SAPIENZA
UNIVERSITÀ DI ROMA

Deep sleep and its cognitive effects

Slow wave activity and the learning and sleep cycle

Dipartimento di Psicologia

Dottorato di Ricerca in Neuroscienze del Comportamento – XXXIII Ciclo

Candidate

Elena Pastorelli
ID number 662116

Thesis Advisor

Prof. Stefano Ferraina

Co-Advisors

Prof. Paolo del Giudice
Prof. Pier Stanislao Paolucci

June 2021, revised after comments from reviewers

Deep sleep and its cognitive effects
Ph.D. thesis. Sapienza – University of Rome

This thesis has been typeset by LaTeX and the Sapthesis class.

Author's email: elena.pastorelli@uniroma1.it

Abstract

Sleep is definitely an essential cerebral state for all living beings. The occurrence of sleep passed through the evolutionary sieve and is widespread in animal species. Sleep is known to be beneficial to cognitive and mnemonic tasks, while chronic sleep deprivation is detrimental. Indeed, young humans spend most of their time sleeping and, at the same time, they are the subjects that learn at faster rates. In adults, sleep deprivation is detrimental for cognition and it is one of the worst tortures that can be inflicted. Despite the importance of the phenomenon in relation with the cognitive effects, a complete understanding of its functions and underlying mechanisms is still lacking.

Recent studies focused the attention on a specific phase of sleep in relation to cognition and learning: the deepest of physiological non-REM sleep stage in which specific phenomena like Slow Oscillations (SO) are the default emergent activity of the cortical network and are observed as an alternation between *Down states* (characterized by nearly silent neurons) and *Up states* (with intense neuronal activity), occurring at a frequency in the range $[0.5, 4]$ Hz (*delta* band).

Among the contribution of this work, there is the building of fundamental instruments for the study of this phenomenon and of its interplay with learning, also thanks to the experimental context in which I am involved: the European *Human Brain Project* (HBP) and *The Motor Control and Cognition Lab* (MCC_lab) directed by Prof. Stefano Ferraina at Sapienza University. Basically, we developed an analysis tool for the extraction of salient feature (velocity, direction, planarity of the slow waves and other specific features) from the experimental recordings *in vivo* acquired with different methodologies (electrophysiology, calcium imaging, etc.), and simulations tools able to reproduce the behaviour of large portion of the cerebral cortex.

To complement this activity, we built theoretical models and simulation tools able to reproduce the slow wave activity and their effect on learning, demonstrating

that such a cerebral state has a fundamental role in memory consolidation. The simulated outputs of these theoretical models are then analyzed with the same instruments used to analyze *in vivo* recorded data, allowing a faster and simpler comparison between experimental and synthetic data.

In the end, I also had the chance to analyze signals recorded *in vivo* on monkeys during learning tasks, more specifically during the countermanding and the transitive inference tasks. I had not the possibility to verify the effect of sleep on these tasks, due to the impossibility of setting-up new experiments on animals in the last period, but the activity done, besides reporting interesting results, was preparatory to the study of the role of sleep on learning process and memory consolidation. Moreover, it gave me the possibility to work on data directly acquired in an experimental lab, that I consider a complementary activity essential for my scientific background.

Contents

Chapter 1

Scientific question: the cognitive effects of sleep

1.1 Introduction

This work inserts in a larger research plan that aims to the creation of data-driven models able to express the cognitive effects of deep-sleep, its interplay with awake learning and the details of the spatio-temporal features of cortical NREM, at large scale and high-resolution. In the future, we plan to extend the current study to the activity and effects of REM sleep, to reproduce the specific effects and characteristics over a larger set of accessible brain-states. To reach this goal, several aspects must be faced in different directions, in order to reach a "whole system" that empowers the study of interest.

One of the first objectives to be afforded was the creation of an environment allowing for large-scale simulations of cortical activity (e.g. the whole mouse hemisphere), able to reproduce different states of such a complex system. In particular, the large-scale model, made up of millions of neurons and billions of synapses, must be able to reproduce the behaviour of the cortex both during sleep (more specifically the deep-sleep phase of slow waves considered as the major responsible for memory consolidation) and during awake-like states.

On the other hand, the study of the interplay between deep-sleep and learning also requires the use of models able to learn. These models must be provided of

synaptic plasticity, that is quite expensive in terms of computational resources, about four times more expensive and above all with a requirement in terms of simulation length that approach millions of seconds. Due to this, plasticity was not used in the setup of the large-scale model previously described. In addition, a simplified small model, made up of a very limited number of neurons, allows for a deeper control of the learning process. These were the main reasons that drove the choice to study this aspect on a simplified thalamo-cortical model. The further step will be that to use the knowledge acquired in creating large-scale simulations to increase the computational power of the learning model itself.

On top of this, tools are required to verify the results of simulations and their relevance in comparison with experiments: with the mean of an iterative process of simulation and results comparison, parameter refinement of the used models will convey toward a final refined system able to accurately reproduce physical phenomena.

The models above described have been calibrated using data collected in rodents under anesthesia, a proxy for NREM sleep that could however introduce significant difference compared to physiological sleep in humans.

To overcome this limit, the work should have included, according to the original plan, a thorough analysis of electro-physiological data acquired on monkeys during learning tasks and involving sleep, aiming to confirm (or reject) the findings obtained with the above rodent-based models and simulations. Unfortunately, due to the COVID-19 contingency, the last part of my activity could only be considered a sort of "methodological" section, in which I could experiment the analysis of *in vivo* data recorded in the past years by *The Motor Control and Cognition Lab*, without any interaction with sleep. I clearly understand that this section could be considered a bit disconnected from the rest of the work, but represented for me a really fruitful and formative activity, preparatory for the study of sleep effects on cognition tasks in the next future.

1.2 Sleep: a phenomenon essential for cognition and learning

Human brains spend about one-third of their life-time sleeping. Sleep is present in every animal species that has been studied (see Tononi and Cirelli 110). This happens notwithstanding two negative aspects: the danger caused by sleep, that diminishes the capability to defend from predators and other threats, and the reduction of time available for activities targeting immediate rewards (e.g. hunting or gathering food). Having survived the evolutionary selection in all species, sleep must therefore provide strong advantages. Another notable fact is that newborns' human brains occupy the majority of their time asleep, nevertheless they learn at a very fast rate. The reader can see, for instance, Tononi et al. 111, containing a compendium of motivations for the study of sleep, also in relation to consciousness. Moreover, even if occasional awakening does not seriously impair brain biology and cognitive functions, in the long term chronic deprivation of sleep produces measurable effects on cognition, mood and health (see Killgore 53). Experimental studies like Tononi and Cirelli 110 and Watson et al. 118 investigated the effects of sleep on firing rates and synaptic efficacies. In 110, Tononi and the other authors formulated hypotheses about homeostatic processes occurring during sleep. A possible motivation for the brain entering in the sleep activity would be to set a better energy consumption regime during next wakefulness cycle. This might be obtained by reducing firing rates and the amplitude of evoked post-synaptic potentials. It could be reached by pruning non necessary synapses and by reducing synaptic weights within the limits imposed by the conservation of adequate coding for memories. In 118 Watson et al. propose a novel intriguing experimental evidence. They used large-scale recordings to examine the activity of neurons in the frontal cortex of rats and observed that neurons with different pre-sleep firing rate are differentially modulated by different sleep substates (REM, non-REM and micro arousal). Sleep activity such as slow waves activity and sharp-waves ripples have been shown to be beneficial for memory consolidation (see

Walker and Stickgold 116 and Jadhav et al. 51) and task performances optimization (Smulders et al. 100).

The impact that sleep has on learning and cognition, drove my work on the study of this phenomena. In particular, I concentrated on the specific phase of slow oscillations and on their effects on awake performance. At the neural population level, in both deep-sleep and anesthesia, slow oscillations emerge in the deepest stage of NREM sleep (stages IV), from the alternation between transients of high neural firing (UP states) and transients of near silence (DOWN states) (see Steriade et al. 105), occurring at a frequency in the range $[0.5, 4]$ Hz (*delta* band).

In a sleeping brain, SO are not an isolated dynamical pattern. Indeed, Up and Down states involves large parts of the cortical tissue and propagate on the cortical surface as travelling waves. Observing the sleeping brain at a macroscopic scale, a spatial component appears, that transforms SO in a phenomenon synchronized in space and time as traveling waves (slow wave activity, SWA). Slow waves has been observed both in *in vivo* and *in vitro* (cortical slices) with speeds in the order of few tens of millimiters per second (ranging from 10mm/s to 30mm/s depending on the animals, the kind of sleep, in case from the used anaesthetic, etc).

In the following I try to analyze, both with models and data analysis, the effects of deep-sleep-like slow oscillation activity on learning and memory consolidation.

1.3 Analysis of experimental data and simulations of cortical Slow Waves

Today, the richness of data and methodologies, used for analysis in the neurophysiological field, constitute a key point for the study of brain dynamics and their function. A landscape of different recording techniques spotlights on different aspects of neural activity. Recorded data vary among them for several aspects: kind of signals (e.g. electrophysiological, optical imaging, magnetic resonance, single unit), time resolution (from milliseconds to seconds), space resolution (from millimeters to the whole brain), and so on (Chen et al. 19).

From one side, this facilitates scientific progress by enabling a variety of new methodological approaches and analytical findings. On the other hand, this diversity makes it more difficult to relate complementary yet different approaches, to make the results comparable, and to incorporate them into models, such that the heterogeneity of analytical methods, tools, and even terminologies makes it difficult to form a coherent understanding. In other words, making scientific progress by building a cumulative understanding based on existing results is every day more challenging.

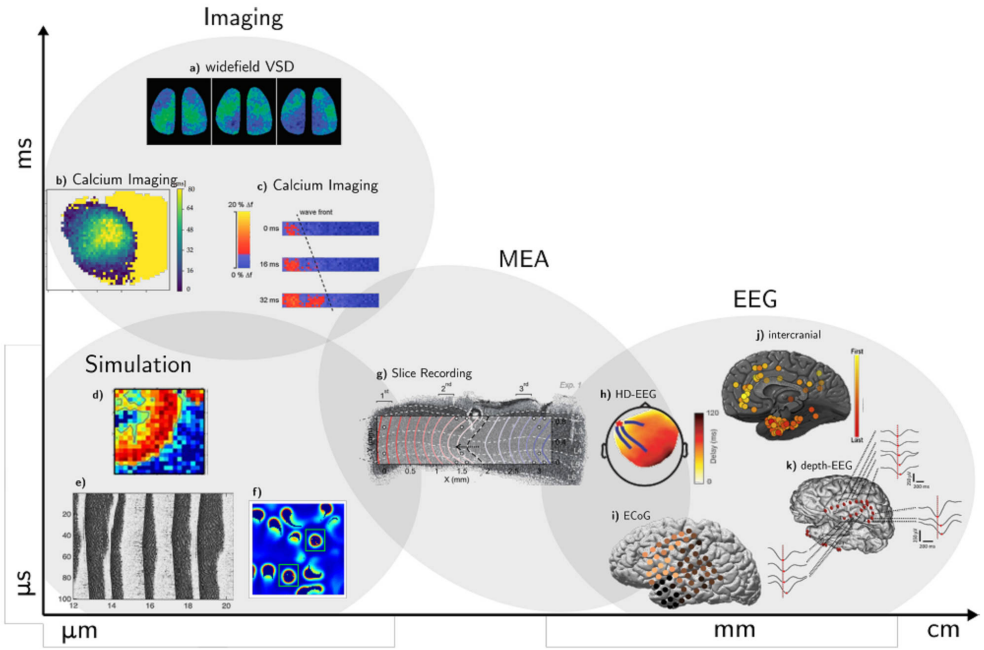

Figure 1.1. Multi-scale and multi-methodology recording of a the slow waves phenomenon. Different kind of recording produce different readings of the same phenomenon in space resolution, ranging from micron to few centimeters, and in time resolution, from microseconds to milliseconds and up. Adapted from Gutzen et al. (43)

.

The slow waves study is an example of this scenario. The phenomenon can be observed during deep-sleep and anesthesia with various measurement techniques and in a wide variety of species (Adamantidis et al. 2, Alkire et al. 3), from mouse to primates, up to the humans (see Figure 1.1 adapted from Gutzen et al. 43). At the same time, several models can be used to replicate this phenomenon. This poses a requirement of cross-domain comparison in order to reach a comparability, needed

both for models calibration and validation and for the quantification of experimental variability.

In our work we thought that the use of a unified software analysis pipeline (De Bonis, et al. 25) could be a boost in the direction of improving the understanding of neurological phenomena and on the comparison of results coming from different acquisition methodologies. Moreover, it can give a help also in the direction of reproducibility of data analysis for all the scientific experiments who gives full access to experimental results (see Plesser, 83).

The construction of such an analysis tool is not simple, when we take into account the different data acquisition methods and also the different kind of observables we want to extract from them. In this context, modularity is of crucial importance to allow the user to *build* its own analysis workflow and to interface with different input data.

In this optic, the analysis pipeline I collaborated to build during this work, under the coordination of Jülich Research Center (FZ Jülich) and the National Institute of Nuclear Physics (INFN), can be seen as a workflow that combines multi-scale, multi-species, multi-methodology experimental data sets or simulation models and can be applied to both experimental data and simulation results (see Gutzen et al. 43, and software references 119, 120, 121).

1.4 Modeling of SWA and asynchronous AW-like states

The multiscale organization of the brain is multifaceted as despite an invariance of its structural properties, like the connectivity between its components: the ongoing collective dynamics of this complex system can display disparate and qualitatively different states. A sleeping brain, for instance, expresses at the single-neuron and local network levels slow oscillations (SO) of activity which at a macroscopic scale appears to be synchronized in space and time as traveling waves (slow wave activity, SWA). The same brain during quiet wakefulness (QW) shows an asynchronous distributed activity which, even in the absence of interactions with the environment, displays

a multiscale correlation continuously changing in time, eventually highlighting the existence of competing subnetworks.

Several computational models are available to provide the community with tested and proven open-source models to be used in new studies and implementations. Just to cite some of them, the Potjans-Diesmann model (84) describes a cortical microcircuit containing two cell types (excitatory and inhibitory) distributed in four layers, representing the cortical network below a surface of 1 mm^2. This model has been used in a huge number of studies, also adapting it to specific scientific purposes (for example using different classes of neurons models). Another example is the Virtual Brain (94), a neuroinformatic platform based on the Jirsa model, that tries to simulate the brain organization on the macroscopic level of detail. This tool is based on the idea of taking advantage of available functional and structural brain data generated by imaging techniques such as MRI, functional MRI and trans-cranial magnetic stimulation. Virtual Brain applications span from the research study (for instance allowing the monitoring of the dynamics of communications between different brain regions) to more concrete medical applications, also at the scale of the human brain.

The purpose of my work is much more limited of those of Virtual Brain, both in terms of area and complexity of rhythms. Nevertheless it was not less challenging, because the transitions between SWA and QW like those naturally expressed during the wake-sleep cycle can be considered a demanding task for parallel simulation of large multi-scale models of the brain. During SWA, part of the network is synchronously active with a huge rate of exchanged spikes while the rest of the system is almost silent. On the other hand, during QW an asynchronous state involves homogeneously the whole network. This poses a question whether computational power and inter-process communication load of simulation platforms are fully exploited in both these brain states.

To address this issue, we run a bunch of simulations of spiking neuron networks organized in two-dimensional grids of modules, each aiming at modelling a cortical

column, including up to 50G synapses connecting 46M point-like neurons (Leaky Integrate and Fire with Spike Frequency Adaptation) distributed over a large set of MPI processes. Just changing a few parameters in the neuron description (mainly the parameter that express the *fatigue* or the *adaptation* of the neuron), slow oscillation and transition to the awake-like state can be modulated.

Distributed simulations have been executed on the NEST platform both under SWA- and QW-like brain states. For comparison and validation, the same scaling analyses have been performed on the proprietary mixed time and event driven DPSNN engine (Distributed Simulator of Plastic Spiking Neural Networks) running the same large-scale spiking neuron network models (79).

The use of DPSNN allows for faster simulations of the columnar model when compared with the simulations of the same model with the most flexible and general-purpose NEST engine. There is a widely felt need for versatile, general-purpose neural simulators that offer a user-friendly interface for designing complex numerical experiments and provide the user with a wide set of models of proven scientific value. This boosted a number of initiatives (notably the NEST initiative, now central to the European Human Brain Project). However, such flexibility comes at a price. Performance-oriented engines, missing all the layers required for offering user generality and flexibility, contain the bare minimum code. In the case of DPSNN, this resulted in higher simulation speed, reduced memory footprint, and diminished initialization times, es extensively reported in Chapter 3 (see also Pastorelli et al., 79).

As a last point, let me address the acceleration obtained with the use of GPU instead of simple CPU. Over the past decade there has been a growing interest in the development of parallel hardware systems for simulating large-scale networks of spiking neurons. Compared to other highly-parallel systems, GPU-accelerated solutions have the advantage of a relatively low cost and a great versatility, thanks also to the possibility of using the CUDA-C/C++ programming languages. In this context has been developed NeuronGPU (see Golosio et al. 38 and 39), a GPU

library for large-scale simulations of spiking neural network models, written in the C++ and CUDA-C++ programming languages, based on a novel spike-delivery algorithm. This library includes simple LIF (leaky-integrate-and-fire) neuron models as well as several AdEx (adaptive-exponential-integrate-and-fire) neuron models with current or conductance based synapses, user definable models and different devices.

1.5 Awake-sleep cycles in incremental learning tasks

This section summarizes the scientific question of my thesis activity: the key aspect of the role of sleep during learning. Sleep is essential in all animal species, and it is believed to play a crucial role in memory consolidation as described also in Walker ans Stickgold (116), or in Jadhav et al. (51), in the creation of novel associations, as well as in the preparation of tasks expected during the next awake periods. Among the multiple effects of sleep on the brain and body, we focus here on the consolidation of learned information, as studied also from Buzsáki in the past years (14). Homeostatic processes could normalize the representation of memories and optimize the energetic working point of the system by recalibrating synaptic weights (Tononi and Cirelli 110) and firing rates (Watson et al. 118). Specifically, Watson et al. show that fast-firing pyramidal neurons decrease their firing rates over sleep, whereas slow-firing neurons increase their rates, resulting in a narrower population firing rate distribution after sleep. Also, sleep should be able to select memories for association, promoting higher performance during the next awake phases (see Smulders et al. 100). Indeed, in Capone et al. (15) we demonstrate the beneficial effects of sleep-wake phases involving homeostatic and associative processes in a visual classification task.

In this work, we show the interesting effects of deep-sleep-like slow oscillation activity on a simplified thalamo-cortical model which is trained to encode, retrieve and classify images of handwritten digits. During slow oscillations, spike-timing-dependent-plasticity (STDP) produces a differential homeostatic process. It is

characterized by both a specific unsupervised enhancement of connections among groups of neurons associated to instances of the same class (digit) and a simultaneous down-regulation of stronger synapses created by the training. This hierarchical organization of post-sleep internal representations favours higher performances in retrieval and classification tasks. The mechanism is based on the organizing principle of the cortex described by Larkum (58): the interaction between top-down cortico-thalamic predictions and bottom-up thalamo-cortical projections during deep-sleep-like slow oscillations. Indeed, when learned patterns are replayed during sleep, cortico-thalamo-cortical connections favour the activation of other neurons coding for similar thalamic inputs, promoting their association.

After this first promising work, the activity proceeded on the same line with the development of a refined thalamo-cortical model. Increasing experimental evidence is mounting for both the role played by the combination of bottom-up (perceptual) and top-down/lateral (contextual) signals and for the beneficial effects of sleep as key components of many high-level cognitive functions in the brain. In the following, I give an overview of some aspects, driven from experimental observations, that we have taken as fundamental building blocks for the construction of this model.

It is known that the cortex follows a hierarchical structure (see Barone et al. 4); starting from this, Larkum et al. (58) propose an associative mechanism built-in at a cellular level into the pyramidal neuron, exploiting the cortical architectural organization. Long-range connectivity in the cortex follows the basic rule that sensory input (i.e., the feed-forward stream) terminates in the middle cortical layers, whereas information from other parts of the cortex (i.e., the feedback stream) mainly projects to the outer layers. This also applies to projections from the thalamus, a structure that serves both as a gateway for feed-forward sensory information to the cortex and a hub for feedback interactions between cortical regions. Indeed, only 10% of the synaptic feedback inputs to the apical tuft come from nearby neurons, and the missing 90% arise from long-range feedback connections. This feedback information stream is vitally important for cognition and conscious perception, and this picture

leads to the suggestion that the cortex operates via an interaction between feed-forward and feedback information. Larkum et al. highlight that, counter-intuitively, distal feedback input to the tuft dendrite could dominate the input/output function of the cell: short high-frequency bursts would be produced on a combination of distal and basal input. As a consequence, although small (under-threshold) signals contribute only to their respective spike initiation zones, the fact that input has reached the threshold in one zone is quickly signaled to other zones. This provides the possibility for a contextual prediction: the activity in the apical tuft of the cell can lower the activity threshold driven by the basal region, the target of the specific nuclei in the thalamus that projects there the perceptual and feed-forward streams. In summary, this mechanism is ideally suited to associating feed-forward and feedback cortical pathways. Thus, they propose a conceptual interpretation of these biological pieces of evidence: the feedback signal aims at predicting whether a particular pyramidal neuron could or should be firing. Moreover, any neuron can fire only if it receives enough feed-forward input. Resulting from this interpretation, the internal representation of the world by the brain can be matched at every level with ongoing external evidence via a cellular mechanism, allowing the cortex to perform the same operation with massively parallel processing power.

Inspired by those biological mechanisms, we exploited the combination of context and perception in a thalamo-cortical model based on a soft winner-take-all circuit of excitatory and inhibitory spiking neurons. After calibrating this model to express awake and deep-sleep states with features comparable with biological measures, we demonstrate the model capability of fast incremental learning from few examples, its resilience when proposed with noisy perceptions and contextual signals, and an improvement in visual classification after sleep due to induced synaptic homeostasis and association of similar memories.

1.6 Preparing the analysis of complex cognitive tasks in monkeys

The brain exhibits capabilities of fast incremental learning from few noisy examples, as well as the ability to associate similar memories in autonomously created categories and to combine contextual hints with sensory perceptions. Together with sleep, these mechanisms are thought to be key components of many high-level cognitive functions, as demonstrated by the model above described.

At this point the idea was to study the same phenomena on experimental recordings, analyzing the traces recorded on monkeys to find hints of the interplay between sleep and the cognitive tasks. Unfortunately, this kind of recordings, reporting the effect of sleep and in particular of slow wave activity on learning and cognition, were not available for non-human primates at the moment of my thesis. Anyhow, to take confidence with experimental data recorded during cognitive tasks on monkey, I started my activity analyzing the data already available in the *The Motor Control and Cognition Lab*, mainly with the experimental recordings of transitive inference and countermanding tasks.

The complexity and multidimensionality of cortical activity indisputably stands out during learning. Neural representations are modulated in the brain behavioural and cognitive tasks (Mattia et al. 65). The idea behind this research work was to understand how such complexity change as a function of a specific behavioural context and as a function of the received stimulus. Often recent studies adopted dimensionality reduction to analyze neural populations and to extract features from them (see Cunningham et al. 23 and Gao et al. 31). With dimensionality reduction methods, data can be reduced to a low-dimensional space, allowing a faster and easier insight into the dynamic of neural correlates.

Following this new trend, I concentrated my attention on the study of the dimensionality variability in learning tasks. In some sense, the neuronal dimensionality has been taken as a measure of neuronal complexity along a behavioural cognitive task as in Gao et al. (31). In our case, for example, some indications could arise

from the characterization of this measure for different difficulty levels of the tasks, such as different symbolic distances in transitive inference (Brunamonti et al. 10) or stop signal and short delays in countermanding experiments (Mirabella et al. 69).

I performed all this kind of analysis relaying on the classical Principal Component Analysis (PCA) to obtain dimensionality reduction, after having take into consideration also other methodologies. In particular I considered the one proposed by Laio (29), and I conducted a comparison between these two methods, as reported in section 5.4, concluding that PCA, with the opportune caveats, still remain the most reliable method.

1.7 Structure of the thesis

In Chapter 2 I briefly present an adaptable and reusable analysis pipeline to support quantitative statistical comparisons on the level of slow wave characteristics. The key objective of the pipeline is to bring together existing methods, standards, and tools in a flexible and modular manner in order to serve the requirements of a wide range of datasets and research questions. This analysis tool represent a promising item for the statistical study of identified waves, propagation velocities and local excitability, not only in mouse SWA but also on different rhythms in monkeys.

Simulation methodologies and their support on the reproducibility of the different rhythms of cerebral cortex are presented in Chapter 3. This chapter underline the difficulties and the problematic afforded to set-up high-resolution, wide-field simulations using different approaches (both mean-field and spiking), as well as the applied accelerations techniques.

Chapter 4 describes the thalamo-cortical models able to classify images extracted from the MNIST set of handwritten digits. The models are based on the organizing principle of the cerebral cortex proposed by Larkum (58). They show capabilities of fast incremental learning from few noisy examples and the ability to associate similar memories in autonomously-created categories, combining contextual hints

with sensory perceptions. Together with sleep, these mechanisms are thought to be key components of many high-level cognitive functions.

Finally, Chapter 5 reports the activity done on the analysis of recorded experimental data coming from two different kind of experiments conducted on the MCC_lab, the countermanding and the transitive inference. The focus of the activity is on the study of dimensionality and its variability over trials, in addition to some preliminary analysis describing a methodology prepared to be applied on data that will be hopefully available in the future, accounting of the sleep phase and its interaction with behavioural tasks.

1.8 Personal contribution

These thesis reports results related to both the analysis of neural correlates on non-human primates performing behavioural tasks recorded in the *The Motor Control and Cognition Lab* and cortical simulation and data analysis carried out as participant in the European project *Human Brain Project* (HBP) in the framework of the HBP APE Research Group of INFN Sezione di Roma.

In particular I contributed to the developments of models reproducing slow wave activity and awake-like state, and on the high-resolution large-scale simulation engine able to reproduce this kind of cerebral activity as described in Chapter 3 and published in 79 and 78. The analysis I performed on simulations for real-time execution was published in 99, while the acceleration studies using GPU platform in 38 and 39. The preliminary results of these publications was also presented as oral contribution at the *NEST Conference 2017 - A Forum for Users & Developers, 19-20 December2017, Jülich, Germany.*

For what concerns models, in addition to the one above mentioned that is able to reproduce the slow waves traveling on the whole hemisphere of a mouse at biological neural and synaptic resolution, I also participated in the development of the thalamo-cortical models described in Chapter 4 that demonstrate the deep-sleep effects on learning and cognition and published in 15 and 37. These models was preliminary

presented as a poster at the *Human Brain Project Summit 2018, 15-18 October 2018, Maastricht*, together with the simulation results above described, and published in 15 and 37.

In addition, I also contributed to the development of the analysis tools described in Chapter 2, an aim toward the FAIR principles that in these last years are animating the scientific community and presented with my contribution at the *Bernstein Conference 2020* with the poster 43. I also participate to present the preliminary results obtained with the developed analysis pipeline applied to experimental and simulated data to the *SP3-based meeting: Scientific highlights and collaborations across HBP - 2018*.

The activity reported in section 5 constitute my personal contribution to the analysis of already recorded data on monkeys, in particular during transitive inference tasks. The data acquisition was done in the past years in *The Motor Control and Cognition Lab* (MCC_lab) directed by Prof. Stefano Ferraina at Sapienza University.

Chapter 2

Analysis of SWA in neural correlate through an innovative tool

The phenomenology of slow wave activity (SWA) covers a crucial role in this work. Therefore, the analysis tools able to extract the salient features of this phenomenon constitute a prerequisite, in order to study the effects of SWA on learning and cognition.

Several tools are already available in the neuroscience community, but quite often they are dedicated to analyzing specific input data, and quite rarely provide a complete workflow adaptable to different cases.

More in general, in the last years, there is a strong tendency among scientific community, to move toward a *"FAIR"* approach to science. FAIR is an acronym that stands for Findability, Accessibility, Interoperability, and Reuse of scientific data and their management. These principles have the aim to facilitate knowledge discovery by assisting humans and machines enhancing their ability to automatically find and reuse scientific data, in addition to supporting data integration and analysis throughout associated algorithms and workflows.

In this context, the team of work I am involved in, decided to set up a dedicated tool inspired to the basic FAIR principles (see Gutzen et al. 43). The tool is manly focused on the study of slow waves, but at the same time it could be general enough

to be reused in different context, foreseeing a modularity and a flexibility for what concern the input data that could be analyzed, the flux of analysis to be followed, the features that need to be extracted.

The tool has been applied to the data we had in our Lab, in particular to some ECoG datasets provided by IDIBAPS (Institut d'Investigacions Biomèdiques Agustí Pi i Sunyer), and to optical data constituted by calcium imaging datasets, acquired from mice, by LENS (European Laboratory for Non-Linear Spectroscopy).

To complement this work, we also set up theoretical models expressing slow waves, that are able to demonstrate that such waves play a crucial role on the learning process (see Chapter 4). Our analysis tool, thanks to its flexibility and reusability, has been also applied to synthetic data obtained from simulations, allowing us better calibration of detailed models of brain dynamics and function.

2.1 An analysis pipeline for slow waves activity

The Slow Waves Analysis Pipeline (SWAP) is a reusable and adaptable tool, able to analyze diverse datasets of slow rhythms in the cerebral cortex. The main investigation focus is the analyses of the features and cognitive effects of slow wave activity (SWA) observed in the cerebral cortex during deep sleep and anesthesia. More in detail, SWAP can be used to extracts key spatial-temporal characteristics from experimental or simulated slow waves. SWA can be acquired with multiple experimental methodologies (e.g. using micro-ECoG arrays or wide-field Calcium imaging techniques) at local and multi-areal spatial resolution; when applied to simulation results similar features should be extracted, to enable a quantitative comparison between synthetic and experimental data, fostering a better calibration of simulations.

For this reason, a common framework has been established for experimental and simulated data, addressing the efforts towards the design and implementation of a set of analysis tools that integrate and expand already existing more general tools widely used in the neuroscience community.

Finally, there is experimental evidence that wave-like propagation can be observed across several frequency bands (from delta to gamma) in electro-physiological data (see for instance Wu et al. 125, Sato et al. 95 and Muller et al. 73). Thus, despite the scientific question that triggered this work was the SWA, such that the pipeline in its current formulation has been designed and developed specifically for addressing slow rhythms in the cortex, we are evaluating a further generalization of the tools, in the direction of a "WAP", Wave Analysis Pipeline, aiming at disentangling the different components that appear to be present when looking at the collective cerebral signals expressed across brain states.

2.1.1 Reproducibility and modularity

The pipeline is built to be adaptable to different input data types and different analysis methods to address various scientific questions, trying to follow the FAIR principles. For this reason, it is structured as a modular system executing the processing and the analysis in order to satisfy multiple requirements. At a very high level of description, the input data may differ in terms of spatial or temporal resolution, scale, signal type, etc. Disregarding if the data are obtained via EEG, electrode arrays, imaging techniques, or even simulations, the aim is that to converge towards a common description of the phenomenon of interest. Once reduced the input data to a common language, the pipeline can be used selecting a variety of specific analysis to derive the characterization metrics according to the phenomenon the user needs to analyze (see Figure 2.1)

The pipeline is structured into a series of sequential stages, where each stage is a self-consistent logical unit with a specific purpose, well defined inputs and outputs. Each stage is further segmented into blocks. A block defines a concrete action to be performed on the data, i.e., it implements a method. Similar to stages, blocks have specific inputs and outputs by which they are chained together, but in contrast to stages, block are modular and are not necessarily executed in a predefined sequence.

At the current level of implementation, the SWAP pipeline is based on five stages,

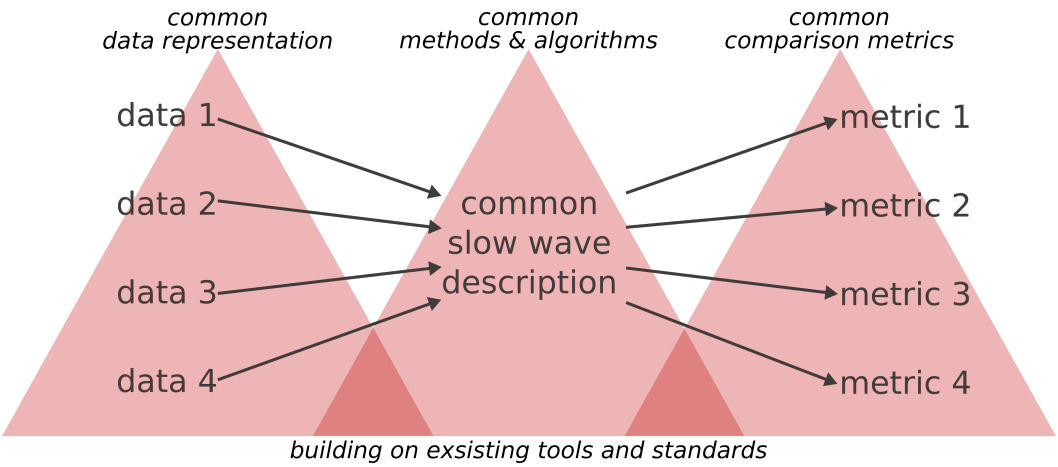

Figure 2.1. SWAP: a modular pipeline approach

each one with its specific purpose as described here in the following.

1. ***Data Entry***: the first stage is used to load datasets and, when available, metadata. The most important feature addressed by this stage is the conversion of the data into a common standardized format, the NEO data format (32). In addition, a check is done whether data and metadata conform to the pipeline requirements.

   The main block implemented in this stage is the ***data_curation***. The input data, coming from different kind of recordings, such as EcOG, optical acquisition, or simulations, are converted in a common format, the NEO data format. Since this point on, all the data result uniform to the same format.

2. ***Processing***: during this second stage, data are prepared for the analysis. Here the selection of the blocks that will be executed in accordance with the data type and the analysis objectives is essential.

   Among the main blocks that can be selected, the user can chose to apply different type of filters (***frequency_filter***), normalization (***normalization***), spatial smoothing (***spatial_downsampling***), detrending (***detrending***) and the selection of the region of interest (ROI, ***roi_selection***).

3. ***Trigger Detection***: this stage implements a trigger for the detection of the

transition times from DOWN to UP states (up-ward transitions) and from UP to DOWN states (down-ward transitions), on the data already pre-processed during previous stages. The trigger strongly depends on the dataset, the processing and the detection method selected by the user.

Different methodologies can be used to detect the phase transitions, such as ***threshold***, ***hilbert_phase***, or ***minima***. Moreover, if needed, the ***remove_short_states*** block can be used to removes short UP and/or DOWN states.

4. ***Wavefront Detection***: at this point, the data have been reduced to a common description of waves. The detection methods can operate on trigger times, grouping them into individual wavefronts, agnostic about the type of the original data.

   In this stage two blocks are mainly notable: the ***clustering*** that groups triggered events by spatial and temporal distance, and the ***optical_flow*** that calculates vector velocity field.

5. ***Wave Characterization***: this is the final stage used to calculates the characteristic measures related to a user-defined specific task or to a given scientific question. The result can be for instance used for the comparison between different datasets or between methods of analysis, for model validation, method benchmarking, and so on.

   Some of the blocks already implemented for the analysis of slow waves are the ***direction*** and the ***velocity_planar*** respectively for the interpolation of the directions of planar waves and for the interpolation of waves propagation velocity.

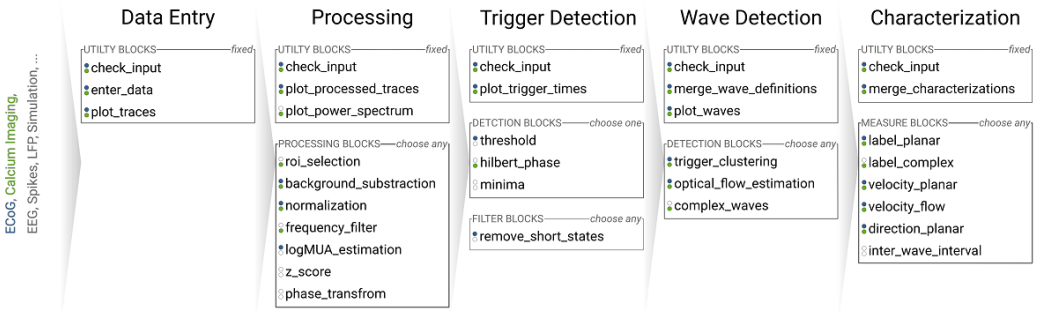

Figure 2.2. Stages of the SWAP pipeline: five successive stages containing each a set of block performing specific analysis on the input data

2.2 An analysis pipeline extracting local oscillation features

The Slow Oscillation Analysis Pipeline (SOAP) has been developed to investigate the local oscillations of the multi-unit activity signal between Up and Down states, i.e. without including the information on spatio-temporal coherence among the recording sites and thus the signal propagation. This tool, strictly related to SWAP, focuses on the characteristic of SO at each pixel or electrode, producing an output that complements and enriches the information of SWAP.

Basically, SOAP results can be obtained re-using the first three stages of the SWAP pipeline, with this also demonstrating the modularity and reusability of SWAP. In particular the blocks of the ***Trigger Detection*** stage can be used to this purpose. Depending on the user-specific analysis, the blocks of ***threshold***, ***hilbert_phase***, or ***minima*** can be used as alternative methodologies.

The idea consists in understanding where the wave is at a specific instant on the selected point, disregarding the spatial component of the wave. In some sense, it can be considered as a local analysis of a macro-phenomena, where the spatial components are not taken into account. Obviously, if the area under study is little enough that its behaviour can be considered synchronous (all the neurons of the system are simultaneously in UP or in DOWN state), also larger portions of cortex can be considered, as in optical imaging analysis where the whole voxel is analyzed, or when the distance of electrical micro-arrays is so little that the area have the

same behaviour. In this case the scientific question is slightly different from the complete analysis of SWA done by the SWAP pipeline, but its reconfigurability and modularity make it perfectly suitable to be re-used also as with a different goal, the local study of Slow Oscillations.

2.3 Characterizing the spatio-temporal feature of the waves

The phenomena of cortical Slow Waves is observed in all mammals in a state of deep sleep or under anaesthesia (Sanchez-Vives, et al 92, Nghiem et al.,76) and is characterized by a large-scale collective activation of groups of neurons, with a characteristic undulatory space-time pattern. The propagation throughout the cortex modulates the spiking frequency of the underlying neurons populations. This results in a transition of the involved neurons from a state of low spiking activity (down-state), to a state of more intense activity (up-state) while waves travel on the cortex.

With the last stage of the pipeline, traveling waves already identified in the previous stages are analyzed according to the scientific question of the user. In this version of the SWAP pipeline specific characteristics are extracted, such as planarity, velocity and direction of the wave. But thanks to its intrinsic modularity, the tool can be expanded to address the identification of additional features that can be gathered from complex wave patterns, such as phase, source or sink location, excitability, just as a few examples.

In the following section an example of this feature extraction is reported, in the specific starting from optical imagining in half hemisphere of cortical mouse.

2.3.1 Analysis of high-resolution whole cortical hemisphere optical GECI recordings of SWA in mouse

Up to few years ago, electroencephalography (EEG) and electrocorticography (ECoG), as well as more invasive techniques like extracellular and intracellular

recording both in vivo and in vitro, have been the standard methods for collecting slow wave data (anchez-Vives et al. 90, Steriade et al. 106, Capone et al. 16). Even though electrophysiological recordings have been largely used to characterize brain states, they are limited in spatial resolution and cannot target specific neuronal populations. Recently, large-scale optical imaging techniques coupled with functional indicators overcame these restrictions and will give new important insight on the SW phenomenon. Among them, GECIs represents an interesting technique enabling the visualization of fluctuations in calcium concentration, which represents an indirect reporter of neuronal spiking activity (see Grienberger et at. 42).

In detail, cortex-wide mesoscopic optical imaging, combined with fluorescent indicators of activity, provided new insight on the spatio-temporal propagation pattern of the brain activity. The combination of wide-field fluorescence microscopy and a transgenic mouse model expressing a calcium indicator (GCaMP6f) in excitatory neurons is becoming an innovative methodology for examining the neuronal activity in anesthetized and awakening mice, as described in detail in Resta et al. (86).

Among the first results obtained applying the analysis pipeline to the calcium imaging data, there is the computation of the waves velocity, the firing rate distribution, the preferential propagation directions and the waves origin points on the half hemisphere of the mouse cortex. In Figure 2.3 the main results of the analysis are reported.

Panel A represents the histogram of the estimated average velocity of the traveling waves on the whole mouse hemisphere. To obtain the average speed of each wave, the velocity is calculated on different points. Knowing the passage time function of the wave T(x,y), representing the time at which a the (x,y) position has been reached by the wave during its propagation, the speed of a wave on a specific point can be defined as the inverse of the module of the function's gradient:

$$v(x, y) = \frac{1}{\nabla T(x, y)} \tag{2.1}$$

Computing the gradient and taking its module, we obtain:

$$v(x,y) = \frac{1}{\sqrt{(\frac{\partial T(x,y)}{\partial x})^2 + (\frac{\partial T(x,y)}{\partial y})^2}} \tag{2.2}$$

Panel B reports the firing rate distribution calculated on the whole mouse hemisphere. On panel C, propagation directions are reported in polar coordinate. The result is obtained starting from the same data of wave velocity, but this time the velocity of each single pixel is mediated trough a Gaussian filter $w_{(\mu,\sigma)}$ centered in the pixel involved in the transition and with $\sigma = 2$ (see 2.3). In this way, noise is removed from this velocity estimation with a weight decreasing according to the Gaussian sigma (2 pixels in the example).

$$\theta(x,y) = tan^{-1}\left(\frac{\langle v_y \rangle_{w_{(y,\sigma)}}}{\langle v_x \rangle_{w_{(x,\sigma)}}}\right) \tag{2.3}$$

Finally, panel D represent the main origin points of slow waves. More the point tends to yellow, more it is source of slow waves. The results reported are mediated on seven acquisitions of about 40 seconds each, on a mouse anesthetized with isoflurane.

2.4 Integration in the EBRAINS platform of the Human Brain Project

The analysis pipeline discussed in this chapter is part of a work done in cooperation with some partners of the Human Brain Project (HBP). In particular, the development has been curated by a team of my lab (INFN) and a team of the Jülich Research Centre, in collaborations with other HBP partners mainly responsible for providing the data to be analyzed (IDIBAPS, LENS). The final goal is that to make the tool publicly available to the neuroscience community, in order to facilitate the investigation of SWA, and in the future also of other brain rhythms, on a variety of different kind of recorded and synthetic input data. The distribution of SWAP will be done throughout the EBRAINS platform, a shared digital brain research infrastructure for the European Community, that is being set up in the framework

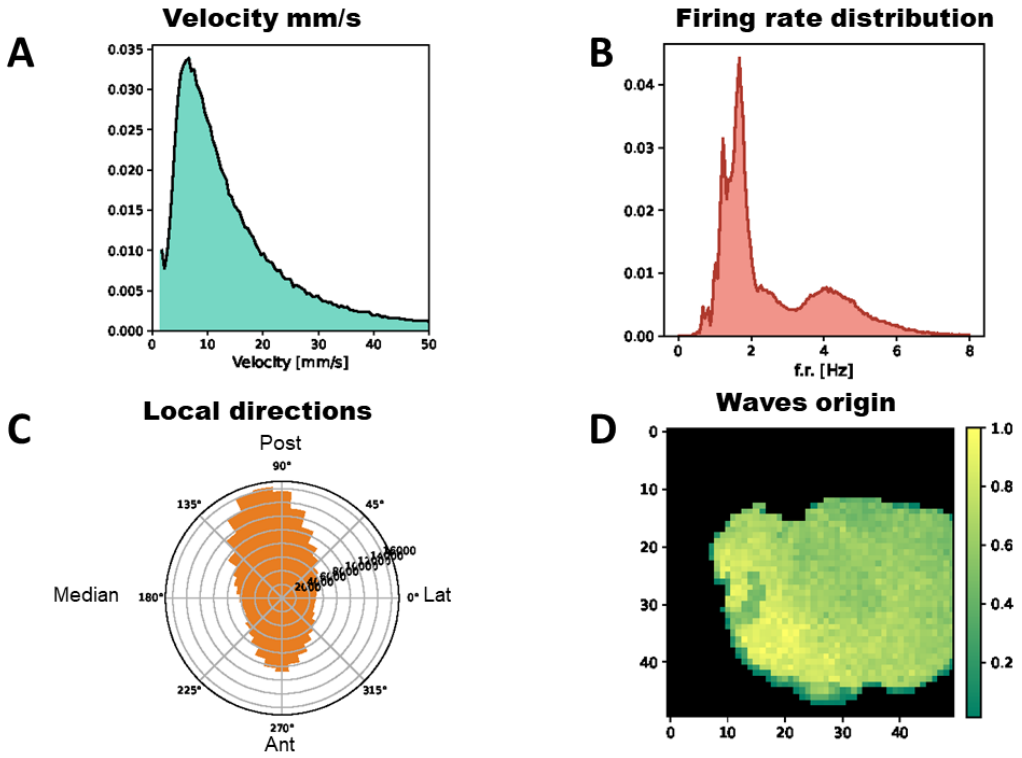

Figure 2.3. Waves analysis. Panel A: histogram of all the velocity points calculate on the whole mouse hemisphere. Panel B: firing rate distribution. Panel C: preferential propagation directions of the waves calculated taking into account the contribution of surrounding pixels. Panel D: representation of the main points of waves origin.

of the EU-funded Human Brain Project.

EBRAINS is a platform providing tools and services which can be used to address challenges in brain research and brain-inspired technology development. Its components are designed to facilitate researchers activity and, above all, with their contribution. The tools assist scientists to collect, analyze, share, and integrate brain data, and to perform modelling and simulation of brain function, to accelerate the effort to understand human brain function and disease.

Accessing the EBRAINS platform, a set of data, atlases, models and software are made available to the research community following the FAIR principles that are nowadays beginning more and more relevant. A large part of EBRAINS infrastructure is constituted by the Knowledge Graph, an online solutions based on these FAIR principles to facilitate sharing and access of research data, computa-

tional models and software. A preliminary demonstration to the general public of the application of the SWAP pipeline (March 2020 release) can be found at https://wiki.ebrains.eu/bin/view/Collabs/slow-wave-analysis-pipeline, which also contains pointers to the GitHub repository.

Chapter 3

High-resolution wide-field simulations of SWA

The use of models for the study of physical systems is an acquired and well-established methodology by which the scientist aim to demonstrate the behaviour of complex systems. In a closed-loop fashion, that can be covered in both directions, scientists start from theory, derive the equations describing the fundamental behaviour of the system, then build a model to confirm the supposed theory and compare the results with experimental data. On the contrary, starting from the experimental data, models can be built deriving parameters from observations, then a theory can be formulated that describe the observed experimental phenomena (see Figure 3.1 for a graphical explanation). Anyhow, the built of models is one of the central points of this methodology, already extensively used since many years in several fields of study, from engineering to meteorology, from automotive to aerospace.

The human brain is another complex system that can benefit from modelling and simulation: the models can help scientists to better understand the mechanism underlying specific problems, to test hypothesis and to find solutions sometimes hardly manageable directly using experimental results.

The trend in building models of neuronal activity is strongly emerging in the last two decades, as also demonstrated by the wide variety of neural simulator today available (NEST 34, 52, NEURON 47, 17, GENESIS 124, BRIAN 41, 107), Virtual

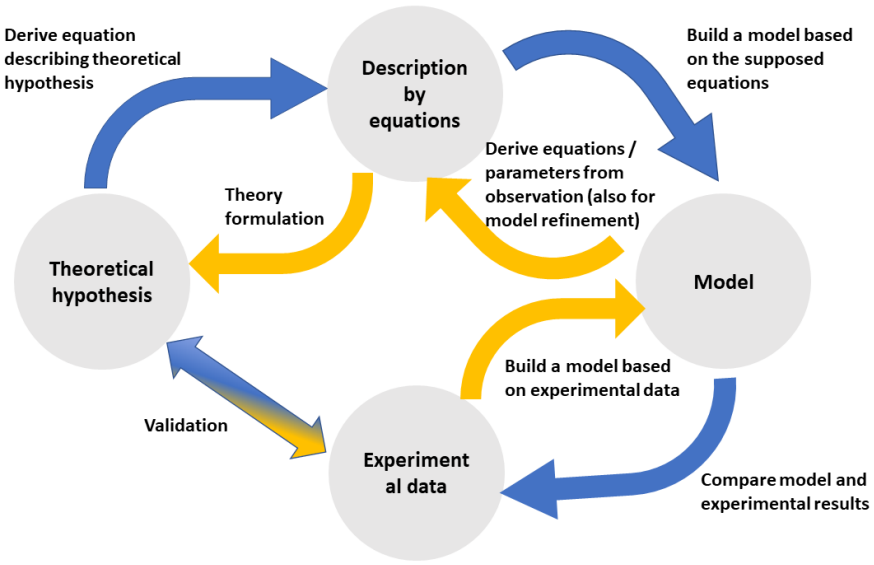

Figure 3.1. Loop from theory to experiments and back

Brain 94), or even specific accelerated systems (Modha et al. 70, Izhikevich et al. 50, Nageswaran et al. 75), or neuromorphic platforms (SpiNNaker 30, BrainScaleS 96, TrueNorth 67).

Also in my work, in order to answer the scientific question about the effect of SWA on learning with the analysis of recorded data, not yet available at the beginning of the activity, I started working on models and on different kind of simulators that could be in any case useful to replicate the phenomena.

The goal was that to simulate the whole cortical mouse hemisphere activity starting from connectivity and other parameters inferred from wide-field optical imaging data of SWA, provided by LENS, to obtain a model with biological neural and synaptic resolution. The result obtained is depicted in Figure 3.2, that represent a short segment of a movie in which a traveling wave crosses the mouse cortex. The whole movie can be downloaded from GitHub at the address https://github.com/epastorelli/MouseHemisphereMovie.

The following sections of this chapter describe in detail the model of the whole mouse hemisphere here briefly mentioned, and the results of the SWAP analysis pipeline applied on the simulation outputs of this model. In order to reach this

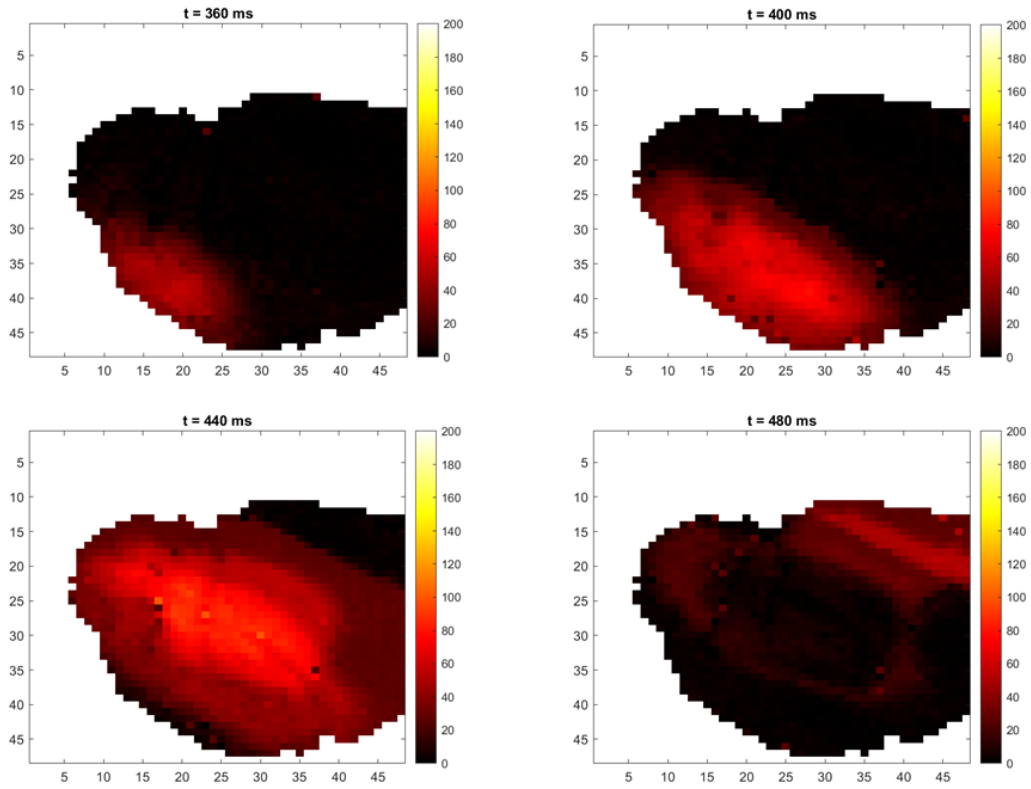

Figure 3.2. Spiking simulation of SWA: four snapshots taken at a distance of 40 ms that clearly show the slow wave traveling across the cortical surface of the whole mouse cortex

objective, several steps was required. First of all, I worked on the reproduction of large-scale non-plastic simulations, I mean disregarding the synaptic plasticity, that will be introduced in successive steps to study learning and cognition with the model. Plastic large-scale high-resolution simulation are an ambitious goal that can not be reached before having a solid background in being able to support simulations with millions of neurons and billions of synapses. Indeed, plasticity requires a very high computational cost, that can be roughly estimated in a 4x of the original cost. Also, the optimization techniques described in the last sections of this chapter are focused on obtaining effective instruments to reach this goal.

Part of the material reported in this chapter has been extracted from the publication of Pastorelli et al. (79) published in the journal Frontiers in Systems Neuroscience. All the pictures extracted by my previous publications are systemati-

cally indicated.

3.1 Spiking simulation at hemisphere scale at biological neural and synaptic resolution

Thanks to the simulation methodology to which I contributed during these years, I was able to reproduce a large-scale high-resolution spiking simulation of the model of the mouse half hemisphere at biological neural and synaptic density. The final aim is that to provide accurate simulations of SWA and AW-like states that could be used by the scientific community for the study of cortical rhythms and their interaction with consciousness and learning.

Spiking simulations methodology, described in detail in the following sections of this chapter, were refined and calibrated in terms of input parameters and biological plausibility. This was possible using a specific procedure for the extraction of the required parameters from experimental recordings, using a likelihood maximization technique that progressively increases the match between experimental data and the results of mean-field simulations.

The model is based on a columnar network of Adaptive exponential integrate-and-fire model neurons (Adex) described by Brette and Gerstner (7) composed of two populations of neurons, excitatory and inhibitory, respecting the usual proportion of 4:1 (exc:inh). The functional connectivity and other parameters for simulation of SWA activity of a whole hemisphere of the wild type mouse have been inferred from optical GECI calcium imaging GCaMP-6f recording (field of view: 4x5 mm^2, pixel resolution: 50x50 μm^2, temporal resolution: 40ms) provided by the HBP partner LENS. See for details Celotto et al. (18). Mainly three parameters are inferred from simulated data and injected into a spiking simulation:

1. lateral connectivity
2. mean external firing rate
3. neuronal adaptation parameter

In the right panel of Figure 3.3 an example of lateral connectivity is reported for eleven neurons. As can be seen, not all the neurons present the same connectivity, also visible from the distribution of the intensity around them. The left panel of the same figure reports instead a snapshot of a wave taken at 480 ms after the simulation starts. A more explicit example of the resulting simulation can be seen in Figure 3.2, where four snapshot of a simulation are reported, showing the traveling wave on the mouse hemisphere at a distance of 40 ms each from another, while the movie downloadable from https://github.com/epastorelli/MouseHemisphereMovie represent 10 second of simulated slow wave activity.

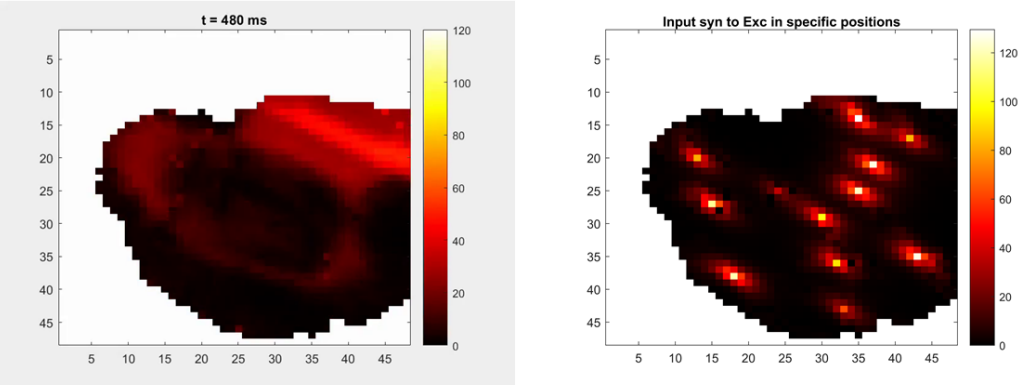

Figure 3.3. Spiking simulation of SWA. Left panel: instantaneous snapshot of wave. Right panel: example of inferred connectivity for eleven neurons scattered on the mouse cortical surface

3.2 Analysis of simulation outputs and comparison with experimental data

The output of the above described simulation of whole mouse hemisphere have been analyzed with the SWAP analysis pipeline. Figure 3.4 reports the main results of this analysis.

In detail, panel A shows the histogram of the estimated velocity calculated point by point on the whole mouse hemisphere using equation 2.2. The histogram has the same qualitative aspect of the one produced analyzing the calcium imaging data

reported in Figure 2.3, even if the exact match from a quantitative point of view has not been reached, yet. Some further calibration in simulation parameters is still needed.

Another interesting output of the analysis of simulated data is reported in Panel B and describes the firing rate derived from the inter-spike intervals calculated on all waves and in all cortical points. The range of frequency is qualitatively compatible with the biological one (delta band).

Panel C describes the propagation directions of the waves in polar coordinate calculated according to equation 2.3. As can be seen, most of the waves propagate diagonally in an antero-posterior way, toward a north-east direction. With less frequency, in the simulation there are also waves propagating in the opposite direction (postero-anterior waves, from north-east to south-west) as often happens in biology (Massimini et al. 61, Murphy et al. 74). Comparing this plot with the one available for experimental data of calcium imaging (see panel right of Figure 2.3), it is noteworthy the similarity between experimental and simulated results.

Panel D reports the main origin points of all the waves during the whole simulation. As can be seen from the picture, the bottom-left side of the cortex is the one with the higher number of waves generated (antero-posterior waves). In minor number, the other source point can be individuated in the higher part of the mouse cortex (postero-anterior waves).

As a preliminary comparison between the analysis of simulated and experimental data, it is evident that the simulation parameters still need calibration, but the macroscopic results are already in qualitative agreement with those coming from the experiments. In addition, the simulation results here analyzed are relative to a single simulation of ten seconds length. For a fair comparison between Figures 3.4 (simulation) and 2.3 (experiment) more run should be done with a longer simulation time. The activity is already proceeding in this direction, with the aim of finding the best tuning of all the parameters.

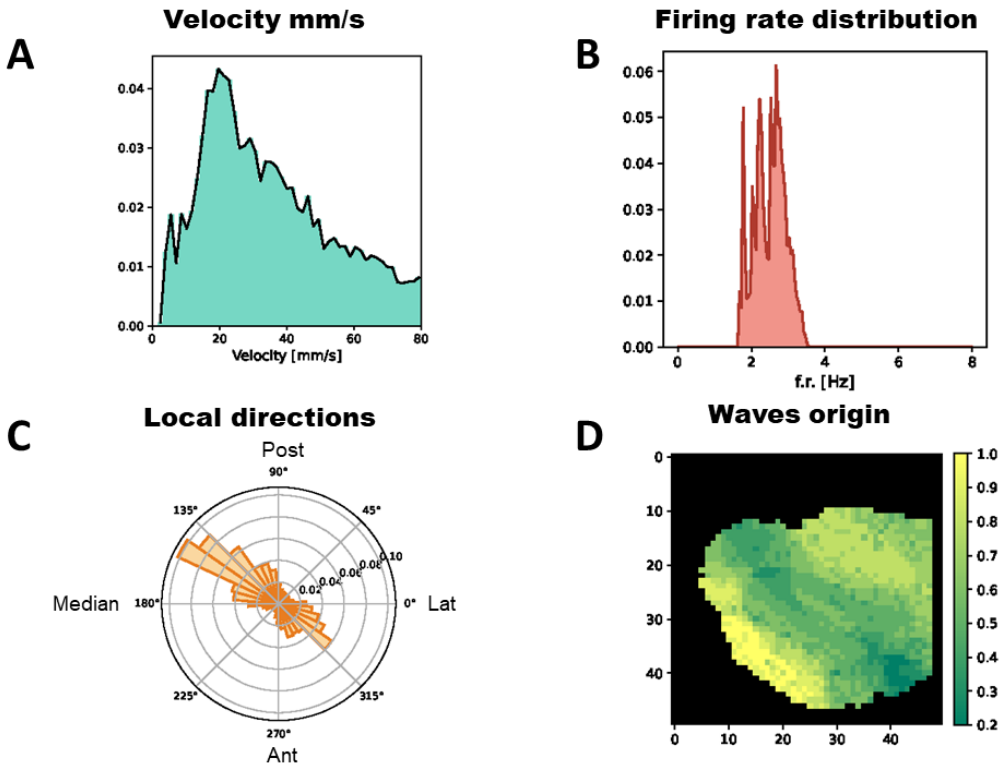

Figure 3.4. Analysis on spiking simulation output of the mouse whole hemisphere. Panel A: histogram of all the velocity points calculated on the whole mouse hemisphere. Panel B: firing rate distribution. Panel C: preferential propagation directions of the waves. Panel D: wave origin points on the mouse cortex

3.3 Anatomical priors: exponential decay of lateral connectivity in a columnar architecture

A simulation including a few tens of billions of synapses is what is required to simulate the activity of one cm^2 of cortex at biological resolution (e.g. 54K neuron/mm^2 and about 5K synapses per neuron in the rat neocortex area as described by Schnepel et al. in 97). The capability to scale a problem up to such a size allows simulating an entire cortical area. In my work I focuses on the possibility to build models able to implement connectivities pointed out in recent studies reporting about long range intra-areal lateral connectivity in many different cerebral areas, from cat primary visual cortex (Stepanyants et al. 103), to rat neocortex (Schnepel et al. 97 and Boucsein et al. 6), just as examples. For instance, in rat neocortex, the impact of lateral connectivity on the pyramidal cells in layer 2/3 and

layer 6a, results in ∼75% of incoming remote synapses to neurons of these layers.

Previous studies considered intra-areal synaptic connections dominated by local connectivity: e.g. Schuz in 98 estimated at least 55% the fraction of local synapses, reaching also a ratio of 75%. Such shorter-range lateral connectivity has often been modeled with a distance dependent Gaussian decay (see Potjans and Diesmann, 84) $B \cdot exp(\frac{-r^2}{2\sigma^2})$, where r stands again for distance between neurons, σ^2 is the variance that determines the lateral range and B fixes the total number of projections.

Longer-range intra-areal connectivity can be modeled by a distance-dependent exponential decay of the probability of synaptic connections between pairs of neurons: i.e. $A \cdot exp(\frac{-r}{\lambda})$, where r stands for the distance between neurons, λ is the exponential decay constant and A is a normalization factor that fixes the total number of lateral connections. Decay constants in the range of several hundred microns are required to match experimental results.

At the beginning of my activity, I have taken into account both kind of connectivity. In one of my works (78) I deeply studied the intrinsic technical difficulties in the implementation of the two modalities. Exponentially decaying lateral connectivity (longer-range) have been compared to a Gaussian connectivity decay (shorter-range), analyzing the scaling and the memory usage of the used simulator.

At the end of these studies, the model, as well as the spiking large-scale custom simulator adopted to accommodate it, where both optimized to obtain the maximum efficiency specifically in case of the selected long-range lateral connectivity, being this the most prominent feature according to the more recent studies, as reported above. To complete the overview of the used models, also mean-field simulations have been done inside the team as described in the following section 3.4 to accelerate the process of study of SWA and above all to define a strategy for the refinement of spiking neural network models parameters (see 3.1).

3.4 Mean-field columnar simulation of whole cortical hemisphere in mouse

For the study of the slow waves phenomenon and the comparison with recorded data, one of the faster approach is that to set up mean-field simulation able to reproduce the whole cortical mouse hemisphere.

The brain is a highly complex system composed of millions of neurons interconnected by billions of synapses and displaying extremely complicated firing patterns. Mean-field approaches aim at simplifying such dynamics by focusing of the collective activity and on the correlation patterns of neurons grouped together, in this way bridging microscopic (single-neuron spikes) and macroscopic scales (instantaneous spike rate of a neuronal ensemble).

The model we presented in our work on Frontiers in System Neuroscience in 2019 (79) is able to reproduce both the slow wave synchronous and an awake-like asynchronous activities. A prerequisite was the elaboration of a dynamical mean-field description for the simulation of the two cortical states, formulated by one of the authors of our work (Cristiano Capone), that represents the starting point for the definition of the parameters for the spiking simulations.

The assumption is that the inputs received by the different neurons in a module are independent but are conditioned by the same (and possibly time-dependent) mean and variance of the input synaptic current, as described in Brunel and Hakim (11) and Mattia and Del Giudice (63).

The gain function ϕ_i for this kind of neurons was firstly found by Ricciardi in the 1977 (88) as the reciprocal of the first-passage time (FPT) of the membrane potential to the firing threshold, for an integrate-and-fire neuron with stationary white noise input current, in the diffusion approximation. In detail, the transfer function elaborated to fit our work was the following:

$$\phi_i(\mu, \sigma^2) = \frac{1}{\sigma^2} \int_{-\infty}^{V_\theta} dV \int_{max(V,V_r)}^{V_\theta} du\, e^{-\frac{1}{\tau_m \sigma^2} \int_V^u [f(v)+\mu\tau_m] dv} \tag{3.1}$$

where μ and σ are the mean and variance of the input current, V_θ is the voltage spiking threshold of the neuron, V_r is the voltage reset value after a spike, τ_m is the membrane time constant and $f(v) = -(v(t) - E_l)$, with E_l being the reversal potential.

Spike frequency adaptation is an important ingredient for the occurrence of slow oscillations in the mean-field dynamics, as demonstrated by the works of Gigante et al. (35) and Capone et al. (16).

The mean-field dynamics for the average activity $\nu_i, i = \{F, B, I\}$ is determined by the gain functions ϕ_i as follows:

$$(F,B): \begin{cases} \dot{\nu_i} &= \frac{\phi_i(\overrightarrow{\nu},\overrightarrow{c})-\nu_i}{\tau_E} \\ \dot{c_i} &= -\frac{c_i}{\tau_c} + \alpha_c \nu_i \end{cases} ; \quad (I): \ \dot{\nu_i} = \frac{\phi_i(\overrightarrow{\nu}) - \nu_i}{\tau_I} \tag{3.2}$$

where τ_E and τ_I are phenomenological time constants for excitatory and inhibitorory neurons respectively. The interplay between the recurrent excitation embodied in the gain function, and the activity-dependent self-inhibition, is the main driver of the alternation between a high-activity (Up) state and a low-activity (Down) state (see Mattia and Sanchez-Vives (64) and Capone et al. (16)).

In the mean-field description of the neural module, synaptic connectivity is chosen to match the total average synaptic input that the neurons would receive inside the multi-modular network. The fixed points of the dynamics expressed by equations (3.2) can be analyzed using standard techniques described by Strogatz (108). The nullclines of the system (where $\dot{\nu} = 0$ or $\dot{c} = 0$) cross at fixed points that can be predicted to be either stable or unstable. In the simulations described here, the strengths of recurrent synapses $J_{t,s}$ connecting source population s and target population t, and of external synapses $J_{t,ext}$, is randomly chosen from a Gaussian distribution with mean $J_{t,s}$ and variance $\Delta J_{t,s} = 0.25 \times J_{t,s}$.

We relied on mean-field analysis to identify neural parameters setting the network's modules in SW or AW dynamic regimes. Panel A of Figure 3.5 shows an

example of nullclines for the mean-field equations (3.2) of a system displaying SW. The black S-shaped line is the nullcline for the rate ν, while the red straight line is the one for the fatigue variable c (for details see Mattia and Sanchez-Vives (64) and Capone et al. (16)). The stable fixed point, at the intersection of the nullclines, has a low level of activity and is characterized by a small basin of attraction: the system can easily escape from it thanks to the noise, and it gets driven towards the upper branch of the ν nullcline from which, due to fatigue, it is attracted back to the fixed point, thus generating an oscillation (see panel B of Figure 3.5).

Network parameters can also be set in order to have an asynchronous state, mainly by setting a lower Foreground-to-Foreground (FF) synaptic efficacy, which generates a more linear ν nullcline close to the fixed point (see panel C of Figure 3.5). In this case the basin of attraction of the fixed point is larger, and oscillations do not occur, resulting in a stationary asynchronous state (see panel D of Figure 3.5), in which neural activity fluctuates around the mean-field fixed point.

3.5 Preparing the technology for large scale columnar spiking simulation

At the large scale, the neural dynamics of the cerebral cortex results from an interplay between local excitability and the pattern of synaptic connectivity. This interplay results in the propagation of neural activity. A case in point is the spontaneous onset and slow propagation of low-frequency activity waves during the deep stages of natural sleep or deep anesthesia (Destexhe and Contreras (27), Hobson el al. (48), Sanchez-Vives and Mattia (91), Reyes-Puerta et al. (87)).

The brain in deep sleep expresses slow oscillations of activity at the single neuron and local network levels which, at a macroscopic scale, appear to be synchronized in space and time as traveling waves (slow wave activity, SWA). The “dynamic simplicity” of SWA is increasingly being recognized as an ideal test bed for refining and calibrating network models composed of spiking neurons. Understanding the dynamical and architectural determinants of SWA serves as an experimentally

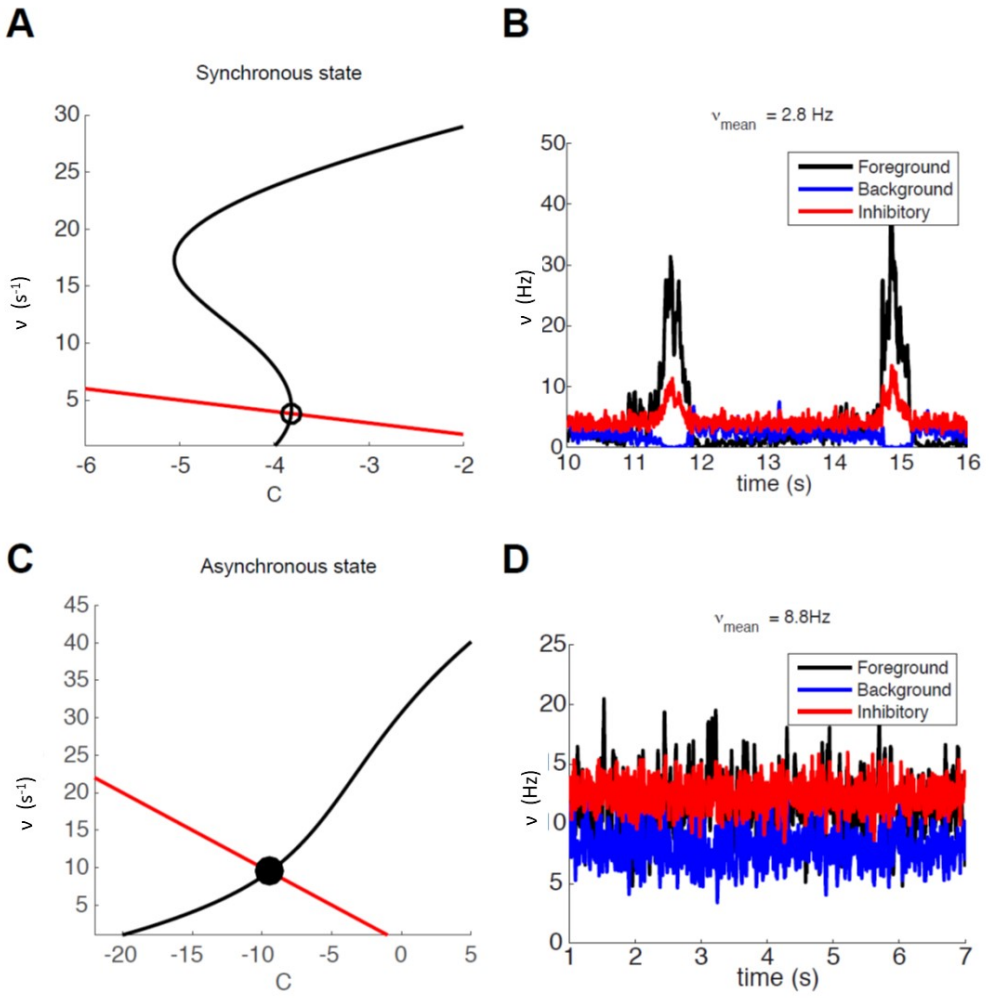

Figure 3.5. Dynamical representation of the SW and AW states. Panel A: phase space representation of mean-field analysis with a weakly stable fixed point. Panel B: firing-rate time course for an example module, for foreground, background, and inhibitory sub-populations (respectively in black, blue, and red) in the SW state. Panel C: phase space representation of mean-field analysis with a stable fixed point at a high level of activity. Panel D: Firing-rate time course for an example module, for foreground, background, and inhibitory sub-populations (respectively in black, blue, and red) in the AW state. Adapted from Pastorelli et al. (79).

grounded starting point to tackle models of behaviorally relevant, awake states (Han et al. (44), Luczak et al. (60), Curto et al. (24)). A critical juncture in such a logical sequence is the description of the dynamic transition between SWA and asynchronous, irregular activity (AW, asynchronous wake state) as observed during fade-out of anesthesia, for instance, the mechanism of which is still a partially open

problem (Solovey et al. (101), Curto et al. (24), Steyn-Ross et al. (102)). To help determine the mechanism of this transition, it may be of interest to identify the factors enabling the same nervous tissue to express global activity regimes as diverse as SWA and AW. Understanding this repertoire of global dynamics requires high-resolution numerical simulations of large-scale networks of neurons which, while keeping a manageable level of simplification, should be realistic with respect to both nonlinear excitable local dynamics and to the spatial dependence of the synaptic connectivity, as well as the layered structure of the cortex (Hill and Tononi (46), Bazhenov et al. (5), Potjans and Diesmann (84), Krishnan et al. (56)).

Notably, efficient brain simulation is not only a scientific tool, but also a source of requirements and architectural inspiration for future parallel/distributed computing architectures, as well as a coding challenge on existing platforms. Neural network simulation engine projects have focused on: flexibility and user friendliness, biological plausibility, speed and scalability (e.g. NEST 34, 52, NEURON 47, 17, GENESIS 124, BRIAN 41, 107). Their target execution platforms can be either homogeneous or heterogeneous (e.g. GPGPU-accelerated), high-performance computing (HPC) systems, such as the ones described by Modha 70, Izhikevich 50, Nageswaran 75, or neuromorphic platforms, for either research or application purposes (e.g. SpiNNaker 30, BrainScaleS 96, TrueNorth 67).

From a computational point of view, SWA and AW pose different challenges to simulation engines, and comparing the simulator performance in both situations is an important element in assessing the general value of the choices made in the code design. During SWA, different and limited portions of the network are sequentially active, with a locally high rate of exchanged spikes, while the rest of the system is almost silent. On the other hand, during AW the whole network is homogeneously involved in lower rate asynchronous activity. In a distributed and parallel simulation framework, this raises the question of whether the computational load on each core and the inter-process communication traffic are limiting factors in either cases. We also need to consider that activity propagates for long distances across the modeled

cortical paths, therefore the impact of spike delivery on the execution time depends on the chosen connectivity.

In preparation of this, part of my activity was concentrated in achieving a fast and flexible simulator: the Distributed and Plastic Spiking Neural Networks (DPSNN) engine, able to run models of SWA as well as of AW-like activity (79). Early versions of the simulator, described in 77 by Paolucci et al., originated from the need for a representative benchmark developed to support the hardware/software co-design of distributed and parallel neural simulators. DPSNN was then extended to incorporate the event-driven approach of Mattia and Del Giudice (62), implementing a mixed time-driven and event-driven strategy similar to the one introduced by Morrison et al. (71). In the following chapters, the simulator is described in more details and compared with the state-of-the-art available at that moment.

3.5.1 Simulations at very large scale on CPU-based HPC

Cortical synapses organization supports a range of dynamic states on multiple spatial and temporal scales, from synchronous slow wave activity (SWA), characteristic of deep sleep or anesthesia, to fluctuating, asynchronous activity during wakefulness (AW). Such dynamic diversity poses a challenge for producing efficient large-scale spiking models that embody realistic metaphors of short- and long-range synaptic connectivity. In fact, during SWA and AW, different spatial extents of the cortical tissue are active in a given timespan and at different firing rates, which implies a wide variety of loads of local computation and communication. A balanced evaluation of performance and robustness should therefore include tests of a variety of cortical dynamic states.

The designed spiking model (79) is able to express both SWA and AW-like activity using neural populations organized as bidimensional grids of modules (mimicking cortical columns) as in Capone et al. (16) which reflects the modular organization of the cortex.

With this model, using available simulators, both DPSNN and NEST, I explored

networks up to 192×192 modules, each composed of 1,250 integrate-and-fire neurons with spike-frequency adaptation, and exponentially decaying inter-modular synaptic connectivity with varying spatial decay constant.

For each population in the modular network, specific parameters for the neural dynamics can be defined, as well as specific intra- and inter-columnar connectivity and synaptic efficacies. Connectivity among different populations can be modeled with specific laws based on distance-dependent probability, specific to each pair of source and target subpopulation. By suitably setting the available interconnections between different populations, cortical laminar structures can also be potentially modeled.

Each local module is always composed of $K = 1250$ neurons, further subdivided into subpopulations. We implemented a ratio of 4:1 between excitatory and inhibitory neurons (K_I=250 neuron/module); in each module, excitatory neurons (E) were divided into two subpopulations: K_F=250 (25%) strongly coupled "foreground" neurons (F), having a leading role in the dynamics, and K_B=750 (75%) "background" neurons (B) continuously firing at a relatively low rate. Populations on the grid are connected one to each other through a spatial connectivity kernel. The probability of connection from excitatory neurons decreases exponentially with the inter-module distance d:

$$C_{ts\lambda}(d) = C^0_{ts\lambda} \times exp(\frac{-d}{\lambda}) \,. \tag{3.3}$$

More specifically, d is the distance between the source ($s = \{F, B\}$), and target ($t = \{F, B, I\}$) module, and λ a characteristic spatial scale of connectivity decay. d and λ are expressed using inter-modular distance units (imd). For the simulations here described, the translation to physical units sets imd in the range of a few hundred micrometers. Simulations are performed considering different λ values (0.4, 0.5, 0.6, 0.7) imd, but $C^0_{ts\lambda}$ is set so as to generate the same mean number of projected synapses per neuron ($M_t = 0.9 * K_t, t = \{F, B, I\}$) for all λ values. Connections originating from inhibitory neurons are local (within the same local

module) and also in this case $M_t = 0.9 * K_t, t = \{F, B, I\}$. All neurons of the same type (excitatory/inhibitory) in a population share the same mean number of incoming synapses. The connectivity has open boundary conditions on the edges of the two-dimensional surface.

Synaptic efficacies are randomly chosen from a Gaussian distribution with mean J_{ts} and SD ΔJ_{ts}, chosen in different experiments to set the system in different working regimes and simulate different states. The procedure for the selection of the efficacies is based on a mean-field method described above. Each neuron also receives spikes coming from neurons belonging to virtual external populations, collectively modeled as a Poisson process with average spike frequency ν_{ext} and synaptic efficacy J_{ext}. Excitatory neurons are point-like leaky integrate-and-fire (LIF) neurons with spike frequency adaptation (SFA), as described in Gigante et al. (35) and in Capone et al. (16). SFA is modeled as an activity-dependent self-inhibition, described by the fatigue variable $c(t)$. The time evolution of the membrane potential $V(t)$, and $c(t)$, of excitatory neurons between spikes is governed by:

$$\begin{cases} \dot{V} &= -\frac{V-E}{\tau_m} - g_c \frac{c}{C_m} + \sum J_i \delta(t - t_i - \delta_i) + \sum J^i_{ext} \delta(t - t_i^{poiss}) \\ \dot{c} &= -\frac{c}{\tau_c} + \alpha_c \sum_k \delta(t - t_{\rm sp}^{(k)}) \end{cases} \tag{3.4}$$

τ_m is the membrane characteristic time, C_m the membrane capacitance, and E the resting potential. SFA is not considered for inhibitory neurons; that is, in (3.4), the second equation and the $g_c \frac{c}{C_m}$ term in the first are absent. Incoming spikes, generated at times t_i, reached the target neuron with delay δ_i and provoked instantaneous membrane potential changes of amplitude J_i. Alike, external stimuli produce a J^i_{ext} increment in the membrane potential, with t_i^{poiss} representing the spike times generated by a Poisson distribution of average ν_{ext}. An output spike at time $t_{\rm sp}^{(k)}$ was triggered if the membrane potential exceeded a threshold V_θ. On firing, the membrane potential was reset to V_r for a refractory period τ_{arp}, whereas c was increased by the amount α_c. Once the network connectivity and the neural

dynamics have been defined, synaptic efficacies and external stimuli can be set to determine the dynamical states accessible to the system, by means of mean-field theory.

3.5.2 Spiking simulation speed and scaling measures

The model above described has been used to test and improve the DPSNN simulator (see more details on Paolucci et al. 77 and on Pastorelli et al. 79 and 78). The benchmark measures assessing simulator scalability in terms of run time, initialization time, and memory usage herein described, were performed on the Galileo server platform provided by the CINECA consortium. Galileo is a cluster of 516 IBM nodes, where each node includes 16-cores, distributed on two Intel Xeon Haswell E5-2630 v3 octa-core processors clocked at 2.40 GHz. The nodes are interconnected through an InfiniBand network. Hyperthreading is disabled on all cores, therefore the number of MPI processes launched during each run exactly matches the hardware cores, with a maximum of 1,024 hardware cores (or equivalently MPI processes) available for any single run, due to a specific configuration of the server platform, in which the maximum-allowed partition of the server usable for a single run includes 64 nodes.

DPSNN execution times for SW and AW state simulations are reported in Figure 3.6. Panels A and B show example snapshots of the simulated network activity in the SW and AW states, respectively; panels C and D show corresponding power spectra of the network activity, confirming the main features predicted by theory for such states (such as the low-frequency power increase for SW, the high-frequency asymptote proportional to the average firing rate, the spectral resonances related to SFA, and delays of the recurrent synaptic interaction). Panels E and F represent strong scaling for SW and AW: the wall-clock time required to simulate 1s of activity, *vs* number of processes, for different network sizes, up to 46 Million neurons and 70 Giga synapses. In the same plots, weak scaling behavior is measured by joining points referring to a four-fold increase in both the network size and the number of

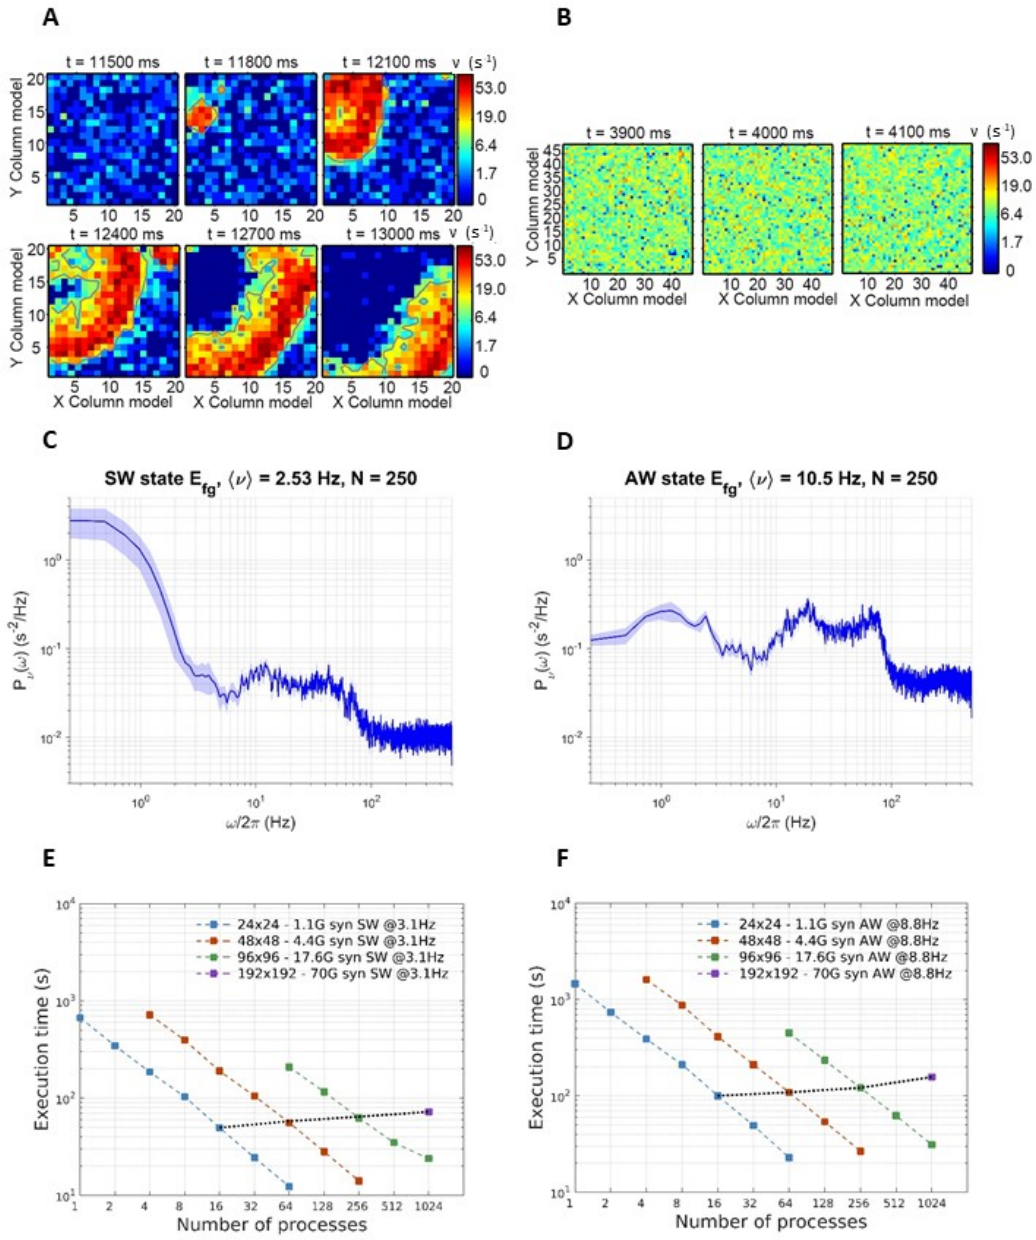

Figure 3.6. Simulation time scaling and phenomenological behavior. Panel A: time consecutive snapshots of the activity distribution in space, showing the wavefront propagation during a simulation expressing SW states. Panel B: consecutive snapshots of the whole network activity in an asynchronous state. Panels C and D: power spectra of network activity respectively in SW (C) and AW (D) states. Panels E and F: scaling of wall-clock execution time for 1 s of SW (E) and AW (F) simulated activity. In both SW and AW states, the scaling has been measured on different network sizes. Adapted from Pastorelli et al. (79).

processes.

Simulation speeds are measured using the equivalent synaptic events per second, defined as the product of the total number of synapses and the number of spikes occurred across the whole simulation, divided by the elapsed execution time. This way, a comparison of the simulation cost among different problem sizes and hardware/software resources (core/processes) can be captured in a single graph. Departures from the ideal scaling behavior (linearly decreasing strong scaling plots, horizontal weak scaling plots) are globally captured in panel A of Figure 3.7, for both SW and AW states and for all the problem sizes. For the simulations reported in this study, the simplified synaptic-events-per-second metric is a good approximation of a more complex metric separating recurrent and external events (Poisson noise), as demonstrated by panel B of Figure 3.7: looking in more detail, the simulation of recurrent synaptic events is slower than that of external events that are locally generated by the routine that computes the dynamics of individual neurons. In our configurations there is about one order magnitude more recurrent synaptic event than external.

3.5.3 Efficiency for real time simulations

Fast simulation of spiking neural network models plays a dual role: (i) it contributes to the solution of a scientific grand challenge — i.e. the comprehension of brain activity — and, (ii) by including it into embedded systems, it can enhance applications such as autonomous navigation, surveillance and robotics, requiring real-time performances. Moreover, real-time simulation of neural networks will be essential for understanding the mechanisms underlying the cognitive functions of the brain. Indeed, brain simulations should be embedded in complex environments, e.g. robotic platforms interacting with the world in real-time, which makes requirements on power consumption so much tighter. Therefore, cortical simulations assume a driving role in shaping the architecture of either specialized and general-purpose multi-core/many-core systems to come, standing at the crossroads between em-

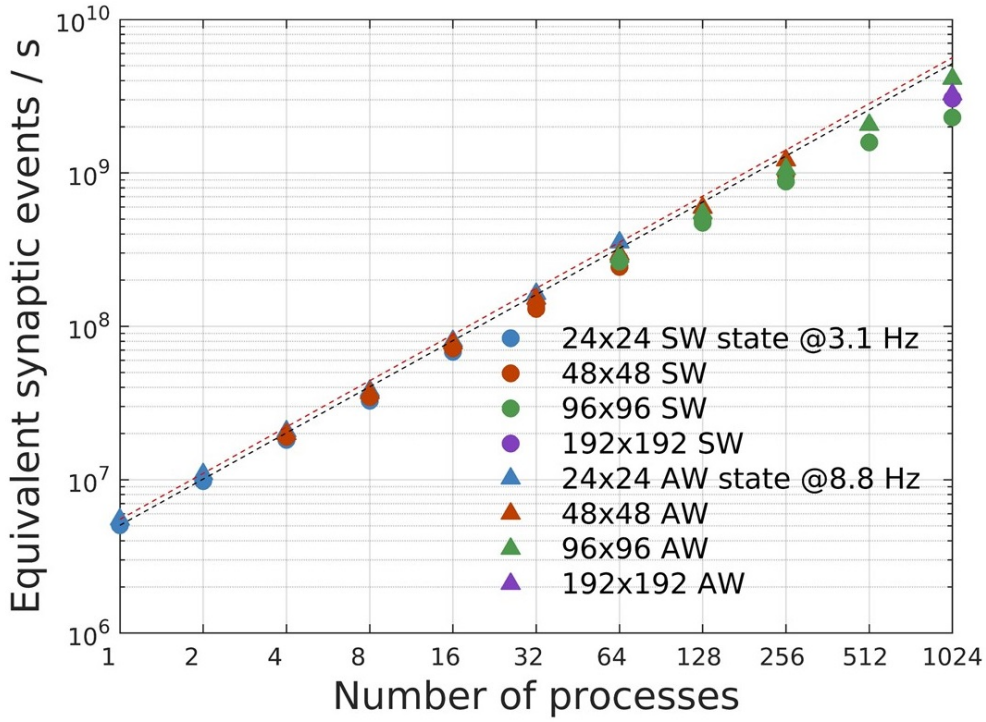

Figure 3.7. Simulation speed-up. Speed-up of the total number of equivalent simulated synaptic events per wall-clock second evaluated for SW (circles) and AW simulations (triangles) on the four used network configurations. Dashed lines stand for ideal speed-up of the simulated synaptic events for SW (black) and AW (red). Adapted from Pastorelli et al. (79).

bedded and HPC. See, for example Merolla et al. (67), describing the TrueNorth low-power specialized hardware architecture dedicated to embedded applications, and Stromatias et al. (109) discussing the power consumption of the SpiNNaker hardware architecture, based on embedded multi-cores, dedicated to brain simulation. Worthy of mention are also the work of Gewaltig and Diesmann (34) and the one of Modha et al. (70) as examples of approaches based on standard HPC platforms and general-purpose simulators.

I would also mention, even if it's not the focus of this thesis, the activity done to measure the power consumption and energy-to-solution for real-time cortical simulation on DPSNN and the profiles of the relative scaling concerning computation, communication and synchronization (Simula et al. 99). Specifically, we perform several neural simulations to compare the performances of ARM and Intel-based

multi-core platforms, with further focus on the possible impact of the usage of off-the-shelf vs. custom networking components.

In this domain, being "real-time", under a "soft" assumption, means a working point for the application such that the total wall-clock time for running it is not greater than the total simulated time, a condition necessary, but not sufficient, for robotics applications and embedding HPC simulations into virtual or real world environments, that would impose more stringent "hard" constraint to be satisfied at the scale of each step, lasting at most a few tens of milliseconds each. The aim of this work is to identify the obstacles that impede reaching the real-time target for large neural networks. We performed a set of strong scaling tests on neural networks of increasing size. For all network, we simulated 10 s of neural activity.

Figure 3.8 shows the runtimes for three neural network sizes. They should all be able to run in real-time if the scaling valid for larger configurations applied. Indeed, the 20480 neurons configuration reaches real-time (9.15 seconds to simulates 10 seconds of activity). The network with 20480 neurons reached its maximum speed when distributed on 32 processes (Figure 3.8). Communication and synchronization are the main obstacles against scaling (see Figure 3.8). For the 20480 neuron configuration they block a further acceleration over 32 processes and start impeding the scaling toward real-time of larger neural networks. The profiling of the computation and communication components demonstrated the critical impact of interconnect on the scaling, limiting the size of the network that can be simulated in real-time.

In conclusion, the design of low-latency, energy-efficient interconnects supporting collective communications is of primary importance to enable a time- and energy-efficient exchange of neural spikes; this is expected to not only make cortical simulations possible at a larger scale but also push their use in embedded systems where it is often precluded by tight real-time constraints and limited power budget.

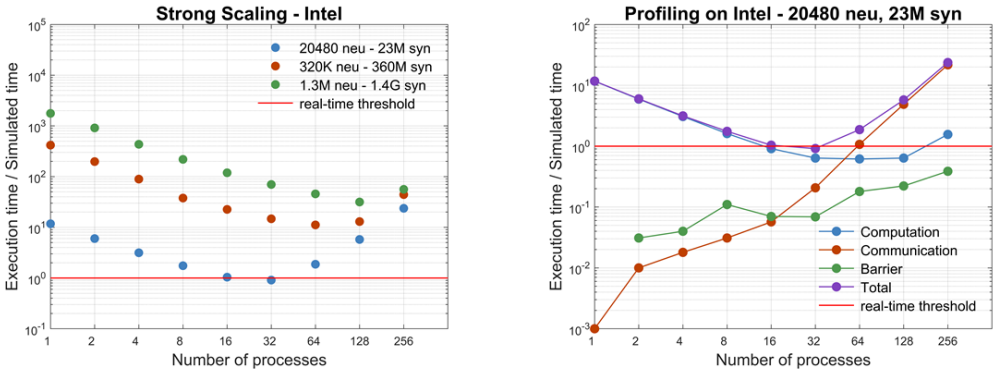

Figure 3.8. Real time simulations. Left panel: Strong scaling of different problem sizes on an Intel-based platform. The red line is the threshold to be reached for soft real-time execution. Right panel: detail of the analysis on the Intel-based platform, in which is evident that communication and synchronization are the main obstacles against scaling. Adapted from Simula et al. (99).

3.5.4 Acceleration throughout a GPU-based HPC

It is evident, from what exposed also in section (3.5.3), that the cost of large-scale spiking simulations is really high in terms of resources and computational time: the simulation of brain activity at the level of signals produced by individual neurons is extremely demanding, even if it is limited to relatively small regions of the brain. Therefore, there is a growing interest in the development of high-performance hardware and software tools for efficient simulations of large-scale networks of spiking neuron models.

At the moment, this represents a fertile field of research. In recent decades, the use of highly parallel hardware systems has been investigated for simulating large-scale networks of spiking neurons. Such systems include custom made neuromorphic very-large-scale-integration (VLSI) circuits (Indiveri et al. 49), field programmable gate arrays (FPGAs) (Wang et al. 117) and systems based on graphical processing units (GPUs) (Brette and Goodman 8, Sanders and Kandro 93, Garrido et al. 33, Vitay et al. 115, Chou et al. 20). Compared to other highly parallel systems, the latter have the advantage of a relatively low cost, a great versatility, thanks also to the possibility of using the CUDA (Compute Unified Device Architecture) platform,

and a sustained technological development driven by the consumer market. General purpose computing on graphical processing units (GPGPU) is widely employed for massively parallel computing. GPGPUs can significantly reduce the processing time compared to multi-core CPU systems for tasks that require a high degree of parallelism, because a single GPU can perform thousands of core computations in parallel.

Following this trend, the NeuornGPU library (38) has been developed: a GPU library for simulation of large-scale networks of spiking neurons, written in the C++ and CUDA-C++ programming languages. NeuronGPU can simulate networks of any neuron model and synaptic current model whose dynamics can be described by a system of ordinary differential equations (ODEs). The network can be easily described by the user: the neuron parameters, connection weights and synaptic delays can be initialized either using fixed values or through arrays or probability distributions. Neuron groups can be connected either using predefined connection rules (one-to-one, all-to-all, fixed indegree, fixed outdegree, fixed total number) or by user-defined connections. In addition to the standard synapse model, nearest-neighbor spike-timing-dependent-plasticity (STDP) is also available (Morrison et al. 72). Different types of devices can be simulated, including Poisson signal generators, spike generators, multimeters and parrot neurons.

Several tests have been conducted on this library and comparisons with the state-of-the-art reported good results in terms of accuracy and velocity of execution. NeuronGPU was recently proposed for being integrated within the NEST neural simulator (39).

Chapter 4

Models of cycles of incremental learning and sleep

The scientific question at the basis of my thesis was the study of the effects of deep-sleep-like activity on learning. For the models of cycles of incremental learning and sleep, the choice to focus exclusively on slow oscillation, neglecting the spacial structure of slow waves widely discussed in previous chapter, is due to a stepwise approach. The interaction of sleep and learning must be based on a system provided of plasticity, totally absent in the simulations presented in the previous chapter, as already underlined. Now plasticity (together with the effect of deep-sleep on it) becomes the central node of the study. The idea was that to start with a plastic system and its interplay with slow oscillations, to move in future steps toward the application of a spatial structure to the whole system. In this perspective, the knowledge acquired on how to speed-up large-scale simulations, as described in Chapter 3, represents an essential building block for the extension of this model toward biological dimension and plausibility (left-right visual areas, different decodification layers to improve models accuracy, etc.).

To approach the problem, I started using a simplified thalamo-cortical model trained to encode, retrieve and classify images of handwritten digits (Capone, et al., 15). During slow oscillations, spike-timing-dependent-plasticity (STDP) produces a differential homeostatic process. It is characterized by both a specific unsupervised

enhancement of connections among groups of neurons associated to instances of the same class (digit), as described in the left panel of Figure 4.1, and a simultaneous down-regulation of stronger synapses created by the training (right panel of the same Figure). The starting hypotheses is that this hierarchical organization of post-sleep internal representations favours higher performances in retrieval and classification tasks. The supposed mechanism should be based on the interaction between top-down cortico-thalamic predictions and bottom-up thalamo-cortical projections during deep-sleep-like slow oscillations. Indeed, when learned patterns are replayed during sleep, cortico-thalamo-cortical connections favour the activation of other neurons coding for similar thalamic inputs, promoting their association.

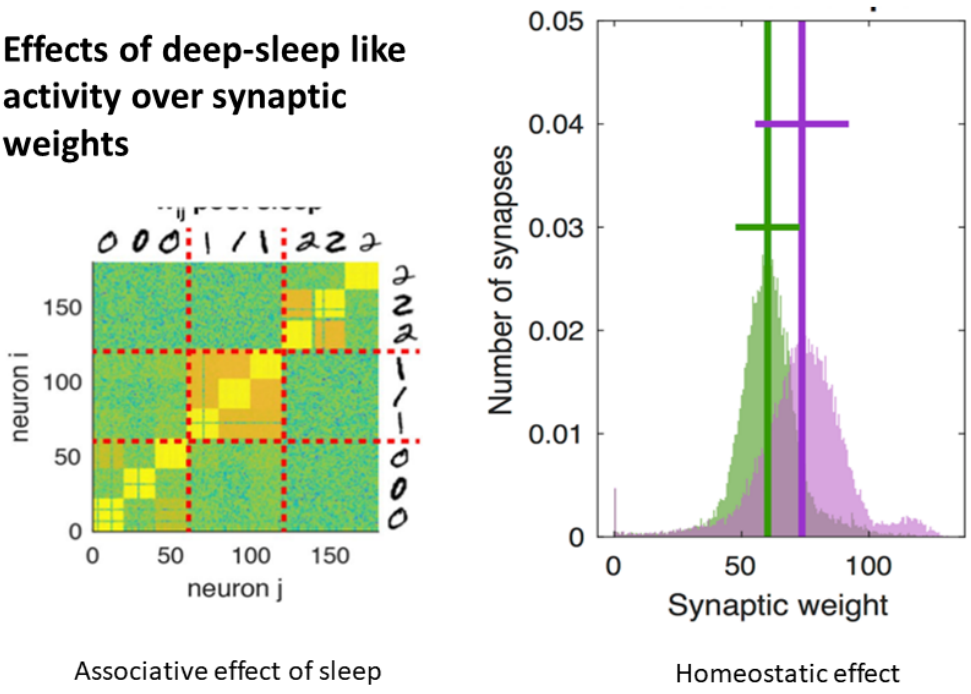

Figure 4.1. Effect of deep-sleep on synapses: associative and homeostatic effects. Left panel: Synaptic weights matrix of the recurrent connectivity of excitatory population after the occurrence of sleep-like activity. The creation of associations among synapses of different digit start emerging. Right panel: pre-sleep synaptic weights (violet), and post-sleep (green) synaptic weights.

The model has been then improved as described in Golosio et al. (37) with parameter calibration thanks to an iterative comparison with experimental results. As

a consequence, also contextual signal has been more correctly calibrated, producing a very important consequence: the memory of learned digits during the awake phase are more equally represented, while, on the contrary, allowing for continuous cycling of learning-sleep-learning no possible in the previous model. This enables a huge advantage in terms of training velocity and precision (see section 4.5). Moreover, the model demonstrates resilience when subjected to noisy perceptions with better performances than other AI algorithms as for instance the K-nearest-neighbour (Knn) algorithms, that is one of the most common and performant classification methods (see Figure 4.2).

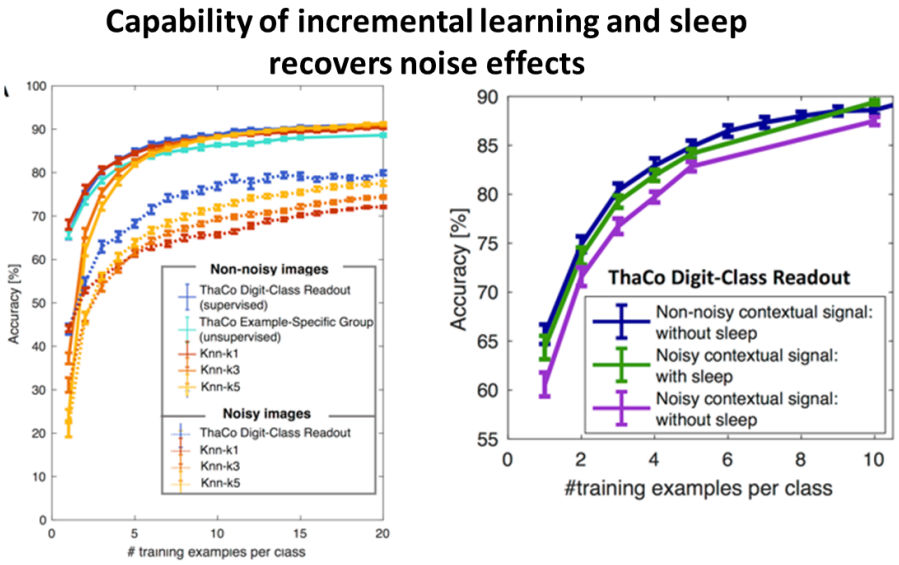

Figure 4.2. Sleep mitigation of incremental learning in case of noisy signals for the improved model (Thaco). Left panel: comparison of the average accuracy of the proposed thalamo-cortical spiking model compared to artificial K-nearest-neighbour incremental algorithms in absence of noise (solid lines) and with noisy inputs (dotted lines). Right panel: Sleep mitigation of the effects of noisy contextual signals on classification performances. Adapted from Golosio et al. 37

In the following of this chapter, the thalamo-cortical model (and its subsequent evolution we named *Thaco*) is described in detail, giving a deep insight of it and of the obtained results. The main source of this material has been extracted by some of my previous articles published during the Phd period with my group of work in the framework of the HBP APE Research Group: the paper Capone et al.

(15) published in Sci. Rep. in 2019, and the arXive Golosio et al (37), accepted for publication in PLOS Computational Biology.

4.1 Data driven thalamo-cortical model

A few computational models have been developed to investigate the interaction of sleep-like activity and plasticity. In González-Rueda et al. (40), the authors showed that Up states specifically mediate synaptic down-scaling with beneficial effect on signal to noise ratio received by post-synaptic neurons. The effect of thalamo-cortical sleep on pre-stored time sequences is explored by Wei in 122 and 123.

In this work, also reported in Capone et al. (15), we focus on sleep mediated memory association and its implications on cognitive tasks performances. The presented minimal thalamo-cortical model, after being trained on handwritten characters in unsupervised mode, is induced to express sleep-like dynamics. We measured its effects on the classification accuracy (defined as the percentage of correct digit identified by the model over the whole set of proposed digits), the structure of the synaptic matrix and the firing rate distributions.

Slow oscillations (SO) are considered the default emergent activity of the cortical network (Sanchez-Vives, et al, 92) and are observed during the deepest of physiological non-REM sleep stages as an alternation between *Down states* (characterized by nearly silent neurons) and *Up states* (in which a subset of neurons goes in a high firing rate regime) occurring at a frequency in the range $[0.5, 4]$ Hz (*delta* band). We set the model SO at a comparable frequency. SO activity is expected to play two complementary roles, which are separately mediated by Up states and Down states. Down states would play a purely biological function, with a lower immediate impact on cognitive performance. The role of Down states, with a majority of neurons put in a silent state for a long fraction of deep sleep time (several hundreds of ms), would serve a restoration purpose, enabling periodic biological maintenance and recovery, as it happens in the whole body when at rest (Vyazovskiy and Harris, 114). Our modeling and investigation are focused only on the effects mediated by the Up states

dynamics. During sleep, external perceptions are, at least, strongly attenuated, and the majority of the motor system is blocked (Bucci and Grasso, 13). For this reason, in our model the local interaction between cortex and thalamus is crucial during sleep rather than contextual signal coming from other cortical modules and sensory input coming from thalamic pathways.

In order to make a biologically realistic learning protocol and to implement the role of the context in the learning phase, we took inspiration from the "organizing principle" of Larkum (58) for the Cerebral Cortex, which describes specific computational strategies implemented by the neuronal structure of Layer 5. The architecture is grounded on the separation of intra-areal and inter-areal contextual information (reaching the apical dendrites of pyramidal neurons) from the feed-forward flow of area specific information (targeting its basal synapses). Cellular mechanisms, like Ca^{++} spikes, promote the detection of coincidence between contextual and specific activity. High-frequency bursts of spikes are emitted when the coincidence is detected. Relying on these observations we introduced in our model external stimuli mimicking contextual information which changes the effective firing threshold of specific subsets of neurons during the presentation of examples in the training phase. For each example, a vector of features is projected toward the cortical network by thalamic neurons. Due to the change in the perceptual effective firing threshold, spike-timing-dependent-plasticity (STDP) creates stronger bottom-up (thalamo-cortical) and top-down (cortico-thalamic) connections between a subset of cortical neurons and the thalamic neurons.

4.1.1 Network architecture

The network has been designed to be a minimal thalamo-cortical model as described by Destexhe in 26 (see Figure 4.3), composed of two populations of cortical neurons (one excitatory cx and one inhibitory in) and two thalamic populations (one excitatory tc and one inhibitory re). The spiking simulation of this model was performed using the NEST simulation engine, the high-performance general purpose

simulator developed by the NEST Initiative, release 2.12 (57).

The role of cortical inhibitory neurons in our model is to provide a shared inhibition supporting a winner-take-all mechanism. Indeed, when a part of the whole network responds to a specific input, the rest of the network is inhibited. A classical choice to approximate a biological cortical network is to set a 4:1 ratio between excitatory and inhibitory neurons. Keeping a fixed excitatory:inhibitory ratio is not critical for the model here presented. In our runs, the ratio varied from 4:1 to 1:1 because for simplicity we kept the number of inhibitory fixed in all simulations, while we increased the number of excitatory with to the number of training examples.

The connection probability is $p = 1.0$ for the populations connected by the arrows in Figure 4.3. The learning mechanism is allowed by symmetric ($\alpha = 1.0$) spike-timing-dependent-plasticity (STDP) present in the $cx \rightarrow cx$, $cx \rightarrow tc$ and $tc \rightarrow cx$ connections which shapes the weights structure.

The *contextual signal* is a Poissonian train of spikes which mimics a contextual signal coming from other brain areas and selectively facilitates neurons to learn new stimuli. The *top-down prediction* is the signal flowing through $cx \rightarrow tc$ connections, predicting the thalamic configuration which activated a specific cortical activity pattern.

We tested the role and the mechanisms of the occurrence of SO in such thalamo-cortical network model which was previously trained to learn and recall images. In this case we used hand-written digits examples extracted from the MNIST dataset, pre-processed through the application of the histogram of oriented gradients (HOG) algorithm, that allows the transformation of each image in a sequence of bits (in our case we used a 28x28 pixel image, histograms applied on a 14x14 cell size with stride equal to 7 and 9 bins per histogram, with a total binary encoding of 4 bit/feature). The network model included thalamic relay (tc) and reticular (re) neurons in the thalamus, as well as pyramidal neurons (cx) and inhibitory interneurons (in) in the cortex (Figure 4.3, panel A) following a standard minimal structure for thalamico-cortical models (Destexhe, 26). Panel B shows an example of activity time-course in

the *cx* and *tc* populations during the training phase, the retrieval phase, and the early stage of sleep phase.

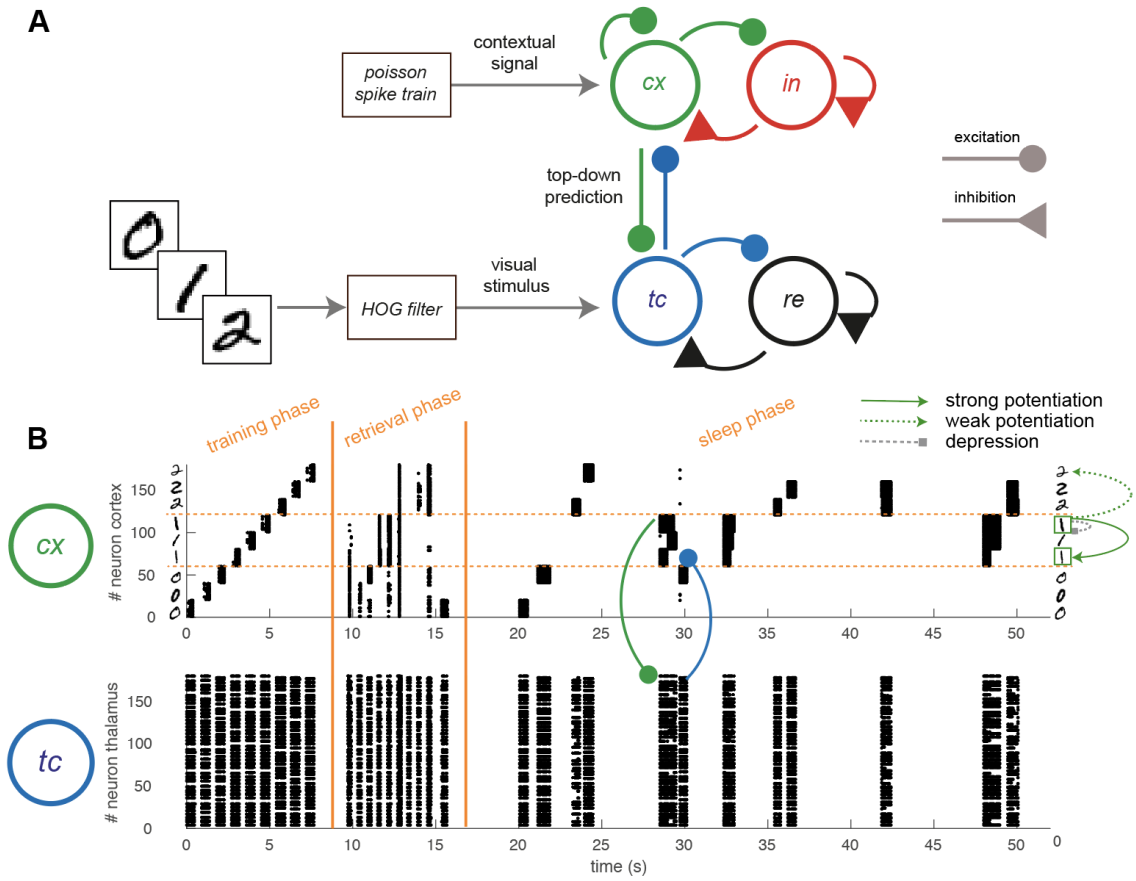

Figure 4.3. Thalamo-cortical model and protocol description. Panel A: sketch of the structure of the simplified thalamo-cortical model considered, which is composed of an excitatory and an inhibitory population both for the cortex (*cx*, *in*) and for the thalamus (*tc*, *re*). Connectivity structure is represented by solid lines. The visual input is fed into the model through the thalamic population, mimicking the biological visual pathways. In the training phase a lateral stimulus enhances a specific subset of *cx* neurons to preferentially represent the stimulus. Panel B: activity produced during training phase, pre-sleep retrieval and the first 40*s* of SO activity in the *cx* (top) and *tc* (bottom) populations. Only first 180 neurons in *tc* population are shown for visual purposes. In the training 3 instances of 3 classes of digits (0,1,2) are learned by the network. In the replay during sleep, thalamo-cortical connections promote the activation of neurons coding for similar patterns of activity, causing the potentiation of cortico-cortical connections between neurons representing digits of the same class. A general depression reduces the largest synaptic weights. Post SO retrieval is not shown. Adapted from Capone et al. 15.

4.1.2 Training and pre-sleep retrieval

During the training, 9 different images were presented to the network, in a first set of runs: 3 instances for each class of digit, for a total of 3 different classes (e.g. 0,1,2). In a second group of runs, 30 examples per digit constituted the training set. For each image an external stimulus (*contextual signal*) induced a different subset of *cx* neurons to code for that specific image, with STDP shaping the intra-cortical, the thalamo-cortical and the cortico-thalamic connectivity. In order to adopt the Larkum prescription (58), the parameters are set to make the cortical neurons fire during the training only if they receive both sensory and contextual stimuli. This training procedure works even in the extreme case where only one neuron is used to code each digit example. However for one or a few neurons, self sustained oscillations would not be well defined. For this reason, we chose a population of 20 cortical neurons for each newly presented example, but we verified that the same qualitative results are obtained also for the cases in which each example is coded by a different number of cortical neurons (for example we tested it for 10, 15 and 25 neurons).

After the training, images were presented again (retrieval phase) without the external stimulus. Thanks to the synaptic reinforcement of a specific group of neurons induced by STDP during the training phase, the population of *cx* neurons responding to that specific image is the same as the one selected in the training phase by the external stimulus, demonstrating the success of the retrieval. In a few words, just reading the firing rate of the different *cx* population (each one coded for a specific image during the training phase), we get the answer of the model to the presented stimulus.

4.1.3 Induction of Slow-oscillation

A non-specific stimulus at low firing rate was provided to cortical neurons only, while model parameters are modulated eliciting the spontaneous occurrence of cortically generated Up states and of thalamo-cortical SO (see Figure 4.3, panel B).

We relied on the framework of mean-field theory to obtain a model displaying

different dynamical regimes. An oscillatory regime, closely resembling Slow Oscillations observed in deep sleep and anesthetized states can be induced by introducing a relatively strong recurrent excitation and spike frequency adaptation, as described in Gigante et al. (35) and Capone et al. (16). We tuned the parameters of the network to make it display SO frequency (between $0.25Hz$ and $1.5Hz$) and Up state duration (a few hundred of milliseconds) comparable with experimental observations in deep sleep recordings (see Contreras and Steriade 22, Steriade et al 104).

The activity (Figure 4.3, panel B) during the initial stages of simulated SO displays that Up states are independently sustained by populations coding for different memorized images. However, thanks to the cortico-thalamo-cortical pathways each population tends to recruit other populations sharing similar thalamic representations. Indeed, when a population initiates an Up states, it activates thalamic patterns similar to the ones responsible for the activation of the population itself during the retrieval phase. This is what we call *top-down prediction*. The thalamus, in turn, activates all the other populations which coded for a similar thalamic input. Across the sleep period, thanks to the cortico-cortical plasticity, the co-activation of populations coding for similar *predicted* synaptic input becomes a more and more prominent feature.

4.2 Effects of deep-sleep activity on synaptic weights and on firing rates

We first considered the case with the training set composed of 3 classes and 3 examples per class. After the training stage the system undergoes a $600s$ period of sleep. During this stage, the activation of groups of neurons associated to different training examples induces, through STDP, not only a synaptic pruning but also the creation of stronger synapses between groups of neurons that share enough commonality in the features they received during the training.

This effect can be noticed comparing the structure of the synaptic matrix before and after sleep as in Figure 4.4: panels A and B report the change from the initial

flat structure (which reflects the individual training examples) towards a hierarchical structure (embedding the categories of learned digits).

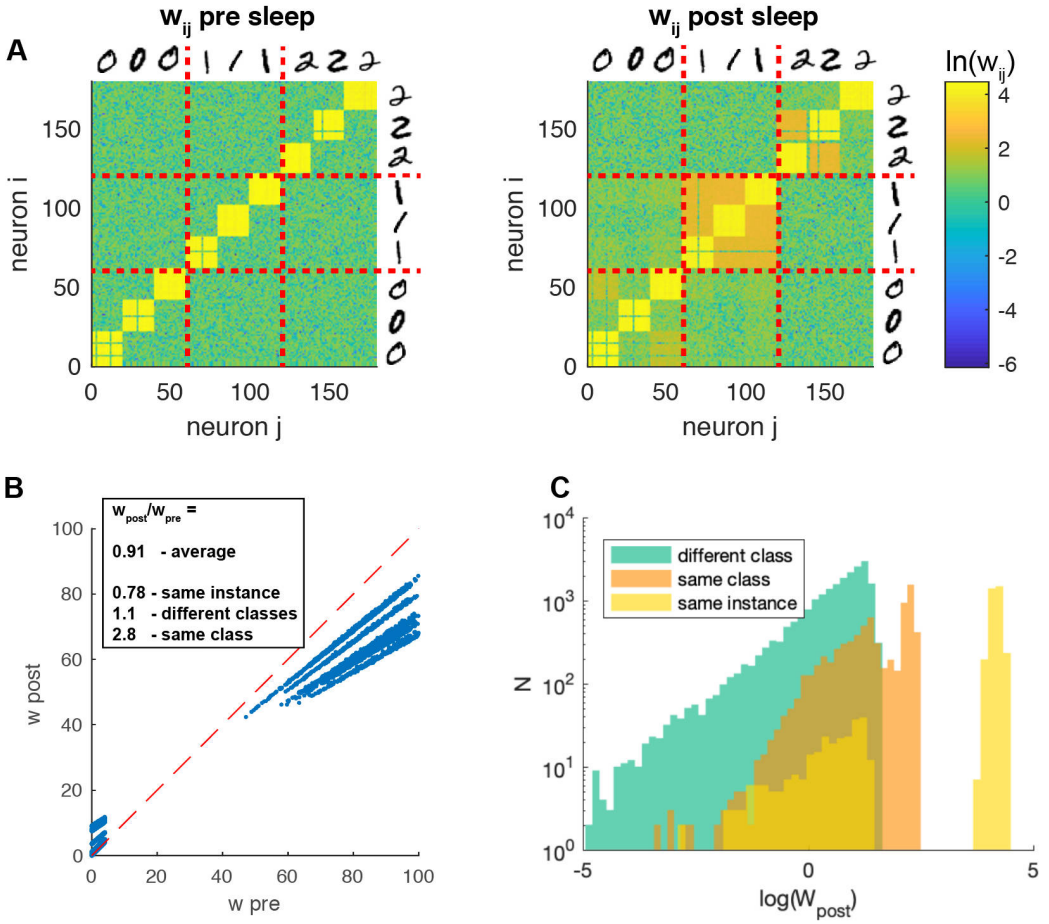

Figure 4.4. SO effects on connectivity structure. Panel A: synaptic weights matrix of the recurrent connectivity of *cx* population, before (left) and after (right) the occurrence of sleep-like activity. The yellow squares represent high weights emerged between neurons encoding the visual input related to the same object (single instance of 0,1,2 ... image). Red solid lines separate the neurons encoding visual inputs related to different classes of objects (0,1,2 ...). Panel B: scatter-plot of the same synaptic weights before and after sleep. Panel C: synaptic weights after sleep, separated in three groups, synapses between neurons encoding the same object (yellow), the same class (but not the same object, orange) and different classes (green). Adapted from Capone et al. 15.

Indeed we can observe that novel synapses are created among examples in the same digit category (Figure 4.4, panel A-left). At the same time the system undergoes a down-scaling of the strongest synapses, those linking neurons coding for the same image (Figure 4.4, panel A-right). The differential effect on synapses can also be observed in Figure 4.4, panel C, where the histogram of synaptic weights after-sleep

is reported. The synapses between patterns encoding for different learned images of the "same class" and those between "different classes" (respectively orange and green distribution) are originally drawn from the same distribution by definition, while after sleep they are clearly differentiated.

The change in structure of the synaptic weights matrix modifies the activity expressed by the network in the retrieval phase. This can be appreciated looking at the difference of the correlations between groups of neurons before and after sleep (see Figure 4.5, panel A).

Figure 4.5, panel B reports the difference among correlations evaluated after and before sleep. It shows decorrelation (blue squares) of populations encoding different classes, and correlation (red regions) of the ones coding the same class. Such information is reported also in Figure 4.5, panel C, showing the correlation changes for populations in the same class (blue) and in different classes (green). Such effect might provide benefits in retrieval and classification tasks.

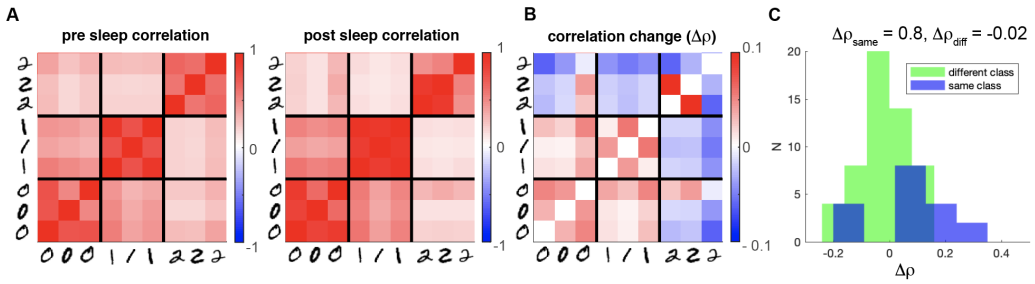

Figure 4.5. SO effects on internal representation. Panel A: activity correlation between all pairs of populations representing the single images before (left) and after (right) sleep. Panel B: correlation difference between after and before sleep. Panel C: histogram of correlation differences for populations encoding the same class (blue) and different classes (green). Adapted from Capone et al. 15.

The consistency of such result is testified by Figure 4.6. There, the same simulation is performed for different training sets (different examples of 0,1,2 digits). All simulations show that the synapses between neurons in the same class are the more reinforced (Figure 4.6, panel A, orange versus green) and that their internal representation has an increased correlation (Figure 4.6, panel B, blue versus green).

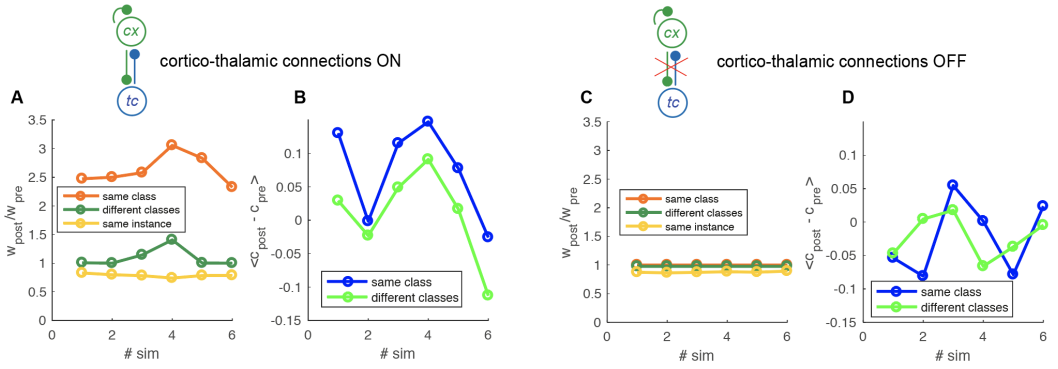

Figure 4.6. Analysis of populations: synaptic weights and comparison between correlations with and without *cortico-thalamic predictions*. Panel A: average ratio between weights post- and pre- sleep for each simulation (top). The different categories are separated in different colors. Yellow: synapses connecting neurons coding for the same image, orange: different image of the same class of digits and green: different classes. Panel B: the average change in correlation between post- and pre- sleep for each simulation (top) and histogram of the distribution over all the simulations (n=6, bottom). Blue: same class, green: different classes. Panels C and D: as in A-B but in absence of cortico-thalamic connections. Adapted from Capone et al. 15.

4.3 Cognitive effects of sleep – awake cycles

In our model we observed that sleep induces both the association of patterns encoding learned images belonging to the same category and a differential synaptic down-scaling. This is also reflected in a differential modulation of firing rates, producing observations similar to the ones reported by Watson et al. (118). Such effect is probably related to energetic optimization in biological networks, that has also beneficial effects on the performances of our network in the image recognition task.

We propose that such effect is due to the interplay between cortico-thalamic predictions and thalamo-cortical connections. In other words when a group of neurons undergoes an Up state it formulates a prediction in the thalamus by activating a thalamic pattern similar to the one received during training. In turn, the thalamus projects to the cortex and activates those populations trained for similar input patterns. This mechanism promotes the connections between populations of neurons coding for the images of the same class through STDP. To prove this, we reproduced the same experiment switching off the cortico-thalamic prediction. Results are

reported in Figure 4.6, panels C and D. It is evident that there is no sign of the preferential association observed in the control condition, nor in the synaptic structure (panel C), neither in the internal representation (panel D).

4.4 Performance improvements in post-sleep classification tasks

Finally, we evaluated the effect of sleep-like activity on the performance of a set of classification task trials. Networks were exposed to example and test images drawn from all the ten classes of digits in the MNIST dataset. Each simulation trial used a different test set of 250 images, and a set of 3 training examples per digit (30 training instances in total, also randomly extracted for each classification trial). Each network was exposed to the training examples using the same protocol discussed above for the simpler retrieval task. For each test image, the classification was determined looking for the class of the neuron responding with the higher firing rate. We note that class labels were used only during classification and not during the training that was completely unsupervised.

We observed a net increase in the classification accuracy across the sleep period. Panel A of Figure 4.7 reports in blue the time course of accuracy increase as a function of the sleep time. After $3000s$ of sleep, the improvement was on average $6.0\% \pm 0.5\%$ (accuracy rose from 58.0% to 64.0%, average performed over 24 simulations). In absence of thalamic feedback the improvement is significantly lower (red line in panel A), proving that the memory association due to the cortico-thalamo-cortical interaction is beneficial to performance in a classification task.

Average weights evolution as a function of sleep time are reported in panel B. Synapses between groups of neurons encoding for different instances of the same digit class (yellow solid line) were on average strongly potentiated, much more than the ones connecting training examples belonging to different classes (green solid line). Synapses interconnecting neurons representing individual training instances were down-scaled (orange solid line). When the thalamo-cortical feedback was switched

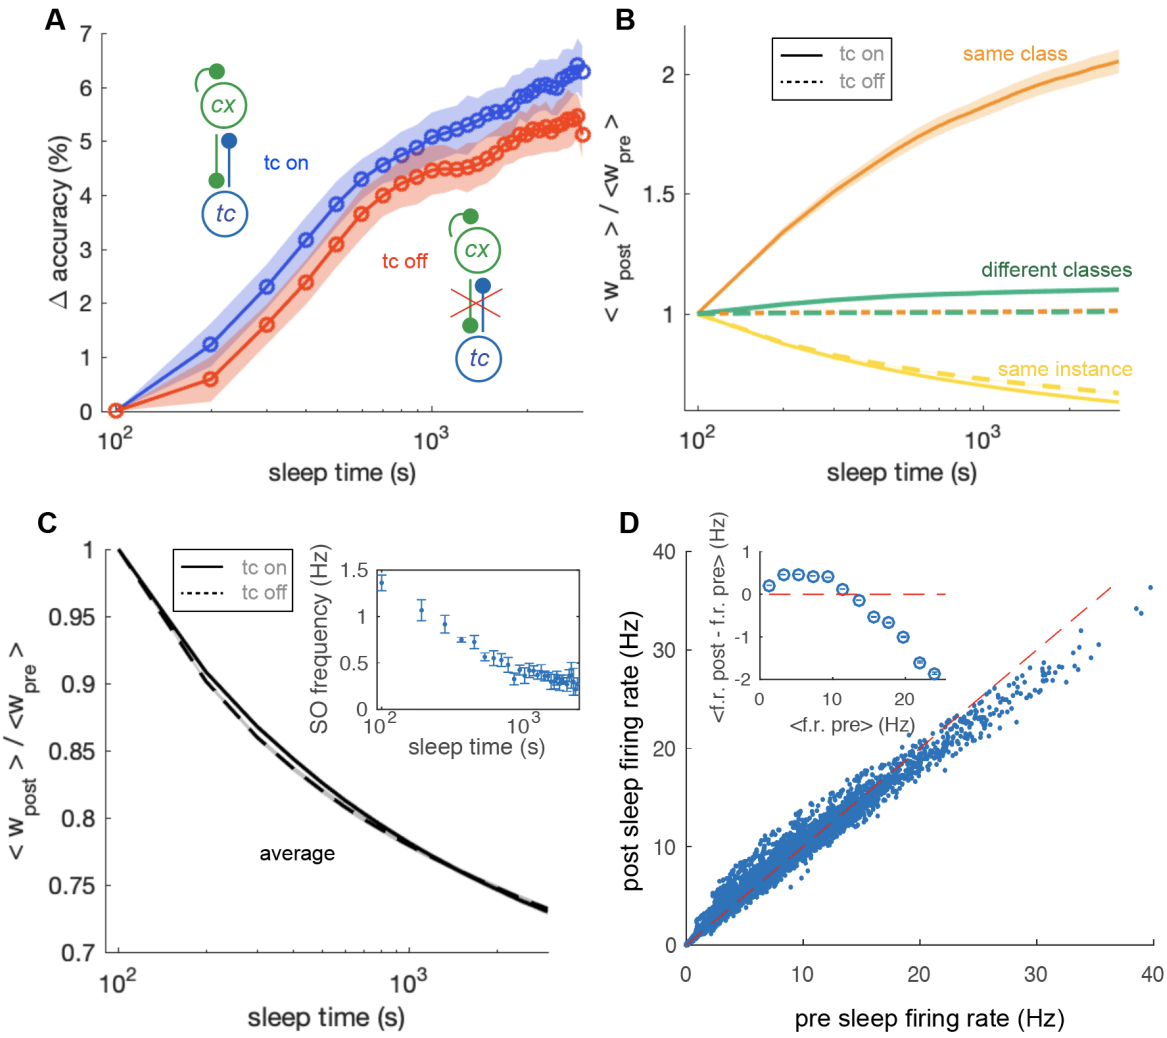

Figure 4.7. Sleep effects on a classification task. Panel A: change in classification accuracy across over 30 sleep epochs (100s each). Blue and red are respectively the conditions in which thalamus is on and off. The improvement in accuracy is averaged over 30 simulation trials. SEM is reported in the shading. Panel B: average synaptic potentiation and depression over 30 sleep epochs. The colors indicate connections between neurons coding the same instance (yellow), different instances of the same class (green) and instances of different classes (orange). Dashed and solid lines represent the comparison between the conditions in which thalamus is on and off. Panel C: average synaptic depression over all the synapses. (inset) Average decrease of SO frequency across sleep time, average over 4 simulations. Panel D: scatter of single neurons activity in 8 simulations averaged over time in a classification task before and after 3000s of sleep. Inset, average difference of activity after and before sleep as a function of activity before sleep. Adapted from Capone et al. 15.

off (same colors, dashed lines) this differential effect did not happen. This is the same effects already reported for the case with 9 training examples (Figure 4.4,

simpler retrieval task).

Such differential mechanism occurred together with a general synaptic depression (see panel C, thalamus on and thalamus off in solid and dashed line respectively). We notice that the average down-scaling is very similar for simulations executed in absence and presence of thalamic feed-back, allowing for a fair comparison of the two conditions. As a consequence of such general synaptic depression the SO frequency decreases over time (inset in panel C) consistently with experimental observations of Hobson el al. (48).

We observe that an optimal range of SO frequencies is important to obtain the reported results. Indeed an extremely low Up state occurrence would make weaker the specific cortico-thalamo-cortical association. On the other hand a very high Up state occurrence frequency would increase the probability to randomly associate different classes of digits. Both scenarios would impair the positive effect of the sleep period on network performances.

Firing rates also underwent a differential modulation. Neurons with pre-sleep low activity became more active after sleep (and *vice versa*). Panel D of the same Figure 4.7 displays a scatter of time averaged single neuron activity during the classification task, before and after the 3000s sleep period (data are drawn from 8 simulations). The distribution of individual firing rates is rotated respect to the bisector line (red dashed line). The average difference of activity after and before sleep is positive for low values of pre-sleep firing rates and negative for high pre-sleep rates (inset in panel D). This prediction is similar to what observed by Watson et al. in 118, strengthening the biological plausibility of our model.

4.5 Model refinements

The previously described thalamo-cortical spiking model has been improved and refined with the purpose of demonstrating a link among two phenomena that we believe to be essential for the brain capability of efficient incremental learning from few examples in noisy environments (Golosio et al. 37). Grounded in two

experimental observations – the first about the effects of deep-sleep on pre- and post-sleep firing rate distributions (Watson et al. 118), the second about the combination of perceptual and contextual information in pyramidal neurons (Larkum 58) – our model joins these two ingredients. This new model, named ThaCo, alternates phases of incremental learning, classification and deep-sleep. Memories of handwritten digit examples are learned through thalamo-cortical and cortico-cortical plastic synapses. In absence of noise, the combination of contextual information with perception enables fast incremental learning. Deep-sleep becomes crucial when noisy inputs are considered. We observed in ThaCo both homeostatic and associative processes: deep-sleep fights noise in perceptual and internal knowledge and it supports of categorical association of examples belonging to the same digit class, through reinforcement of class-specific cortico-cortical synapses. The distributions of pre-sleep and post-sleep firing rates during classification change in a manner similar to those of Watson experimental observation. These changes promote energetic efficiency during recall of memories, better representation of individual memories and categories and higher classification performances.

4.5.1 Thaco main improvements

The main differences between Thaco and the previous simplified thalamo-cortical model, above described, can be principally ascribed to the trend of bringing it as near as possible to biological principles. Almost all parameters had been adjusted to reach a behaviour similar to that observed in neuro-phisiological experiments, as the ones reported by Watson et al. in 118. This was achieved with an iterative method based on simulations and comparison of results with the experimental data (mainly firing rate of neural populations during awake and deep-sleep states), that allowed to calibrate the synthetic networks parameters.

A more accurate calibration of the contextual signal impinged a collateral, really important effect: the memories are equally represented in the network, allowing for continuously cycling of incremental learning, no possible in the previous simplified

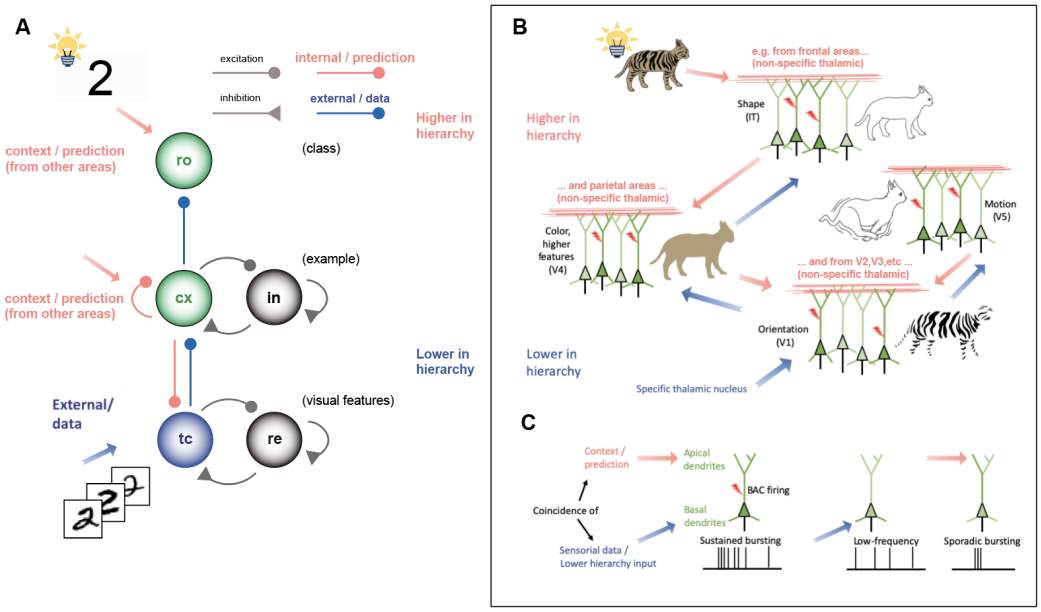

Figure 4.8. Thalamo-cortical spiking model (ThaCo). Panel A: scheme of the Thalamo-cortical spiking model (ThaCo). Input images, passed through a filter (HOG) are projected (blue arrow) to thalamic excitatory neurons (*tc*), mimiking the mechanism of the retinal *visual stimulus*. Thalamic neurons stimulate cortical excitatory neurons (*cx*) with a *perceptual feedforward* excitation (blue). Cortico-cortical and cortico-thalamic are considered as *top-down prediction* connections (red). (Red arrows - context/prediction) Currents coding for higher abstraction features incoming from other cortical areas. Cortical inhibitory neurons (*in*) arbitrate competition among cortical groups in a soft Winner-Take-All mechanism (WTA). Inhibitory reticular neurons (*re*) control the thalamic firing rate. Panel B: a cellular mechanism for associating feed-forward and feedback signals. Low-level features are encoded in primary sensory regions and this signal propagates up the visual hierarchy (e.g. striate cortex (V1) sensitive to orientation, V4 sensitive to colour, V5 sensitive to motion, and inferior temporal (IT) cortex sensitive to shapes and objects). Higher-level areas provide feedback information (context or expectation) to lower areas.The ThaCo model presented in this paper is a single area model and the contextual signal is assumed to collect during training the knowledge carried by all other areas in the hierarchy (see red arrows in panel A). Panel C: conceptual representation of the back-propagation activated calcium (BAC) firing hypothesis supporting efficient binding of features and recognition. Pyramidal neurons receiving predominantly feed-forward information are likely to fire steadily at low rates, whereas the simultaneous presence of contextual and perceptual streams changes the mode of firing to bursts (BAC firing). This coincidence mechanism is mimicked in our ThaCo model. Panels B and C are used to compare the architecture of our model with the biological principles described by Larkum in 58.Adapted from Golosio et al. 37.

model. The cycles of incremental learning demonstrated to perform better than other Artificial Intelligence algorithms (for example Knn) when the input signal is presented in absence of noise.

In addition, in this new model we studied the beneficial effects of sleep on a noisy signal, introducing a *'Salt and Pepper'* noise in the contextual signal: slow oscillations can compensate this noise through homeostasis, equalizing synaptic weights and creating beneficial associations that improve classification performance.

Improvements were also implemented on the coding of thalamic input features. A pre-processing HOG algorithm is applied to the MNIST handwritten digits images to make them suitable to be encoded into the neural circuits, as in the previous simplified thalamo-cortical model. The innovation introduced in Thaco was a fuzzy-logic inspired approach that codes the HOG output into 6 truth values, each one stimulating a specific thalamic neuron, instead of the 4 values binning previously used. This creates a smoother representation of the HOG information and reduces the loss of information, possible whenever real values are discretized.

4.5.2 Thaco architecture

ThaCo exploits the combination of contextual and perceptual signals to construct a soft Winner-Take-All mechanism (WTA) capable of fast learning from few examples in a synaptic matrix shaped by spike-timing-dependent plasticity (STDP). We set the network parameters to induce the creation of soft WTA mechanisms by emulating the organizing principle of the cortex described by Larkum as depicted in Figure 4.8.

The proposed ThaCo circuit is organized into three layers, as shown in Fig. 4.8, panel A: an input layer, the *thalamus*, which consists of an excitatory population (*tc*) whose firing rate is under the control of a *reticular* inhibitory fully-connected population (*re*); the *cortex*, consisting of an excitatory population (*cx*) and an inhibitory population (*in*), also fully connected; a *readout* layer, to which the cortex is also fully connected, composed of a subgroup of neurons associated to each class.

The phases of the training and testing protocol is described in Figure 4.9. During

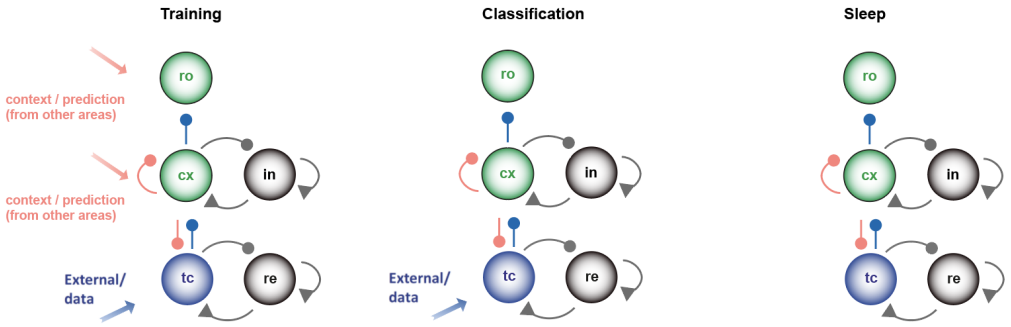

Figure 4.9. Phases in Thalamo-cortical spiking model. During training (left), the injection of *contextual signal*, plays the role of internal prediction and increases the perceptual threshold of a subset of cortical neurons. The simultaneous presence of *perceptual* and *contextual* promotes a high firing rate in such neurons, mimicking the BAC mechanism.In the classification phase (centre) the contextual signal is turned off. In the sleep phase (right), the sensory pathways are turned off, and all the activity is generated spontaneously. Adapted from Golosio et al. 37.

the training phase, for each example that have to be learned, groups of excitatory neurons in *tc* are facilitated through the presentation of a contextual signal, leading to an enhancement of connections between the cortical neurons trained over the presented example and with the subgroup of readout neurons associated to the correct class. The simultaneous stimulation from perceptual and contextual signals emulates, as a first approximation, the organizing principle of the cerebral cortex, as described by Larkum. This approach approximates the effects of the dendritic apical amplification mechanism at the cellular level. It is worth noting that this is the only stage when a category-specific (rather than an example-specific) signal is given to the network: protocols concerning this third layer are *supervised* training protocols, whereas those for the other layers can be referred as *unsupervised* training protocols. During the classification phase, on the other hand, the contextual signal is sent neither to the cortical neurons nor to the readout ones. We infer the network answer to the classification task in two different ways: first, unsupervised, taking the class of the example over which the most active subgroup of cortical neurons has been trained; second, supervised, taking the class associated to the most active subgroup of readout neurons. Specifically, the readout layer performs the integration of signals coming from the subgroups of cortical neurons trained over different

examples belonging to the same class for a more detailed representation of the learning process.

4.5.3 Comparison with biological observations

ThaCo has been calibrated to express deep-sleep-like activity and to induce modifications to the distributions of pre- and post-sleep firing rates comparable to biological measures like those carried out by Watson for an investigation of the deep-sleep effects on learning and classification. Specifically, with ThaCo we investigated several aspects and learning capabilities: 1- incremental learning from few examples; 2- resilience to noise when trained over degraded-quality examples and asked to classify corrupted images; 3- comparison with the performances of knn algorithms; 4- the ability to fight noise in the contextual signal thanks to the introduction of a biologically-plausible deep-sleep-like state, inducing beneficial homeostatic and associative synaptic effects.

As a first step, we compared the ThaCo model behaviour with biological observations on the changes of firing rate distributions in awake, sleep and post-sleep phases (see Fig. 4.10). Indeed, since one of the goals of this work is to implement a biologically-plausible model which is capable to display different *cognitive states*, the comparison with experimental outcomes become an important issue to asses its plausibility. We compare the network behaviour of the model during three simulated phases (pre-sleep awake-like, deep-sleep-like and post-sleep awake-like, see Figure 4.10) with those observed in mice by Watson et al. in 118. The previous thalamo-cortical model described in section 4.1 was already able to express the transition between states, such as sleep-like slow oscillations activity and awake-like classification, but in the improved Thaco model, we refined its parameters to make it more biologically plausible using as calibration tool the comparison with experimental observations about differential changes of firing rates.

After the validation obtained against experimental results, we demonstrate the capability of the model for incremental learning, i.e. with additional learning

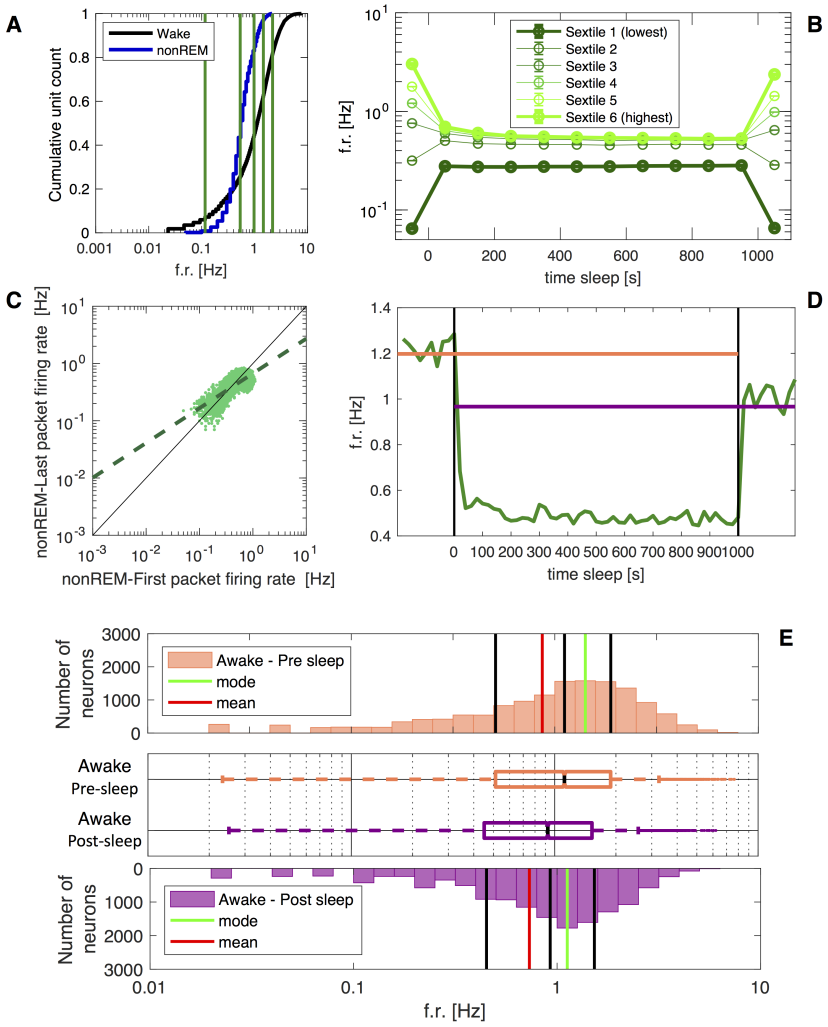

Figure 4.10. Sleep-like features. Panel A: state-wise differences of average firing rate, to be compared with Figure 2A by Watson et al. 118. Cumulative distribution of the firing rates of individual cortical neurons (log scale); note the brain-state dependent differences (colour). Vertical lines separate neurons sorted by AWAKE firing rates into six subgroups (sextiles) with an equal number of elements. Panel B: firing rate changes across sleep in each of the six groups defined by the awake firing rates, to be compared with Figure 3B by Watson et al.. High firing rate neurons show decreasing activity; low firing rate neurons do not increase their activity over sleep. Panel C: opposite modulation of neurons of different firing rates, to be compared with Figure 3D by Watson et al.. Comparison of individual neuron firing rates during the first and last packet of sleep. The regression line is significantly different from unity, showing that high and low firing rate neurons are oppositely modulated over sleep. Panel D: cortical neuron population mean firing rate changes across sleep, to be compared with Figure 3B by Watson et al.. Panel E: awake firing rate distribution of cortical neurons pre-sleep (upper plot) and plot-sleep (lower plot). Solid lines depict descriptive statistics parameters: Q1, 25% quartile; Q2, 50% quartile (median); Q3, 75% quartile. Middle plot: boxplots of the distributions. The central mark indicates the median, and the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. Adapted from Golosio et al. 37.

examples presented to the trained network. In addition, we also tested Thaco with noisy signals, obtained injecting a *'Salt and Pepper'* noise (density = 0.2) into the unprocessed MNIST images.

In conclusion, in order to be compliant with biological rhythms, we first verified that the proposed network, derived from the thalamo-cortical spiking model proposed in 4.1, is able to reproduce the experimental measures of neuronal firing rates during awake and deep-sleep states performed by Watson et al 118. Then we verified that the proposed model is capable of fast incremental learning from few examples and of alternating several learning-sleep phases. At the end we assessed its resilience when subjected to noisy perceptions with better performances compared to other common classification algorithms.

Chapter 5

Preparing the study of cognitive tasks in monkeys

Among the objectives of my activity there was the study of neural correlates in non-human primates during cognitive tasks and motor behaviour, in case of controlled answers. The presented work is a preliminary data analysis, mainly focalized on the methodology to produce the tools useful to work on electrophysiological data also in different conditions, for instance to study the effect of sleep (and more specifically, of slow waves) on the learning process and memory consolidation in behavioural tasks, which constituted a specific interest among my activities in the last years.

As a first step in this direction, to take confidence with the experimental data, I started analysing data collected by researchers of the *The Motor Control and Cognition Lab* (MCC_lab) directed by Prof. Stefano Ferraina at Sapienza University in the past years. I want to underline that the work here reported is based on data previously published and curated by the staff of the *The Motor Control and Cognition Lab*, that performed animal care, housing and experimental procedures conformed to European and Italian laws on the use of nonhuman primates in scientific research.

The recordings were acquired on two rhesus monkeys chronically implanted in the left dorsal premotor (PMd) cortex with a 96-channel microelectrode array (Blackrock Microsystems, Salt Lake City, Utah). Monkeys were tested both on transitive inference (TI) tasks and on countermanding (CM) tasks while neural

activity was recorded extracellularly. From the recorded signal, the unfiltered Local Field Potentials (uLFP), I extracted the Multi-Unit Activity (MUA) that represents an aggregate measure of the neuronal spiking activity in the local neighborhood of the recording electrode of about 200 μm, providing a good estimate of the firing rate of the neurons surrounding the electrode tip (Mattia and Del Giudice, 63). Starting from the raw recorded signal, I extracted the MUA according to the methodology set-up by M. Mattia et al. and well described in 89. For each channel/electrode, the Power Spectral Density (PSD) of the signal is computed. The MUA signal is calculated as the average power of the normalized spectra filtered in the frequency band 0.2–1.5 kHz. Then the MUA is logarithmically scaled to balance the large fluctuations due to the spikes of the neurons surrounding the single electrode. At the end, the log(MUA) time signal is smoothed using a moving average based on a sliding window of 80 ms. The MUA calculated with this methodology constitutes the basis of all the analysis described in this chapter.

One prospective focus of this research is the study of the transitions between brain states, such as from deep anaesthesia to awake state, which entail changes in the "complexity" of neural activity. On the other hand, during the performance of a task, neural activity has been observed to ‘loose complexity’ when a task-dependent decision is taken, or upon processing a sensory stimulus (often observed as a collapse of inter-trial variability). Therefore, important changes in the complexity of neural dynamics are observed both in spontaneous neural activity and in task-related activity, and suitable definitions of complexity are wanted. Among the many proposals in this direction, the dimensionality of neural activity is arguably a useful complexity measure; I focus in the following on simple dimensionality measures, as defined in 1 by Abbott, Rajan and Sompolinsky and exploited by Recanatesi in 85 and Mazzucato 66, and test them on neural activity recorded during CM and TI tasks. A summary description of these tasks is reported in section 5.1 and 5.2 respectively.

In addition, I also analyzed transitive inference experiments to find specific

neural correlates that start acting in different conditions, for instance something that can help distinguishing when the correct answer is on left or right, or if there are differences in neural correlates for different symbolic distances.

The results of these studies are reported in the following sections.

5.1 The countermanding task (CM)

The countermanding paradigm is particularly suitable for the study of the ability to suppress an impending action, or in other words of an inhibitory control of the actions (see Logan and Cowan 59 or Hanes and Schall 45). In experiments done in the MCC_lab (Pani et al. in 82 and 81), the focus of the study done using this paradigm was the inhibitory control of reaching arm movements in monkeys.

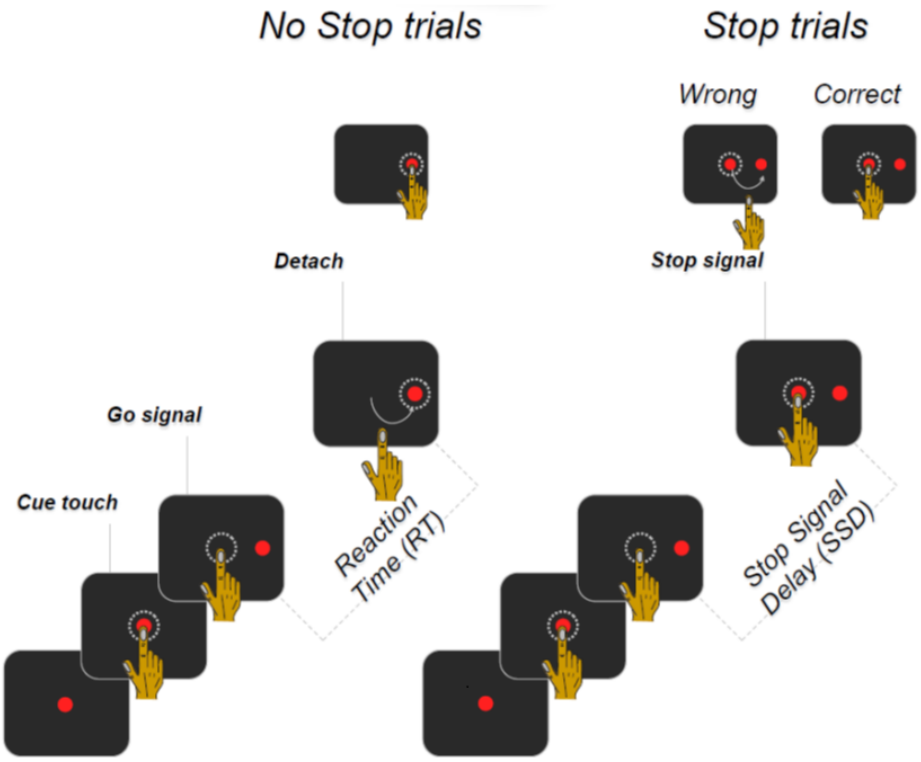

Figure 5.1. Countermanding task: the two different protocols of *no stop* trials and *stop trials* are described in the picture. Adapted from countermandings works performed at the MCC_lab.

Each session of the experiment is constituted of several trials in which a central

cue appears on a screen. The animal has to touch it until a new target appears either on the left or on the right of the same screen. When the central cue disappears (*go* signal), the monkey have to reach the new spot and receive the reward. In a fraction of trials, after a variable delay, the central cue reappears signaling a *stop* in the action execution. If the monkey aborts the movement the trial is considered correct, otherwise it is a wrong one. The protocol is described in Figure 5.1.

5.2 The transitive inference task (TI)

Transitive inference (TI) is an inductive reasoning task that allows to derive relations between items that have not been explicitly compared before. Generally speaking, the TI task demonstrates the ability to support deductions like:

$$(B > C) \wedge (C > D) \implies (B > D) \tag{5.1}$$

i.e., if item B is greater in rank of item C and item C is greater in rank of item D, then B is superior in rank to D. TI has been used in several research studies on humans and non-humans subjects, to understand various aspects of inferential problem solving (see for instance Bryant and T. Trabasso 12 and Brunamonti et al. 9, 10).

Figure 5.2 represents a sketch of the experiment set up in the MCC_lab for the recording data that I used in my thesis. At the beginning of the daily experimental session, a set of six items was randomly selected out of a set of about 80 items consisting of abstract images. The six selected stimuli were randomly organized in an always novel serial hierarchy at every daily experimental session. During the experiment, the animal is required to learn the relationship among these six elements. At the beginning, the monkey is posed in front of a screen where, after the beginning of the task, a central cue appears. The monkey touches the cue. Then two of the selected symbols appear on the screen (*target on*), one to the left, the other to the right. As soon as the central cue disappears, the monkey starts moving the arm

toward one of the two symbols (*detach*). The correct choice happens if the selected item is the higher in rank. In that case, the animal receives a reward.

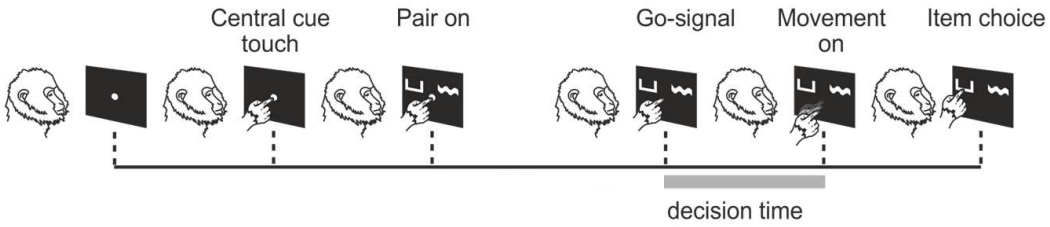

Figure 5.2. Transitive inference task: the animal is in front of a screen pressing the central cue when it appears. Then two items are proposed. At the *go* signal (cue off), the monkey selects one of the symbols and receives a reward if the selected one is greater in the established ranking. Adapted from Mione et al., 68.

.

Each session of the experiment evolves through three phases:

1. ***Learning phase***: two different procedures can be adopted: 1- a sequential learning procedure and 2- a chain-linking learning procedure. In the first one, the first pair of symbols, adjacent in rank (i.e.: with symbolic distance equal to one), is presented to the animal, that has to identify the correct one by trial and error (all correct answers are rewarded). The winning item can be positioned either on the left or on the right side of the screen. This phase goes on until the percentage of correct responses is at least 80% for that pair. Then the next pair is presented following the same procedure. Instead, in the chain-linking learning procedure, the animal is trained, with the same methodology, only on the first three items (e.g. $A > B > C$), then on the last three items ($D > E > F$) until the correct response level of 80% is reached, as in previous case. At the end the chain is closed instructing the animal to learn the missing relation among the six elements ($C > D$).

2. ***Consolidation of learning phase***: all five pairs already known by the animal are presented in random order. Note that all the pairs are always at symbolic distance one (e.g. $A - B$, $B - C$, $C - D$, etc.). This phase ends only when the monkey reaches the percentage of correct responses >60%. For the

sequential learning procedure this phase is performed after the whole sequence has been learned (i.e. at the end of the previously described learning phase), while in chain-linking learning procedure it is done just before the chain closing (i.e., before the animal is instructed to recognise the relation between the two linking items $C > D$)

3. ***Testing phase***: here both learned and novel pairs are presented, and the animal is asked to infer the hierarchy of symbols, generalizing the basic rules already acquired for consecutive pairs of symbols. The presented symbols are already known by the monkey during the learning phase, but the symbolic distance at which they are presented could range from one to the maximum allowed (i.e. "five" for the experiments I analyzed).

As a matter of fact, at the end of the experiment, the animal should be able to recognize the higher ranked symbol when compared with any other symbol on which it has been trained that day, either if it is presented on left or on right side of the screen and whichever is the symbolic distance between the two presented symbols. For more details see also the papers of Brunamonti et al. (9 and 10) that describe the exact setup used to record the data I analyzed.

5.3 Dimensionality

In neuroscience, dimensionality provides a representation of how compact the measure of the network activity could be, or, in other words, how many degrees of freedom can be used to explore a phenomenon. The activity of a huge number of neurons can be seen as a set of points in the firing rate space. The dimension of this space is in general equal to the neuronal ensemble size, but usually the effective activity can be localized in a smaller subset of them. Dimensionality reduction acts in this sense: it creates a new space that reflects the essential features of the original one, but with a reduced dimension and, in general, it helps in removing noise or information that are irrelevant to the task.

Roughly speaking, dimensionality is the minimal number of dimensions that is necessary to provide an accurate description of the neural dynamics. It exploits the redundancy of the input data finding a smaller set of new variables, each being a combination of the input ones, basically containing the same information.

A lower dimensional data representation with respect of the original one, in which the noise and other redundant information are removed, just retaining the signal of interest, can be essential in understanding hidden structures and patterns. In this sense, it can be viewed as a method for latent feature extraction. Several techniques are available for dimensionality reduction. The paper of Cunningham and Yu (23) represents a comprehensive review of many of them.

5.4 Dimensionality over trials evolution

Dimensionality could not be considered constant during the execution of the whole trial. Conceivably it changes around specific events that occur along it. For instance, in case of countermanding tasks, the *go*, *stop*, or *reward* signals are specific points that change the value of dimensionality (see 5.5.1). The same happens in transitive inference tasks (see 5.7). In this sense, the trials can be considered time series whose analysis can be used to infer properties describing the underlying structure of the system.

Principal component analysis (PCA) is one of the most simple and widely used dimensionality reduction methods, but as other many linear methods commonly used, it disregards the temporal structure hidden in the data, treating data at different time steps as independent static samples (Clark et al. 21, Recanatesi et al. 85). I searched for other approaches that could give more reliable and complete results for the calculation of this value for data displaying high variability due to their temporal structure.

One methodology used to calculate dimensionality variability along trials is described by Recanatesi in 85 and Mazzucato in 66. After the application of PCA to the data coming from the acquisition on several electrodes, the value of dimensionality

is calculated as the weighted measure of the number of axes explored by the "cloud" generated by the distribution of the neural population across many trials:

$$D = \frac{(\sum_{i=2}^{N} \lambda_i)^2}{\sum_{i=2}^{N} \lambda_i^2}. \tag{5.2}$$

where λ are the eigenvalues of the covariance matrix.

For completeness, in this work I also took into account an extension of linear analysis methodology. In particular I considered the calculation of dimensionality reduction proposed by Laio et al. (29). The authors propose a new estimator for the intrinsic dimensionality using only the distance of the first and the second nearest neighbor of each point in the sample. They affirm that the proposed method is theoretically exact in uniformly distributed datasets and provides consistent measures in general. Moreover, they say it also returns good results also in presence of high-dimensional noise and seems to require a low computational cost.

First of all, I compared the calculus of dimensionality reduction using the two approaches (Laio methodology and equation 5.2 applied to the eigenvalues of the covariance matrix obtained with PCA) on synthetic noiseless data to facilitate the comparison. The data series was built mixing two signals with different dimensionality. The first signal (S1) is a hyper-plane of dimension two, while the second signal (S2) is a hyper-sphere of dimension N. The data series is built including S1 and S2 with different percentage throughout a sigmoid function that, at the beginning, produces the full inclusion of the S1 signal without any contribution of S2, then it smoothly increases the contribution of S2 and correspondingly reduces the contribution of S1 until it fully disappears. The sigmoid function is used to weight the value of the dimension and move from the minimum to the maximum value in a smooth way, ruling the transition from lower to higher dimensionality. For the signal S2, the hyper-sphere, I used two values of dimensionality: a hyper-sphere of dimension 10 and hyper-sphere of dimension 96.

The results are reported in Figure 5.3, respectively in panel C and D. Panel A

and B give a 3D projection respectively of the hyper-plan and of a generic hyper-sphere. It is evident as the Laio methodology works well for low dimensionalities, but presents some instability for higher dimensions, that is not really desirable when the dimensionality reduction has to be applied to a signal with a starting dimension equal to 96 (the channel number of the Utah array used to acquire the electrophysiological signal during the experiments).

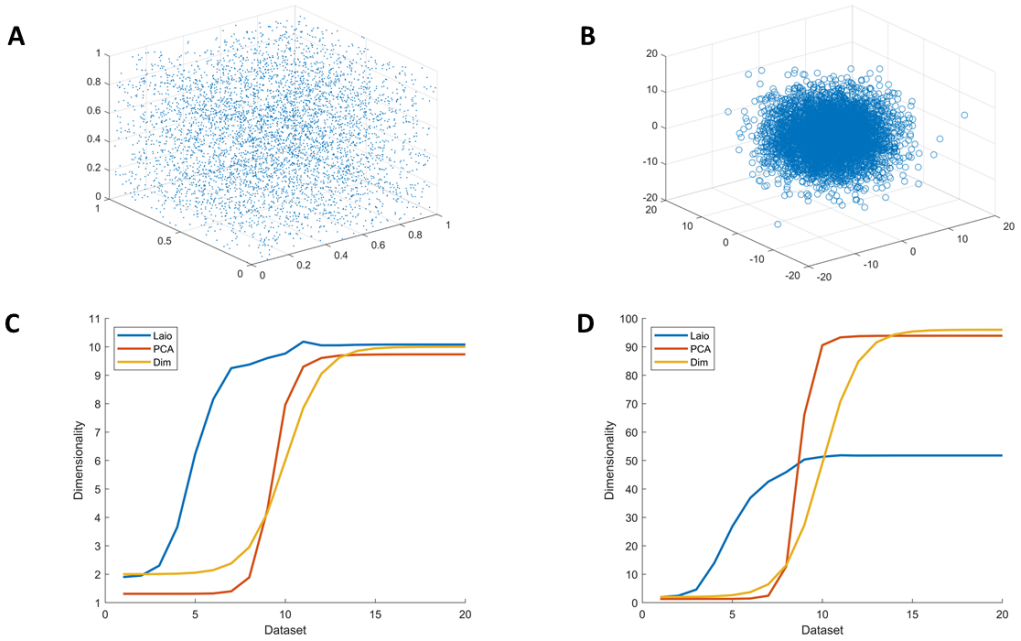

Figure 5.3. Comparison of dimensionality reduction using Laio methodology or the PCA and equation 5.2, for a different number of elements in a non-stationary time series. Panel A: 3D projection of the starting element of the non-stationary series represented by a 2-dim hyper-plan. Panel B: 3D projection of the hyper-sphere of dimension 96, representing the last element of the non-stationary series. Panel C and D: calculation of dimensionality using PCA and Laio for the time series in the range 2 to 10 (C) or 2 to 96 (D).

In a second experiment I compared the behaviour of the two methods with noisy signals variable in time. In the specific example reported in 5.4, I created a hyper-plan and after I added a white Gaussian noise to it (see panel A of 5.4). Then I moved in time randomly changing the noise at each time step. Panel B of the same figure reports the dimensionality calculated with equation 5.2 (based on the eigenvalues of the covariance matrix obtained with the PCA method, and always resulting around the value of 2, 1.98 in this case) and that calculated with Laio

method, that is highly variable and presents a mean well above the value of 2 (23.08 in this case). Both dimensionality calculations have been computed using a sliding window of size equal to 3 milliseconds.

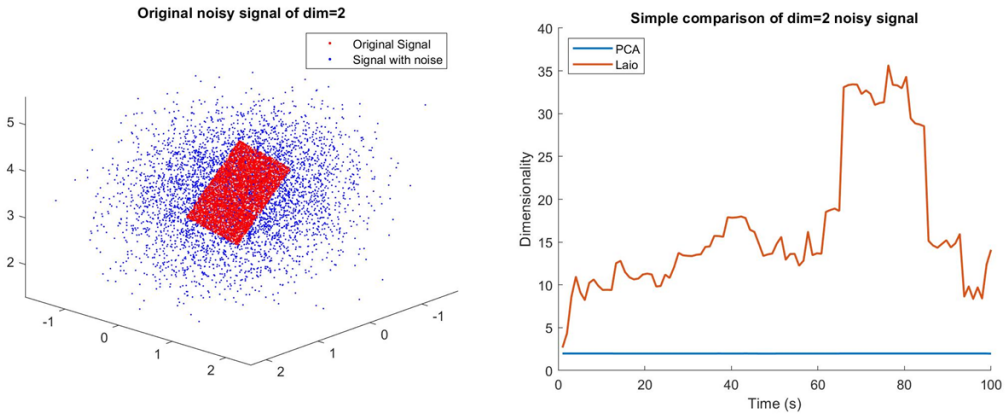

Figure 5.4. Comparison of PC and Laio methodology for a different noisy signals: Panel A: noisy signal produced adding white Gaussian noise to a hyper-plan. Panel B: results obtained applying PCA and Laio dimensionality reduction methods. In both methods, a sliding window of 3 ms has been used for the computation.

At the end I can conclude that the methodology proposed by Laio is highly sensible to the signal noise, as in the case of neural signals that generally present a high degree of noise, both for the possible problems during the acquisition phase, and for the high variability of the signal itself. Moreover, the application to an original space presenting high dimensionality does not produce stable results.

As a matter of fact, possibly due to this high sensibility to noise, the application of the Laio methodology to the transitive inference tasks, for the reduction of the original space created by the 96 electrodes of the Utah array, and the creation of a less-dimensional space for the analysis of the neural dynamics for different symbolic distances, does not produce outstanding results.

In conclusion, I decided to use the PCA and equation 5.2, with the caveat to apply it to trial-averaged signals and to work on small time windows, in order to reduce variability due to noise retaining the variance due to neural dynamics.

Table 5.1. CM trials per kind and per monkey

	Trials			
Monkey	no-stop	stop-correct	stop-wrong	Total
Monkey 1	689	87	89	865
Monkey 2	221	38	20	279

5.5 Dimensionality along trials in CM tasks

I analyzed the data coming from two countermanding experiments each one done on a single monkey. According to the applied protocol, the trials are divided in *no-stop* trials, *stop-correct* trials and *stop-wrong* trials (the one in which the monkey receives the *stop* signal but continues the arm movement). Table 5.1 summarize the number of trials available for the three kind of tasks.

In a first measurement, I analyzed only the *no stop* and *stop wrong* trials, aligning all of them to the *detach* signal. I started from the raw signal registered in PMd cortex of monkeys at the MCC_lab using the 96 channels Utah array. From this signal, I extracted the MUA spectral estimate as described in the introduction of this Chapter 5 for each channel. Then I calculate the dimensionality reduction using the PCA method.

First of all, in the new PCA space, I subtracted the MUA spatial mean, that represent the projection of the data on the first principal component (PC1), that substantially would dominate the analysis. Therefore, because I'm interested in analyzing the dimensionality of the subspace resulting orthogonal to PC1, before proceeding with the analysis I needed to remove the spatial mean of the obtained MUA signal.

The results of the dimensionality calculation for monkey 1 are reported in Figure 5.5: in this graph, only trials involving the monkey arm movement are taken into consideration, both the correct ones (*no stop*) and the wrong ones (stop wrong). In case of *no stop* trials the dimensionality starts increasing at *detach* signal and remains almost stable until reward (basically the end of task). On the contrary, in

stop wrong trials it decreases and remains at much lower values after the stop signal. The dimensionality collapse immediately after the monkey received the *go* signal represents a loss in complexity, and can be considered as a result of the fact that the decision about the movement has been taken, and the motor plan is already formed. Then, after the *touch* signal in case of correct *no stop* trials, the dimensionality rises again more or less at the values it presents at the beginning of the task.

Quite different is the condition for *stop wrong* trials, in which the decrease of this value is still visible at the reaction time, but then it remains at low values (better visible in Figure 5.6), because probably the monkey has already realized the error.

As said before, all the trials are aligned to the *detach* signal. The additional depicted lines represent the occurrence of other events happening during the trial. Considering that the different events occur with a little variability for each trial, to better show this happening, I traced larger lines, that show the first and the last occurrence of each single event. In same cases the lines are smaller (the occurrence is almost the same in all the trials), in other cases they are larger (the difference among trials is more pronounced).

In an additional analysis I considered also the *stop correct* trials, starting from a different alignment given by the *go* signal. In this case results are reported in Figure 5.6. Also in this graph, the additional significant events are reported, with the same consideration of line width as in Figure (5.5). Dimensionalitity between correct trials (*no stop* and *stop correct*) are quite similar, even if the moment in which dimensionality rises again happens in two different moments: after the *stop* signal for *stop correct* trials, and after *detach* for *no stop* ones, both representing the moment in which the motor plan can be considered concluded for that specific task. On the contrary, as in the previous plot, the case of *stop wrong* trials presents a collapses after the *detach*, reaching completely different values well below the ones of the other cases, as if the missed aborted movement is still keeping the animal in a low complexity neural state.

In conclusion, the temporal profile of the dimensionality here analyzed is very

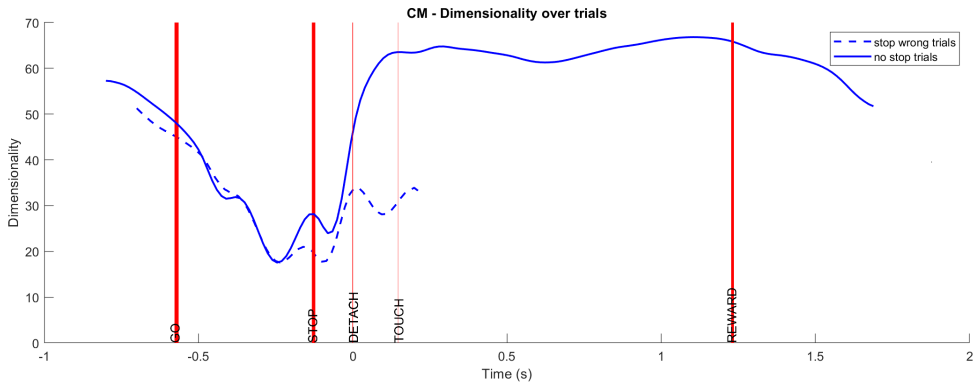

Figure 5.5. Dimensionality over trials in CM tasks. Variation of dimensionality during two different kind of tasks: the *no stop* trials and the *stop wrong* ones. Is evident how much dymensionality is reduced in case of wrong trials with respect to the others. The width variability of the lines marking the different events on the plot are due to their different happening within the various trials.

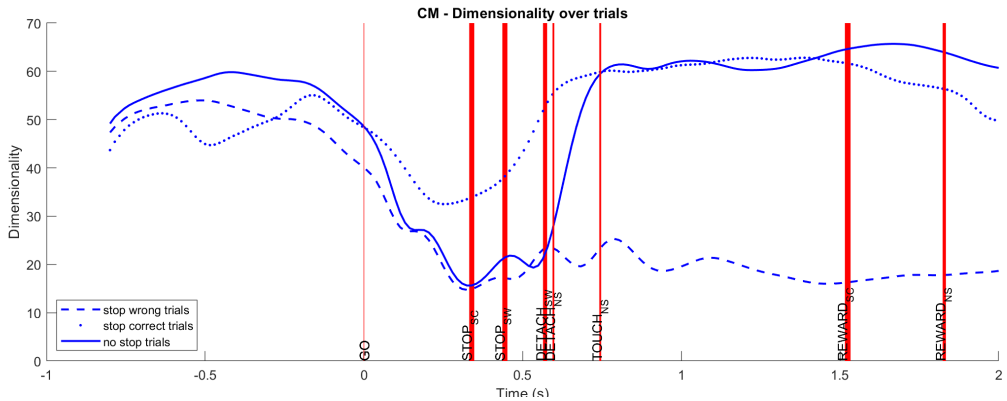

Figure 5.6. Dimensionality over trials in CM tasks. Variation of dimensionality during three different kind of tasks: *no stop*, *stop correct* and *stop wrong* tasks. In this case the alignment of the tasks is at the *go* signal. The width variability of the lines marking the different events on the plot are due to their different happening within the various trials.

similar for *no stop* and *stop wrong* trials until the movement decision is being taken, and value of dimensionality significantly decrease before the movement initiation. At the same time, the value definitively change among the two cases after the *Touch* event, when probably the monkey realize the error (it received the *stop* signal but doesn't abort the movement). On the contrary, in *stop correct* trials, the value of dimensionality is still high when the monkey receive the *stop* signal, coherently with the idea that it is able to inhibit the movement because it is not ready for the movement, yet. In other words, the movement inhibition is possible because the

motor plan has not been completed at that moment.

A further consideration is about the time of arrival of the *stop* signal: from Figure 5.6 it is evident that the *stop correct* trials receive the *stop* signal well before the *stop wrong* ones. Probably this could give the monkey more time for the decision and for the completion of the motor plan itself.

5.5.1 Dynamic of movement inhibition in CM tasks

I also analyzed the dynamic of movement inhibition in countermanding tasks, with the aim to find differences among the different kind of trials. I started from the raw signal registered in PMd cortex of monkeys at the MCC_lab using the 96 channels Utah array. As in previous analysis, I extracted from it the MUA spectral estimate for each channel. Then I calculate the dimensionality reduction using the PCA method.

In the new PCA space, after the subtraction of the MUA spatial mean, I analyzed the trajectory of the dynamics in the three different kind of countermanding tasks, as reported in Figure 5.7. It is clear from it that there is a strong similarity between the *no stop* and *stop wrong* trajectories: when in the *stop wrong* case the animal doesn't perform the required stop, it continues the movement, more or less as in the other case. Different is the dynamic of *stop correct* case: as soon as the stop signal is perceived, the movement is aborted and the dynamic in the PCA space exactly reflect this behaviour.

In figure 5.7, the different trajectories are reported with a sequence of dots, sampled at constant time. Anyhow, in the graph the dots appear at different distances one from the other (within the same kind of trials). This give a clear hint of the velocity with which the trajectory is covered during the movement: where the dots appear more distant, the velocity is higher than where the dots are denser.

An additional comment on the trajectories concerns the stereotyped PCA space where each trajectory tends. It seems that going toward the *reward* state, the dynamics of each kind of trial (*stop correct and no stop*), move toward a similar

localized PCA space, while the *stop wrong* trials have a different, but always stereotyped end point. This is confirmed also from the cluster analysis done by Baglietto et al. in collaboration with the MCC_lab.

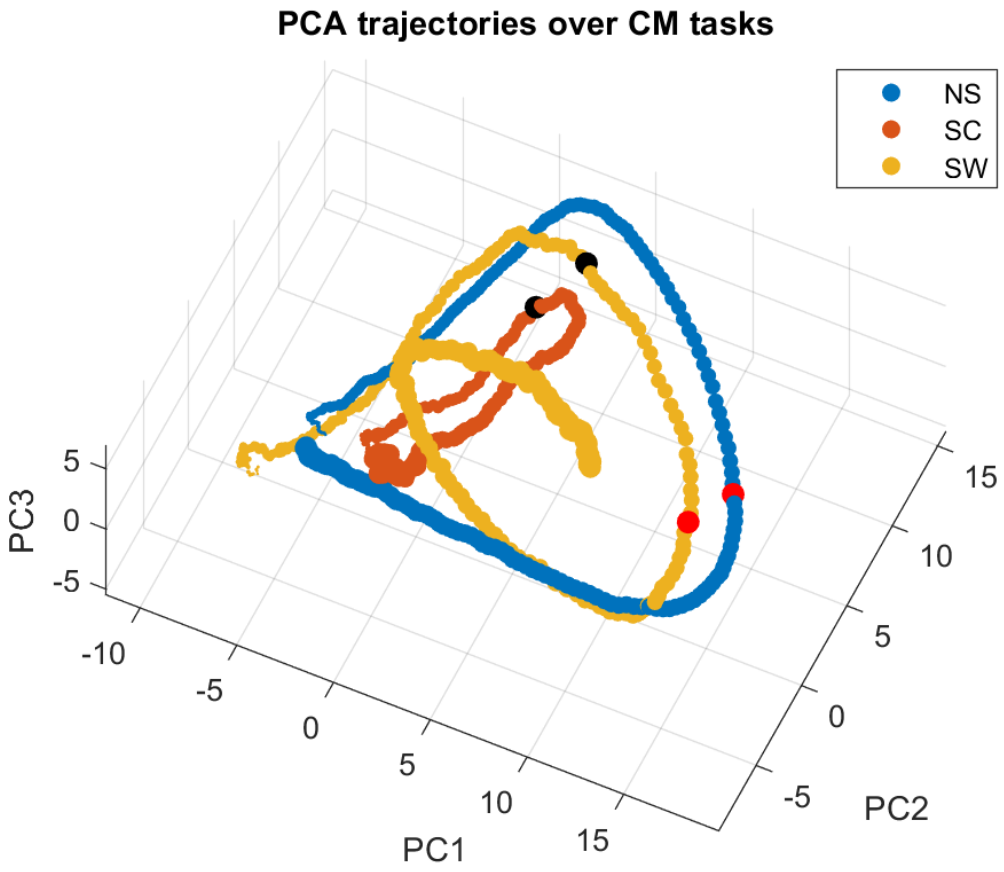

Figure 5.7. Dynamic of movement in CM tasks. PCA trajectories in case of three different CM tasks: *no stop* (NS), *stop correct* (SC) and *stop wrong* (SW). The trajectory start is placed at smaller dots and proceed toward larger ones. The black dots represent the arrival of the *stop* signal for *stop correct* and *stop wrong* trials, while the red dots represent the movement initiation for *no stop* and *stop wrong* trials.

The results described in sections 5.5 and 5.5.1 are consistent with those obtained by Ferraina, Del Giudice et al., and presented with the poster "Attractors and Flows in the Neural Dynamics of Movement Control" at the Organization for Computational Neuroscience CNS 2019 in Barcelona.

5.6 Analysis of TI tasks

To check the flexibility and general applicability of the analysis tools illustrated above, and in view of further study to be developed, I next considered the transitive inference tasks, that, from the perspective of MUA recorded signal, is just beginning to be studied in the MCC_lab. In this section I will first illustrate a summary of behavioural data collected from two monkeys performing the TI task; I will then describe the time course of the recorded MUA in different conditions of the task, and I will check to what extent it can be usefully captured by dimensionality reduction; finally (similarly to the analysis of the CM task) I will inspect the time course of a measure of the dimensionality of the neural activity during the correct and error trials.

5.6.1 Behavioural phenomenology

TI experiments were conducted on two monkeys ("Monkey 1" and "Monkey 2" in the following). For a set of six visually presented items (chosen to be equally visually salient and such as to avoid visual bias, and renewed every day) an arbitrary rank order is defined: labelling the visual stimuli as A,B,C,D,E,F, such order implies A>B>C>D>E>F. For each trial, a pair of symbols are presented on a screen. The monkey has to perform a reaching movement towards the higher-rank symbol of the pair, independently of whether it appears on the left or right on the screen. Training proceeds along two alternative protocols (for details see Mione et al. 68 and Brunamonti et al. 10, 9): either the monkey is trained on all pairs of symbols which are adjacent in the rank order, or training initially involves two 3-symbols subsequences (A,B,C; D,E,F), which are later 'chained'. The recordings analyzed in the present work refer to the *test phase* (test protocol is the same for the two training modes), in which the monkey is presented with two symbols that can be non-adjacent in the rank order (e.g. A-C, F-B), and their distance in ranking is termed 'symbolic distance' (SD), such as for instance A-D have SD=3. The successful monkey is able to infer that F<B based on the learnt rank relations between adjacent pairs, though

Table 5.2. Summary of TI tasks for right and left winning target position

	Trials				
Monkey	Correct right	Correct left	Wrong right	Wrong left	Total
Monkey 1	155	148	8	12	399
Monkey 2	167	170	26	22	450

pairs with SD>1 were never seen during training. It is intuitive and verified by data that, the difficulty of a trial decreases with SD.

Table 5.2 summarizes the number of test trials performed by each monkey, counting separately correctly and wrongly performed trials, and trials in which the winning symbol was presented on the left or on the right on the screen. Though the requested direction of movement is known to imply differences in the cortical activities associated with the preparation and execution of the movement, in the analysis we did not focus on this aspect and we lumped together right and left trials.

Table 5.3 summarizes the performances of the two monkeys for the five SDs.

From table 5.3 it appears evident that, though performances of Monkey 1 are systematically higher than those of Monkey 2, in both cases the success rate increases (while the failure rate decreases) with increasing SD, as anticipated.

The level of difficulty of the trials is also expected to affect the reaction times (defined as the difference between the time of *go* signal and the moment when the monkey detaches the finger from the screen to reach the chosen target - *detach* event); also, according to standard evidence in decision making experiments, reaction times are expected to be different between correct and wrong trials. This is in fact observed, as reported by Mione et al. in 68.

5.6.2 Analysis of MUA signal

I analyzed the MUA signal in all different cases reported in table 5.3. For each monkey I analyzed both correct and wrong trials, as well as different time alignments of the trials, taking either the *target on* or the *detach* events as reference times. Of course, once aligned to one event, trials show a spread in time of the other events.

Table 5.3. Summary of TI tasks for each SD

	SD1		
Monkey	Correct	Wrong	% wrong on total
Monkey 1	52	9	15%
Monkey 2	57	19	25%
	SD2		
Monkey	Correct	Wrong	% wrong on total
Monkey 1	64	4	6%
Monkey 2	56	22	28%
	SD3		
Monkey	Correct	Wrong	% wrong on total
Monkey 1	61	3	5%
Monkey 2	76	4	5%
	SD4		
Monkey	Correct	Wrong	% wrong on total
Monkey 1	60	2	3%
Monkey 2	73	0	0%
	SD5		
Monkey	Correct	Wrong	% wrong on total
Monkey 1	66	0	0%
Monkey 2	75	1	1%

For each case I calculated the time course of the mean MUA over trials; I also represented the time course of the MUA as a rasterplot of color-coded MUA values across trials (sorted by reaction times), separately for each SD. Correct trials are aligned to *target on* in Figure 5.8, and to *detach* in Figure 5.9.

Although the time course of the average MUA (top-left panel in Figure 5.8) shows little variations for different SDs, it can be noted that the MUA ramps across the decision period are shifted in time for different SDs, easier trials (larger SD) being typically anticipated with respect to harder trials (smaller SD). This is tentatively interpreted as a neural correlate of the different reaction times observed for different SDs. From the MUA rasters it can be seen that, across the time interval from the *go* signal to the completion of the movement, the MUA peak visible in the plot of average MUA is in fact differently spread across trials depending on SD: the upward and downward 'fronts' of the peak are pretty aligned in time for larger SDs, while they exhibit more (skewed) variability for smaller SDs. This is consistent with the intuition that the time needed for the planning and execution of the movement in easily decided cases is more reproducible across trials.

In case of wrong trials, the much smaller number of cases makes of course the plots noisy and of hard interpretation, providing a very few statistics to be analyzed. The only aspect that clearly emerge from the analysis of the data, is that the reaction time for wrong trials (0.46 s) is greater than that for correct ones (0.40 s) with a difference of about 60 ms, showing that the reaction time is correlated with the performance, as a confirmation of what reported by Mione et al. in 68.

5.6.3 PCA trajectories over TI tasks

Analogously to what I reported for the analysis of the CM task, as a preparation for the analysis of dimensionality of the neural activity, I turned to the representation of the MUA dynamics in the PCA space.

In the (by now standard) application of PCA dimensional reduction to the analysis of neural data, a few choices depend on the quantity and quality of available

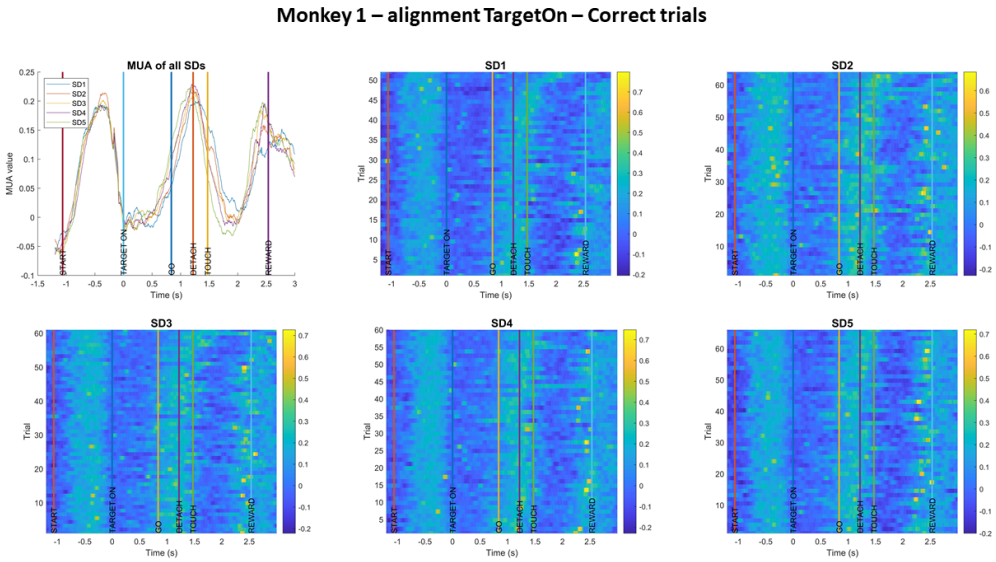

Figure 5.8. Mean MUA for all SDs and rasterplots for each SD - Target On alignment - Correct trials. The mean MUA representation shows the MUA signal for the different SDs. A peak is evident at the *detach* moment when the animal start the movement. Also remarkable is the ramp-up of the MUA immediately before the *detach*: the lines are ordered according to the SD value, hint that the velocity in reaching the *detach* moment is inversely proportional to the SD. Also in the rasterplots, even if less evident, the reaction time is more remarkable for smaller SDs.

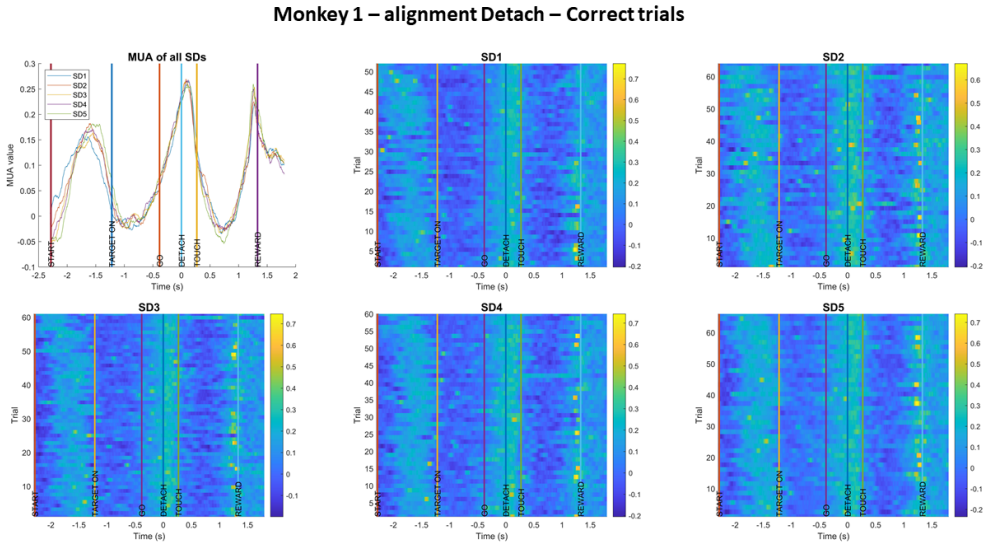

Figure 5.9. Mean MUA for all SDs and rasterplots for each SD - Detach alignment - Correct trials. This panel is analogous to the previous one, but here the alignment is to the *detach* event. Of course the ramp up of the MUA before the *detach* is less appreciable, but the reaction time variation for the different SDs is still notable.

data, and the scientific question asked. In the case examined here (both for the CM task and for the TI task) the feature space is defined by the recording channels, while for the samples different choices are available: for example, all trials (or all trials of a given type) can be concatenated so that for each channel the sample data are the MUA across all the chosen time points of all trials; or, averaging across trials, the sample data are then the average MUA across the chosen trial time interval. In the first case, the variability accounted for by PCA include temporal, inter-trial variations, which are averaged in the second case; on the other hand, the second option can fruitfully eliminate noise, deemed detrimental for the analysis. Once a choice is made, a commonly used representation is, as in the CM case, obtained by projecting the data, in time, onto the first three PC in order to draw 3-d trajectories in the PCA space.

In the case studied here, according to the first approach the PCA components accounting for at least the 80% of the variability in the neuronal activity turn out to be 12 in case of Monkey 1 and 34 for Monkey 2. Indeed, the first three PC explain 63% of the variance in case of Monkey 2, and only the 34% for Monkey 1. Instead, averaging the signal over trials, the first three PCA components account for more than 86% for both monkeys. Therefore, for the 3-D trajectories to 'live' in the space maximally accounting for data variability, I decided to average the signal on the trials.

In the study of the PCA trajectories for transitive inference tasks, the main discriminant factors that could give some hints were related to the analysis for different symbolic distances and the comparison between correct and wrong trials, for the same SD. To obtain these plots, I calculated the PCA over the whole set of trials under analysis (for instance, in case of different SDs I calculated the PCA considering the space given by all five symbolic distances, obviously separated for each monkey). In this way I obtained a sort of normalization in the PCA space, common to all the SD of that specific animal. Then I projected the MUA signal over the PCA components and plotted a 3d graph of the obtained trajectories using

only the first three components. Data were smoothed by a Gaussian filter of 100 ms and then z-scored, to obtain a sort of normalization, in which the data set has mean 0 and standard deviation 1, even if it still retains the shape properties of the original data set.

The plots of the trajectories in all the subsequent figures are done using the *target on* event as task alignment and as starting point of the trajectories (black dot). The trajectories course stops one second after the occurring of the *detach* signal represented by a green dot. The *go* signal is also indicated, using a red dot.

Figure 5.10 reports the plot of all the trajectories (one for each SD) for correct and wrong trials, for Monkey 1 and Monkey 2. It is evident the regularity and homogeneity of the lines in correct trials, while the wrong ones do not have a common trace. In addition, for the wrong cases, the trajectories became more variable and unpredictable for growing SDs. This is mainly due to the fact that the number of wrong trials for higher SDs is really small, sometimes between 1 and 3 trials, and it cannot be mediated as for correct ones.

Figures 5.11 and 5.12 report instead the comparison of correct and wrong trials for each SD, respectively for Monkey 1 and Monkey 2. An aspect that emerges from all the comparisons, is that the wrong trials start diverging from the correct ones since the beginning, and increase with the SD: the more SD is high in rank (4 or 5 in the studied cases), the more trajectories differ. For the same reason exposed before, statistic speculation on these cases with really small number of trials are very difficult.

5.7 Dimensionality along trials in TI tasks

Analogously to what done on CM tasks, I used the same methodology to calculate the dimensionality expressed by transitive inference tasks, according to equation 5.2. The used data are those analyzed in the previous sections, recorded during the test phase of TI tasks both on Monkey 1 and on Monkey 2 (see section 5.2 for details).

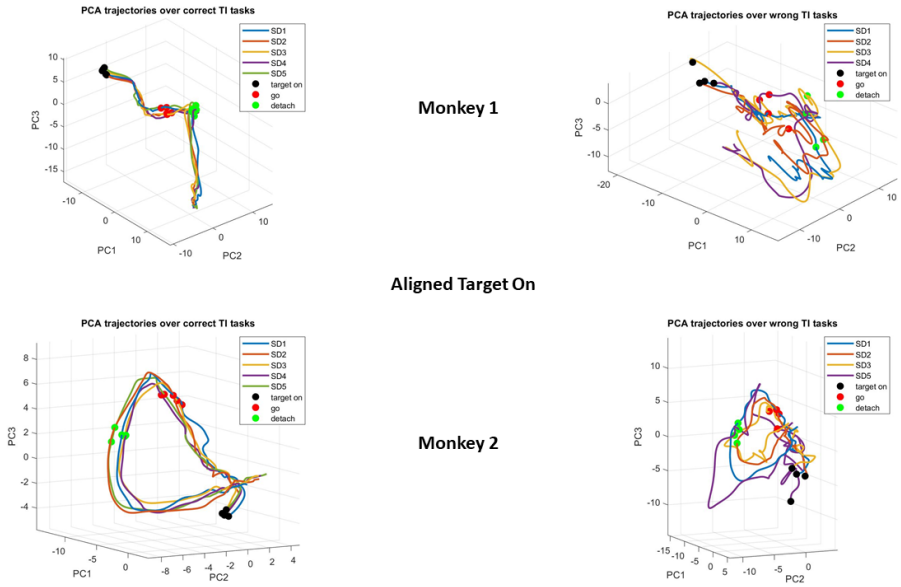

Figure 5.10. TI PCA trajectories of correct and wrong trials for the two monkeys. The trials starts at the *target on* event (black dot) and terminate one second after the *detach* (green dot). In addition also the *go* signal is reported (red dot)

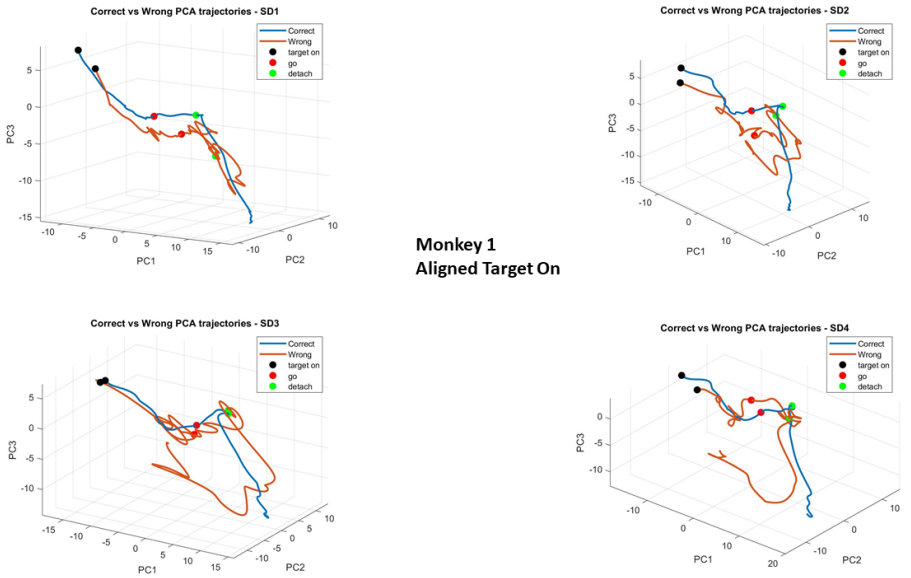

Figure 5.11. Comparison of correct vs wrong TI PCA trajectories for each SD in Monkey 1. The trials starts at the *target on* event (black dot) and terminate one second after the *detach* (green dot). In addition, also the *go* signal is reported (red dot)

The tasks are aligned to the *target on* signal and the dimensionality is calculated with 5.2, using a sliding window of 50 ms, with an overlap between a window and the following of 25 ms. For this analysis, PCA was computed on the concatenated

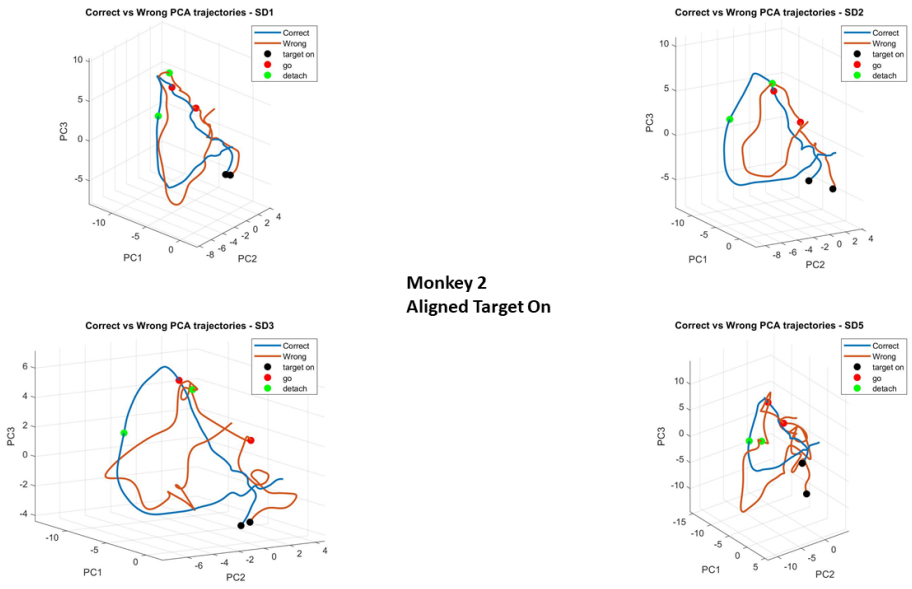

Figure 5.12. Comparison of correct vs wrong TI PCA trajectories for each SD in Monkey 2. The trials starts at the *target on* event (black dot) and terminate one second after the *detach* (green dot). In addition, also the *go* signal is reported (red dot)

trials, no average across trials was performed.

The dimensionality for the two monkeys is very different during the trial time course. One of the few common points is that a collapse of dimensionality can be observed close to *touch*, at the end of the movement. Moreover, in the time interval of interest, dimensionality in wrong trials is lower compared to that of correct ones; at the moment one can only speculate that correct trials imply the simultaneous, ongoing processing of alternatives the monkey has difficulty in choosing from, and this situation can plausibly entail higher complexity/dimensionality.

Figure 5.14 can be considered an expansion of the previous one, underlining the details for each SD. No relevant features can be extracted from it, but can be considered a complement to the previous figure.

In conclusion, the behaviour of the two monkeys is quite different under several aspect, as the previous showed analysis demonstrated. It would be interesting, during a future expansion of this work, to verify if this diversity can be attributed to differences that can be found session per session also with the same monkey, or if they are the hint of an intrinsic systematic difference in the kind of cortical

processing schema observable in monkeys.

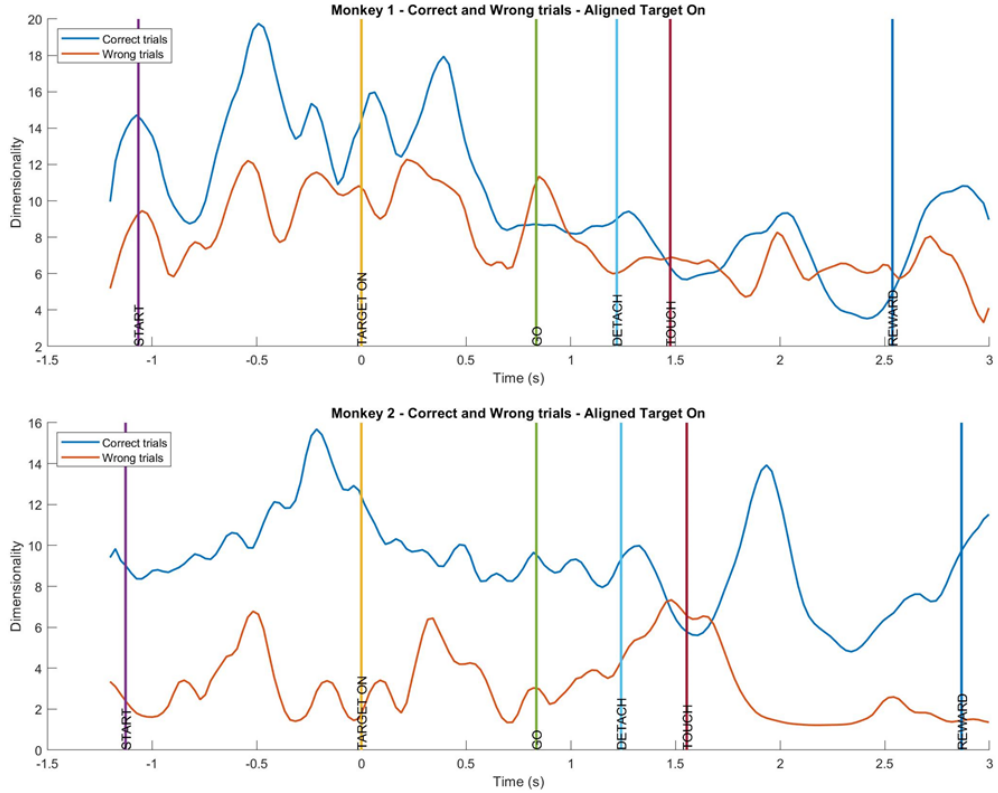

Figure 5.13. Dimensionality over trials in TI tasks. The figure shows the trend of dimensionality for correct and wrong trials, both for Monkey 1 and Monkey 2. The presented alignment is to the *target on* event. In both monkeys the wrong trials dimensionality remains well below the dimensionality of correct ones, symptom that the complexity of the neural representation is more elevated when the animal correctly complete the task.

5.8 Preliminary analysis on the learning phase

Before concluding this activity, I started a very preliminary analysis on the learning phase applied to the two animals, Monkey 1 and Monkey 2. As described in section 5.2, the animal must learn the predefined ranking order of six randomly chosen items that are showed on a touch-screen, both on right or left side. Naming these items with alphabetical letters, just to simplify the comprehension of the task design, the order should for instance be A > B > C > D > E > F. In the two tasks I analyzed, the chain-linking learning procedure had been used to instruct the animal to recognize the items and their ranking order. The experiment is structured in the

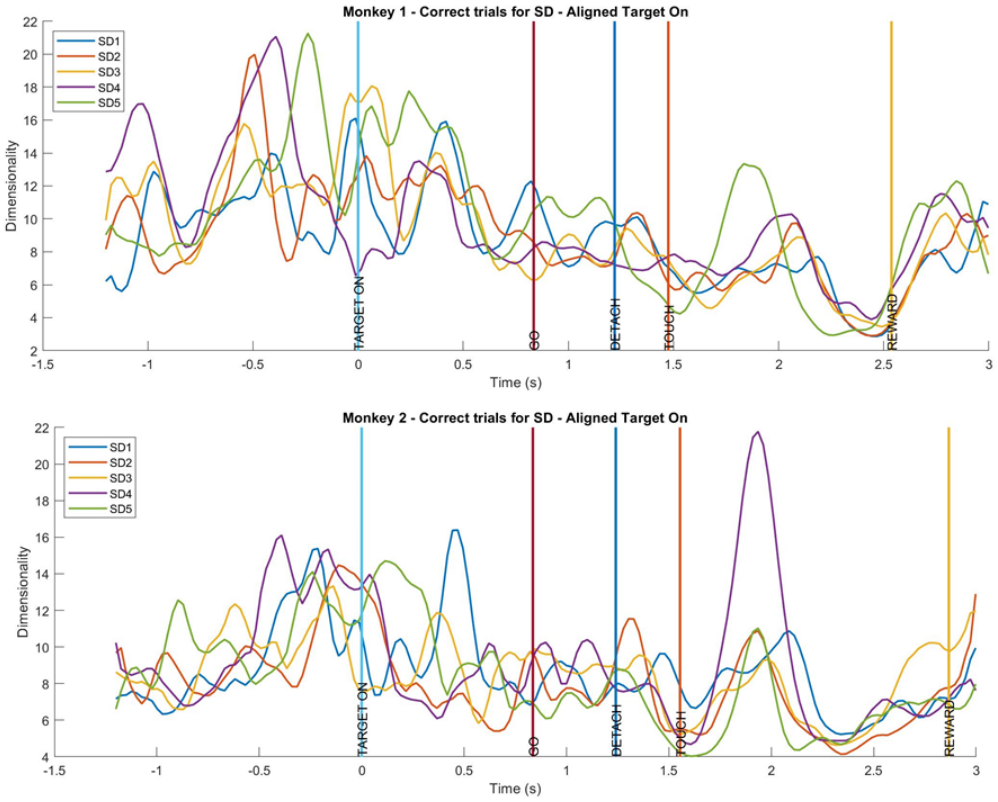

Figure 5.14. Dimensionality over correct trials in TI tasks calculated for each SD. In this picture the detail about the different SDs is reported. Also in this case the alignment is to the *target on* event. No specific differences among SDs can be extrapolated from this plot.

following way:

1- all the adjacent pairs of symbols are presented to the animal, with the winning target both on right or on left side of the screen, until it reaches about 80% of performances: at the beginning only the first pair (A-B), then the second (B-C), and so on, until the required level of performance is reached; the only pair never presented is the central one (C-D);

2- the already known pairs are presented in random order and, again, on right or left side of the screen, until about 60% of performances is reached;

3- the central pair is presented (C-D) until also in this case the required level of performance is reached.

I focused on this last phase of learning, looking for evidence in the MUA profile analysis accounting for two specific cases: a comparison of the signal when the

monkey performs a correct or a wrong trial, but always moving to the same direction; a comparison when the monkey performs correct trials with a movement toward right or toward left. Figure 5.15 is related to this two questions using the results acquired on Monkey 1.

More specifically, panel A shows the comparison between correct and wrong tasks, where the monkey always moves toward left direction. The two signals start diverging immediately after the target presentation, and the spread become more prominent after the *detach* event, sign that the error causes a change in the MUA profile with respect to the correct phase. Panel B shows instead the comparison between correct trials that entail or right or left movements. In this case the signals are quite similar, symptom that the direction of movement has not a relevant impact on the MUA profile.

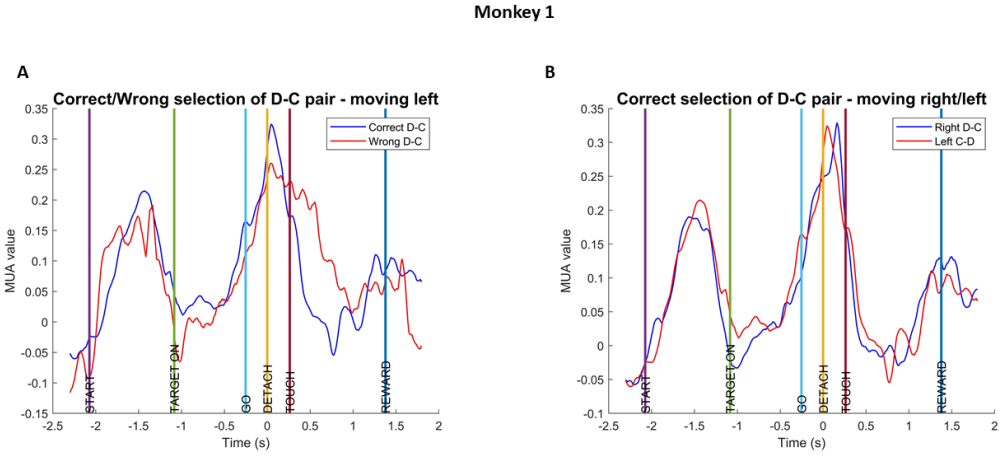

Figure 5.15. Analysis of learning phase - CD pair. Panel A reports the MUA signal calculated for wrong and correct trials, where the movement is always done toward left. In panel B the comparison between movements toward right and toward left for correct trials is reported.

During the analysis of this learning phase I noticed a strange phenomenon, mainly present in Monkey 2: around the target presentation, strong oscillations were clearly visible in the MUA signal. I then decided to analyze all the trials clustered

according to the following three conditions: all trials preceded by a correct one, all those preceded by a wrong one, all those preceded by two wrong ones. In Monkey 2, all the trials preceded by a wrong trial, even if sometimes they are concluded in the correct way, present oscillations. Of course, this happens also for those preceded by two wrong trials. Instead, the ones preceded by correct ones (independently if they are correct or wrong) do not present this phenomenon. On Monkey 1 this kind of oscillations are hardly visible, just before the target presentation.

At the moment, the reason of these oscillations is not clear, I can only give an interpretation, absolutely preliminary, that it could be a behavioural reflex on the MUA signal, strictly connected with the subject. Of course, this topic, such as the analysis on the learning phase, needs to be examined more in depth, first of all using a larger datasets, but also using different approaches (for instance, clustering analyses) at the moment well beyond the possibility of my work.

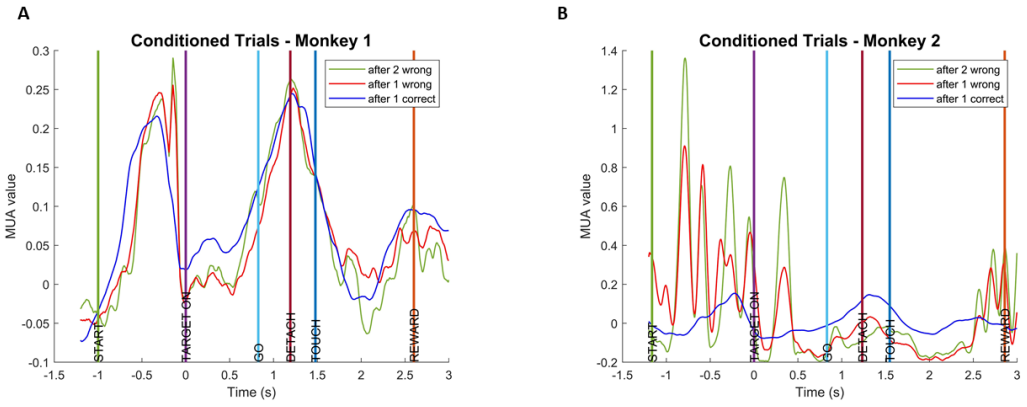

Figure 5.16. Analysis of learning phase - Conditioned trials. This figure reports the MUA signal for Monkey 1 (panel A) and Monkey 2 (panel B) calculated for all the trials of learning phase, both correct and wrong, collected using the following three conditions: all trials preceded by a correct one, all those preceded by a wrong one, all those preceded by two wrong ones. Different behaviour in the two monkeys are evident in the phase preceding the target presentation, with Monkey 2 MUA subjected to strong oscillations in case of trials preceded by errors.

5.9 The COVID-19 impact on data acquisition

The impact of the SARS-CoV-2 (Covid-19) coronavirus pandemic on my research activity was not negligible. The Covid-19 has slowed a large part of laboratory activity mainly for what concern data collection, face to face discussion with colleagues, students and supervisors. Moreover, and it is not a minor point, there was a strong decrease of academic conferences, symposia and meetings, notwithstanding a few of them took place remotely as part of videoconferences.

In my specific case, if the course of my last PhD year would have been normal, it is possible that some additional recording could be available for the analysis before the end of this thesis and could be included in my work. Of course, it is not to be excluded that some of them could be structured in order to study the effect of sleep on monkey behavioural tasks.

5.10 Future work: analysis of sleep effects on complex learning tasks on monkeys

Up to now I studied the effects of sleep on simulated simple tasks. Next step will be that of study these effect on data recorded during more complex tasks in animals. Specific experiments should be organized to confirm (or reject) what already find in simulations.

One of the first points is surely the recording on naturally asleep (or if no possible anesthetized) monkey and the analysis of the data using the SWAP analysis pipeline described in Chapter 2. I expect to find slow waves parameters like the ones recorded on mouse and rodents. This would be mainly a confirmation of the flexibility of our analysis tool, also because it has been never tested on animals with a more complex brain than mice.

More interesting, instead, would be the recording of different sessions of the same animal, on different days, to check the effect of sleep on firing rate. Using the transitive inference task, for instance, with the same elements in two subsequent tests interleaved by a sleep period, I expect that the firing rate should decrease

during the trials, while performances should increase. Moreover, also the synaptic weight should reduce, but it could be more difficult to be recorded.

For the moment, the recording during sleep has not been possible because it requires a wireless implant on the animal, to let the monkey free during the sleep phase, at the same time allowing the acquisition. The second type of experiment is instead more feasible, but not possible at the moment of the thesis, because the entire experiment set-up was not available. As soon as possible, an analysis in this direction would be for me of great interest.

Chapter 6

Conclusions

In this work I studied the effect of deep sleep slow oscillation on learning and cognition. Sleep is surely recognized as a fundamental element in this kind of activity, but till now we can affirm that a complete understanding of its functions and of the underlying mechanisms is still lacking.

In this study, I used models that try to exploit this mechanism, with the aim of comparing the results of simulated interaction of slow waves with brain states. In details two models were built, a large-scale model, able to reproduce the behaviour of the phenomena on a portion of cerebral cortex at biological resolution, and a simplified thalamo-cortical model that is more focused on the effects of slow waves on learning and cognition.

For what concern the first model, thanks to a preparatory work done in the past years, I was able to develop a methodology that allows for large-scale high-resolution spiking simulation. The result was a parallel distributed neural simulator, with emphasis on the robustness of its performance and scaling with respect to quite different collective dynamical regimes. This simulation engine has been used to simulate large-scale networks including up to 46 million point-like spiking neurons interconnected by 70 billion instantaneous current synapses (79). The final goal was that to be able of execute accurate simulations of SWA and AW-like states that could provide useful models to the scientific community for the study of cortical rhythms and their interaction with consciousness and learning. Specific accelerations

and tricks of simulations were applied to provide a fast simulation methodology for the study of this occurrence.

The main result obtained with this model was the simulation of the whole cortical mouse hemisphere activity starting from connectivity and parameters inferred from experimental data of wide-field optical imaging:

https://github.com/epastorelli/MouseHemisphereMovie.

The experimental data was recorded at the LENS laboratory (European Laboratory for Non-Linear Spectroscopy) by means of the GECIs methodology that enables the visualization of fluctuations in calcium concentration, considered an indirect reporter of neuronal spiking activity.

The results obtained with this kind of simulation were analyzed with the SWAP analysis pipeline, to extract the main features of the synthetic slow wave activity and compare them with the same observables extracted from experimental data recordings. The comparison evidenced a qualitative similarity, but not yet a quantitative match. This means that further iterations on parameter extraction are needed to better calibrate simulations.

The second model I worked on was a simplified thalamo-cortical spiking model built to test the role and the mechanisms of the occurrence of SO in a network of spiking neurons previously trained to learn and recall images extracted from the MNIST dataset. Several versions of this model are available. The last and more complete one exploits the combination of context and perception and is based on a soft winner-take-all circuit of excitatory and inhibitory spiking neurons. The simultaneous stimulation from perceptual signals and contextual signals emulates, as a first approximation, the organizing principle of the cerebral cortex, as described by Larkum (58). This approach approximates the effects of the dendritic apical amplification mechanism at the cellular level. The model is compliant with biological rhythms, as resulted from the comparison with the experimental measures of neuronal firing rates during awake and deep-sleep states performed by Watson et al. (118).

The proposed model is capable of fast incremental learning from few examples

(its performances are comparable to those expressed by Knn, of rank increasing with the number of examples) and of alternating several learning-sleep phases; moreover, it demonstrates resilience when subjected to noisy perceptions with better performances than Knn algorithms, known to be the fastest and more accurate in this kind of tasks. The crucial ingredient, which mostly differentiates this approach from other works, is the introduction of a contextual signal which drives the training procedure, making it similar to a target-based approach, enabling huge advantages in terms of training velocity and precision. Such mechanism, as said before inspired by the work done by Larkum, suggests that the activity of a neuron is amplified when it receives a coincidence of signals from both lower and higher levels of abstraction. This allows the recruitment of new neurons to learn novel examples through the incremental building of a soft-WTA mechanism.

Moreover, the predictions of our model are also a first step toward the reconciliation of recent experimental observations about both an average synaptic down-scaling effect (the SHY, synaptic homeostatic hypothesis by Tononi and Cirelli 112) and a differential modulation of firing rates (Watson et al. in 118) induced by deep-sleep, which is believed to be a default state mode for the cortex (Sanchez-Vives, et al. in 92). In this sense, this work hints to a careful balance between architectural abstraction and experimental observations as a valid methodology for the description of brain mechanisms and of their links with cognitive functions.

All this work should be complemented with the analysis of experimental data recorded on non-Human primates in the *The Motor Control and Cognition Lab* (MCC_lab) directed by Prof. Stefano Ferraina at Sapienza University, that could confirm or reject what found on-silico. Up to the end of this activity, I could only analyze experimental data recorded during awake state in monkeys, without the possibility to understand the sleep, or better the slow wave effects on the learning process. In this sense, the last part of the presented work can be considered a methodological section in which I had the possibility to set up the instruments needed to work on neurophysiological data, to be later applied when data including

sleep effects on learning in behavioural tasks will be available for analysis.

The work done on these recordings was anyhow very interesting and definitely useful to introduce me into the world of electro-physiology. In this framework I stated working on a measure that could give hints on the representation of the neural activity, the dimensionality. As extensively described by Mazzucato et al. in 66, dimensionality is able to capture the essential information of the neural network circuits and the salient features about the relation among neurons. A high value of dimensionality represents a low correlations among the neural ensemble that can be explained as a widespread firing rate of all the neurons under analysis. On the contrary, a low dimensionality value can be seen as a collapse in this space of correlation, giving a hint of ensemble activity strongly uniform, generally in coincidence of specific stimuli or activities performed. As a matter of facts, I retrieved similar results analyzing the data coming from the MCC_lab.

My activity on slow waves and their interaction with learning and memory consolidation will proceed also after the end of this PhD Thesis work. I hope I could also work on newly acquired data on CM or TI experiments, when there will be the possibility to study recordings involving sleep. It will be interesting to verify if the results obtained with models and simulations could find some correspondence on the experimental data, as I could do only for mice, up to now. Moreover, with new sleep acquisitions on monkeys, there will also be the possibility to use the SWAP analysis pipeline, with the opportune extensions of the already existing modules, with the hope to bring a contribution in enlightening this intriguing phenomenon.

Chapter 7

Acknowledgment

This work was carried out with support from the European Union Horizon 2020 Research and Innovation program under the FET Flagship Human Brain Project (grant agreement SGA3 n. 945539 and grant agreement SGA2 n. 785907) I Want to thank for their support also all the researchers of *The INFN HBP_APE Lab* and of *The Motor Control and Cognition Lab* (MCC_lab) directed by Prof. Stefano Ferraina at Sapienza University. A special thanks to the professors that supported and helped me in this activity: Prof. Stefano Ferraina, Prof. Paolo Del Giudice and Prof. Pier Stanislao Paolucci.

Chapter 8

Publications and dissemination of the activity

1. E. Pastorelli et al., "*Scaling of a Large-Scale Simulation of Synchronous Slow-Wave and Asynchronous Awake-Like Activity of a Cortical Model With Long-Range Interconnections*", Frontiers in Systems Neuroscience, Vol. 13, pages 33, 2019, https://www.frontiersin.org/article/10.3389/fnsys.2019.00033, 10.3389/fnsys.2019.00033

2. F. Simula et al., "*Real-Time Cortical Simulations: Energy and Interconnect Scaling on Distributed Systems,*" 2019 27th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP), Pavia, Italy, 2019, pp. 283-290, doi: 10.1109/EMPDP.2019.8671627

3. C. Capone et al. "*Sleep-like slow oscillations improve visual classification through synaptic homeostasis and memory association in a thalamo-cortical model.*" Sci Rep 9, 8990 (2019). https://doi.org/10.1038/s41598-019-45525-0

4. E. Pastorelli et al., "*Gaussian and Exponential Lateral Connectivity on Distributed Spiking Neural Network Simulation.*", 2018 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP), Cambridge, 2018, pp. 658-665, doi: 10.1109/PDP2018.2018.00110.

5. B. Golosio et al., "*Thalamo-cortical spiking model of incremental learning*

combining perception, context and NREM-sleep-mediated noise-resilience", arXiv:2003.11859v1, formally accepted for publication in PLOS Computational Biology

6. R. Gutzen et al., "*Building adaptable and reusable pipelines for investigating the features of slow cortical rhythms across scales, methods, and species*", 2020, Bernstein Conference, online, Germany, 29 Sep 2020 - 29 Oct 2020, doi: 10.12751/NNCN.BC2020.0030 (Preprint)

7. B. Golosio et al., "*A new GPU library for fast simulation of large-scale networks of spiking neurons*" arXiv:2007.14236 (Preprint)

8. R. Gutzen et al., "*Modular SWAP analysis pipeline from Bernstein Conference*", (in progress)

9. Models:

   https://kg.ebrains.eu/search/instances/Model/cea3e597-2fbd-4022-bbfc-a64b9fa49d68

   https://kg.ebrains.eu/search/instances/Model/97670076281ccbdc38ea2c2d76a64e64

   https://kg.ebrains.eu/search/instances/Model/fa08c511f9b444922b0975f538b10abd

Bibliography

[1] L.F. Abbott, K. Rajan, and H. Sompolinsky. (2011). “Interactions between intrinsic and stimulus-evoked activity in recurrent neural networks,” in The Dynamic Brain: An Exploration of Neuronal Variability and its Functional Significance, eds D. L. Glanzman and M. Ding (New York, NY: Oxford University Press), 65–82.

[2] A.R. Adamantidis et al. (2019) "Oscillating circuitries in the sleeping brain." Nature Reviews Neuroscience 1-17. doi: 10.1038/s41583-019-0223-4

[3] M.T. Alkire et al. (2008) "Consciousness and anesthesia." Science 322.5903: 876-880. doi: 10.1126/science.1149213

[4] P. Barone et al. (2000) “Laminar distribution of neurons in extrastriate areas projecting to visual areas v1 and v4 correlates with the hierarchical rank and indicates the operation of a distance rule,” The Journal of neuroscience : the official journal of the Society for Neuroscience 20, 3263–81

[5] M. Bazhenov et al. (2002). "Model of thalamocortical slow-wave sleep oscillations and transitions to activated states." J. Neurosci. 22, 8691–8704. doi: 10.1523/JNEUROSCI.22-19-08691.2002

[6] C. Boucsein et al. (2011) “Beyond the cortical column: Abundance and physiology of horizontal connections imply a strong role for inputs from the surround,”, Frontiers in Neuroscience, vol. 5, p. 32.

[7] R. Brette and W. Gerstner. (2005). Adaptive Exponential Integrate-And-Fire Model As An Effective Description Of Neuronal Activity. Journal of neurophysiology. 94. 3637-42. 10.1152/jn.00686.2005.

[8] R. Brette and D.F.M. Goodman. (2012) "Simulating spiking neural networks on GPU." Computation in Neural Systems, 23(4):167–182.

[9] E. Brunamonti et al. (2011) "Gaze modulates non-propositional reasoning: further evidence for spatial representation of reasoning premises." Neuroscience 173:110 –115. doi: https://doi.org/10.1016/j.neuroscience.2010.11.011.

[10] E. Brunamonti et al. (2016). "Neuronal modulation in the prefrontal cortex in a transitive inference task: Evidence of neuronal correlates of mental schema management". Journal of Neuroscience 27 January 2016, 36 (4) 1223-1236; doi: https://doi.org/10.1523/JNEUROSCI.1473-15.2016

[11] N. Brunel and V. Hakim. (1999). "Fast global oscillations in networks of integrate-and-fire neurons with low firing rates." Neural computation 11, 1621–1671

[12] P. Bryant and T. Trabasso. (1971) "Transitive Inferences and Memory in Young Children." Nature 232, 456–458 (1971). https://doi.org/10.1038/232456a0

[13] A. Bucci and M. Grasso. (2017) "Sleep and dreaming in the predictive processing framework." In Metzinger, T. K. & Wiese, W. (eds) Philosophy and Predictive Processing, chap. 6, https://doi.org/10.15502/9783958573079 (MIND Group, Frankfurt am Main, 2017).

[14] G. Buzsáki. (2015) “Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning,” Hippocampus 25, 1073–1188, https://onlinelibrary.wiley.com/doi/pdf/10.1002/hipo.22488.

[15] C. Capone, et al. Sleep-like slow oscillations improve visual classification through synaptic homeostasis and memory association in a thalamo-cortical model. Sci Rep 9, 8990 (2019). https://doi.org/10.1038/s41598-019-45525-0

[16] C. Capone et al. "Slow waves in cortical slices: How spontaneous activity is shaped by laminar structure." Cereb. Cortex 2019, 29, 319–335

[17] N.T. Carnevale and M.L. Hines. (2006). 2The NEURON Book." Cambridge, UK:Cambridge University Press.

[18] M. Celotto et al. "Analysis and Model of Cortical Slow Waves Acquired with Optical Techniques." Methods Protoc. 2020, 3, 14.

[19] R. Chen et al. "Neural recording and modulation technologies" Nat Rev Mater 2, 16093 (2017). https://doi.org/10.1038/natrevmats.2016.93

[20] T. Chou et al. (2018) "CARLsim 4: An open source library for large scale, biologically detailed spiking neural network simulation using heterogeneous clusters." In 2018 International Joint Conference on Neural Networks (IJCNN).

[21] D.G. Clark et al. (2019) "Unsupervised Discovery of Temporal Structure in Noisy Data with Dynamical Components Analysis". arXiv:1905.09944

[22] D. Contreras and M. Steriade. (1995) "Cellular basis of eeg slow rhythms: a study of dynamic corticothalamic relationships." Journal of Neuroscience 15, 604–622

[23] J. Cunningham, J, B. Yu. (2014) "Dimensionality reduction for large-scale neural recordings." Nat Neurosci 17, 1500–1509. https://doi.org/10.1038/nn.3776

[24] C. Curto et al. (2009). "A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex." J. Neurosci. 29, 10600–10612. doi: 10.1523/JNEUROSCI.2053-09.2009

[25] G. De Bonis, et al. (2019) "Analysis Pipeline for Extracting Features of Cortical Slow Oscillations", Frontiers in Systems Neuroscience, Vol. 13, pages 70, doi: 10.3389/fnsys.2019.00070

[26] A. Destexhe. (2009) "Self-sustained asynchronous irregular states and up–down states in thalamic, cortical and thalamocortical networks of nonlinear integrate-and-fire neurons. Journal of computational neuroscience 27, 493

[27] A. Destexhe and D. Contreras (2011). “The fine structure of slow-wave sleep oscillations: from single neurons to large networks,” in Sleep and Anesthesia, Chapter 4, ed A. Hutt (New York, NY: Springer New York), 69–105.

[28] L. de Vivo et al. (2017) "Ultrastructural evidence for synaptic scaling across the wake/sleep cycle." Science 355, 507–510, https://doi.org/10.1126/science.aah5982

[29] E. Facco et al. (2017) "Estimating the intrinsic dimension of datasets by a minimal neighborhood information." Sci Rep 7, 12140. https://doi.org/10.1038/s41598-017-11873-y

[30] S.B. Furber et al. (2013). "Overview of the SpiNNaker system architecture." IEEE Trans. Comput. 62, 2454–2467. doi: 10.1109/TC.2012.142

[31] P. Gao et al. (2017) "A theory of multineuronal dimensionality, dynamics and measurement" bioRxiv 214262; doi: https://doi.org/10.1101/214262

[32] S. Garcia et al. "Neo: An Object Model for Handling Electrophysiology Data in Multiple Formats." Frontiers in Neuroinformatics 8, 10. doi: 10.3389/fninf.2014.00010 (2014).

[33] J.A. Garrido et al. (2011) "Event and time driven hybrid simulation of spiking neural networks. In Advances in Computational Intelligence", pages 554–561. Springer Berlin Heidelberg.

[34] M.-O. Gewaltig and M. Diesmann. (2007). "NEST (NEural Simulation Tool)." Scholarpedia 2:1430. doi: 10.4249/scholarpedia.1430

[35] G. Gigante, M. Mattia and P. Del Giudice, (2007). "Diverse population-bursting modes of adapting spiking neurons." Phys. Rev. Lett.98, 148101. doi: 10.1103/PhysRevLett.98.1481

[36] J.S. Goldman et al. "Bridging Single Neuron Dynamics to Global Brain States", Frontiers in Systems Neuroscience, Vol. 13, pages 75, 2019, https://www.frontiersin.org/article/10.3389/fnsys.2019.00075, 10.3389/fnsys.2019.00075

[37] B. Golosio et al. (2020) "Thalamo-cortical spiking model of incremental learning combining perception, context and NREM-sleep-mediated noise-resilience", arXiv:2003.11859v1, formally accepted for publication in PLOS Computational Biology

[38] B. Golosio et al. (2021) "Fast Simulations of Highly-Connected Spiking Cortical Models Using GPUs". Frontiers in Computational Neuroscience, Vol. 15, pages 13, https://www.frontiersin.org/article/10.3389/fncom.2021.627620, doi: 10.3389/fncom.2021.627620

[39] B. Golosio et al. (2020) "Toward a possible integration of NeuronGPU in NEST." In NEST Conference 2020, page 7

[40] A. González-Rueda et al. (2018) "Activity-dependent downscaling of subthreshold synaptic inputs during slow-wave-sleep-like activity in vivo." Neuron 97, 1244–1252

[41] D. Goodman and R. Brette. (2009). "The brian simulator." Front. Neurosci. 3:26. doi: 10.3389/neuro.01.026.2009

[42] C. Grienberger et at. (2012) "Imaging Calcium in Neurons." Neuron 2012, 73, 862–885.

[43] R. Gutzen et al. "Building adaptable and reusable pipelines for investigating the features of slow cortical rhythms across scales, methods, and species",

2020, Bernstein Conference, online, Germany, 29 Sep 2020 - 29 Oct 2020, doi: 10.12751/NNCN.BC2020.0030 (Preprint)

[44] F. Han et al. (2008). "Reverberation of recent visual experience in spontaneous cortical waves." Neuron 60, 321–327. doi: 10.1016/j.neuron.2008.08.026

[45] D.P. Hanes & J.D. Schall (1996) "Neural control of voluntary movement initiation" Science. 1996 Oct 18;274(5286):427-30. doi: 10.1126/science.274.5286.427

[46] S.L. Hill and G. Tononi (2005). "Modeling sleep and wakefulness in the thalamocortical system." J. Neurophysiol. 93, 1671–1698. doi: 10.1152/jn.00915.2004

[47] M.L. Hines and N.T. Carnevale, (1997). "The neuron simulation environment." Neural Comput. 9, 1179–1209.

[48] J.A. Hobson et al. (2002) "The cognitive neuroscience of sleep: neuronal systems, consciousness and learning. Nat. Rev. Neurosci. 3, 679–693. doi: 10.1038/nrn915

[49] G. Indiveri et al. (2011) "Neuromorphic silicon neuron circuits." Frontiers in Neuroscience, 5, 2011

[50] E.M. Izhikevich and G.M. Edelman. (2008). "Large-scale model of mammalian thalamocortical systems." Proc. Natl. Acad. Sci. U.S.A. 105, 3593–3598. doi: 10.1073/pnas.0712231105

[51] S.P. Jadhav et al. (2012) “Awake hippocampal sharp-wave ripples support spatial memory,” Science 336, 1454–1458, https://science.sciencemag.org/content/336/6087/1454.full.pdf.

[52] J. Jordan et al. (2018). "Extremely scalable spiking neuronal network simulation code: from laptops to exascale computers." Front. Neuroinform. 12:2. doi: 10.3389/fninf.2018.00002

[53] W.D. Killgore. (2010) “Effects of sleep deprivation on cognition,” (Elsevier, 2010) pp. 105 – 129.

[54] J.C. Knight and T. Nowotny. (2018) "GPUs outperform current HPC and neuromorphic solutions in terms of speed and energy when simulating a highly-connected cortical model." Frontiers in Neuroscience.

[55] D. Kobak et al. (2016) "Demixed principal component analysis of population activity in higher cortical areas reveals independent representation of task parameters". arXiv:1410.6031, doi: 10.7554/eLife.10989

[56] G.P. Krishnan et al. (2016). "Cellular and neurochemical basis of sleep stages in the thalamocortical network." eLife 5, 1–29. doi: 10.7554/eLife.18607

[57] S. Kunkel et al. (2017) Nest 2.12.0, https://doi.org/10.5281/zenodo.259534 (2017).

[58] M. Larkum. (2013) "A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex." Trends in Neurosciences 36, 141–151, https://doi.org/10.1016/j.tins.2012.11.006

[59] G.D. Logan & W.B. Cowan (1984) "On the ability to inhibit thought and action: A theory of an act of control." Psychological Review, 91(3), 295–327. doi: https://doi.org/10.1037/0033-295X.91.3.295

[60] A. Luczak et al. (2009). "Spontaneous events outline the realm of possible sensory responses in neocortical populations." Neuron 62, 413–425. doi: 10.1016/j.neuron.2009.03.014

[61] M. Massimini et al. (2004) "The sleep slow oscillation as a traveling wave." J Neurosci. 2004 Aug 4;24(31):6862-70. doi: 10.1523/JNEUROSCI.1318-04.2004

[62] M. Mattia and P. Del Giudice. (2000). "Efficient event-driven simulation of large networks of spiking neurons and dynamical synapses." Neural Comput. 12, 2305–2329. doi: 10.1162/089976600300014953

[63] M. Mattia and P. Del Giudice. (2002). "Population dynamics of interacting spiking neurons." Phys. Rev. E 66:051917. doi: 10.1103/PhysRevE.66.051917

[64] M. Mattia and M.V. Sanchez-Vives. (2012). "Exploring the spectrum of dynamical regimes and timescales in spontaneous cortical activity." Cognit. Neurodyn. 6, 239–250. doi: 10.1007/s11571-011-9179-4

[65] M. Mattia et al. (2013) "Heterogeneous Attractor Cell Assemblies for Motor Planning in Premotor Cortex." Journal of Neuroscience 3 July 2013, 33 (27) 11155-11168; doi: 10.1523/JNEUROSCI.4664-12.2013

[66] L. Mazzucato et al. (2016) "Stimuli Reduce the Dimensionality of Cortical Activity". Front. Syst. Neurosci., 17 February 2016 doi: 10.3389/fnsys.2016.00011

[67] P.A. Merolla et al. (2014). "A million spiking-neuron integrated circuit with a scalable communication network and interface." Science 345, 668–673. doi: 10.1126/science.1254642

[68] Mione et al. (2020) "Dorsal Premotor Cortex Neurons Signal the Level of Choice Difficulty during Logical Decisions" Cell Reports 32, 107961, https://doi.org/10.1016/j.celrep.2020.107961

[69] G. Mirabella et al. (2011) "Neural correlates of cognitive control of reaching movements in the dorsal premotor cortex of rhesus monkeys". Journal of Neurophysiology 2011 106:3, 1454-1466

[70] D.S. Modha et al. (2011). "Cognitive computing." Commun. ACM 54, 62–71. doi: 10.1145/1978542.1978559

[71] A. Morrisonet al. (2005). "Advancing the boundaries of high-connectivity network simulation with distributed computing." Neural Comput. 17, 1776–1801. doi: 10.1162/0899766054026648

[72] A. Morrison, M. Diesmann, and W. Gerstner. (2008) "Phenomenological models of synaptic plasticity based on spike timing." Biological Cybernetics, 98(6):459–478

[73] L. Muller et al. (2018) "Cortical travelling waves: mechanisms and computational principles." Nat Rev Neurosci 19, 255–268 (2018). https://doi.org/10.1038/nrn.2018.20

[74] M. Murphy et al. (2009) “Source modeling sleep slow waves.” Proceedings of the National Academy of Sciences of the United States of America vol. 106,5 (2009): 1608-13. doi:10.1073/pnas.0807933106

[75] J.M. Nageswaran et al. (2009). "A configurable simulation environment for the efficient simulation of large-scale spiking neural networks on graphics processors." Neural Netw. 22, 791–800. doi: 10.1016/j.neunet.2009.06.028

[76] T.A.E. Nghiem et al. (2018) "Cholinergic switch between two different types of slow waves in cerebral cortex.2 bioRxiv

[77] P.S. Paolucci et al. (2013). "Distributed simulation of polychronous and plastic spiking neural networks: strong and weak scaling of a representative mini-application benchmark executed on a small-scale commodity cluster." arXiv:1310.8478.

[78] E. Pastorelli et al. (2018) "Gaussian and Exponential Lateral Connectivity on Distributed Spiking Neural Network Simulation," 2018 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP), Cambridge, 2018, pp. 658-665, doi: 10.1109/PDP2018.2018.00110.

[79] Pastorelli Elena et al. (2019) “*Scaling of a Large-Scale Simulation of Synchronous Slow-Wave and Asynchronous Awake-Like Activity of a Cortical Model With Long-Range Interconnections*“, Frontiers in Systems Neuroscience, Vol. 13, pages 33, https://www.frontiersin.org/article/10.3389/fnsys.2019.00033, 10.3389/fnsys.2019.00033

[80] K. Pearson (1901) "On lines and planes of closest fit to systems of points in space." The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 2(11):559–572.

[81] P. Pani et al. (2018) "Visual salience of the stop signal affects the neuronal dynamics of controlled inhibition".Sci Rep 8, 14265. https://doi.org/10.1038/s41598-018-32669-8

[82] P. Pani et al. (2019) "Neuronal population dynamics during motor plan cancellation in non-human primates". bioRxiv 774307; doi: https://doi.org/10.1101/774307

[83] H.E. Plesser. (2018). "Reproducibility vs. Replicability: A Brief History of a Confused Terminology", Frontiers in Neuroinformatics, vol. 11, pages 76, https://www.frontiersin.org/article/10.3389/fninf.2017.00076, doi: 10.3389/fninf.2017.00076

[84] T.C. Potjans and M. Diesmann. (2014). "The cell-type specific cortical microcircuit: Relating structure and activity in a full-scale spiking network model," Cerebral Cortex, vol. 24, no. 3, pp. 785–806

[85] S. Recanatesi et al. (2019) "Dimensionality in recurrent spiking networks: Global trends in activity and local origins in connectivity." PLoS Comput Biol 15(7): e1006446. DOI: 10.1371/journal.pcbi.1006446

[86] F. Resta et al. (2020). "Study of Slow Waves (SWs) propagation through wide-field calcium imaging of the right cortical hemisphere of GCaMP6f mice" [Data set]. EBRAINS. doi: 10.25493/3E6Y-E8G

[87] V. Reyes-Puerta et al. (2016). "Propagation of spontaneous slow-wave activity across columns and layers of the adult rat barrel cortex in vivo." Brain Struct. Funct. 221, 4429–4449. doi: 10.1007/s00429-015-1173-x

[88] L.M. Ricciardi. (1977). "Diffusion Processes and Related Topics in Biology." Berlin; Heidelberg; New York, NY: Springer-Verlag Berlin Heidelberg.

[89] M. Ruiz-Mejias et al. (2011). "Slow and fast rhythms generated in the cerebral cortex of the anesthetized mouse." J. Neurophysiol. 106, 2910–2921. doi: 10.1152/jn.00440.2011

[90] M.V. Sanchez-Vives et al., "Cellular and network mechanisms of rhythmic recurrent activity in neocortex." Nat. Neurosci. 2000, 3, 1027–

[91] M.V. Sanchez-Vives, and M. Mattia (2014). "Slow wave activity as the default mode of the cerebral cortex." Arch. Ital. Biol. 152, 147–155. doi: 10.12871/000298292014239

[92] M.V. Sanchez-Vives, et al. "Shaping the Default Activity Pattern of the Cortical Network." Neuron 2017, 94, 993–1001, [PubMed]

[93] J. Sanders and E. Kandrot. (2010) "CUDA by Example: An Introduction to General-Purpose GPU Programming." Addison-Wesley, Upper Saddle River, NJ.

[94] P. Sanz-Leon et el. (2013) "The Virtual Brain: a simulator of primate brain network dynamics". Frontiers in Neuroinformatics 7:10. doi: 10.3389/fninf.2013.00010

[95] T.K. Sato et al. (2012) "Traveling waves in visual cortex." Neuron, Volume 75, Issue 2, Pages 218-229, https://doi.org/10.1016/j.neuron.2012.06.029.

[96] S. Schmitt et al. (2017). “Neuromorphic hardware in the loop: training a deep spiking network on the BrainScaleS wafer-scale system,” in 2017 International Joint Conference on Neural Networks (IJCNN) (Anchorage, AK: IEEE), 2227–2234. doi: 10.1109/IJCNN.2017.7966125

[97] 1034 P. Schnepel et al., “Physiology and impact of horizontal connections in rat neocortex,”, Cerebral Cortex, vol. 25, no. 10, pp. 3818–3835, 2015.

[98] A. Schuz et al., “Quantitative aspects of corticocortical connections: A tracer study in the mouse,” Cerebral Cortex, vol. 16, no. 10, pp. 1474–1486, 2006.

[99] F. Simula et al. "Real-Time Cortical Simulations: Energy and Interconnect Scaling on Distributed Systems," 2019 27th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP), Pavia, Italy, 2019, pp. 283-290, doi: 10.1109/EMPDP.2019.8671627

[100] F. Smulders et al. (1997) “The effects of sleep loss on task performance and the electroencephalogram in young and elderly subjects,” Biological Psychology 45, 217 – 239 (1997), mental Resources: Intensive and Selective Aspects.

[101] G. Solovey et al. (2015). "Loss of consciousness is associated with stabilization of cortical activity." J. Neurosci. 35, 10866–10877. doi: 10.1523/JNEUROSCI.4895-14.2015

[102] M.L. Steyn-Ross et al. (2013). "Interacting Turing-Hopf instabilities drive symmetry-breaking transitions in a mean-field model of the cortex: a mechanism for the slow oscillation. Phys. Rev. X 3:21005. doi: 10.1103/PhysRevX.3.021005

[103] A. Stepanyants et al. (2009) “The fractions of short- and long-range connections in the visual cortex,”, Proceedings of the National Academy of Sciences, vol. 106, no. 9, pp. 3555–3560.

[104] M. Steriade et al. (1991) "Network modulation of a slow intrinsic oscillation of cat thalamocortical neurons implicated in sleep delta waves: cortically induced synchronization and brainstem cholinergic suppression." Journal of Neuroscience 11, 3200–3217 (1991).

[105] M. Steriade et al. (1993) "The slow ($<$ 1 Hz) oscillation in reticular thalamic and thalamocortical neurons: scenario of sleep rhythm generation in interacting thalamic and neocortical networks", 1993 Journal of Neuroscience, 13(8):3284-3299.

[106] M. Steriade et al., "A novel slow ($<$ 1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components.", J. Neurosci. 1993, 13, 3252–3265.

[107] M. Stimberg et al. (2014). "Equationoriented specification of neural models for simulations." Front. Neuroinform. 8:6. doi: 10.3389/fninf.2014.00006

[108] S.H. Strogatz. (2018). "Nonlinear Dynamics and Chaos With Student Solutions Manual: With Applications to Physics, Biology, Chemistry, and Engineering." Boca Raton, FL: CRC Press.

[109] E. Stromatias et al. (2013) “Power analysis of large-scale, real-time neural networks on spinnaker,” in The 2013 International Joint Conference on Neural Networks (IJCNN), pp. 1–8

[110] G. Tononi and C. Cirelli. (2014) “Sleep and the price of plasticity: From synaptic and cellular homeostasis to memory consolidation and integration,” Neuron 81, 12–34.

[111] G. Tononi et al. (2018) "Center for sleep and consciousness", Research Center, http://centerforsleepandconsciousness.med.wisc.edu

[112] G. Tononi and C. Cirelli. (2020) “Sleep and synaptic down-selection,” European Journal of Neuroscience 51, 413–421, https://onlinelibrary.wiley.com/doi/pdf/10.1111/ejn.14335.

[113] N. Tort-Colet. (2019) “Attractor competition enriches cortical dynamics during awakening from anesthesia,” bioRxiv (2019), 10.1101/517102, https://www.biorxiv.org/content/early/2019/01/10/517102.full.pdf

[114] V.V Vyazovskiy and K.D. Harris. (2013) "Sleep and the single neuron: the role of global slow oscillations in individual cell rest." Nature Reviews Neuroscience 14, 443–451, https://doi.org/10.1038/nrn3494

[115] J. Vitay et al. (2015) "ANNarchy: a code generation approach to neural simulations on parallel hardware." Frontiers in Neuroinformatics.

[116] M.P. Walker and R. Stickgold. (2006) “Sleep, memory, and plasticity,” Annual Review of Psychology 57, 139–166, pMID: 16318592, https://doi.org/10.1146/annurev.psych.56.091103.070307

[117] R.M. Wang et al.(2018) "An FPGA-based massively parallel neuromorphic cortex simulator." Frontiers in Neuroscience.

[118] B. Watson et al. (2016) "Network homeostasis and state dynamics of neocortical sleep." Neuron 90, 839–852, https://doi.org/10.1016/j.neuron.2016.03.036

[119] https://wiki.ebrains.eu/bin/view/Collabs/slow-wave-analysis-pipeline

[120] https://github.com/INM-6/wavescalephant

[121] https://github.com/APE-group/wavescalephant

[122] Y. Wei et al. (2016) "Synaptic mechanisms of memory consolidation during sleep slow oscillations." Journal of Neuroscience 36, 4231–4247

[123] Y. Wei et al. (2018) "Differential roles of sleep spindles and sleep slow oscillations in memory consolidation. PLoS computational biology 14, e1006322

[124] M.A. Wilson et al. (1989). “Genesis: a system for simulating neural networks,” in Advances in Neural Information Processing Systems 1, ed D. S. Touretzky (San Francisco, CA: Morgan-Kaufmann), 485–492.

[125] J.Y. Wu et al. (2008) "Propagating waves of activity in the neocortex: What they are, what they do." Neuroscientist. 2008 Oct; 14(5): 487–502. doi: 10.1177/1073858408317066

Models:

https://kg.ebrains.eu/search/instances/Model/
cea3e597-2fbd-4022-bbfc-a64b9fa49d68

https://kg.ebrains.eu/search/instances/Model/
97670076281ccbdc38ea2c2d76a64e64

https://kg.ebrains.eu/search/instances/Model/

fa08c511f9b444922b0975f538b10abd