenin-Whitehead K*, Jung JS*, Tan MW*, Jiang J, Yamada D, Hige T, van Rossum MCW, Lin AC *equal contribution, alphabetical order

Are the brain's circuit architectures and synaptic plasticity rules in some sense 'optimal'? If so, in what sense, or in what contexts? We address these questions using olfactory associative memory in the fruit fly Drosophila. Flies can learn to associate a particular odour with a reward (e.g., food) or punishment (e.g., shock) and thereafter approach or avoid the trained odour. These associative memories are stored in Kenyon cells in the mushroom body, by weakening synapses from odour-responsive Kenyon cells onto mushroom body output neurons (MBONs) that lead to incorrect actions (e.g., odour+punishment weakens KC->Approach synapses). Why weaken incorrect actions rather than strengthening correct actions? Notably, synaptic depression is also used for learning in the vertebrate cerebellum, which has a remarkably similar architecture to the insect mushroom body, suggesting that using depression may be functionally advantageous.

We show both analytically and using simulations that depression outperforms potentiation for discriminating odours with overlapping KC representations, under a particular condition: if behaviour depends on the relative, not the absolute, difference between Avoid vs. Approach MBON activities. To test whether behaviour depends on the relative difference, we measured aversive learning for a range of odour concentrations and punishment intensities, in an individual-fly T-maze. We automatically tracked the flies' decisions to enter or leave the side with the punished odour, and from the statistical distributions of these stochastic decisions, we inferred the mean and variance of the flies' underlying preference for/against the odour. Bayesian modelling indicated that the data best fit a model where behaviour depends on the relative, not the absolute, difference between Avoid and Approach. These results suggest that flies learn by synaptic depression because, in the mushroom body, it is computationally superior to synaptic potentiation.

2. Mechanisms of Memory Retrieval in a CA1 Hippocampal Microcircuit

Andreakos N, Shigang Y, Cutsuridis V.

Theta rhythm in the hippocampus has been linked to memory formation, encoding and retrieval. In this study, we extended a computational model of region CA1 of the rat hippocampus to decipher the biophysical mechanisms of memory retrieval. Our model incorporated a variety of theta-modulated excitatory and inhibitory interneurons waxing and waning at different phases of theta. We scaled feedforward and feedback excitatory and inhibitory synapses in excitatory and inhibitory cells in the microcircuit model to test how memory recall performance (RP) is affected.

Results showed that scaled feedforward excitatory synapses in bistratified cells were the most effective way to reduce spurious activity and improve recall quality. RP decreased when more memory engram cells were involved in decoding more complex memories, and when the interference between memories was increased. Greater context complexity reduced RP particularly when stored memories were recalled in less familiar contexts. On the contrary, the total number of stored memories and network size had minimal impact on RP, with a slight improvement observed only with larger networks. Reducing inhibitory synapses in proximal dendrites impaired RP, while the absence of distal inhibition had no effect.

These findings offer insights into how synaptic, cellular, and network mechanisms support memory retrieval. Efficient retrieval is essential for decision-making, and understanding these mechanisms could lead to better strategies for improving memory and treating memory-related disorders such as Alzheimer's.

Acknowledgement: This work was supported by the European Union's Horizon 2020 Research and Innovation Programme under the Marie-Sklodowska Curie grant no778062 ULTRACEPT.

3. A half-centre oscillator encodes sleep need in the Drosophila brain

Hasenhuetl PS, Sarnataro R, Vrontou E, Rorsman HO, Talbot CB, Brain R, Miesenböck G (PSH and RS contributed equally)

Homeostatic regulation of sleep and wakefulness is essential for survival, but the underlying mechanisms are poorly understood. In Drosophila, a central role in sleep homeostasis is attributed to neurons projecting to the dorsal fan-shaped body (dFB), which store sleep need via their intrinsic excitability. Here, we test if and how dFB neurons incorporate this intrinsic account of sleep drive into their network activity. In vivo dFB neuron ensemble activity is rhythmic, with a peak at 0.2-1 Hz (measured via calcium imaging). The amplitude of these slow oscillations increases with sleep drive, and their optogenetic replay using a dFB neuron-restricted driver line promotes sleep. While this corroborates analogies to mammalian slow-wave sleep, our experiments also reveal an important distinction: dFB neurons receive and can modulate arousal signals (via mutually-inhibitory connections with arousal-mediating dopamine neurons), but their oscillations persist during arousal. Moreover, slow-wave activity in the dFB neuron ensemble originates from a subset of dFB neurons, which are anticorrelated between hemispheres and rely on mutually-inhibitory connections. Accordingly, blocking dFB neuronal output disrupts oscillations and sleep. Lastly, these output synapses undergo homeostatic plasticity upon sleep deprivation, which might be used to maintain robust dFB neuron network activity. Our experiments provide direct functional evidence for the involvement of dFB neural activity in the control of sleep in Drosophila and reveal that these neurons harbour a half-centre oscillator that generates a neural code for sleep need.

4. Computational mechanisms underlying how humans adapt their choices to the average effort of the environment

Scholey EV, Mehta NM, Apps MAJ

When you move to a new city, how do you decide which route to take to work? Should you go via the steep path in front of you, or should you spend a bit longer to go via an easier route? To make these decisions effectively, we have to learn whether it's 'worth it' to exert ourselves given how effortful the alternative options in the environment are, otherwise known as the opportunity costs. Reinforcement learning based models have captured the computational mechanisms underlying how the average reward rate influences opportunity costs. But a similar account of the opportunity costs of effort is absent. Here, we offer a computational account of how effort-based decisions are modulated by opportunity costs (as defined by the average effort rate in an environment), and how this average effort is learned. Human participants performed a novel sequential task, where they decided whether to spend time exerting different levels of physical effort for reward, or forgo the current option and wait for potentially less effortful alternatives. Participants completed this task in both easy environments (a high proportion of low effort offers) and hard environments (a high proportion of high effort offers). Across three studies, we found that participants were less willing to exert moderate effort in easier environments with higher opportunity costs. Model comparison revealed that a time-based reinforcement learning model that tracked the average effort rate could best capture these effects. Preliminary fMRI results suggest that fronto-basal ganglia circuits may be involved in tracking both offer values and environmental statistics related to effort demands. These findings provide a possible computational framework for a previously unexplored but fundamental aspect of effort-based decisions.

5. Autapses improve sensitivity to dynamic inputs in biological neurons.

McSweeney L, Vasilaki E, Toutounji H.

Neurons communicate by releasing neurotransmitters across synapses which influence the activity of downstream neurons. Synapses between a neuron and itself have been found in high proportions in specific brain regions and neuron types. These so-called 'autapses' give the neuron a delayed self-feedback after every action potential. Why these types of synapses exist, and what they might contribute to neural dynamics and computation is largely unknown. Previous theoretical studies investigating the neurocomputational properties of autapses have used neuronal models which aren't known to produce autapses. In addition, they used biologically implausible autaptic delays on the order of tens of milliseconds. Here we develop a simple model of autaptic feedback using biologically plausible parameters and implemented in a spiking neuron model fit to electrophysiological data. By constraining our simulations to biologically plausible parameters, we begin to understand what role autapses may play in biological computations. Our autaptic neuron model can induce surprising changes to the neuron's firing behaviour. Specifically, we find through bifurcation analysis that autapses induce a transition towards bursting which may have implications for information transfer and network synchronization. Additionally, as revealed through inspecting their responses to injected currents, autapses improve neuronal sensitivity to time-varying inputs. Our work is a first step in understanding how autaptic neurons influence computations within heterogeneous neural networks made up of neurons with a variety of firing properties

contributing to a deeper understanding of biological neural networks, and the improved design and efficiency of artificial neural networks.

6. Theoretical Frameworks for Credit Assignment via Behavioral Timescale Synaptic Plasticity

Cone I, Clopath C, Costa RP

Behavioral Timescale Synaptic Plasticity (BTSP) is a type of synaptic plasticity in which large, Ca2+ "plateau potentials" in the dendrites of hippocampal pyramidal cells drive the formation of hippocampal fields. BTSP differentiates itself from typical forms of plasticity by a) learning correlations of the timescales of seconds and b) rapidly changing network activity in only a few trials. These two key features of BTSP run counter to the traditional dynamics of learning in neural networks, whereby small, spatiotemporally precise changes gradually accumulate to optimize network activity and function. We propose a generalized BTSP rule (gBTSP) and examine how it can be used to learn in networks, in both unsupervised and supervised contexts. Further, since BTSP has been observed to trigger plasticity in inputs that are both feedforward (CA3→CA1) and recurrent in nature (CA3→CA3), we examine the function of gBTSP in both feed-forward and recurrently connected networks. We show that unsupervised learning with gBTSP leads to competitive learning, learning place field maps (in the feed-forward case), and attractive memory networks (in the recurrent case). For the supervised case, we derive how a plateau event reduces the task error, allowing us to find "when" and "where" a plateau event should occur to optimize a given loss. We show that networks learning via supervised gBTSP can solve tasks such as trajectory matching or delayed-non-match-to-sample. However, we find that rapid, BTSP-mediated activity changes become more difficult to recapitulate with increasing network depth/recurrence. We propose that certain conditions on neural architectures and tasks are required to support rapid, few-shot changes in neural activity. It remains an open question how complex, highly recurrent models of population-level hippocampal learning (i.e. the formation of "cognitive maps") might be constructed such that they support the few-shot learning observed via BTSP.

7. Dopamine D1 receptor expression in prefrontal parvalbumin neurons increases distractibility in marmosets compared with macaques

Ivanov TG, Joyce MKP, Krienen FM, Mitchell JF, Ma S, Inoue W, Nandy AS, Datta D, Duque A, Arellano J, Gupta R, Gonzalez-Burgos G, Lewis DA, Sestan N, McCarroll SA, Martinez-Trujillo J, Froudist-Walsh S, Arnsten AFT

Marmosets and macaques are common non-human primate models of cognition, yet marmosets appear more distractible and perform worse in cognitive tasks. The dorsolateral prefrontal cortex (dlPFC) is pivotal for sustained attention, and prior research in macaques suggests that dopaminergic modulation and inhibitory parvalbumin (PV) neurons could influence distractor resistance. Here we compare the two species using a visual fixation task with distractors, perform molecular and anatomical analyses in dlPFC, and link functional microcircuitry with cognitive performance using computational modeling. We show that marmosets are more distractible than macaques, and that marmoset dlPFC PV neurons contain higher levels of dopamine D1 receptor (D1R) transcripts and protein, similar to levels in mice. Our modeling

indicates that higher D1R expression in marmoset dlPFC PV neurons may increase distractibility by making dlPFC microcircuits more vulnerable to disruptions of their task-related persistent activity, especially when dopamine is released in dlPFC in response to unexpected salient stimuli.

8. Communication versus computation: The hidden costs, shaping the brain's architecture

Fakhar K, Akarca D, Luppi A, Oldham S, Hadaeghi F, Vertes P, Hilgetag C, Astle D.

Large-scale human brain networks exhibit complex topological characteristics, likely reflecting a balance among competing objectives such as minimizing wiring cost and maximizing communication efficiency. Interestingly, computational modelling has suggested that the connectivity of the brain is biased towards enhanced communication rather than a minimized wiring cost. Yet, the relationship between such communication efficiency and the functional capacity of the brain, e.g., to solve computational problems, remains unclear. To address this question, we used a game-theoretical framework in which individual brain regions establish connections only if it improves their signalling efficiency, given the wiring cost. We show that, firstly, complex network architectures naturally emerge from these local interactions, capturing some hallmarks of the brain. Secondly, resulting networks have both superior communication and reduced wiring cost compared to empirical brain networks. However, these optimal networks exhibited diminished memory capacity relative to empirical networks. Our findings suggest that efficient communication does not necessarily translate to improved computation. Instead, functional capacity may have played an essential role in shaping brain network architecture, potentially even at the expense of communication efficiency.

9. Tracking the Brain's Path from Perception to Categorization of Mathematical Objects

Karami A, Debray S, Valerio D, Caute M, Dehaene S

The ability to recognize mathematical objects as belonging to specific categories—such as identifying "one" and "two" as integers, or "triangle" and "rectangle" as geometric shapes—is a fundamental aspect of human cognition. In our study, we recorded brain activity using MEG while participants performed a non-mathematical stimulus detection task. During the task, they were presented with mathematical objects from three categories: integers, fractions, and geometric shapes.

These objects were shown in two formats: a verbal format (e.g., "one," "two," "rectangle") and a visual format (e.g., "1," "2," "□"). We used Representational Similarity Analysis (RSA) to model the contributions of low-level visual features and category-level representations. Our preliminary, time-resolved RSA results indicate that low-level visual features are processed early for both formats, while category-level information—reflecting semantic processing—emerges later in time.

10. Synaptic strength fluctuations from a model of stochastic gene expression in neurons

O'Donnell C, Veltz R, Senkevich O

Recent experimental results suggest that synapses fluctuate in size on time scales of hours-days, even in the absence of electrical activity. Understanding these fluctuations will be critical for understanding how information is stored in the brain. However, the source and general properties of these fluctuations are unknown. In this study, we show how such fluctuations can be explained by a simple stochastic model of gene expression in a single neuron that includes gene activation/deactivation, mRNA production/degradation and transport, and protein production/degradation and transport.

We characterised this model both analytically and via detailed numerical simulations and found several interesting phenomena: 1) synaptic protein fluctuations were highly correlated, both spatially – across different neuronal compartments – and across time; 2) there are multiple timescales of fluctuations, with different spatial eigenmodes corresponding to different processes in the model; 3) the apparent magnitude and correlation of the synaptic fluctuations strongly depend on the observation timescale; 4) due to resource sharing, synapses show heterosynaptic plasticity – plasticity at one synapse causes knock-on plasticity in its dendritic neighbours. Overall, this work provides evidence that stochastic gene expression may be the dominant driver of synaptic strength fluctuations in neurons and makes several concrete predictions for future experiments.

11. Long delays reduce the need for precise weights in spiking neural networks

Pengfei Sun, Jascha Achterberg, Dan F. M. Goodman, Danyal Akarca'

Recent work has shown that the performance of spiking neural networks (SNNs) on temporally complex tasks improves significantly when axonal delays are treated as learnable parameters. This raises an important question: If temporal delays improve a network's computational capacity, how precise do synaptic weights need to be? In this work, we investigate the relationship between delay-based computation and weight precision by combining quantized synaptic weights with a range of learnable delays on a challenging neuromorphic audio task. Our results reveal that short delays contribute little to performance, whereas medium to long delays are critical. Building on this insight, we introduce a learnable thresholding mechanism to suppress short delays that can be effectively compensated for by weights. These findings suggest that delays can reduce the burden on weight precision, highlighting a promising direction for energy-efficient SNN design and offering new perspectives on the role of delay in biological and neuromorphic computation.

12. Computational Roles of Higher Order Thalamocortical Feedback in Context-Dependent Sensory Processing

McCarthy PT, Bruno RM, Costa RP

Thalamic areas can be coarsley grouped into first and higher order nuclei. First order thalamic nuclei are driven by sensory organs and relay sensory information to primary sensory cortex. Higher order thalamic nuclei are driven by the primary area of their respective sensory

modality and send projections back to primary cortex as well as higher sensory areas, but the function of these thalamocortical connections is unknown. Despite many theories having been proposed, most are not computationally well-defined and little work has been done to operationalise them as mathematical models. Here, we model different computational variants of these connections and investigate their role in simple context-dependent image classification tasks. We find that in models with driving (additive) thalamocortical feedback, thalamic population representations are dominated by stimulus identity, optimised for sensory discrimination, whilst in models with modulatory (multiplicative) feedback, they are dominated by cue identity, optimised for context representation.

13. Category-biased patches encircle core domain-general regions in the human lateral prefrontal cortex

Assem M, Shashidhara S, Glasser M, Duncan J

The fine-grained functional organization of the human lateral prefrontal cortex (PFC) remains poorly understood. Previous fMRI studies delineated focal domain-general, or multiple-demand (MD), PFC areas that co-activate during diverse cognitively demanding tasks. While there is some evidence for category-selective (face and scene) patches, in human and non-human primate PFC, these have not been systematically assessed. Recent precision fMRI studies have also revealed sensory-biased PFC patches adjacent to MD regions. To investigate if this topographic arrangement extends to other domains, we analyzed two independent fMRI datasets (n = 449 and n = 37) utilizing the high-resolution multimodal MRI approaches of the Human Connectome Project (HCP). Both datasets included cognitive control tasks and stimuli spanning different categories: faces, places, tools and body parts. Contrasting each stimulus category against the remaining ones revealed focal interdigitated patches of activity located adjacent to core MD regions. The face and place results were robust, replicating across different executive tasks, experimental designs (block and event-related) and at the single subject level. In one dataset, where participants performed both category and sensory tasks, place patches overlapped with visually biased regions, while face patches were positioned between visual and auditory biases. Our results paint a refined view of the fine-grained functional organization of the PFC, revealing a recurring motif of interdigitated domain-specific and domain-general circuits. This organization offers new constraints for models of cognitive control, challenging traditional views that distinct broad PFC regions support separate cognitive control processes. We propose that through this distributed overlapping architecture, the PFC can dynamically integrate multiple streams of information, supporting the construction of flexible cognitive control models that guide adaptive behaviour.

14. Biological deep learning for simulating the role of SST Interneurons in cognitive resilience

Duisberg BFE, Marmouset-De La Taille I, Greedy W, Underwood S, Carlyle B & Costa RC

Two major snRNA-seq studies have shown that SST interneurons specifically are lost along the AD disease trajectory (Gabitto et al., 2024, Mathys et al., 2023). We have previously shown that levels of SST-14 peptide in the angular gyrus correlate with cognitive performance independent of AD pathology (Morgan & Carlyle, 2024). To explore why SST is a key

predictor of cognitive decline, we used a state-of-the-art cortical network model that approximates both deep learning and a large number of neuroscience studies (BurstCCN, Greedy et al. 2022).

We simulated progressive SST-cell degeneration by permanently removing SST-like inhibitory connections in BurstCCN, and measured "cognitive performance" as the network's ability to learn and generalize a digit classification task (MNIST). For comparison, we also assessed the impact of removing pyramidal-like excitatory connections.

The loss of SST-like connections led to a substantially greater reduction in test accuracy than the loss of pyramidal-like connections, highlighting the importance of SST-mediated inhibition in network function. This decline was primarily due to disrupted feedback input to the apical dendrites of pyramidal neurons, impairing hierarchical processing. Notably, our simulation results aligned with snRNA-seq data from post-mortem brains of AD patients, showing a correlation between SST cell abundance and cognitive resilience across individuals with varying levels of neuropathology.

Our findings suggest that SST interneuron loss has a disproportionately large impact on cognitive function, both in artificial and biological systems. The model predicts that this is due to the key role that SST interneurons play in propagating task-error signals throughout the cortex. These results not only support the hypothesis that SST interneurons play a critical role in cognitive resilience, but also provide deeper insight into Alzheimer's disease by linking cortical circuit physiology to cognitive dysfunction.

15. Overtraining Enhances Behavioural Flexibility on a Serial Reversal Learning Task: A Reinforcement Learning Perspective

Maggi S, Renstrom J, Grasmeder Allen R, Juty J and Bast T

Behavioural flexibility—the ability to adaptively shift behaviour when environmental contingencies change—is essential for survival and impaired in various psychiatric conditions. Flexibility of a response is typically thought to decrease with overtraining. Contrasting with this idea, reversal learning, which requires switching responses when reward contingencies are reversed, can be enhanced by overtraining, which is referred to as the overtraining reversal effect (ORE).

We investigated this phenomenon in a lever-press reversal learning paradigm with male Lister hooded rats, comparing normally trained (NT) rats that reversed immediately after reaching criterion with overtrained (OT) rats receiving 100 additional trials before reversal across ten serial reversals.

OT rats showed enhanced flexibility in later reversals with reduced trials to criterion. OT rats also exhibited increased perseverative errors but decreased regressive errors. A Bayesian trial-by-trial strategy analysis revealed that OT rats reduced engagement with task-irrelevant cues and achieved criterion through rapid shifts in response strategy rather than gradual evidence accumulation seen in NT rats.

To interpret these behavioural patterns, we built a reinforcement learning model with dynamic learning rates (α), exploitation parameters (β), and meta-learning rates (η). Parameter fitting

revealed that OT rats displayed higher α, β, and η than NT rats. This explains the seemingly contradictory findings: higher β accounts for greater perseveration when committed to a response strategy, while higher α enables rapid value update when feedback contradicts expectations, facilitating abrupt strategy shift.

Additional analysis revealed that the ratio between overtraining duration and initial learning requirements determines whether overtraining enhances or impairs flexibility, potentially establishing a quantifiable boundary between these behavioural states and resolving contradictions in the literature.

16. Semantic Enrichment of Episodic Memories through Bidirectional Replay

Albesa-González A, Clopath C

Traditional complementary learning systems (CLS) models highlight a predominantly unidirectional information flow from a fast-learning episodic system (medial temporal lobe, MTL) to a slower-learning semantic system (neocortex, CTX). However, episodic memories are not only sensorial, and instead embed a substantial semantic understanding of the episode. Furthermore, MTL regions such as the entorhinal cortex and the hippocampus have been shown to encode semantic representations of the environment, as evidenced by concept cells in humans or splitter cells in rodents. Yet, existing models typically overlook how semantic knowledge could reciprocally influence episodic memory encoding.

Here, we introduce a model of cortical-subcortical interactions that addresses these limitations by implementing, on top of bidirectional communication, bidirectional learning. Our framework explicitly models a three-phase learning loop: initial semantic abstraction from sensory episodes, semantic reorganization of MTL (subcortical) representations based on cortical replay, and semantic refinement via replay of sensory-semantic episodes. Crucially, our model distinguishes two different subcortical subpopulations—dense (sensory-based) and sparse (semantic-specific)—facilitating the dynamic restructuring of episodic encodings based on semantic information. Notably, this is achieved using only Hebbian and homeostatic plasticity and no input-target distinction, in contrast with models employing Contrastive Hebbian Learning or backpropagation.

The resulting semantic-enriched episodic encodings significantly enhance generalization and compositional capabilities, allowing efficient encoding of never-seen episodes composed of familiar semantic elements. Our model extends classic CLS theories by elucidating the critical role of reciprocal semantic-episodic interactions in fostering flexible and robust memory systems.

17. Emergence of Auditory Receptive Fields Based on Surprise at Multiple Timescales

Yashaswini, Dash Sneha, Bandyopadhyay Sharba

Understanding how sensory systems efficiently encode natural stimuli is a fundamental challenge in neuroscience. While the efficient coding hypothesis explains many aspects of sensory processing, its role in processing surprising auditory inputs remains unclear. This study

investigates how auditory receptive fields emerge through synaptic adaptation to surprise across timescales.

We present a computational framework modeling the development of auditory neural receptive fields via unsupervised learning. A single-layer network's synaptic weights adapt to auditory inputs to maximize activations for surprising events while minimizing overall activity. The weights are adjusted using three factors (α, β, γ) and the gradient of the L1-norm of activations. An autoregressive generative model (MelNet), trained on LibriSpeech, provides the joint probability distribution to calculate surprise, defined as the negative log probability of time-frequency bin energy conditioned on previous time steps and other frequency channels. We find learning to be fast, with robust convergence of weights using random speech samples.

The model learns tuning characteristics similar to auditory cortical neurons. Low and high saturation thresholds are derived from the distribution of surprise values: when surprise falls between these thresholds, weights increase by α, mimicking Long-Term Potentiation, while low surprise reduces weights via β, akin to synaptic depression. The γ factor ensures neural stability. To generalize across loudness (20–100 dB SPL), cochleagrams from intensity-scaled waveforms were input to MelNet for surprise calculation. The updated network weights result in biologically plausible frequency response areas-narrowly tuned at low SPLs and broader at high SPLs. This framework shows how efficient coding and surprise signals together shape biologically plausible receptive fields without supervision, revealing a possible fundamental principle for sensory representation.

18. Modeling Neural Adaptation to Tendon Transfer Surgery

Kawakita G, Gallego J

Motor adaptation has been mostly studied in the context of learning to counteract changes in the environment or external forces applied to the body. In comparison, the challenge of adapting to changes in the neuromusculoskeletal system remains poorly understood.

Previous studies investigated this question through tendon transfer (TT) surgeries in non-human primates. A recent study showed that multiple processes underpin regaining limb control following TT surgery in monkeys (Philipp et al., 2024). In experiments swapping finger flexor and extensor muscles, an initial fast stage—marked by a rapid swap of muscle activity profiles—is followed by a slower fine-tuning stage that ultimately leads to a reversion to the original muscle activity patterns. The results suggest that different learning processes with distinct timescales underpin adaptation to changes in the neuromusculoskeletal system. However, the underlying computational mechanisms remain unclear.

We used musculoskeletal modeling and reinforcement learning (RL) to investigate the mechanisms underlying TT adaptation. Using MyoSuite, a biologically accurate musculoskeletal modeling framework, we studied how agents learned to counteract the effects of TT surgeries in a hand model performing motor tasks. We applied a TT surgery that swaps flexor and extensor tendons after the initial task learning and studied how the agent adapted.

Our next step will be to train RL agents using both model-free and model-based algorithms. We hypothesize that while both methods will lead to learning, only the model-based approach may exhibit a two-stage adaptation process, mirroring the fast-slow dynamics observed in the previous study. By comparing the adaptation dynamics of these agents to the changes in behavior and muscle activity observed in the previous experimental studies, our work aims to elucidate the algorithms that underpin learning to adapt to persistent changes in the neuromusculoskeletal system.

19. Adaptive learning in cortical networks: a computational model of Cholinergic neuromodulation

Filipovica M, Kermani Nejad K, Greedy W, Zhu HW, Mellor J, Ponte Costa R

The cholinergic system plays a crucial role in learning, memory, attention, and adaptive plasticity. Acetylcholine (ACh) has been linked to experience-dependent plasticity and the expansion of receptive fields for behaviorally relevant stimuli. Its importance for maintaining cognitive functions is also reflected in its association with cognitive decline and neurodegenerative disorders, such as Alzheimer's Disease (AD), where the degeneration of the cholinergic system is a key pathological marker. We use a computational model to investigate whether some of the cognitive deficits observed in AD could stem from impaired cholinergic modulation and the resulting disruptions in plasticity and neural representations. Using biologically plausible neural network models and adaptive learning algorithms, the proposed framework modulates learning based on the history of error signals in cortical ensembles. The model further examines the spatial specificity of cholinergic modulation, testing whether local, circuit-level modulation is necessary to support both rapid learning and resilience to neurodegeneration. Results show that locally modulated networks form sparser, more selective and more robust representations, while global or non-adaptive modulation leads to more fragile representations that are more likely to be disrupted by neurodegeneration. Together, these findings provide a computational perspective on the role of cholinergic modulation in cortex, suggesting that disruptions to the cholinergic system could impair both learning efficiency and network robustness, offering a novel theoretical link between cholinergic function and cognitive decline in AD.

20. Balancing external-internal learning: a theory of serotonin

Iris Marmouset-de la Taille, Simon J. B. Butt, Florencia Iacaruso, Rui Ponte Costa

Despite research demonstrating the involvement of serotonin in learning processes, its precise role remains unclear due to seemingly contradictory results of the effects of this neuromodulator on behavioral inhibition, cognitive flexibility, as well as reward and aversive learning. Here, we propose an integrated framework describing serotonin as a controller determining the relative weight given to learning driven by internal and external processes. We hypothesize that, through its diverse receptor types and brain-wide projections, serotonin gates out external input to promote internally-driven learning, that is, the reorganization of existing mental models of the world through associative and generative processes. To test this

hypothesis, we developed a computational model incorporating serotonin as a gating factor between current sensory input and memory in a cortico-striatal network driving motor-related decisions during a two-armed bandit task. Our self-adapting model tracks the dynamics of serotonin signaling throughout different task phases, revealing how changes in contingencies lead to a shift of the model's reliance on current input and internal memory representations. Moreover, by manipulating our model we can test if artificially altering serotonin levels shifts the learning strategies, and the behavioral outcomes associated with them. By reconceptualizing the role of serotonin in learning, we open new avenues for understanding and treating disorders involving learning disruption and serotonergic dysfunction, particularly those characterized by imbalances between internally and externally guided brain states such as depression or schizophrenia.

21. Balancing stability and plasticity through gain neuromodulation and inhibitory gating

Rodriguez-Garcia A, Ramaswamy S

Biological and artificial learning systems alike confront the plasticity–stability dilemma. In the brain, neuromodulators such as acetylcholine and noradrenaline relieve this tension by tuning neuronal gain and inhibitory gating, balancing segregation and integration of circuits. Fed by dense cholinergic and noradrenergic projections from the ascending arousal system, layer-5 pyramidal neurons in the cerebral cortex offer a relevant substrate for understanding these dynamics. When distal dendritic signals coincide with back-propagating action potentials, calcium plateaus turn a single somatic spike into a high-gain burst, and interneuron inhibition sculpts the output. These properties make layer-5 cells gain-tunable amplifiers that translate neuromodulatory cues into flexible cortical activity. To capture this mechanism we developed a two-compartment Izhikevich model for pyramidal neurons and single-compartment somatostatin (SOM) and parvalbumin (PV) interneurons, linked by Gaussian connectivity and spike-timing-dependent plasticity (STDP). The soma and apical dendrite are so coupled that somatic spikes back-propagate, while dendritic plateaus can switch the soma from regular firing to bursting by shifting reset and adaptation variables. We show that stronger dendritic drive or tighter coupling raise gain by increasing the likelihood of calcium-triggered somatic bursts. In contrast, dendritic-targeted inhibition suppresses gain, while somatic-targeted inhibition raises the firing threshold of neighboring neurons, thus gating neurons output. Notably, bursting accelerates STDP, supporting rapid synaptic reconfiguration and flexibility. When translated to artificial neural systems, this suggests that brief gain pulses driven by neuromodulators could serve as an adaptive two-timescale optimization mechanism, effectively decomposing weights into fast and slow components to enable flexible learning with reduced interference during continual learning.

22. Selective inhibitory neuronal contributions to local field potentials in flexible perceptual decisions: A mean-field model prediction

Azimi A, Lenfesty B, Wong-Lin KF

Previous experimental and theoretical studies on perceptual decision-making demonstrate that local selective inhibitory neurons in parietal cortex are crucial for circuit stabilization and

enhancing competition between pro- and anti-evidence neuronal populations (Najafi et al., 2020; Roach et al., 2023). Additionally, parietal cortical local field potentials (LFPs) reflect the summed pro- and anti-evidence signals (Bollimunta and Ditterich, 2012). However, these studies used simple tasks, leaving LFP computation in more complex, flexible perceptual decision tasks unclear.

To address this, we developed two mean-field computational models: one with implicit non-selective inhibitory populations (an extension of Wong and Wang, 2006) and another with explicit selective inhibitory populations in a perceptual decision context abstracted from motor effectors (Shushruth et al., 2022). Each model includes a neuronal group for evidence integration and another for action selection. In the task, non-human primates discriminated visual motion direction (leftward or rightward) and, after a delay, reported decisions via saccades to one of two randomly presented targets (blue for leftward; yellow for rightward). Higher motion coherence (percentage of dots moving coherently) yielded easier decisions.

Models were fitted to monkey's choice accuracy and response time (go-RT), and simulated firing-rate activity during action selection matched experimental data (Shushruth et al., 2022). The implicit non-selective inhibitory model showed increased LFP activity with higher motion coherence, whereas the explicit selective inhibitory model exhibited the opposite trend, with lower LFP activity at higher coherence. Moreover, only the explicit model showed a distinct dip in activity at choice target onset. These results predict clear differences in LFP signatures for decisions abstracted from effectors. Future work will compare these predictions with empirical data.

23. Data-driven biophysically detailed computational modeling of neuronal circuits with the NeuroML standard and software ecosystem

Sinha A, Gleeson P, Ponzi A, Ray S, Panagiotou S, Marin B, Silver RA

Computational models are essential for integrating multiscale experimental data into unified theories and generating new testable hypotheses. Realistic models that include biological intricacies of neurons are critical tools for gaining a mechanistic understanding of neuronal processes. Their complexity and the disjointed landscape of software for computational neuroscience, however, makes model construction, fitting, simulation, and re-use and dissemination a considerable challenge. Here, we present NeuroML and show that it accelerates modelling workflows and promotes FAIR (Findable, Accessible, Interoperable, Reusable) and open computational neuroscience [1].

NeuroML is an established standardised language that provides a simulator independent model representation and accompanying ecosystem of compliant tools that support all stages of the model life cycle: creating, validating, visualising and analysing, simulating, optimising, sharing, and re-using models. It provides a curated set of building blocks for constructing new models and remains extensible by allowing the definition of new model entities in the Low Entropy Modelling Specification (LEMS) language [2]. Further, NeuroML also links to other neuroscience initiatives (PyNN, SONATA[3]), systems biology standards (SBML, SED-ML) and machine learning/AI formats (Model Description Format [4]) to promote interoperability. Finally, a large archive of published standardised models supports re-use of existing models.

We demonstrate how NeuroML supports the model life cycle by presenting a number of published NeuroML models in different species (C. elegans, rodents, humans) and different brain regions (cortex, cerebellum), highlighting their scientific contributions. We also list resources on using NeuroML and existing models.

References
https://doi.org/10.7554/eLife.95135
https://doi.org/10.3389/fninf.2014.00079
https://doi.org/10.1371/journal.pcbi.1007696
https://doi.org/10.1007/s10827-024-00871-5

24. Structure-function relationships in connectome-based echo state networks

McAllister J, Houghton C, Wade J, O'Donnell C

The Echo State Network (ESN) framework is an efficient recurrent neural network paradigm of importance for both neuroscience and machine learning. It is quick and efficient to train, and also has been suggested as a possible model of brain function. It is not fully known how network structure influences ESN functionality, dynamics, and robustness. We used biological networks to study this, compared with randomly initialised ESNs.

In both biological and artificial neural network contexts, neurons and clusters demonstrate functional specificity. We asked if synapse-resolution connectome-based ESNs demonstrate functional specificity/generality at the neural level. We used the larval Drosophila melanogaster connectome, which exhibits a hierarchical modular structure according to type, class and function. We built connectome-based ESNs from this and compared neural specificity metrics between connectome and equivalent random networks across tasks in memory, decision-making, and time-series prediction. Connectome ESNs contain smaller subsets of task-selective neurons, while random networks exhibit more distributed, “general” groups. We tested the interpretation of these metrics by systematically pruning nodes based on their measure of engagement, finding that connectome ESNs maintain performance more robustly across pruning. Finally, we investigated structural features of the networks, uncovering correlations between task-relevance and characteristics such as recurrence, node degree, and biological cell-type annotations. We find that correlations are consistent across connectome and conventional ESNs, but that connectome ESNs exhibit stronger correlations.

These findings indicate that biologically-inspired connectivity can enable sparsely selective and compact neural networks, which may reduce energy consumption and optimise robustness. Relationships between node features and task engagement suggest a way to initialise better performing, more efficient and robust ESNs.

25. Dimensional Reconfiguration of Functional Network Manifolds Underlies Working Memory Performance in Neurofibromatosis Type 1

Khanbeigi AA, Litwińczuk MC, Garg S, Parkes LM, Madadi-Asl M, Lea-Carnall CA

Introduction

Neurofibromatosis type-1(NF1) is a genetic condition linked to working memory deficits, potentially due to disrupted network dynamics [1]. Insights from dynamical systems theory can shed light on network dynamics. It has been shown that system states evolve toward an "attractor manifold," whose intrinsic dimension (ID) reflects system dynamics [2]. Here, we use local PCA to estimate the intrinsic dimension of the network manifolds in working memory fMRI data from adolescents with NF1.

Methods

fMRI data were acquired (3T Philips Achieva, TR = 2 s, TE = 12/35 ms) from 44 NF1 participants during 6 minutes of working memory task. Time series from 300 cortical ROIs (averaged over voxels of ROI, Schaefer-300 parcellation [3]) were assigned to 7 networks, split by hemisphere [4]. Preprocessing (SPM12) included dual-echo extraction and averaging (DEToolbox), slice-time correction and realignment to the first image, motion correction (ART), segmentation (DARTEL), MNI normalization, and denoising (Conn toolbox). ID was computed for 14 networks and compared with accuracy and response time of working memory performance.

Results and Conclusion

We show ID in LH Dorsal Attention Network is strongly correlated with better performance in working memory. The higher ID indicates more diverse patterns of activation and locally less constrained state space of fMRI activity.

[1] Ibrahim, Amira FA, et al. "Spatial working memory in neurofibromatosis 1: Altered neural activity and functional connectivity." NeuroImage: Clinical 15 (2017): 801-811.
[2] Dudkowski, Dawid, et al. "Hidden attractors in dynamical systems." Physics Reports 637 (2016): 1-50.
[3] Schaefer, Alexander, et al. "Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI." Cerebral cortex 28.9 (2018): 3095-3114.
[4] Yeo, BT Thomas, et al. "The organization of the human cerebral cortex estimated by intrinsic functional connectivity." Journal of neurophysiology (2011).

26. Layer-specific input-output functions can shape representational geometry

Willard KJ, Ponte Costa R, Bruno RM

The neocortex is divided into distinct layers, where the cells in each layer have unique physiological properties. One of these properties is the input-output (IO) function, or f-I curve, which characterises the firing rate of a neuron as a function of injected current. These IO functions vary across cortical layers. Cells in superficial layer L2/3 fire only at higher current levels and with lower rates than cells in deeper layers including L5/6 (Zhao et al., 2016). Here, we ask how the cellular IO function impacts representations at the network level within cortical layers. First, we confirmed the layer-specific differences previously observed in the IO function by performing in vivo whole-cell patch-clamp in the barrel cortex of mice. Now, we

are using abstractions of these recorded IO functions as activation functions in computational models. This allows us to test how these biologically-inspired activation functions change the network representations using geometrical methods (e.g., Alleman et al., 2024). The models are trained to group noisy patterns following simple classification tasks, approximating how mice discriminate shapes using patterns of whisker-mediated tactile inputs during barrel-cortex dependent sensory behaviour (Rodgers et al., 2021). Early results suggest that L5/6-like functions better represent the precise pattern of inputs, whereas L2/3-like functions represent global stimulus class. This difference in network representation seems to hold across simple models with a single hidden layer as well as more biological models with a structure that mimics cortical connectivity. Overall, we conclude that IO functions can impact network representations within cortical layers, which may allow the layers to specialise during sensory discrimination in vivo.

27. Mice wiggle a wheel to boost the salience of low visual contrast stimuli

Ghani, N. and The International Brain Laboratory

From the Welsh tidy mouse to the New York City pizza rat, movement reveals rodent intelligence. We show that head-fixed mice develop an active sensing strategy while performing a visual perceptual decision-making task (The International Brain Laboratory, 2021). Akin to humans shaking a computer mouse to find the cursor on a screen, we demonstrate that mice wiggle the wheel that controls the movement of a stimulus to boost low visual contrast salience. Moreover, mice wiggle the wheel at a frequency (11.5 $\pm$ 2.6 Hz) that maximizes temporal contrast sensitivity (Umino et al, 2018). With the "old method of watching and wondering about behavior," we reveal that mice exploit that it is easier to see something moving than something stationary by wiggling (Tinbergen, 1973).

28. Cortically-Embedded RNNs for integration of cortex-wide neuroscience data into recurrent neural network models

Sevenster E, Thrivikraman A, Davies G, Klatzmann U, Pedamonti D, Froudist-Walsh S,

Current state-of-the-art recurrent neural network models can capture complex neural dynamics during the performance of higher cognitive tasks. However, they largely overlook anatomy, limiting their ability to make species-specific and anatomically-precise predictions for experimentalists. Cortex-wide dynamical models increasingly integrate anatomical features including connectivity, dendritic spines and receptors, but are incapable of solving most cognitive tasks. Here, we introduce Cortically-Embedded Recurrent Neural Networks (CERNNs), which embed artificial neural networks into a species-specific cortical space, facilitating direct comparisons to empirical neuroscience data across the entire cortex and allowing the incorporation of biologically-inspired constraints. We trained CERNNs, with macaque or human anatomy, to perform multiple cognitive tasks (e.g. working memory, response inhibition). CERNNs were trained with different architectural constraints and biologically-inspired loss functions. We evaluated CERNNs on (1) task performance, (2) alignment of connectivity with the macaque mesoscopic connectome, and (3) task-evoked activity patterns. The best performing models penalised both long-distance connections and deviations from empirical spine density. These results suggest that distributed cognitive

networks may arise naturally as the brain attempts to solve complex tasks under wiring constraints with systematic gradients of single neuron properties. More broadly, CERNNs constitute a framework by which artificial neural networks can be integrated with cortex-wide neuroanatomy, physiology and imaging data to produce anatomically-specific testable hypotheses across species.

29. Cerebellum drives hippocampal networks for stable goal-driven navigation

Pedamonti D, Pemberton J, Costa RP

The hippocampus is critical for goal-directed spatial navigation, yet its interactions with other brain regions, such as the cerebellum, remain largely unexplored. We introduce a systems-level reinforcement learning model that integrates cerebellar and hippocampal modules to investigate how task-specific cerebellar signals influence spatial learning.

Comparing the cerebello-hippocampal model with a hippocampus-only baseline in navigation tasks shows that cerebellar input consistently improves goal-oriented learning by allowing the agent to follow shorter and smoother trajectories while stabilising emerging place and border cells during training. The same predictive input enables generalisation to novel starting locations and larger mazes. Lesioning the cerebellar module impairs learning during the test phase and produces irregular trajectories similar to those observed in biological systems with cerebellar deficits.

Our results indicate that task-specific predictions from the cerebellum accelerate and stabilise hippocampal computation, providing a mechanistic bridge between inter-regional dynamics and behavioural outcome in spatial navigation.

30. Distributional dopaminergic codes for economic decision-making

Salmasi M, Dolan R

Effective economic decision-making depends on the brain's capacity to integrate uncertainty about the state of the world and reward/loss outcomes. Experimental evidence suggests that the brain uses a representation of uncertainty to facilitate probabilistic decision-making. Various theories, such as expected utility and prospect theory, have been proposed to model the economic preferences of the agents. However, full comprehension of what specific preference and risk measures the brain employs has yet to be achieved. Moreover, while progress in neuroeconomics has been substantial, critical gaps persist in understanding how dopaminergic system encodes and processes preference and risk. In this study, we develop a framework to explore these issues. We hypothesize that reward distributions are encoded by distributed distributional codes (DDC), wherein distributions are represented by the expected values of a set of dopaminergic encoding functions. We first show that the existing recordings from VTA dopamine neurons of mice are in broad agreement with this hypothesis. We then propose a DDC network that, by learning and mapping the weighted sums of DDC values, can approximate a wide range of preference/risk measures, including expected utility, conditional value at risk (expected tail gain), and compound utility-risk measures. The network provides a biologically plausible implementation of economic decision processes in the dopaminergic

system, where preference measures can be easily computed by mapping parallel readouts of dopamine neurons. Overall, this framework enables inference of decision policies and preference measures during economic decision-making and offers a mechanism through which dopaminergic system can compute associated probabilistic measures.

31. Connectome-constrained recurrent neural networks

Rovný M, Akarca D, Achterberg J, Duncan J, Astle D

Recent advancements in computational neuroscience have focused on bridging the gap between artificial neural networks and biological neural systems. While many methods focus on individual neurons or small assemblies, approaches capturing dynamics at larger scales remain relatively rare. It is this scale where our project weighs in, introducing a novel method for modelling regional and whole-brain function based on the underlying structure.

Building upon the framework of spatially embedded recurrent neural networks (Achterberg & Akarca et al., 2023) we implement a flexible regularisation method for constraining the weights within artificial neural networks based on target topology. This allows for the replication of brain-like structural connectivity at varying levels of granularity: global measures (e.g. average path length, global transitivity), distribution comparisons (e.g. degree or edge weight distributions), and matrix distances (e.g. adjacency or communicability matrices).

Results demonstrate structural and functional divergence depending on target topology. Fully trained networks are capable not only of performing tasks, their weight matrices show clear differences in trajectories and endpoints across targets. This translates to structural differences between networks based on empirical connectomes, generated adjacency matrices, and individual graph theoretical measures. Importantly, the differences can be observed not only in the structural properties of the networks but also in their functional organisation, with networks achieving equivalent accuracy yet differing in localisation of task processing on functional readouts.

Our approach opens new avenues for investigating how select aspects of (brain) topology impact computational efficiency, potentially offering insights into the organisational principles underlying both artificial and biological neural systems.

32. Spike timing in vivo for identified transcriptomic cell types

Shinn M, Zhou B, Prankerd I, Maat C, Bourdenx M, Nicoloutsopoulos D, Tilbury R, Shuker P, Bugeon S, Harris KD

The brain contains an enormous diversity of neurons that differ widely in structure and function. Recent advances in transcriptomics show that neuron types defined by gene expression encapsulate differences in anatomy and physiology, ranging from morphology to connectivity motifs. But do these neurons play different roles in neural circuits? Despite many technical innovations, it is still difficult to record from and identify multiple types of neurons simultaneously in vivo. This is especially true when using electrophysiology to observe

individual spikes, which have limited means for neuron identification but are essential to understanding the flow of information through neural circuits.

Here, we present preliminary data of multiple simultaneously-recorded neurons with identified transcriptomic cell types at single-spike resolution. We record from layer 5 of the mouse visual cortex using Neuropixels probes while the mouse views a battery of visual stimuli. To identify neuron types, we perform two-photon imaging in the same tissue, and identify the cell types using in situ transcriptomics (CoppaFISH). Then, we link neurons between electrophysiology and two-photon imaging using functional fingerprinting with a novel stimulus for eliciting distinct visual responses ("zebra noise"). This allows individual spikes to be recorded in transcriptomically identified neurons.

Spike timing is critical to answer questions about neuronal oscillations, synchrony, and the hierarchical flow of information through a neural circuit. Understanding the role of different neuron types reveals how the modular architecture of the brain serves as a scaffold for neuronal computations.

33. Higher-order interactions reveal circuit motifs across brain regions and behavioural states

Rashid Shomali S, Rasuli SN, Shimazaki H, Sadeh S

Neuronal activities display pairwise and higher-order interactions (HOIs) that can illuminate underlying circuit connectivity. While modern neurotechnology enables extensive long-term recordings, analysing these datasets for HOIs presents significant computational and statistical challenges. We introduce a robust and efficient method for calculating cross-validated HOIs from large-scale spiking activities. Using this approach, we quantified statistically significant HOIs from large-scale recordings of spiking activity across multiple datasets. Our analysis revealed consistent patterns of HOIs across different animals and visual areas. By projecting HOIs onto the analytical plane of triple-wise versus pairwise interactions that links circuit motifs to interactions, we identified a specific region corresponding to excitatory inputs to pairs as the dominant circuit motif. Furthermore, we found that HOIs can distinguish between stationary and running states, suggesting a circuit motif involving lateral inhibition. These results were consistent across individual animals and different datasets.

By performing numerical simulations of spiking neurons and mathematical analysis, we confirmed that the suggested circuit motifs can indeed underlie the HOI patterns. Notably, balanced spiking networks with random connectivity and varying degrees of recurrent excitation and inhibition failed to reproduce the experimentally obtained HOI patterns. However, similar HOI patterns emerged when random balanced networks were equipped with neuronal clusters with a Mexican hat connectivity profile: stronger recurrent excitatory connectivity within clusters and stronger inhibitory connections between clusters.

Our work provides a systematic analysis of HOIs in diverse datasets and demonstrates that HOIs can effectively reveal the circuit motifs underlying neural dynamics across different brain regions and consequently uncover the state of animals.

34. Theta sweeps along the Papez's circuit

Ji Zilong, Burgess Neil

Recent findings demonstrate that grid cells in the medial entorhinal cortex (MEC) generate an internal positional signal that sweeps outward from an animal's current location into the surrounding environment. These theta sweeps, determined by directional activity in upstream head direction cells, alternate in a stereotyped fashion between leftward and rightward directions across successive local field potential (LFP) theta cycles (Vollan et al, Nature, 25). We have modelled the generation of theta sweeps along the Papez circuit, encompassing bidirectional sweeps in head direction cells (Ji et al, 25, Hippocampus), left-right sweeps in grid cells (Ji et al, 25, Current Biology), and forward-backward sweeps in place cells (Chu* & Ji* et al, 24, eLife). These models provide mechanistic accounts for many empirical observations and yield numerous testable predictions for future experimental validation.

Across all models, three common features emerge: continuous attractor-like dynamics, theta modulation, and firing rate adaptation. Based on this convergence, we predict that a neural system exhibiting these three properties will intrinsically generate theta sweeps. We further explore the functional role of theta sweeps and propose that they constitute a form of subliminal mental exploration, enabling rapid coverage of unvisited space and supporting the efficient encoding of novel environments into spatial memory (Marshall* & Ji* et al, 25, CCN). These internally generated theta sweeps are also readily integrated with top-down modulation, accounting for goal-directed patterns observed in planning and decision-making tasks (Ormond & O'Keefe, 22, Nature).

Together, this body of work delineates a comprehensive computational framework for theta sweeps along the Papez's circuit, advancing our understanding of spatial cognition and memory, and offering insights into how these processes may be disrupted in disorders such as Alzheimer's disease.

35. The impact of transcranial ultrasound induced heating on thalamic neurons and networks

Kour R, Jameel A, Smith J, Bain P, Nandi D, Jones B, Quest R, Gedroyc W, Borisyuk R, Yousif N

The ventral intermediate nucleus of the thalamus (VIM) is the typical target for ablation in the treatment of essential tremor. Recently, high-intensity focused ultrasound (FUS) has become a non-surgical approach to achieve such ablation by heating tissue. During a FUS procedure, test sonications are used to heat the tissue to sub-ablative temperatures to assess targeting accuracy and potential side effects. This effect of temperature increase on the VIM is less well understood, and therefore the occurrence of side effects cannot be managed. We use biophysical models of thalamic neurons and networks to investigate how incremental temperature increases, approaching ablative levels, influence neuronal activity.

Temperature values were informed by magnetic resonance thermography data collected during a clinical HIFU procedure, which showed target site temperatures reaching up to 62°C. We adapted existing neuron models to incorporate temperature-dependent changes in both

ionic channel gating and membrane capacitance. Simulations were performed using the NEURON environment, first examining one neuron with a single compartment, which has previously been shown to exhibit a range of firing frequencies with varying channel conductance. We also considered a network of two excitatory and two inhibitory thalamic neurons.

We found that neuronal firing becomes progressively suppressed as temperature increases, with complete silencing observed above 40°C, regardless of the state of the model. In the network model, despite individual variability in excitatory and inhibitory firing patterns, overall activity also ceased at elevated temperatures.

These results suggest that sub-ablative heating alone can significantly influence neural activity, supporting clinical observations of tremor reduction during test sonications. Further work will expand this network approach to explore downstream effects, potentially improving the safety and efficacy of ultrasound-based tremor therapies.

36. Cortico-basal ganglia dynamics underlying skilled locomotion

Esparza-Iaizzo M, Lazar I, Fortunato C, Safaie M, Gallego JA

Locomotion is a fundamental behaviour exhibited by all vertebrates across the animal kingdom. Despite seeming like an automatic behaviour, it requires the coordinated activation of muscles through spinal circuits (central pattern generators; CPGs) which are partly driven by supraspinal inputs. While extensive research has characterized CPGs, how cortical and subcortical structures coordinate their activity to modulate spinal circuits during skilled locomotion is not well defined.

To address this question, we designed a novel task in which head-fixed mice running on a large spherical treadmill counteract rapid, unpredictable perturbations, applied from 12 different positions. This allows us to elicit a broad range of behavioural responses, determined by the actuator's location and the timing of the perturbation relative to the animals' ongoing gait cycle. To study the role of cortico-basal pathways in sensorimotor integration, we are recording from limb-specific subregions of sensorimotor cortices, downstream basal ganglia projections, and relay centers in the motor thalamus projecting back to cortical areas, using Neuropixel probes. Simultaneously, we are tracking whole-body 3D kinematics via markerless pose estimation.

Our first findings show within-region differences in the covariation patterns underlying unperturbed running before and after the start of the perturbations, as latent dynamics obtained with Principal Component Analysis were highly distinguishable across conditions. This indicated a potential change in control policy. Moreover, responses that reflected perturbation direction were observed across the thalamo-cortico-basal ganglia regions at similar latencies. Future work will focus on defining sensorimotor transformation across these pathways. By analyzing the geometry and dynamics of neural population activity across regions and relating it to kinematic responses, we aim to describe the contribution of supraspinal centers to skilled locomotion.