UK Neural Computation 2019

Programme

July 02-03 2019

University of Nottingham, UK

Schedule

Tuesday 2nd		
11:45 – 12:50	Registration & lunch	Lunch at 12
12:50 – 13:00	Welcome	
13:00 – 13:30	Tim Behrens	The Tolman-Eichenbaum Machine
13:30 – 14:00	Claudia Clopath	Inhibitory microcircuits for top-down plasticity of sensory representations
14:00 – 14:30	Benjamin Evans*	Building biological constraints into convolutional neural networks for classification overcomes biases within datasets
14:30 – 15:00	Coffee	
15:00 – 15:30	Barbara Webb	Neural mechanisms of insect navigation
15:30 – 16:00	Brendan Bicknell*	Learning the principles of single-neuron computation
16:00 – 16:30	Chris Summerfield	Structure learning and the parietal cortex
16:30 – 18:30	Posters A & reception	
Wednesday 3rd		
9:30 – 10:00	Tim Vogels	Motor primitives in time and space by targeted gain modulation in recurrent cortical networks
10:00 – 10:30	Mat Evans*	Limits on the capacity for independent neural population codes in somatosensory cortex
10:30 – 11:00	Rosalyn Moran	The Free Energy Principle plays Doom: a comparison with reward-based decision making in artificial intelligence environments
11:00 – 11:30	Coffee	
11:30 – 12:00	Matthias Hennig	Sloppy spaces in neural encoding and plasticity
12:00 – 12:30	Shirley Mark*	A mechanistic account of transferring structural knowledge across cognitive maps
12:30 – 13:00	Andrew Barron	The major transitions in the evolution of cognition
13:00 – 14:00	Lunch	*Overlaps with start of poster session*
13:30 – 15:30	Posters B & coffee	
15:30 – 16:00	Athena Akrami	Past matters: Sensory stimulus history in the neuronal populations and its impact on working memory decision tasks
16:00 – 16:30	Cian O'Donnell	The rules of synaptic plasticity are stochastic
16:30 – 17:15	Simon Laughlin (Plenary)	What is neural computation and where does it live?

* Selected talk from abstracts
All talks are in the lecture theatre A42
All food and drink will be served in room A48
Poster sessions are in the main corridors

Chair's note

Welcome to UK Neural Computation 2019!

Our community is thriving of late. We've major investment in computational neuroscience here at the University of Nottingham, and notable new strengths through recruitment at Sheffield, Bristol and other institutions. The AI explosion has seen a concomitant burst of interest in theories of how the brain learns, remembers, and computes. And with the climate crisis we should all be mindful of our carbon footprint, for which a strong local meeting is one small contribution.

Thus it seemed the time was ripe to launch a national meeting for the UK community working on the computational side of neuroscience. Be it modelling, maths, data science, theory, machine-learning, and all combinations thereof.

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 functional 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

School of Psychology, the University of Nottingham

I'd like to especially thank my co-organisers at the University of Sheffield for great assistance throughout

Hannes Saal (Abstracts Chair)

Robert Schmidt

Stuart Wilson

Thanks

The conference organisers would like to thank:

Silvia Maggi & Mat Evans for assistance

Charline Tesserau & Francois Cinotti for volunteer duties

Hannes Saal & Francois Cinotti: programme book compilation and checking

Dominic Henry (Nottingham Store): registration

Melissa Wheatley and the Nottingham Conferences team: event co-ordination

Steve Coombes for advice

Sponsors

We are grateful for the support of the Medical Research Council (UK), and of the University of Nottingham's Precision Imaging Beacon

MRC Medical Research Council

The University of Nottingham

Contents

Talk abstracts

Tuesday 2nd

Tim Behrens

The Tolman-Eichenbaum Machine

Flexibly generalising knowledge from one task to another is a hallmark of intelligent behaviour but is notoriously difficult to achieve in artificial neural networks. The hippocampal-entorhinal system is essential for generalisation in both spatial and non-spatial tasks. Famously, in space, cells in these regions display a buffet of distinct bespoke firing patterns thought useful for representing maps. Existing theories extend some of these representations to non-spatial situations within a single task. These frameworks, however, do not explain generalisation, and account for only a subset of recorded cell types. Here we provide a mechanistic understanding of the hippocampal role in generalisation, and a unifying principle underlying many different entorhinal and hippocampal cell-types. We propose medial entorhinal cells form a basis for describing structural knowledge, and hippocampal cells-types reflect conjunctions of this basis with sensory knowledge. Adopting these principles, we introduce the Tolman-Eichenbaum machine (TEM), a model that learns and generalises abstract structural knowledge. After learning, TEM entorhinal cells display diverse properties resembling apparently bespoke cell types, such as grid cells, band cells, border cells and object vector cells. TEM hippocampal cells resemble place cells and landmark cells and remap between environments. TEM predicts, however, that hippocampal remapping is not random as previously believed. Rather structural knowledge is preserved across environments. We demonstrate this structural transfer over remapping empirically in simultaneously recorded place and grid cells.

Claudia Clopath

Inhibitory microcircuits for top-down plasticity of sensory representations

Animals learn better when it matters for them. For example, they learn to discriminate sensory stimuli when they receive a reward. As a result of learning, neural responses to sensory stimuli are adjusted even in the first processing stages of sensory areas (Goltstein et al. 2013, Khan et al. 2018, Poort et al. 2015). It is thought that behaviourally relevant contexts, such as rewards, trigger an internal top-down signal available to these early sensory circuits. This could be mediated by cholinergic inputs from the basal forebrain for example (Letzkus et al. 2011). One challenge remains: contextual signals are typically present for a short time only, but synaptic changes require time to be expressed. How can these time scales be bridged? We hypothesise that interneuron circuits, which recently emerged as key players during learning and memory, bridge the timescales. We investigate how temporary top-down modulation by rewards can interact with local excitatory and inhibitory plasticity to induce long-lasting changes in sensory circuitry. We propose that learning can happen in two stages: 1) Unspecific top-down signals rapidly induce an inhibitory connectivity structure between different interneuron types. 2) The inhibitory structure induces changes in sensory representations by guiding excitatory plasticity between pyramidal cells. Using a computational model of layer 2/3 primary visual cortex, we demonstrate how inhibitory microcircuits could store information about the rewarded stimulus to guide long-term changes in excitatory connectivity in the absence of further reward. We make specific testable predictions in terms of activity of different neuron types. The suggested two-stage plasticity mechanism in canonical cortical microcircuits could be conserved across different modalities.

Gaurav Malhotra, Benjamin Evans and Jeffrey Bowers

Building biological constraints into convolutional neural networks for classification overcomes biases within datasets

Part of the appeal of deep convolutional networks is their ability to learn on raw data, obviating the need to hand-code the feature space. It has been demonstrated that when networks perform "end-to-end" learning, they develop features in early layers that not only lead to a good classification performance but also resemble the representations found in biological vision systems. These results have been used to draw various parallels between deep learning systems and human visual perception. In this study, we show that end-to-end learning in standard convolutional neural networks (CNNs) trained on a modified CIFAR-10 dataset are found to rely upon idiosyncratic features within the dataset. Instead of relying on abstract features such as object shape, end-to-end learning can pick up on low-level and spatially high-frequency features, such as noise-like masks. Such features are extremely unlikely to play any role in human object recognition, where instead a strong preference for shape is observed. Through a series of empirical studies, we show that these CNNs cannot overcome such problems merely through regularisation methods or more ecologically plausible training regimes. However, we show that these problems can be ameliorated by forgoing end-to-end learning and processing images with Gabor filters in a manner that more closely resembles biological vision systems. These results raise doubts over the assumption that simply learning end-to-end in "vanilla" CNNs leads to the emergence of similar representations to those observed in biological vision systems. By adding more biological input constraints, we show that deep learning models can not only capture more aspects of human visual perception, but also become more robust to idiosyncratic biases within training sets.

Barbara Webb

Neural mechanisms of insect navigation

Insect navigation has been a focus of behavioural study for many years, but the neural mechanisms are largely unknown. We have used computational modelling to bridge this gap, by relating the computational requirements of navigational tasks to the type of computation offered by invertebrate brain circuits. We have shown that visual memory of passing views could be acquired by associative learning in the mushroom body neuropil, and allow insects to recapitulate long routes. We have also proposed a circuit in the central complex neuropil that integrates sky compass and optic flow information on an outbound path and can thus steer the animal directly home. The models are strongly constrained by neuroanatomy, and are tested in realistic agent and robot simulations.

Brendan Bicknell and Michael Hausser

Learning the principles of single-neuron computation

The transformation of spatiotemporal patterns of synaptic input to action potential output is central to computation in the brain. At the level of single neurons, this operation is shaped by the elaborate morphology of the dendritic tree and a suite of biophysical mechanisms that confer nonlinear processing capabilities. Although extensively studied, important questions remain unanswered as the interactions between input statistics and dendritic biophysics are difficult to infer through current techniques. To address these challenges, we introduce a method for exploring the biophysics of computation by interrogating the dynamics of a model neuron as it learns to perform a task. We augment a detailed model of a layer 2/3 pyramidal neuron with a set of variational equations that report precisely how the somatic voltage at a given point in time depends on the history of synaptic input. In this we account for the influence of morphology, interactions between excitatory and inhibitory inputs, and other features such as global coupling from back-propagating action potentials. We use this as the basis for a learning algorithm with which we train the model neuron to discriminate noisy patterns of synaptic input. We find that as the computational load on the neuron increases (i.e. the number of patterns to be discriminated) synaptic weight distributions become increasingly sparse, and form a representation of the probability of observing a preferred pattern. This suggests a general role for dendritic processing in supporting computation in the presence of noise, and provides quantitative insight into which features of spatiotemporal input statistics may be used to carry salient information.

Chris Summerfield

Structure learning and the parietal cortex

Deep learning models have been proposed as a theory of computation in the ventral visual stream, but current networks suffer from data inefficiency, lack of robustness, and failures of generalisation. In the primate, visual processing is shared between dorsal and ventral pathways. I will argue that the parietal cortex is key for understanding relational structure in spatio-topically organised visual scenes (gestalt perception). Visual structure learning may be a key building block for understanding more abstract concepts, such as proximity, similarity and magnitude. I will describe experiments using human brain imaging which suggest that the parietal cortex learns low-dimensional relational information about otherwise high-dimensional visual objects. I will argue that structure learning in primates is grounded in the actions that we take to orient towards objects, coded in parietal areas such as LIP. I will draw connections to theories of spatial encoding and structure learning in the hippocampal-entorhinal system.

Wednesday 3rd

Tim Vogels

Motor primitives in time and space by targeted gain modulation in recurrent cortical networks

Animals perform an extraordinary variety of movements over many different time scales. To support this diversity, the motor cortex (M1) exhibits a similarly rich repertoire of activity. Although recent neuronal network models capture many qualitative aspects of M1 dynamics, such as complex multiphasic activity transients, they can generate only a few distinct movements. Additionally, it is unclear how M1 efficiently controls movements over a wide range of speeds and shapes. Here we demonstrate that simple modulation of neuronal input-output gains in recurrent network models with fixed connectivity can substantially and predictably affect downstream muscle outputs. Consistent with the observation of diffuse neuromodulatory projections to motor areas, our results suggest that a relatively small number of modulatory control units can provide sufficient flexibility to adjust high-dimensional network activity on behaviourally relevant time scales. Such modulatory gain patterns can be obtained through a simple reward-based learning rule. Novel movements can also be assembled from previously learned primitives, thereby facilitating fast acquisition of hitherto untrained muscle outputs. Moreover, we show that it is possible to separately change movement speed while preserving movement shape, thus enabling efficient and independent movement control in space and time. Our results provide a new perspective on the role of neuromodulatory systems in controlling recurrent cortical activity and suggests plasticity of single-neuron excitability as an important substrate of learning.

Mathew Evans, Rasmus Petersen and Mark Humphries

Limits on the capacity for independent neural population codes in somatosensory cortex

A fundamental concept in sensory systems is population coding: representation of information by the joint activity of a neuron group. The number of independently-active populations in sensory cortex sets limits on an animal's capacity for coding. But unknown is the capacity of independent populations in sensory cortex. Here we use a comprehensive sample of task-relevant neurons to ask how many independent neural populations lie within a defined cortical region, and whether those populations carry independent codes. To address these questions we extended our previous work on unsupervised spectral clustering (Humphries J.Neuroscience 2011) to include generative null models of expected correlation structure, enabling accurate detection of independent dimensions of activity that exceed those expectations.

We applied this framework to large-scale recordings (up to 2000 simultaneously recorded neurons) from superficial primary somatosensory cortex (L2/3 S1) in awake behaving mice (Peron et al. Neuron 2014). Mice were trained to perform a sensory discrimination task with a single whisker, thus the ~6000 neurons of a 'barrel' column forms a functional bottleneck for the task (O'Connor et al. J.Neuroscience 2010).

We found recordings always contained independently active neural populations beyond those expected by generative models. The number of independent populations scaled approximately linearly with the number of simultaneously recorded cells, predicting an upper limit of ~50 independent populations in L2/3. Classic dimension reduction using principal components analysis estimated up to ten-fold greater dimensionality. PCA dimensionality estimates were not linearly related to those from the generative model, underlining the importance of appropriate null models. The leading neural dimension in every animal jointly encoded touch and whisking amplitude, and had sustained activity over the delay period between stimulus offset and reward delivery. Our results suggest the potentially limited capacity for coding by independent populations in L2/3 of S1 is offset by their multiplexing sensory and task information.

Rosalyn Moran

The Free Energy Principle plays Doom: a comparison with reward-based decision making in artificial intelligence environments

Under Active Inference (Friston 2009), decisions – such as that to move ones' eyes - is driven by the imperative to minimise a bound on surprise known as the Free Energy. In the context of partially observable Markov decision processes (POMDPs), a model-based framework in which we can cast naturalistic decision-making tasks; the Free Energy of a policy (a sequence of actions) can be understood as a drive to both minimize cost (maximise the likelihood of achieving a goal) while maximising the information return from a given set of actions. This scheme has been used to model decision making in tasks such as 'the urn task' and also in reading.
In my talk I will introduce the technical framework of Free Energy minimization in the context of online gaming environments (designed to test artificial intelligence algorithms) and present data from decision-making simulations. Specifically I will present the game 'Doom' and compare agents trained under Active Inference to agents trained to maximise reward. Linking these simulations to putative neurobiological substrates I will describe the potential links from brain to computation.

Matthias Hennig

Sloppy spaces in neural encoding and plasticity

Synaptic connectivity constrains the state space accessible to network of neurons, and neural plasticity shapes this space during development as well as in an experience-dependent manner. Experiments show that this subspace can be stable over a long period of time, while the strength of individual synapses may fluctuate considerably. Moreover, synaptic plasticity is commonly assumed to be strictly local, and is unclear how the effect of a synaptic modification on the functional integrity of neural circuits can be locally evaluated. Here we use energy-based models of neural circuits to address this question in three scenarios, the encoding of sensory information, storing of new memories in a recurrent network, and pruning of excess synaptic connections. We find a highly anisotropic synaptic parameter space, with stiff directions where small changes lead to large functional changes and sloppy directions that do not affect the model behaviour. Interestingly, the relative importance of each connection can be evaluated locally at each synapse. We show that this can be utilised to suppress response variability in an encoding scenario, enables continual learning, and effective developmental pruning.

Shirley Mark, Tomas Parr, Steve Kennerley and Tim Behrens

A mechanistic account of transferring structural knowledge across cognitive maps

Animals can learn abstract generalizable knowledge from different environments and transfer this knowledge to similar circumstances. This ability has been characterised several decades ago, yet its underlying neuronal and computational mechanisms are unknown. It has been suggested that upon entering a new environment, or encountering a new task, a cognitive map that represents the relationships between elements in this environment or task is being learned. We suggest that generalization across environments occurs by representing sets of prior knowledge about the probable relationships between elements in different environments and tasks that have been encountered previously. Frequently the relationship between elements in the environment will follow some pattern, or structure. For example, in some environments elements can be clustered while in others they can be organised along a tree. An important question concerns how animals transfer such structural knowledge between different cognitive maps and use it to efficiently construct a new cognitive map to guide their decisions. We present a computational architecture that supports the flexible construction of a cognitive map within a novel environment, based on an explicit representation of a basis set for structural knowledge. These basis sets represent structural knowledge in a compressed manner and are disentangled from sensory stimuli. We show that this type of representation allows inference of important states and drawing inferences about routes that have not been yet taken). In line with our model, we show that participants who have a correct structural prior are able to make better inferences about unobserved routes compared to participants with the wrong prior and are able to infer better important task states. Therefore supporting the idea that this abstract relational code can be acquired and generalised across different cognitive maps.

Andrew Barron

The major transitions in the evolution of cognition

Recent findings in comparative cognition seem to have confused rather than clarified our understanding of the evolution of cognition. In a world of algebraic bees and smart slime moulds is there any pattern to the evolution of cognitive capacity? Maynard Smith and Szathmary famously provided a framework for understanding the grand scheme of biological evolution by positing a few major transitions - such as the origins of chromosomes and multicellular life – which enabled radically different forms of life, new evolutionary options and increased evolvability. Here we propose that the evolution of cognition can also be comprehended as a series of major transitions: each transition being a qualitative change in the structure of information flow within systems. These transitions enabled new types of cognitive capacity while transforming the scope of existing cognitive abilities. Here we present each transition in term of system organization. We discuss the capacities enabled by each transition, and the consequences of this perspective for our understanding of the evolution of cognition and the diversity of animal intelligences.

Work done in collaboration with Colin Klein

Athena Akrami

Past matters: Sensory stimulus history in the neuronal populations and its impact on working memory decision tasks

Individuals adapt to the statistical properties of the environment, by forming appropriate "priors" through experience. But very little is known causally and mechanistically about which brain regions are necessary for prior stimulus history effects, or about how prior stimulus information is represented at the level of individual neurons and networks. We addressed these issues by developing a rodent model of an auditory "parametric working memory" (PWM) task, the sequential comparison of two graded stimuli separated by a delay period of a few seconds, with simultaneous recordings in posterior parietal cortex (PPC). Using a PWM allowed us to address the representation and use of priors via the well-documented phenomenon "contraction bias", the progressive shift of the first stimulus representation towards the center of a prior distribution built from the past experience. Using the large data made possible by our high-throughput semiautomated training protocols, we developed powerful statistical models to quantify behavior and to link brain activity to behavior. We characterized how and on what timescales the prior distribution is built and how it interacts with other contributing factors in order to form the final decision. Through a series of causal experiments, we have discovered a crucial history-dependent contribution of the rat posterior parietal cortex (PPC) to working memory behaviors and specifically its impact on the "contraction bias", via the memory of previous sensory events. Although previously it has been shown that the activity of neurons in PPC are influenced by previous animals' choice, this is the first time that the contribution of PPC in the history of sensory input is documented.

Cian O'Donnell

The rules of synaptic plasticity are stochastic

To understand how brain learning works we'd like to model the rules of synaptic plasticity, then study their properties in neural circuit simulations. Traditional plasticity models are either 1) phenomenological, abstracting away the molecular machinery of synapses, or 2) highly detailed biochemical models. The former type are likely too simple to capture key properties of synapses, but the latter type are too complicated to analyse mathematically or even to simulate in neural circuit models. We instead aimed to build a new model at an intermediate level of detail, based on physiological data from rodent hippocampal synapses. We found that the model was surprisingly stochastic, with a pattern that matches systematic variability across protocols in electrophysiological data.

Joint work with Romain Veltz, Yuri Rodrigues, and Hélène Marie.

Simon Laughlin

What is neural computation and where does it live?

At first sight this question is redundant. We already have a perfectly good answer: neural computation is information processing by neurons and glia, performed in brains. I advocate placing neural computation in a broader context. I will argue from first principles, and by giving examples, that neural computation should be viewed as a specialised form of biological information processing, largely performed by macromolecular circuits. This wider perspective has the advantage of greatly increasing the numbers of ways in which neurons can implement a particular computation, but it also raises the alarming prospect of being lost in a much larger space of all plausible models. However, because this space is reduced by biophysical, cell biological and evolutionary constraints, we should accept the challenge of taking a broader view of how neurons and glia compute.

List of posters

Poster session A

1. **Effect of gamma oscillations in the rat basal ganglia on spiking and reaction times**
   Chua A, Leventhal D, Mallet N, Berke JD, Schmidt R

2. **Long-range hippocampal-frontal cell assemblies dynamically integrate and segregate information during working memory**
   Domanski A, Kucewicz M, Tricklebank M, Robinson E, Durstewitz D, Jones M

3. **Combining connectomics and modelling to understand innate olfactory behaviour in Drosophila.**
   Bates A, Drummond N, Schlegel P, Frechter S, Roberts R, Tamimi I, Dhawan S, Huoviala P, Love F, Rubin G, Costa M, Bock D, Jefferis G

4. **Theta oscillations gate the transmission of reliable sequences in the medial entorhinal cortex**
   Balachandar AN, Assisi C

5. **Spontaneous spiking homeostasis as a ROS-regulating ATP-excursion**
   Chintaluri C, Vogels T

6. **Neural Activity Reveals Interactions Between Episodic and Semantic Memory Systems During Retrieval**
   Weidemann C, Kragel J, Kahana M

7. **SpikeInterface: A standardized framework for sorting, analysis, and evaluation of extracellular recordings**
   Hurwitz C, Buccino A, Magland J, Gerlei K, Nolan M, Hennig M

8. **Sensory Processing and Categorization in Cortical and Deep Neural Networks**
   Pinotsis D

9. **Hybrid synaptic plasticity modelling via integration of electrical and biochemical dynamics via KappaNEURON**
   Linkevicius D, Sterratt DC

10. **The Synfire/deSync Model: Deciphering Episodic Content from Cortical Alpha Oscillations**
    Parish G, Michelmann S, Hanslmayr S, Bowman H

11. **MIIND: A simulator for neural populations**
    Osborne H, Deutz L, Lai YM, de Kamps M

12. **Modelling of retinal mosaic formation using an agent-based model**
    de Montigny J, Sernagor E, Bauer R

13. **Interacting cortical subnetworks for first-order and higher-order sensory processing**
    van Rheede J, Buchan M, Gothard G, Akerman C

14. **A Neurally-Informed Modelling Approach for Investigating Cross-Modal Associations in Perceptual Decision-Making**
    Bolam J, Ince R, Boyle S, Delis I

15. **Spectral theory of non-stationary dynamics of integrate-and-fire neurons**
    Deutz L, Osborne H, de Kamps M

16. **Constrained plasticity can compensate ongoing drift in neural populations**
    Rule ME, Loback AR, Raman DV, Harvey CD, O'Leary TS

17. **Non-synaptic interactions in the neural encoding of odorants: a good start is half the battle**
    Pannunzi M, Nowotny T

18. **Weight-dependent BCM plasticity leads to a rich family of receptive fields**
    van Rossum M, Froc M, Williamson O

19. **Variability and regulation of reinforcement learning processes in rats**
    Cinotti F, Girard B, Khamassi M

20. **Generalisation of Frequency Mixing and Temporal Interference through Volterra Expansion**
    Perez Nieves N, Goodman D

21. **Simple vs complex: how complexity affects causal inference in visual perception**
    Gekas N

22. **Robust reconstruction of mouse poses during visual experiments**
    Storchi R, Cootes TF, Killick R, Lucas RJ

23. **Firing rate of neurons with dendrites, soma and axon in the fluctuation-driven, low-rate limit**
    Gowers R, Timofeeva Y, Richardson M

24. **Deconstructing the protomap versus protocortex hypotheses via computational modelling**
    James S, Whiteley D, Brooke J, Wilson S

25. **Degenerate Optimal Boundaries for Multiple-Alternative Decision Making**
    Baker S, Griffith T, Lepora N

26. **A biologically grounded model of the hypothalamic control of motivated behaviour**
    Jimenez-Rodriguez A, Schmidt R, Gilmour W, Wilson S

27. Frequency-sensitivity and magnitude-sensitivity in decision-making: a model-based study
Bose T, Bottom F, Reina A, Marshall J

28. Does structure in neural activity match anatomical structure?
Delaney T, O'Donnell C

29. The Interplay Between Somatic and Dendritic Inhibition Promotes the Emergence and Stabilization of Place Fields
Pedrosa V, Clopath C

30. Prediction of Behavioral Responses Using Bayesian Model from Multi-Unit Spike Recordings in Sensorimotor Cortex of Rats
Öztürk S, Güçlü B

Poster session B

1. **Computational model of sodium Nav1.7 channel activation in nociceptive dorsal root ganglion neurons**
   Capurro A, Thornton J, Armstrong L, Sernagor E

2. **Neural processing of sentences**
   Burroughs A, Houghton C, Kazanina N

3. **The statistics of optimal decision-making: Exploring the relationship between signal detection theory and sequential analysis**
   Griffith T, Baker S, Lepora N

4. **Learning, extinction and re-learning: a model for Drosophila olfactory plasticity**
   Menzat B, Londsdale R, Felsenberg J, Waddell S, Vogels T

5. **Persistence of sensory representation during neurodevelopment**
   Mizusaki BEP, Kourdougli N, Suresh A, Portera-Cailliau C, O'Donnell C

6. **Efficient replay memory through sensory adaptation**
   Macmillan C, Chua R, Costa RP

7. **Degeneracy in neuromodulatory neural circuits: serotonin-dopamine interaction in conditioning tasks**
   Behera C, Joshi A, Wang D, Sharp T, Wong-Lin K

8. **Towards understanding one-shot place learning in spatial navigation: A reinforcement learning approach**
   Tessereau C, O'Dea R, Bast T, van Rossum M, Coombes S

9. **Codependent plasticity: learning with interacting synapses**
   Agnes EJ, Vogels TP

10. **Modelling the role of substance P in reinforcement-based sequence learning**
    Favila Vazquez N, Gurney K, Overton PG

11. **Multiphysics of neural signals**
    Wang H

12. **Energy efficient synaptic plasticity and network learning**
    Li HL, van Rossum M

13. **Modelling Spinal Circuits using MIIND**
    Osborne H, Deutz L, de Kamps M

14. **Recurrent inhibition and inhibitory plasticity outperform alternative mechanisms for sparse coding in a model of the Drosophila mushroom body**
    Bennett J, Vogels T, Nowotny T

15. **Influence of Prediction Error on Memory Encoding and Parahippocampal and Medial Prefrontal Connectivity**
    Khoo J, Derrfuss J, Danielmeier C

16. **Nonlinear scaling of resource allocation in sensory bottlenecks**
    Edmondson LR, Jimenez Rodriguez A, Saal HP

17. **A topological insight on neuronal morphologies**
    Kanari L, Hess K, Markram H

18. **Discovering The Building Blocks of Hearing: A Neuro-Inspired, Data-Driven Approach**
    Weerts L, Clopath C, Goodman D

19. **Robust learning with subtractive recurrent neural networks**
    Llera Montero M, Malhotra G, Bowers J, Costa RP

20. **Fiber Tract Stimulation using a Phenomenological Application of tFUS in NEURON Leads to Synaptic Potentiation**
    Alkhawashki M

21. **Effective connectivity extracts clinically relevant prognostic information from resting state activity in stroke**
    Adhikari M, Griffis J, Siegel J, Deco G, Insabato A, Gilson M, Corbetta M

22. **Neural Topic Modelling**
    Hathway P, Goodman D

23. **Modelling tactile responses from the sole of the foot**
    Kazu Siqueira R, Katic N, Bent L, Strzalkowski N, Raspopovic S, Saal HP

24. **The dynamics of explore-exploit decisions reveals a signal-to-noise mechanism for random exploration**
    Feng S, Wilson R, Wang S

25. **Integrating classifiers and electrophysiology to better understand hearing loss**
    Smith S, Wallace M, Berger J, Sumner C

26. **A Bayesian evidence accumulation approach to inferring learning strategies**
    Maggi S, Hock R, Moran P, Bast T, Humphries M

27. **The effect of the recovery variable parameters on oscillating Izhikevich networks**

Berg WP, Onken A

28. Modelling the optimal integration of navigational strategies in the insect brain
Sun X, Mangan M, Yue S

29. An Attentional Inhibitory Feedback Network for Multi-label Classification
Chu Y, Luk W, Goodman D

30. Phase-Amplitude Coupled DCM and Time-Frequency Analysis of EEG Data from AD Patients
Tyrer A, Gilbert J, Fagerholm E, Adams S, Bankole A, Gilchrist I, Moran R

Poster abstracts

Poster session A

(1)

Effect of gamma oscillations in the rat basal ganglia on spiking and reaction times

Chua A, Leventhal D, Mallet N, Berke JD, Schmidt R

The basal ganglia (BG) display prominent oscillatory activity in multiple frequency bands including theta, beta, and gamma. Beta oscillations in the BG may be related to cue utilization (Leventhal et al., 2012), and gamma oscillations in the ventral striatum seem to be related to rewards (van der Meer et al., 2010). While slow gamma oscillations (50Hz) in the ventral striatum might originate from the piriform cortex (Carmichael et al., 2017), the role and origin of gamma oscillations in the dorsolateral striatum and other basal ganglia subregions are less clear. Here, we investigated the LFP and activity of single units recorded in the BG of rats performing a stop-signal task. We found that gamma oscillations were present in most subregions of the BG, but were most prominent in the globus pallidus (GP). Contrary to previous studies conducted in the ventral BG, we did not find slow gamma oscillations when animals approached the reward location in the dorsolateral striatum and GP. Our preliminary analyses also suggested that striatal medium spiny neurons and fast-spiking interneurons were locked to different phases of slow gamma oscillations in the dorsolateral striatum. Similarly, arkypallidal and prototypical neurons were locked to different phases of slow gamma oscillations in GP. This is important because it means that the preferred gamma phase of a GP neuron could be used as a method to distinguish arkypallidal and prototypical neurons. Lastly, gamma and beta power in the GP correlated with reaction time briefly before the Go cue, supporting that these oscillations affect the utilization of cues for behaviour. We conclude that gamma oscillations in the BG play a role for the temporal organization of spiking activity in the striatum and GP.

(2)

Long-range hippocampal-frontal cell assemblies dynamically integrate and segregate information during working memory

Domanski A, Kucewicz M, Tricklebank M, Robinson E, Durstewitz D, Jones M

Information integration between functionally specialised brain areas is central to flexible cognition. One proposed signature of network-level interaction supporting this process is the transient activation of ‘cell assemblies’, whose constituent neurons are identifiable against wider network activity by virtue of their co-fluctuating spiking. However, assemblies’ roles in information transfer within and between circuits remains enigmatic. For example, do assemblies form between neurons sharing the same task-related information – thus strengthening their individual representations - or rather, between neurons encoding different features, thereby binding multiple information streams into a unified percept? Do these rules differ between brain areas and are they critical for task performance?

We addressed these questions during interactions between hippocampal dCA1 and mPFC circuits, both critically implicated in spatial working memory, using simultaneous multiple single unit recordings from rats trained to perform a Delayed Non-Match-To-Sample task. Firing rates of dCA1 and mPFC units provided different time-varying information profiles about task cue identity with specific vulnerability to increasing delay length emerging in mPFC. Assemblies were identified from the pooled spike trains using a novel factor analysis-based method. Compared to assembly non-members, mPFC neurons forming inter-area assemblies with dCA1 neurons showed both a stronger representation of the cue identity around the sample phase and a specific response-preparatory elevation of cue identity recall preceding the choice.

Spike trains of joint dCA1-mPFC assembly-forming units in both areas also robustly discriminated between sample and choice task phases and were uniquely modulated by a 5Hz rhythm, which was attenuated on error trials. Synthetic cell assemblies constructed from our data demonstrate a near-optimal representation of cue information by fewer than 5 neurons, recapitulating the regional and temporal dynamics of information representation.

Our findings unveil rules for the dynamic segregation and integration of information during decision-making by rhythmically modulated cell assemblies linking dCA1 and mPFC.

(3)

Combining connectomics and modelling to understand innate olfactory behaviour in Drosophila.

Bates A, Drummond N, Schlegel P, Frechter S, Roberts R, Tamimi I, Dhawan S, Huoviala P, Love F, Rubin G, Costa M, Bock D, Jefferis G

The ability to fully characterise the connectivity of a neural circuit can both generate new insight, and informed and testable hypotheses regarding information processing and behavioural output. Whole brain serial section electron microscopy (EM) data enables circuit reconstruction at synaptic resolution. Drosophila melanogaster has ~50 olfactory receptor neuron types projecting axons to ~50 corresponding glomeruli in the antennal lobe (AL). Projection neurons (PNs) then relay olfactory information to higher brain areas, with a majority projecting to the lateral horn (LH), a neuropil associated with innate behaviours. Understanding how olfactory information is integrated within the LH requires a detailed understanding of what information is transmitted by PNs, and how this is represented within this neuropil. Using serial section EM data (Zheng et al, 2018) we have reconstructed the full complement of olfactory PNs. These reconstructions have revealed a previously unobserved diversity, and increased the number, of excitatory and inhibitory multiglomerular PNs, which receive input in multiple glomeruli. We also observe a structural bias in the sampling of the odour space represented within the AL by these neurons. Further, we have reconstructed a set of third order LH neurons postsynaptic to PNs. We have combined these PN and LH neuron reconstructions with >400 whole cell patch clamp recordings of LH neurons, obtained while providing a live fly with a battery of 36 odours (Frechter et al, 2019). Using a simple artificial neural network (ANN) we investigate the importance of different types of neuronal interactions within the LH in order to predict the recorded odour response. Predicted and observed topology and edge weights are compared between neuronal reconstructions, and those obtained from the ANN, which show parallels.

(4)

Theta oscillations gate the transmission of reliable sequences in the medial entorhinal cortex

Balachandar AN, Assisi C

Stability and precision of sequential activity in the entorhinal cortex is crucial for encoding spatially guided behavior and memory. These sequences are driven by constantly evolving sensory inputs and persist despite a noisy background. In a realistic computational model of a medial entorhinal cortex (MEC) microcircuit, we show that intrinsic neuronal properties and network mechanisms interact with theta oscillations to generate reliable outputs.
In our model, sensory inputs activate interneurons near their most excitable phase during each theta cycle. As the inputs change, different interneurons are recruited and postsynaptic stellate cells are released from inhibition. This causes a sequence of rebound spikes. The rebound time scale of stellate cells, due to an h-current, matches that of theta oscillations. This fortuitous similarity of time-scales ensures that stellate spikes get relegated to the least excitable phase of theta and the network encodes the external drive but ignores recurrent excitation. In contrast, in the absence of theta, rebound spikes compete with external inputs and disrupt the sequence that follows. Further, the same mechanism where theta modulates the gain of incoming inputs, can be used to select between competing inputs to create transient functionally connected networks. Our results concur with experimental data that show, subduing theta oscillations disrupts the spatial periodicity of grid cell receptive fields. In the bat MEC where grid cell receptive fields persist even in the absence of continuous theta oscillations, we argue that other low frequency fluctuations play the role of theta.

(5)

Spontaneous spiking homeostasis as a ROS-regulating ATP-excursion

Chintaluri C, Vogels T

Reactive oxygen species (ROS) is largely generated during mitochondrial respiration. It's well known to be toxic to many cellular processes, and it's generated under two conditions. Firstly, under heavy mitochondrial respiratory events, when the rate of ATP production is high, a portion of the electrons in the electron transport chain naturally produce what we term "exhaustive" ROS. This can happen in neurons that are ATP-depleted and nearing exhaustion, for instance when firing rapidly or undergoing excitotoxicity. On the other hand, in neurons with low ATP demand, e.g., during sensory deprivation, electrons lingering in the transport chain jump the production line to generate ROS directly. In this case, a cell is ATP-rich and devoid of any ATP consuming processes, and therefore subjected to what we call "energetic" ROS. We postulate that neurons which are synaptically quiescent, are ATP-rich, presumably in anticipation of energy-demanding recovery from spiking and thus especially vulnerable to this energetic-ROS condition. Neurons can thus experience high levels of ROS in highly active periods as well as during prolonged quiescence.

Intracellular ROS, in addition to being toxic, also mediates complex changes in ion channel function, but it's unclear how these changes are orchestrated. Here, we show through modelling that energetic-ROS facilitates changes in the ion channels to favour a spontaneous spiking event. Such action potentials elicit ATP consumption due to Na+/K+ ion gradient maintenance. This momentarily restores ATP production and reduces intracellular ROS. For quiescent neurons, energetic ROS may thus mediate fast homeostatic increases in spontaneous activity, potentially with many downstream effects. Intriguingly, exhaustive-ROS can trigger complementary changes in the family of ion channels that lead to the minimisation of spurious spikes. From our model emerges a comprehensive picture of an intrinsic, ROS-mediated spike rate homeostasis that can account for the fast timescales of spike regulation.

(6)

Neural Activity Reveals Interactions Between Episodic and Semantic Memory Systems During Retrieval

Weidemann C, Kragel J, Kahana M

Whereas numerous findings support a distinction between episodic and semantic memory, it is now widely acknowledged that these two forms of memory interact during both encoding and retrieval. The precise nature of this interaction, however, remains poorly understood. To examine the role of semantic organization during episodic encoding and retrieval, we recorded intracranial encephalographic signals as 69 neurosurgical patients studied and subsequently recalled categorized and unrelated word lists. Applying multivariate classifiers to neural recordings, we were able to reliably predict encoding success, retrieval success, and temporal and categorical clustering during recall. By assessing how these classifiers generalized across list types, we identified specific retrieval processes that predicted recall of categorized lists and distinguished between recall transitions within and between category clusters. These results particularly implicate retrieval (rather than encoding) processes in the categorical organization of episodic memories.

(7)

SpikeInterface: A standardized framework for sorting, analysis, and evaluation of extracellular recordings

Hurwitz C, Buccino A, Magland J, Gerlei K, Nolan M, Hennig M

Recent breakthroughs in microelectronics have enabled high precision extracellular recording of thousands of neurons both in vitro and in vivo. While the increased data volume and complexity offers unprecedented opportunities for understanding brain function, it also heightens the need for standardized, reproducible analysis techniques. In particular, inferring the activity of single neurons from raw recorded signal, a blind source separation problem called spike sorting, is an essential preprocessing step in many electrophysiological studies, yet remains largely non-standardized across labs. Since sorting errors can affect subsequent analysis, it is important to establish community-accepted standards by validating and comparing preexisting spike sorting algorithms. Traditionally, running and comparing multiple spike sorting software packages is a laborious process, complicated by file format and algorithm differences. To enable straightforward validation and comparison, we developed SpikeInterface (https://github.com/SpikeInterface), a simple framework for extracting and analyzing relevant information from both raw and spike-sorted extracellular datasets of any established file format. SpikeInterface was designed to standardize how data is retrieved from files, rather than how it is stored, allowing users to access, sort, and analyze extracellular datasets with the same tools, regardless of the underlying file format. This design paradigm also allows for wrapping preexisting spike sorting algorithms within SpikeInterface, making the process of running and comparing multiple spike sorters accessible. In this paper, we demonstrate the ease-of-use and benefits of SpikeInterface by analysing tetrode recordings from the entorhinal cortex of awake behaving mice. We sort the data with three popular spike sorting algorithms and perform an analysis of spatial firing properties, comparing the functional differences between agreed upon units. Taken together, we present an easy to use framework for improving the reliability, reproducibility, and accessibility of analyzing extracellular recordings.

(8)

Sensory Processing and Categorization in Cortical and Deep Neural Networks

Pinotsis D

Many recent advances in artificial intelligence (AI) are rooted in neuroscience. For example, a large body of AI work has used ideas from neuroscience to build neural networks that can perform sensory information processing and reinforcement learning tasks. However, ideas from more complicated paradigms like perceptual decision making tasks are less used. Although automated decision making systems are ubiquitous in modern applications (driverless cars, pilot support systems, medical diagnosis algorithms etc.) achieving human-level performance is still a challenge. Humans can effortlessly accomplish decision making tasks that are deemed difficult for artificial systems. Thus, understanding complex decision making dynamics in the brain and modeling them using deep neural networks could open new avenues to tackle these difficulties. Here we used multivariate methods and deep recurrent neural networks to model some of the complex neural interactions during sensorimotor decision making in primate brain. We investigated how brain dynamics flexibly represented and distinguished between sensory processing and categorization in two different sensory domains: motion direction and color. We found that representations changed depending on context. Selectivity of each brain area depended not only on the stimulus represented but also on the domain of categorization. We trained deep recurrent neural networks with monkey LFP recordings and found that they could process sensory information and perform categorization in the motion and color domains similarly to the animals. By comparing brain dynamics with network predictions, we found that computations in different brain areas also changed flexibly between the two tasks. Color computations appeared to rely more on sensory processing, while motion computations more on categorization. Overall, our results shed light to the biological basis of categorization and differences in selectivity and computations in different brain areas.

(9)

Hybrid synaptic plasticity modelling via integration of electrical and biochemical dynamics via KappaNEURON

Linkevicius D, Sterratt DC

Synaptic plasticity is a fundamental property of the brain, essential for learning and memory formation. Understanding of the biophysical mechanisms of synaptic plasticity is still lacking, in part, due to a complex interplay between factors such as intracellular biochemical networks, neurotransmitters, membrane receptors and dendritic morphology. Further understanding of synaptic plasticity requires integration of these various levels into coherent and consistent models. We combine various experimental literature and models into to a bottom-up model of synaptic plasticity by using KappaNEURON - a novel software package. This approach allowed us to show how various experimental synaptic plasticity protocols manifest themselves at the subcellular level and to make novel experimental predictions.

KappaNEURON is a hybrid simulator which couples NEURON 7.4 and SpatialKappa 2.1.4 in a consistent and accurate manner. KappaNEURON has the advantage of allowing biochemical interactions to be specified in the rule-based Kappa language, which is particularly suited to describing protein-protein interactions compactly. Relevant biochemical reactions collected from literature were translated into SpatialKappa's rule-based format, with free or under-constrained reaction rates being tuned to match physiological data. The resulting biochemical spine model, embedded in a CA1 pyramidal cell model, includes calcium influx through NMDA receptors and voltage-gated calcium channels and Ca2+-dependent downstream protein cascades.

The simulations incorporating a single biochemically complex spine, in the context of a detailed electrical model of a neuron, allow for deeper understanding of naturally occurring plasticity, as well as pathological plasticity mechanisms, such as in Alzheimer's or Parkinson's. Novel proteomics studies, if extended to hippocampus in health and disease, would provide invaluable protein-protein interaction and genetic information, which might be translatable to parameter changes or model extensions, representing gene level phenomena.

(10)

The Synfire/deSync Model: Deciphering Episodic Content from Cortical Alpha Oscillations

Parish G, Michelmann S, Hanslmayr S, Bowman H

A review of findings from the oscillations and memory field has argued that cortical alpha oscillations regulate our online experience, where neural desynchronisations are thought to signal information flow (Hanslmayr, 2012). In support of this, a previous model has shown how desynchronisations caused by differential rates of neuronal firing can indicate when an item is being actively represented (Parish, 2018). Here, we present a novel computational model that expands this notion with the addition of a complimentary timing mechanism, thus enabling the emergence of qualitative temporal information during the active representation of a pattern of items. As such, we recreate findings from a recent episodic memory paradigm, where a unique temporal signature of a memory can be identified from a period of desynchronisation (Michelmann, 2016). This is achieved through the analysis of the phase-reset patterns of cortical alpha oscillations, such that a representational-similarity-analysis (RSA) can subsequently dynamically detect the re-instatement of content specific memories. To replicate these findings, we expand the use of feed-forward computational solutions in the encoding of time (Itskov, 2011), providing a neural substrate mechanism that might enable the contextualisation of multiple and simultaneous hierarchical neural time-keepers (Mauk, 2004). We also instantiate a cortical network that is entrained in an alpha frequency through recurrent excitatory-inhibitory interactions, where event-driven activation desynchronises on-going oscillatory activity in a phase-reset. Each event is then bound as a discrete moment in time, similar to other models (Bowman, 2007), in a complimentary learning systems (CLS) framework (O'Reilly, 2011). By recalling events in temporal order, we show how one can dynamically detect the encoding and retrieval of unique temporal signatures within the brain. We hope to provide computational evidence for the findings from Michelmann, (2016) and for the hypothesis that cortical desynchronisations can indicate information flow and can be used to decipher information content.

(11)

MIIND: A simulator for neural populations

Osborne H, Deutz L, Lai YM, de Kamps M

MIIND is a neural simulation development environment which utilises a geometric population density technique to model a network of interacting neural populations (de Kamps, Lepperød & Lai, 2019). The density technique builds a geometric mesh or grid which describes the deterministic dynamics of a neuron model. Assuming input spikes are Poisson distributed, MIIND then simulates the probabilistic behaviour of a population of such neurons. By describing the evolution of a probability density function across state space, metrics such as the average firing rate can be derived which show good agreement with direct simulation techniques. As would be expected, reducing the number of cells in the mesh, improves simulation times but increases the derived error. However, even with very granular meshes, the general behaviour of the population is preserved, so MIIND can be used to quickly prototype complex networks of populations while significantly reducing the pain of parameter tweaking. MIIND has recently been used to simulate a spinal interneuron and motor neuron population network model. The model supports experimental evidence of the proprioceptive effect on muscle recruitment during an isometric leg task and implicates spinal circuitry as a site of so called muscle synergy encoding. MIIND has also been used to reduce a Hodgkin Huxley style bursting neuron model to a simple two dimensional system. Using this reduced model as a basis, an example of a complex central pattern generator circuit can be reproduced in MIIND, displaying experimentally confirmed behaviours in fictive cat locomotion. This presentation introduces the MIIND population density technique and its usage, and demonstrates examples of its utility for supporting experimental investigations.

De Kamps, M., Lepperød, M. and Lai, Y.M., 2019. Computational geometry for modeling neural populations: From visualization to simulation. PLoS computational biology, 15(3), p.e1006729.

(12)

Modelling of retinal mosaic formation using an agent-based model

de Montigny J, Sernagor E, Bauer R

Individual retinal cell types exhibit semi-regular spatial patterns called retinal mosaics. These mosaics, enabling uniform sampling of visual information, are formed to varying degrees across cell subtypes. Retinal ganglion cells (RGC), amacrine, horizontal and photoreceptor cells are known to exhibit such layouts. In addition to cell body mosaics, dendritic arbours also form mosaics - dendrites of homotypic cells exhibit avoidance while dendrites of different types overlap.

Mechanisms responsible for the formation of such organised structures are still not well understood and they follow three main theories. (1) Homotypic cells prevent nearby cells from adopting the same type. (2) Cell tangential migration (CM), with homotypic cells repulsing each other. (3) Cell death (CD), with specific cell types (mainly RGCs) exhibiting high rates of apoptosis, increasing spatial regularity.

Here, we use BioDynaMo, an agent-based simulation framework, to build a detailed and mechanistic computational model of mosaic formation in 3D physical space. In particular, we investigate the implications of the three theories and their combinations, then compare them to experimental data we obtained in mouse.

We show that the CM mechanism yields the most regular mosaics. Moreover, the longest cellular migration distance achieved in these simulations is in agreement with experimental observations. We also found that CD can create regular mosaics only if the death rate is kept between 25 to 55%. After half of cells have died, CD appears to have a negative impact on RI. This result matches our experimental findings, reporting a decrease of RI after half of cell have died.

(13)

Interacting cortical subnetworks for first-order and higher-order sensory processing

van Rheede J, Buchan M, Gothard G, Akerman C

The neuronal circuitry of the cortex underlies the mammalian ability to show flexible and adaptive behaviour in complex, dynamic environments. This requires the cortex to process incoming stimuli simultaneously with their contextual information and associated learning signals. In line with this, anatomical and functional experiments have distinguished 'first-order' and 'higher-order' inputs to cortex, originating from different thalamic nuclei. First-order information is relayed from the sensory periphery and drives basic sensory responses, while higher-order information originates from cortex and can modulate the local sensory response or act as a learning signal to gate sensory plasticity. However, is not known how the circuitry of the cortex allows for the convergence of first-order and higher-order signals, or how these information streams interact.

We have identified two intermingled subnetworks of cortical neurons that are derived from different cortical progenitor cell types, and established that these differ in the amount of first-order and higher-order input they receive. To dissect the cortical integration of these different types of information, we used optogenetic interference during electrophysiological recordings in mouse barrel cortex, measuring each subnetwork's contribution to neural activity driven by whisker stimuli. We report that driving activity in either subnetwork elicits a different profile of excitation across cortical input, processing and output layers. Moreover, the first-order and higher-order subnetworks have different effects on the neuronal response to the sensory stimulus. Notably, the higher order subnetwork appears to modulate long-term potentiation (LTP) of the whisker response, and selective activation of this subnetwork at a physiological whisking frequency may be sufficient to drive LTP.

Our findings support a model of sensory cortex in which parallel but interacting cortical subnetworks are biased to process first-order vs. higher-order information. The conditions for eliciting activity in the higher-order circuit may be of particular relevance for understanding what drives learning in sensory cortex.

(14)

A Neurally-Informed Modelling Approach for Investigating Cross-Modal Associations in Perceptual Decision-Making

Bolam J, Ince R, Boyle S, Delis I

Our brains process multisensory information by formulating cross-modal associations between stimulus properties from different modalities. Task-dependent manipulations of such properties have been found to modulate cross-modal associative effects on behavioural performance. For example, humans associate high-pitch auditory tones with small visual objects, and vice versa, which is known as the auditory pitch-visual size association effect. Congruent presentations of such cross-modal stimuli have been shown to increase accuracy and reduce reaction times (RTs) in perceptual decision-making tasks. However, which decision-related processes are dynamically modulated by cross-modal associations and their neural correlates remain elusive predominantly due to the mixed selectivity of the corresponding neural responses. Here, we sought to investigate the neural representations of cross-modal associations in an auditory pitch-visual size association task. We used a modified version of the Implicit Association Task to obtain measurements of participants' behavioural performance i.e. choice accuracy and RTs, and coupled it with electroencephalography (EEG) to measure the underlying neural activity corresponding to manipulations induced by stimulus features, sensory modality, and congruency. To identify the brain activity patterns that underpin task performance, we employed demixed Principal Components Analysis (dPCA), a supervised dimensionality reduction technique, that decomposed the recorded EEG activity into task-relevant components. This approach yielded distinct stimulus-encoding, modality-encoding and congruency-encoding EEG components consistently across participants. We then integrated these neural components into a Bayesian modelling framework termed the Hierarchical Drift Diffusion Model (HDDM) to assess which decision-related process they may be implicated in. Our preliminary results suggest that the congruency-encoding component significantly modulates early sensory processing, but not late decision-related, mechanisms. Ultimately, our neurally-informed modelling approach demixes the selectivity of EEG signals to identify the neural correlates of latent perceptual and cognitive processes that underlie perceptual decision-making.

(15)

Spectral theory of non-stationary dynamics of integrate-and-fire neurons

Deutz L, Osborne H, de Kamps M

Understanding the behaviour of large recurrent networks of spiking neurons is one of the major challenges in computational neuroscience. A first step towards understanding the dynamical properties of such networks is to determine the location and stability of equilibria and how they depend on the connectivity profile and single neuron properties. Integrate-and-fire neuron models

are widely used for this purpose because they are simple enough to be studied analytically while still being able to capture important dynamical features observed on a single neuron level and in a network context. In the diffusion approximation, the synaptic input is replaced by Gaussian white noise which gives rise to a Fokker-Planck equation (FPE) describing how neurons respond on average to a given input.
Even for simple models, solving the FPE for time-varying inputs is a difficult task. If the temporal modulations of the input are small compared to its stationary baseline, then perturbation theory can be used. In first order approximation, the neuron can be described as a linear filter often referred to as the transfer function. It determines the gain and the phase shift of the neuron response for given frequency components in the input. However, the first order results are restricted to dynamical regimes close to stationary state and can not account for strong temporal modulations in the input. Here, we propose a method to extend the known first order results systematically to higher orders using the method of spectral decomposition which expands the neuron dynamics in terms of eigenfunctions of the Fokker-Planck operator. The higher order corrections could be used to establish a mapping between spiking and rate-based neuron models capable of describing the rate dynamics in non-stationary regimes.

(16)

Constrained plasticity can compensate ongoing drift in neural populations

Rule ME, Loback AR, Raman DV, Harvey CD, O'Leary TS

Recent experiments reveal that neural populations underlying behavior reorganize their tunings over days to weeks, even for routine tasks. How can we reconcile stable behavioral performance with ongoing reconfiguration in the underlying neural populations? We examine drift in the population encoding of learned behaviour in posterior parietal cortex of mice navigating a virtual-reality maze environment. Over five to seven days, we find a subspace of population activity that can partially decode behaviour despite shifts in single-neuron tunings. Additionally, directions of trial-to-trial variability on a single day predict the direction of drift observed on the following day. We conclude that day-to-day drift is concentrated in a subspace that could facilitate stable decoding if trial-to-trial variability lies in an encoding-null space. However, a residual component of drift remains aligned with the task-coding subspace, eventually disrupting a fixed decoder on longer timescales. We illustrate that this slower drift could be compensated in a biologically plausible way, with minimal synaptic weight changes and using a weak error signal. We conjecture that behavioral stability is achieved by active processes that constrain plasticity and drift to directions that preserve decoding, as well as adaptation of brain regions to ongoing changes in the neural code.

(17)

Non-synaptic interactions in the neural encoding of odorants: a good start is half the battle

Pannunzi M, Nowotny T

In many insect species, olfactory receptor neurons (ORNs) are housed in hair-like sensilla in a stereotypical manner. Each sensillum contains two or more ORNs of different types. ORNs within the

same sensillum interact, without synaptic connection, mostly inhibiting each other. As suggested previously, these non-synaptic interactions (NSIs) could be important for insects' capability to resolve concentration ratios or the timing between different odorants at high resolution.
Here we test with a computational model the hypothesis that NSIs could improve the spatiotemporal resolution of odour recognition in mixed odour plumes: If a single source emits an odour mixture (multiple odorants), odorants arrive in close synchronization, NSIs take effect and responses of all ORNs in a sensillum are diminished. If separate sources emit odorants, their concentrations are less correlated, and NSIs have almost no effect, resulting in larger ORN responses. This fast mechanism could provide an essential advantage because behavioural relevant temporal and spatial scales of plumes can be very fine, on the order of tens of milliseconds and tens of millimetres.
First, we generated a model of the early olfactory system of insects - a simple circuit model with two ORNs and their corresponding projection neurons (PNs) and local neurons (LNs) in the antennal lobe (AL) - and reproduced the responses to simple odour stimuli reported in the literature. Second, we analyzed the advantages of having NSIs (compared to not having them) for detecting target odours within complex mixtures.
Therefore NSIs and lateral inhibition may implement two different functions: NSIs have a high spatiotemporal resolution and they generate selective inhibition between ORNs; LN networks take effect later to decorrelate PN activities and normalize them with respect to the average input from ORNs.

(18)

Weight-dependent BCM plasticity leads to a rich family of receptive fields

van Rossum M, Froc M, Williamson O

Despite the success of artificial supervised learning, the rules that underlie unsupervised plasticity are still not well known. The Bienenstock Cooper Munro model of unsupervised synaptic plasticity is one of the leading models for receptive field development in sensory cortices. In the BCM model the amount of
increase (LTP) or decrease (LTD) of the synaptic strength is independent of the synaptic strength itself. In experiments, however, strong synapses are found to be easier to depress than weak ones, whereas for potentiation there is no such effect. This 'soft-bound' plasticity naturally caps the unlimited run-away
plasticity often present in models. Here we explore the effect of including the experimental phenomenon of weight-dependence on the outcome of BCM learning.
We simulate the common situation where a set of input patterns is repeatedly provided to a single neuron and study the possible steady states of the weights. We find that weight-dependence leads to additional fixed points in the learning dynamics. This is surprising as weight-dependence usually leads to a reduction of fixed points. When the neuron has just two inputs, inclusion of weight-dependence
increases the number of fixed points from 2 to 4. The 2 new fixed points are more robust noise. More generally, the number of fixed points grows exponentially in the number of inputs, so that for a realistic number of inputs there is an enormous diversity in the number of stable weight

configurations. The emerging receptive fields are selective but less so than regular BCM. Other properties of the BCM rule, such as the need for a rapidly updating threshold, as well as
the response to monocular deprivation experiments are largely unaffected by weight-dependence. In conclusion we find that this more realistic BCM model substantially enriches its behavior compared to vanilla BCM, while maintaining many of its characteristic properties.

(19)

Variability and regulation of reinforcement learning processes in rats

Cinotti F, Girard B, Khamassi M

In this work, I will discuss two main topics : firstly metalearning, or regulation of learning and decision-making parameters ; secondly, inter-individual variability in the strategies used in a simple Pavlovian conditioning experiment. In both cases, I will adopt a behavioural standpoint using reinforcement learning algorithms to model experimental data while also attending to related dopaminergic functions in the rat brain.
In order to adapt to a changing environment, an organism should be capable of regulating its own behavioural parameters, a process known as metalearning. In particular, regulating the tendency to explore potentially better options or the rate at which action values are updated might constitute two powerful metalearning strategies. We will start our investigation of metalearning by assessing how different metalearning models can explain long-term changes of rat behaviour in a three-armed bandit task. Additionally, it was recently proposed (Humphries et al. 2012) that dopamine plays a modulatory role on the exploration-exploitation trade-off, and we verified this claim by careful modelling of an experimental task devised to test this idea.
Finally, in a Pavlovian conditioning task in which the appearance of a lever predicts food delivery, it is well known that two kinds of behaviour can appear in a rat population. On the one hand, so-called sign-trackers become strongly attracted to the lever which they will approach and nibble, while goal-trackers will prefer to immediately go to the site of reward delivery. In parallel, there are differences in the associated dopamine signals, sign-trackers presenting a classical reward prediction error pattern of signalling, contrary to goal-trackers. A model attempting to explain this behaviour in terms of competing model-free and model-based learning strategies was previously proposed and we studied its response to varying inter-trial intervals.

(20)

Generalisation of Frequency Mixing and Temporal Interference through Volterra Expansion

Perez Nieves N, Goodman D

It has been recently shown that it is possible to use sinusoidal electric fields at kHz frequencies to enable focused, non-invasive, neural stimulation at depth by delivering multiple electric fields to the brain at different frequencies (f1 and f2) that are themselves too high to recruit effective neural firing, but whose difference f1-f2 is low enough to drive neural activity. This is called temporal interference (TI)(Grossman,2017).

However, for frequencies below 150Hz generates activity at the difference and sum of the frequencies due to the non-linearity of the spiking mechanism in neurons, which is called frequency mixing (FM)(Haufler,2019). However, it remains unknown why at higher frequencies we still see activity at the difference with no activity present at any other frequency.

We propose modelling the non-linearity using a Volterra expansion. First, we show that any order P non-linear system stimulated by N sinusoids outputs a linear combination of sinusoids at frequencies given by the linear combinations of the original frequencies with coefficients m,n=±{0,1,...,P} and amplitudes depending on the P-dimensional Fourier transform of the Pth-order kernel of the expansion.

We simulate a population of LIF-neurons stimulated by sinusoidal currents at f1 and f2 and record the average firing rate. For low frequencies, we see all combinations nf1+mf2 as in Haufler,2019. For high frequencies, we only find f1-f2 as in Grossman,2017.

Computing the 2D Fourier transform of the second order kernel shows that for high frequencies, only coefficients corresponding to f1-f2 are high enough to generate a response in the network, thus explaining TI stimulation. For low stimulation frequencies, all coefficients are high enough to produce a response at all nf1+mf2.

We generalised previous experimental and theoretical results on TI and FM. Understanding TI stimulation mechanisms will facilitate clinical adoption, help develop improvement strategies and may reveal new computational principles of the brain.

(21)

Simple vs complex: how complexity affects causal inference in visual perception

Gekas N

An important aspect of perceptual processing is causal inference, i.e. the process of inferring whether or not an event A is caused by another event B. An observer infers that there is a hidden variable (e.g. the true orientation of a grating) based on an observed variable (a noisy measurement). It has been suggested that observers have a tendency to prefer simple models that are consistent with the data over complex models (Gershman & Niv, 2013) but the properties of this “simplicity bias” are still unclear. Here, we use psychophysical and computational methods to investigate the relation between statistical complexity and causal inference. Observers are presented with a coloured grating and are asked to report the orientation or the colour of the grating. The frequency of each of the two features of the stimuli follow two bimodal distributions, which interact to produce certain combinations of features more frequently than others. The observers learn to associate specific colours with specific orientations, and vice versa, and infer a number of underlying causes and their associated parameterisations in regards to the sensory data. The distance between the means of the bimodal distributions is systematically manipulated, and a Bayesian model is used to show that inference of a simple or complex model of the stimulus statistics correlates directly with the distance between the means of the distributions approximating a logistic function.

(22)

Robust reconstruction of mouse poses during visual experiments

Storchi R, Cootes TF, Killick R, Lucas RJ

Quantifying changes in mouse postures as a function of changes in the visual scene is a critical step for understanding visual processing and learning. However a robust method to detect postures from specified body landmarks in freely moving animals is currently missing. Here we propose a solution to this problem by combining convolutional neural networks, 3D reconstruction from multiple camera views and active shape models.

First we use convolutional networks to estimate body landmarks from individual cameras. From these landmarks we then generate a set of candidate 3D reconstructions of the mouse postures. Finally these postures are scored and regularised by using active shape models to generate a single robust estimate of the mouse posture.

We apply this algorithm to a dataset of innate and conditioned behavioural responses to a variety of behaviourally salient stimuli. Our results show that the proposed algorithm, combined with time series segmentation and classification allows for extraction of a remarkable diversity of visually evoked behaviours.

(23)

Firing rate of neurons with dendrites, soma and axon in the fluctuation-driven, low-rate limit

Gowers R, Timofeeva Y, Richardson M

Neurons in-vivo are subject to synaptic input with highly variable pulse strength and arrival time. Integration of this stochastic synaptic drive has been modelled extensively since the late 1960s, with significant analytical progress made in how class-specific membrane and synaptic properties affect neuronal integration. Due to tractability, most analytical work approximates neurons as electrotonically compact or comprising a small number of discrete connected compartments.

We consider continuous models of neurons comprising dendrites, soma and axon with spatially-distributed fluctuating synaptic drive. We demonstrate that in the fluctuation-driven, low firing-rate limit, the steady-state firing rate can be approximated using a level-crossing approach. This low firing-rate limit is applicable to pyramidal neurons, for which average firing rates has been observed experimentally to be low compared to the reciprocal of the membrane time constant.

We apply this approach to very simple neuronal morphologies, first demonstrating that certain dendritic morphologies have firing-rate functions of the input drive that are independent of the electrotonic length, yet distinct from the functions found for point-like models. Second, when an axon is added, we demonstrate that the firing rate varies non-monotonically with the axonal radius, with the peak firing rate corresponding to a radius similar to that found for pyramidal cells. Third, we observe that adding dendrites with fluctuating drive does not always increase the firing rate of the neuron. Finally, we show that soma size and the position of spike-initiation in the axon alter both of these non-monotonic relationships, with a larger soma resulting in the firing rate being maximised for a higher number of dendrites.

Though applied to "toy models" of spatio-temporal synaptic integration, these initial findings show it is possible to obtain analytical results in the low-rate limit that capture the effects of spatially-distributed synaptic drive.

(24)

Deconstructing the protomap versus protocortex hypotheses via computational modelling

James S, Whiteley D, Brooke J, Wilson S

Patterns of cortical arealization self-organize through mechanisms of embryonic and early postnatal development that are highly robust, and which are coordinated by genetic processes that are highly conserved within and between mammalian species. The processes by which cortical arealization patterns emerge are often described from one of two broad perspectives. From the "protomap" perspective, the area fate of cortical cells is determined a priori, through the interplay of genetic influences. From the "protocortex" perspective, cortical cells grow into an undifferentiated tissue, and their area fate is later established as the outcome of cooperative and competitive mechanisms between cortical cells that may be differentially driven by afferent, i.e., thalamocortical, input. Adopting a view of cortical development from either of these extremes can be productive, but neither can account satisfactorily for all of the available data. Real arealization patterns almost certainly reflect a combination of afferent-driven competitive processes within the cortex (protocortex) as well as a more precise coordination under genetic control (protomap), and the interaction of these processes is likely to reflect dynamics that occur across a range of spatial (from cells to tissues) and temporal (developmental and evolutionary) scales. Understanding the potentially complex interaction of these complicated multiscale processes is a key challenge for the study of cortical evolution and development. We are using computational modelling to formalise the distinction between these processes, and to facilitate a systematic investigation of their relative contribution to cortical patterning. Specifically, we present i) testable predictions about the effect of boundary shape on spontaneous pattern formation within cortical fields, ii) a method for identifying gene interaction networks that generate protomaps consistent with those observed in gene knock-out experiments, and iii) a reaction-diffusion style model of competitive interactions between thalamocortical projections under the guidance of molecular signalling cues.

(25)

Degenerate Optimal Boundaries for Multiple-Alternative Decision Making

Baker S, Griffith T, Lepora N

Integration to threshold models of 2-choice perceptual decision making have guided and developed our understanding of human and animal behaviour and underlying neural processes for decades. While such models seem to apply equally well to multiple-choice decision making, they have yet to converge to a general normative framework, and the implications for and of thresholds characteristics has not been explored. Here, we develop such a framework using sequential Bayesian inference and the conceptualisation of decision making as a particle diffusing in n-D. We show that optimal decision thresholds comprise a degenerate set of complex structures and speed-accuracy tradeoffs, contrary to current 2-choice results. Such thresholds support both stationary and collapsing boundaries as optimal strategies for decision-making, both of which result from stationary, complex threshold representations. Our findings thus expand the much-needed theory for multiple-choice decision making, provide a characterisation of optimal decision thresholds under this framework, and contributes to the debate between stationary vs dynamic decision thresholds for optimal decision making.

(26)

A biologically grounded model of the hypothalamic control of motivated behaviour

Jimenez-Rodriguez A, Schmidt R, Gilmour W, Wilson S

The hypothalamus is an evolutionarily conserved brain structure that plays a central role in the initiation and regulation of goal-directed behaviour in all vertebrate species. The lateral hypothalamus has been a particular focus for recent research, due to its involvement in functional pathways for the representation and management of internal states, and for its interactions with other motivation-related structures, such as the ventral tegmental area (VTA) and the nucleus accumbens (NA). With a view to modelling the neural bases of motivated behaviour, we here consider the functional role of the lateral hypothalamus in managing the energetic costs of autonomic and behavioural homeostasis. We present a computational model of hypothalamic circuitry, which includes a representation of the interaction between lateral hypothalamus and dopamine signalling in VTA and NA. We demonstrate that the model can be used to control the regulatory behaviour of simple simulated agents, and account for a range of data from experiments that link hypothalamus activation to motivated behaviour. To our knowledge, this is the first model of the control of goal-oriented behaviour that is constrained by anatomical and physiological descriptions of hypothalamic circuitry.

(27)

Frequency-sensitivity and magnitude-sensitivity in decision-making: a model-based study

Bose T, Bottom F, Reina A, Marshall J

We theoretically study decision-making behaviour in a model-based analysis
related to binary choices with pulsed stimuli. Assuming a strong coupling between external stimulus and its internal representation, we argue that the frequency of external periodic stimuli represents an important degree of freedom in decision-making which may modulate behavioural responses. We consider various different stimulus conditions, including varying overall magnitudes and magnitude ratios as well as varying overall frequencies and frequency ratios, and different duty cycles of the pulsed stimuli. Decision time distributions, mean decision times and choice probabilities are simulated and compared for two different models -- a leaky competing accumulator model and a diffusion-type model with multiplicative noise. Our results reveal an interplay between the sensitivity of the model systems to both frequency and magnitude of the stimuli. In particular, we

find that periodic stimuli may shape the decision time distributions resulting from both models by resembling the frequencies of the pulsed stimuli.

We obtain significant frequency-sensitive effects on mean decision time and choice probability for a range of overall frequencies and frequency ratios. Our simulation analysis makes testable predictions that frequencies comparable with typical sensory processing and decision-making timescales may influence choice and response times in perceptual decisions.

(28)

Does structure in neural activity match anatomical structure?

Delaney T, O'Donnell C

Information in the brain is carried in correlated network activity. Previous research has established that these correlations play a crucial role in representing sensory information. For example, the onset of visual attention has been shown to have a greater effect on the correlations in the macaque V4 than on the firing rates in that region. In order to understand the representation of sensory information we must understand the interactions between neurons.

Because of limitations in recording technology almost all research has explored correlations between neurons within a given brain region. Relatively little is known about correlations between neurons in different brain regions. However, the recent development of 'Neuropixels' probes has allowed extracellular voltage measurements to be collected from multiple brain regions simultaneously, routinely, and in much larger numbers than traditional methods. In this project we used a publicly-available Neuropixels dataset to analyse correlations between different brain regions.

Using two probes, spiking activity was simultaneously collected from over 800 neurons in the brain of an awake mouse for a period of 84 minutes. During this period, the mouse and was shown various visual stimuli. The 800 neurons were distributed across 5 different brain regions: V1, hippocampus, thalamus, motor cortex, and striatum. Using these data, we measured single neuron firing rates across all regions. We also measured pairwise correlations and mutual information between neurons within the same region, and between neurons in different regions. We compared the distribution of these measurements across regions.

We then used spectral clustering to cluster the neurons based on the pairwise correlations and mutual information. We compared the structure of these cluster networks to the anatomical structure of the cells.

(29)

The Interplay Between Somatic and Dendritic Inhibition Promotes the Emergence and Stabilization of Place Fields

Pedrosa V, Clopath C

The hippocampus plays a key role in the encoding of spatial representation. During the exploration of novel environments, place fields are rapidly formed in hippocampal CA1 neurons. We propose a computational model of the hippocampal CA1 network to describe the formation, dynamics and stabilization of place fields. We show that although somatic disinhibition is sufficient to form place cells, dendritic inhibition along with synaptic plasticity is necessary for stabilization. Our model suggests that place fields stabilize through exploration due to large excitatory synaptic weights and large dendritic inhibition. We show that the interplay between somatic and dendritic inhibition balances the increased excitatory weights, so that place cell firing rate returns to baseline levels after exploration. Our model suggests that different types of interneurons are essential to unravel the mechanisms underlying place field plasticity. Finally, we predict that artificial induced dendritic events can shift place fields even after place field stabilization.

(30)

Prediction of Behavioral Responses Using Bayesian Model from Multi-Unit Spike Recordings in Sensorimotor Cortex of Rats

Öztürk S, Güçlü B

Operational performance of cortical neuroprostheses directly affects the quality of life for the patients with tetra/paraplegia, amputees or different types of neurological disorders. To be able to feed the most naturally encoded signal back into the cortex, spiking activity has to be decoded into its input/output components in natural life, as stimulus into the black box and behavioral response from it. For this study, 2 rats were trained to perform yes/no task based on a mechanical stimulation applied on hind paws (frequency: 40-Hz sinusoidal vibrotactile stimuli, duration: 0.5 s, amplitude: 200 m). After stabilized learning on ROC, they are implanted with 16-channel microwire array electrodes and multiunit spike data were collected during behavioral decision making task. After sorting spikes with K-means (clusted Lratio < 0.5), average spike rates for pre-stimulus (Rb) and the duration of stimulus (Rd) were calculated. For Rb and Rd comparisons, expectation of stimulus were observable by inhibition based manner, as significant decrease of rate even in stimulus-off trials. In general, more units were recruited during Rd on hit and false alarm trials than miss and correct rejection trials, which advocates incorrect trials for the sake of animal. There are significant correlations between Rd and false alarm rate (Rat #1: r=0.49, p=0.047), and between Rd-Rb and hit rate (Rat #2: r=0.58, p=0.038) which may lead the model to predict behavioral output. By feeding the Bayesian Model with prior stimulus condition (on/off) and two dimensional feature set (Rb and Rd) for average spike rate, psychophysical prediction accuracies were generated for 2 rats with logistic regression (stimulus-on trials, Rat #1: 73%, Rat #2: 75%; stimulus-off trials, Rat #1: 65%, Rat #2: 80%). This model will run on trial by trial basis in real-time to mimic daily needs of neuroprosthetics.

Poster session B

(1)

Computational model of sodium Nav1.7 channel activation in nociceptive dorsal root ganglion neurons

Capurro A, Thornton J, Armstrong L, Sernagor E

Chronic pain is a global healthcare problem, affecting near 14 million people in the UK [1]. Given its significant human and socio-economic burden, there is an urgent need to develop more precise and effective forms of medication for pain management. In this study we focus on the Nav1.7 voltage gated sodium channel, which acts an as amplifier of the receptor potential in the nociceptive neurons of the dorsal root ganglion (DRG) [2, 3]. We present a conductance based model of the DRG neurons and explore its behaviour for different parameters related to the voltage dependence of the Nav1.7 channel activation. This refers to a pathological condition known as erythromelalgia [4], caused by mutations that increase the excitability of sensory neurons by decreasing the level of depolarization needed to activate the channel. We use the computational model to analyse the effects of channel blockers used for the treatment of chronic pain. Finally, we relate the results of the numerical simulations with experimental recordings of human induced pluripotent stem cell derived DRG neurons performed using multielectrode arrays.

[1] Van Hecke et al. (2013) Br J Anaesth 111 (1): 13-18.
[2] Rush et al. (2007) J Physiol 579 (1): 1-14.
[3] Dib-Hajj et al. (2013) Nat Rev Neurosci 14: 49-62.
[4] Cao et al (2016) Sci Transl Med 8, 335ra56335ra56.

(2)

Neural processing of sentences

Burroughs A, Houghton C, Kazanina N

How the human cognitive system is able to comprehend language has been a matter of recent debate. The brain may generate meaning by using learned grammatical rules to decompose sentences into a hierarchy of syntactic structures. On the other hand, it may rely on simpler, statistical methods that rely on sequential processing. Cortical activity tracks speech at the rate of syllable, phrase and sentence presentation. This has been used as evidence for the language systems ability to follow a rule-based grammar during sentence comprehension. However, the same results could be generated by a language system that relies only upon sequences of word representations, without the need for hierarchical phrase-structure building. Our goal is to shed light on the extent to which hierarchical structure plays a role in language processing during sentence comprehension. We use human EEG to record from participants as they listen to streams of four word sentences. We confirm that cortical activity does synchronise with the rate at which syllables, phrases and sentences are presented. Sentence stimuli that contain both appropriate grammar and semantic information generate a full brain response at the rate of phrase presentation. This response is reduced if either grammar or semantic information is absent. In sentences that lack both correct grammar and meaningful semantics, the response at the phrasal rate disappears. This suggests that both syntax and semantics are important contributors during sentence processing, as opposed to just syntax alone. In conditions where word order is manipulated, cortical activity still entrains to the frequency at which sentences were presented, probably due to the repetition of words from the same lexical category at the same position within each sentence. Our results are therefore in support of a language system that relies on sequential, as opposed to hierarchical rule-based, processing.

(3)

The statistics of optimal decision-making: Exploring the relationship between signal detection theory and sequential analysis

Griffith T, Baker S, Lepora N

Perceptual decision-making is the process by which animals use sensory information to make inferences about the world and select actions. Sometimes actions result in rewards that are subjectively experienced and evaluated by the animal. The typical goal of the animal is to maximize their reward by selecting optimal actions.

Perceptual decision-making is difficult in practice because real perceptual systems, inferences, action selections and reward outcomes are all subject to uncertainty. Understanding how animals overcome this challenge is critically important, since doing so will provide an explanation of animal behaviour, as well as a design basis for autonomous decision-making systems.

The leading perceptual decision-making models for binary-choice tasks in psychology and neuroscience are signal detection theory (SDT) and the sequential probability ratio test (SPRT). Although SDT and SPRT are routinely used, it may not be widely understood how closely related they are. In addition, despite these models' successes in explaining behavioural data, our understanding of how they could be implemented in biological hardware is incomplete.

Here, we review the SDT and SPRT algorithmic approaches to binary-choice perceptual decision-making. In doing so, we make clear that SDT and SPRT differ only in their sampling procedures, and so can be viewed as static and dynamic variants from the same family of hypothesis tests. We also present methods for finding the map between decision-task cost structures and optimal decision-parameters for SDT and SPRT. Further analyses of a special case show that animals can make optimal decisions mechanistically by testing linear combinations of raw sensory observations against optimal thresholds. The optimal thresholds encode three important model parameters, the priors, the error costs and the sampling cost. This is suggestive of a plausible biological implementation for the probabilistic SDT and SPRT methods of binary-choice perceptual decision-making.

(4)

Learning, extinction and re-learning: a model for Drosophila olfactory plasticity

Menzat B, Londsdale R, Felsenberg J, Waddell S, Vogels T

The fruit fly can form sparse representations of various odours and learn odour-specific associations such as the history of positive or negative experiences, in a brain area called the mushroom body.

Here, we introduce a four-layer spiking model of the fruit fly's olfactory system to investigate the mechanisms behind sparse odour representation and associative learning. We tune our model to elicit realistic firing activity and show that the interplay between multiple neural mechanisms (lateral inhibition in the antennal lobe and feedback inhibition in the mushroom body) produces temporal and spatial sparseness of odour representations. We test the robustness of our model with a dataset of olfactory receptor neurons responses to 110 different odours and show that our model network can form sparse, decorrelated representations of the odours in the dataset.
To study the plasticity mechanism behind aversive learning and extinction (odour w/o expected shock) as seen in recent experimental results (Felsenberg et al., Cell, 2018), we introduce a dopamine-modulated learning rule. High levels of dopaminergic activity paired with an odour trigger depression of synapses between odour coding Kenyon Cells and Mushroom Body Output Neurons (KC-MBONs). We show that a depression-only learning rule is sufficient to train the odours in our dataset, and we can classify them as either naive, aversive or extinguished, depending on training protocol. We also add a mechanism for strengthening KC-MBON synapses at low levels of dopaminergic activity to enhance the fly's ability to re-learn aversive associations after extinction.
Our model predicts that KC-MBON synapses are depressed during learning and extinction and potentiated during re-learning. Our spiking model of the fly's brain provides a mechanistic explanation of how aversive odour associations can be acquired, extinguished and re-learned.

(5)

Persistence of sensory representation during neurodevelopment

Mizusaki BEP, Kourdougli N, Suresh A, Portera-Cailliau C, O'Donnell C

Many studies of psychiatric disorders pinpoint dysfunctional neuronal components, but it is still not well understood how these are integrated into information representation at the level of cortical circuits. In this work we focus on Fragile X Syndrome, a neurodevelopmental disorder related to autism spectrum disorder, using a mouse model that has the FMR1 gene knocked-out. One common symptom of Fragile X is sensory hypersensitivity. A recent hypothesis is that this could be caused by inconsistent stimuli representation, due to circuit hyperexcitability. We performed in vivo calcium imaging of layer 2/3 neurons of the somatosensory cortex in wild type (WT) and Fmr1 KO mice aged P14 and P19, using GCaMP6s. These ages represent the end of a critical period during development and peak FMRP expression in WT mice. We recorded both spontaneous activity and evoked activity elicited by whisker stimulation at two different frequencies. We re-identified individual neurons recorded at both time points (average of 20 cells), which enabled us to assess individual neuron contributions to the population representation. We characterized the sensory representation, and asked how much it changed between recording sessions, by comparing neuronal firing rates, pair-wise correlations and statistical models of the joint population activity. Finally, we tested whether the representation was more or less stable in FMR1 KO vs WT animals.

(6)

Efficient replay memory through sensory adaptation

Macmillan C, Chua R, Costa RP

The cortex adapts quickly to repetitive stimuli. Such rapid adaptation suggests that dissimilar (or novel) experiences are more likely to be retained in memory. In contrast, recent deep reinforcement learning models rely on storing every single input into a hippocampal-like episodic memory. Here, inspired by rapid forms of synaptic plasticity a key neural basis of sensory adaptation we propose a reinforcement learning algorithm -- filtered experience replay -- in which only dissimilar enough inputs are stored into the replay memory. We show that our method leads to a more efficient memory representation (reduced memory load), as similar inputs tend to be discarded. In addition, a model in which experience discards vary as the agent gradually learns to explore its environment achieve similar performance to standard replay memory methods.
This gradual change in adaptation is akin to the experimentally observed modifications of short-term plasticity over development and learning, and suggests an important role for this phenomenon in systems-level learning. Finally, our results suggest that filtered experience replay learns a better predictive model of the environment by relying less on consecutive inputs. Overall, our work shows how systems models of memory and learning can shed light on the function of synaptic plasticity and sensory adaptation.

(7)

Degeneracy in neuromodulatory neural circuits: serotonin-dopamine interaction in conditioning tasks

Behera C, Joshi A, Wang D, Sharp T, Wong-Lin K

Degeneracy is the ability to perform same function or yield a similar output despite being structurally different. One of the key advantages of degeneracy in neural circuits is the stability and robust maintenance of network function or cognition in the face of changes in electrophysiological activity. It is well known that neural circuit configurations can be altered via neuromodulators, leading to cognitive, behavioural task switching with the same set of neurons. Still, it is not clear whether the source of neuromodulators have neural circuits that can be degenerated. In this study, we addressed this issue by focusing on modelling the interactions of two key neuromodulators serotonin and dopamine. Based on our previous mean-field modelling approach, we developed a neural circuit model consisting of the ventral tegmental area (VTA) and dorsal raphe nucleus (DRN) in rewarding-/aversive-based conditioning tasks. We restricted the activities of the model to some prescribed activity profiles based on the observations in recent optogenetic and electrophysiological studies during conditioning tasks. Then, the structure of the DRN-VTA network is systematically altered. Each version of the DRN-VTA model always consists of the same afferent input structure and neuronal populations. We prescribed a maximal deviation range on the simulated activity profile and the desired profile at every time bins. Only network architectures that lie within this prescribed tolerance range are considered. We identified twelve different network architectures that can produce the same activity profile. Dynamical systems theory is applied to evaluate the stability of the DRN-VTA networks. Then, we simulated D2 receptor agonist effects in our models with increasing dosage level. We distinguished the sub-groups of network architectures based on their differential responses to the agonist. Overall, this work suggests the plausibility of degeneracy in the DRN-VTA neuromodulatory circuit and has important implications on the stable and robust maintenance of neuromodulation.

(8)

Towards understanding one-shot place learning in spatial navigation: A reinforcement learning approach

Tessereau C, O'Dea R, Bast T, van Rossum M, Coombes S

Animals, including humans, can perform very fast, 1-trial, learning of new places [1,2,6]. The underlying mechanisms are the focus of ongoing research in experimental and computational neuroscience. We will present a new approach to modelling spatial learning and navigation of an agent in a circular open field maze with distal cues. It comprises a combination of Temporal Difference (TD) learning, adapted from [5], and the Successor Representation (SR) [3,4]. TD models estimate a value function and a policy (best action to perform) in each state (place in the maze). In contrast, the SR provides a topological encoding of spatial information that allows to perform rapid path planning to novel goal locations in a familiar environment. Model performance is compared to one-trial place learning performance by rats and humans in the watermaze or watermaze analogues, respectively, when new goal locations need to be learned repeatedly [1,2,6]. We confirm and extend previous findings [5] that TD models can learn goal locations, although with only a gradual improvement in performance that does not capture the flexibility required to perform one-shot place learning, shown by rats and humans [1,2,6]. The SR is able to simulate direct navigation to any goal locations in an open field environment similar to the watermaze, but its exact computation is very expensive [3].
We are currently investigating the combination of the two approaches, that will be employed to investigate one-shot place learning shown by both rats and humans [1,2,6].

[1] Bast, T., Wilson, I. A., Witter, M. P., and Morris, R. G. (2009). PLoSBiology.
[2] Buckley, M.G. and Bast, T., 2018. Hippocampus.
[3] Corneil, D. S. and Gerstner, W. (2015). Advancesin Neural Information Processing Systems.
[4] Dayan, P. (1993). NeuralComputation.
[5] Foster, D.J., Morris, R. G., and Dayan, P. (2000), Hippocampus.
[6] Steele, R.J. and Morris, R.G.M. (1999). Hippocampus.

(9)

Codependent plasticity: learning with interacting synapses

Agnes EJ, Vogels TP

Synaptic plasticity drives learning. Many forms of plasticity have been observed and tested in models, but we are far from understanding all the brain's learning capabilities. Recent experimental work has shed light on the fact that synapses of different neurons and types directly interact with each other during learning. E.g., inhibitory synapses change according to the excitatory-to-inhibitory input ratio, and remain static when excitatory synapses are silent. Plasticity at an excitatory connection, on the other hand, seems to depend on the recent history of its excitatory neighbours, while the direction of changes depends on the level of inhibitory input. Here we combine this experimental evidence in a model of codependent synaptic plasticity in a simple feedforward (single postsynaptic neuron) and a recurrent network of spiking neurons. We show that when synapses interact through codependent plasticity, excitatory currents are naturally normalised and precisely balanced by inhibitory inputs. Due to the dependency on input currents, our model does not require a fixed-point for neuronal activity (firing-rate homeostasis). In the feedforward case, receptive field plasticity is quick and long-lasting when inhibition tightly controls excitatory plasticity and both excitatory and inhibitory plasticity act in synergy, similar to what has been reported in auditory cortex. The combination of excitatory and inhibitory codependent plasticity in recurrent networks allows for learning of a stable state of high-conductance in which neurons receive strong excitatory and inhibitory currents. This creates a stable activity regime at baseline that can be transiently disrupted by specific inputs to generate rich and complex spatiotemporal dynamics that eventually returns to baseline, such as observed in motor cortex dynamics. Codependent synaptic plasticity may be a missing ingredient towards understanding how our brains quickly and robustly learn long-lasting memories.

(10)

Modelling the role of substance P in reinforcement-based sequence learning

Favila Vazquez N, Gurney K, Overton PG

Several lines of evidence suggest that the striatum is one of the fundamental structures involved in sequence learning. Substance P is a neuropeptide abundant in the striatum known to interact with dopamine which is, itself, instrumental in producing changes in cortico-striatal synapses as new motor patterns are learned. We sought to investigate substance P role in sequence learning in an operant conditioning experiment, followed by computational modelling to achieve a mechanistic understanding of substance P's role. Rats (n = 20) were trained to perform a two action sequence in a Skinner box with two levers, for at least 25 sessions and until they had reached stable behaviour for five days. Then, half of the rats were switched to learn a new sequence and half were kept doing the same sequence. In the first three days of this second phase, rats were injected (i.p.) with a substance P (neurokinin 1) antagonist (L-733,060: 2mg) or with vehicle. Rats injected with the substance P antagonist learned the new sequence faster and performed the previous sequence less than vehicle rats. The performance of the rats that kept doing the same sequence was not affected. Using a temporal difference reinforcement learning model with an actor-critic paradigm we were able to simulate learning of a two action sequence. Several biologically constrained hypotheses about the role of substance P were tested in the model, and we were able to replicate the effects of blocking substance P in both experiments by manipulation of the state value learning rate. It has been suggested that striosomes encode state values and, interestingly, it has been reported that substance P boosts dopamine release only in striosomes. Thus, it is possible that by blocking substance P we indirectly affected dopamine, and thus, the state value learning rate, as suggested by our modelling results.

(11)

Multiphysics of neural signals

Wang H

A Multiphysics neural model is proposed by accounting two reactive components in neural systems: the inductive coil structure of the myelin sheath and the piezoelectric effect of the plasma membrane. The model reveals the interaction of electrical, magnetic and mechanical mechanisms in the generation of the neural signal and its transit along axonal tracts. A novel neural circuit with an RLC configuration is presented and is shown to provide an excellent fit to neural stimulation experiments. With this idea of the presence of reactive components in neuronal models, a number of phenomena can be explained, such as the opposite spiraling orientations of adjacent myelin sheaths, the experimental observations from magnetic nerve stimulation, the mechanism of acoustic nerve stimulation and the mechanical wave accompanying the action potential.

(12)

Energy efficient synaptic plasticity and network learning

Li HL, van Rossum M

Many aspects of the brain's design can be understood as the result of evolutionary drive towards efficient use of metabolic energy. In addition to the energetic costs of neural computation and transmission, experimental evidence indicates that synaptic plasticity is metabolically demanding as well. Yet, how this might have shaped the rules underlying neural plasticity is not known. Here we show that standard learning procedures to train neural networks would require extremely large amounts of energy when naively implemented in biology. We propose that this can be avoided by precisely balancing labile forms of synaptic plasticity with more stable forms. This algorithm, termed synaptic caching, can boost energy efficiency considerably. Apart from relevance for energy efficient neuromorphic designs, our results yield a novel interpretation of the multiple forms of neural synaptic plasticity observed experimentally.

(13)

Modelling Spinal Circuits using MIIND

Osborne H, Deutz L, de Kamps M

MIIND is a neural simulation development environment which utilises a geometric population density technique to model a network of interacting neural populations (de Kamps, Lepperød & Lai, 2019). The density technique builds a geometric mesh or grid which describes the deterministic dynamics of a neuron model. Assuming input spikes are Poisson distributed, MIIND then simulates the probabilistic behaviour of a population of such neurons. By describing the evolution of a probability density function across state space, metrics such as the average firing rate can be derived which show good agreement with direct simulation techniques. As would be expected, reducing the number of cells in the mesh, improves simulation times but increases the derived error. However, even with very granular meshes, the general behaviour of the population is preserved, so MIIND can be used to quickly prototype complex networks of populations while significantly reducing the pain of parameter tweaking. MIIND has recently been used to simulate a spinal interneuron and motor neuron population network model. The model supports experimental evidence of the proprioceptive effect on muscle recruitment during an isometric leg task and implicates spinal circuitry as a site of so called muscle synergy encoding. MIIND has also been used to reduce a Hodgkin Huxley style bursting neuron model to a simple two dimensional system. Using this reduced model as a basis, an example of a complex central pattern generator circuit can be reproduced in MIIND, displaying experimentally confirmed behaviours in fictive cat locomotion. This presentation introduces the MIIND population density technique and its usage, and demonstrates examples of its utility for supporting experimental investigations.

(14)

Recurrent inhibition and inhibitory plasticity outperform alternative mechanisms for sparse coding in a model of the Drosophila mushroom body

Bennett J, Vogels T, Nowotny T

Sparse coding is commonplace in sensory systems, and benefits reward learning by reducing the overlap between neural representations of different stimuli that cue different rewards. In Drosophila melanogaster, ~2000 Kenyon cells (KCs) in the mushroom body (MB) exhibit sparse responses to olfactory stimuli as a result of recurrent inhibition from a single neuron, the APL (anterior paired lateral neuron; Lin et al. 2014). KCs also drive MB output neurons, which bias behaviour, via synapses that undergo reward-modulated plasticity, such that flies can learn reward-seeking behaviours (Owald & Waddell, 2015). Here, we present a spiking neuron model of the MB to investigate how recurrent inhibition from the APL, in conjunction with inhibitory synaptic plasticity (ISP; Vogels et al. 2011), achieves a sparse code that is useful for learning.

We compare the APL+ISP model with two alternative models of sparse coding in the MB: 1) feedforward inhibition (FF) (Luo et al. 2010, Kee et al. 2015) with ISP, and 2) using an adaptive threshold (AT) that recruits a certain percentage of the most active KCs (Luo et al. 2010, Peng & Chittka 2017). We demonstrate that: 1) the AT approach is very sensitive to noise in the inputs as compared with the APL+ISP and FF+ISP models; 2) variability in responses between single KCs is greater in the AT model than in the APL+ISP and FF+ISP models; 3) variability in KC population responses between stimuli is greater in the FF+ISP model than in the APL+ISP and AT models.

We conclude that recurrent inhibition from the APL, in conjunction with inhibitory plasticity at APL-KC synapses, yields more consistent responses between KCs, between stimuli, and in the presence of noise. Consequently, reward-predicting responses in MB output neurons, and the resulting fly behaviours, are more accurate and reliable.

(15)

Influence of Prediction Error on Memory Encoding and Parahippocampal and Medial Prefrontal Connectivity

Khoo J, Derrfuss J, Danielmeier C

Learning and decision-making depend on the ability to encode memories. As a result of learning, we might expect a reward in a given context. Reinforcement learning theory predicts that the unexpected absence of such rewards elicts a negative prediction error (PE). Negative PEs should reduce our expectations for future reward. On the other hand, the PE might also increase the probability for recognizing the context that was not rewarding.
In the human brain, negative PEs have been typically associated with reductions in dopaminergic activity and activations of the posterior medial frontal cortex. Successful recognition memory, on the other hand, has typically been associated with parahippocampal activations. We aim to investigate whether these systems interact depending on the size of the PE.
We scanned participants using fMRI while they performed a probabilistic reversal learning task. In this task, the reward associated with a given choice changes repeatedly over the course of the experiment, and participants learn which category was currently rewarded. 24 hours later, a surprise recognition memory test involving items previously seen was administered.
We compared the fits of four variants of model-free, reinforcement learning models on the learning task. Among the models we tested, the best account for participant behaviour in response to reward stochasticity in this task was a model which included a variable learning rate that is modulated by a surprise factor.
We performed model-based fMRI analysis using parameters extracted from the computational models showing activation of the medial prefrontal cortex in response to PE. Also, parahippocampal activations were observed during encoding of subsequently remembered items. In addition, a psycho-physiological interaction (PPI) analysis will be conducted to investigate how functional connectivity between parahippocampal and medial prefrontal areas is modulated by prediction errors.

(16)

Nonlinear scaling of resource allocation in sensory bottlenecks

Edmondson LR, Jimenez Rodriguez A, Saal HP

In many sensory systems, information transmission is constrained by a neural bottleneck, where the number of output neurons is vastly smaller than the number of input neurons. For example, there are many more photoreceptors in the retina than there are retinal ganglion cells in the optic nerve. Efficient coding theory has been used to predict how the brain should allocate its limited resources in these scenarios by removing redundant information. Previous work has typically assumed that receptors are uniformly distributed across the sensory sheet, when in reality these vary in density, often by orders of magnitude: in vision, the density of cones varies dramatically between the central fovea and the periphery; in touch, mechanoreceptors are many times more densely packed in the fingertips than in the palm. How, then, should the brain efficiently allocate output neurons when the density of input neurons is nonuniform?

Here, we show analytically and in extensive computer simulations that resource allocation scales nonlinearly in efficient coding models that maximize information transfer, when inputs arise from separate regions with different receptor densities. Importantly, the proportion of neurons allocated to a given input region changes depending on the size of the bottleneck, and thus cannot be predicted from input density or region size alone: increasing the number of output neurons might increase or decrease a region's representation, depending on the bottleneck's initial size. Finally, over a wide range of bottleneck sizes neurons are allocated to a given region at a constant rate, slowly changing its proportional representation. Our results imply that both over- or under-representation of sensory input regions can arise in efficient coding models and might crucially depend on the neural resources available.

(17)

A topological insight on neuronal morphologies

Kanari L, Hess K, Markram H

The morphological diversity of neurons supports the complex information-processing capabilities of biological neuronal networks. A major challenge in neuroscience has been to reliably describe neuronal shapes with universal morphometrics that generalize across cell types and species. Inspired by algebraic topology, we have developed a topological descriptor of trees that couples the topology of their complex arborization with their geometric structure, retaining more information than traditional morphometrics.

The topological morphology descriptor (TMD) has proved to be very powerful in separating neurons into well-defined groups on morphological grounds. The TMD algorithm led to the discovery of two distinct morphological classes of pyramidal cells in the human cortex that also have distinct functional roles, suggesting the existence of a direct link between the anatomy and the function of neurons. The TMD-based classification also led to the objective and robust morphological clustering of rodent cortical neurons.

The TMD of neuronal morphologies is also essential for the computational generation (i.e., synthesis) of dendritic morphologies. Our results demonstrate that a topology-based synthesis algorithm can reproduce both morphological and electrical properties of reconstructed biological rodent cortical dendrites. Since the topology-based synthesis is generalizable to a wide variety of different dendritic shapes, it is suitable for the generation of unique neuronal morphologies to populate the digital reconstruction of large-scale, physiologically realistic networks.

(18)

Discovering The Building Blocks of Hearing: A Neuro-Inspired, Data-Driven Approach

Weerts L, Clopath C, Goodman D

Our understanding of hearing and speech recognition rests on controlled experiments requiring simple stimuli. However, these stimuli often lack the characteristics of complex sounds such as speech. We propose an approach that combines neural modelling with machine learning to determine auditory features that can be extracted by networks that are compatible with known auditory physiology. Our approach bridges the gap between detailed neuronal models that capture specific auditory responses, and research on the statistics of real-world speech data and speech recognition. Importantly, our model can capture a wide variety of well-studied features and allows us to unify concepts from different areas of hearing research.

We introduce a feature detection model with a modest number of parameters that is compatible with auditory physiology. We show that this model is capable of detecting a range of features such as amplitude modulations (AMs) and onsets. In order to objectively determine relevant feature detectors within our model parameter space, we use a simple classifier that approximates the information bottleneck. By analysing the performance in a classification task, our framework allows us to determine the best model parameters as well as their neurophysiological and psychoacoustic implications.

We show that in the context of a phoneme classification task, some model variants improve classification performance compared to the original signal. By analysing properties of high performing variants, we rediscover several proposed mechanisms for robust speech processing, such as formant detection and AM sensitivity. Moreover, our analysis hints at less-established ideas such as the importance of onset suppression. Finally, we show how our approach can easily be extended to more complex environments.

Our approach has various potential applications. Firstly, it could lead to new, testable experimental hypotheses for understanding hearing. Moreover, promising features could be directly applied as a new acoustic front-end for speech recognition systems.

(19)

Robust learning with subtractive recurrent neural networks

Llera Montero M, Malhotra G, Bowers J, Costa RP

The relative impact of subtractive and divisive gating is a long standing question in neuroscience. Such forms of gating are also key components of machine learning architectures such as long short-term memory networks (LSTMs) which are suited for contextual problems such as natural language processing. Recently, we proposed a mapping between LSTMs and the structure of cortical microcircuits (Costa et al. NeurIPS 2017) where the key difference is that networks operate with subtractive (subLSTM) rather than multiplicative gates as found in vanilla LSTMs.

Here we compared these two architectures over a set of sequential and contextual tasks. Our preliminary findings suggest that there are qualitative and quantitative differences between these models. In general, LSTMs seem to achieve better absolute training/testing error than subLSTMs. However, subLSTMs shows faster learning in a wide range of conditions. In addition, in some cases the validation error of LSTMs was unstable, regularly producing rapid fluctuations in error during learning, which is less prevalent in subLSTMs. When analysing the learned representations, we found that LSTMs also behave differently to subLSTMs while processing sequential tasks. The activation pattern for subLSTMs is akin to the typical transient sensory responses observed across the cortex while LSTMs exhibit a slow ramping-up activity. Moreover, lower dimensional analysis using tSNE and PCA suggests that the activation patterns for both models are equally well clustered according to stimulus category. Overall our findings suggest that subtractive gating provides more robustness during learning across a range of tasks and better captures neuronal dynamics observed in the cortex.

(20)

Fiber Tract Stimulation using a Phenomenological Application of tFUS in NEURON Leads to Synaptic Potentiation

Alkhawashki M

Aims: To establish that the stimulation of a fibre tract would lead to long-term potentiation (LTP) at the synapses regardless of the occurrence of postsynaptic spiking and, therefore, that it provides a viable target for neuromodulation using focused ultrasound (fUS).
Motivation: Transcranial Focused ultrasound (tFUS) is a promising method for the non-invasive neuromodulation of fibre tracts due to its ability to safely reach deep brain structures with high spatial resolution. Currently, multiple experimental studies using tFUS, but no clinical trials, have been conducted on humans, none of which explicitly targeted fibre tracts.
Methods: I created a two-cell network of biological neurons with myelinated axons in silico using the NEURON Simulation Environment. I stimulated them at several points along the axon using protocols based on those used in experimental studies with pulsed ultrasound (US) and pulsed focused ultrasound (fUS).
Results: Repeated stimulation of a myelinated axon exhibited synaptic potentiation accompanied by either repeated postsynaptic spiking or excitatory postsynaptic potential (EPSP).
Conclusion: Targeting fibre tracts for neuromodulation with tFUS is feasible. Experimental and clinical studies on animals and humans are needed to study the effects of fibre tract stimulation with tFUS in vivo on intact brains and to establish approaches to treatment.

(21)

Effective connectivity extracts clinically relevant prognostic information from resting state activity in stroke

Adhikari M, Griffis J, Siegel J, Deco G, Insabato A, Gilson M, Corbetta M

The objective of this study was to predict cognitive performance on behavioral tests from the resting-state (RS) fMRI data in a cohort of 30 healthy individuals and 100 first-time stroke patients. We compared the predictive accuracy of pairwise functional connectivity (FC) and whole-brain linear effective connectivity (EC) derived from RS-fMRI data. Behavioral tests were performed in language, motor, memory and attention domains at acute and two chronic stages (1-2 weeks, 3 months and 1 year post stroke onset). Subsequently, we trained a classifier to distinguish healthy from individuals with stroke using FC or EC links as predictors. We also identified biomarker sub-networks that support this classification and investigated their longitudinal evolution. We observed that accuracies of both FC and EC in distinguishing healthy individuals from patients were significantly higher than the chance level and were not significantly different from each other at all three stages. FC biomarkers mostly included links within resting state networks (RSNs) in line with previous research. EC biomarkers, on the other hand, included across RSNs links. We then considered prognosis of cognitive deficits based on FC or EC as a task involving classification of patients with either zero, single or multiple deficits in performance score of seven factors derived from within domain factor analysis. Here, EC turned out to be a better prognostic feature than FC. Further, classification of patients according to number of deficits was more accurate when we considered EC features at the acute stage in contrast to EC features at concurrent stages. Our results demonstrate that the second-order statistics of resting state activity in stroke patients at the acute stage has pertinent information for clinical prognosis that can be extracted by a whole-brain model derived effective connectivity.

(22)

Neural Topic Modelling

Hathway P, Goodman D

New data analysis methods are needed to process the ever-increasing number of neurons recorded by multi-electrode arrays. We propose Neural Topic modelling (NTM) - an unsupervised, scalable and interpretable tool which can be applied across different spatial and temporal scales.
NTM is based on Latent Dirichlet Allocation, a method routinely used in text mining to find latent topics in texts. The spike trains are converted into the presence or absence of discrete events (e.g. neuron 1 has a higher firing rate than usual), which we call "neural words". The recording is split into time windows that reflect stimulus presentation ("neural documents"). NTM results in a number of topics - probability distributions over words - which best explain the given occurrences of neural words in the neural documents.
We applied NTM to an electrophysiological dataset of mouse visual cortex neurons while sparse spatial noise was presented (recorded from a Neuropixel electrode). A majority of the topics identified by NTM either show a clear preference for distinct receptive fields (a small region of the visual field to which the words in the topic responded preferentially or non-preferentially) or have a similar brightness-sensitivity or reflect proximity on the recording electrode. Receptive fields were defined as weighted mean probabilities of the appearance of topic words given the stimulus location.
NTM receives no knowledge about the cortex topography nor about the spatial structure of the stimuli, but is nonetheless able to recover these relationships. Choosing neural activity patterns as neural words that are relevant to the brain makes the topics interpretable by both the brain and researchers. The combination of scalability (due to conversion of spikes into events), applicability across temporal and spatial scales and the biological interpretability of NTM sets this approach apart from other machine learning approaches to neural data analysis.

(23)

Modelling tactile responses from the sole of the foot

Kazu Siqueira R, Katic N, Bent L, Strzalkowski N, Raspopovic S, Saal HP

Cutaneous feedback from the foot sole is crucial for gait and balance control. Mechanoreceptive responses arise from a complex interplay of the mechanics of the skin tissue of the foot and the intrinsic firing properties of the receptors themselves. Predicting neural firing patterns that emerge from stimulation of the foot sole would allow the study of population-level responses and help to deliver improved biomimetic feedback in neuroprosthetic applications.
Here, we present a computational model that simulates neural spiking responses from the combined population of mechanoreceptive afferents innervating the sole of the foot. Our work takes into account the unique mechanics of the skin tissue in the foot (including depth and hardness) through a

simple mechanical model of tissue elasticity. Individual mechanoreceptor responses are then simulated using integrate-and-fire models that are fit to different mechanoreceptor types. The model includes all four relevant afferent classes (FAI, FAII, SAI and SAII) that signal skin indentation, slip, vibration, and stretch. Single-unit recordings of afferent firing characteristics were used for validation and fitting. The work builds on previous research in our group on mechanotransduction in the hand.

We used the model to estimate receptive field sizes and thresholds (absolute and tuning) across afferent types and show that they match their empirical counterparts closely. We also quantified how skin mechanics affect the tactile responses across different regions of the foot.

In future work, we plan to use the model to improve somatosensory feedback in prosthetic devices through neural interfaces with lower limb nerves.

(24)

The dynamics of explore-exploit decisions reveals a signal-to-noise mechanism for random exploration

Feng S, Wilson R, Wang S

When choosing a degree to pursue in college, should you exploit Psychology, which you know and love, or explore Machine Learning, which you know less about? Exploiting Psych is the way to a good grade, but exploring Machine Learning could be the path to a better life. As with choosing a field of study, making optimal explore-exploit decisions is hard --- a difficulty that arises because of the fundamental computational properties of explore-exploit problems: to make an optimal decision we need to consider all possible futures, but simulating all possible futures is beyond what any brain (or computer) can do.

Work in Machine Learning has shown that one way to overcome this computational complexity and make good explore-exploit choices in practice, is to choose randomly some of the time. Moreover, recent work in Psychology suggests that people actually use such “random exploration,” becoming more random in their decisions when they have more opportunity to explore. Here we investigate how this adaptive behavioral variability for random exploration could be controlled by the brain.

In particular, we model the explore-exploit choice using a drift diffusion process in which decisions are made by accumulating a noisy value signal towards one of two bounds: one for explore and one for exploit. In this model, behavioral variability is controlled by two parameters: the signal-to-noise ratio of the accumulation process (“drift”) and the separation between the decision bounds (“threshold”). By fitting people’s choices and reaction times in an explore-exploit task we find that while, statistically, both drift and threshold change as people have more opportunity to explore, numerically, it is the change in drift that has the main effect on random exploration. This suggests that random exploration is primarily driven by changes in the signal-to-noise ratio with which value information is represented in the brain.

(25)

Integrating classifiers and electrophysiology to better understand hearing loss

Smith S, Wallace M, Berger J, Sumner C

A moderate hearing loss can pose major challenges for speech identification, principally in noisy environments. It is not yet clear how the neural representation of speech changes following hearing loss, and whether it is in fact worse. Here, in order to quantify this, we present a framework that integrates a Bayesian classifier with neural data.

Neural responses to the presentation of vowel-consonants (VCs) overlapping one another, with onset asynchronies between 0 ms and ± 262.5 ms, were recorded from the midbrain of anaesthetised guinea pigs. Animals either had normal hearing (NH) or a moderate high frequency hearing loss imposed (HL). A naïve Bayesian-inspired classifier was implemented to predict auditory perception. The classifier was trained and tested on neural data recorded from either the NH or HL animals. The stimuli used to train the model were VCs in quiet, but it was also tested with prior knowledge of overlapping VCs.

The classifier was set to predict VCs in quiet with a 95% accuracy. When presented with the overlapping VC stimuli the classifier's performance reduced to as low as 53%. When this classifier was instead fed the HL data, identification of overlapping VCs reduced further (up to 20% worse than NH). Despite predicted errors mainly resulting from consonant confusions it was vowel confusions that became more prominent following hearing loss. Classifier performance improved when tested with prior knowledge of overlapping VCs, particularly for the HL data.

Overall, this work offers evidence for a degraded representation of speech in complex acoustic backgrounds at the midbrain level, following hearing loss. The principle of applying a machine classifier to the neural coding of speech appears to be a promising method for understanding the real-world problems associated with hearing loss.

(26)

A Bayesian evidence accumulation approach to inferring learning strategies

Maggi S, Hock R, Moran P, Bast T, Humphries M

Learning by trial-and-error is a fundamental method of problem-solving, to learn the action that leads to a reward. Typically, researchers focus on when participants (whether humans or rodents) acquire a defined action-outcome association. But there are many strategies participants could use to learn a new association; and often multiple strategies for achieving the same outcome, in which case it is crucial to determine what they have learnt compared to what they were expected to learn. To address these issues, we built a probabilistic approach to estimate trial-by-trial use of strategies by a participant, based on Bayesian accumulation of evidence – the participant's choices - over time. This approach allows us to estimate the probability of occurrence of each strategy (implicit or explicit) and its uncertainty on a trial-by-trial basis. We can then infer the most likely strategy being used at each point in time during training and infer when the learning of an association occurs. Testing this model on empirical data of rats learning new task rules revealed that, before learning, animals alternate between a few different strategies, showing a higher tendency to explore. By contrast, when the learning starts to emerge the animals tend to exploit more. In addition, we have tested a range of implicit and explicit strategies in order to explore the rat's perception of the

environment structure and find cases where the rat appears to have learnt a task rule but is implementing an alternative strategy. Our evidence accumulation approach infers behavioural strategies during learning and could thus form the basis for further investigation of the neural mechanisms of the identified strategies.

(27)

The effect of the recovery variable parameters on oscillating Izhikevich networks

Berg WP, Onken A

Recent studies have corroborated that the brain organises dynamically into functional groups, and that neural oscillation is a central component in this process. How this is mediated, however, remains unclear. One approach to illuminating this emergence is by computational modelling of biologically realistic, yet computationally tractable network models. In a recent paper, Oliveira et al. (2019: Neurocomputing, 337) analyse Izhikevich spiking neural networks, containing a mixture of excitatory and inhibitory neurons. They find that the resulting network oscillation rate can be predicted by a heuristic function, which may be utilised in constructing models. As stated originally by Izhikevich (2003: IEEE Transactions on Neural Networks, 14(6)), when the value determining the sensitivity to subthreshold oscillations increases towards a specific value, the neuron reaches its bifurcation point. Increasing the value thereafter entrains regular firing without any external stimulus, due to subthreshold oscillations. In our current work, we demonstrate that the firing rates in similar network models rely heavily on the subthreshold oscillations of the single neurons. We demonstrate that there is a deterministic and approximately linear relationship between the regular firing rate and the aforementioned value. Interestingly, we observe that this relationship also holds for network models consisting of excitatory and inhibitory neurons. Further, we show that there is an exponential relationship between the sensitivity variable and its time scale, given constant firing rates. Our results clarify the effect of specific parameters on oscillatory behaviour in the Izhikevich model. These findings may be used in order to more accurately characterise oscillating network models firing at specific rates.

(28)

Modelling the optimal integration of navigational strategies in the insect brain

Sun X, Mangan M, Yue S

Insect are expert navigators capable of searching out sparse food resources over large ranges in complex habitats before relocating their often hidden nesting sites. These feats are all the more impressive given the limited sensing and processing available to individual animals. Recently, significant advances have been made in identifying the brain areas driving specific navigational behaviours, and their functioning, but an overarching computational model remains elusive. In this study, we present the first biologically constrained, computational model that integrates visual homing, visual compass and path integration behaviours. Specifically, we demonstrate the challenges faced when attempting to replicate visual navigation behaviours (visual compass and visual homing) using the known mushroom body anatomy (MB) and instead propose that the central complex (CX) neuropil may instead compute the visual compass. We propose that the role of the mushroom body (MB) is to modulate the weighting of the path integration and visual guidance systems depending on the current context (e.g. in a familiar or unfamiliar visual surrounding). Finally, we demonstrate that optimal integration of directional cues can be achieved using a biologically realistic ring attractor network.

(29)

An Attentional Inhibitory Feedback Network for Multi-label Classification

Chu Y, Luk W, Goodman D

It's not difficult for people to distinguish the sound of a piano and a bass in a jazz ensemble even if these combinations have never been experienced before. However, these multi-label recognition tasks remain challenging for current machine learning and computational neural models. The first challenge is to learn to generalize along with the combinatorial explosion of novel combinations, in contrast to brute-force memorization. The second challenge is to infer the multiple latent causes from mixed signals.

Toward this direction, we present a new attentional inhibitory feedback model here. It outperforms feedforward-only baseline network with fewer neurons, in an overlapping-handwritten-digits recognition task that is designed to be easy for humans but challenging for machine learning models. Our simulation results also provide new understanding of feedback guided synaptic plasticity and complementary learning systems theory.

The proposed model has a feature encoder built on a multi-layer fully connected neural network. Each encoder neuron receives an inhibitory feedback connection from a corresponding attentional neural network. During recognition, an image is first fed through the encoder, yielding a first guess. Then, based on the most confidently recognized digit, the attention module feeds back a multiplicative inhibitory signal to each encoder neuron. In the following time step, the image is processed again, but by the modulated encoder, resulting in a second recognition result. This feedback loop can carry on several times.

In our model, attention modulates the effective plasticity of different synapses based on the predicted contributions. While the attention networks learn to select more distinctive features, the encoder learns better with synapse-specific guidance from attention. We also found intriguing dynamics during the co-learning process among attention and encoder networks, suggesting further links to neural development phenomena and memory consolidation in the brain.

(30)

Phase-Amplitude Coupled DCM and Time-Frequency Analysis of EEG Data from AD Patients

Tyrer A, Gilbert J, Fagerholm E, Adams S, Bankole A, Gilchrist I, Moran R

Alzheimer's disease (AD) is the most common cause of dementia in the aging population and covers approximately two-thirds of all dementia cases. AD patients initially display a deficiency in explicit (recognition) memory, while implicit (priming) memory is preserved during the early stages of the disease and begins to decline in mid to late stages. Medial temporal lobe atrophy is a strong indicator of developing AD in non-demented individuals and has been implicated in explicit memory processing. Oscillatory power at low frequencies, including theta (4 – 8 Hz), has been found to significantly increase in scalp recordings of temporal and frontal regions during successful episodic memory recall.

Dynamic Causal Modelling (DCM) is an established computational method which aims to make inferences about changes in effective connectivity, i.e. the strength of connections between sources of activity and the influences of one neuronal source over another in the context of an experimental task, by utilising neuroimaging data such as EEG data.

We interrogated EEG time-frequency data in bilateral temporo-frontal cortical networks in AD patients (n=21), non-AD dementia patients (n=16) and age-matched healthy controls (n=21), during two behavioural tasks. These tasks investigated visual priming and recognition, in which novelty is manipulated, to examine implicit and explicit memory. We found increased gamma activity in AD patients compared to controls, and in controls' novel trials compared to repeated trials indicating a potential prediction error. We then applied two dynamic causal models (DCMs) to test the slow amplitude envelope of signal transmission through the network, and for phase-amplitude coupling to each individuals' recordings. Parameters were analysed using Parametric Empirical Bayes (PEB) which aims to test for significant covariation in parameter estimates for group-level covariates. PEB analysis of slow envelope DCMs revealed a loss of temporo-frontal connectivity specifically in AD patients during the recognition phase, suggesting significant dropout in memory circuit connectivity in AD. Phase-amp DCMs revealed strong correlations between oscillator amplitudes in bilateral top-down and bottom-up connections in AD patients, which did not appear in controls or non-AD patients.

Our results show increased gamma activity in temporal sources in novel compared to repeated trials, and in AD patients compared to controls which may indicate a prediction error in recognition memory tasks. Also, controls show increased temporal theta activity compared to AD patients, which has been linked to recall and recognition memory. PEB analysis found that left hemisphere connectivity between frontal and temporal regions are compromised in AD, suggesting explicit memory circuit dropout.