



HAL
open science

Surfing cognition like a traveling wave: from unconscious learning to predictive coding

Andrea Alamia

► **To cite this version:**

Andrea Alamia. Surfing cognition like a traveling wave: from unconscious learning to predictive coding. Cognitive science. Université Toulouse Paul Sabatier, 2024. tel-04403904

HAL Id: tel-04403904

<https://hal.science/tel-04403904>

Submitted on 19 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License



UNIVERSITÉ
TOULOUSE III
PAUL SABATIER



Université Paul Sabatier – Toulouse III
Ecole Doctorale ED-MITT

Mémoire présenté en vue de l'obtention de Habilitation
à Diriger des Recherches

Surfing cognition like a traveling wave: from unconscious learning to predictive coding

Andrea Alamia

Defense date: 17th / 01 / 2024

Jury members: Dr. Rufin VanRullen – Université Toulouse III (President)

Prof. Dr. Peggy Seriés – University of Edimburgh (Rapporteure)

Dr. Frédéric Chavane – Université Aix-Marseille (Rapporteur)

Prof. Dr. Martin Vinck – Max Planck Institute (Rapporteur)

Centre de Recherche Cerveau et Cognition

(CerCo)

CNRS UMR 5549



Table of contents

Summary	3
Chapter I : Overview and Introduction	4
A multidisciplinary path in Cognitive Neuroscience	4
From Experimental to Computational Neuroscience and Machine Learning	4
Unconscious processes.....	5
Predictive Coding	6
Brain oscillations.....	8
Chapter II : Previous Research	9
Scientific activity during my doctorate (2014-2017)	9
Unconscious learning	9
Demonstrating unconscious learning.....	9
Unconscious learning and visual attention	11
Unconscious learning modulates motor response.....	13
Pupil response to unconscious surprisal	13
Sequence learning.....	15
The role of Broca's area in sequence learning	15
Other studies on chunking.....	15
Scientific activity during my post-doctorate (2017-2021)	17
Predictive Coding and traveling waves	17
A relatively simple computational model.....	18
Quantifying traveling waves in EEG data	20
The role of traveling waves in visual cognition	22
Artificial grammar and neural networks	25
Visual reasoning, brain oscillations, and neural networks.....	27
Chapter III: First independent scientific activity and future directions	31
First steps as Principal Investigator (2021 – 2023)	31
Finalizing ongoing projects post-covid pandemic.....	31
Machine learning and Neuroscience	33
ERC Starting Grant 'OSCI-PRED' (2023 – 2028)	36
Introduction and project motivation.....	36
Main goal and overall structure of the project.....	38
WP 1 - Combining Predictive Coding with cross-frequency coupling.	38
WP 2 - Going deeper: characterizing laminar dynamics.....	41
WP 3 - Integrating information across cortical regions: the role of the Pulvinar.....	45
WP 4 - Attention and conscious perception modulate oscillatory traveling waves.....	47
Chapter IV: Conclusions and long term perspective	51
Predictive Processes and Oscillatory Traveling waves: a unifying framework for cortical functions?	51
Traveling waves and Predictive Coding: from psychedelics to schizophrenia patients	52
The computational role of traveling waves	54
Brain-inspired models to improve AI state-of-the-art	55
Long-term perspective and conclusion	58
New directions: Computational Psychiatry	58
Conclusions - Beyond EEG recordings: intracortical recordings and animal models.	60
References	61

Summary

This manuscript summarizes my previous work in cognitive neuroscience and delineates the short-term steps and the global vision for my future research interests. The overarching goal of my research career is to investigate human cognition using different tools and methods. Inter- and multi- disciplinary are two key aspects that characterize my scientific journey, and it's essential to keep these elements in mind when navigating through this manuscript. Several topics and methods sparked my scientific curiosity, from unconscious learning to brain oscillations, from computational modeling to neural networks. Besides my neverending interest in human cognition, the pole star that has always set the course is the quest to find a comprehensive framework to understand the human mind. I started my journey during my phd when I explored unconscious and sequence learning, running several experiments on healthy humans (often medicine students). After establishing a methodologically sound framework to investigate unconscious processes, I explored how these influence eye movements, pupil size, and EEG recordings. Next, I set the course toward more computational shores, investigating how predictive coding could give rise to neural oscillations and traveling waves during my postdocs. Moving through differential equations and neural networks, I compared the performance of models with humans in different tasks, such as visual reasoning or artificial grammar learning. In the next years, combining all I learned along the way, I will dive into brain dynamics at different scales, understanding whether predictive coding could be the critical framework for understanding brain dynamics or, at least, traveling waves. In the long term, I plan to set sail toward more clinical-oriented applications, exploring the fascinating new world of computational psychiatry.

Chapter I : Overview and Introduction

A multidisciplinary path in Cognitive Neuroscience

This manuscript summarizes my scientific work from the beginning of my Ph.D. until my first years as an independent researcher. Over nine years, I had the chance to dive into several topics in Cognitive Neuroscience, exploring different directions and using various experimental techniques and computational approaches. One of the main characteristics of my work is the interplay between experimental and computational methods, and the structure of this manuscript mirrors such an interdisciplinary trajectory. Below, I introduced the research topics I investigated over the years, from experimental psychology to neurophysiology and artificial intelligence. In this manuscript, I followed a chronological order, tracking my progression over the years in different domains. Even though the results may seem patchy and erratic, it's possible to find the 'fil rouge' that ties together all of my previous work, which is the ambitious quest to improve our understanding of human cognition.

From Experimental to Computational Neuroscience and Machine Learning

During my doctoral studies (2013-2017), under the supervision of Etienne Olivier and Alexandre Zénon, I carried out my research using different experimental techniques. During this period, I learned methods in psychophysics, Transcranial Magnetic Stimulation (TMS), eye movements, pupillometry, and electrophysiology (EEG). After completing my graduate studies in 2017, I integrated these experimental methods with more computational approaches. During my two postdocs (2017-2019 & 2019-2021), working with Rufin VanRullen and Thomas Serre, I learned new techniques, such as modeling using differential equations (ODEs) and machine learning optimizations, specifically neural networks. Currently, I'm applying both computational and experimental methods to tackle my scientific investigations, pursuing a strongly interdisciplinary approach.

Irrespective of the experimental or computational methods, during my scientific career, I had the opportunity to focus on various topics within Cognitive Neuroscience. The first part of Chapter II describes the work during my doctoral studies about unconscious processes and implicit sequence learning, which is learning new contingencies without being aware of them. In this domain, I have explored how unconscious learning influences behavioral performances as well as eye movements, pupil responses, and electrophysiological signals. After my doctoral studies, I realized that I needed a general framework to not only interpret my results but also describe cognitive processes from a broader perspective. Eventually, I came across the

framework that would shape my first postdoc and the following of my career: Predictive Coding (and, more broadly, Predictive processes and the Bayesian framework). The second part of Chapter II describes the work I performed during my postdocs with Rufin VanRullen. In this period, I started investigating the role of brain oscillations as a functional mechanism pivotal in synchronizing brain activity across brain regions. In particular, I got interested in the relationship between Predictive Coding and brain oscillations, especially when considering their temporal and spatial components, i.e., as traveling waves. In parallel with these studies about Predictive Coding and brain oscillations, I renewed my interest in Computational Neuroscience and Machine Learning by leading two studies using Neural Networks to investigate cognitive processes and human behavior. In both studies, I compared human performance with feedforward and recurrent networks to understand the computational mechanisms involved in visual reasoning (first study) and artificial grammar learning (second study). Finally, chapter III details my ongoing research, in which I have the chance to integrate the experimental and computational approaches developed during my career. The main goal of my current work is to address the ambitious question of whether Predictive Processes and oscillatory traveling waves may provide a framework to understand brain dynamics at different scales and their relation to various cognitive functions (perception, attention, etc.).

Given the variety of subjects approached during my journey, I'll briefly introduce below the three main topics that have characterized my scientific interests and that recur through this manuscript: unconscious processes, predictive coding, and brain oscillations.

Unconscious processes

Unconscious processing occurs in everyday life, influencing several aspects of our behavior and choices. For example, learning our mother tongue is primarily based on unconscious mechanisms, as we acquire grammatical rules without being able to describe them precisely. In modern psychology, the conscious-unconscious dichotomy has always played an essential role in the scientific and clinical domains. Since Freud, many scientists have tried to unveil the mental processes that escape the scrutiny of consciousness. For a long while, the term 'unconscious processing' has been related to subliminal learning (Greenwald et al., 1995), in which stimuli presented for a few dozen milliseconds and not perceived consciously influence participants' behavior. In this domain, scientists were interested in understanding how consciously perceived information would affect perceptual processing and decision-making (Bargh & Morsella, 2008). Regarding supraliminal stimuli, some studies investigated unconscious processing in automatic and habitual behavior (Dickinson, 1985; Ouellette & Wood, 1998), whereas others focused on the linguistic domain, particularly grammar learning. Arthur Reber initially started this line of research in 1967 with a seminal study about the implicit learning of artificial grammars (Reber, 1967a), and it introduced the term 'implicit learning' for

the first time, paving the way for further studies in this domain beyond the linguistic framework. Following those results, several other paradigms have been proposed to demonstrate the occurrence of unconscious learning (Cleeremans et al., 1998). To this day, the literature about unconscious processing is exceptionally vast, and it embraces many different approaches and methods developed in neuropsychology and cognitive sciences over the last decades. Nevertheless, the very existence of unconscious learning was questioned on the ground of methodological and theoretical considerations (Newell & Shanks, 2014; Shanks & Stjohn, 1994). The first part of my doctorate work addressed these criticisms, as detailed in the first part of chapter II.

Despite the fundamental scientific interests, the studies of unconscious processes also have an essential impact on the studies of clinical populations. For example, several studies reported no difference between amnesic patients and healthy participants in paradigms based on implicit priming effects (Warrington & Weiskrantz, 1968, 1982) and complex implicit learning paradigms (Channon et al., 2002). Similar studies reveal a lack of difference between patients affected by severe prosopagnosia and healthy controls in tasks involving implicit facial memory (de Haan et al., 1987) and no difference in the implicit processing of visually presented words between dyslexic and aphasic patients and healthy controls (Shallice & Saffran, 1986). Lastly, a study published in 1987 (Nissen & Bullemer, 1987) reported no impairment in implicit sequence learning tasks in patients affected by Alzheimer's syndrome compared to healthy participants. All in all, the clinical perspective indicates that conscious and unconscious processes possibly rely on different brain regions involving distinct neuronal mechanisms. In a study during my postdoc, I further tested this intriguing hypothesis by comparing human performance in implicit and explicit learning tasks (artificial grammar learning) with feedforward and recurrent neural networks (Alamia, Gauducheau, et al., 2020), as detailed in the second part of chapter II.

Predictive Coding

Predictive Coding is one of my main scientific interests, as it provides a comprehensive framework to investigate cognitive processes. From an evolutionary perspective, it aims at reducing signal redundancy and optimizing computational efficiency (Huang & Rao, 2011), as our surroundings are mostly constant and predictable, as objects have finite physical dimensions and tend to persist over time. These regularities cause a significant redundancy in the natural signals that the brain samples from the environment, thus encouraging the emergence of optimal information-theory strategies to optimize sensory processing, such as predictive processes. Its fundamental concept postulates that natural neural networks learn the environment's statistics to predict regularities and reduce uncertainty over time. In a

nutshell, the key idea is that a hierarchical system (e.g., the visual system) learns an internal representation of the world (i.e., a model) from which it generates predictions to anticipate incoming signals. This approach implies that perception is a *generative* process that can be formalized appropriately within the framework of Bayesian theory (Colombo & Seriès, 2012; De Lange et al., 2018; Vilares & Kording, 2011). Broadly, this probabilistic framework postulates that the brain has an internal model of the world that represents the external environment and the agent itself (via interoceptive and proprioceptive systems). Such a model constitutes, in Bayesian terms, the *Prior* from which generating predictions to interact with the environment and its own state efficiently. Sensory evidence defines the likelihood function, which, combined with the Prior, generates the Posterior distribution, which ultimately represents the updated internal model. Crucially, the Bayesian formulation describes how the brain operates in a condition of uncertainty, as the Prior and the environmental information (the Likelihood of the sensory data) are optimally integrated accounting for their uncertainty (Knill & Pouget, 2004). Interestingly, dysfunctions in properly integrating the two sources of information, i.e., the internal Prior and the external sensory evidence, may produce pathological behaviors related to psychiatric diseases, such as schizophrenia and autism (Angeletos Chrysaitis & Seriès, 2023; Sterzer et al., 2018a). The last chapter of this work will elaborate more along these lines within the new emerging domain of Computational Psychiatry (Series, 2020).

Besides the relevance for understanding cognitive models, PC has important implications for neural systems and electrophysiology. In one of the first implementations of PC, Rao and Ballard (Rao & Ballard, 1999) proposed a hierarchical model in which the neural activity of higher regions represents *predictions* about the activity of the lower areas. The difference between the prediction and the actual activity of the lower region is defined as a *prediction error*, and it is sent from the lower to the higher region to correct the next prediction. Remarkably, when implementing these dynamics, they demonstrated that artificial neurons developed receptive fields whose properties were very similar to those in real neurons in visual area V1, illustrating the biological plausibility of such a theory. Their seminal study boosted a very fruitful line of research that has endorsed PC as one of the working principles of the brain. This success benefited Neuroscience as well as other domains, such as Machine Learning and Artificial Intelligence, as it inspired a novel approach to designing brain-inspired architectures (Choksi et al., 2020; Lotter et al., 2017; O'reilly et al., 2021) (as detailed in the second part of chapter II). Finally, when interacting with a natural environment, it is necessary to predict not only *what* is going to happen, but also *when* (Giraud & Arnal, 2012). Recently, theoretical reviews have proposed brain oscillations as a suitable candidate for both functions in cortical circuits, effectively enabling predictive coding dynamics in the brain (Arnal et al.,

2011; Bastos, Litvak, et al., 2015). This leads to the third and last topic essential to my scientific interests and this manuscript: brain oscillations.

Brain oscillations

In the last decades, several studies have pointed at Brain Oscillations as a pivotal computational tool in several cognitive functions. Perception, memory and even consciousness may result from networks' synchronization (Boly et al., 2011; Fell & Axmacher, 2011). Several factors influence the amplitude and the topography of brain rhythms, thus suggesting their involvement in several tasks. Generally, cortical oscillations are thought to play a role in neuronal communication and synchronization of brain regions (Fries, 2015), motor coordination and execution (Brittain & Brown, 2014), perception and information processing (e.g., feature binding, discrete perception, (Eckhorn et al., 1988; VanRullen, 2016), etc.). As mentioned above, brain oscillations have also been recently related to Predictive Coding, associating prediction and prediction-error to different frequency bands based on experimental evidence. Precisely, fast gamma oscillations (>30Hz) have been related to local cortical processes, and besides characterizing a wide range of cognitive functions (Lundqvist et al., 2016; Z. G. Zhang et al., 2012), they proved to match sensory expectations. Consequently, they are thought to reflect prediction errors, corroborated by the fact that the regular repetition and the unexpected omission of a stimulus decrease and increase gamma oscillations' activity (Fujioka et al., 2009; Iversen et al., 2009). On the other hand, alpha-beta band oscillations (~8-30Hz) have been related to top-down activity, possibly carrying predictions from higher to lower brain regions. This interpretation has also received experimental support (Haegens et al., 2011; Samaha et al., 2015; van Pelt et al., 2016). However, new experimental and theoretical work reinterprets the association between frequency bands and predictive coding processes, suggesting that oscillations play a role in sustaining information processing and improving its efficiency (Vinck et al., 2023). I'll discuss more on this in the later section, 'Predictive Processes and Oscillatory Traveling waves: a unifying framework for cortical functions?'. All in all, several hypotheses have been proposed to clarify the interplay between Brain Oscillations and Predictive Coding. However, a functional model combining the key experimental evidence at different levels (i.e. macro and mesoscopic) is still lacking. Arguably, one of the next ambitious goals in Cognitive Neuroscience is to propose such a general and coherent scheme, combining Predictive Coding with oscillatory brain dynamics. In Chapter III's section 'ERC Starting Grant 'OSCI-PRED' (2023 – 2028)', I will describe my proposal to tackle this exciting question.

Chapter II : Previous Research

Scientific activity during my doctorate (2014-2017)

MAIN PAPERS OF THE SECTION	<p>Alamia A., de Xivry J.J., Anton E., Olivier E., Cleeremans A., Zenon A. (2016) "Unconscious associative learning with conscious cues". <u>Neuroscience of Consciousness</u> 1-10.</p> <p>Alamia A., Zenon A. (2016) "Statistical Regularities Attract Attention when Task-Relevant". <u>Frontiers in Human Neuroscience</u>, 10:42.</p> <p>Alamia A., Solopchuk O., D'Ausilio A., Van Bever V., Olivier E., Zenon A. (2016) "Disruption of Broca's Area Alters Higher-order Chunking Processing during Perceptual Sequence Learning". <u>Journal of cognitive neuroscience</u>. Vol 28, N°3, p.402-417.</p> <p>Alamia A., Solopchuk O., Olivier E., Zenon A. (2016) "Non-parametric algorithm to isolate chunks in response sequences". <u>Frontiers in Behavioral Neuroscience</u>, 10:177.</p> <p>Alamia A., Solopchuk O., Zenon A. (2018) "Strong conscious cues suppress preferential gaze allocation to unconscious cues". <u>Frontiers in Human Neuroscience</u>, 12:427.</p> <p>Alamia A., VanRullen R., Pasqualotto E., Mouraux A., Zenon A. (2019) "Pupil responds to unconscious surprisal". <u>Journal of Neuroscience</u>, 3010-18</p> <p>Alamia A., Zenon A., VanRullen R., Duque J.,Derosiere G..(2019) "Unconscious perceptual cues drive oscillatory activity in the motor cortex during action selection ". <u>Neuroimage</u>, 186, 424-436</p>
----------------------------	--

Unconscious learning

Demonstrating unconscious learning

The main focus of my doctoral studies was on unconscious learning. The primary motivation for investigating unconscious processes is that it is one of the key experimental ways to study human consciousness. Typically, it is possible to frame the current literature on unconscious learning along two dimensions (Figure 2.1, from Alamia, Orban de Xivry, et al., 2016): the first one determines whether the stimuli used during learning are supraliminal or subliminal, whereas the second dimension characterizes the complexity of the rules or associations to be learned. More specifically, the term “complex rule” is used here to refer to task structures that are composed of a large number of contingencies, such as in sequence learning tasks (in which numerous transitions between successive button presses have to be learned), or artificial grammar tasks (in which a set of probabilistic rules drive the generation of grammatical strings). In opposition, simple rules can be defined as task structures composed of a small number of contingencies. Notably, there is not a precise separation between these two rule classes, which instead represent two extremes of a continuum. The hallmark of implicit learning is the use of supraliminal stimuli to induce learning of abstract, complex rules (top left corner in Figure 2.1) (Reber, 1967).




Figure 2.1 – A) Schematic representation of unconscious learning paradigm. The x-axis represents the stimulus perceptibility (subliminal or supraliminal) and the y-axis represents the rules complexity (simple–complex rules). Our study is the first one to propose a supraliminal task with simple rules, thus addressing the main criticisms proposed by Shanks and colleagues about the existence of unconscious learning. B) Experimental design of the first study: following a fixation cross displayed for 600 ms, a patch of moving dots was displayed for 300 ms. The participants had 500 ms to provide a response indicating the motion direction. Unbeknownst to the participants, two out of the three possible colors were always associated with a given direction/response, while one was uninformative.

However, in a seminal study published in 1994, followed by another one in 2014 (Newell & Shanks, 2014; Shanks & Stjohn, 1994), Shanks and colleagues cast a shadow on the very existence of unconscious processing, raising significant methodological concerns over the previous studies. Importantly, they suggested a set of 4 criteria to test the level of awareness reliably:

- 1- The *sensitivity criterion* regards the sensitivity of the measures of awareness;
- 2- the *information criterion* suggests that the measure of awareness should probe the same information as the experimental task;
- 3- the *immediacy criterion* imposes that the tests should be simultaneous (or follow immediately) to the experimental task;

- 4- the *relevance criterion* suggests that the measure of awareness should avoid irrelevant information.

In their studies, they concluded that “*convincing evidence in favor of the existence of unconscious learning was still lacking*”. Starting from this context, the first goal of my thesis was to settle this controversy by demonstrating the existence of unconscious learning by addressing all Shanks' criteria. The main issue relates to the vast complexity of the rules learned implicitly (e.g., an artificial grammar composed of several rules) and the fact that the experimental tests are not sensitive enough to detect the actual amount of explicit and implicit knowledge learned by the participants. The critical intuition –and the main challenge- of our study (Alamia, Orban de Xivry, et al., 2016) was to elicit unconscious learning with a simple experimental design, thus addressing the most stringent criteria. We designed a motion discrimination task in which some stimulus colors were associated with the motion direction, thus cueing the correct response. Remarkably, participants exploited the color information but proved to be truly unaware of the association. We demonstrated the lack of awareness with three independent tasks (a questionnaire, a familiarity, and a generative task), fulfilling all Shanks' criteria. Our study was the first to propose a robust and straightforward behavioral framework leading to implicit learning, proving its behavioral relevance and providing a reliable framework to study it thoroughly.

Unconscious learning and visual attention

Once we had a solid experimental framework to investigate unconscious processes, we explored its relationship with several cognitive functions, taking advantage of different experimental techniques. We first examined the interaction between **visual attention** and unconscious learning before considering its **electrophysiological correlates**, specifically in the motor cortex (see below). To address the first point, I performed two studies in which I explored how unconscious information influences eye movements (Alamia et al., 2018; Alamia & Zénon, 2016). In one study (Alamia et al., 2018), we exploited a modified version of the motion discrimination task (figure 2.2A), in which we displayed two patches of dots moving coherently in the same direction (either left or right). Participants could freely move their eyes to fixate on either of the two patches and were instructed to report the motion direction while we recorded eye movements. Importantly, patches could have two coherence levels, either high or low (easy and hard to perceive, respectively), and, unbeknownst to the participants, their color was associated with the motion direction (as in the previous study). In any given trial, either patch could have high or low coherence, and its color could be informative or not, thus having an explicit cue (the coherence) and an implicit one (the color). First, we showed that participants unconsciously learned the association between color and motion direction; then, we demonstrated that the unconscious learning of the color-motion association

influenced attentional allocation (i.e., eye movements) but that this effect was present only for low coherence levels, i.e., for difficult trials. In other words, our results reveal that conscious and unconscious sources of information influence attentional allocation and suggest a selection process that weights cues in proportion to their reliability. The second study found similar results in the statistical learning domain, specifically in the implicit learning of statistical regularities. We showed participants two moving stimuli (Figure 2.2B), changing color in the middle of their trajectory. Importantly, the changes were driven by a simple Markovian process, which determined the color transition probabilities. Participants were instructed to press a button when a given color was on the screen. Our results showed that participants tracked stimuli when these were predictive of the upcoming target, that is, when the stimuli were task-relevant (Alamia & Zénon, 2016). In conclusion, our two studies demonstrated that eye movements are affected by unconscious information and statistical learning. Next, we investigated whether we could observe neurophysiological correlates of unconscious learning via EEG recordings.




Figure 2.2 – A) Experimental Design of the experiment in Alamia, Solopchuk & Zenon, 2018. The figure shows a representative trial composed of three parts: fixation cross, stimulus presentation and response time. Participants were instructed to report the motion direction of the patches (which was always consistent between stimuli). B) Experimental Design of the experiment in Alamia & Zenon 2016. The upper part is a schematic representation of a whole trial, while the lower part of the picture represents the successive stages of a block. The dashed line represents the range of dot positions in which a change of color may occur. Colors changed following the transition probabilities as shown in the inset. Two colors were always associated to each other (conditional probability = 1, predictable colors) while the remaining colors all shared a conditional probability equal to 0.33 (unpredictable colors).

Unconscious learning modulates motor response

To explore the **electrophysiological signature** of unconscious learning, I performed one study in collaboration with Dr. Gerard Derosi re and Prof. Julie Duqu e, (Universit  Catholique de Louvain, Belgium), in the context of a collaboration about the role of motor cortex in implicit learning (Alamia, Z non, et al., 2019; Derosi re et al., 2015, 2017). In previous studies, using brain stimulations (Transcranial Magnetic Stimulation, more specifically cTBS, which disrupts the regular activity of the targeted brain area), we showed that the primary motor cortex plays an important role in the implementation of implicit rules, crucially redefining the role of such region from simple motor actuator to a more prominent involvement in resolving motor decisions (Derosi re et al., 2015, 2017). From these premises, we investigated the electrophysiological recordings related to motor components while participants were involved in the motion discrimination paradigm presented above (in which the color-motion associations were unconsciously learned). Our results revealed that implicit cues modulate the oscillatory contralateral activity in the beta range (16-25Hz), known to be involved in motor preparation, and this modulation had a behavioral effect, influencing the speed of the motor response (i.e., reaction times). In addition to replicating our previous results, we provided new evidence that implicit and explicit sources of information are integrated into the motor cortex during decision-making, thus expanding and integrating our previous studies.

Pupil response to unconscious surprisal

The last study within the unconscious learning framework played an essential role in my transition from cognitive to computational neuroscience, and it sparked my interest in Predictive processes in general and Predictive Coding in particular. In this study (Alamia, VanRullen, et al., 2019), we showed for the first time that unconscious learning of statistical regularities can be tracked by pupillary responses, such as a dilation induced by arousing events (e.g., unconscious surprise). Participants were asked to stare at a stream of letters, presented at 1Hz, to detect a specific target (i.e., the letter 'A,' which rarely occurred). Unbeknown to the participants, the four letters (excluding the target) which composed the stream follow some statistical rules, as shown in Figure 2.3A. Specifically, each letter was likely to precede two letters but very rarely a third one (a rare transition occurring only 5% of the time). After confirming thoroughly that the rules were truly unconscious (once again by

meeting all Shanks' criteria via different tasks that probe participants' knowledge of the rules – see above), we showed that the pupil dilates more in response to unconsciously surprising events (figure 2.3 B). Moreover, we replicated these results in a series of 2 additional experiments: in the first one, we demonstrated the crucial role of attention in the occurrence of learning, as we did not observe the pupillary dilation if participants were performing another task that oriented attention away from the letters (i.e., detecting changes in the fixation cross color). In the second one, we revealed via EEG recordings the electrophysiological correlates related to the rare transition, as shown in Figure 2.3 C. Importantly, this project, which concluded my doctoral studies on unconscious learning, introduced me to the idea of predictions and prediction-errors, and more broadly to the framework that will determine most of my scientific interests: *Predictive Coding*.

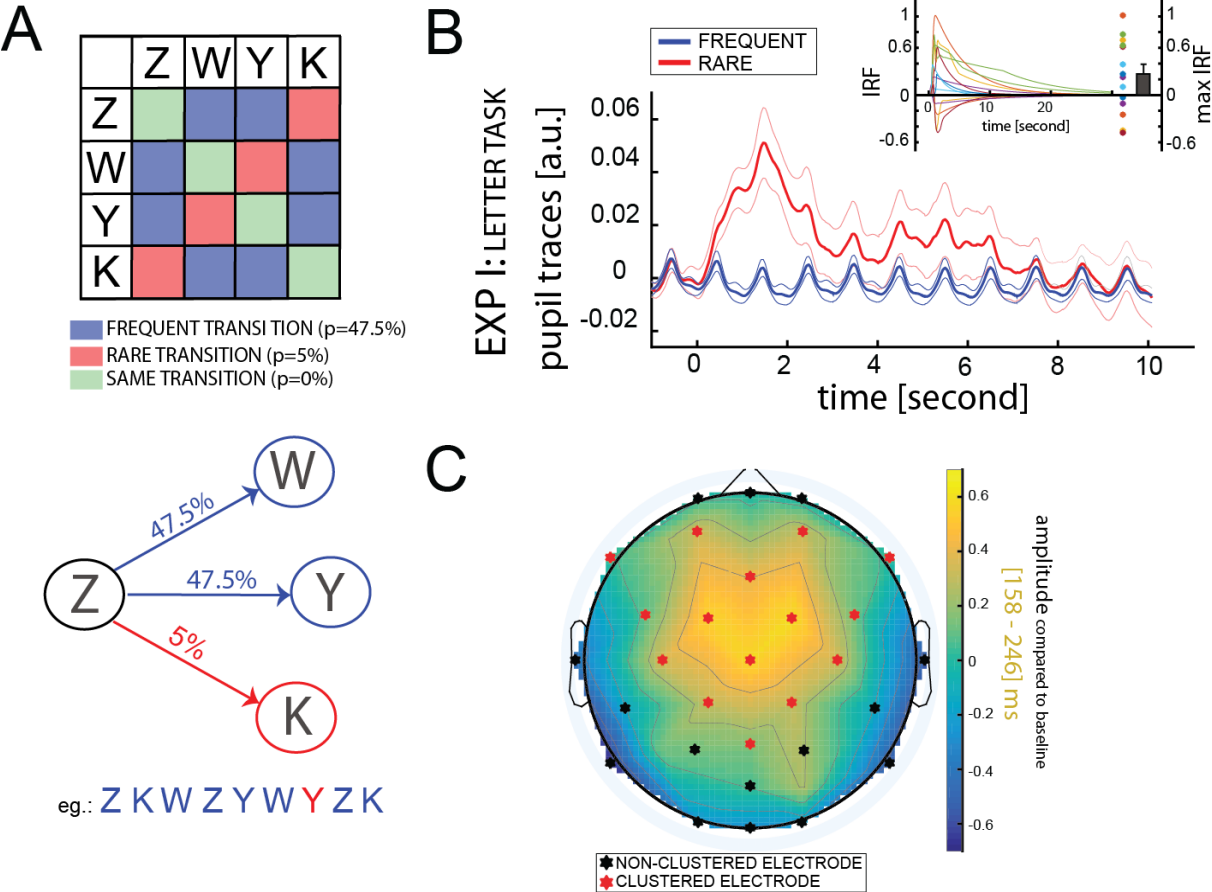


Figure 2.3 – A) Experimental Participants attended to a stream of letters. The transitions between letters followed a Markovian process, with frequent and rare transitions. B) Pupillary response is larger after the rare transition (in red). C) EEG evoked activity correlates with the pupillary response.

In the next chapter, I will outline the studies I carried out during my postdoc, investigating the link between oscillatory activity and predictive coding. However, before that, in the next section I will briefly overview another line of research I carried out during my graduate studies about

sequence learning in collaboration with another fellow Ph.D. student, Oleg Solopchuk, and my supervisor, Etienne Olivier.

Sequence learning

In this section, I will briefly present my works in *Sequence Learning*, a line of research I developed in parallel with *Unconscious Learning* during my doctoral studies, which proved decisive in learning new technical skills (i.e., TMS), establishing collaborations, and discovering new topics. In this domain, together with my colleague and friend Oleg Solopchuk, we investigated the cognitive and neural correlates of chunking. These information-processing mechanisms consist of grouping consecutive items in a sequence, for example, to facilitate its memorization (as when memorizing a phone number, we tend to decompose it in small chunks) but also to execute complex motor actions. Importantly, chunking is a cognitive process related to many cognitive functions, from language to body movements.

The role of Broca's area in sequence learning

Our main study investigated the role in sequence learning of the Broca's area, a brain region generally involved in language (Alamia, Solopchuk, et al., 2016). Specifically, we applied Transcranial Magnetic Stimulation to our participants via continuous Theta Burst Stimulation (cTBS), a sequence disrupting the targeted area's activity. Participants performed a perceptual sequence learning task in two sessions, in which we either targeted the poster part of Broca's area or a control site (figure 2.4A). Importantly, the sequence of items had a specific structure, which allowed participants to group it (or chunk it) at different hierarchical levels (figure 2.4A). We developed a framework to analyze reaction times by applying a clustering method borrowed from network analysis (Mucha et al., 2010; Wymbs et al., 2012), namely the *Louvain method for community detection* (Blondel et al., 2008), which allowed us to measure the strategy used by participants during the task. Remarkably, our results revealed that the disruption of Broca's area increased the processing time of higher levels but not the lower ones, thus modifying the participants' chunking strategy. These results contribute to understanding the role of this brain region in processing hierarchical structures, such as language and complex motor actions.

Other studies on chunking

Besides investigating the role of Broca's area in chunking, we also explored the contribution of the Supplementary Motor Area on sequence representation, using a similar approach based on cTBS stimulation (Solopchuk et al., 2017), and we discussed the potential role of the dorsal premotor cortex in sequence learning (Solopchuk, Alamia, & Zénon, 2016). We also design an experiment to test the role of chunking in working memory (Solopchuk, Alamia, Olivier, et al., 2016), correlating chunking strategies with the performance in other tasks involving working

memory processes. Lastly, we developed a non-parametric algorithm to reliably identify chunks in reaction times series (Alamia, Solopchuk, Olivier, et al., 2016), which was used in previous studies to quantify chunking strategies. Besides this work on chunking and sequence learning, during my doctoral path, I had the chance to collaborate with the group of Valery Legrain, also at the Institute of Neuroscience in Brussels, where I contributed to the data analysis and the implementation of Bayesian methods to adapt task difficulty, within the context of visual perception and nociceptive stimulation. This work, which I won't present here as quite tangential to my career, nevertheless turned into a few publications and enriching collaborations (Filbrich, Alamia, Blandiaux, et al., 2017; Filbrich, Alamia, Verfaillie, et al., 2017; Filbrich et al., 2018; Vanderclausen et al., 2017). All in all, these studies, complementary to my main line of research in Unconscious learning, proved very important to learn new experimental skills and theoretical frameworks in Cognitive Neuroscience.

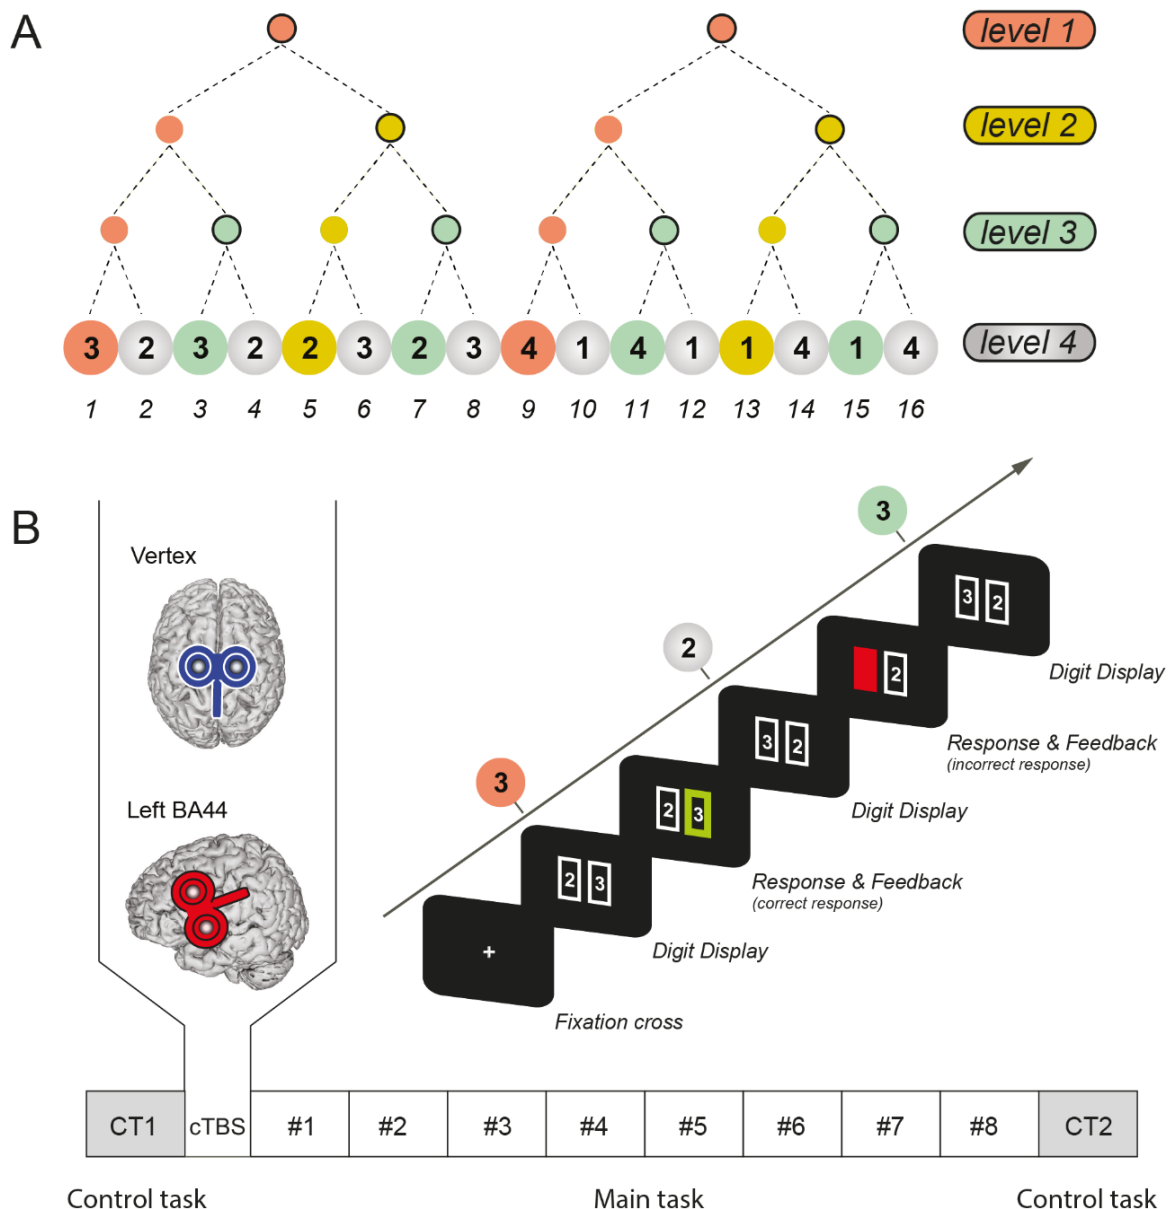


Figure 2.4 – A) The sequence to be learned contained 16 elements and was characterized by a particular relational structure, so it could be chunked consistently by all participants according to four distinct levels. B) Time course of a sequence and design of the experiment. Each sequence started with the display of a fixation cross for 3000 msec, followed by the display of a first pair of stimuli (digits in a rectangle). The participant had to select the digit belonging to the sequence by pressing the appropriate response key (right index or right middle finger). This was followed by feedback consisting of a green rectangle surrounding the selected digit for a correct response or a red mask for an incorrect one. Then the next pair of digits was displayed. Each block comprised five sequence repetitions. The lower part of the figure illustrates the experimental design. The experiment started with the first control task (CT1), which had the same design as the main task except that the four digits were replaced by four letters (A, B, C, and D) and there was no sequence to learn. It was followed by the application of cTBS either over either the left BA 44 or the vertex (control group). Then, the participants performed the main task (8 blocks x 5 sequences), and the control task was repeated (CT2) at the end of the experiment.

Scientific activity during my post-doctorate (2017-2021)

MAIN PAPERS OF THE SECTION	<p>Alamia A., VanRullen R. (2019) "Alpha oscillations and travelling waves: signatures of predictive coding?" PLOS Biology, 17(10), e3000487</p> <p>Alamia A., Luo C., Ricci M., Kim J., Serre T., VanRullen R. (2021) " Differential involvement of EEG oscillatory components in sameness vs. spatial-relation visual reasoning tasks" eNeuro 10.1523/ENEURO.0267-20.2020</p> <p>Vaishnav M., Cadene R., Alamia A., D Linsley, VanRullen R., Serre T.(2022) "Understanding the computational demands underlying visual reasoning" Neural Computation</p> <p>Luo C., VanRullen R., Alamia A. (2021) "Conscious perception modulates perceptual echoes" Neuroscience of Consciousness</p> <p>Alamia A., Gauducheau V., Paisios D., VanRullen R. (2020) "Comparing feedforward and recurrent neural network architecture with human behavior in Artificial Grammar Learning" Scientific Reports 10(1), 1-15</p> <p>Alamia A., Timmermann C., Nutt D.J., VanRullen R., Carhart-Harris R. (2020) " DMT alters cortical travelling waves" eLife 9:e59784</p> <p>Timmermann C., Roseman L., Haridas S., Rosas FE, Luan L., Kettner H, Martell J., Erritzoe D., Tagliazucchi E., Pallavicini C., Girn M., Alamia A., Leech R., Nutt DJ., Carhart-Harris R. (2023) "Human brain effects of DMT assessed via EEG-fMRI" Proceedings of the National Academy of Science 120, 13 e2218949120</p> <p>Choksi M., Mozafari M., O'May C., Ador B., Alamia A., VanRullen R. (2021) "Predify: augmenting deep neural networks with brain-inspired predictive coding dynamics" Advances in Neural Information Processing System 34, (NeurIPS)</p> <p>VanRullen R., Alamia A. (2021) "GattaNet: Global agreement for convolutional neural networks" ICANN2021</p> <p>Pang Z., Alamia A., VanRullen R. (2020) "Turning the stimulus on and off changes the direction of alpha travelling waves" eNeuro, 7(6), ENEURO.0218-20.2020</p>
----------------------------	---

Predictive Coding and traveling waves

After my doctoral studies, I started looking for a general framework to interpret my results from a broader perspective. I felt that most of the studies in Cognitive Sciences, from Experimental Psychology to Electrophysiology, were but a small brick of a more extensive construction that I couldn't perceive in its entirety. Such as the perhaps unfulfilled promise of String theory in physics, I was looking quite ambitiously for a view in Cognitive Neuroscience that could integrate all of the plethora of results and findings from different studies (and possibly, differently than String theory, be tested experimentally). Eventually, I came across the framework of Predictive Coding (or, more broadly, Bayesian inference), which seemed a promising starting point for such an ambitious quest. Accordingly, during my first postdoc, the primary goal was to design a computational model based on Predictive Coding principles able to explain some fundamental physiological observations, such as brain oscillations. I got

particularly interested in neural oscillations as they proved a pivotal computational tool in many cognitive functions, as demonstrated by several studies in the last decades. Perception, memory, and even consciousness may result from networks' synchronization (Boly et al., 2011; Fell & Axmacher, 2011). Several factors influence brain rhythms' amplitude and topography, suggesting their involvement in several tasks. Generally, cortical oscillations play a role in neuronal communication and synchronization of brain regions (Buzsáki & Draguhn, 2004; Fries, 2015; Vezoli et al., 2021; Vinck et al., 2023), motor coordination and execution (Brittain & Brown, 2014), perception and information processing, as in feature binding or discrete perception, etc. (Eckhorn et al., 1988; VanRullen, 2016). As mentioned in the first chapter, brain oscillations have also been related to Predictive Coding, associating prediction and prediction-error to different frequency bands, and several hypotheses have been proposed to clarify the interplay between Brain Oscillations and Predictive Coding. However, a functional model reconciling distinct experimental evidence at different scales (i.e., macro and mesoscopic) is still lacking. Arguably, one of the next ambitious goals in Cognitive Neuroscience is to propose such a general and coherent scheme, combining Predictive Coding with oscillatory brain dynamics. In my first postdoc with Rufin VanRullen, we partially contributed to this challenging question, implementing a relatively simple model based on Predictive Coding principles that could explain the emergence of oscillatory dynamics.

A relatively simple computational model

Given a hierarchical system, such as a sensory system, the critical idea in Predictive Coding is that the higher level predicts the lower level activity, and the difference between the prediction and the actual activity (i.e., the prediction error) is used to update and correct the next prediction. Ideally, given a stable and predictable environment, the higher-level activity tracks the changes in the stimuli, eventually converging to a stable interpretation of the world in which the prediction error is null. However, our study's fundamental intuition was considering biologically plausible constraints when modeling how information propagates between adjacent brain regions (figure 2.5A). More specifically, if we consider a dynamic environment and the time delay that the information takes to propagate between levels, both the prediction and the prediction error will constantly be late in updating each other, eventually evolving into a constant overshooting. This simple dynamic generates oscillatory activity, whose frequency depends on the temporal delay and neuronal time constants (intuitively, shorter delays result in higher frequencies, see Figure 2.5B). Below are the equations describing this simple model, governed by a system of first-order differential equations:

$$\frac{dy_L}{dt} = \frac{1}{\tau} \cdot x_L(t-\Delta T) + \frac{1}{\tau_D} \cdot (y_{L+1}(t-\Delta T) - y_L(t)) \quad (1)$$

$$x_L(t) = y_{L-1}(t) - y_L(t - \Delta T) \quad (2)$$

Equation 1 determines each region's neuronal activity $y(t)$, whose dynamic is driven by two factors: the bottom-up residual x_L (defined in equation 2) and the difference between its representation and the top-down predictions (which, in Bayesian terms, can be considered a form of prior). In the equations above, L denotes the levels, and ΔT represents the temporal communication delay: more specifically, the prediction from the L level will have a temporal delay of ΔT before reaching the $L - 1$ level. In our study, we assumed the temporal delay to be the same in the forward and backward directions (note that it's straightforward to demonstrate that similar oscillatory dynamics would be found for symmetric and asymmetric delays, e.g., with $\Delta T = 12\text{ms}$, or with $\Delta T_{\text{forward}} = 16\text{ms}$ and $\Delta T_{\text{backward}} = 8\text{ms}$). Besides ΔT , two other parameters are crucial in our implementation: τ and τ_D , which describe neuronal integration and decay's temporal dynamics (time constants). Each time constant weights the residual computed from the lower-level and the prediction from the higher level. Importantly, when the prior y_{L+1} is set to 0, the second term in equation (1) acts as a decay term, ensuring that the prediction $y_L(t)$ returns to zero in the absence of inputs. Considering additionally that higher-level brain signals typically vary slower than low-level input signals (Gauthier et al., 2012; Kiebel et al., 2008; McKeeff et al., 2007; Murray et al., 2014), we set the time constant τ_D an order of magnitude larger than τ (Alamia & VanRullen, 2019a). Importantly, this simple two-layer version of the model implementing predictive coding dynamics could produce oscillatory activity, whose frequency depends on the model's parameters (see Figure 2.5B). Next, we decided to extend the model to several levels, and we introduced a prior activity representing an internal model of the world which constantly generates predictions about the environment (and, for simplicity, shares the same statistics as the input signal, i.e., white noise in our simulations). Surprisingly, such extension from two to a multi-level representation leads to observing an oscillatory front of activity propagating through the levels as a traveling wave, as shown in figure 2.5C (Alamia & VanRullen, 2019a; Muller et al., 2018). In this multi-level version of the model, we observed that the oscillatory activity presents a phase shift between levels, representing oscillatory traveling waves that propagate through regions, in agreement with recent studies investigating cortical traveling waves (Muller et al., 2014, 2018; H. Zhang et al., 2018). Importantly, these simulations generate new hypotheses about the propagation of these oscillatory waves, which arise naturally from our models' dynamics. In particular, our model formulates precise predictions about the direction of propagation of the waves: the waves move forward and backward, driven respectively by the input or the prior, thus reflecting the 'cognitive state' of the model (i.e., whether the system is processing a stimulus or not). To verify these predictions, we analyzed two datasets of EEG recordings in healthy human subjects. In the first one, participants attended a visual input; in the second, they kept their eyes closed for one minute.

Remarkably, we observed a significant percentage of forward waves in the first condition and backward waves in the other, thus confirming the predictions of our model. All in all, we showed that a simple Predictive Coding model could explain the generation of alpha-band oscillations and that these oscillations propagate as traveling waves backward or forward (top-down or bottom-up, respectively), depending on the cognitive state of the system (with or without perceptual inputs).

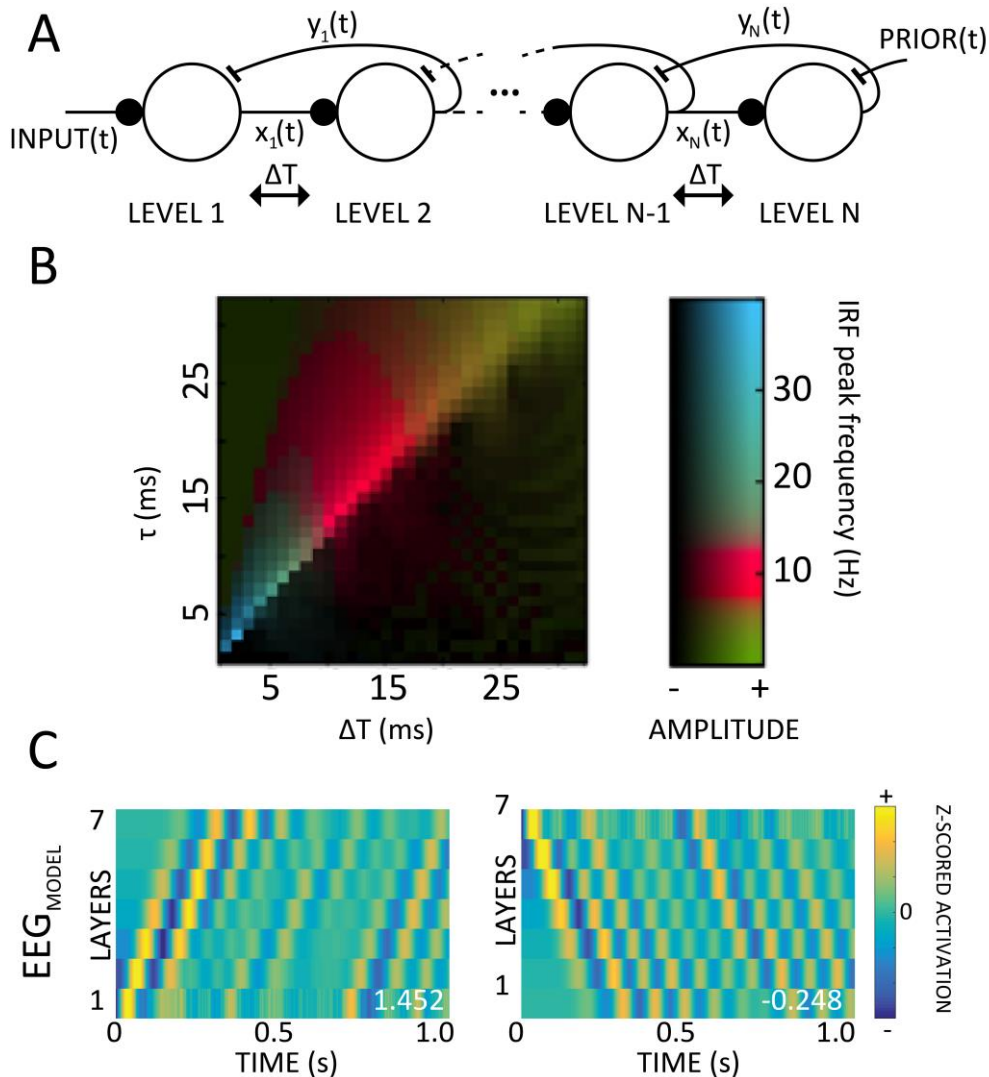


Figure 2.5 – A) Multilevel version of the model: the same parameters ($\Delta T = 12$ ms and $\tau = 20$ ms) are used throughout. The model is fed either a time-varying input (left) or a time-varying prior signal (right), reflecting top-down expectations computed in other parts of the brain. B) Systematic exploration of these two parameters suggests that alpha reverberation is a robust phenomenon (red colors) within a biologically plausible range of values. C) Two-dimensional maps with only input (left column) or prior signals (right column): traveling waves are visible in the raw prediction signals (considered as a proxy for the EEG). All the values have been z-scored level-wise for visualization purposes only.

Quantifying traveling waves in EEG data

This exciting result led us to investigate further the functional role of traveling waves in different cognitive contexts. To carry out these studies, we devised a new method to quantify

traveling waves in EEG recordings in a simple but efficient way, considering an array of electrodes. In all studies, we measured the amplitude of waves propagating forward or backward, as shown in figure 2.6A. Specifically, we apply a moving time window over the EEG signals, lasting 1000ms, with an overlap of 500ms. We combine the signals from 7 electrodes for every time window generating a 2D map having time and electrodes as axis (see figure 2.6A). In all experiments, the waves are computed along the midline of electrodes, using the electrodes from Oz to Fz, except in one recent study investigating visual attention and lateralization, in which we considered distinct, lateralized lines of electrodes (see section ‘Finalizing ongoing projects post-covid pandemic’ for details). Next, we compute 2D Fourier Transformation (2D-FFT) for each map. Importantly, given the property of the 2D-FFT, the upper- and lower-left quadrants represent the power of forward (FW) and backward (BW) traveling waves, respectively (see Figure 2.6B for a schematic representation and an intuition about 2D-FFT properties). Since the 2D-FFT is symmetrical around the origin, the lower-left and upper-right quadrants contain the same information. From quadrants of the same side, we select the maximum value for each temporal frequency, thus obtaining a spectrum representing the raw amount of FW and BW waves for each time window (it is possible, as we did in the first studies, to focus on specific frequency band, such as alpha). To compute a baseline, we run the same procedure after shuffling the electrodes’ order, thereby disrupting spatial information (including the waves’ directionality) while retaining the same overall spectral power. In other words, the surrogate measures reflect the amount of waves expected solely due to the temporal fluctuations in the signal. After having computed the maximum values for the FW and BW waves of the surrogate 2D-FFT spectra one hundred times (and averaging the 100 values), we calculate the net amount of FW and BW waves in decibels (dB) by applying the following formula:

$$WdB = 10 * \log_{10} \frac{W}{W_{SS}}$$

where W represents the maximum value extracted for each quadrant (i.e., forward FW or backward BW), and W_{SS} is the respective surrogate value. Importantly, this value -expressed in decibels- represents the net amount of waves against the null distribution, that is, against the hypothesis of having no FW or BW waves when shuffling the electrodes. It is informative when compared to zero to assess the significance of waves. A direct comparison between FW and BW waves in each time-bin is not readily interpretable, as it is possible to simultaneously record waves propagating in both directions—as observed during visual stimulation epochs. Compared to other methods for quantifying waves in EEG recordings, this approach provides a more robust and reliable estimation, as several electrodes are considered together when quantifying the waves’ amount, even though it has a significant limitation due to the arbitrary

choice of the propagation axis. However, in all the experiments, we were primarily interested in the anterior-posterior axis, predominantly related to visual processing, thus making the 2D-FFT method the ideal choice.

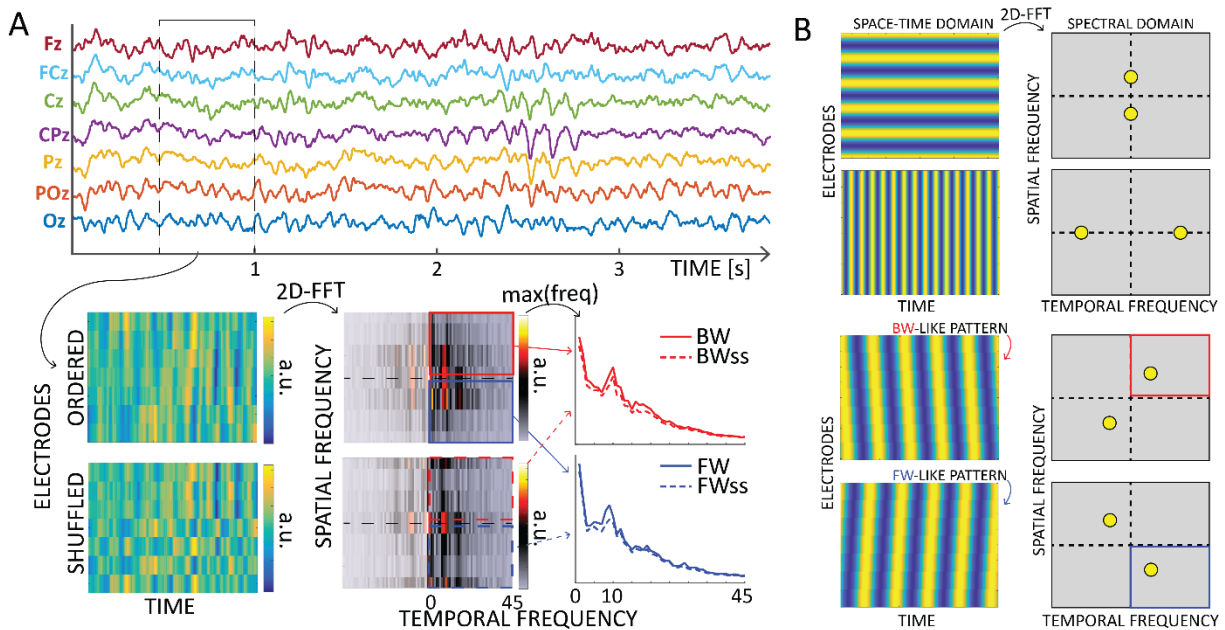


Figure 2.6 A) Schematic of the waves' quantification method. After defining time windows over each electrode line, we computed 2D Fourier transformation to quantify the amount of forward (in blue) and backward (in red) waves. From the upper and lower quadrants of the 2D-FFT spectra, we consider the maximum value over spatial frequencies, providing a 1D spectrum of forward and backward waves in the temporal frequency domain. After shuffling the electrodes' order, the same procedure provides a surrogate measure, used as a baseline. Notably, such surrogate distribution captures the $1/f$ trend and the alpha-band peak, accounting for these factors in the final waves' quantification. B) The 2D-Fast Fourier Transform (2D-FFT) decomposes an image (e.g., a space-time representation of an EEG signal) into its spectral components. The upper part shows the decomposition of a 2D sinusoid propagating along the vertical or horizontal axis of the image. The corresponding peaks are found on the axis in the spectral domain, and their position depends on the frequency of the oscillations. The lower part of the figure shows how the spectra change when the oscillations propagate with a backward- or forward- like pattern. Importantly, the spectral peaks rotate in two of the four quadrants depending on the direction, providing a reliable measure of forward or backward waves in the image.

The role of traveling waves in visual cognition

This section briefly presents our results when quantifying traveling waves in different experimental conditions. This research mainly occurred during my second postdocs, from 2019 to 2022. First, we investigated the role of alpha-band traveling waves during visual stimulation (Pang et al., 2020). In this study, I co-supervised the graduate student, Zhaoyang Pang, involved in this study. Participants attended a visual stimulus for 5 seconds, followed by 5 seconds of resting state (i.e., blank screen), while we recorded EEG signals (Figure 2.7B). We then quantified the moment-to-moment progression of traveling waves. As shown in Figures 2.7C and D (left panel), we observed alpha-band waves propagating from occipital to frontal

regions during visual processing, while waves propagating in the opposite direction dominated in the absence of visual inputs. Interestingly, only during visual stimulation, both forward and backward waves co-occur on average but correlated negatively, revealing an alternation of forward and backward waves over time during visual perception. These results confirmed our previous findings based on the computational model, demonstrating that the alpha-band traveling waves' direction of propagation depends on the cognitive state of the system, specifically on the presence of visual stimulation.

In another study, we explored the role of alpha-band traveling waves in conscious visual perception using a binocular rivalry design (Luo et al., 2021). This study was performed by Canhuang Luo, a former Ph.D. student in the team, while I was involved as a senior author supervising the project. In this work, we investigated the relationship between conscious perception and the generation and amplitude of perceptual echoes, that is, alpha-band oscillations reflecting visual processing (Vanrullen & MacDonald, 2012). Perceptual echoes are computed by cross-correlating the EEG signal with the visual stimulus (e.g., a white noise sequence of luminance changes), thus obtaining the impulse response function (IRF), which quantifies how and in which frequency band the visual system reverberates the sensory information. Using a binocular rivalry setup based on dichoptic mirrors (Figure 2.7A), we asked participants to continuously report which of two colored Gabor patches with different orientations they consciously perceived while recording their EEG signals. Crucially, we could estimate the perceptual echoes generated by the perceived (dominant) and non-perceived (suppressed) stimulus due to the random changes in the luminance of each patch. We first compared the alpha power generated by each patch with a control condition (participants performed the same experiment but with monocular vision, that is only one patch was displayed during the experiment without eliciting any rivalry). Interestingly, we observed that the alpha power related to the consciously perceived stimulus was comparable with the control one and higher than the one induced by the suppressed stimulus. In addition, we found that both perceived and suppressed echoes propagated as a travelling wave from the posterior to frontal brain regions. These results hint at the relation between travelling waves and conscious perception; further studies will model and interpret these results in the framework of Predictive Coding (see 'WP 4 - Attention and conscious perception modulate oscillatory traveling waves.' section).

Concerning the Predictive Coding interpretation of traveling waves, we performed a further study in collaboration with Dr. Robin Carhart-Harris (from the Psychedelic Research Center at the Imperial College in London during the study). Interestingly, they proposed a model based on Predictive Coding to interpret the effect of psychedelics on the brain (Carhart-Harris & Friston, 2019). According to their model, psychedelics reduce the 'precision-weighting of

priors', thus altering the balance of top-down versus bottom-up information passing. In other words, psychedelic drugs clutter one's model of the world (encoded in the prior), generating an altered state of mind and sensory hallucinations (more on the Bayesian interpretation of traveling waves in the next chapter). Remarkably, the results confirmed our predictions: although participants had their eyes closed, after DMT injection, EEG recordings revealed a spatiotemporal pattern of traveling waves similar to that elicited by visual stimulation (figure 2.7D), consisting of an increase of forward waves and a decrease of backward waves. According to the predictions of the proposed model based on Predictive Coding, DMT induces a reduction in the precision-weighting of priors, as reflected in a decrease of top-down waves (encoding predictions), and a consequent rise of prediction-errors, as supposedly revealed by the forward waves. In the following part of this chapter, I'll briefly present additional studies I carried out during my postdoc, which were unrelated to the traveling waves line of research, but compared neural networks' performance with human behavior.




Figure 2.7 A) Binocular rivalry design. Participants stared at the screen through a set of dichoptic mirrors that projected the left and right side of the screen to the left and right eye, respectively. Two stimuli, placed on the two sides of the screen, were Gabor patches of different colour and orientation, either red or green. Participants reported which patch they perceived by moving a joystick to either side, each one associated with a stimulus (pseudo-randomly between participants, consistent across blocks and sessions). B) Experiment design of the Pang et al. (2020) study. Schematic diagrams of a subblock, composed of three identical trials, for a 30-s-long time course. The luminance changed randomly from 0 to 255 on each screen refresh. In each trial, luminance sequences were displayed for the first 5 s (stimulus-on period), followed by a 5 s blank screen (stimulus-off period). C) A spectral analysis on TW comparing visual stimulation (ON) and closed eyes (OFF) reveals a rich dynamic in the waves propagating from occipital to frontal regions (FORWARD) and vice versa (BACKWARD). D) The two plots show the amount of forward and backward alpha-band waves during visual stimulation (left panel) and after the injection of the psychedelics drug DMT (right panel, DMT is injected at minute 5, as shown by the dashed vertical line). Although participants had closed eyes, the waves' profile elicited by DMT injection is remarkably similar to those observed during visual stimulation.

Artificial grammar and neural networks

In parallel with the studies about Predictive Coding, Oscillations, and Travelling Waves, during my first and second postdocs, I carried out two studies using Neural Networks to investigate brain processes and human behavior. The first study (Alamia, Gauducheau, et al., 2020) aligned with my previous work exploring Unconscious Learning. The main goal was to compare human behavior and performance in an Artificial Grammar Learning task with the performance of two different neural network architectures. This task is generally used to investigate implicit learning and language acquisition, and its computational mechanisms are still controversial (Tecumseh Fitch & Friederici, 2012). In one version of the task, the goal is to be able to classify sequences generated according to some rules (i.e. a grammar) and sequences that violate such rules. Previous experimental studies demonstrated that human subjects can learn artificial grammar without explicit knowledge of the rules and after little exposure (Reber, 1967). Our working hypothesis was that we could correlate human learning (i.e., performance over time) with recurrent but not feedforward neural networks, corroborating the hypothesis that recursion plays a crucial role in language (Corballis, 2014). Interestingly, we aimed to generalize these findings by replicating the experiment over four types of grammars, spanning three levels of the Chomsky hierarchy (see figure 2.8A), which determines the *computational complexity* of each grammar. One important novelty of our study was that we tested both humans and neural networks with the same amount of training examples (around 500), which proved to be challenging for the network training but provided ground for a fair comparison. As shown in figure 2.8B, our results show that both architectures can learn the grammars, but that recurrent networks perform closer to humans than feedforward ones, for all the grammars we tested. Remarkably, our results suggested that explicit learning is best modeled by recurrent architectures, whereas feedforward networks better capture the dynamics involved in implicit learning. This conclusion is similar to the one drawn in other studies investigating visual processing (Lamme & Roelfsema, 2000), in which feedforward and recurrent architectures have been related to unconscious and conscious processes, respectively.

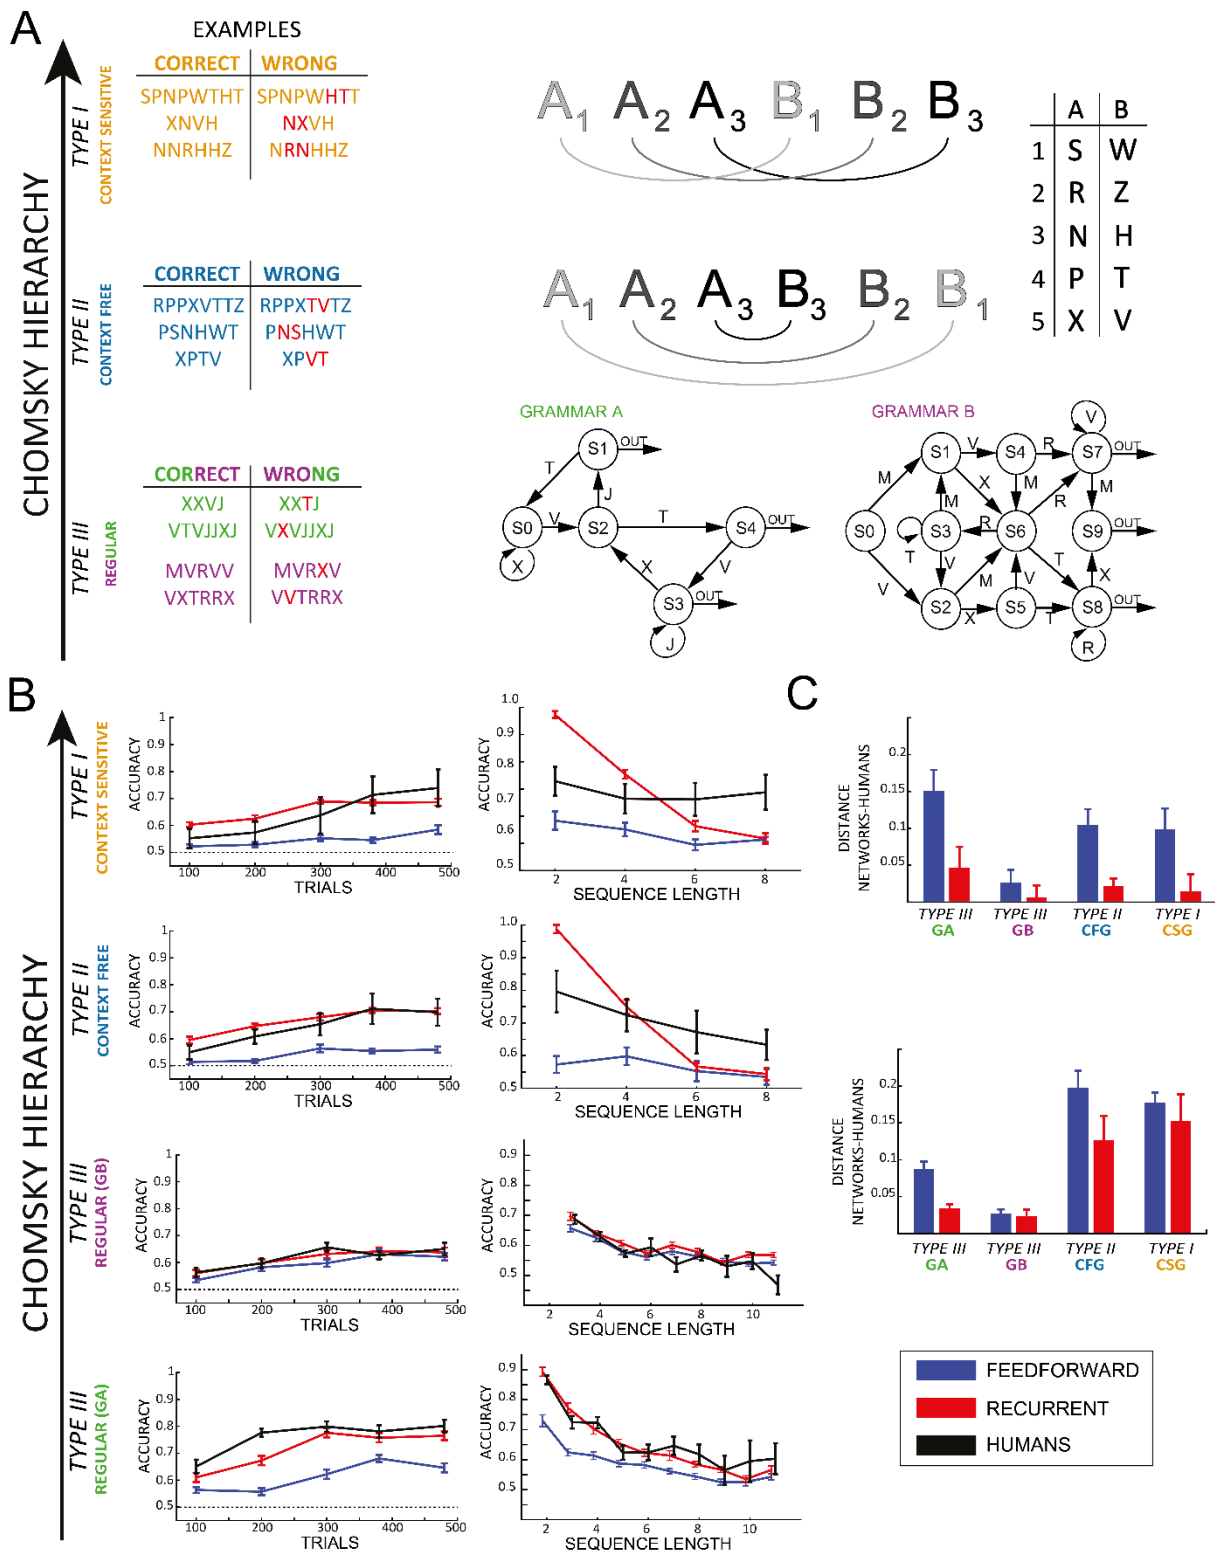


Figure 2.8 - A schematic representation of the 4 grammars employed in this study, arranged according to the Chomsky hierarchy. Type I or **Context Sensitive** grammar (CS, in orange), a translation was applied from the first to the corresponding second half of the string. For example, considering the first half as 'PPN' and referring to the same pairing as in the picture, the corresponding mirrored version would be 'HTT' in the CF, and 'TTH' in the CS grammar. The incorrect sequences were obtained by switching only two non-identical letters within one of the two halves (in red in the table). In type II or **Context Free** grammar (CF, in blue), the second half of the sequence mirrored the first half, applying the respective pairing to each letter (e.g. if A3 was the letter N, then the letter B3 was H). Note that a correct string in the context free grammar is considered as incorrect in the context sensitive, and vice

versa. **Regular** grammars are defined by a scheme organized in directionally connected nodes (type III, the lowest in green and purple). For both grammars A and B, each correct sequence was generated starting at node S0 of each respective scheme and then randomly choosing an arc departing from that node and recording its letter. This procedure was iterated until one of the arc labeled as 'out' was reached, thus terminating the string. Incorrect sequences had only one inexact item, being substituted by another letter from the proper vocabulary but violating grammatical rules (in red in the table). **B**) Results over trials and string lengths for humans (in black) feedforward (in blue) and recurrent (in red) networks. For humans, each bin is an average of over 40 trials (20 trials before and after, respectively, except the last bin, which includes the last 40 trials of the experiment). Each row represents a grammar, ordered according to Chomsky's hierarchy. **C**) The plots show the distance between humans' performance and FF (in blue) and RR (in red). Each distance represents the area measured between the human and the network curves. Except for grammar B, RR networks are significantly closer to human performance as a function of training time and sequence length.

Visual reasoning, brain oscillations, and neural networks

In the previous study about artificial grammar learning, we compared human and neural network performance in the same task to characterize the computational mechanisms involved in grammar learning in humans. Vice versa, in this project, we started with some compelling results from Machine Learning to draw and test specific predictions about cortical computations in humans. In a recent study, Kim and colleagues (Kim et al., 2018) tested feedforward architectures on different visual learning tasks (using the SVRT dataset (Fleuret et al., 2011)). They demonstrated that feedforward networks could easily solve problems involving *spatial relations* (SR), such as judging whether three items are on the same line or whether the first is contained in the second one. Still, they fail to identify whether two items are identical or different (*same-different task*, SD). This conclusion led us to hypothesize that different computational mechanisms are needed to perform those two visual reasoning tasks in the human brain. In collaboration with the Thomas Serre lab at Brown University (who co-supervised my second postdoc), we designed a study to test the hypothesis that the visual system indeed relies on different computational mechanisms to perform the two tasks and that these mechanisms would be reflected by some biological markers, namely brain oscillations (Alamia, Luo, et al., 2020). First, we replicated Kim's and colleagues' results on our own set of stimuli. Specifically, we trained and tested feedforward networks on two tasks using the same set of stimuli (hexominoes, see figure 2.9A): in one case, the networks learned to judge whether two items were the same or not (same-different task), whereas in the other case it classified whether they were more horizontally or vertically aligned (spatial relations task). Figure 2.9B shows that the network could be successfully trained in both conditions, but it could generalize to the test set only in the SR condition, confirming the results of Kim and colleagues. We additionally trained a "Siamese" forward network, in which each item was provided separately to the network (in other words, solving the selective attention problem via inductive bias): in this case, the network could also perform the SD task, demonstrating that the main limitation concerns the spatial localization of the stimuli when presented at the same time. We then decided to

replicate these results while recording EEG signals from human participants, while they were performing both SD and SR tasks. In both cases we used the same stimuli as for the neural networks, with an adapted experimental design (in particular, the items were shown for 30 ms, and participants could provide an answer after 1000ms, receiving feedback at each trial). As hypothesized, our results revealed a significant difference between the two conditions, specifically in the beta range (15-25 Hz) in the occipital brain regions (figure 2.9D). Specifically, we observed increased beta-band activity when performing the SD over the SR condition, possibly reflecting the fundamental involvement of recurrent mechanisms implementing cognitive functions such as working memory and attention. In another study, in which I contributed as co-author, we further investigated the processes underlying visual reasoning from a machine learning perspective (Vaishnav et al., 2022). In this study, led by the Ph.D. student Mohit Vaishnav and supervised by Thomas Serre, we assessed the ability of deep convolutional networks to solve 23 distinct visual reasoning problems (specifically the SVRT challenge). As in our previous work, our results first reveal a novel taxonomy of visual reasoning tasks, primarily explained by two main types of relations, that is, same-different and spatial-relation judgments. Then, we equipped the network with a spatial and feature-based attention system and evaluated their ability to solve the SVRT challenge. We found that attentional systems help to solve the hardest tasks more efficiently and that the improvements in individual tasks partially explained the novel taxonomy.

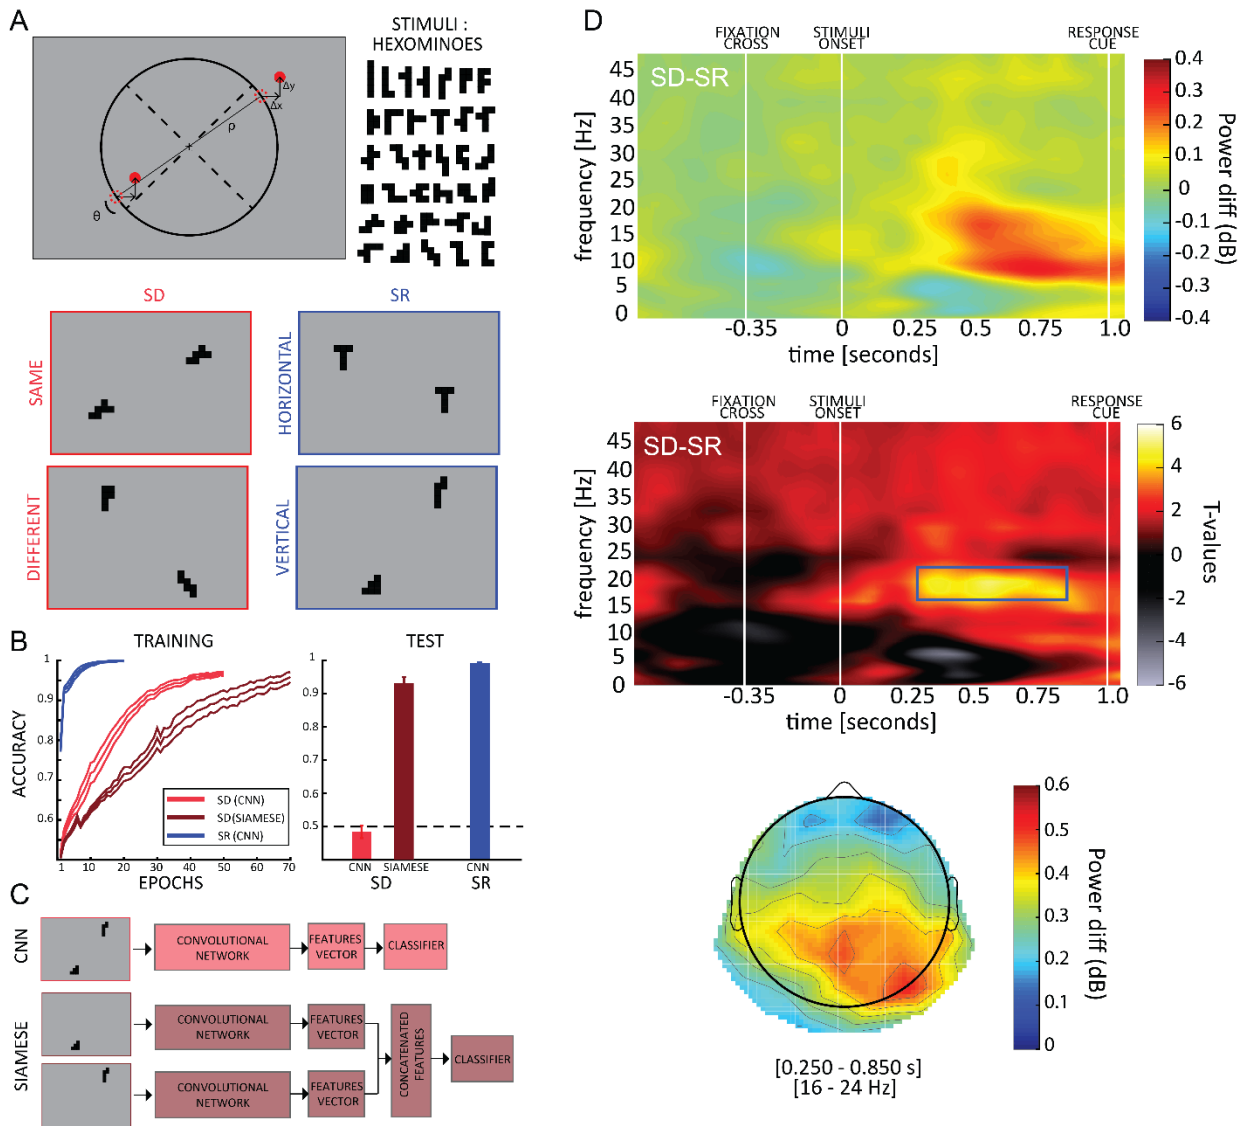


Figure 2.9 – A) The stimuli were the same in the simulations and in the human experiments. The items were displayed at opposite sides of the screen (either 45° and 225° or -45° and -225°). Both item positions were jittered by a random amount in both the x- and y-axes (Δx and Δy in the picture) to make the task non-trivial for human participants (i.e., preventing participants from performing the SR task considering only the position of one item, thus ignoring the SR between the two items). The items used are hexominoes (right panel). The four subplots show some examples of stimuli position for the SD task (left column) and spatial relation task (SR, right column). For the sake of illustration, the ratio between the screen and hexominoes size has been modified (stimuli here look bigger than in the real experiment). B, C) Accuracy of the CNN network on the SD (light red) and SR (blue) tasks, and of a Siamese network trained on the SD task (dark red). The Siamese network mimics segmentation in a feedforward network, by separating the items in two distinct channels of the network (see C). The left panel shows the training curves for each network (accuracy over epochs during training); we stopped the training when the validation accuracy reached 90%. In the right panel, we show the training accuracy at the last epoch and the test accuracy. The latter was evaluated using novel items never used for training, and it reveals that the CNN seems to only learn the required rule for the SR but not for the SD task, as shown in a previous study. Conversely, the Siamese network (CNN with segmentation) can solve the SD task, demonstrating that segmentation can allow the CNN to successfully accomplish this task. In both panels we show average values \pm SE over 10 repetitions using different random initializations. D) The difference between SD and SR power spectra is shown in the first panel. White lines indicate the onset of the fixation cross, the stimuli and the response cue. B, The second panel shows the corresponding t values (when testing the difference against zero). We observed a significant region in the low β band (16–24

Hz), between 250 and 950 ms after stimulus onset. C, The topography of the significant time-frequency window reveals the involvement of occipital-parietal regions.

Overall, this work provided a better understanding of visual reasoning and produced predictions about the involvement of feature-based and spatial attention mechanisms in different reasoning problems, besides illustrating the potential benefits of incorporating brain-inspired mechanisms in artificial networks to achieve human-level ability in visual reasoning.

Chapter III: First independent scientific activity and future directions

MAIN PAPERS OF THE SECTION	Alamia A., VanRullen R. (2023) "A Traveling Waves Perspective on Temporal Binding" <u>Journal of Cognitive Neuroscience</u> 1-9
	Alamia A., Lucie Terral, Malo Renaud D'ambra., VanRullen R. (2023) "Distinct role of forward and backward alpha-band waves in spatial attention" <u>eLife</u> 12 e85035
	Alamia A., Mozafari M., Choksi M., VanRullen R. (2023) "On the role of feedback in visual processing: a predictive coding perspective" <u>Neural Networks</u> 157, 280-287
	Van Bree S., Alamia A., Zoefel B. (2022) "Oscillations or not- why we can and need to know" <u>European Journal of Neuroscience</u>
	Choksi M., Mozafari M., O'May C., Ador B., Alamia A., VanRullen R. (2021) "Predify: augmenting deep neural networks with brain-inspired predictive coding dynamics" <u>Advances in Neural Information Processing System</u> 34, (NeurIPS)
VanRullen R., Alamia A. (2021) "GattaNet: Global agreement for convolutional neural networks" <u>ICANN2021</u>	

First steps as Principal Investigator (2021 – 2023)

Finalizing ongoing projects post-covid pandemic

As for everyone else, the Covid pandemic impacted considerably the ongoing scientific activity, especially experimental studies. In the year 2021, I started the position as CNRS researcher at CerCo. Despite the inability to run experiments and the limitations due to the sanitary situation, I carried on my research using previously collected datasets and focusing on computational work, following my growing interest in Machine Learning. Regarding the experimental work, I had the opportunity to supervise two master students, who contributed to collecting and analyzing the data right before the onset of the pandemic. In this study, we further investigated the role of traveling waves in cognitive functions, focusing specifically on how selective attention modulates traveling waves (Alamia, Terral, Renaud D'ambra, et al., 2023). The study's main goal was to pinpoint and disentangle the distinct roles of alpha-band oscillations. On the one hand, alpha oscillations have been associated with inhibitory functions, showing, for example, that visual attention increases alpha-band power in the hemisphere ipsilateral to the attended location. On the other hand, perceptual echoes (as introduced above) revealed that alpha oscillations positively correlate with visual perception. These two main findings hint at different processes underlying alpha band dynamics. In this study, we aimed to experimentally disentangle the two functionally distinct alpha-band oscillations using an approach based on traveling waves. We analyzed three EEG datasets (one recorded in the lab and two publicly available, figure 3.1) in which participants performed a visual selective attention task, and our results revealed two distinct processes. First, top-down alpha-band waves propagating from frontal to occipital regions increase ipsilateral to the attended location,

and their spectral power correlate positively with alpha-band power in frontal and occipital regions. However, only during visual stimulation we also observed alpha-band waves propagating in the opposite direction (from occipital to frontal areas) and contralateral to the attended location. The presence of forward exclusively during visual stimulation suggested the presence of a separate mechanism related to visual processing and independent from the top-down, inhibitory one. All in all, the results of this study reveal two processes reflected by different propagation directions and confirm the importance of considering oscillations as traveling waves when characterizing their functional role.




Figure 3.1 - Each column in the figure represents a different EEG dataset involving experiments with visual stimulation (left and middle columns) and without visual stimulation (right column). In the upper panels, the net amount of forward (blue) and backward (red) waves is represented along different lines of electrodes, normalized to the midline. The left and central panels reveal an increase (decrease) of forward (backward) waves contralateral to the attended location when participants attended to visual stimulation. The right column shows that when participants attended an empty screen (data from (Foster et al., 2017)), only backward waves were modulated by visual attention, and no effect was observed in the forward waves without visual stimulation. Error bars represent standard errors of the mean. The middle row shows schematic representations of the screen during the tasks: the central panel illustrates the task from (Feldmann-Wüstefeld & Vogel, 2019), where D and T stand for Distractor and Target, respectively. In the task from (Foster et al., 2017), the screen was empty, as the eight circles were not displayed during the task but here illustrate the stimulus positions (Foster et al., 2017). The lower panels represent the lines of electrodes in all datasets.

In addition to this study, I had the opportunity to write two commentary papers during these years. In the first one, in collaboration with Benedikt Zoefel, we argued about the importance of identifying and considering true oscillations as an essential mechanism involved in cognition (Bree et al., 2022). In another work with Rufin VanRullen, we defended the thesis that oscillations are best understood when considering their spatial component (that is, as traveling waves), and we disentangle different cases in relation to temporal binding (Alamia & VanRullen, 2023). Both works are in line with my ongoing scientific interest, which tries to underpin the (oscillatory) dynamics and mechanisms at the heart of cognitive processes. Besides these works in cognitive neuroscience, in parallel, I also pursued some work in machine learning, as described in the next section.

Machine learning and Neuroscience

In my scientific quest, the road that links Cognitive Neuroscience and Machine Learning is a two-way street, where round trips are the rule rather than the exception. In other words, neural networks are a valuable tool to investigate and test specific hypotheses about brain dynamics, and on the other hand, brain-inspired methods could improve and provide more robust algorithms. This mutual benefit is illustrated in one of my recent studies (Alamia, Mozafari, et al., 2023). In this work, we tested the functional role of top-down connections within the predictive coding framework in object recognition under noisy conditions and adversarial attacks. Specifically, we implemented deep convolutional networks (i.e., ResNet and EfficientNet), including Predictive Coding dynamics, using a toolbox we developed in a previous study (Choksi et al., 2021). The activity in each layer is then described by the following equation (and figure 3.2 A):

$$m_i(t+1) = \boldsymbol{\mu} m_i(t) + \boldsymbol{\gamma} F(m_{i-1}(t+1), \theta_i^{ff}) + \boldsymbol{\beta} B(m_{i+1}(t), \theta_{i+1}^{fb}) - \boldsymbol{\alpha} \nabla \text{MSE}(B(m_i(t), \theta_i^{fb}), m_{i-1}(t))$$

where feedback connections θ_i^{fb} were trained for the reconstruction of clean images, whereas forward connections θ_i^{ff} were trained for classification. The last term is the prediction-error as quantified by the gradient of the Mean Squared Error between the predicted and the actual activity (Rao & Ballard, 1999). Importantly, a specific hyperparameter (in bold in the equation) modulates each term's functional relevance (i.e., weight), namely the memory, forward, backward, and prediction error. After a first feed-forward pass that initializes the network (i.e., without memory or top-down connections), the activity is then determined at each time step by the contribution of each term. Interestingly, we first showed that such a network implementing

Predictive Coding dynamics has better accuracy than a purely forward network with the same number of parameters (figure 3.2B-D, right column).




Figure 3.2 – A) Predictive Coding dynamics. Architecture of the shallow model, composed of three convolutional layers and two fully connected ones. Each layer's activity is a combination of four

terms at each time-step, shown in different colors. Each term is modulated by a specific hyper-parameter. B) Values of hyper-parameters and accuracy of the deep predictive coding networks for PResNet18 with hyper-parameters that are trained on CIFAR100-C images, and are shared across layers. (C) PResNet18 and (D) PEffNetB0 with hyper-parameters that are trained respectively on CIFAR100-C and ImageNet under Gaussian and Salt & Pepper noise, and are separate for each layer. Plots in the first column show the hyper-parameters as a function of the layers (PCoder) from input to top layer under medium noise level. The circles indicate layer with maximum feedback error. In middle columns, relative values of hyper-parameters are plotted across noise levels. In case of separate hyper-parameters, the PCoder with maximum feedback error is shown. For each noise level, the accuracy difference to the first time-step (i.e. feedforward backbone) is depicted in the last column. Error bars show standard error of the mean (SEM) over 19 CIFAR100-C noise types. In all cases, the networks achieve accuracy gain by utilizing more feedback and forward error as the noise severity increases.

Specifically, we showed that such a network is more robust to noise and adversarial attacks. More importantly, however, we also demonstrated the computational role of each term in different experimental situations, modulating the noise of the input. After having trained the network parameters θ_i^{fb} and θ_i^{ff} , we then optimized and interpreted the hyper-parameters controlling the network's recurrent dynamics. In other words, we let the optimization process determine whether and how much each term (i.e., top-down connections and predictive coding dynamics) is functionally beneficial with different noisy conditions. Across different model depths and architectures (3-layer CNN, ResNet18, and EfficientNetB0) and against various types of noise (CIFAR100-C), we find that the network increasingly relies on top-down predictions as the noise level increases (figure 3.2 B-D). Interestingly, we also observed that when we train a separate hyperparameter for each layer in deeper networks, this effect is most prominent at lower layers (figure 3.2D). All in all, the results of this study proved beneficial for both Neuroscience and Machine Learning: on the one hand, we provide novel insights by confirming the functional role of feedback connections in sensory systems; on the other hand, we demonstrate how deep neural networks can be improved in terms of robustness using brain-inspired dynamics, such as Predictive Coding.

On a similar note, another example of how neuroscience can inform machine learning is a recent work realized with Rufin VanRullen, who is the first and leading author of the study (VanRullen & Alamia, 2021a). In this study, we equipped a convolutional network with attentional mechanisms loosely inspired by biological systems, that is, a distinct and unified network that receives inputs from and exerts modulatory influence on the entire hierarchy of visual regions. Here, we tested networks with such an attentional system, in which each spatial position produces a key-query vector pair for every network layer, and all queries are then pooled into a global attention query. Eventually, the match between each key and the global attention query modulates the network's activations—emphasizing or silencing the locations that agree or disagree (respectively) with the global attention system. We proved that this brain-inspired Global Attention Agreement network (GAttANet) improves accuracy over the corresponding baseline for various convolutional backbones (from a simple 5-layer toy model

to a standard ResNet50 architecture) and in different datasets (CIFAR10, CIFAR100, Imagenet-1k). This work was inspired by an ongoing project, which describes a similar architecture implemented with complex-valued units, in which the global system is used to synchronize different neurons in the network. This project will be described in detail in a following section.




Figure 3.3 –The GAttANet architecture from (VanRullen & Alamia, 2021a). Each layer’s activation, of a classic convolutional network, is projected into the global attention system. According to the blackboard theory, here, as in the Pulvinar, a global saliency map is created via global agreement (using a novel architecture known in machine learning as ‘transformer(Vaswani et al., 2017)’). Such an attention agreement map directly modulates the layer’s activations on the next time step, thus closing the feedback loop.

ERC Starting Grant ‘OSCI-PRED’ (2023 – 2028)

In the following five years (2023-2028), my scientific activity will be focused on carrying out the project I proposed in the ERC Starting Grant ‘OSCI-PRED,’ which officially started in March 2023. Based on the proposal, the following part describes the motivation and main goals of the project, divided into four different working packages. This project follows my previous work on traveling waves and predictive coding, tackling these questions from an interdisciplinary angle, combining experimental and computational methods, as for most of my previous work.

Introduction and project motivation

The human brain has remarkable processing capabilities unmatched by artificial models: around 10^{11} neurons, organized in anatomically and functionally distinct regions, allow very complex sensory-motor behaviors, such as playing a tennis match or performing a virtuosic piano sonata. One of the most exciting questions in neuroscience is understanding how the

brain integrates and coordinates the activity between distinct brain regions. Several studies point to oscillations as the critical mechanisms regulating communication between different areas (Engel & Singer, 2001; Singer, 2013). For example, the synchronous activity allows to dynamically modify the communication pattern between brain regions, irrespective of the connections' anatomical structure, and flexibly adapt to the current task. Experimental evidence confirmed that oscillatory dynamics modulate the effective connectivity of populations in the cortex (Aertsen et al., 1989; Sohal et al., 2009), shape plasticity between neurons (Huerta & Lisman, 1995), and play an active, causal role in several cognitive functions (Gulbinaite et al., 2014; Koshy et al., 2020). Yet, most research on neural oscillations investigated their functional role considering zero-lag synchrony, namely when oscillations are aligned in time in different cortical regions. Recent studies (Alamia & VanRullen, 2023; Davis et al., 2020; Muller et al., 2018; Pang et al., 2020) support the critical idea that oscillatory dynamics are best understood when interpreted as Oscillatory Traveling Waves (OTW) propagating across cortical regions, suggesting that OTW are more than just a by-product of interactions between brain regions, and may play an essential role in yet-to-investigate cognitive functions. However, what are the mechanisms generating them, and what is their functional role remain poorly understood. In one study presented above (Alamia & VanRullen, 2019a), we demonstrated that oscillations can be generated via the interactions between different brain regions within the Predictive Coding (PC) framework. Additionally, several studies –including my own described above– point to a tight relationship between oscillatory processes and Predictive Coding. For instance, when interacting with a dynamic environment, it is necessary to predict *what* will happen and *when* (Giraud & Arnal, 2012), and brain oscillations may be a suitable candidate for both functions (Arnal et al., 2011; Bastos, Litvak, et al., 2015). All in all, one of the key goals of OSCI-PRED is to combine these two main fields in Neuroscience: Predictive Coding and Brain Oscillations, to propose a unified and multi-scale model to address the ambitious question of whether predictive coding can be the unified framework that describes the role of oscillatory waves in cognition. If successful, OSCI-PRED will advance the interpretation of oscillations as traveling waves based on their spatiotemporal features rather than only their temporal component. Additionally, OSCI-PRED will move the field toward a more integrative and holistic interpretation of brain dynamics: the core idea is to apply the same principles (based on PC) at different scales to explain distinct cognitive functions. Presumably, Predictive Coding won't be the 'theory of everything' in Neuroscience. However, OSCI-PRED's results will still provide an essential understanding of at which scale these mechanisms are relevant to understanding brain dynamics: it could be the case that PC is crucial at the scale of cortical regions but less so at the level of each cortical column or vice versa. Ultimately, one of the main ideas of this project is to direct cognitive neuroscience towards a path based on solid theoretical models that can overcome an experimental approach

based on correlational studies. As in modern theoretical physics, in which the Higgs Boson was theorized before being observed, this project aims to push Cognitive Neuroscience towards an approach theory-driven, in which computational and theoretical simulations can generate precise predictions about the relationship between neural mechanisms and cognitive functions even before collecting the data. Arguably, the ultimate and most compelling ambition in cognitive neuroscience should be to provide a coherent and holistic interpretation of brain dynamics and cognitive functions.

Main goal and overall structure of the project

OSCI-PRED's primary goal is to investigate the functional role of oscillatory traveling waves by proposing a multi-scale computational model grounded in the well-defined framework of Predictive Coding. The proposal combines a biologically plausible model with carefully designed EEG experiments to assess the model's predictions and specific hypotheses related to OTW and Predictive Coding. In addition, the proposed modeling framework can be effectively applied to deep learning architectures to improve the current models with brain-inspired implementations based on PC and oscillatory dynamics. Overall the project comprises four Working Package (WP), each investigating the role and the mechanisms involved in the propagation of the traveling waves from different perspectives. Each WP tests specific hypotheses, spanning from different scales (macroscale in WP1 and mesoscale in WP2) and combining computational modeling (WP1, WP2, and WP3) with EEG experiments (WP2 and WP4). In the **first WP**, we test the hypothesis that OTWs synchronize the activity of different brain regions (macroscale) via multiplexing mechanisms. In the **second WP**, we test how the cortical column (mesoscale) can generate oscillations at various frequency bands and optimally integrate top-down and bottom-up information within a PC framework. We will also test experimentally (via EEG recordings) how OTW with distinct frequency bands can be interpreted in light of PC. In **WP3**, we'll investigate the role of a subcortical structure, the Pulvinar, in generating and propagating OTW through the cortex. In both WP2 and WP3, I plan to extend the proposed architectures to deep learning implementations, incorporating brain-inspired dynamics (i.e., PC and oscillations) in current deep learning models. Lastly, **WP4** tests the functional role of OTW experimentally in high-level cognitive functions –such as attention and conscious perception- based on the predictions from WP1 and some specific hypotheses about OTW and PC. Below a summary of the 4 WPs composing OSCI-PRED and a detailed explanation.

WP 1 - Combining Predictive Coding with cross-frequency coupling.

Cross-frequency coupling is one of the main oscillatory mechanisms involved in cortical processes, which describes how different frequencies coordinate and produce temporal synchronization over time. This mechanism enables *multiplexing*, consisting of the

combination of slow (alpha/theta, [4-12Hz]) and fast (gamma range, [$>30\text{Hz}$]) oscillations, where the slow rhythms act like a metronome to synchronize the local gamma-band activity of distinct regions, involved in features representation (fig.3.4A).




Figure 3.4. **A.** Sensory information is hierarchically processed by different cortical regions when attending to a visual stimulus. Multiplexing describes how the fast oscillations, which extract sensory features at a local level, are coordinated by slower oscillations, which act as a pacemaker to synchronize their activity over time. **B.** Multiplexing, or cross-frequency coupling, has been proposed in two versions. In the first one, named theta-gamma mode, a different visual item is processed in each cycle of the fast oscillations, thus allowing an explorative mode. In the second one, called Communication-Through-Coherence (CTC), only one visual item is processed at each fast cycle, resulting in a focused or selective mode. Computational simulations showed that the amplitude of the slow oscillation allows switching from one mode to the other, reconciling the two theories and providing flexibility to the system (McLelland & VanRullen, 2016).

The interplay between fast and slow frequencies effectively modifies the cortical connectivity between regions, allowing the orchestration of overlapping neuronal ensembles over time, thus generating unequivocal and distinct representations. Several theories have been proposed to describe this process in the literature. One theory, named ‘theta/gamma discrete phase code’ (Lisman & Jensen, 2013), aimed at explaining the ordering of items in working memory and items saliency and was initially developed from experimental observations in rodents’ hippocampus. It suggests that neurons participating in a given ensemble fire together during one gamma cycle, thus encoding several items during the slow-wave cycle (fig.1B, left panel). Other theories proposed that interareal connectivity requires high *coherence*, such as “Communication Through Coherence” (CTC). This metric quantifies the frequency and amplitude of synchronous activity between pools of neurons in different regions (Fries, 2005, 2015). The main intuition is that effective communication between areas is achieved by selecting a specific ensemble of neurons with optimal gamma-phase relation (fig.3.4B, right panel). Coherence-based connectivity also received experimental support (Fries, 2015; Michalareas et al., 2016; Rohenkohl et al., 2018), even though recent work proposed an interesting alternative, in which the synaptic activity and the connectivity between areas are responsible for driving inter-areal coherence, thus determining the optimal communication between areas (Schneider et al., 2021; Vinck et al., 2023). Generally, different theories relate multiplexing to two main modes: theories such as CTC describe a selective mechanism in

which only one stimulus is processed during each slow-wave cycle, whereas other approaches, such as the theta/gamma code, characterize an exploratory mode as several items are processed per each cycle (fig.3.4B). Remarkably, these opposite approaches have been recently reconciled in a computational model that integrates both mechanisms (McLelland & VanRullen, 2016). These simulations demonstrate that tuning one parameter modulating the global inhibition of the slow oscillation produces the switch between the two modes. However, such a proposed model doesn't provide a biological interpretation of the generation of the slow rhythms representing global inhibition. In the first working package, we aim to investigate whether forward and backward alpha-band waves, generated via PC mechanisms, could participate in multiplexing dynamics, coordinating the fast gamma-band activity of different cortical regions across the visual system, as described in cross-frequency coupling mechanisms. Here, I test this hypothesis from a computational point of view, assessing whether an architecture based on PC can generate and sustain such functional interplay between slow and fast oscillations.

Multiplexing: Gamma-band surfing Alpha-band oscillations. As mentioned above, multiplexing describes the coupling of fast and slow oscillations: the local, fast rhythms are coordinated by the phase of the slower oscillation, which acts as a pacemaker to synchronize the activity over different brain regions. This interplay between frequencies orchestrates overlapping neuronal ensembles, allowing unequivocal and distinct representations to occur over time. The architecture proposed in this WP is similar to previous work (McLelland & VanRullen, 2016), in which each brain region is composed of a pool of neurons, implementing a Pyramidal Inter-neuronal Network Gamma mechanism (PING) which generates fast, local activity (Buzsáki & Wang, 2012). Here we aimed to combine this mechanism with PC dynamics. According to PC principles, each region's activity aims to 'explain away' the activity in the lower region. In other words, higher levels predict the lower levels' activity, and the difference between the prediction and the actual activity (the prediction error) is used to update and correct the next prediction. Considering the temporal delay between levels (ΔT in fig.2.5A), such interplay between predictions and prediction-errors generates oscillatory dynamics whose frequency depends on the value of the temporal delay. Our previous work demonstrated that such oscillatory activity propagates as a traveling wave through brain regions. This WP aims to show that such alpha-band oscillation acts as a pacemaker to synchronize the activity of faster gamma-band oscillations, which are locally generated via PING mechanisms (fig. 3.5).

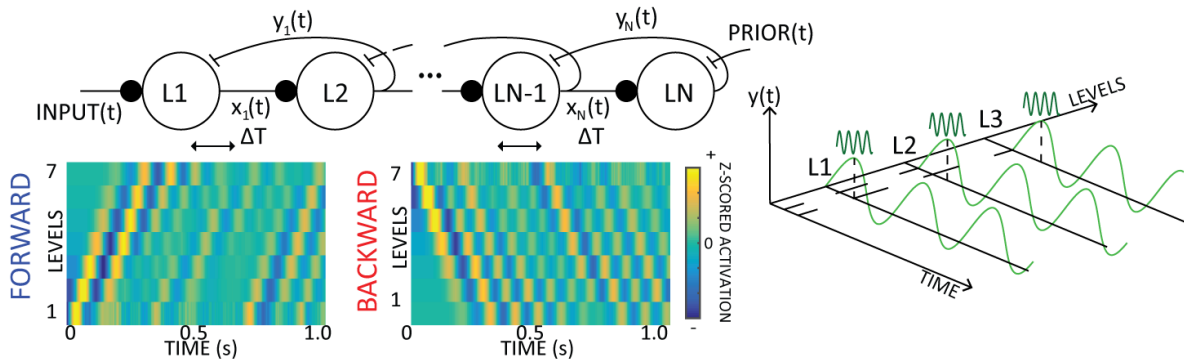


Figure 3.5 -Computational models and experimental predictions. – Schematic representation of the model based on PC principles generating oscillatory traveling waves. The lower left plot shows the activity over time, and different levels: the slow oscillations propagate as a wave with a phase shift (dashed lines) and synchronize the activity of faster oscillations in each level (L1, ..., LN). The panel on the right represents the slow oscillations that propagate from forward or backward depending on the cognitive state of the input (adapted from our previous study (Alamia & VanRullen, 2019a))

Thus, the key objective of WP1 is to reconcile in a unique architecture the fast gamma-band rhythms generated in each region via the PING mechanism (or other models of fast oscillations, such as those proposed by Izhikevich, 2007) with the slow alpha-band oscillations rising from PC dynamics between areas. In our implementation, we will demonstrate that the fast activity is phase-locked to the slower oscillations, and such slow-wave propagates through the hierarchical model, forward or backward, depending on the cognitive state of the system (e.g., stimulus' presence).

WP 2 - Going deeper: characterizing laminar dynamics.

In the second working package, we aim to expand the model described in WP1, shifting the focus from global to local dynamics. Each region of the visual system encodes and integrates relevant visual features, such as spatial orientation, phase, and frequency. Yet, an outstanding question is whether there is a basic unit/module of cortical computation. Several studies identified this anatomical and functional module in the cortical column and described its activity via the canonical microcircuits. From a functional point of view, each cortical column codes for a specific visual feature and has its visual receptive field. Each brain region is composed of several cortical columns encoding particular visual features. Interestingly, its organization has been shown to be more efficient than a randomly connected neuronal network (Haeusler & Maass, 2007). Notably, the laminar interaction described by the canonical microcircuits can be elegantly framed in the light of Bayesian processing (George & Hawkins, 2009), thus consequently supporting a PC interpretation. According to electrophysiological recordings and anatomical tracing, each column in the sensory cortex receives input projections in the granular layer (i.e., L4 - see Figure 3.6A). From a PC interpretation (Bastos et al., 2012), these inputs are mostly feedforward Prediction Errors (PE) coming from lower regions or the actual sensory information in the case of the first column (in the human brain, perceptual inputs are

preprocessed by a series of pre-cortical nuclei, e.g., the Lateral Geniculate Nucleus, LGN). From layer L4, PE travel to supragranular layers (L2/3), which encode the expectations about the specific visual features (Haeusler & Maass, 2007). From layers 2/3 depart strong feedforward connections to the infragranular layers (L5/6), which generate the predictions sent to lower regions and the granulate layer L4. Figure 3.6A shows a summary of these connections. In sum, the goal of WP2 is to expand the architecture described in WP1 by introducing microcircuit dynamics in each brain region and modeling the cortical column organization. This modeling approach will incorporate for the first time in a unique model strong evidence regarding cortical processes and brain oscillations (Arnal et al., 2011; Bastos, Litvak, et al., 2015; Bastos, Vezoli, et al., 2015; Giraud & Arnal, 2012). In addition, it will allow us to test three hypotheses about the role of the cortical column in the generation, propagation, and integration of oscillatory dynamics through different cortical regions.

Hypothesis I: The emergence of different rhythms in each cortical layer is a consequence of PC. The canonical microcircuit can be elegantly described in the light of Predictive Coding (Bastos et al., 2012; G. B. Keller & Mrcic-Flogel, 2018), producing compelling hypotheses about its oscillatory dynamics. Indeed, several experimental observations revealed that cortical layers have different spectral properties: consistent empirical evidence has shown that the superficial layers (i.e., L2/3) have strong gamma-band oscillations (Buffalo et al., 2011; van Kerkoerle et al., 2014; van Pelt et al., 2016), whereas deeper layers (L5/6) show prominent beta-band power (Michalareas et al., 2016). This difference in the spectral profile can be explained in the light of PC formulation (Bastos et al., 2012), and our simulations aim to demonstrate it empirically. Specifically, predictions are the results of the integration over time of PE. The natural effect of such accumulation is eventually the suppression of high-frequencies (the same effect is a characteristic of all Bayesian filtering schemes, including PC (K. Friston, 2008)). Here, I plan to test this hypothesis by implementing PC dynamics in the cortical columns, as shown in Figure 3.6. As mentioned above, the differential equations driving the laminar dynamics follow PC principles (Rao & Ballard, 1999), similar to those outlined in my previous work (Alamia & VanRullen, 2019a). A systematic parametric space analysis will detail the system's behavior (i.e., time constant and temporal delays between regions), eventually focusing on biologically plausible values. Importantly, each columnar layer entails distinct excitatory and inhibitory neurons whose functional role rests on PC principles, as shown in Figure 3.6 (superficial layers encoding PE, deeper layer encoding predictions). Besides demonstrating that PC mechanisms generate different spectral profile in each layer, our model advances further the interpretation that oscillations in each layer are traveling waves. This idea generates a novel hypothesis (i.e., hypothesis II), which will be tested in EEG recordings, as detailed below.

Hypothesis II: Gamma-band OTW as a neural signature of Prediction Errors. Next, we aim to broaden our understanding of TW's functional roles in EEG recordings, considering the results of our simulations. As mentioned above, previous studies (Bastos et al., 2012; Michalareas et al., 2016) indicate a relationship between bottom-up prediction errors (PE) and fast gamma oscillations and between alpha/low-beta bands rhythms and top-down predictions. Notice that not all gamma oscillations are related to PEs: fast oscillations encoding sensory features – such as those described in WP1- are not directly related to PEs. Furthermore, recent theoretical and experimental work questioned the relationship between prediction-errors and gamma oscillations, specifically in lower visual regions, such as V1 (Vinck et al., 2023; Vinck & Bosman, 2016). Alternatively, they propose that gamma synchronization reflects the extent to which the (visual) input can be predicted from the surround (i.e., outside of the receptive field). In other words, Vinck and colleagues reverse the relationship, proposing that gamma oscillations reflect predictions rather than prediction-errors. One possible interpretation that reconciles the two view is that gamma oscillations reflect low-level predictions at lower level in the visual system but reflects precision-weighted prediction-errors at higher level in the visual hierarchy: this would explain why white noise (Jia et al., 2013) or random motion (Kruse & Eckhorn, 1996) don't induce increase in gamma oscillations, as both have low precision despite being unpredictable, thus not generating reliable (or informative) prediction-errors. Whether gamma-oscillations reflect predictions or prediction-error, we will test these supposedly exclusive hypotheses by quantifying participants' OTW in EEG recording during an oddball paradigm experiment, in which an unexpected event violates a series of predictable ones (such as a change in tone in a rhythmic repetition). We expect that the unanticipated event lead to an increase (or decrease) of bottom-up (i.e., forward) TW in the gamma-band frequency (Fig. 3.6B), and we aim at localizing such transient of gamma oscillations in the lower or higher regions of the visual hierarchy. Possibly, such unpredictable events elicit an increase in prediction errors (Wacongne et al., 2012), reflected in a rise in forward OTW in the gamma band in higher visual areas, followed by an increment in top-down TW in the alpha/beta band, related to the update in the top-down predictions.

Hypothesis III: The laminar dynamic ideally integrates different frameworks describing functional connectivity. The visual system includes several top-down and bottom-up connections, linking cortical regions throughout the visual hierarchy. Several frameworks have been proposed to explain the functional connectivity within these streams. For example, PC suggests that top-down feedback carries predictions, whereas bottom-up connections propagate prediction-errors, which are used to update the upcoming predictions. Alternatively, some authors proposed that temporal differences of top-down predictions may replace bottom-up prediction-error units with the mediation of thalamic nuclei (especially the Pulvinar, see WP3, and O'reilly et al., 2021). Finally, other authors proposed the reversed dynamic, where

predictions flow in a bottom-up fashion (rather than top-down) and are updated based on the error propagated from top-down connections (Heeger, 2017; Sabour et al., 2017). In sum, different approaches have been proposed to describe the functional role of bottom-up and top-down connections. Here, I plan to test the hypothesis that the laminar dynamic connectivity (i.e., the canonical microcircuits) is computationally ideal for efficiently integrating top-down, bottom-up, and lateral connections, reconciling potentially different frameworks in one anatomical and functional structure. The anatomical organization in different layers allows each column to disambiguate between endogenous top-down and exogenous bottom-up signals precisely. Specific populations carry out this differentiation in distinct layers that receive each stream separately and are characterized by different spectral profiles (see hypothesis I above). Notably, the cortical column connectivity provides a framework that is suitable for various proposals: as shown in Figure 3.6C, the connectivity between layers is compatible with a PC implementation (as described in (Bastos et al., 2012)), but also with an implementation proposing the opposite dynamics (as described above, and more precisely, in (Heeger, 2017)). Both frameworks postulate the computation of prediction errors, defined as a difference between the expected and the actual activity. Moreover, Figure 3.6C proposed an architecture suitable for deep learning (as shown in our previous work (Alamia, Mozafari, et al., 2023; Choksi et al., 2020)) that reconciles the two approaches via a structure inspired by the cortical column, which allows a direct calculation of E_{TD} and E_{BU} , i.e., top-down and bottom-up errors, respectively. The network's training would be based simply on optimizing a loss function defined as the sum of the two errors. Interestingly, the role of the bottom-up and the top-down errors may be tuned by an attentional mechanism, allowing the model to switch between modes, depending on whether the task at hand required more contextual top-down information or bottom-up sensory processing (similarly to the role of the global inhibition in WP1). In this line, our recent study (described above) showed that in deep networks (i.e., ResNet and EfficientNet) embedded with PC dynamics, the optimization processes privileged top-down feedback when processing noisy images (Alamia, Mozafari, et al., 2023).




Figure 3.6 – The cortical column and the canonical microcircuits. **A.** A 2-level predictive coding model comprises a lower and a higher region to the left. $X(t)$ represents prediction errors (PE), $y(t)$ the predictions. A schematic representation of the columns composing each brain region and the canonical micro-circuit describing their dynamics is shown to the right. Each brain region comprises several cortical columns composed of three functional layers: L2/3, L4, and L5/6. The connection between the layers follows PC principles, in which the PE in superficial layers (L2/3) is used to update predictions in deep layers (L5/6). **B.** Experimental predictions about the amount of oscillatory waves in different frequency bands. Our model would predict an increase of forward gamma-band oscillations (>25Hz, highlighted in dark green) and a subsequent decrease in backward alpha-band oscillations (~10Hz, highlighted in light green) after the onset of a surprising or unexpected event. **C.** A schematic representation of a neural network in which a cortical column inspires each layer structure (notice that the layers' order in the architecture does not correspond to their anatomical counterpart). At each layer, top-down (in green) and bottom-up (in blue) predictions are compared to the current activity (in the middle) to compute the respective errors. W_F and W_B represent the synaptic connections that are trained by optimizing the energy function defined in red as the sum of the two errors.

WP 3 - Integrating information across cortical regions: the role of the Pulvinar.

In WP1 and WP2, as well as in most studies investigating computational models of cognitive functions, the main focus is on cortical processes and the role of feedback and feedforward connections between cortical regions. However, increasing experimental evidence points to sub-cortical structures –such as the thalamus- as important factors to consider when modeling cognitive processes. Besides the cortical dynamics, thalamic nuclei, such as the Pulvinar and the reticular nuclei, generate oscillatory dynamics, especially in the alpha rhythms. It is logical to consider the possible role of the cortical-thalamic loop in generating and coordinating oscillatory traveling waves through the cortex. This WP will focus specifically on one thalamic nucleus, namely the Pulvinar. Previous studies demonstrated the critical role of the Pulvinar in several cognitive and perceptual tasks, coordinating the activity along feedforward and

feedback pathways (Jaramillo et al., 2019). It is crucially involved in spatial attention (Saalmann et al., 2012; Zhou et al., 2016), in visual working memory (Rotshtein et al., 2011), and its neural activity correlates with confidence during decision-making tasks (Komura et al., 2013). Interestingly, the Pulvinar is reciprocally and topographically connected with several cortical regions, including visual areas at different hierarchical levels. However, a global theory describing the role of the Pulvinar is still lacking. In the early 90s, David Mumford proposed a compelling hypothesis about the role of the Pulvinar in perception, namely the '*active blackboard*' theory (Mumford, 1991), which aligns with current Predictive Coding theories. This theory recently gained renewed interest in several computational and experimental studies (Ketz et al., 2015; O'reilly et al., 2021). The core idea is that each cortical region is a generative model whose activity represents the sensory input. Each cortical area projects its specific representation of the sensory data into the Pulvinar, which serves as a '*blackboard*' to integrate all cortical projections. In other words, the Pulvinar acts as a central integration state, which combines all cortical representation and feeds back in the cortex the most likely interpretation of the data. Additionally, the interaction between cortical regions and the Pulvinar is modulated by alpha oscillations, as demonstrated by electrophysiological recordings (Halgren et al., 2019; Lopes da Silva et al., 1980). In this WP, we aim to test the hypothesis that cortical-thalamic interactions are responsible for generating oscillatory alpha-band traveling waves in the cortex, which coordinate gamma-band activity in each cortical region (see WP1). Remarkably, previous computational simulations demonstrated that the slow-oscillatory activity of one area (i.e., the Pulvinar) could modulate the communication between two different areas (i.e., two cortical regions), and the phase of such slow oscillations determines the directionality of interareal communication (Quax et al., 2017). In the first step of this WP, I plan to achieve this dynamic within the Predictive Coding framework. Specifically, I will implement a model in which PC mechanisms generate slow oscillations due to the delayed communication between each cortical region and the Pulvinar (i.e., within the thalamic-cortical loop, Alamia & VanRullen, 2019). I expect that such cortico-thalamic interactions will generate OTW, propagating forward or backward depending on the timing of the cortical-pulvinar interaction: either the top-down prior or the sensory bottom-up activity will lead the phase difference, thus defining the direction of the waves' propagation. In case the model confirms such hypotheses, I will combine in a unique model the cortical dynamics described in WP1 with the cortico-thalamic detailed above (see Figure 3.7A) to test the role and the interactions between oscillatory waves generated by these distinct mechanisms (cortical vs. cortico-thalamic) (Halgren et al., 2019). In our simulations, the Pulvinar has the role of integrating information from different cortical regions, combining all representations to provide feedback to each cortical area, based on the most likely interpretation of the data. Interestingly, this architecture share remarkable similarities with the GattaNet implementation described above, in which the global attention space plays

the role of the Pulvinar in combining the sensory information and estimating saliency maps (fig.3.7B). In this WP we aimed at going further by introducing in such deep learning architecture 1) predictive coding dynamics, as we proposed in (Alamia, Mozafari, et al., 2023; Choksi et al., 2020), and 2) oscillatory activity, which would coordinate the network activity over time-steps. In sum, this WP's goal is to integrate the framework proposed in WP1 and WP2 by implementing the thalamic-cortical loop. Importantly, this model doesn't aim at precisely describing the details of the anatomical connectivity and simulating the functional activity of the Pulvinar faithfully. Instead, we plan to test whether the Pulvinar orchestrates cortical activity via oscillatory traveling waves, and whether this dynamic can be described within PC principles.

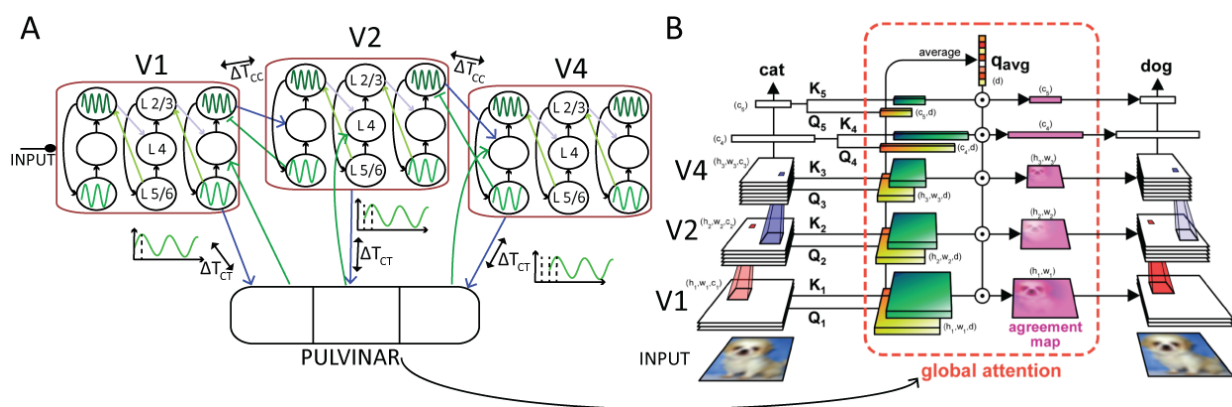


Figure 3.7 – Cortical-thalamic loop and the blackboard theory. A. An illustrative representation of the entire model, showing cortical-cortical and cortical-thalamic loop. All interactions are based on PC mechanisms, generate a rich spectrum of oscillations in the alpha and low beta band (Alamia & VanRullen, 2019a). A pool of cortical columns characterizes each cortical region, as described in WP2, generating faster, local, gamma-band oscillations. **B.** The GAttANet architecture (from (VanRullen & Alamia, 2021b)). Each layer's activation, of a classic convolutional network, is projected into the global attention system. According to the blackboard theory, here, as in the Pulvinar, a global saliency map is created via global agreement (using a novel architecture known in machine learning as 'transformer' (Vaswani et al., 2017)). Such an attention agreement map directly modulates the layer's activations on the next time step, thus closing the feedback loop.

WP 4 - Attention and conscious perception modulate oscillatory traveling waves.

As shown in previous WPs, especially the first one, interpreting the functional role of OTW from a PC perspective generates novel hypotheses. The goal of WP4 is to test them experimentally. Specifically, I plan to investigate how cognitive functions, such as attention and conscious perception, modulate the amount of forward and backward OTW in different frequency bands. In the following, I'll describe three hypotheses I plan to test via EEG recordings in healthy human participants. These hypotheses are derived from the computational model detailed in

previous WPs, or are formulated to provide additional experimental evidence about the functional role of OTW within a PC framework.

Top-down alpha-band inhibition modulates the exploration-exploitation trade-off. As described at the beginning of this section, the multiplexing mechanisms can integrate two working modes: an explorative and a selective one. In the former, distinct items are processed in each slow cycle, whereas in the latter, the same item is processed at each cycle. Interestingly, the global inhibition, which in our implementation based on PC principles corresponds to high-level predictions (i.e., the priors), regulates the switch between the two modes (McLelland & VanRullen, 2016). Remarkably, the PC framework provides a meaningful interpretation of such two modes' switch: higher global inhibition implies prior with large values (or -in agreement with a probabilistic interpretation- a small variance, Feldman & Friston, 2010), and this drives the model in a selective mode; vice versa, when the priors play a minor role in perception, the emergent mode becomes the exploratory one. This conclusion agrees with recent studies showing that the magnitude of the prior does indeed play a role in influencing the exploration-exploitation dilemma also in sensory systems (Feldman & Friston, 2010; K. Friston et al., 2015; Steyvers et al., 2009; Summerfield & Egner, 2009). This appealing account of how prior knowledge plays a role in tuning different modes aligns with previous studies (De Lange et al., 2018; Kok et al., 2013). Here, we plan to test this prediction in a simple experimental paradigm (figure 3.8A). Participants keep fixating on a cross in the center of the screen. In one condition (focused condition), one arm of the cross cues participants' attention to a specific position out of the four corners of the screen. In the other condition (neutral condition), participants pay attention to all four positions. After a prestimulus interval between 1400ms and 2000ms, four Gabor patches appeared at the four positions. Participants are instructed to report the orientation of the Gabor patch indicated by the fixation cross during the stimulus presentation. The target patch can be the same as the one indicated by the cue (valid condition) or a different one (invalid condition). We will assess attention manipulation by comparing the participants' performance between conditions: we expect a better accuracy in the valid than in the neutral condition, and we expect the worst accuracy in the invalid condition. Importantly, we will measure traveling waves in the focus and neutral conditions during the prestimulus phase (when participants pay attention to either one specific location or several ones). This will allow us to contrast the effect of attention on traveling waves: our model would predict an increase of top-down alpha-band traveling waves in the focused condition (i.e., attention focused on a single, selected item), as opposed to the neutral one (i.e., attention diffuse over several items). Traveling waves will be measured in the alpha-band range [8-13Hz] along the electrodes midline, as well as in the contra- and ipsi- lateral hemisphere, to assess a potential lateralization effect as in our previous study (Alamia, Terral, d'Ambra, et al., 2023).

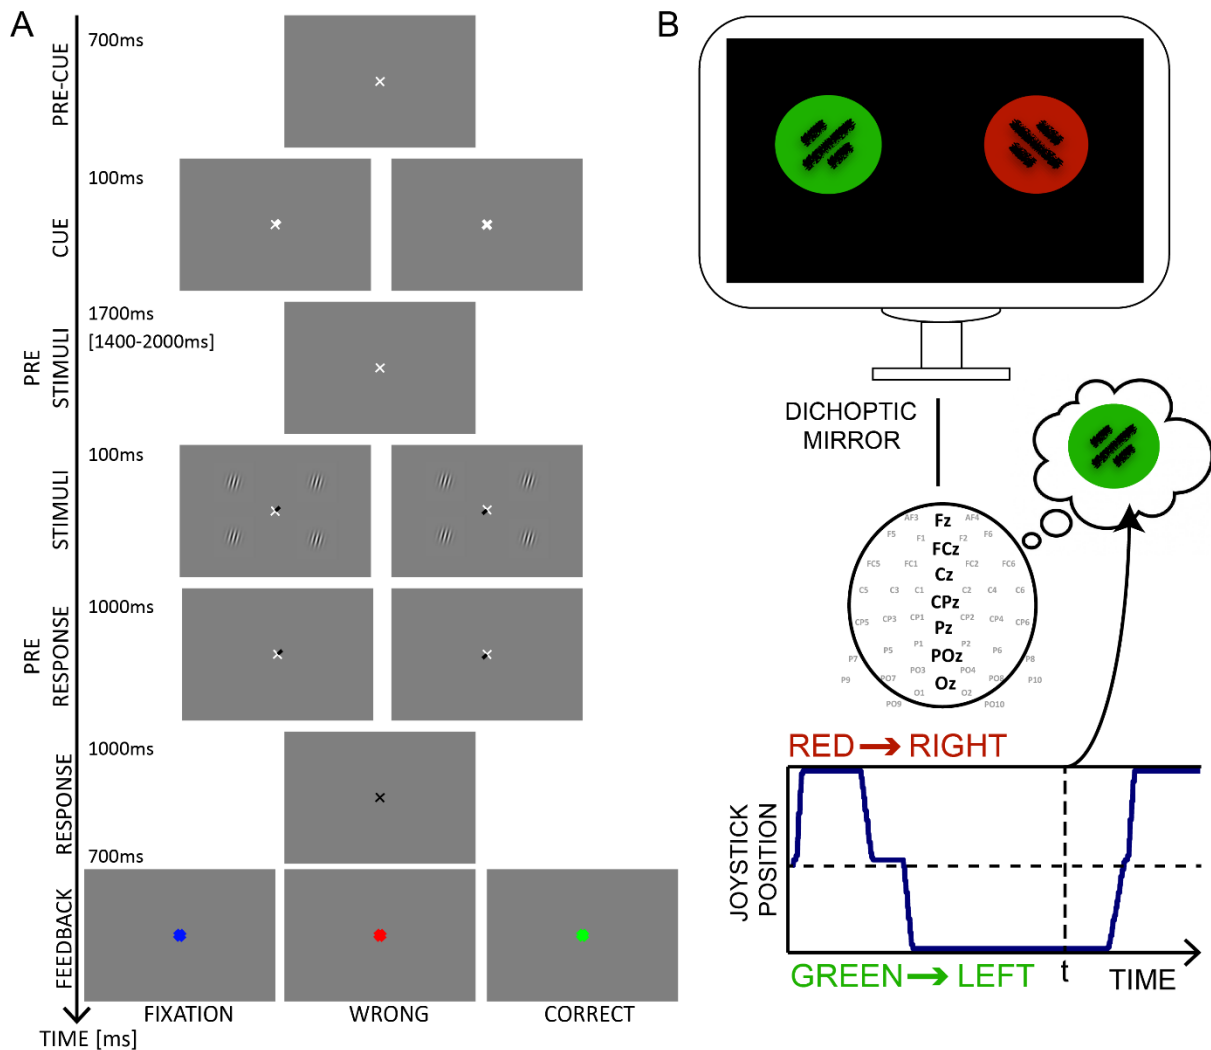


Figure 3.8. - Visual attention, conscious perception, and traveling waves. A. Schematic representation of the experimental design to test the focused and neutral condition. Keeping fixation throughout each block, participants will report the orientation of one of four Gabor patches. We manipulate their attention: in one condition they'll attend only one location, whereas in the other condition they'll attend all of them. Our model would predict a difference in the amount of backward waves in the alpha-band, considered a potential marker indicating switches between focused and diffused attention. **B.** Experimental design for the binocular rivalry experiment - from(Luo et al., 2021). Participants will stare at the screen through a set of dichoptic mirrors that projected the left and right sides of the screen to the left and right eye, respectively. Two Gabor patches of different colors and orientations are placed on the two sides of the screen. Participants reported which patch they perceived by moving a joystick to either side, each associated with a stimulus (pseudo-randomly between participants, consistent across blocks and sessions).

Conscious perception of bistable stimuli depends on top-down oscillatory mechanisms. In this last experiment, we plan to investigate the role of oscillatory waves on conscious perception. Specifically, we plan to test participants in two conditions, using bistable perception. In the first experiment, we will test participants using a Binocular Rivalry (BR) design (see Figure 3.8B). Specifically, we will show each eye a differently colored and oriented Gabor patch. We will use a set of dichoptic mirrors to induce the rivalry between the two stimuli, as already done in our previous work (Luo et al., 2021). Participants will attend the patches for 30 seconds,

continuously reporting the one consciously perceived. We will record the EEG signals and the amount of waves along the midline of electrodes. In the Predictive Coding framework, our compelling hypothesis is that an increase in FW waves, associated with forward prediction-errors (specifically in the gamma band, see WP2), will predict the moment of the perceptual switch. According to this interpretation, the suppressed stimulus, not consciously perceived, will generate an increasing amount of prediction errors over time –as it is not ‘explained away’ by higher regions- thus inducing the switch in perception. We also expect that top-down waves will predict participants’ perceptual switches, suggesting that BW waves represent predictions about the environment and are related to conscious perception (in line with PC principles). In the second experiment, we will test a different set of stimuli, bistable images (e.g., the Necker cube). Similarly, we plan to relate the amount of waves preceding the switch in perception. Unlike in the BR condition, with bistable images, the non-perceived interpretation supposedly doesn’t generate a prediction error (as in the case of the suppressed stimulus in the BR condition). Accordingly, in this experiment, we expect to observe a modulation of top-down waves that predicts the changes in perception but no modulation in the number of forward waves. Both experiments will help us elucidate the relation between traveling waves and conscious perception, grounding the results in the theoretical framework of Predictive Coding.

Chapter IV: Conclusions and long term perspective

Predictive Processes and Oscillatory Traveling waves: a unifying framework for cortical functions?

In the last few decades, the field of Neuroscience has witnessed an extraordinary development in quantitative approaches and experimental methods, producing a large amount of data and recordings from different animal models and scales. In parallel with this growth, computational models (and, more broadly, Artificial Intelligence) also underwent an impressive revolution led by deep learning models, which have de facto worked out several problems that a few years ago seemed out of reach in different fields, such as computer vision or language generation. Arguably, the combination of these two impressive boosts, in both Neuroscience and Artificial Intelligence, set the stage **for investigating Cognition with more holistic and comprehensive theories** that could span across scales and degrees of abstraction. Ideally, such an approach aims at going beyond the mere description of a phenomenon, providing a mechanistic explanation able to bridge and unify disparate experimental frameworks (Levenstein et al., 2023). OSCI-PRED's ambition is to take one step in this direction, testing the compelling yet disputed hypothesis that Predictive Coding could provide the basis for such a general scheme.

In this manuscript, as in most of my recent publications on this topic, I used the term *Predictive Coding* loosely to refer to a more general scheme at times labeled *Predictive Processes* or, even more broadly, *Bayesian inference*. Predictive Coding, as commonly defined (Bastos et al., 2012; Huang & Rao, 2011; Rao & Ballard, 1999; Shipp, 2016), assumes a series of neural elements that are not strictly necessary to implement Bayesian inference (Heilbron & Chait, 2018), such as prediction-error neurons, or specific oscillatory laminar profiles. One of OSCI-PRED's goals is to test from a computational and experimental point of view which elements and ingredients are necessary to describe different oscillatory dynamics (e.g., traveling waves). On top of this, an exciting notion emerging when generalizing PC to the Bayesian inference framework is that the brain implements **probabilistic coding** (Findling et al., 2023; Walker et al., 2020; Zemel et al., 1998). Specifically, Bayesian models propose that sensory variables are encoded as distributions, representing their expected value (i.e., the mean) and their uncertainty (i.e., its variance). Let's now consider this from a Predictive Coding and traveling waves perspective. Prior distributions generate predictions in hierarchically higher brain regions, whereas prediction-errors update such priors based on the difference between

the prior and the likelihood distribution describing the sensory evidence (Figure 4.1A). From this, it follows that there should be different patterns of traveling waves depending on the precision (i.e., variance) of the distribution: more precise priors generate stronger predictions and, in turn, stronger backward waves, whereas less specific prior information generates inaccurate predictions, hence higher prediction-errors, reflected by stronger forward waves (Figure 4.1B summarizes this). This hypothesis provides an explicit link between traveling waves and PC processes, which is worth testing experimentally (for example, by directly manipulating the expectations/priors and observing the changes in backward and forward traveling waves – see working package two of OSCIPRED and Figure 3.6B). As described below, we have successfully tested such hypothesis indirectly in two experimental studies.

Traveling waves and Predictive Coding: from psychedelics to schizophrenia patients

In the first study, which has already been described in a previous chapter, we modulated the prior precision via pharmacological manipulation. As detailed above, a recently proposed model (Carhart-Harris & Friston, 2019) suggests that psychedelic drugs relax the precision of the high-level prior, thus eliciting a decrease in top-down waves and a consequent increase of forward ones. As shown in Figure 2.7D, we were able to confirm such predictions experimentally (Alamia, Timmermann, et al., 2020), observing a decrease (increase) in the alpha-band top-down (bottom-up) traveling waves after the intake of a serotonergic psychedelics drug (i.e., N,N-Dimethyltryptamine, DMT). However, one may wonder whether it would be possible to observe the symmetric effect, i.e., an increase in alpha-band top-down traveling waves due to more precise prior (Figure 4.1B). Some recent studies proposed that schizophrenia patients have a dysfunctional updating of their cognitive world model, usually described within the framework of Bayesian inference and predictive coding (Corlett et al., 2009; Krystal et al., 2017). Specifically, it was suggested that schizophrenia patients have stronger, i.e., more precise, priors (Corlett et al., 2019; Friston, 2005; Powers et al., 2017), contributing more to perception than the sensory evidence. In collaboration with Michael Herzog and his team in Lausanne, I had the opportunity to analyze EEG recordings in schizophrenia patients. In line with our hypothesis, our results reveal a substantial increase in top-down and a decrease in bottom-up alpha waves in schizophrenia patients (N=121) compared to healthy participants (N=96) during resting states (Figure 4.1C, D). Altogether, our results support the hypothesis that schizophrenia patients have more precise priors (i.e., with less variability) than the healthy participants (Corlett et al., 2019; Friston, 2005; Powers et al., 2017), which in turn causes larger predictions and a reduction in prediction-errors. Accordingly, in terms of traveling waves, we observed a higher amount of waves propagating in the top-down direction due to stronger priors and attenuation of bottom-up waves (Figure 4.1C, D).

Additionally, we found a positive correlation between backward waves in the beta and gamma bands and positive symptoms assessment (SAPS), which quantifies the symptoms related to hallucinations, delusions, and aberrancies in perception. Such correlation may suggest further evidence corroborating the relationship between top-down traveling waves and prior belief. Altogether, both the pharmacological manipulation with psychedelics and the EEG recordings in schizophrenia patients provide compelling but indirect evidence that oscillatory traveling waves may indeed reflect predictive processes.




Figure 4.1. – Bayesian interpretation of TW, and Schizophrenic patients results. **A.** In the Bayesian perspective, predictions are generated by prior distributions in higher brain regions, prediction-errors are computed to update the prior based on the sensory evidence (i.e., the likelihood). **B.** Considering backward (BW) and forward (FW) waves as proxies of predictions and prediction-errors, respectively (Alamia & VanRullen, 2019b), one would expect different patterns of traveling waves depending on the precision of the prior: more precise prior (rightmost panel) generate stronger predictions (stronger backward waves), whereas less precise prior (leftmost pattern) generates inaccurate predictions, hence higher prediction-errors, reflected by stronger forward waves. **C.** Row power for each spectral band in the midline electrodes (x-axis), for both the patients and control group in the two datasets. Each color represents a different frequency band. Error bars represent standard errors. **D.** The left panels illustrate

the spectra for both forward (blue) and backward (red) waves for the two groups in both datasets; the right panels show the mean for each frequency bands. Error bars represent standard errors.

The computational role of traveling waves

As discussed in the 'ERC Starting Grant 'OSCI-PRED' (2023 – 2028)' section, I plan to test the relationship between traveling waves and predictive coding directly by manipulating the prior variance experimentally and computationally. Both approaches will be crucial to test this hypothesis precisely. Regarding the experimental approach, modifying a stimulus's context and sensory expectations will allow us to directly manipulate its uncertainty and, supposedly, the waves modulating its representation. However, it's crucial to consider the idea of hierarchy inherent to predictive coding systems (Huang & Rao, 2011; Murray et al., 2014; Sterzer et al., 2018b), as different hierarchical levels represent separately low and high stimulus' features. Accordingly, modulating prior uncertainty at distinct hierarchical levels may generate distinct symptomatic behaviors. For example, visual illusions (Lhotka et al., 2023) or implicit perceptual learning (Valton et al., 2019) may rely on relatively lower-level priors, whereas temporal expectations may involve higher-order processing (Seymour et al., 2004; Visalli et al., 2019). Disentangling the different hierarchical priors will allow us to formulate precise experimental predictions and assess what information is carried by traveling waves. However, other experimental works remind us of the broader range of functions associated with traveling waves, which go beyond the predictive coding dynamic: previous studies highlighted the involvement of waves in gating sensory perception (Davis et al., 2020), as well as illusory motion (Chemla et al., 2019), and waves are most likely involved in synchronizing activity within regions to integrate sensory information, leveraging horizontal and lateral connectivity (Chavane et al., 2022; Liang et al., 2017; Sato et al., 2012). An exciting line of research will be to understand the links between these different functions and to understand to which extent traveling waves are a general mechanism involved in cortical computations.

On top of the experimental approach, it is possible to investigate the relationship between traveling waves and predictive processes from a computational perspective. Recent work is already investigating their functional role in different models. For example, Benigno et al., 2023 showed that traveling waves could be instrumental in predicting complex and naturalistic input (without implementing a predictive coding dynamic explicitly), thus playing an essential role in synchronizing spatiotemporal structure over cortical maps. Similarly, two other studies independently showed that traveling waves emerge spontaneously with topographically organized (i.e., local) connectivity and conduction delays in biologically plausible spiking networks (Davis et al., 2021), as well as in deep neural networks (T. A. Keller & Welling, 2023). Concerning my future work in this direction, it is possible to introduce oscillatory dynamics in deep neural networks by adding temporal delays on top of a predictive coding scheme, as in

our previous study (Alamia & VanRullen, 2019b). Considering such multi-layer networks, it's then possible to manipulate at each layer the strength (i.e., the precision, from a probabilistic point of view) of each term involved in its activity, including the prior (or feedback) activity (similarly to what was described above, and in Alamia, Mozafari, et al., 2023). Using this method, we could test how the oscillations propagate through the network as a function of the prior strength. Besides testing the relationship between traveling waves and predictive coding, this approach will provide one way to introduce oscillatory dynamics in neural networks. Other interesting paths have been proposed to introduce oscillations in neural networks, such as using Complex-valued activity (Löwe et al., 2022; Reichert & Serre, 2013; Trabelsi et al., 2017). This line of research, where I'm actively involved and that's currently explored in the team (in collaboration with Rufin VanRullen and Thomas Serre), is an exciting venue for ongoing and future research, as described in the next paragraph.

Brain-inspired models to improve AI state-of-the-art

All current machine learning approaches train the model's parameters via backpropagation, one of the most common techniques based on propagating the gradient of the error of a given loss function. However, the neuronal spiking activity –which produces the oscillatory dynamics in the models described above- is not differentiable and thus not amenable to being trained via backpropagation. Although some solutions have been recently proposed by computing surrogate gradients (Huh & Sejnowski, 2018; Neftci et al., 2019), another promising approach involves using Complex units (Hirose, 2012). Besides being differentiable and then suitable for backpropagation, one of the critical properties of complex activation functions is that their polar form is readily interpretable in terms of temporal oscillations. Previous work implementing artificial neural networks with complex units has shown promising results despite not being motivated by a biological plausible implementation (Hirose, 2012; Trabelsi et al., 2017), with the notable exception of (Löwe et al., 2022; Reichert & Serre, 2013) (notice that in Reichert & Serre, 2013 complex units were added to the network after the training, introducing an oscillatory dynamics that however was not learned). One research direction I'm currently pursuing in collaboration with Rufin VanRullen and Thomas Serre is to represent neurons' activity in deep neural networks in the Complex domain. One way of implementing this is to characterize the activity of each neuron with three parameters, representing frequency, phase, and module of the complex unit. This approach allows training the connections between neurons via a learning rule.

As mentioned in a previous section, recent studies showed that purely feedforward networks struggle to perform same-different tasks (Kim et al., 2018), in which the model has to report

whether two or more items are identical. We demonstrated that oscillatory dynamics are instrumental in the human brain to carry out the task successfully (Alamia, Luo, et al., 2020). This visual reasoning task is an excellent benchmark to test and demonstrate that oscillations can be a helpful tool to implement effective recurrent processes and propose a novel, biologically inspired architecture that can reach state-of-the-art results in visual reasoning tasks.

In collaboration with Rufin VanRullen, we recently developed an architecture that takes inspiration from the mathematical formalism of Complex-valued numbers, which -as mentioned above- are differentiable and can be interpreted as oscillations. Specifically, we introduce an additional 'phase' layer in the network, which is paired with each layer in the model (named 'PhaseNet' and shown in Figure 4.2A). The objective is to achieve a dynamic reminiscent of the "binding by synchrony" mechanism in the brain (Engel et al., 2001; Singer, 2007), in which neurons coding for the same object or feature become more synchronized over time, enhancing grouping and segmentation. Similarly, the 'phase' layer in the network aims to improve the neurons' activity in the convolutional and dense layers coding for the same object, thus favoring grouping and classification. More precisely, each phase value φ evolves over time steps according to the dynamics generated by the Kuramoto equation (Cumin & Unsworth, 2007; Kuramoto, 1975), which leads to the synchronization of phases having positive coupling r_{ij} over time steps (see Figure 4.2A, lower panels):

$$\varphi_i(t+1) = \varphi_i(t) + \lambda \frac{\sum_{j=1}^N r_{ij} \sin(\varphi_j(t) - \varphi_i(t))}{\sum_{j=1}^N |r_{ij}| + \varepsilon}$$

In the current implementation, the Kuramoto couplings r_{ij} are computed in the 'attentional space,' an n-dimensional space inspired by the recent 'transformer' architecture (Vaswani et al., 2017), in which features from every layer are embedded (via linear transformation) and compared at every time step:

$$r_{ij} = \langle W_q m_i, W_k m_j \rangle$$

Importantly, as shown in Figure 4.2A, the phases synchronize over time steps in each convolutional layer, highlighting different image features. We expect that such synchronization will improve the network's performance compared to a forward network with the same number of parameters, as the phases will modulate the activity of the network as described in the following equations (for convolutional and dense layers, respectively):

$$C_{L+1} = \int [1 + \alpha \cos(\varphi_{L+1}(x) - \varphi_L(x - \tau))] * C_L(x - \tau) * K(\tau) d\tau$$

$$d_i^{L+1} = \sum_i \sum_k d_k^L W_{k,i} (1 + \alpha \cos(\varphi_k^L - \varphi_k^{L+1}))$$

Additionally, we aim to leverage these dynamics as a top-down attentional mechanism in which the network can focus on specific items of the image to improve visual tasks' robustness, such

as image classification or visual reasoning tasks. Preliminary results show promising improvements, suggesting that phases could benefit current deep neural networks. In the long term, the goal is to equip deep neural networks with biological-plausible dynamics, combining PC architectures (previous work) and oscillatory dynamics, to address some of the main problems in computer vision. Indeed, most networks proved remarkably vulnerable to different issues, among which adversarial attacks in object recognition (Goodfellow et al., 2014), suggesting the need for more robust and efficient architectures, possibly inspired by brain-inspired mechanisms. Combining Predictive Coding and oscillatory dynamics will prove crucial in creating more robust than state-of-the-art models in computer vision

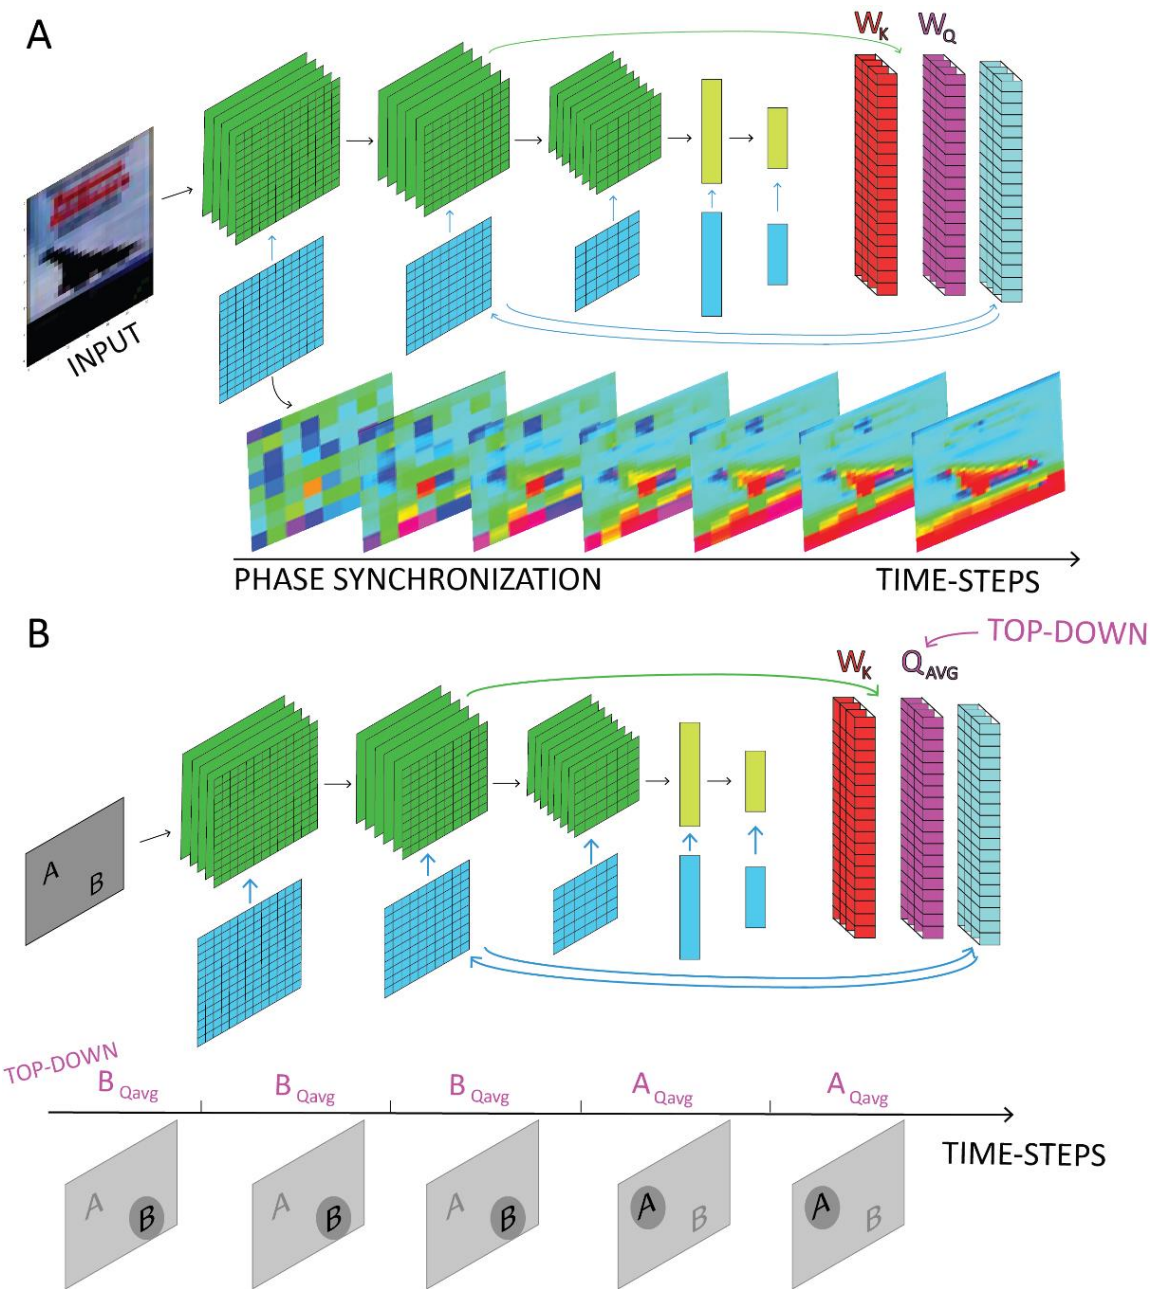


Figure 4.2. – PhaseNet architecture. A) The model is composed of three convolutional and two dense layers. Each layer's features are embedded in the 'Attention Space', similar to the attentional mechanisms used in current transformer architectures. The phases follow a dynamic inspired by the Kuramoto equation, which models the dynamic of a large set of oscillators (e.g. neurons), describing their synchronization over time given some assumptions (i.e. all oscillators are identical and weakly coupled, and their interaction depends on the sinusoid of their phase difference). Lower panels show the phases synchronized to different features in the image over time steps. B) Same architecture as above, but in this implementation, the couplings are determined by the dot product between the keys and a fixed queries, which represent a 'top-down' (or endogenous) attentional system. Given several stimuli in the figure (for example the letter 'A' and 'B'), the phases will synchronize over time to the item encoded in the queries. This mechanism will act as an endogenous attentional system.

Long-term perspective and conclusion

New directions: Computational Psychiatry

Understanding the relationship between traveling waves and predictive coding could provide a fundamental understanding of brain dynamics and cognitive processes. Besides, it could provide helpful insights into several clinical applications, such as in the emerging field of computational psychiatry. The compelling intuition behind this exciting new field, which greatly resonates with the scientific vision I attempted to describe above, is to provide a mechanistic understanding of psychiatric disorders beyond the 'mere' identification of diagnostic biomarkers (which remains a useful and precious diagnostic tool). In other words, computational psychiatry grants the tools to investigate cognition in specific pathological conditions to provide therapeutic support while revealing brain dynamics. For example, let's consider once again the case of schizophrenia patients. As mentioned earlier, schizophrenia is a severe mental disease that affects about one percent of the world's population (Krystal et al., 2017; McCutcheon et al., 2020; Series, 2020; Wang & Krystal, 2014). Schizophrenia patients reveal abnormal features in several behavioral traits, including perception, cognition, and personality. Such a broad range of symptoms hints that schizophrenia has a multifactorial cause with a complex presentation. In particular, patients may present positive or negative symptoms, which may be diagnostic of the disease. Positive symptoms comprise delusions and hallucinations, which can concur with disorganized speech and thoughts, whereas negative symptoms relate to catatonic behavior and diminished emotional expressions. Both positive and negative symptoms can co-exist and waver over time. Despite almost a century of research on its etiology, we still lack a clear understanding of its causes. Remarkably, a (computational) model of the disorder grounded in a theoretical framework such as the one provided by the Bayesian one could help elucidate its origins and causes. Furthermore, understanding such pathology within a specific framework will be also beneficial to define and understand brain dynamics in healthy conditions. Previous works have shown how top-down processes and prior expectations positively relate to positive symptoms such as hallucinations

(Corlett et al., 2019; Powers et al., 2017). In healthy subjects, experimental work demonstrated that expectations modulate perception (Kok et al., 2012; Summerfield & de Lange, 2014): the conditioned stimulus induces the perception of the unconditioned stimulus, even when this is not actually present (Ellson, 1941; Kafadar et al., 2022). Such an effect is more prominent in hallucinating schizophrenia patients (Kot & Serper, 2002). However, other studies provide evidence against the hypothesis that schizophrenia patients have stronger prior. In particular, experimental works demonstrated that patients are less sensitive to visual context (Choung et al., 2022) and visual illusions effects (Dakin et al., 2005), which are mediated by top-down priors (Geisler & Kersten, 2002). How to reconcile such supposedly contradictory findings? Some authors proposed a more nuanced framework to account for such contradictory results (Corlett et al., 2019; Sterzer et al., 2018b), leveraging the crucial idea of hierarchy inherent to predictive coding systems (as mentioned in a previous paragraph). Specifically, they proposed a differential modulation of priors at different hierarchical levels: visual illusion may rely on relatively lower-level priors, which affect visual perception specifically, whereas schizophrenia patients may have impairments in higher-level priors, involved in higher-order functions. For example, schizophrenia patients proved more sensitive than control to the triple flash illusion (Norton et al., 2008), which involves higher-order processing such as temporal expectations. Experimental work on various illusory measures explained by the predictive coding theory revealed a weak correlation (Lhotka et al., 2023), providing evidence against a common prior for perceptual and higher phenomena but supporting the notion of different hierarchical priors. According to these considerations, we could surmise that traveling waves measured via EEG recordings indeed reflect higher-order prior, and schizophrenia patients have indeed more precision, specifically in higher-order priors (Figure 4.1D). More generally, neural oscillations play a crucial role in coordinating neural activity, and their synchronization may be a core pathophysiological mechanism involved in schizophrenia (Uhlhaas & Singer, 2006, 2010). Differences between schizophrenia patients and healthy control have been observed in all frequency bands in different tasks or experimental conditions: theta-band related to working memory (Schmiedt et al., 2005), alpha and beta band regarding long-range connectivity (Liddle et al., 2016; Nikulin et al., 2012), higher frequency gamma-band (Hirano et al., 2015; Lee et al., 2003; Sun et al., 2011) and low-high frequency coupling (e.g., theta-gamma multiplexing, Barr et al., 2017; Kirihara et al., 2012). The differences between healthy populations and schizophrenia patients (and, more broadly, psychiatric patients) goes far beyond neural oscillations. As mentioned above, schizophrenia is a very heterogeneous condition, which manifests in several phenotypes. The appeal of computational psychiatry is to provide the mean to ground in a theoretical framework all these disparate evidence (Valton et al., 2017), to establish a comprehensive understanding of distinct psychiatric diseases and, ultimately, human cognition.

Conclusions - Beyond EEG recordings: intracortical recordings and animal models.

In this thesis, I provided an overview of my previous work on understanding cognitive functions in the human brain. Most of the work I presented here relies on computational models or human recordings to achieve this goal (from behavioral performance to pupillometry to EEG recordings). Since my postdoc years until now, I have mostly based my work on human EEG recordings. EEG has several interesting advantages, such as the efficiency and easiness of running and testing several participants in a relatively short amount of time, not to mention the very high temporal resolution, which is ideal for investigating brain oscillations. All in all, it's the ideal tool to investigate oscillatory dynamics in different experimental designs to test various hypotheses. However, it becomes necessary to leverage other tools to understand in depth the neural dynamics involved in cognitive functions and validate the predictions of the models I plan to implement and explore at different scales (from cortical columns to the cortical-thalamic connections). For example, I foresee in the long-term future collaborating more closely with other scientists involved in intra-cortical recordings of epileptic patients. In line with this direction, I'm currently setting up a collaboration at CerCo with Dr. Leila Reddy, who is willing to share her valuable expertise to investigate traveling waves in sEEG recordings. Furthermore, I'm also excited to be involved in a collaboration with Dr. Frederic Chavane from the Institute de Neurosciences de la Timone (INT, Marseille). Together, we plan to investigate traveling waves in voltage-dye recordings in Marmosets to assess the characteristics of spontaneous and evoked traveling waves in these recordings. All in all, besides analyzing 'better' data in terms of scale and explanatory power, I'm excited to be able to establish scientific collaborations with inspiring and talented people. During my career, I had the opportunity to be mentored by exceptional scientists, both from a scientific and a human point of view. In the future, I wish to be able to give back to the people who will cross my path the same excitement and inspiration I received during these years.

References

- Aertsen, A. M. H. J., Gerstein, G. L., Habib, M. K., & Palm, G. (1989). Dynamics of neuronal firing correlation : Modulation of « effective connectivity ». *Journal of Neurophysiology*, 61(5). <https://doi.org/10.1152/jn.1989.61.5.900>
- Alamia, A., Gauducheau, V., Paisios, D., & VanRullen, R. (2020). Comparing feedforward and recurrent neural network architectures with human behavior in artificial grammar learning. *Scientific Reports*, 10(1), 1-15. <https://doi.org/10.1038/s41598-020-79127-y>
- Alamia, A., Luo, C., Ricci, M., Kim, J., Serre, T., & VanRullen, R. (2020). Differential involvement of EEG oscillatory components in sameness vs. Spatial-relation visual reasoning tasks. *eneuro*, ENEURO.0267-20.2020. <https://doi.org/10.1523/eneuro.0267-20.2020>
- Alamia, A., Mozafari, M., Choksi, B., & VanRullen, R. (2023). On the role of feedback in image recognition under noise and adversarial attacks : A predictive coding perspective. *Neural Networks*, 157, 280-287. <https://doi.org/10.1016/j.neunet.2022.10.020>
- Alamia, A., Orban de Xivry, J.-J., San Anton, E., Olivier, E., Cleeremans, A., Zenon, A., de Xivry, J.-J. O., San Anton, E., Olivier, E., Cleeremans, A., & Zenon, A. (2016). Unconscious associative learning with conscious cues. *Neuroscience of Consciousness*, 2016(1), 1-10. <https://doi.org/10.1093/nc/niw016>
- Alamia, A., Solopchuk, O., D'Ausilio, A., Van Bever, V., Fadiga, L., Olivier, E., Zénon, A., D'Ausilio, A., Van Bever, V., Fadiga, L., Olivier, E., & Zénon, A. (2016). Disruption of Broca's Area Alters Higher-order Chunking Processing during Perceptual Sequence Learning. *Journal of Cognitive Neuroscience*, 28(3), 402-417. https://doi.org/10.1162/jocn_a_00911
- Alamia, A., Solopchuk, O., & Zénon, A. (2018). Strong Conscious Cues Suppress Preferential Gaze Allocation to Unconscious Cues. *Frontiers in Human Neuroscience*, 12(October), 1-9. <https://doi.org/10.3389/fnhum.2018.00427>
- Alamia, A., Terral, L., d'Ambra, M. R., & VanRullen, R. (2023). Distinct roles of forward and backward alpha-band waves in spatial visual attention. *Elife*, 12, e85035.
- Alamia, A., Terral, L., Renaud D'ambra, M., & Vanrullen, R. (2023). Distinct roles of forward and backward alpha-band waves in spatial visual attention. *eLife*, 12:e85035.
- Alamia, A., Timmermann, C., Nutt, D. J., VanRullen, R., & Carhart-Harris, R. L. (2020). DMT alters cortical travelling waves. *ELife*, 9, e59784. <https://doi.org/10.7554/eLife.59784>
- Alamia, A., & VanRullen, R. (2019a). Alpha oscillations and traveling waves : Signatures of predictive coding? *PLOS Biology*, 17(10), e3000487. <https://doi.org/10.1371/journal.pbio.3000487>
- Alamia, A., & VanRullen, R. (2019b). Alpha oscillations and traveling waves : Signatures of predictive coding? *PLOS Biology*, 17(10), e3000487. <https://doi.org/10.1371/journal.pbio.3000487>
- Alamia, A., & VanRullen, R. (2023). A Traveling Waves Perspective on Temporal Binding. *Journal of Cognitive Neuroscience*, 1-9. https://doi.org/10.1162/jocn_a_02004
- Alamia, A., VanRullen, R., Pasqualotto, E., Mouraux, A., & Zenon, A. (2019). Pupil-Linked Arousal Responds to Unconscious Surprisal. *The Journal of Neuroscience*, 39(27), 5369-5376. <https://doi.org/10.1523/JNEUROSCI.3010-18.2019>
- Alamia, A., & Zénon, A. (2016). Statistical Regularities Attract Attention when Task-Relevant. *Frontiers in Human Neuroscience*, 10(February), 1-10. <https://doi.org/10.3389/fnhum.2016.00042>
- Alamia, A., Zénon, A., VanRullen, R., Duque, J., & Derosiere, G. (2019). Implicit visual cues tune oscillatory motor activity during decision-making. *NeuroImage*, 186, 424-436. <https://doi.org/10.1016/j.neuroimage.2018.11.027>
- Angeletos Chrysaitis, N., & Seriès, P. (2023). 10 years of Bayesian theories of autism : A comprehensive review. *Neuroscience & Biobehavioral Reviews*, 145, 105022. <https://doi.org/10.1016/j.neubiorev.2022.105022>
- Arnal, L. H., Wyart, V., & Giraud, A. L. (2011). Transitions in neural oscillations reflect prediction errors generated in audiovisual speech. *Nature Neuroscience*, 14(6), 797-801. <https://doi.org/10.1038/nn.2810>

- Bargh, J. A., & Morsella, E. (2008). The Unconscious Mind. *Perspectives on Psychological Science*, 3(1), 73-79. <https://doi.org/10.1111/j.1745-6916.2008.00064.x>
- Barr, M. S., Rajji, T. K., Zomorodi, R., Radhu, N., George, T. P., Blumberger, D. M., & Daskalakis, Z. J. (2017). Impaired theta-gamma coupling during working memory performance in schizophrenia. *Schizophrenia Research*, 189, 104-110. <https://doi.org/10.1016/j.schres.2017.01.044>
- Bastos, A. M., Litvak, V., Moran, R., Bosman, C. A., Fries, P., & Friston, K. J. (2015). A DCM study of spectral asymmetries in feedforward and feedback connections between visual areas V1 and V4 in the monkey. *NeuroImage*, 108, 460-475. <https://doi.org/10.1016/j.neuroimage.2014.12.081>
- Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76(4), 695-711. <https://doi.org/10.1016/j.neuron.2012.10.038>
- Bastos, A. M., Vezioli, J., Bosman, C. A., Schoffelen, J. M., Oostenveld, R., Dowdall, J. R., DeWeerd, P., Kennedy, H., & Fries, P. (2015). Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*, 85(2), 390-401. <https://doi.org/10.1016/j.neuron.2014.12.018>
- Benigno, G. B., Budzinski, R. C., Davis, Z. W., Reynolds, J. H., & Muller, L. (2023). Waves traveling over a map of visual space can ignite short-term predictions of sensory input. *Nature Communications*, 14(1), Article 1. <https://doi.org/10.1038/s41467-023-39076-2>
- Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks [Physics and Society; Statistical Mechanics; Computers and Society; Data Structures and Algorithms]. *Journal of Statistical Mechanics: Theory and Experiment*, 10, P10008. <https://doi.org/10.1088/1742-5468/2008/10/P10008>
- Boly, M., Garrido, M. I., Gosseries, O., Bruno, M. A., Boveroux, P., Schnakers, C., Massimini, M., Litvak, V., Laureys, S., & Friston, K. (2011). Preserved feedforward but impaired top-down processes in the vegetative state. *Science*, 332(6031), 858-862. <https://doi.org/10.1126/science.1202043>
- Bree, S., Alamia, A., & Zoefel, B. (2022). Oscillation or not—Why we can and need to know (commentary on Doelling and Assaneo, 2021). *European Journal of Neuroscience*, 55(1), 201-204. <https://doi.org/10.1111/ejn.15542>
- Brittain, J. S., & Brown, P. (2014). Oscillations and the basal ganglia : Motor control and beyond. In *NeuroImage* (Vol. 85, p. 637-647). <https://doi.org/10.1016/j.neuroimage.2013.05.084>
- Buffalo, E. A., Fries, P., Landman, R., Buschman, T. J., & Desimone, R. (2011). Laminar differences in gamma and alpha coherence in the ventral stream. *Proceedings of the National Academy of Sciences*, 108(27), 11262-11267. <https://doi.org/10.1073/pnas.1011284108>
- Buzsáki, G., & Draguhn, A. (2004). Neuronal Oscillations in Cortical Networks. *Science*, 304(5679), 1926-1929. <https://doi.org/10.1126/science.1099745>
- Buzsáki, G., & Wang, X. J. (2012). Mechanisms of gamma oscillations. In *Annual Review of Neuroscience* (Vol. 35, p. 203-225). <https://doi.org/10.1146/annurev-neuro-062111-150444>
- Carhart-Harris, R. L., & Friston, K. J. (2019). REBUS and the Anarchic Brain : Toward a Unified Model of the Brain Action of Psychedelics. *Pharmacological Reviews*, 71(3), 316-344. <https://doi.org/10.1124/pr.118.017160>
- Channon, S., Shanks, D., Johnstone, T., Vakili, K., Chin, J., & Sinclair, E. (2002). Is implicit learning spared in amnesia ? Rule abstraction and item familiarity in artificial grammar learning. *Neuropsychologia*, 40(12), 2185-2197. [https://doi.org/10.1016/S0028-3932\(02\)00037-4](https://doi.org/10.1016/S0028-3932(02)00037-4)
- Chavane, F., Perrinet, L. U., & Rankin, J. (2022). Revisiting horizontal connectivity rules in V1 : From like-to-like towards like-to-all. *Brain Structure and Function*, 227(4), 1279-1295. <https://doi.org/10.1007/s00429-022-02455-4>
- Chemla, S., Reynaud, A., Volo, M. di, Zerlaut, Y., Perrinet, L., Destexhe, A., & Chavane, F. (2019). Suppressive Traveling Waves Shape Representations of Illusory Motion in Primary Visual

- Cortex of Awake Primate. *Journal of Neuroscience*, 39(22), 4282-4298.
<https://doi.org/10.1523/JNEUROSCI.2792-18.2019>
- Choksi, B., Mozafari, M., Biggs O' May, C., ADOR, B., Alamia, A., & VanRullen, R. (2021). Predify : Augmenting deep neural networks with brain-inspired predictive coding dynamics. *Advances in Neural Information Processing Systems*, 34, 14069-14083.
<https://proceedings.neurips.cc/paper/2021/hash/75c58d36157505a600e0695ed0b3a22d-Abstract.html>
- Choksi, B., Mozafari, M., Biggs O'May, C., Ador, B., Alamia, A., & VanRullen, R. (2020). Brain-inspired predictive coding dynamics improve the robustness of deep neural networks. *NeurIPS 2020 Workshop SVRHM*.
- Choung, O.-H., Gordillo, D., Roinishvili, M., Brand, A., Herzog, M. H., & Chkonia, E. (2022). Intact and deficient contextual processing in schizophrenia patients. *Schizophrenia Research: Cognition*, 30, 100265. <https://doi.org/10.1016/j.scog.2022.100265>
- Cleeremans, A., Destrebecqz, A., & Boyer, M. (1998). Implicit learning : News from the front. *Trends in cognitive sciences*, 2(10), 406-416.
- Colombo, M., & Seriès, P. (2012). Bayes in the Brain—On Bayesian Modelling in Neuroscience. *The British Journal for the Philosophy of Science*, 63(3), 697-723.
<https://doi.org/10.1093/bjps/axr043>
- Corballis, M. C. (2014). Recursive cognition as a prelude to language. In *Language and Recursion* (Vol. 9781461494, p. 27-36). https://doi.org/10.1007/978-1-4614-9414-0_3
- Corlett, P. R., Frith, C. D., & Fletcher, P. C. (2009). From drugs to deprivation : A Bayesian framework for understanding models of psychosis. *Psychopharmacology*, 206(4), 515-530.
<https://doi.org/10.1007/s00213-009-1561-0>
- Corlett, P. R., Horga, G., Fletcher, P. C., Alderson-Day, B., Schmack, K., & Powers, A. R. (2019). Hallucinations and Strong Priors. *Trends in Cognitive Sciences*, 23(2), 114-127.
<https://doi.org/10.1016/j.tics.2018.12.001>
- Cumin, D., & Unsworth, C. P. (2007). Generalising the Kuramoto model for the study of neuronal synchronisation in the brain. *Physica D: Nonlinear Phenomena*, 226(2), 181-196.
<https://doi.org/10.1016/j.physd.2006.12.004>
- Dakin, S., Carlin, P., & Hemsley, D. (2005). Weak suppression of visual context in chronic schizophrenia. *Current Biology*, 15(20), R822-R824. <https://doi.org/10.1016/j.cub.2005.10.015>
- Davis, Z. W., Benigno, G. B., Fletteman, C., Desbordes, T., Steward, C., Sejnowski, T. J., H. Reynolds, J., & Muller, L. (2021). Spontaneous traveling waves naturally emerge from horizontal fiber time delays and travel through locally asynchronous-irregular states. *Nature Communications*, 12(1), Article 1. <https://doi.org/10.1038/s41467-021-26175-1>
- Davis, Z. W., Muller, L., Martinez-Trujillo, J., Sejnowski, T., & Reynolds, J. H. (2020). Spontaneous travelling cortical waves gate perception in behaving primates. *Nature*, 587(7834), 432-436.
<https://doi.org/10.1038/s41586-020-2802-y>
- de Haan, E. H., Young, A., & Newcombe, F. (1987). Face recognition without awareness. *Cogn Neuropsychol*, 4(4), 385-415. <https://doi.org/10.1080/02643298708252045>
- De Lange, F. P., Heilbron, M., & Kok, P. (2018). How Do Expectations Shape Perception ? Perceptual Consequences of Expectation. *Trends in Cognitive Sciences*, xx, 1-16.
<https://doi.org/10.1016/j.tics.2018.06.002>
- Derosiere, G., Zénon, A., Alamia, A., & Duque, J. (2017). Primary motor cortex contributes to the implementation of implicit value-based rules during motor decisions. *NeuroImage*, 146, 1115-1127. <https://doi.org/10.1016/j.neuroimage.2016.10.010>
- Derosiere, G., Zenon, A., Alamia, A., Klein, P.-A., & Duque, J. (2015). Contribution of primary motor cortex to perceptual and value-based decision processes. *Frontiers in Neuroscience*, 9.
<https://doi.org/10.3389/conf.fnins.2015.89.00056>
- Dickinson, A. (1985). Actions and Habits : The Development of Behavioural Autonomy. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 308(1135), 67-78.
<https://doi.org/10.1098/rstb.1985.0010>

- Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., & Reitboeck, H. J. (1988). Coherent oscillations : A mechanism of feature linking in the visual cortex? *Biological Cybernetics*, 60(2), 121-130. <https://doi.org/10.1007/bf00202899>
- Ellson, D. G. (1941). Hallucinations produced by sensory conditioning. *Journal of Experimental Psychology*, 28(1), 1-20. <https://doi.org/10.1037/h0054167>
- Engel, A. K., Fries, P., & Singer, W. (2001). Dynamic predictions : Oscillations and synchrony in top-down processing. *Nature Reviews Neuroscience*, 2(10), 704-716. <https://doi.org/10.1038/35094565>
- Engel, A. K., & Singer, W. (2001). Temporal binding and the neural correlates of sensory awareness. In *Trends in Cognitive Sciences* (Vol. 5, Numéro 1, p. 16-25). [https://doi.org/10.1016/S1364-6613\(00\)01568-0](https://doi.org/10.1016/S1364-6613(00)01568-0)
- Feldman, H., & Friston, K. J. (2010). Attention, uncertainty, and free-energy. *Frontiers in human neuroscience*, 4(December), 215. <https://doi.org/10.3389/fnhum.2010.00215>
- Feldmann-Wüstefeld, T., & Vogel, E. K. (2019). Neural Evidence for the Contribution of Active Suppression During Working Memory Filtering. *Cerebral Cortex*. <https://doi.org/10.1093/cercor/bhx336>
- Fell, J., & Axmacher, N. (2011). The role of phase synchronization in memory processes. In *Nature Reviews Neuroscience* (Vol. 12, Numéro 2, p. 105-118). <https://doi.org/10.1038/nrn2979>
- Filbrich, L., Alamia, A., Blandiaux, S., Burns, S., & Legrain, V. (2017). Shaping visual space perception through bodily sensations : Testing the impact of nociceptive stimuli on visual perception in peripersonal space with temporal order judgments. *PLOS ONE*, 12(8), e0182634. <https://doi.org/10.1371/journal.pone.0182634>
- Filbrich, L., Alamia, A., Verfaillie, C., Berquin, A., Barbier, O., Libouton, X., Fraselle, V., Mouraux, D., & Legrain, V. (2017). Biased visuospatial perception in complex regional pain syndrome. *Scientific Reports*, 7(1), Article 1. <https://doi.org/10.1038/s41598-017-10077-8>
- Filbrich, L., Halicka, M., Alamia, A., & Legrain, V. (2018). Investigating the spatial characteristics of the crossmodal interaction between nociception and vision using gaze direction. *Consciousness and Cognition*, 57, 106-115. <https://doi.org/10.1016/j.concog.2017.11.011>
- Findling, C., Hubert, F., Laboratory, I. B., Acerbi, L., Benson, B., Benson, J., Birman, D., Bonacchi, N., Carandini, M., Catarino, J. A., Chapuis, G. A., Churchland, A. K., Dan, Y., DeWitt, E. E., Engel, T. A., Fabbri, M., Faulkner, M., Fiete, I. R., Freitas-Silva, L., ... Pouget, A. (2023). *Brain-wide representations of prior information in mouse decision-making* (p. 2023.07.04.547684). bioRxiv. <https://doi.org/10.1101/2023.07.04.547684>
- Fleuret, F., Li, T., Dubout, C., Wampler, E. K., Yantis, S., & Geman, D. (2011). Comparing machines and humans on a visual categorization test. *Proceedings of the National Academy of Sciences of the United States of America*, 108(43), 17621-17625. <https://doi.org/10.1073/pnas.1109168108>
- Foster, J. J., Sutterer, D. W., Serences, J. T., Vogel, E. K., & Awh, E. (2017). Alpha-Band Oscillations Enable Spatially and Temporally Resolved Tracking of Covert Spatial Attention. *Psychological Science*, 28(7), 929-941. <https://doi.org/10.1177/0956797617699167>
- Fries, P. (2005). A mechanism for cognitive dynamics : Neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, 9(10), 474-480. <https://doi.org/10.1016/j.tics.2005.08.011>
- Fries, P. (2015). Rhythms for Cognition : Communication through Coherence. In *Neuron* (Vol. 88, Numéro 1, p. 220-235). <https://doi.org/10.1016/j.neuron.2015.09.034>
- Friston, K. (2008). Hierarchical models in the brain. *PLoS Computational Biology*, 4(11). <https://doi.org/10.1371/journal.pcbi.1000211>
- Friston, K. J. (2005). Hallucinations and perceptual inference. *Behavioral and Brain Sciences*, 28(6), 764-766. <https://doi.org/10.1017/S0140525X05290131>
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, 6(4), 187-224. <https://doi.org/10.1080/17588928.2015.1020053>

- Fujioka, T., Trainor, L. J., Large, E. W., & Ross, B. (2009). Beta and gamma rhythms in human auditory cortex during musical beat processing. *Annals of the New York Academy of Sciences*, 1169, 89-92. <https://doi.org/10.1111/j.1749-6632.2009.04779.x>
- Gauthier, B., Eger, E., Hesselmann, G., Giraud, A.-L., & Kleinschmidt, A. (2012). Temporal Tuning Properties along the Human Ventral Visual Stream. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 32(41), 14433-14441. <https://doi.org/10.1523/JNEUROSCI.2467-12.2012>
- Geisler, W. S., & Kersten, D. (2002). Illusions, perception and Bayes. *Nature Neuroscience*, 5(6), Article 6. <https://doi.org/10.1038/nn0602-508>
- George, D., & Hawkins, J. (2009). Towards a mathematical theory of cortical micro-circuits. *PLoS Computational Biology*, 5(10). <https://doi.org/10.1371/journal.pcbi.1000532>
- Giraud, A.-L., & Arnal, L. H. (2012). Cortical oscillations and sensory predictions. In *Trends in Cognitive Sciences* (Vol. 16, Numéro 7, p. 390-398). <https://doi.org/10.1016/j.tics.2012.05.003>
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 3(January), 2672-2680. https://doi.org/10.3156/jsoft.29.5_177_2
- Greenwald, A. G., Klinger, M. R., & Schuh, E. S. (1995). Activation by Marginally Perceptible (« Subliminal ») Stimuli : Dissociation of Unconscious From Conscious Cognition. *Journal of Experimental Psychology: General*, 124(1), 22-42. <https://doi.org/10.1037/0096-3445.124.1.22>
- Gulbinaite, R., van Rijn, H., & Cohen, M. X. (2014). Fronto-parietal network oscillations reveal relationship between working memory capacity and cognitive control. *Frontiers in Human Neuroscience*, 8. <https://doi.org/10.3389/fnhum.2014.00761>
- Haegens, S., Handel, B. F., & Jensen, O. (2011). Top-Down Controlled Alpha Band Activity in Somatosensory Areas Determines Behavioral Performance in a Discrimination Task. *Journal of Neuroscience*, 31(14), 5197-5204. <https://doi.org/10.1523/jneurosci.5199-10.2011>
- Haeusler, S., & Maass, W. (2007). A statistical analysis of information-processing properties of lamina-specific cortical microcircuit models. *Cerebral Cortex*, 17(1), 149-162. <https://doi.org/10.1093/cercor/bhj132>
- Halgren, M., Ulbert, I., Bastuji, H., Fabó, D., Eross, L., Rey, M., Devinsky, O., Doyle, W. K., Mak-McCully, R., Halgren, E., Wittner, L., Chauvel, P., Heit, G., Eskandar, E., Mandell, A., & Cash, S. S. (2019). The generation and propagation of the human alpha rhythm. *Proceedings of the National Academy of Sciences of the United States of America*, 116(47), 23772-23782. <https://doi.org/10.1073/pnas.1913092116>
- Heeger, D. J. (2017). Theory of cortical function. *Proceedings of the National Academy of Sciences of the United States of America*, 114(8), 1773-1782. <https://doi.org/10.1073/pnas.1619788114>
- Heilbron, M., & Chait, M. (2018). Great Expectations : Is there Evidence for Predictive Coding in Auditory Cortex? *Neuroscience*, 389, 54-73. <https://doi.org/10.1016/j.neuroscience.2017.07.061>
- Hirano, Y., Oribe, N., Kanba, S., Onitsuka, T., Nestor, P. G., & Spencer, K. M. (2015). Spontaneous Gamma Activity in Schizophrenia. *JAMA Psychiatry*, 72(8), 813-821. <https://doi.org/10.1001/jamapsychiatry.2014.2642>
- Hirose, A. (2012). Complex-valued neural networks. *Studies in Computational Intelligence*, 400, 1-214. <https://doi.org/10.1007/978-3-642-27632-3>
- Huang, Y., & Rao, R. P. N. (2011). Predictive coding. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(5), 580-593. <https://doi.org/10.1002/wcs.142>
- Huerta, P. T., & Lisman, J. E. (1995). Bidirectional synaptic plasticity induced by a single burst during cholinergic theta oscillation in CA1 in vitro. *Neuron*, 15(5), 1053-1063. [https://doi.org/10.1016/0896-6273\(95\)90094-2](https://doi.org/10.1016/0896-6273(95)90094-2)
- Huh, D., & Sejnowski, T. J. (2018). Gradient descent for spiking neural networks. *Advances in Neural Information Processing Systems, 2018-Decem*, 1433-1443.

- Iversen, J. R., Repp, B. H., & Patel, A. D. (2009). Top-down control of rhythm perception modulates early auditory responses. *Annals of the New York Academy of Sciences*, 1169, 58-73. <https://doi.org/10.1111/j.1749-6632.2009.04579.x>
- Izhikevich, E. M. (2007). *Dynamical Systems in Neuroscience*. MIT Press.
- Jaramillo, J., Mejias, J. F., & Wang, X. J. (2019). Engagement of Pulvino-cortical Feedforward and Feedback Pathways in Cognitive Computations. *Neuron*, 101(2), 321-336.e9. <https://doi.org/10.1016/j.neuron.2018.11.023>
- Jia, X., Xing, D., & Kohn, A. (2013). No consistent relationship between gamma power and peak frequency in macaque primary visual cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 33(1), 17-25. <https://doi.org/10.1523/JNEUROSCI.1687-12.2013>
- Kafadar, E., Fisher, V. L., Quagan, B., Hammer, A., Jaeger, H., Mourgues, C., Thomas, R., Chen, L., Imtiaz, A., Sibarium, E., Negreira, A. M., Sarisik, E., Polisetty, V., Benrimoh, D., Sheldon, A. D., Lim, C., Mathys, C., & Powers, A. R. (2022). Conditioned Hallucinations and Prior Overweighting Are State-Sensitive Markers of Hallucination Susceptibility. *Biological Psychiatry*, 92(10), 772-780. <https://doi.org/10.1016/j.biopsych.2022.05.007>
- Keller, G. B., & Mrosovsky, T. D. (2018). Predictive Processing : A Canonical Cortical Computation. In *Neuron* (Vol. 100, Numéro 2, p. 424-435). <https://doi.org/10.1016/j.neuron.2018.10.003>
- Keller, T. A., & Welling, M. (2023). *Neural Wave Machines : Learning Spatiotemporally Structured Representations with Locally Coupled Oscillatory Recurrent Neural Networks*. <https://openreview.net/forum?id=5tJSt3kn4s>
- Ketz, N. A., Jensen, O., & O'Reilly, R. C. (2015). Thalamic pathways underlying prefrontal cortex-medial temporal lobe oscillatory interactions. In *Trends in Neurosciences* (Vol. 38, Numéro 1, p. 3-12). <https://doi.org/10.1016/j.tins.2014.09.007>
- Kiebel, S. J., Daunizeau, J., & Friston, K. J. (2008). A hierarchy of time-scales and the brain. *PLoS Computational Biology*, 4(11). <https://doi.org/10.1371/journal.pcbi.1000209>
- Kim, J., Ricci, M., & Serre, T. (2018). Not-So-CLEVR: Learning same-different relations strains feedforward neural networks. *Interface Focus*, 8(4). <https://doi.org/10.1098/rsfs.2018.0011>
- Kirihara, K., Rissling, A. J., Swerdlow, N. R., Braff, D. L., & Light, G. A. (2012). Hierarchical Organization of Gamma and Theta Oscillatory Dynamics in Schizophrenia. *Biological Psychiatry*, 71(10), 873-880. <https://doi.org/10.1016/j.biopsych.2012.01.016>
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain : The role of uncertainty in neural coding and computation. *Trends in Neurosciences*. <https://doi.org/10.1016/j.tins.2004.10.007>
- Kok, P., Brouwer, G. J., van Gerven, M. A. J., & de Lange, F. P. (2013). Prior expectations bias sensory representations in visual cortex. *Journal of Neuroscience*, 33(41), 16275-16284. <https://doi.org/10.1523/JNEUROSCI.0742-13.2013>
- Kok, P., Jehee, J. F. M., & de Lange, F. P. (2012). Less Is More : Expectation Sharpens Representations in the Primary Visual Cortex. *Neuron*, 75(2), 265-270. <https://doi.org/10.1016/j.neuron.2012.04.034>
- Komura, Y., Nikkuni, A., Hirashima, N., Uetake, T., & Miyamoto, A. (2013). Responses of pulvinar neurons reflect a subject's confidence in visual categorization. *Nature Neuroscience*, 16(6), 749-755. <https://doi.org/10.1038/nn.3393>
- Koshy, S. M., Wiesman, A. I., Proskovec, A. L., Embury, C. M., Schantell, M. D., Eastman, J. A., Heinrichs-Graham, E., & Wilson, T. W. (2020). Numerical working memory alters alpha-beta oscillations and connectivity in the parietal cortices. *Human Brain Mapping*, 41(13), 3709-3719. <https://doi.org/10.1002/hbm.25043>
- Kot, T., & Serper, M. (2002). Increased susceptibility to auditory conditioning in hallucinating schizophrenic patients : A preliminary investigation. *The Journal of Nervous and Mental Disease*, 190(5), 282-288. <https://doi.org/10.1097/00005053-200205000-00002>
- Kruse, W., & Eckhorn, R. (1996). Inhibition of sustained gamma oscillations (35-80 Hz) by fast transient responses in cat visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 93(12), 6112-6117.

- Krystal, J. H., Murray, J. D., Chekroud, A. M., Corlett, P. R., Yang, G., Wang, X.-J., & Anticevic, A. (2017). Computational Psychiatry and the Challenge of Schizophrenia. *Schizophrenia Bulletin*, 43(3), 473-475. <https://doi.org/10.1093/schbul/sbx025>
- Kuramoto, Y. (1975). Lecture Notes in Physics. *International Symposium on Mathematical Problems in Theoretical Physics*, 39, 420.
- Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. In *Trends in Neurosciences* (Vol. 23, Numéro 11, p. 571-579). [https://doi.org/10.1016/S0166-2236\(00\)01657-X](https://doi.org/10.1016/S0166-2236(00)01657-X)
- Lee, K.-H., Williams, L. M., Breakspear, M., & Gordon, E. (2003). Synchronous Gamma activity : A review and contribution to an integrative neuroscience model of schizophrenia. *Brain Research Reviews*, 41(1), 57-78. [https://doi.org/10.1016/S0165-0173\(02\)00220-5](https://doi.org/10.1016/S0165-0173(02)00220-5)
- Levenstein, D., Alvarez, V. A., Amarasingham, A., Azab, H., Chen, Z. S., Gerkin, R. C., Hasenstaub, A., Iyer, R., Jolivet, R. B., Marzen, S., Monaco, J. D., Prinz, A. A., Quraishi, S., Santamaria, F., Shivkumar, S., Singh, M. F., Traub, R., Nadim, F., Rotstein, H. G., & Redish, A. D. (2023). On the Role of Theory and Modeling in Neuroscience. *The Journal of Neuroscience*, 43(7), 1074-1088. <https://doi.org/10.1523/JNEUROSCI.1179-22.2022>
- Lhotka, M., Ischebeck, A., Helmlinger, B., & Zaretskaya, N. (2023). No common factor for illusory percepts, but a link between pareidolia and delusion tendency : A test of predictive coding theory. *Frontiers in Psychology*, 13, 1067985. <https://doi.org/10.3389/fpsyg.2022.1067985>
- Liang, H., Gong, X., Chen, M., Yan, Y., Li, W., & Gilbert, C. D. (2017). Interactions between feedback and lateral connections in the primary visual cortex. *Proceedings of the National Academy of Sciences*, 114(32), 8637-8642. <https://doi.org/10.1073/pnas.1706183114>
- Liddle, E. B., Price, D., Palaniyappan, L., Brookes, M. J., Robson, S. E., Hall, E. L., Morris, P. G., & Liddle, P. F. (2016). Abnormal salience signaling in schizophrenia : The role of integrative beta oscillations. *Human Brain Mapping*, 37(4), 1361-1374. <https://doi.org/10.1002/hbm.23107>
- Lisman, J. E., & Jensen, O. (2013). The Theta-Gamma Neural Code. In *Neuron* (Vol. 77, Numéro 6, p. 1002-1016). <https://doi.org/10.1016/j.neuron.2013.03.007>
- Lopes da Silva, F. H., Vos, J. E., Mooibroek, J., & van Rotterdam, A. (1980). Relative contributions of intracortical and thalamo-cortical processes in the generation of alpha rhythms, revealed by partial coherence analysis. *Electroencephalography and Clinical Neurophysiology*, 50(5-6), 449-456. [https://doi.org/10.1016/0013-4694\(80\)90011-5](https://doi.org/10.1016/0013-4694(80)90011-5)
- Lotter, W., Kreiman, G., & Cox, D. (2017). Deep predictive coding networks for video prediction and unsupervised learning. *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*.
- Löwe, S., Lippe, P., Rudolph, M., & Welling, M. (2022). *Complex-Valued Autoencoders for Object Discovery* (arXiv:2204.02075). arXiv. <https://doi.org/10.48550/arXiv.2204.02075>
- Lundqvist, M., Rose, J., Herman, P., Brincat, S. L. L., Buschman, T. J. J., & Miller, E. K. K. (2016). Gamma and Beta Bursts Underlie Working Memory. *Neuron*, 90(1), 152-164. <https://doi.org/10.1016/j.neuron.2016.02.028>
- Luo, C., VanRullen, R., & Alamia, A. (2021). Conscious perception and perceptual echoes : A binocular rivalry study. *Neuroscience of Consciousness*, 2021(1). <https://doi.org/10.1093/nc/niab007>
- McCutcheon, R. A., Reis Marques, T., & Howes, O. D. (2020). Schizophrenia—An Overview. *JAMA Psychiatry*, 77(2), 201-210. <https://doi.org/10.1001/jamapsychiatry.2019.3360>
- McKeeff, T. J., Remus, D. a, & Tong, F. (2007). Temporal limitations in object processing across the human ventral visual pathway. *Journal of Neurophysiology*, 98(1), 382-393. <https://doi.org/10.1152/jn.00568.2006>
- McLelland, D., & VanRullen, R. (2016). Theta-Gamma Coding Meets Communication-through-Coherence : Neuronal Oscillatory Multiplexing Theories Reconciled. *PLoS Computational Biology*, 12(10). <https://doi.org/10.1371/journal.pcbi.1005162>

- Michalareas, G., Vezoli, J., van Pelt, S., Schoffelen, J. M., Kennedy, H., & Fries, P. (2016). Alpha-Beta and Gamma Rhythms Subserve Feedback and Feedforward Influences among Human Visual Cortical Areas. *Neuron*, *89*(2), 384-397. <https://doi.org/10.1016/j.neuron.2015.12.018>
- Mucha, P. J., Richardson, T., Macon, K., Porter, M. a, & Onnela, J.-P. (2010). Community structure in time-dependent, multiscale, and multiplex networks. *Science (New York, N.Y.)*, *328*(5980), 876-878. <https://doi.org/10.1126/science.1184819>
- Muller, L., Chavane, F., Reynolds, J., & Sejnowski, T. J. (2018). Cortical travelling waves : Mechanisms and computational principles. In *Nature Reviews Neuroscience* (Vol. 19, Numéro 5, p. 255-268). <https://doi.org/10.1038/nrn.2018.20>
- Muller, L., Reynaud, A., Chavane, F., & Destexhe, A. (2014). The stimulus-evoked population response in visual cortex of awake monkey is a propagating wave. *Nature Communications*, *5*. <https://doi.org/10.1038/ncomms4675>
- Mumford, D. (1991). On the computational architecture of the neocortex—I. The role of the thalamo-cortical loop. *Biological Cybernetics*, *65*(2), 135-145. <https://doi.org/10.1007/BF00202389>
- Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., & Wang, X. J. (2014). A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience*, *17*(12), 1661-1663. <https://doi.org/10.1038/nn.3862>
- Neftci, E. O., Mostafa, H., & Zenke, F. (2019). Surrogate Gradient Learning in Spiking Neural Networks : Bringing the Power of Gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, *36*(6), 51-63. <https://doi.org/10.1109/MSP.2019.2931595>
- Newell, B. R., & Shanks, D. R. (2014). Unconscious influences on decision making : A critical review. *Behavioral and Brain sciences*, *37*, 1-19. <https://doi.org/10.1017/S0140525X12003214>
- Nikulin, V. V., Jönsson, E. G., & Brismar, T. (2012). Attenuation of long-range temporal correlations in the amplitude dynamics of alpha and beta neuronal oscillations in patients with schizophrenia. *NeuroImage*, *61*(1), 162-169. <https://doi.org/10.1016/j.neuroimage.2012.03.008>
- Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning : Evidence from performance measures. *Cognitive Psychology*, *19*(1), 1-32. [https://doi.org/10.1016/0010-0285\(87\)90002-8](https://doi.org/10.1016/0010-0285(87)90002-8)
- Norton, D., Ongur, D., Stromeyer, C., & Chen, Y. (2008). Altered 'three-flash' illusion in response to two light pulses in schizophrenia. *Schizophrenia Research*, *103*(1), 275-282. <https://doi.org/10.1016/j.schres.2008.03.002>
- O'reilly, R. C., Russin, J. L., Zolfaghar, M., & Rohrllich, J. (2021). Deep predictive learning in neocortex and pulvinar. *Journal of Cognitive Neuroscience*, *33*(6), 1158-1196. https://doi.org/10.1162/jocn_a_01708
- Ouellette, J. A., & Wood, W. (1998). Habit and Intention in Everyday Life : The Multiple Processes by Which Past Behavior Predicts Future Behavior. *Psychological Bulletin*, *124*(1), 124-154. <https://doi.org/10.1037/0033-2909.124.1.54>
- Pang, Z., Alamia, A., & Vanrullen, R. (2020). Turning the stimulus on and off changes the direction of α traveling waves. *eNeuro*, *7*(6), 1-11. <https://doi.org/10.1523/ENEURO.0218-20.2020>
- Powers, A. R., Mathys, C., & Corlett, P. R. (2017). Pavlovian conditioning–induced hallucinations result from overweighting of perceptual priors. *Science*, *357*(6351), 596-600. <https://doi.org/10.1126/science.aan3458>
- Quax, S., Jensen, O., & Tiesinga, P. (2017). Top-down control of cortical gamma-band communication via pulvinar induced phase shifts in the alpha rhythm. *PLoS Computational Biology*, *13*(5). <https://doi.org/10.1371/journal.pcbi.1005519>
- Rao, R., & Ballard, D. (1999). Predictive coding in the visual cortex : A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79-87. <https://doi.org/10.1038/4580>
- Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, *6*(6), 855-863. [https://doi.org/10.1016/S0022-5371\(67\)80149-X](https://doi.org/10.1016/S0022-5371(67)80149-X)
- Reichert, D. P., & Serre, T. (2013). *Neuronal Synchrony in Complex-Valued Deep Networks*.

- Rohenkohl, G., Bosman, C. A., & Fries, P. (2018). Gamma Synchronization between V1 and V4 Improves Behavioral Performance. *Neuron*, *100*(4), 953-963.e3. <https://doi.org/10.1016/j.neuron.2018.09.019>
- Rotshtein, P., Soto, D., Grecucci, A., Geng, J. J., & Humphreys, G. W. (2011). The role of the pulvinar in resolving competition between memory and visual selection : A functional connectivity study. *Neuropsychologia*, *49*(6), 1544-1552. <https://doi.org/10.1016/j.neuropsychologia.2010.12.002>
- Saalmann, Y. B., Pinsk, M. A., Wang, L., Li, X., & Kastner, S. (2012). The pulvinar regulates information transmission between cortical areas based on attention demands. *Science*, *337*(6095), 753-756. <https://doi.org/10.1126/science.1223082>
- Sabour, S., Frosst, N., & Hinton, G. E. (2017). Dynamic routing between capsules. *Advances in Neural Information Processing Systems, 2017-Decem*, 3857-3867.
- Samaha, J., Bauer, P., Cimaroli, S., & Postle, B. R. (2015). Top-down control of the phase of alpha-band oscillations as a mechanism for temporal prediction. *Proceedings of the National Academy of Sciences*, *112*(27), 8439-8444. <https://doi.org/10.1073/pnas.1503686112>
- Sato, T. K., Nauhaus, I., & Carandini, M. (2012). Traveling Waves in Visual Cortex. *Neuron*, *75*(2), 218-229. <https://doi.org/10.1016/j.neuron.2012.06.029>
- Schmiedt, C., Brand, A., Hildebrandt, H., & Basar-Eroglu, C. (2005). Event-related theta oscillations during working memory tasks in patients with schizophrenia and healthy controls. *Brain Research. Cognitive Brain Research*, *25*(3), 936-947. <https://doi.org/10.1016/j.cogbrainres.2005.09.015>
- Schneider, M., Broggin, A. C., Dann, B., Tzanou, A., Uran, C., Sheshadri, S., Scherberger, H., & Vinck, M. (2021). A mechanism for inter-areal coherence through communication based on connectivity and oscillatory power. *Neuron*, *109*(24), 4050-4067.e12. <https://doi.org/10.1016/j.neuron.2021.09.037>
- Series, P. (2020). *Computational Psychiatry : A Primer*. MIT Press.
- Seymour, B., O'Doherty, J. P., Dayan, P., Koltzenburg, M., Jones, A. K., Dolan, R. J., Friston, K. J., & Frackowiak, R. S. (2004). Temporal difference models describe higher-order learning in humans. *Nature*, *429*(6992), Article 6992. <https://doi.org/10.1038/nature02581>
- Shallice, T., & Saffran, E. (1986). Lexical processing in the absence of explicit word identification : Evidence from a letter-by-letter reader. *Cognitive Neuropsychology*, *3*(4), 429-458. <https://doi.org/10.1080/02643298608252030>
- Shanks, D., & Stjohn, M. (1994). *Characteristics of Dissociable Human Learning-Systems. February 2010*. <https://doi.org/10.1017/S0140525X00035032>
- Shipp, S. (2016). Neural Elements for Predictive Coding. *Frontiers in Psychology*, *7*. <https://www.frontiersin.org/articles/10.3389/fpsyg.2016.01792>
- Singer, W. (2007). Binding by synchrony. *Scholarpedia*, *2*(12), 1657. <https://doi.org/10.4249/scholarpedia.1657>
- Singer, W. (2013). Cortical dynamics revisited. In *Trends in Cognitive Sciences* (Vol. 17, Numéro 12, p. 616-626). <https://doi.org/10.1016/j.tics.2013.09.006>
- Sohal, V. S., Zhang, F., Yizhar, O., & Deisseroth, K. (2009). Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. *Nature*, *459*(7247), 698-702. <https://doi.org/10.1038/nature07991>
- Solopchuk, O., Alamia, A., Dricot, L., Duque, J., & Zénon, A. (2017). CTBS disruption of the supplementary motor area perturbs cortical sequence representation but not behavioural performance. *NeuroImage*. <https://doi.org/10.1016/j.neuroimage.2017.09.013>
- Solopchuk, O., Alamia, A., Olivier, E., & Ze, A. (2016). *Chunking improves symbolic sequence processing and relies on working memory gating mechanisms*. *23*(2016), 108-113. <https://doi.org/10.1101/lm.041277.115>
- Solopchuk, O., Alamia, A., & Zénon, A. (2016). The role of the dorsal premotor cortex in skilled action sequences. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.1199-16.2016>

- Sterzer, P., Adams, R. A., Fletcher, P., Frith, C., Lawrie, S. M., Muckli, L., Petrovic, P., Uhlhaas, P., Voss, M., & Corlett, P. R. (2018a). The Predictive Coding Account of Psychosis. In *Biological Psychiatry* (Vol. 84, Numéro 9, p. 634-643). <https://doi.org/10.1016/j.biopsych.2018.05.015>
- Sterzer, P., Adams, R. A., Fletcher, P., Frith, C., Lawrie, S. M., Muckli, L., Petrovic, P., Uhlhaas, P., Voss, M., & Corlett, P. R. (2018b). The Predictive Coding Account of Psychosis. *Biological Psychiatry*, *84*(9), 634-643. <https://doi.org/10.1016/j.biopsych.2018.05.015>
- Steyvers, M., Lee, M. D., & Wagenmakers, E. J. (2009). A Bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, *53*(3), 168-179. <https://doi.org/10.1016/j.jmp.2008.11.002>
- Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making : Neural and computational mechanisms. *Nature Reviews Neuroscience*, *15*(11), Article 11. <https://doi.org/10.1038/nrn3838>
- Summerfield, C., & Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends in cognitive sciences*, *13*(9), 403-409. <https://doi.org/10.1016/j.tics.2009.06.003>
- Sun, Y., Farzan, F., Barr, M. S., Kirihara, K., Fitzgerald, P. B., Light, G. A., & Daskalakis, Z. J. (2011). Gamma oscillations in schizophrenia : Mechanisms and clinical significance. *Brain Research*, *1413*, 98-114. <https://doi.org/10.1016/j.brainres.2011.06.065>
- Tecumseh Fitch, W., & Friederici, A. D. (2012). Artificial grammar learning meets formal language theory : An overview. In *Philosophical Transactions of the Royal Society B: Biological Sciences* (Vol. 367, Numéro 1598, p. 1933-1955). <https://doi.org/10.1098/rstb.2012.0103>
- Trabelsi, C., Bilaniuk, O., Zhang, Y., Serdyuk, D., Subramanian, S., Santos, J. F., Mehri, S., Rostamzadeh, N., Bengio, Y., & Pal, C. J. (2017). *Deep Complex Networks*. 2016, 1-19.
- Uhlhaas, P. J., & Singer, W. (2006). Neural Synchrony in Brain Disorders : Relevance for Cognitive Dysfunctions and Pathophysiology. *Neuron*, *52*(1), 155-168. <https://doi.org/10.1016/j.neuron.2006.09.020>
- Uhlhaas, P. J., & Singer, W. (2010). Abnormal neural oscillations and synchrony in schizophrenia. *Nature Reviews Neuroscience*, *11*(2), Article 2. <https://doi.org/10.1038/nrn2774>
- Vaishnav, M., Cadene, R., Alamia, A., Linsley, D., VanRullen, R., & Serre, T. (2022). Understanding the Computational Demands Underlying Visual Reasoning. *Neural Computation*, *34*(5), 1075-1099. https://doi.org/10.1162/neco_a_01485
- Valton, V., Karvelis, P., Richards, K. L., Seitz, A. R., Lawrie, S. M., & Seriès, P. (2019). Acquisition of visual priors and induced hallucinations in chronic schizophrenia. *Brain*, *142*(8), 2523-2537. <https://doi.org/10.1093/brain/awz171>
- Valton, V., Romaniuk, L., Douglas Steele, J., Lawrie, S., & Seriès, P. (2017). Comprehensive review : Computational modelling of schizophrenia. *Neuroscience & Biobehavioral Reviews*, *83*, 631-646. <https://doi.org/10.1016/j.neubiorev.2017.08.022>
- Vanderclausen, C., Filbrich, L., Alamia, A., & Legrain, V. (2017). Investigating peri-limb interaction between nociception and vision using spatial depth. *Neuroscience Letters*, *654*, 111-116. <https://doi.org/10.1016/j.neulet.2017.05.060>
- van Kerkoerle, T., Self, M. W., Dagnino, B., Gariel-Mathis, M.-A., Poort, J., van der Togt, C., & Roelfsema, P. R. (2014). Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *Proceedings of the National Academy of Sciences*, *111*(40), 14332-14341. <https://doi.org/10.1073/pnas.1402773111>
- van Pelt, S., Heil, L., Kwisthout, J., Ondobaka, S., van Rooij, I., & Bekkering, H. (2016). Beta- and gamma-band activity reflect predictive coding in the processing of causal events. *Social Cognitive and Affective Neuroscience*, *11*(6), 973-980. <https://doi.org/10.1093/scan/nsw017>
- VanRullen, R. (2016). Perceptual Cycles. In *Trends in Cognitive Sciences* (Vol. 20, Numéro 10, p. 723-735). <https://doi.org/10.1016/j.tics.2016.07.006>
- VanRullen, R., & Alamia, A. (2021a). GAttANet : Global Attention Agreement for Convolutional Neural Networks. In I. Farkaš, P. Masulli, S. Otte, & S. Wermter (Éds.), *Artificial Neural Networks and Machine Learning – ICANN 2021* (p. 281-293). Springer International Publishing. https://doi.org/10.1007/978-3-030-86362-3_23

- VanRullen, R., & Alamia, A. (2021b). GAttANet : Global Attention Agreement for Convolutional Neural Networks. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 12891 LNCS, 281-293. https://doi.org/10.1007/978-3-030-86362-3_23
- Vanrullen, R., & MacDonald, J. S. P. (2012). Perceptual echoes at 10 Hz in the human brain. *Current Biology*, 22(11), 995-999. <https://doi.org/10.1016/j.cub.2012.03.050>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems, 2017-Decem*, 5999-6009.
- Vezoli, J., Vinck, M., Bosman, C. A., Bastos, A. M., Lewis, C. M., Kennedy, H., & Fries, P. (2021). Brain rhythms define distinct interaction networks with differential dependence on anatomy. *Neuron*, 109(23), 3862-3878.
- Vilares, I., & Kording, K. (2011). Bayesian models : The structure of the world, uncertainty, behavior, and the brain. In *Annals of the New York Academy of Sciences*. <https://doi.org/10.1111/j.1749-6632.2011.05965.x>
- Vinck, M., & Bosman, C. A. (2016). More Gamma More Predictions : Gamma-Synchronization as a Key Mechanism for Efficient Integration of Classical Receptive Field Inputs with Surround Predictions. *Frontiers in Systems Neuroscience*, 10. <https://www.frontiersin.org/articles/10.3389/fnsys.2016.00035>
- Vinck, M., Uran, C., Spyropoulos, G., Onorato, I., Broggin, A. C., Schneider, M., & Canales-Johnson, A. (2023). Principles of large-scale neural interactions. *Neuron*, 111(7), 987-1002. <https://doi.org/10.1016/j.neuron.2023.03.015>
- Visalli, A., Capizzi, M., Ambrosini, E., Mazzonetto, I., & Vallesi, A. (2019). Bayesian modeling of temporal expectations in the human brain. *NeuroImage*, 202, 116097. <https://doi.org/10.1016/j.neuroimage.2019.116097>
- Wacongne, C., Changeux, J.-P., & Dehaene, S. (2012). A Neuronal Model of Predictive Coding Accounting for the Mismatch Negativity. *Journal of Neuroscience*, 32(11), 3665-3678. <https://doi.org/10.1523/JNEUROSCI.5003-11.2012>
- Walker, E. Y., Cotton, R. J., Ma, W. J., & Tolias, A. S. (2020). A neural basis of probabilistic computation in visual cortex. *Nature Neuroscience*, 23(1), Article 1. <https://doi.org/10.1038/s41593-019-0554-5>
- Wang, X.-J., & Krystal, J. H. (2014). Computational Psychiatry. *Neuron*, 84(3), 638-654. <https://doi.org/10.1016/j.neuron.2014.10.018>
- Warrington, E. K., & Weiskrantz, L. (1968). A study of learning and retention in amnesic patients. *Neuropsychologia*, 6(3), 283-291. [https://doi.org/10.1016/0028-3932\(68\)90026-2](https://doi.org/10.1016/0028-3932(68)90026-2)
- Warrington, E. K., & Weiskrantz, L. (1982). Amnesia : A disconnection syndrome? *Neuropsychologia*, 20(3), 233-248. [https://doi.org/10.1016/0028-3932\(82\)90099-9](https://doi.org/10.1016/0028-3932(82)90099-9)
- Wymbs, N. F., Bassett, D. S., Mucha, P. J., Porter, M. a., & Grafton, S. T. (2012). Differential Recruitment of the Sensorimotor Putamen and Frontoparietal Cortex during Motor Chunking in Humans. *Neuron*, 74(5), 936-946. <https://doi.org/10.1016/j.neuron.2012.03.038>
- Zemel, R. S., Dayan, P., & Pouget, A. (1998). Probabilistic Interpretation of Population Codes. *Neural Computation*, 10(2), 403-430. <https://doi.org/10.1162/089976698300017818>
- Zhang, H., Watrous, A. J., Patel, A., & Jacobs, J. (2018). Theta and Alpha Oscillations Are Traveling Waves in the Human Neocortex. *Neuron*. <https://doi.org/10.1016/j.neuron.2018.05.019>
- Zhang, Z. G., Hu, L., Hung, Y. S., Mouraux, A., & Iannetti, G. D. (2012). Gamma-Band Oscillations in the Primary Somatosensory Cortex—A Direct and Obligatory Correlate of Subjective Pain Intensity. *Journal of Neuroscience*, 32(22), 7429-7438. <https://doi.org/10.1523/JNEUROSCI.5877-11.2012>
- Zhou, H., Schafer, R. J., & Desimone, R. (2016). Pulvinar-Cortex Interactions in Vision and Attention. *Neuron*, 89(1), 209-220. <https://doi.org/10.1016/j.neuron.2015.11.034>