UCL

2023 UCL Annual NeuroAI Conference

Sainsbury Wellcome Centre

Tuesday 28th November 2023

09:00 - 18:30

Website: https://www.ucl.ac.uk/research/domains/neuroscience/ucl-neuroai

Twitter: @UCL_Neuro_AI

Image credit: Miss Caroline Casey, Institute of Neurology.
Handmade Neural Fireworks. Confccal microscopy
UCL Doctoral School, Research Images as Art Competition entry 2017-2018

Welcome

It is my great pleasure to welcome you to the 4th annual UCL NeuroAI Conference.

The last decade has seen phenomenal advances in the fields of machine learning (e.g. deep learning, reinforcement learning, and AI). While these changes have already had considerable impact on most areas of science they hold a particular resonance for neuroscience. Crucially, AI shares a common lineage with neuroscience and fundamentally both machine learning and the brain employ similar computations to process and compress information. For these reasons AI provides a means to emulate neural functions and the circuits supporting them, providing insights to aid our understanding of the brain and cognition.

Equally, AI tools provide a means to discover, segment, and track distinct neural and behavioural states - yielding more efficient experiments and accelerating the pace of discovery. In turn, this understanding feeds back into the design of more effective AI architectures and models. Essentially, AI problems posed in neuroscience both require and inspire further advances in AI.

The UCL Neuro AI initiative was set up to provide a bridge between researchers working in these areas at UCL, in order to foster dialogue and exchange of ideas between these fields. Among its many activities (see page 3) the annual conference is an important opportunity for these communities to gather together, make new connections and benefit from new insights in one another's fields.

In line with that ethos, today you will hear speakers from industry and academia, from PhD students to Professors, covering topics spanning the scope of NeuroAI from LLMs to mouse behaviour. We hope you enjoy this meeting and are as inspired by the potential of this field as we are.

Finally, we would like to thank Tom Mrsic-Flogel of the Sainsbury Wellcome Centre and Maneesh Sahani of the Gatsby Computational Neuroscience Unit for allowing us to use this venue today. We would also like to thank the UCL Neuroscience Domain and, in particular, its chair, Trevor Smart, for their continued support of NeuroAI, as well as the team in UCL's Research Coordination Office, particularly Gurjit Matharu, who have worked so hard to make this event happen.

Professor Caswell Barry

Professor of Neuroscience and AI, UCL

Introduction to UCL NeuroAI

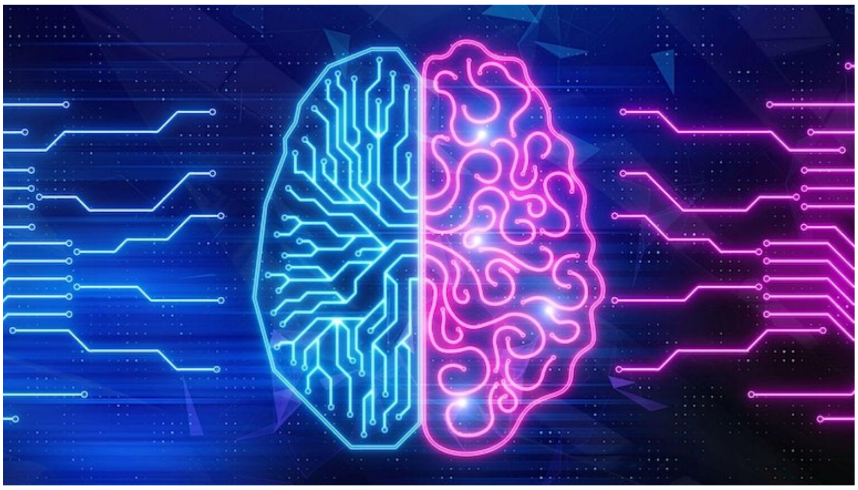

The UCL NeuroAI initiative aims to be a central hub at the university where researchers working in neuroscience, AI, machine learning and related fields can interact and stay updated on the latest advancements in these intersecting areas. Through seminars, meetings and training events, we aim to lower the barriers to the free exchange of ideas between researchers in these rapidly moving scientific fields.

UCL NeuroAI is a core part of the UCL Neuroscience Domain, which is an interdisciplinary research theme spanning across all faculties, departments, divisions and institutes at UCL, to bring together those working in all areas of experimental, clinical and theoretical neuroscience.

Other Activities

Monthly Seminar Series: We host a monthly online seminar series, featuring international speakers from AI and neuroscience, which aims to make leading-edge research accessible to as wide an audience as possible. Recordings of previous talks can also be found on our website (https://www.ucl.ac.uk/research/domains/neuroscience/ucl-neuroai).

NeuroAI Interest Group: A monthly hybrid seminar series focussed on early career researchers presenting and discussing their latest work. If you would like to be included in the mailing list and/or present, please contact w.decothi@ucl.ac.uk.

Training: Organising training in cutting-edge techniques and software at the interface of neuroscience and AI is vital for the exchange of research findings between these fields. We have already conducted a successful one-day tutorial on the behavioural analysis package DeepLabCut at the SWC in May this year. More sessions are in the pipeline, and we welcome suggestions for future topics.

Meet the organising committee

Caswell Barry (chair)

Caswell's goal is to understand the computations instantiated in the brain. To this end his research is focused on memory, in particular memory for places and events. His lab is based in the Research Department of Cell and Developmental Biology and uses a mixture of experimental and computational approaches, including machine learning.

Padraig Gleeson (co-chair)

Padraig is a Principal Research Fellow in the department of Neuroscience Physiology and Pharmacology. He works with Professor Angus Silver on the Open Source Brain platform for model and data sharing in neuroscience, and develops standardised languages for model specification in computational neuroscience and machine learning.

Will de Cothi

Will is a Postdoctoral Research Fellow in Professor Caswell Barry's lab at the department of Cell and Developmental Biology. He is interested in using machine learning and computational modelling to help understand the brain and behaviour.

Clementine Domine

Clementine is a PhD student in Andrew Saxe's lab at the Gastby and the Sainsbury Wellcome Centre and Caswell Barry's lab at UCL. She is studying the theory and computation at the basis of learning and memory.

Sebastian Lee

Sebastian is a PhD student in Andrew Saxe's lab at the Sainsbury Wellcome Centre and Claudia Clopath's lab at Imperial College London. He is interested in the intersection of neuroscience and machine learning, particularly in relation to the paradigms of continual learning and reinforcement learning.

Asaph Zylbertal

Asaph is a Postdoctoral Researcher in Dr Isaac Bianco's lab, based in the Department of Neuroscience, Physiology and Pharmacology at UCL. He is studying how internal state and network dynamics shape behaviour, by combining in vivo experimental approaches in larval zebrafish and computational modelling.

General Information

Joining virtually

If you won't be with us in-person please follow this link to watch the conference live.

Lunches and refreshments

There will be tea, coffee and water available at breaks throughout the day. Lunch will be provided as well as a drinks reception after the event.

Meeting name badges

All participants are required to wear identification badges at all times. If you lose your badge, please go to the registration desk where a new badge will be made for you.

Security and safety

Please do not leave bags and luggage unattended at any time, whether inside or outside the session hall.

How to find us

The address is **25 Howland Street, London, W1T 4JG** (find us on the map).

Mobile phones

As a courtesy to speakers and other participants, all mobile phones and electronic devices must be switched to silent before entering the sessions.

Wi-Fi

Wi-Fi is available for guests throughout the venue. Wi-Fi is available for all participants either via the Eduroam network or using the Guest Wi-fi network. If you have any trouble on the day speak to a member of the registration team.

Programme

08:45 Registration

09:30 Welcome and opening remarks

Professor Caswell Barry, Professor of Neuroscience and AI, UCL

Session 1

Chair: Professor Maneesh Sahani, Professor of Theoretical Neuroscience and Machine Learning, Gatsby Computational Neuroscience Unit, UCL

09:45 Professor Murray Shanahan, Professor in Cognitive Robotics, Imperial College London

Role Play with Large Language Models

10:15 Dr Maria Eckstein, Research Scientist, DeepMind

Predictive and Interpretable: Using Classic Cognitive Models and Artificial Neural Networks to Understand Human Learning and Decision Making

10:45 Refreshment break

11:15 Dr Ann Duan, Senior Research Fellow and Group Leader, Sainsbury Wellcome Centre

Mice Dynamically Adapt to Opponents in Multiplayer Games

11:45 Professor Christopher Summerfield, Professor of Cognitive Neuroscience, University of Oxford

Learning Content and Structure in a Dual-Streams Neural Network

12:15 Lightning Talks

(further details can be found on page 12)

13:15 Networking Lunch and Poster Presentations **(further details can be found on page 16)**

Session 2

Chair: Professor Tom Mrsic-Flogel, Professor of Neuroscience and Director of The Sainsbury Wellcome Centre

14:15 **Professor Ila Fiete, Department of Brain & Cognitive Sciences and Associate Member of the McGovern Institute, Massachusetts Institute of Technology (Remote)**

14:45 **Professor Timothy Behrens, Professor of Computational Neuroscience, University of Oxford**

What Does the Representation of a Schema Look Like?

15:15 **Dr Alexandra Keinath, Assistant Professor, University of Illinois Chicago (Remote)**

Shared and Distinct Representational Structure Links Mouse CA1, Human Spatial Memory, and Models of Cognitive Mapping in Deformed Spaces

15:45 **Refreshment break**

16:05 **Panel discussion**

16:45 **Closing remarks**

17:00 **Networking drinks reception**

18:30 **Close**

Speaker Biographies

Professor Murray Shanahan

Professor of Cognitive Robotics, Imperial College London
Senior Research Scientist, DeepMind

Professor Murray Shanahan is Professor of Cognitive Robotics at Imperial College London and a Senior Research Scientist at DeepMind. He works on artificial intelligence, neurodynamics, and the philosophy of mind. Educated at Imperial College and Cambridge University (King's College), he became a full professor at Imperial in 2006, and joined DeepMind in 2017. He was scientific advisor to the film Ex Machina, and regularly appears in the media to comment on artificial intelligence and robotics. As well as many scientific papers he has published several books, including "Embodiment and the Inner Life" (Oxford University Press, 2010) and "The Technological Singularity" (MIT Press, 2015).

Dr Maria Eckstein

Research Scientist, DeepMind

Dr Maria Eckstein is a research scientist at DeepMind with a PhD in psychology from UC Berkeley. Her work lies at the intersection of AI and cognitive science/neuroscience. She combines tools from artificial intelligence, such as reinforcement learning and neural networks, with those from cognitive psychology and neuroscience, such as controlled lab experiments and moment-to-moment neural recordings. She is particularly interested in questions that lie at the intersection of AI and cognitive science / neuroscience, including learning, decision making, and structured thought.

Speaker Biographies

Dr Ann Duan

Senior Research Fellow and Group Leader,
Sainsbury Wellcome Centre

Dr Ann Duan joined the faculty at the Sainsbury Wellcome Centre in summer 2021. She obtained her PhD in Neuroscience at Princeton University, where she studied prefrontal and collicular contributions to executive functions with Carlos Brody. In 2016, Ann became a Simons Collaboration on the Global Brain postdoctoral fellow in Ning-long Xu's lab at the Institute of Neuroscience in Shanghai, where she used circuit-level tools to investigate cortico-subcortical cooperation during decision formation and maintenance. Her long-term scientific goal centres around extracting general principles of neural computation that give rise to cognition. Her lab at SWC combines theory-motivated behavioural paradigms with modern experimental and analytical tools to conduct large-scale, cellular-resolution investigations of neural circuit mechanisms for flexible decision-making. In particular, she aims to understand how decision-making is modulated by internal and environmental variables, such as risk preference and competitive/cooperative social interaction.

Professor Christopher Summerfield

Professor of Cognitive Neuroscience,
University of Oxford

Christopher Summerfield is Professor of Cognitive Neuroscience at the University of Oxford and a Research Scientist at Deepmind UK. Chris Summerfield was trained in psychology and neuroscience at University College London, Columbia University (New York), and the École normale supérieure (Paris). He is Professor of Cognitive Neuroscience in the department of Experimental Psychology, where he heads a lab focused on understanding the computational mechanisms by which humans make decisions, and how these processes are implemented in the brain. His work, which involves a combination of computer simulations, behavioural testing, and functional brain imaging.

Professor Ila Fiete

Department of Brain & Cognitive Sciences and Associate Member of the McGovern Institute, Massachusetts Institute of Technology

Ila is a fete professor in the Department of Brain and Cognitive Sciences and an associate member of the McGovern Institute at the Massachusetts Institute of Technology (MIT). She was formerly an associate professor in the Department of Neuroscience at the University of Texas at Austin (UT Austin). Fiete's interests are focused on using computational and theoretical tools to better understand the dynamical mechanisms and coding strategies that underlie computation in the brain. Her recent focus is on error control in neural codes, on rules for synaptic plasticity that enable neural circuit organization, and on questions at the nexus of information and dynamics in neural systems, to understand how coding and statistics fundamentally constrain dynamics, and vice-versa.

Professor Timothy Behrens

Professor of Computational Neuroscience, University of Oxford

Tim Behrens is Professor of Computational Neuroscience at Oxford University and UCL, and a Wellcome Trust Research Fellow. He studies how our brains learn and represent knowledge about the world in service of flexible behaviour. His work aims to explain complex human behaviour but is grounded in comparisons with animal models and formal mathematical theory. It has therefore made impact across scales from cells to brain regions and across mammalian species. He has also developed widely used approaches for measuring brain connections non-invasively.

Dr Alexandra Keinath

Assistant Professor of Behavioural Neuroscience, University of Illinois Chicago

Decades of work have revealed rich neural representations of space and episodic memory in the hippocampus and neighbouring cortices of mammals from rodents to humans, with provocative similarities and differences across brain regions and species. Dr Keinath's research takes a multidisciplinary approach to better understand these representations and the functions they instantiate, with a particular emphasis on bridging the gaps between brain regions, species, and levels of explanation. She has led and collaborated on projects leveraging a diverse range of techniques including in vivo electrophysiology in rodents, in vivo calcium imaging in rodents, chemo genetic circuit manipulations in rodents, behaviour in both rodents and humans, functional magnetic resonance imaging in humans, and computational modelling with tests of the specific hypotheses generated by these models in both rodents and humans. While it is uncommon in this field for someone's work to cut across multiple techniques and species, continued advances have led to a mounting need for such work. Now as an Assistant Professor of Behavioural Neuroscience at the University of Illinois Chicago, this need has become the foundational focus of her lab.

Lightning Talks

1. Ms Lauren Bennett, PhD student, Cell and Developmental Biology, UCL

Reinterpreting Boundary Vector Cells as Trajectories in the Successor Representation

In traditional models of the hippocampus, subicular boundary vector cells (BVCs) are sensorily driven and provide input to downstream place cells. However, the subiculum is a major anatomical output of the hippocampus and the existence of vector-trace fields implies that its function is partly mnemonic. Here we present an anatomically consistent, biologically plausible alternative account of BVCs within the paradigm of the hippocampus as a predictive map. Specifically, we demonstrate that BVC-like activity arises from learning a successor representation over environmental states via rodent trajectories. Our model uses place cells as basis features to learn successor representations that recapitulate an animal's proclivity to run along walls to form elongated fields statistically analogous to BVCs. In this way, our work builds upon models of hippocampal place cells as successor representations over predictive future occupancy. Additionally, our model makes distinct empirical predictions, including subicular representations of corners and of learned trajectories independent of boundaries.

2. Professor Rhodri Cusack, Trinity College Institute of Neuroscience

Human Infants are Learning a Foundation Model

Human infants are helpless for a protracted period, long argued to be a result of maternal constraints leading to a short gestation and consequently an immature brain at birth. This is not supported, however, by a growing body of neuroimaging that has found the structure and function of the infant brain to be surprisingly well developed. We hypothesized that instead, human infants are helpless because they are learning a foundation model that will later allow them to rapidly learn a diverse range of tasks. To test this, we acquired the largest cohort of awake fMRI in 2- and 9-month-old infants (N=134), to longitudinally characterise the representational geometry in the infant ventral visual stream, and to compare it to the geometry of DNNs. The geometries were substantially correlated. Early sensory regions were more like early DNN layers and later regions more like later layers, particularly in older infants.

3. Dr George Dimitriadis, Senior Research Fellow, SWC/Gatsby

Curriculum Learning in Animals and Animats

Accumulating evidence show that animals learn complex novel tasks by utilising previously learned and inate strategies and require curricula of incremental complexity to do so. Reinforcement Learning agents on the other hand do better when presented with the full distribution of possible state, action pairs and struggle to transfer learning from one task to the next one in a curriculum driven fashion. We are aiming to compare animal behavioural data with corresponding behaviours from Reinforcement Learning agents on similar curricula based tasks in order to formulate RL models that show similar learning patterns over curricula towards complex task acquisition.

4. Mr William Dorrell, PhD student, Gatsby Unit

Actionable Neural Representations: Normative Theories of Neural Internal Models

The fundamental question in neuroscience is perhaps "why is that neuron firing like that?" - for as long as people have been measuring neurons, they have been proposing theories to explain them. Most such theories are versions of the Efficient Coding Hypothesis: these neurons encode these variables under these biological constraints. Unfortunately, these elegant theories have largely failed to explain non-sensory representations (e.g. grid cells). We'll argue this is because the brain is doing more than efficiently encoding variables, it is computing with them. We will present a theory of neurons computing internal predictive models (if I take this action how will this variable change?). We'll show that it can fit and make predictions about grid cells in the entorhinal cortex, music box representations in the prefrontal cortex, and latent sequence representations in supplementary motor areas; and that it can do this understandably (in contrast to much neural network work).

5. Mr Tom George, PhD Student, Sainsbury Wellcome Centre

Is the Hippocampus a Helmholtz Machine: Bioplausible Substrates for the Wake-sleep Algorithm and Hippocampal Structure Learning

Hippocampus combines internally generated predictions with external acquired stimuli. “Predictor-comparator” models of this ilk account for much of hippocampal function but often at the expense of interpretability and plausibility. So how would a real hippocampus learn? We introduce a new model — the Helmholtz Hippocampus — where neural oscillations coordinate local Hebbian learning rules in a close analog to the wake-sleep algorithm.

6. Dr Marcus Ghosh, Postdoctoral Research Fellow, I-X, Imperial College London

Multimodal Units Fuse-Then-Accumulate Evidence Across Channels

Animals continuously detect information via multiple sensory channels, like vision and hearing, and integrate these signals to realise faster and more accurate decisions. However, despite extensive experimental work the fundamental question of how the brain merges sensory data remains unanswered. The canonical view is that multimodal neurons linearly fuse the evidence accumulated by unimodal areas. However, by working at multiple levels from probabilistic models to artificial and spiking neural networks, we demonstrate that even slight changes to the statistical relationship between input channels cause the canonical algorithm to perform sub-optimally and even to fail in extreme cases. This leads us to propose a novel multimodal algorithm which excels in naturalistic settings, is optimal for a wide class of multimodal problems and is compatible with our current knowledge of multimodal circuits. Our work thus provides new perspectives on multisensory integration, and testable hypotheses at multiple levels: from single neurons to behaviour.

7. Mr Ed Li, Masters Student, Nuffield Department of Clinical Neurosciences, University of Oxford

A Tractable Solution to Imperfect Observation Models: Cut-Posteriors and their Application in Multisensory Integration

In an ideal setting for Bayesian agents, a perfect description of the rules of the environment (i.e., the objective observation model) is available, allowing them to reason through the Bayesian posterior to update their beliefs in an optimal way. But such an ideal setting hardly ever exists in the natural world, so agents have to make do with reasoning about the how they should update their beliefs simultaneously. This work addresses this challenge in both neuroscience and AI by (a) providing a general class of belief updates called cut-posteriors for Bayesian networks and (b) parameterize the space of possible posteriors to enable meta-learning (i.e., choosing the belief update from this space in a principled manner). We apply this proposed framework to model the neuroscience problem of multisensory integration of information in perception, which demonstrates its superiority compared to past models in flexibility and computational efficiency.

8. Mr Samuel Liebana Garcia, PhD Student, Department of Physiology, Anatomy and Genetics, University of Oxford

Striatal Dopamine Reflects Individual Long-term Learning Trajectories

Learning from naïve to expert occurs over long periods of time, entailing extensive changes in behaviour and neuronal signals. The principles governing behavioural and neuronal dynamics during long-term learning remain unknown. We developed a visual decision task for mice that allowed us to study learning trajectories from naive to expert. Mice adopted sequences of strategies that became more stimulus-dependent over time, showing substantial diversity in the strategies they transitioned through and settled on. These transitions were systematic; the initial strategy of naïve mice predicted their strategy several weeks later. Longitudinal imaging of dopamine release in dorsal striatum demonstrated that dopamine signals emerged over learning, reflecting each individual's strategy in using stimuli to make decisions. A deep neural network model trained on the task with reinforcement learning captured behavioural and dopamine dynamics. Analysing the model revealed saddle points where learning trajectories slowed down before strategy transitions. The saddle point analysis accounted for the systematicity and diversity of the animals' trajectories. Together, our results demonstrate that long-term learning involves diverse yet systematic transitions through behavioural strategies, and that dopamine signals as well as teaching signals of deep networks exhibit key characteristics to support such learning.

9. Dr Kryzystof Potempa, Founder and CEO, BRAINCURES

An LTP-gene Powered Biological Intelligence Approach to De-Risked and Accelerated Drug Development

To decode the molecular program behind synaptic plasticity, we first collected gene expression profiles from healthy animals at three different time points that are critical to the establishment of a memory through long-term potentiation. We then applied a systems biology approach to map out how genes implicated in brain function in healthy animals are organized into a molecular corporation comprised of hierarchically interconnected worker-, manager- and director- targets. This novel "biological intelligence" powers the BRAINCURES Discovery Engine to help eliminate preclinical failures and achieve unprecedented clinical success rates.

10. Elizaveta Tennant, PhD Student, Computer Science

Modeling Moral Choices in Social Dilemmas with Multi-Agent Reinforcement Learning

Practical uses of Artificial Intelligence (AI) in the real world have demonstrated the importance of embedding moral choices into intelligent agents. We believe that an interesting and insightful starting point is the analysis of emergent behavior of Reinforcement Learning (RL) agents that act according to a predefined set of moral rewards in social dilemmas. In this work, we present a systematic analysis of the choices made by intrinsically-motivated RL agents whose rewards are based on moral theories. We distinguishing between consequence- and norm-based moral rewards and between single- and mixed-virtue (e.g., multi-objective) methodologies. We evaluate our approach by modeling repeated dyadic interactions between learning moral agents in three iterated social dilemma games. We analyze the emergence of cooperation, defection or exploitation, and the corresponding social outcomes. Finally, we discuss the implications of these findings for the development of moral agents in artificial and mixed human-AI societies.

Poster exhibitors

1. Mr Umar Abubacar, PhD Student, Department of Computer Science

Agent Based Corticogenesis to Functional Neural Networks

The neocortex, pivotal for human cognition, showcases a unique six-layered architecture. Our research employs an agent-based model to recreate the journey from neuroepithelial cells to a sophisticated, interconnected cortical region. Through a gene regulatory network, our model accurately replicates cell differentiation and layer-specific migration. This simulation unravels the autonomous formation of the cortex via gene-type rules and seamlessly blends localized interactions and physical attributes, offering an integrated perspective of the neocortical microcircuit. The resulting network is precisely structured for delving into learning and memory using spiking neural dynamics, highlighting the prowess of high-end agent-based modelling in mirroring the human cortex. Future work aims to hone this adaptable, cross-species model, pushing the frontiers of realistic AI architectures and learning mechanisms.

2. Professor Neil Burgess, UCL Institute of Cognitive Neuroscience & UCL Queen Square Institute of Neurology

A Generative Model of Memory Construction and Consolidation

We present a new computational model explaining how episodic memories are (re) constructed, share neural substrates with imagination, combine unique features with schema-based predictions, and show schema-based distortions that increase with consolidation. In the model hippocampal replay (from an auto associative network) trains generative models (variational autoencoders) to (re)create sensory experiences from latent variable representations in entorhinal, medial prefrontal, and anterolateral temporal cortices via the hippocampal formation. Simulations show effects of memory age and hippocampal lesions in agreement with previous models, but also provide mechanisms for semantic memory, imagination, episodic future thinking, relational inference, and schema -based distortions including boundary extension. The model explains how unique sensory and predictable conceptual elements of memories are stored and reconstructed by efficiently combining both hippocampal and neocortical systems, optimising the use of limited hippocampal storage for new and unusual information. Overall, hippocampal replay training generative models provides a comprehensive account of memory construction, imagination and consolidation.

3. Dr Samuel Failor, Research Associate, Department of Neuromuscular Diseases

Visuomotor Association Orthogonalizes Visual Cortical Population Codes

The brain should be best able to associate distinct behavioural responses to sensory stimuli if these stimuli evoke population firing patterns that are close to orthogonal. To investigate whether task training orthogonalizes population codes in primary visual cortex (V1), we measured the orientation tuning of 4,000-neuron populations in mouse V1 before and after training on a visuomotor task. The effect of task training on population codes could be captured by a simple mathematical transformation of firing rates, which suppressed responses to motor-associated stimuli, but only in cells responding to them at intermediate levels. This transformation orthogonalized the representations of the task orientations by sparsening the population responses to these stimuli. The strength of response transformation varied from trial to trial, suggesting a dynamic circuit mechanism rather than static synaptic plasticity. These results indicate a simple process by which visuomotor associations orthogonalize population codes as early as in primary visual cortex.

4. Mr Sihao Liu, Research Assistant, UCL Department of Cell and Developmental Biology

Transformer Learns Hippocampus CA1 to Septal Circuit

Using simultaneous GCaMP and Neuropixel recordings, we train transformer neural networks on the projections from a number of hippocampal CA1 and lateral septal neurons. We found that with appropriate preprocessing and smoothing the DNN can learn the circuit with high precision, as a comparison to a control group of cortical neurons from other brain regions.

5. Mr Gianvito Losapio, Researcher, Department of Electronics, Information and Bioengineering

Emergent Abilities and Large Language Models: a Computational Perspective

The scientific community has given empirical evidence that Large Language Models (LLMs) improve their performance and sample efficiency over a variety of novel tasks as a function of model scale. The term emergent abilities has been used to describe such an unpredictable phenomenon. The results appear consistent with a hypothesis that language models acquire a range of semantic understanding and logical reasoning skills by increasing the number of parameters. However, that remains yet an intuition. In this work, we formulate a computational model that - drawing inspiration from artificial intelligence and neuroscience - may elicit a more precise, detailed answer. Additionally, we also claim that a new operational definition of intelligence is possible as a specific class of problems in our computational model - with potential implications in future research.

6. Ms Joséphine Raugel, PhD Student, Ecole Normale Superieure - Laboratory of Cognitive and Computational Neuroscience

Inferential Framework Decoded in the Human Brain During Speech Processing

Language is central to human cognition. Yet, its biological and computational bases remain largely unknown. Predictive coding theory is ambiguous about whether surprisal should lead to a (1) sharpening of neural representations to guide attention and prepare update signals or (2) blur them to reduce inconsistencies between inputs and expectations. Here, we put these hypotheses to a test, by decoding two levels of representation of the speech hierarchy, and evaluating how these codes change as a function of their level-specific surprisal. To test the possible presence of an inferential framework in the processing of language, we correlate the decoded representations of words and phonemes with the stimuli posteriors as approximated with Large Language Models (LLM).

7. Mr Songyuan Xiao, Medical Student, University College London Medical School

Behavioural Correlates of Human Hippocampal Sharp-Wave Ripples During Associative Memory Function

In rodents, patterns of hippocampal neural activity observed during active behaviour are later ‘replayed’ during periods of quiet rest or sleep, accompanied by sharp-wave ripple (SWR) events in the local field potential (LFP). SWRs have been extensively studied in rodents, but their role in human memory function is less clear. Here, we aimed to characterise the properties of human hippocampal SWRs and investigate their relationship with memory function. We recorded hippocampal LFP and single-neuron activity from five epilepsy patients performing eleven sessions of an associative memory task. Candidate SWRs showed similar properties to those in rodents, albeit with a lower peak frequency, consistent with previous human data. Interestingly, and in direct contrast to our predictions, we found a significant decrease in the incidence rate of SWRs during memory retrieval compared to all other task periods. As such, our results suggest a role for SWRs in memory formation rather than retrieval.