Release of spiking models of visuo-motor coordination incorporating multiple brain regions, implemented in software, neuromorphic hardware, and neurorobotics (D3.9 – SGA3)

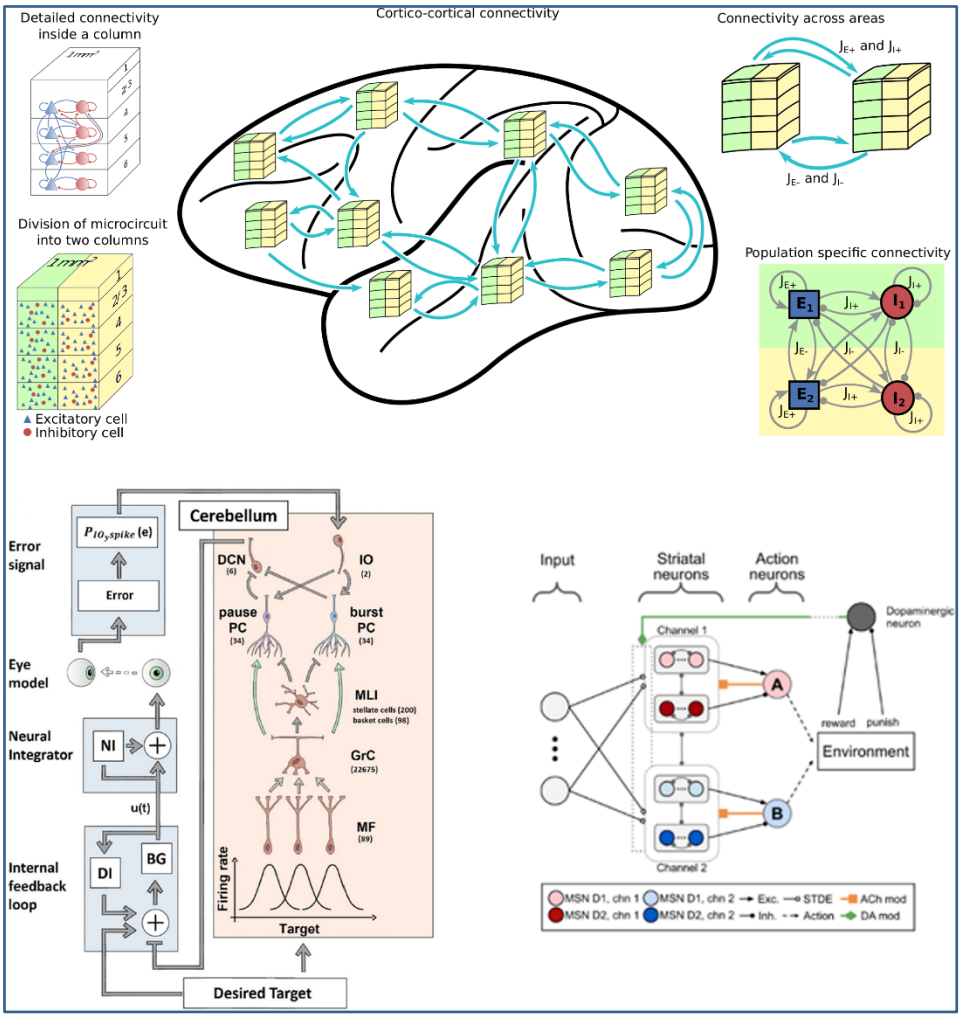

Figure 1: Selection of models of cortex, cerebellum, and basal ganglia.

Top: clustered multi-area model of macaque cerebral cortex, described in Sec. 4.3. Bottom left: Cerebellar control system described in Sec. 2.3. Bottom right: Striatal action selection circuit described in Sec. 3.2.

Project Number:	945539	Project Title:	HBP SGA3

Document Title:	Release of spiking models of visuo-motor coordination incorporating multiple brain regions, implemented in software, neuromorphic hardware, and neurorobotics
Document Filename:	D3.9 (D28) SGA3 M42 SUBMITTED 230915.docx
Deliverable Number:	SGA3 D3.9 (D28)
Deliverable Type:	Other
Dissemination Level:	PU = Public
Planned Delivery Date:	SGA3 M42 / 30 SEP 2023
Actual Delivery Date:	SGA3 M42 / 15 SEP 2023
Author(s):	Sacha VAN ALBADA, JUELICH (P20) Claudia CASELLATO, UNIPV (P70) Egidio D'ANGELO, UNIPV (P70) Simona TRITTO, UNIPV (P70) Mihai PETROVICI, UBERN (P71) Agnes KORCSAK-GORZO, JUELICH (P20) Alvaro GONZÁLEZ REDONDO, UGR (P66) Eduardo ROS, UGR (P66) Mario SENDEN, UM (P117)
Compiled by:	Sacha VAN ALBADA, JUELICH (P20)
Contributor(s):	Jeanette HÄLLGREN-KOTALESKI, KTH (P39), contributed to Section 3 Sten GRILLNER, KI (P37), contributed to Section 3 Jari PRONOLD, JUELICH (P20), contributed to Section 4.3 Anno KURTH, JUELICH (P20), contributed to Section 4.4 Jasper ALBERS, JUELICH (P20), contributed to Section 4.4 Egidio FALOTICO, SSSA (P49), contributed to Section 2.3
WP QC Review:	David GOYER, JUELICH (P20)
WP Leader / Deputy Leader Sign Off:	Sacha VAN ALBADA, JUELICH (P20)
T7.4 QC Review:	N/A
Description in GA:	Models incorporating retina, thalamus, brainstem, basal ganglia, cerebellum, and cerebral cortex able to perform eye, hand, and arm movements, implemented in simulation software (NEURON, NEST, LFPy) and on SpiNNaker, and demonstrated in closed-loop scenarios with neurorobotics
Abstract:	The research described herein concerns the development of computational models that simulate the neural mechanisms involved in visuo-motor coordination. These models incorporate multiple brain regions and are implemented using simulation

	software. The simulations demonstrate their functionality and enable evaluating working hypotheses in closed-loop scenarios with neurorobotics.
Keywords:	Spiking neural network models, visual, motor, visuo-motor interactions, neuromorphic, neurorobotics, biophysics, bio-inspired architectures
Target Users/Readers:	computational neuroscience community, consortium members, neuromorphic engineers, funders, general public, HPC community, neuroimaging community, neuroinformaticians, neuroroboticists, neuroscientific community, neuroscientists, AI researchers, platform users, policymakers, researchers, scientific community, students

Table of Contents

Table of Figures

1. Introduction

The objective of the present work was to develop a collection of biologically based neural network models of various brain regions for two main purposes: 1) bring together biological realism with sensorimotor function to gain neuroscientific insights into sensorimotor, and in particular visuomotor coordination; 2) serve as testbeds for the HBP software and hardware (high-performance computing (HPC), neuromorphic systems) simulation tools and the Neurorobotics Platform (NRP). The results encompass spiking models of cerebellar, basal ganglia, and cortical circuits along with associated nuclei, implemented using EBRAINS tools including the Brain Scaffold Builder (BSB), Snudda, NEST, NEURON, Arbor, SpiNNaker, and the NRP. Through the development of the given models, substantial advances in these tools and platforms were achieved. For example, the large-scale multi-area models of cerebral cortex have served as performance benchmarks for NEST on HPC systems and have directly inspired improvements in the simulation kernel of NEST. In collaboration with UNIMAN, T3.2 partners have further worked on porting a multi-area cortical model to the SpiNNaker neuromorphic hardware as the largest model to be simulated on that platform to date. This work has helped expose and resolve bottlenecks of SpiNNaker relevant for simulating any model using many of its boards. The basal ganglia modelling has made extensive use of Snudda and thereby driven forward its development, while the cerebellar modelling has extensively used the BSB and substantially advanced its development. Virtual robotics applications implemented using the NRP show how knowledge of the brain can help create functional cognitive architectures that simultaneously enable neuroscientific insights. The models that have emerged from this work are made publicly available via the EBRAINS Knowledge Graph (KG)[1]. The work described herein should be of interest to experimental and computational neuroscientists, developers of simulation tools (software and hardware), neuroroboticists, AI researchers, consortium members, scientific funding agencies, policymakers, platform users, students, and the general public. Besides the abovementioned contributions to Computing, it can contribute to AI applications and Robotics by providing ideas on brain-inspired solutions to computational problems, particularly related to sensorimotor coordination. For instance, insights are provided into the role of the neuromodulators dopamine and acetylcholine in action selection and decision-making within the basal ganglia. The work may further inspire applications in Medicine through insights provided into disorders including Parkinson's disease and dystonia.

[1] https://search.kg.ebrains.eu/?category=Model

2. Cerebellar-based systems in sensorimotor tasks

2.1 Cerebellar plasticity and classical eyeblink conditioning

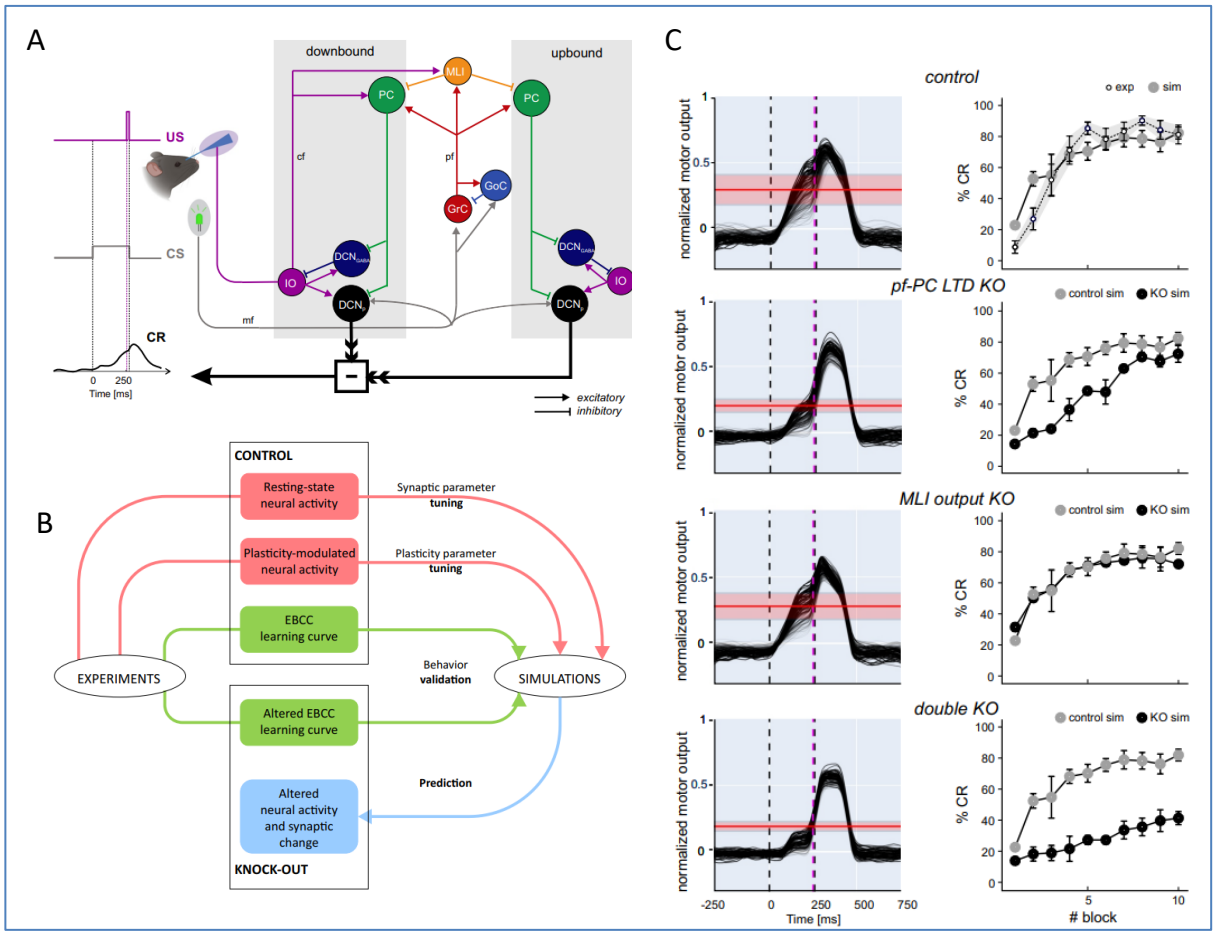

Figure 2: A) Model overview. B) Experiment to simulation. C) Conditional response learning.

According to Marr's motor learning theory, plasticity at the parallel fibre to Purkinje cells synapse (pf-PC) is the main substrate responsible for learning sensorimotor contingencies under climbing fibre control. However, the discovery of multiple forms of plasticity distributed over different cerebellar circuit synapses prompts to remap the cerebellar learning sites. UNIPV simulated classical eyeblink conditioning (CEBC) using an advanced spiking cerebellar model embedding upbound and downbound modules that are subject to multiple plasticity rules (Long-Term Depression and Potentiation – LTD and LTP) (Figure 2A). Simulations show that synaptic plasticity regulates the cascade of precise spiking patterns spreading throughout the cerebellar cortex and cerebellar nuclei. CEBC was supported by plasticity in both the pf-PC synapses and in the feedforward inhibitory loop passing through the molecular layer interneurons, but only the combined switch-off of both sites of plasticity compromised learning significantly. By differentially engaging climbing fibre information and related forms of synaptic plasticity, both modules contributed to generate a well-timed conditioned response, but it was the downbound module that played the major role in in this process. The outcomes of our simulations closely align with the behavioural and electrophysiological phenotypes of mutant mice (Knock-Out, KO) suffering from cell-specific mutations that affect processing of their PC or MLI synapses. Our data highlight that a synergy of bidirectional plasticity

rules distributed across the cerebellum facilitate finetuning of adaptive associative behaviours at a high spatiotemporal resolution.[2]

A similar system architecture was applied to investigate the mechanisms in dystonia. UNIPV modified structural or functional local neural features in the network reproducing alterations reported in dystonic mice. The work suggests that only certain types of alterations, including reduced olivocerebellar input and aberrant PC burst firing, are compatible with the eyeblink conditioning changes observed in dystonia, indicating that some cerebellar lesions can have a causative role in the pathogenesis of symptoms (Figure 3)[3].

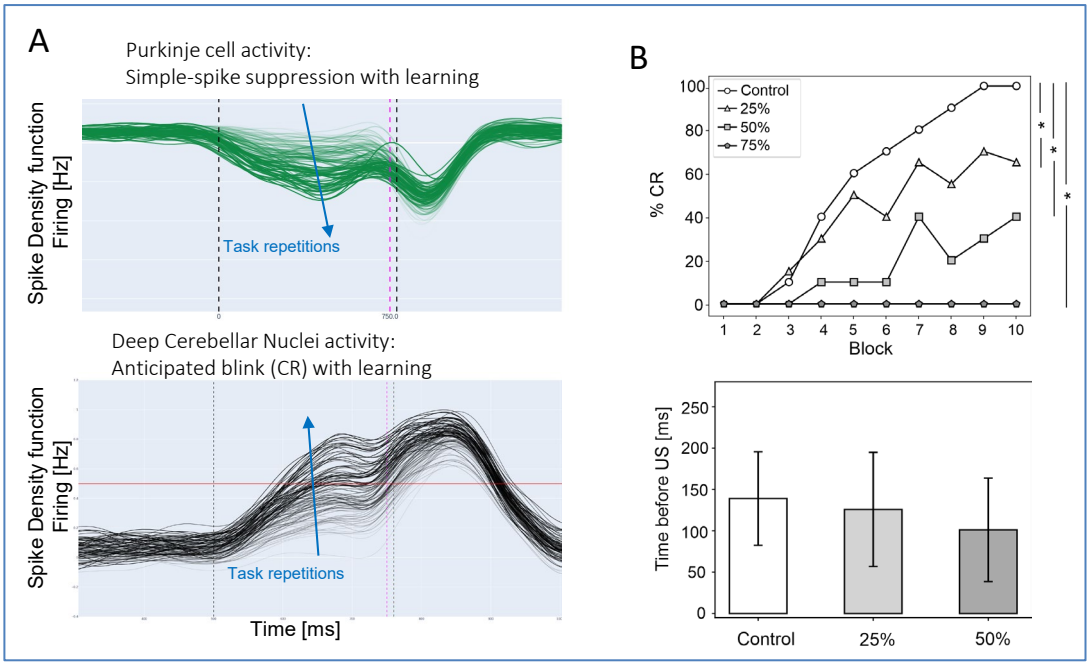

Figure 3: A) Purkinje cell and DCN activity. B) CR dependence on olivocerebellar input.

2.2 Cerebellar STDP in saccadic control

In collaboration with SSSA, UNIPV developed a system able to generate saccadic eye movements, which play a crucial role in visuo-motor control by allowing rapid foveation onto new targets. However, the neural processes governing saccades adaptation are not fully understood. Saccades, due to the short time of execution (20–100 ms) and the absence of sensory information for online feedback control, must be controlled in a ballistic manner. Incomplete measurements of the movement trajectory, such as the visual endpoint error, are supposedly used to form internal predictions about the movement kinematics resulting in predictive control. To characterize the synaptic and neural circuit mechanisms underlying predictive saccadic control, we have reconstructed the saccadic system in a digital controller embedding a spiking neural network of the cerebellum (in NEST) with spike timing-dependent plasticity (STDP) rules driving parallel fibre-Purkinje cell long-term potentiation and depression (LTP and LTD) (Figure 4A). This model implements a control policy based on a dual plasticity mechanism, resulting in the identification of the roles of LTP and LTD in regulating the overall quality of saccade kinematics: it turns out that LTD increases the accuracy by decreasing visual error and LTP increases the peak speed. The control policy also required cerebellar PCs to be divided into two subpopulations, characterized by burst or pause responses. To our knowledge, this is the first model that explains in mechanistic terms the

[2] Mesoscale simulations predict the role of synergistic cerebellar plasticity during classical eyeblink conditioning. A Geminiani, C Casellato, H-J Boele, A Pedrocchi, C I De Zeeuw, E D'Angelo. BioRxiv; doi: 10.1101/2023.06.20.545667; P4054

[3] Geminiani A, Mockevičius A, D'Angelo E, Casellato C. Cerebellum Involvement in Dystonia During Associative Motor Learning: Insights from a Data-Driven Spiking Network Model. Front Syst Neurosci. 2022 doi: 10.3389/fnsys.2022.919761; P3309

visual error and peak speed regulation of ballistic eye movements in forward mode exploiting spike-timing to regulate firing in different populations of the neuronal network. The cerebellar spiking neural network learning by means of dual plasticity rule led to decrease in end foveal error and increase in peak speed across repeated movements (Figure 4B).[4]

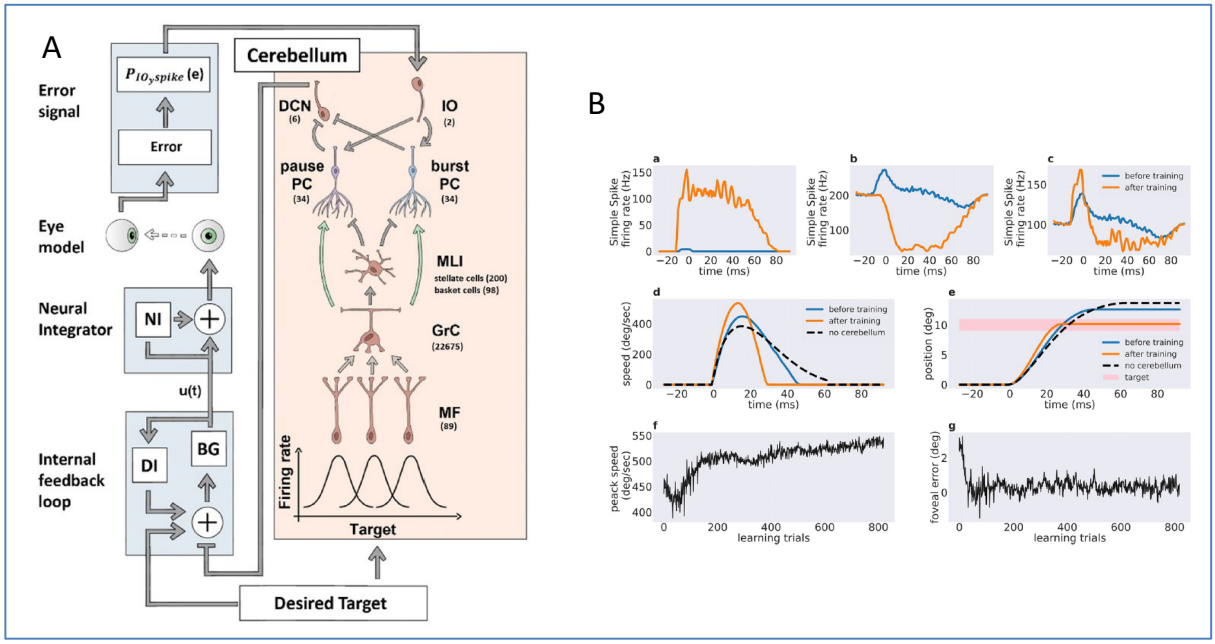

Figure 4: A) Schematic of the control loop. B) Eye movement kinematics.

DI: Displacement integrator; BG: Burst generator.

2.3 Multiscale modelling of the cerebellar involvement in Parkinson's Disease

In collaboration with POLIMI, UNIPV then incorporated the spiking cerebellum into a multiscale computational model of the rodent brain, in which basal ganglia (BG) and cerebellum are the two spiking components while thalamocortical nodes are mass models. Simulations showed that a direct effect of dopamine depletion on the cerebellum must be taken into account to reproduce the alterations of Parkinson's disease (PD) neural activity, particularly the increased beta oscillations widely reported in PD patients. Moreover, dopamine depletion indirectly impacted spike-time-dependent plasticity at the parallel fibre-Purkinje cell synapses, degrading associative motor learning as observed in PD. Overall, these results suggest a relevant involvement of cerebellum in PD motor symptoms (Figure 5)[5].

[4] Fruzzetti L, Kalidindi HT, Antonietti A, Alessandro C, Geminiani A, Casellato C, Falotico E, D'Angelo E. Dual STDP processes at Purkinje cells contribute to distinct improvements in accuracy and speed of saccadic eye movements. PLoS Comput Biol. 2022 doi: 10.1371/journal.pcbi.1010564; P3878

[5] Gambosi B, Sheiban F, Biasizzo M, Antonietti A, D'Angelo E, Mazzoni A, Pedrocchi A. Dopamine-dependent cerebellar dysfunction enhances beta oscillations and disrupts motor learning in a multiarea model. bioRxiv doi: 10.1101/2023.07.18.549459; P4096

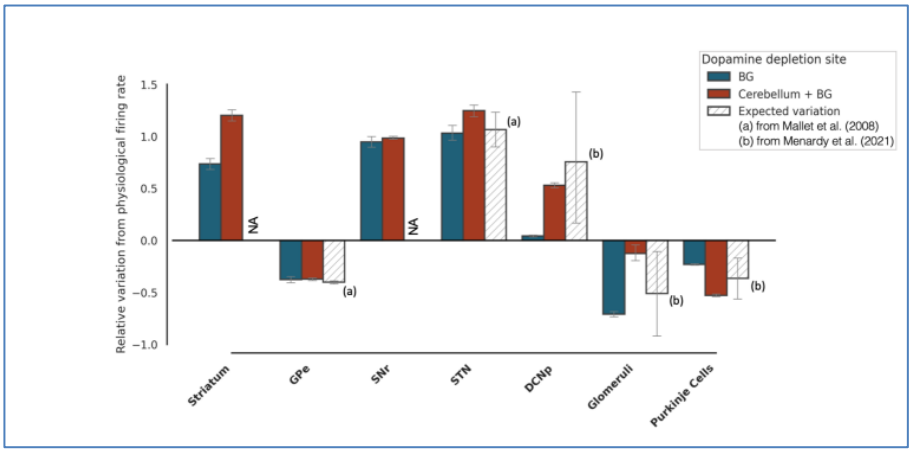

Figure 5: Firing rate change with dopamine depletion only in BG vs. in cerebellum + BG.

The values are computed considering the most severe pathological scenario (depletion level = 0.8). The reported values have been obtained by normalizing the difference between the firing rates in pathological and physiological conditions with respect to the physiological firing rate (0 = no difference, 1 corresponds to a pathological firing rate that is two times the physiological one).

2.4 Neurorobotic whisker system on the NRP

In collaboration with POLIMI, UNIPV further developed a peripheral whisker system (trigeminal ganglion, trigeminal nuclei, facial nuclei, and central pattern generator neuronal populations) connected to an adaptive cerebellar network controller (NEST), exploiting the Neurorobotic Platform (Figure 6). The whole system was able to drive active whisking with learning capability, matching neural correlates of behaviour experimentally recorded in mice[6].

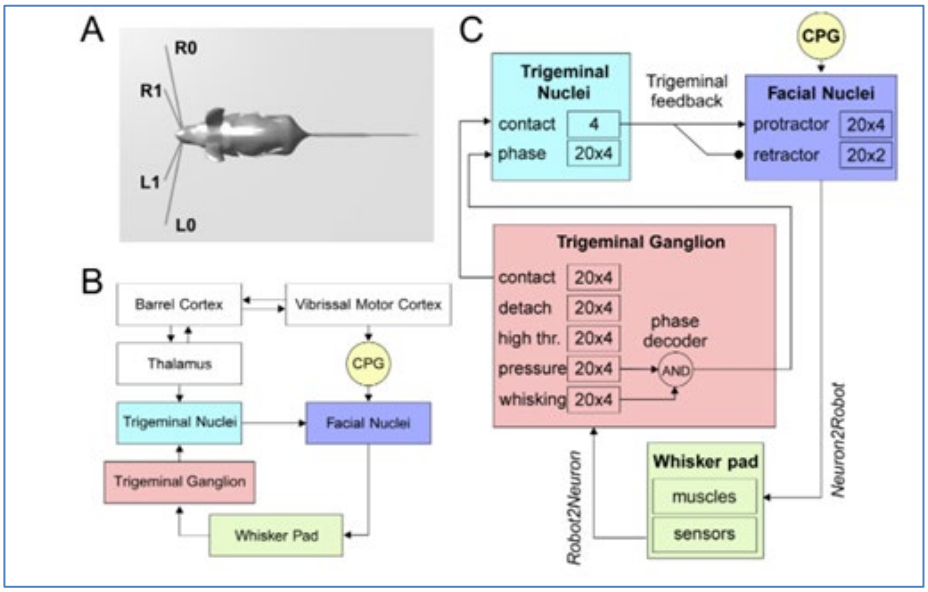

Figure 6: A) Virtual robotic mouse. (B) Rodent whisker system. (C) SNN implementation.

In panel C, numbers in each block represent the size of the neural populations included in that brain region. Arrows represent excitatory connections, circles inhibitory connections.

[6] Antonietti A, Geminiani A, Negri E, D'Angelo E, Casellato C, Pedrocchi A. Brain-Inspired Spiking Neural Network Controller for a Neurorobotic Whisker System. Front Neurorobot. 2022 doi: 10.3389/fnbot.2022.817948; P3741

2.5 Upper limb motor control: spino-cerebellar integration

Different parts of the central nervous system act in the motor control of our body, performing different but intertwined roles and implementing different control mechanisms sustained by specific neural structures. To further study how multiple brain regions coexist in performing motor control, UGR in collaboration with EPFL integrated a cerebellar model and a spinal cord model in a closed loop, operating a musculoskeletal upper limb model consisting of two joints (shoulder and elbow) actuated by eight muscles.

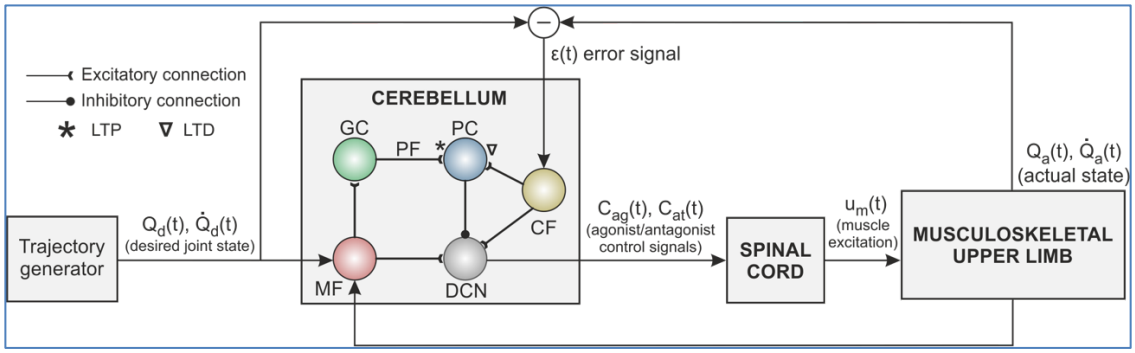

Figure 7: Spino-cerebellar integration in closed-loop control of an upper limb.

The cerebellar model was equipped with STDP at the parallel fibre – Purkinje cell synapses (PF-PC), allowing motor adaptation and learning. The spinal cord model implemented fast and direct control of the upper limb muscles by means of stretch reflex and reciprocal inhibition. The spino-cerebellar integration thus allowed for a hierarchical neural structure, in which the cerebellum provided adaptation to the desired motor task through a gradual learning process (Ito, 2000), and the spinal cord allowed fast muscle reactions and low control primitives (Pierrot-Deseilligny & Burke, 2012). The idea behind the study was to determine whether the fast action of the spinal cord facilitated or hindered the cerebellar motor adaptation.

The spinal cord directly controls and regulates muscle activation; thus, it significantly impacts the body plant dynamics which need to be acquired by higher brain regions (i.e., cerebellum in our case). In our setup, the integrated model faced several repetitions of a set of upper limb motor tasks (2 DOF flexion-extension movements involving the shoulder and elbow joints). Continued exposure to the desired motor tasks allowed the cerebellum to adapt the motor behaviour to minimise the mismatch between the desired and actual state of the upper limb (Medina, 2019). The cerebellar gradual motor adaptation coexisted with the spinal cord fast and direct regulation of muscle activity. Results showed that the spinal cord action facilitated and simplified motor learning in the cerebellum, i.e., the interplay of these two brain regions favours the control of the upper limb dynamics[7] (also cf. Figure 7).

2.6 Further work on cerebellar modelling

Further work by UNIPV and partners brought the Brain Scaffold Builder (BSB)[8] to EBRAINS, interfacing with the simulators NEURON, Arbor, and NEST[9]; cf. also T1.15. Using the BSB and atlas data, a pipeline was developed for detailed in silico cerebellar reconstruction, in collaboration with the

[7] Bruel, A., Abadía, I., Collin, T., Sakr, I., Lorach, H., Luque, N., Ros, E., & Ijspeert, A. (2023). The spinal cord facilitates cerebellar upper limb motor learning and control; inputs from neuromusculoskeletal simulation. bioRxiv, 2023-03, https://doi.org/10.1101/2023.03.08.531839, P4072

[8] https://www.ebrains.eu/tools/bsb

[9] De Schepper, et al. (2022) "Model simulations unveil the structure-function-dynamics relationship of the cerebellar cortical microcircuit", Commun Biol. 5, https://doi.org/10.1038/s42003-022-04213-y; P3729

Marie Skłodowska-Curie project CEN ("Cerebellum and emotional networks"). This pipeline was illustrated with a use case on the declive[10].

At the systems level, UNIPV implemented Bayesian integration in a spiking neural system for sensorimotor control[11] and investigated embodiment during sensorimotor tasks in a NEST-based cerebellar model integrated with further brain regions and a simulated body using PyBullet and MUSIC[12]. A simplified version of the cerebellar circuit simulating the classical delayed eyeblink conditioning protocol was implemented on a neuromorphic low-power microcontroller[13]. Finally, the cerebellar circuit model was ported to SpiNNaker as the first simulation of a large-scale, biophysically constrained cerebellum model performed on neuromorphic hardware[14].

3. Basal ganglia models

3.1 Introduction

The present section details collaborative work by UGR, KTH, and KI to bring functional aspects into spiking models of the basal ganglia and connected brain regions.

3.1.1 Roles of the basal ganglia

Choosing the right action among many available choices represents a primary but also challenging behaviour for animal species. Multiple stimuli spanning different sensory modalities continuously converge to the brain, and adequate responses (taking into account these inputs) need to be decided. The consequences (good or bad) of all these decisions need to be remembered, in order to make better decisions in the future and to be able to avoid fatal mistakes.

The basal ganglia (BG) play a key role in action-selection and reinforcement learning. They are a collection of highly interconnected nuclei located in the deepest part of the brain. They receive action proposals from the cortex and select the most appropriate actions based on past experiences.

Biological studies (Graybiel, 1998; Grillner et al., 2005; Hikosaka et al., 2000) and relevant computational models (Gurney et al., 2001) also have proposed the association between the BG and the action-selection and reinforcement learning. Specifically, the cortex and other brain structures send action proposals to the basal ganglia, which select the appropriate ones to be executed in the current context. Past experiences weigh the selection favouring the ones that resulted better in similar contexts.

3.1.2 Functioning of the basal ganglia

The basal ganglia network consists of several interconnected nuclei (Figure 8). Two of them are an input nucleus called the corpus striatum (STR), with its main cell type being the medium spiny neurons (MSNs), and the output nucleus called substantia nigra pars reticulata (SNr) which projects to the thalamus. A third nucleus, called the substantia nigra pars compacta (SNc), inputs reward information into the network.

[10] HBP Summit in Marseille 2023; E 6419
[11] Grillo M, Geminiani A, Alessandro C, D'Angelo E, Pedrocchi A, Casellato C. Bayesian Integration in a Spiking Neural System for Sensorimotor Control. Neural Comput. 2022. doi: 10.1162/neco_a_01525; P4060
[12] Gambosi et al., NEST Conference 2022; E6693
[13] Gandolfi et al., 2022 J. Neural Eng. 19 036022; P3310
[14] Petruţ A. Bogdan, Beatrice Marcinnò, Claudia Casellato, Stefano Casali, Andrew G.D. Rowley, Michael Hopkins, Francesco Leporati, Egidio D'Angelo, Oliver Rhodes, 2021. Towards a Bio-Inspired Real-Time Neuromorphic Cerebellum, Frontiers in Cellular Neuroscience, Vol. 15, doi: 10.3389/fncel.2021.622870, P2905

The connections within the basal ganglia network form three main pathways: the direct pathway, the indirect pathway, and the hyperdirect pathway. The direct and indirect pathways are traditionally considered to promote or inhibit behaviour based on selective disinhibition or inhibition of the motor thalamus.

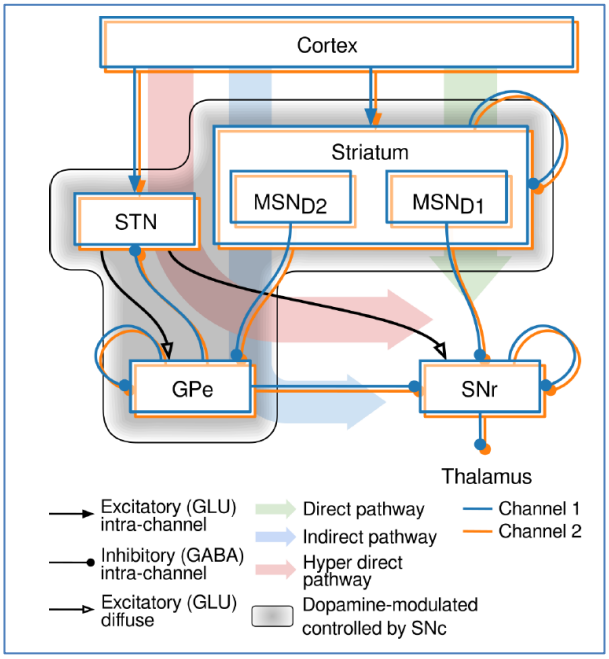

Figure 8: Basal ganglia structure showing direct, indirect, and hyperdirect pathways.

In addition to these broad pathways, important neuromodulators like dopamine (DA) and acetylcholine (ACh) influence different parts of the circuitry. For example, there are dopaminergic projections from the SNc to the MSN and other nuclei with modulatory effects (see shaded box in Fig 7). Also, thalamic projections innervate cholinergic interneurons in the STR, influencing the amount of ACh in this nucleus (Xiao and Roberts, 2021). Phasic DA is generally thought to carry reinforcement-related signals to the STR (Hart et al., 2014), although its full role is still under debate (Berke, 2018). Meanwhile, pauses in ACh seem to define the time window for phasic dopamine to induce plasticity (Reynolds et al., 2022). However, it remains a topic of discussion how these mechanisms combine to enable the striatum to solve action-selection problems.

The BG is well situated for reinforcement learning as it serves as a neural interface between reinforcement signals, primarily through DA, and action representations via cortical input pathways (Mogenson et al., 1980). The STR and its main population of MSNs constitute this interface, where cortical inputs establish plastic synapses modulated by DA. The adjustment of the weights of these cortico-striatal synapses in response to reward signals influences which actions are prioritized in the future (Gurney et al., 2015; Reynolds and Wickens, 2002).

3.2 Computational models of the basal ganglia

3.2.1 Channel structure

Computational models of the basal ganglia have been developed, incorporating the concept of cognitive streams or channels representing potential actions (Gurney et al., 2001; Suryanarayana et al., 2019). The BG are thought to act as an action selection machinery by inhibiting every nonselected action in the thalamus with the SNr, based on their corresponding activity level or salience (Redgrave et al., 1999).

According to recent research, this segregation through the entire cortical-BG-thalamic loop shows a very high specificity to almost neuron-to-neuron level (Foster et al., 2021; Hunnicutt et al., 2016), which could mean that it seems feasible to impact behaviour at different levels of detail.

3.2.2 Action selection

Previous models of the Basal Ganglia (BG) used population models, in which each node symbolises a group of neurons rather than individual ones. These models represent average neuronal activity rather than individual spikes. However, more detailed Spiking Neural Networks (SNNs) represent individual neurons, providing a more precise representation of brain computation at the neuron level.

This specificity is crucial to understand the interaction between various channels during action-selection processes. Burke et al.'s model proposes an explanation for the co-activation of D1 and D2 Medium Spiny Neurons (MSNs), which control behaviourally contrasting pathways. This is achieved by introducing asymmetric lateral connectivity in the striatum and lateral inhibition between neurons of the same receptor type, with stronger D2 to D1 neuron inhibition (Figure 9). This model helps to explain phase-dependent activation synchronization between D1 and D2 MSNs from different channels, which is essential for designing models that can perform action selection.

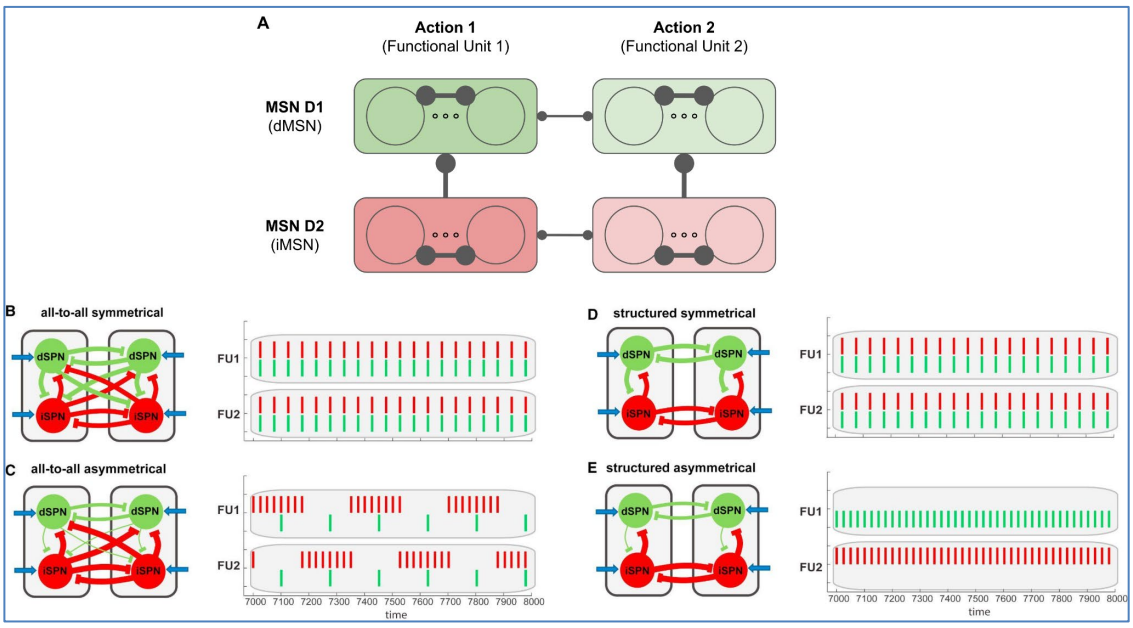

Figure 9: Connectivity model used for channels in the STR.

A) Each column indicates an action channel made up of D1 and D2 MSN subgroups. Inhibition is seen within each subgroup and from D2 to D1 subgroups in the same channel. Lateral inhibition only occurs between different channels of the same type. Except for D2 to D1 inhibition within a channel, all inhibitory synapses are weak. B-E) These are adapted figures from Burke et al. (2017) displaying activity resulting from various connectivity patterns. The connectivity model used in this study matches the structural asymmetrical pattern (E), where lateral inhibition from the active channels limits the activity of other silent channels.

3.2.3 Learning

In SNNs models, a widely used learning rule is spike-timing-dependent plasticity (STDP), a synaptic model with weight adaptation demonstrated in biological systems (Levy and Steward, 1983) and more particularly in the BG (Fino and Venance, 2010). This is a learning rule where the synaptic weight changes based on the relative timing of pre- and post-synaptic spikes. No weight change occurs if the spikes are far apart in time, but as the spikes get closer in time, the weight change becomes more significant. The direction of the weight change typically depends on the order of the spikes: the weight increases if the presynaptic spike comes before the postsynaptic spike and decreases if it comes afterwards. This rule allows neurons to detect and recognize repetitive patterns, useful in unsupervised learning tasks.

In reinforcement learning tasks it is common to learn not just any repetitive patterns, but the ones associated with reward. Reward-modulated STDP (R-STDP) is a modified version of STDP, where the amount of dopamine (DA) influences the weight change, biasing the learning process towards

patterns that correlate with reward signals. It uses eligibility traces, temporary storage of potential synaptic changes, to account for delays between stimuli and rewards.

Spike-Timing-Dependent Eligibility (STDE; Figure 10) is a more flexible model that captures features found in the biological medium spiny neurons (MSN) of the basal ganglia. Different learning kernels can be used depending on the amount and type (reward or punishment) of reinforcement received. Despite some limitations, this model has proven successful in action selection tasks, driven by the timing of input and reward signals.

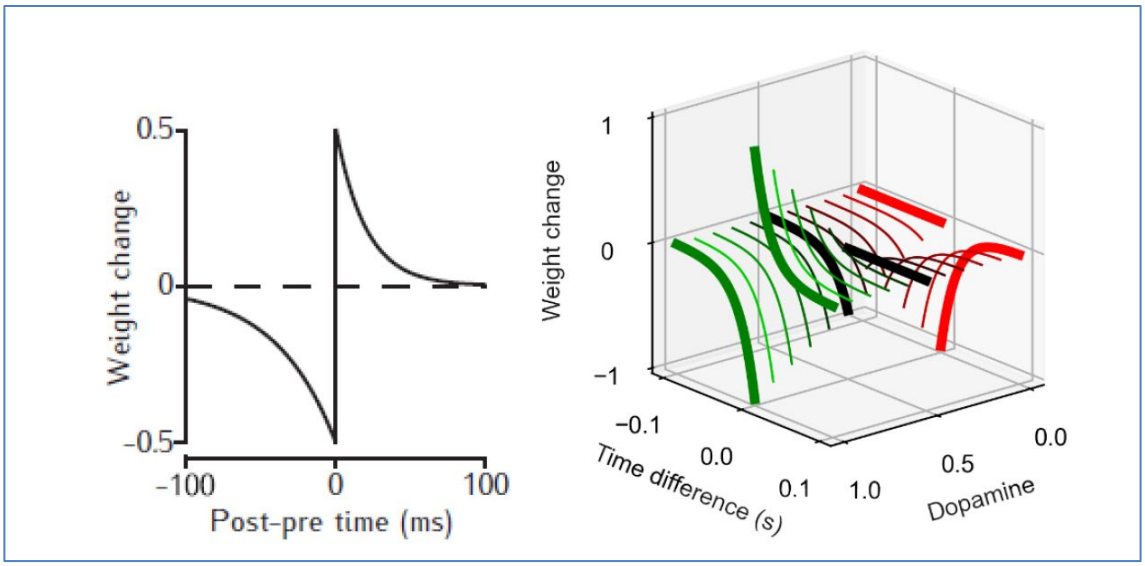

Figure 10: Different kernels used in spike-timing-dependent plasticity-like rules.

(Left) Typical STDP kernel shape, showing the relationship between the relative spike timing and the weight change. (Right) An example of kernel used in STDE learning rule, where the weight change depends not only on the post-pre time difference but also on the available amount of dopamine.

3.2.4 Dopamine and acetylcholine modulation in a reinforcement learning striatal model (recent unpublished work)

In this study, UGR in collaboration with KI and KTH used a modified version of STDE to include the influence of neuromodulator acetylcholine (ACh) in addition to DA. ACh (modulated at STR by thalamic inputs) seems to have an important role regulating learning in MSNs, as ACh pauses define the time window for phasic dopamine to induce plasticity (Reynolds et al., 2022). However, it remains under discussion how these two mechanisms combine to make the striatum able to solve action-selection problems. In the present work, a striatal learning rule is proposed that uses DA as a global reward signal that modulates the kernel of the STDP-like learning rule, and ACh as a local population feedback that signals the responsibility of the recent actions.

UGR, KI and KTH have developed a reinforcement learning computational model of the striatum that demonstrates how the combination of DA and ACh influences action selection and facilitates learning, even in more complex tasks. The model was enriched with lateral connectivity and homeostatic mechanisms, increasing its robustness to parametric changes. The model demonstrated proficiency in recognizing relevant patterns and consistently selecting rewarded actions, while its homeostatic mechanisms facilitated robust learning and recovery from policy changes. Notably, incorporating ACh feedback expedited the learning process as the number of potential actions increased. This study's findings provide a promising basis for future exploration into the intricate learning mechanisms of the brain and the role of neuromodulators therein.

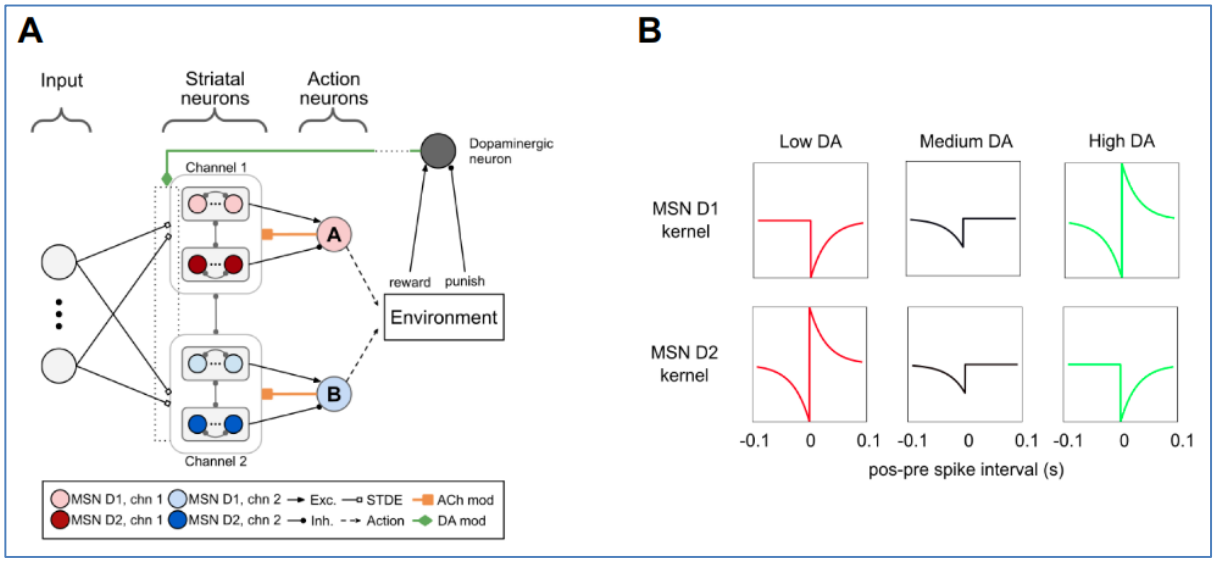

Figure 11: A) Cortico-striatal network for RL task. B) Different STDE learning kernels used.

The network, modelled after[15] and shown in Figure 11A, consists of Leaky Integrate-and-Fire (LIF) neurons organized into channels, with each representing a potential action. Each channel encompasses two MSN populations (D1 and D2 neurons), with lateral inhibition incorporated. Action neurons, which simplify other basal ganglia nuclei by integrating excitatory activity from D1 neurons and inhibitory activity from D2 neurons, select an action if the activity balance between its D1 and D2 neurons leans toward D1, causing the corresponding action neuron to spike.

The environment generates a 200 ms delayed reinforcement signal based on the last action taken and the expected action. A dopaminergic neuron sends a global reward signal to all MSNs. The learning rule uses eligibility traces that decay exponentially with a time constant of 400 ms to store the potential weight changes and apply them according to current DA level, similarly to (Gurney et al., 2015; Izhikevich, 2007). All plastic synapses share a global DA level that decays exponentially with a temporal constant of 20 ms.

Importantly, as a novel addition in this work, action neurons also transmit information about the decision made back to the MSNs, and if the action was taken, ACh levels momentarily dip. Learning is only possible when the ACh level is low in a channel.

To evaluate the robustness of the combined synaptic and homeostatic rules, UGR, KI and KTH simulated an action-selection experiment. In this scenario, the simulated network was presented with multiple possible actions to choose from. There are as many input patterns as possible actions, and these patterns were randomly shown during the simulation 80% of the time, with noise representing the remaining 20%. Each input pattern corresponded to a specific action that, if chosen, would yield a reward. Any other action would result in a punishment. If no action was taken, neither punishment nor reward was given.

This task was tested under various conditions, such as differing numbers of possible actions and with or without the presence of ACh. As the number of potential actions increased, the task's difficulty naturally escalated. This can be seen in Figure 12, where the difficulty of the task increases with the number of possible actions, as it takes longer to achieve high accuracy. The graph also shows that with ACh the model learns much faster with a higher number of actions than the model without ACh. The only situation with no difference is where only two actions are used.

[15] González-Redondo, Á., Garrido, J., Naveros Arrabal, F., Hellgren Kotaleski, J., Grillner, S., Ros, E., 2023. Reinforcement learning in a spiking neural model of striatum plasticity. Neurocomputing 548, 126377. https://doi.org/10.1016/j.neucom.2023.126377; P4039

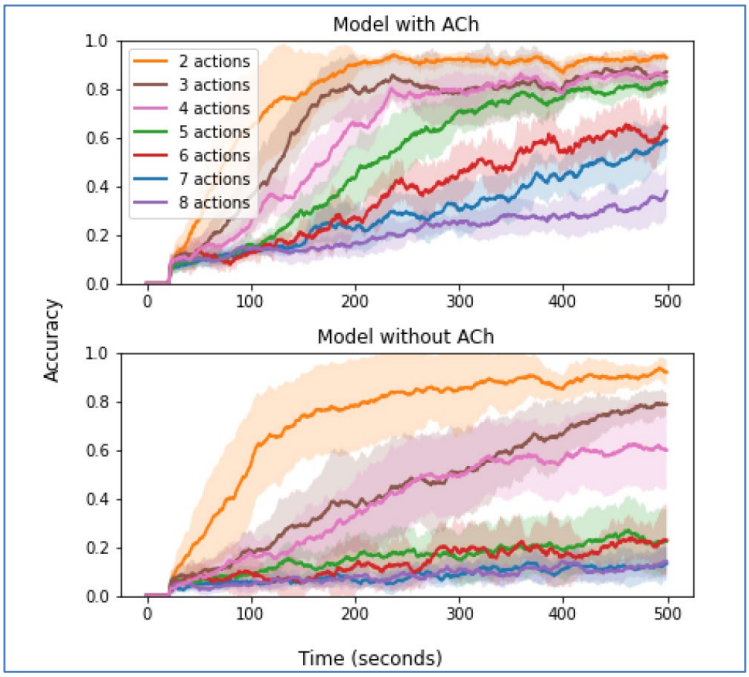

Figure 12: Accuracy evolution with and without ACh solving tasks of different difficulties.

The findings of this study suggest that the developed network model can effectively recognise relevant input patterns and make consistent, rewarding action choices in response to sensory inputs. An interesting observation was that the inclusion of ACh feedback expedited the learning process, particularly as the number of actions increased. In actual brains, ACh is modulated at STR by thalamic inputs. The present model suggests a potential role for this input in facilitating learning by confining it to specific subpopulations within the STR. This finding, along with others, offers an exciting path for further exploration in our understanding of the brain's complex learning mechanisms.

3.2.5 Pro- and anti-saccadic task simulation

UGR, KI and KTH further tested the previous striatal model using an abstract pro- and anti-saccadic task (Figure 13). In this classic task, a subject observes a dot on a screen. The dot's colour indicates the trial type: blue for pro-saccadic and pink for anti-saccadic. After 1 second, another dot of the same colour appears on either side of the screen. The subject must look towards the new target for a pro-saccadic trial, and in the opposite direction for an anti-saccadic trial, with correct actions rewarded.

The stimuli in this experiment are abstracted and represented as sets of input populations, with each population representing a state in the task. There are two trial types (pro and anti), and each trial is subdivided into five steps, resulting in a total of ten input populations. Each population contains ten neurons, with subsequent population steps overlapping by five neurons. These populations represent different states in the task and fire at a rate of 50 Hz when active and 1 Hz at the basal level.

Successful completion of a trial triggers a phasic DA increase, influencing the plasticity of the connectivity between the input and the MSNs. Incorrect actions or failures to respond result in a DA decrease (Figure 14). ACh levels are generally high, dropping locally within a channel when that channel activates sufficiently to take an action.

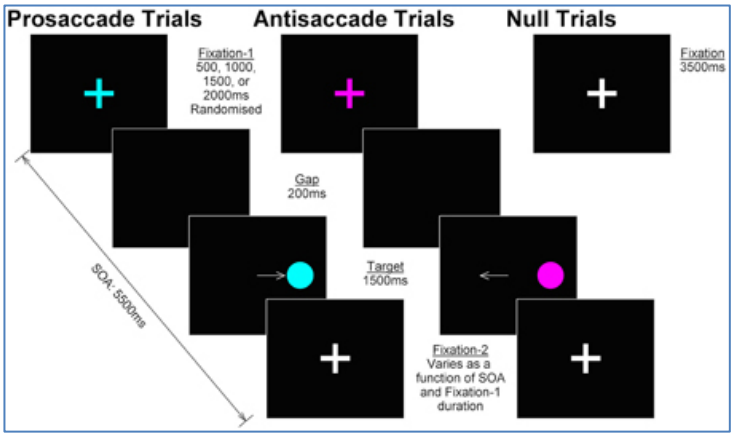

Figure 13: Diagram of the classical pro-saccade (left side) and anti-saccade (right side) tasks.

The colour of the fixation point indicated the trial type (cyan for prosaccades; magenta for antisaccades). The arrow indicates the correct direction of the saccade. Figure adapted from Jamadar et al. (2015).

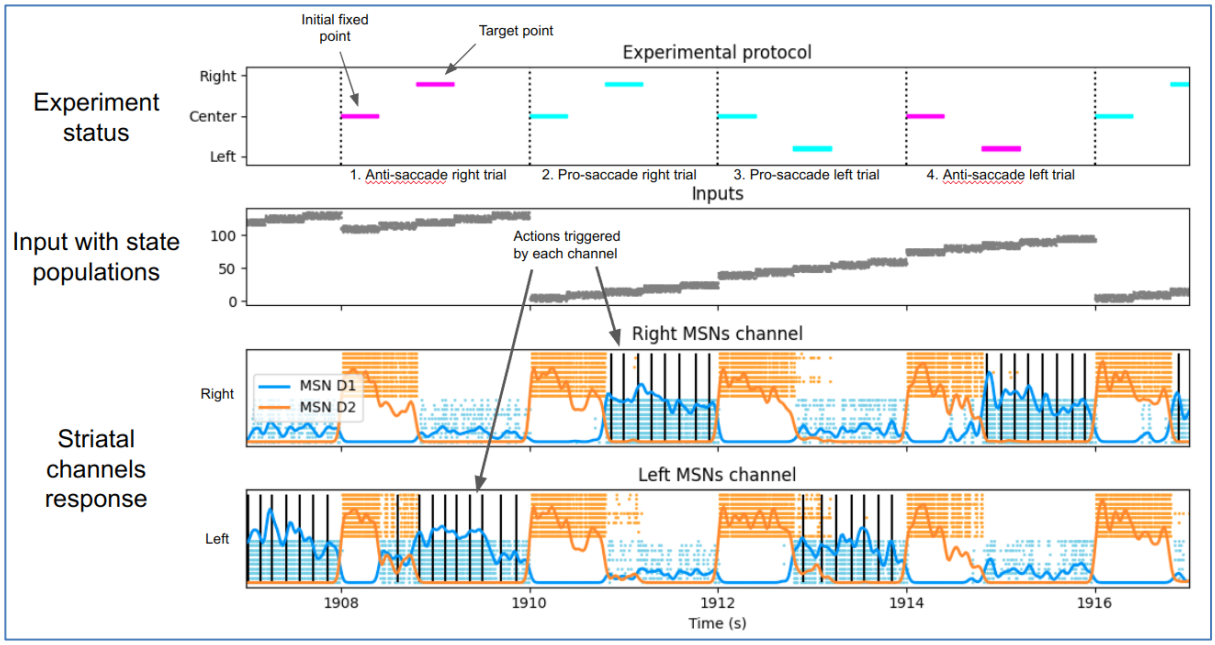

Figure 14: Experimental run example of the full model.

The horizontal axis represents time in all subplots. The top row symbolically displays the experiment status, using cyan and magenta lines for pro- and anti-saccade trials, respectively. The second row presents the spiking activity of the state populations, each trial type comprising five sequentially activated overlapping state populations. Considering the two trial types and two possible stimulus directions, there are four potential sequences and twenty total populations. The last two rows exhibit the activity of each decision channel (left and right) along with their respective MSN D1 and D2 layers. Black vertical bars indicate an action triggered by the channel.

The model's performance was evaluated based on its timely response. Correct responses are scored positively, whereas incorrect responses or non-responses negatively impact the score. The average performance over time is calculated using a moving mean (Figure 14).

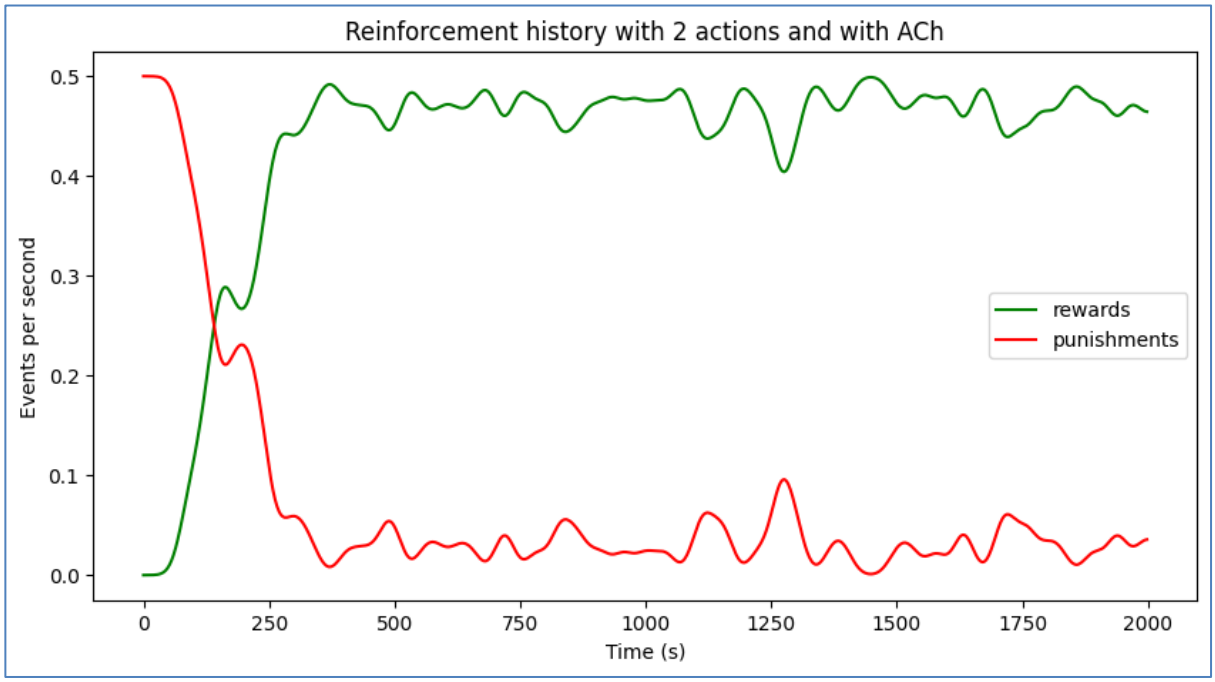

Figure 15: Performance measured as number of rewards/punishments per second.

Each trial lasts two seconds, so the maximum number of rewards is 0.5 per second. The complete network model can learn the task very fast.

The learning of the network in this task was also tested without the use of ACh neuromodulation (Figure 16), which resulted in failure (Figure 17). The model does not converge and is only capable of responding inconsistently to a single direction.

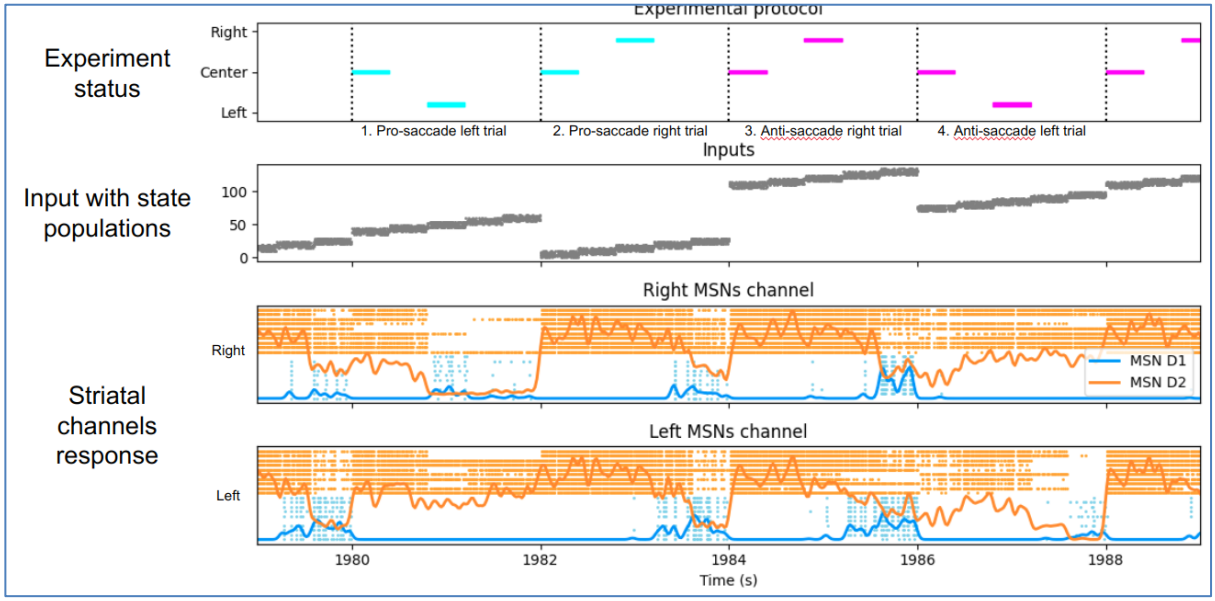

Figure 16: Experimental run example of the ablated model without ACh.

The network is not able to learn the task. As punishments are being received continuously, D2 MSN activity has increased to stop the generation of wrong decisions, effectively making the network silent.

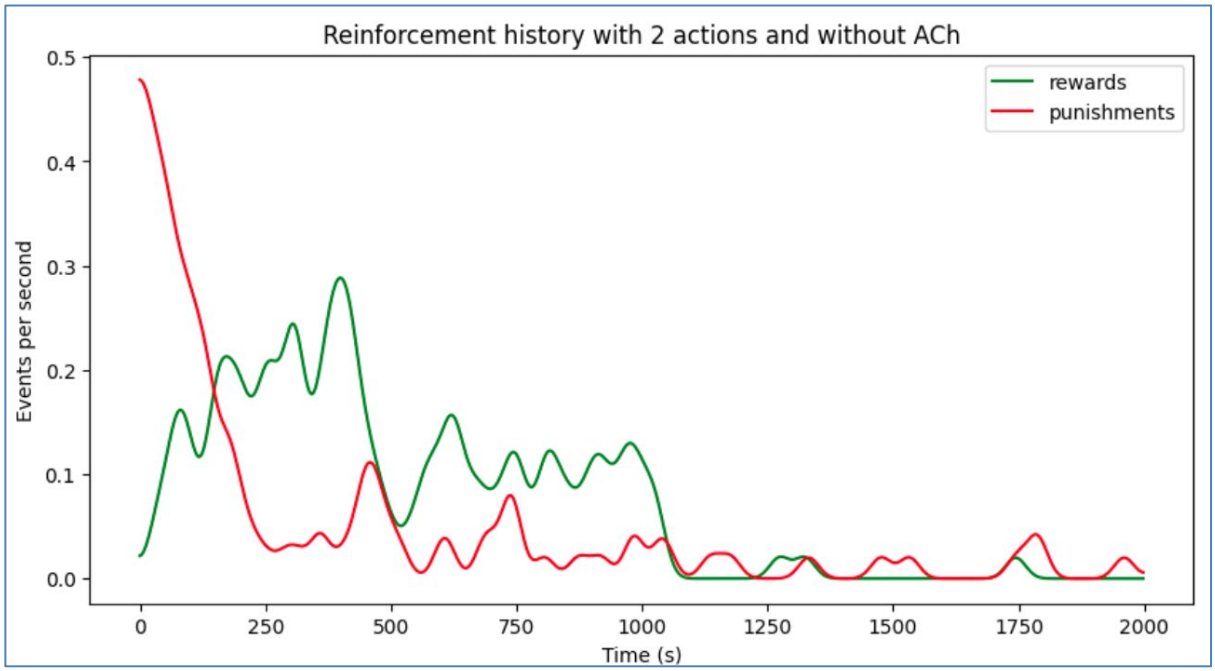

Figure 17: Performance of the ablated model without ACh.

This model fails to consistently learn the task.

4. Models of the cerebral cortex

4.1 A new computational role of cortical oscillations

Evolutionary pressure has endowed the brain with powerful mechanisms to deal with uncertain sensory inputs and perceptual multistability. Could these biological mechanisms be beneficial for neural network models that struggle to switch between interpretations (attractors) of the data? We know, on the one hand, that simulated tempering (Marinari and Parisi, 1992) helps neural networks to regularly escape the modes by flattening the probability landscape with periodical temperature modulations. On the other hand, it is known that rhythmic neural activity oscillations occurring in different frequencies are ubiquitous periodic phenomena in the brain (Buzsáki and Draguhn, 2004). In this line of research, we combined the two ideas and investigated the hypothesis that neural oscillations implement tempering and help the brain mix between interpretations.

In our model, LIF neurons receive functional input from the local network (within functionally adjacent brain areas) and background input from other areas (see Figure 17a). UBERN in collaboration with JUELICH abstract excitatory and inhibitory background inputs as stochastic Poisson processes with characteristic rates (see Figure 18a). In the high-rate regime of the cortex (Destexhe et al., 2003), a neuron's membrane potential is Gaussian distributed, and its response function can be approximated by a logistic function[16] [17]. With increasing rate, the standard deviation of the Gaussian increases, and the response function slope decreases (Figure 17b). An LIF network can then sample from a Boltzmann distribution[15], and a unit's probability of being active given all other states is a logistic function with the inverse Boltzmann temperature as slope. From this, the relationship between the background input rate and the spiking neuron's temperature was derived.

[16] Mihai A. Petrovici, Ilja Bytschok, Johannes Bill, Johannes Schemmel and Karlheinz Meier, 2015, The high-conductance state enables neural sampling in networks of LIF neurons, 24th Annual Computational Neuroscience Meeting: CNS*2015, doi: 10.1186/1471-2202-16-S1-O2, P3570

[17] Mihai A. Petrovici, Johannes Bill, Ilja Bytschok, Johannes Schemmel, and Karlheinz Meier, 2016, Stochastic inference with spiking neurons in the high-conductance state, Physical Review E, Vol. 94, No. 4, doi: 10.1103/PhysRevE.94.042312, P1213

UBERN and JUELICH showed that high temperatures lead to a more uniform network state distribution (Figure 18c). Thus, temperature variation achieves the desired flattening of the distribution. Varying the temperature periodically via the rate (e.g., as a sine wave) induces periodic entropy changes, hence implementing tempering (lower part of Figure 18d).

The mixing problem is pronounced in high-dimensional representations that are, e.g., encoded by training a hierarchical network on realistic visual data. If the model runs in generative mode, it becomes clear how oscillations shape the output into something useful. Under constant-rate noise (Figure 18e, upper), only one image class is produced, while under constant high-rate noise (Figure 18e, middle), the images are a blurry overlap of many different classes. Under tempered noise induced by cortical oscillations, the network produces diverse and clear images (Figure 18d and e, lower). The desired behaviour under tempered noise occurs because high temperatures enable switching the mode, and the subsequent temperature decrease stabilizes the new mode (Figure 18d).

With different statistical measures, it could be shown that tempering greatly improves mixing speed and enables increased target distribution coverage. Mode duration decreases for a broad range of sine wave parameters (Figure 18f). However, representation accuracy, vital for mixing quality, is guaranteed only for a smaller part of the parameter space. The trade-off between these measures revealed that the best tempering effect occurs for slower waves, suggesting links to, e.g., various phases of sleep.

This study thus identifies a new, fundamental computational role of cortical oscillations. It connects them to various phenomena in the brain, such as sampling-based probabilistic inference, memory replay, multisensory cue combination, and place cell flickering.

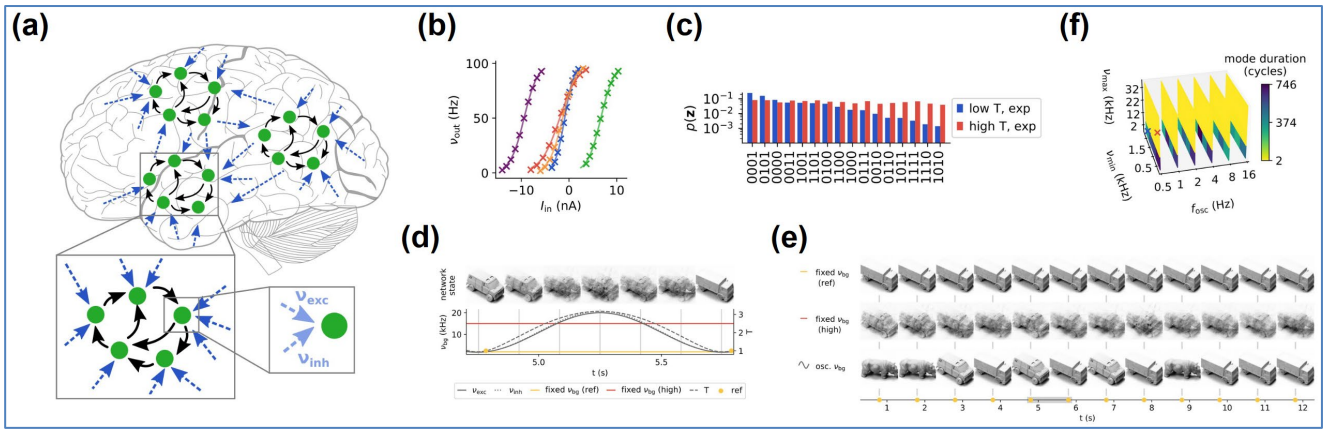

Figure 18: Cortical oscillations implement spike-based tempering.

Background input from other brain areas can be abstracted as Poisson processes with specific rates ν (a). In the high-conductance state, the activation function has a sigmoid shape, the slope of which changes with input rate (b). This observation leads to a temperature definition that has the desired effect of flattening the state distribution for high temperatures (c). Oscillating between high and low temperatures can free the network in high-dimensional data sets from deep local minima (d) and produce simultaneously diverse and clear images in contrast to constant-rate scenarios (e). This happens robustly for an ample oscillation parameter space (f).

4.2 Anomalous phase transitions in spiking neural networks

As a relatively new field of research, theoretical neuroscience has been able to draw on significant inspiration from several other well-established disciplines such as mathematics, physics, and computer science. In particular, statistical physics has played an important role in explaining many aspects of neuronal activity, both for computation and for understanding large-scale ensemble phenomena. In this regard, two of the most influential models have been the Hopfield network and the Boltzmann machine, both inspired by spin glasses and the Ising model of magnetism.

Recently, these models have also been used very successfully to model the activity of spiking neuron ensembles (see Figure 19a). A significant contribution to this body of work has been made from

within the HBP itself [15] [18] [19] [20] [21] [22]. One detail that all these models have in common is that they assume a binary form of the neuronal interaction. While this approximation holds in the more noisy, high-conductance regime of cortical activity, the less noisy regime – present, among others, in cortical down states, has previously not been explored in as much detail. The mathematical framework for spike-based sampling developed in the HBP allowed us to study both regimes in more depth (see also Sec. 4.1.1).

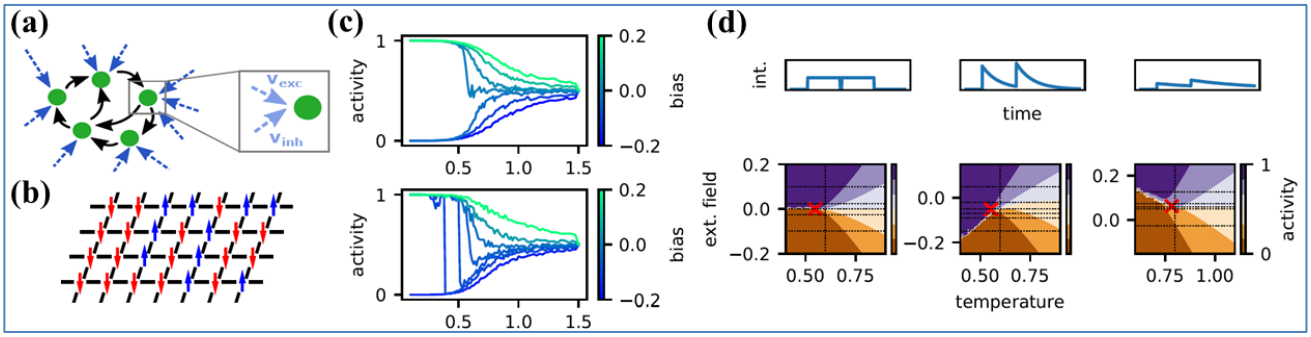

Figure 19: Interaction kernels strongly influence ensemble properties of spiking networks.

Partly adapted from [20]. Cortical neurons exhibit stochastic activity due to background noise (a), similarly to spins in magnetic lattices (b). However, ensemble behaviour depends strongly on PSP shapes (c). For biologically realistic PSP shapes, the phase diagrams of spiking networks are very different from those suggested by classical models such as Hopfield networks or Boltzmann machines (d).

In ensemble theory, the low-noise regime corresponds to a low temperature. As this temperature parameter drops, the exact shape of the synaptic interaction (PSP) between neurons plays an increasingly important role. In classical theories, neuronal lattices (see Figure 19b) undergo a phase transition at the so-called Curie temperature, where they spontaneously fall into either of two states – a highly active, synchronous state, or a state of (near) absolute silence. However, UBERN in collaboration with UHEI has shown that for a more realistic PSP shape, the phase transition is significantly more complicated (see Figure 19c, d). In particular, activity in cortical sub-networks can switch twice as the surrounding cortical activity drops: first, from a balanced state to a low-activity state, and then back up into a highly active state.

This represents a radical shift in our previous understanding of spiking ensembles and can play a significant role in understanding out-of-equilibrium cortical states or pathologies. It could also play an important role in computation: Under a more conventional view, transient states of high activity can help our brain switch between different interpretations of its sensory inputs (see also Sec. 4.1.1). However, our new insights into the low-activity regime suggest that it too might be able to facilitate such a phenomenon. Even more so, the sharp transitions between high and low activity in the low-

[18] Luziwei Leng, Roman Martel, Oliver Breitwieser, Ilja Bytschok, Walter Senn, Johannes Schemmel, Karlheinz Meier, Mihai A. Petrovici, 2018, Spiking neurons with short-term synaptic plasticity form superior generative networks, Scientific Reports, Vol. 8, No. 1, doi: 10.1038/s41598-018-28999-2, P1334

[19] Akos F. Kungl, Sebastian Schmitt, Johann Klähn, Paul Müller, Andreas Baumbach, Dominik Dold, Alexander Kugele, Eric Müller, Christoph Koke, Mitja Kleider, Christian Mauch, Oliver Breitwieser, Luziwei Leng, Nico Gürtler, Maurice Güttler, Dan Husmann, Kai Husmann, Andreas Hartel, Vitali Karasenko, Andreas Grübl, Johannes Schemmel, Karlheinz Meier, Mihai A. Petrovici, 2019, Accelerated Physical Emulation of Bayesian Inference in Spiking Neural Networks, Frontiers in Neuroscience, Vol. 13, doi: 10.3389/fnins.2019.01201, P1360

[20] Dominik Dold, Ilja Bytschok, Akos F. Kungl, Andreas Baumbach, Oliver Breitwieser, Walter Senn, Johannes Schemmel, Karlheinz Meier, Mihai A. Petrovici, 2019, Stochasticity from function — Why the Bayesian brain may need no noise, Neural Networks, Vol. 119, doi: 10.1016/j.neunet.2019.08.002, P1447

[21] Jakob Jordan, Mihai A. Petrovici, Oliver Breitwieser, Johannes Schemmel, Karlheinz Meier, Markus Diesmann, Tom Tetzlaff, 2019, Deterministic networks for probabilistic computing, Scientific Reports, Vol. 9, No. 1, doi: 10.1038/s41598-019-54137-7, P843

[22] Agnes Korcsak-Gorzo, Michael G. Müller, Andreas Baumbach, Luziwei Leng, Oliver J. Breitwieser, Sacha J. van Albada, Walter Senn, Karlheinz Meier, Robert Legenstein, Mihai A. Petrovici, 2022, Cortical oscillations support sampling-based computations in spiking neural networks, PLOS Computational Biology, Vol. 18, No. 3, doi: 10.1371/journal.pcbi.1009753, P2806

noise regime might allow particularly robust switching between different thoughts or memories without the longer transients incurred by slower cortical oscillations.

4.3 Multi-area models of macaque visual cortex

JUELICH developed multi-area models of all vision-related areas in one hemisphere of macaque cortex, based on previous work[23] [24]. The models are implemented in NEST running on HPC systems and contain the full density of neurons and synapses in each local circuit, for a total of around 4 million neurons connected via 24 billion synapses. The usage and scaling of this model was described in a book chapter[25]. In collaboration with UNIMAN, work was done to port the published model to SpiNNaker. JUELICH developed a PyNN version of the model and UNIMAN so far ran the 32 isolated microcircuits of the model as well as a subset of the microcircuits with their full interconnectivity successfully. The model was further implemented on a GPU cluster and the simulation performance was compared with that of NEST running on CPUs[26]. This work contributed to the incorporation of a GPU component into NEST.

In new work, JUELICH incorporated joint clustering of excitatory and inhibitory neurons into the model (Figure 20) to support inter-area propagation. As a positive side effect, JUELICH showed that such clustering leads to more plausible firing rate distributions across all vision-related areas (Figure 21). Furthermore, the clustered model reproduces a reduction in trial-to-trial variability, as measured by the Fano factor, upon stimulation.

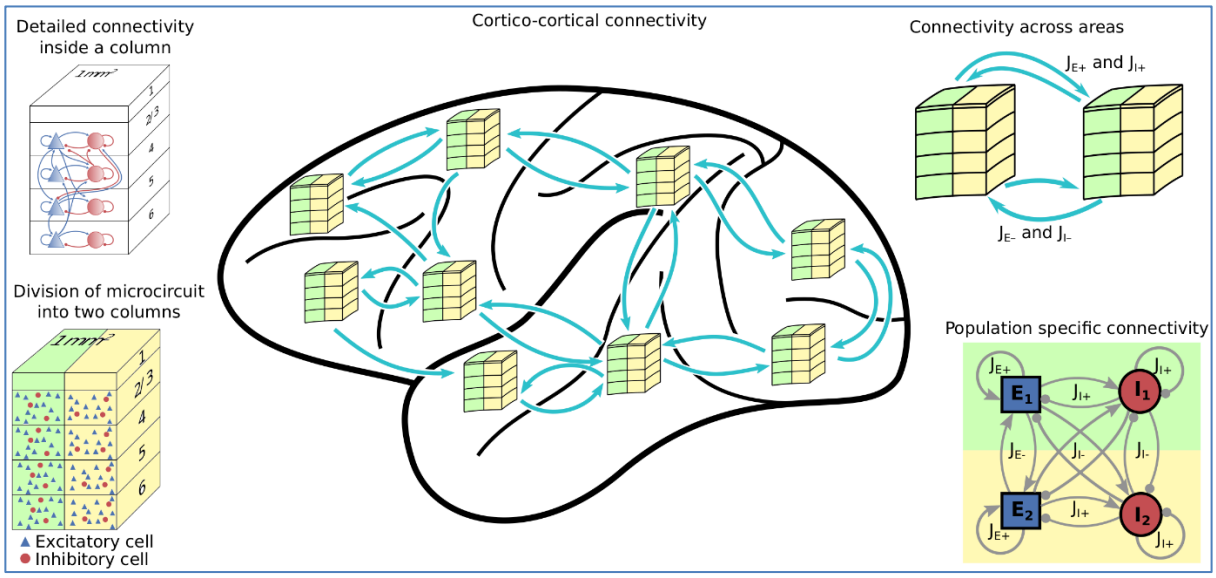

Figure 20: Schematic illustration of the clustered multi-area model of macaque visual cortex.

[23] Maximilian Schmidt, Rembrandt Bakker, Kelly Shen, Gleb Bezgin, Markus Diesmann, Sacha Jennifer van Albada, 2018, A multi-scale layer-resolved spiking network model of resting-state dynamics in macaque visual cortical areas, PLOS Computational Biology, Vol. 14, No. 10, doi: 10.1371/journal.pcbi.1006359, P1457

[24] Maximilian Schmidt, Rembrandt Bakker, Claus C. Hilgetag, Markus Diesmann, Sacha J. van Albada, 2017, Multi-scale account of the network structure of macaque visual cortex, Brain Struct Funct, doi: 10.1007/s00429-017-1554-4, P1036

[25] Sacha J. van Albada, Jari Pronold, Alexander van Meegen, Markus Diesmann, 2021, Usage and Scaling of an Open-Source Spiking Multi-Area Model of Monkey Cortex, Lecture Notes in Computer Science, doi: 10.1007/978-3-030-82427-3_4, P2890

[26] Gianmarco Tiddia, Bruno Golosio, Jasper Albers, Johanna Senk, Francesco Simula, Jari Pronold, Viviana Fanti, Elena Pastorelli, Pier Stanislao Paolucci, Sacha J. van Albada, 2022, Fast Simulation of a Multi-Area Spiking Network Model of Macaque Cortex on an MPI-GPU Cluster, Frontiers in Neuroinformatics, Vol. 16, doi: 10.3389/fninf.2022.883333, P3318

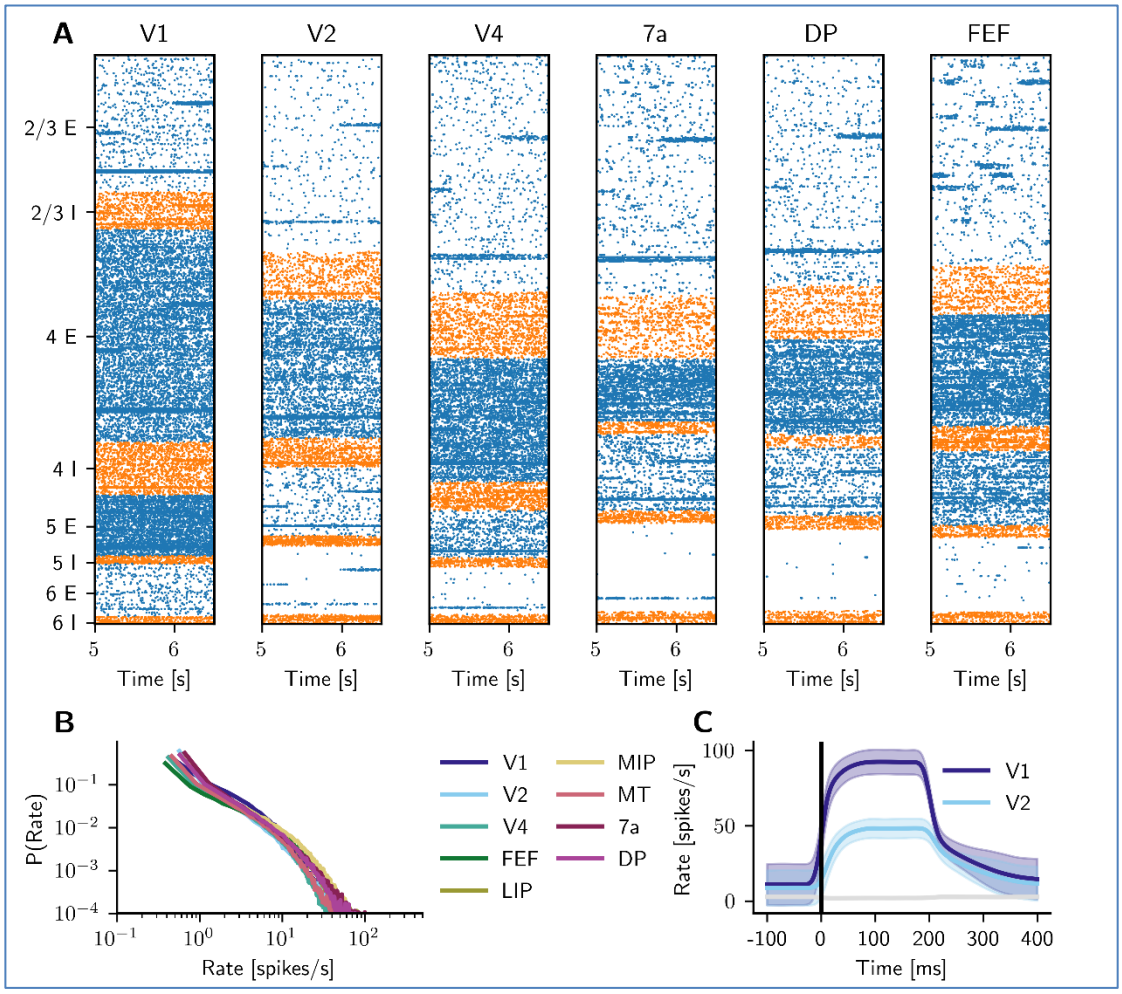

Figure 21: The model supports plausible firing rate distributions and inter-area propagation.

4.4 Spatially resolved large-scale model of macaque V1

In further ongoing work, JUELICH has modelled an expanded version of the primary visual cortex (V1) component of the macaque cortical model, measuring 16 mm2, and incorporated distance-dependent connectivity. The resulting large-scale model is implemented using NEST and run on the JUELICH supercomputers. The connectivity was further refined at the level of cortical layers and populations based on a collation of axonal tracing data (Vanni, 2020), and by incorporating push-pull connectivity (Antolik et al., 2018). The model reproduces various aspects of visual orientation maps (Figure 22) based on the input from effective thalamic units (Sadeh and Rotter, 2014). In future work, the resting-state activity of the model will be explored and made increasingly compatible with available experimental data; and the interactions between V1 and higher cortical areas will be explored.

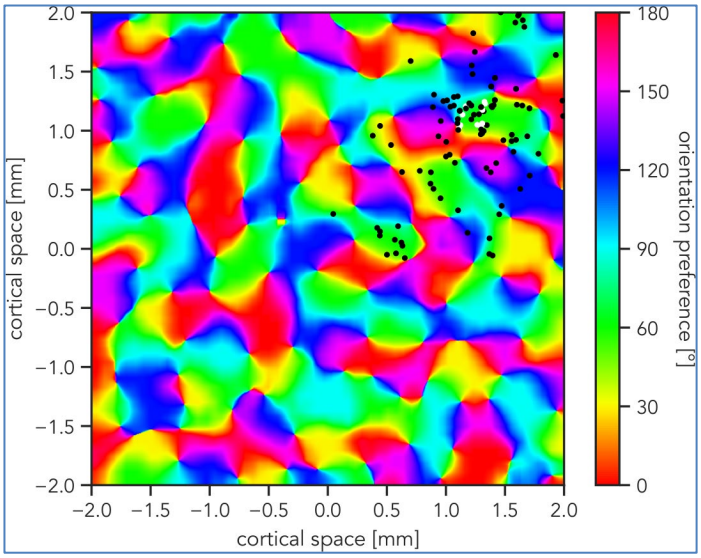

Figure 22: Orientation map of the macaque V1 model.

Dots highlight the elliptical nature of patchy connections. Black dots represent post-synaptic, white dots pre-synaptic neurons.

4.5 Retinal sampling explains aspects of visual cortex

How we perceive the world is an interplay of various neural structures and mechanisms. The retina, with its non-uniform distribution of retinal ganglion cells (RGCs), plays a pivotal role in visual perception. The fovea, the central region of the retina, boasts the highest RGC density, ensuring high-resolution processing of visual stimuli. As one moves away from the fovea, the density of RGCs diminishes, leading to reduced visual acuity in the periphery. This non-uniformity underpins our need for eye movements, focusing on salient regions in the visual field to extract meaningful information. To emulate this eccentricity-dependent distribution of RGCs within computational models, partner UM (P117) recently developed a retinal sampling layer[27] (RSL).

At its core, the RSL is designed to transform uniformly sampled square images into non-uniformly sampled counterparts, reflecting the RGC distribution in the human retina. To map visual field coordinates to ganglion cell coordinates, the RSL integrates and subsequently inverts an empirically derived function that relates ganglion cell density to eccentricity in the visual field.

Convolutional Neural Networks (CNNs) have in recent years become increasingly adopted as a model of sensory neural structures, especially in the domain of vision. However, a generic CNN lacks the biological realism that the human visual system exhibits. The introduction of the RSL into CNN architectures addresses this gap. By pre-processing input images using the RSL before they are fed into a CNN, the network's overall realism can be enhanced. Indeed, when CNNs were augmented with the RSL, they began to display organizational principles reminiscent of the primate visual cortex, such as cortical magnification and eccentricity-dependent receptive field size. Furthermore, these CNNs showcased a radial bias, a phenomenon wherein orientation tuning of neural populations is biased towards radial orientations (i.e., those orientations that if they were extended beyond the receptive field of the population would intersect with the origin). Most notably, the CNN exhibits radial bias only for grating stimuli whose spatial frequency exceeds a certain threshold whereas low spatial frequency stimuli elicit a bias towards orientations that are orthogonal to the radial axis. These findings have led to novel hypotheses about the origin of radial bias that are currently being investigated with ultra-high-field imaging. This underscores the potential of integrating biologically accurate modules like the RSL into deep neural network architectures (also cf. Figure 23).

[27] Danny da Costa, Lukas Kornemann, Rainer Goebel, Mario Senden (2023), Unlocking the Secrets of the Primate Visual Cortex: A CNN-Based Approach Traces the Origins of Major Organizational Principles to Retinal Sampling, BioRxiv, http://dx.doi.org/10.1101/2023.04.25.538251, P3995

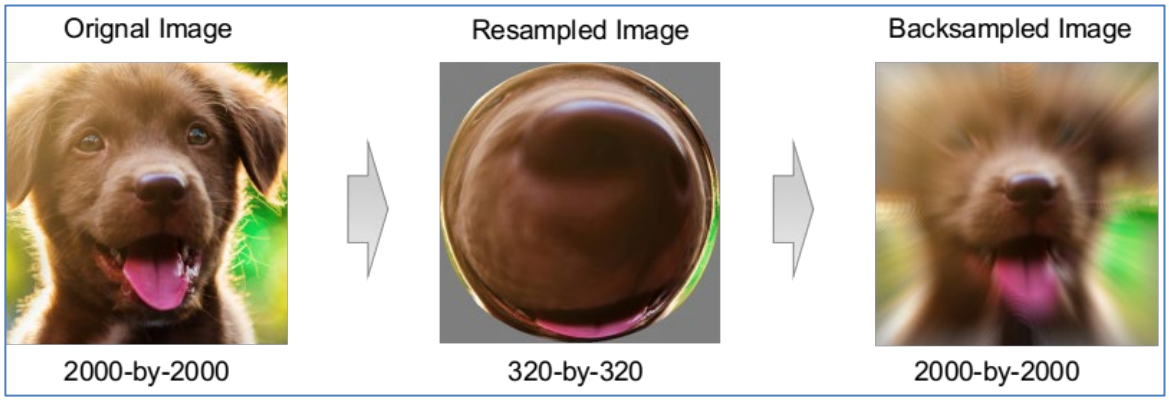

Figure 23: Retinal sampling of a square image.

The input image (left) is sampled according to RGC distributions, resulting in a retinal ganglion cell representation. The image shows an enhanced central region (foveal vision) with progressive loss of information as eccentricity increases (peripheral vision). Inverting the sampling process transforms the retinal ganglion cell representation into a representation similar to human perception. Note that while the distortion caused by retinal sampling may appear extreme, the perceptual image reveals that a significant amount of information is retained.

5. Looking Forward

To obtain a more integrated picture of brain function at the spiking level, it will be important to combine the models described here into more unified architectures, enabling the investigation of coordinated information processing in the various brain regions. Functions covered should go beyond simple visuomotor functions such as eye and arm movements to more closely approximate ongoing sensorimotor coordination in natural environments. Further, the partners will continue to enhance the modelling tools to which the present work has contributed, including the Brain Scaffold Builder and Snudda, extending their functionality, documentation, compatibility with different computing platforms and softwares, and integration into unified workflows.

6. Annex: Non-PLUS References

Antolík, J., Monier, C., Frégnac, Y., & Davison, A. P. (2018). A comprehensive data-driven model of cat primary visual cortex. *BioRxiv, 416156*. doi: https://doi.org/10.1101/416156

Berke, J.D., 2018. What does dopamine mean? Nat. Neurosci. 21, 787–793. https://doi.org/10.1038/s41593-018-0152-y

Buzsaki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. science, 304(5679), 1926-1929.

Destexhe, A., Rudolph, M., & Paré, D. (2003). The high-conductance state of neocortical neurons in vivo. Nature reviews neuroscience, 4(9), 739-751.

Fino, E., Venance, L., 2010. Spike-timing dependent plasticity in the striatum. Front. Synaptic Neurosci. 2.

Foster, N.N., Barry, J., Korobkova, L., Garcia, L., Gao, L., Becerra, M., Sherafat, Y., Peng, B., Li, X., Choi, J.-H., Gou, L., Zingg, B., Azam, S., Lo, D., Khanjani, N., Zhang, B., Stanis, J., Bowman, I., Cotter, K., Cao, C., Yamashita, S., Tugangui, A., Li, A., Jiang, T., Jia, X., Feng, Z., Aquino, S., Mun, H.-S., Zhu, M., Santarelli, A., Benavidez, N.L., Song, M., Dan, G., Fayzullina, M., Ustrell, S., Boesen, T., Johnson, D.L., Xu, H., Bienkowski, M.S., Yang, X.W., Gong, H., Levine, M.S., Wickersham, I., Luo, Q., Hahn, J.D., Lim, B.K., Zhang, L.I., Cepeda, C., Hintiryan, H., Dong, H.-W., 2021. The mouse cortico-basal ganglia-thalamic network. Nature 598, 188–194. https://doi.org/10.1038/s41586-021-03993-3

González-Redondo, Á., Garrido, J., Naveros Arrabal, F., Hellgren Kotaleski, J., Grillner, S., Ros, E., 2023. Reinforcement learning in a spiking neural model of striatum plasticity. Neurocomputing 548, 126377. https://doi.org/10.1016/j.neucom.2023.126377

Graybiel, A.M., 1998. The Basal Ganglia and Chunking of Action Repertoires. Neurobiol. Learn. Mem. 70, 119–136. https://doi.org/10.1006/nlme.1998.3843

Grillner, S., Hellgren, J., Ménard, A., Saitoh, K., Wikström, M.A., 2005. Mechanisms for selection of basic motor programs – roles for the striatum and pallidum. Trends Neurosci. 28, 364–370. https://doi.org/10.1016/j.tins.2005.05.004

Gurney, K., Prescott, T.J., Redgrave, P., 2001. A computational model of action selection in the basal ganglia. I. A new functional anatomy. Biol. Cybern. 84, 401–410. https://doi.org/10.1007/PL00007984

Gurney, K.N., Humphries, M.D., Redgrave, P., 2015. A New Framework for Cortico-Striatal Plasticity: Behavioural Theory Meets In Vitro Data at the Reinforcement-Action Interface. PLOS Biol. 13, e1002034. https://doi.org/10.1371/journal.pbio.1002034

Hart, A.S., Rutledge, R.B., Glimcher, P.W., Phillips, P.E.M., 2014. Phasic Dopamine Release in the Rat Nucleus Accumbens Symmetrically Encodes a Reward Prediction Error Term. J. Neurosci. 34, 698–704. https://doi.org/10.1523/JNEUROSCI.2489-13.2014

Hikosaka, O., Takikawa, Y., Kawagoe, R., 2000. Role of the Basal Ganglia in the Control of Purposive Saccadic Eye Movements. Physiol. Rev. 80, 953–978. https://doi.org/10.1152/physrev.2000.80.3.953

Hunnicutt, B.J., Jongbloets, B.C., Birdsong, W.T., Gertz, K.J., Zhong, H., Mao, T., 2016. A comprehensive excitatory input map of the striatum reveals novel functional organization. eLife 5, e19103. https://doi.org/10.7554/eLife.19103

Ito, M. (2000). Mechanisms of motor learning in the cerebellum. *Brain research*, *886*(1-2), 237-245.), doi: 10.1016/s0006-8993(00)03142-5

Izhikevich, E.M., 2007. Solving the Distal Reward Problem through Linkage of STDP and Dopamine Signaling. Cereb. Cortex 17, 2443–2452. https://doi.org/10.1093/cercor/bhl152

Jamadar, S.D., Johnson, B.P., Clough, M., Egan, G.F., Fielding, J., 2015. Behavioral and Neural Plasticity of Ocular Motor Control: Changes in Performance and fMRI Activity Following Antisaccade Training. Front. Hum. Neurosci. 9.

Levy, W.B., Steward, O., 1983. Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. Neuroscience 8, 791–797. https://doi.org/10.1016/0306-4522(83)90010-6

Marinari, E., & Parisi, G. (1992). Simulated tempering: a new Monte Carlo scheme. Europhysics letters, 19(6), 451.

Medina, J. F. (2019). Teaching the cerebellum about reward. *Nature neuroscience*, *22*(6), 846-848. doi: 10.1038/s41593-019-0409-0

Mogenson, G.J., Jones, D.L., Yim, C.Y., 1980. From motivation to action: Functional interface between the limbic system and the motor system. Prog. Neurobiol. 14, 69–97. https://doi.org/10.1016/0301-0082(80)90018-0

Pierrot-Deseilligny, E., & Burke, D. (2012). *The circuitry of the human spinal cord: spinal and corticospinal mechanisms of movement*. Cambridge University Press, https://doi.org/10.1017/CBO9781139026727.

Redgrave, P., Prescott, T.J., Gurney, K., 1999. The basal ganglia: A vertebrate solution to the selection problem? Neuroscience 89, 1009–1023.

Reynolds, J.N.J., Avvisati, R., Dodson, P.D., Fisher, S.D., Oswald, M.J., Wickens, J.R., Zhang, Y.-F., 2022. Coincidence of cholinergic pauses, dopaminergic activation and depolarisation of spiny projection neurons drives synaptic plasticity in the striatum. Nat. Commun. 13, 1296. https://doi.org/10.1038/s41467-022-28950-0

Reynolds, J.N.J., Wickens, J.R., 2002. Dopamine-dependent plasticity of corticostriatal synapses. Neural Netw. 15, 507–521. https://doi.org/10.1016/S0893-6080(02)00045-X

Sadeh, S., & Rotter, S. (2014). Statistics and geometry of orientation selectivity in primary visual cortex. *Biological cybernetics, 108, 631-653.*

Suryanarayana, S.M., Hellgren Kotaleski, J., Grillner, S., Gurney, K.N., 2019. Roles for globus pallidus externa revealed in a computational model of action selection in the basal ganglia. Neural Netw. 109, 113–136. https://doi.org/10.1016/j.neunet.2018.10.003

Vanni, S., Hokkanen, H., Werner, F., & Angelucci, A. (2020). Anatomy and physiology of macaque visual cortical areas V1, V2, and V5/MT: bases for biologically realistic models. *Cerebral cortex, 30(6), 3483-3517.* https://doi.org/10.1093/cercor/bhz322

Xiao, L., Roberts, T.F., 2021. What Is the Role of Thalamostriatal Circuits in Learning Vocal Sequences? Front. Neural Circuits 15.