Bidirectional interactions between sleep slow waves and synaptic plasticity



Martin Christopher Kahn

Lincoln College

Department of Physiology, Anatomy and Genetics

University of Oxford

Thesis submitted for the degree of

Doctor of Philosophy

Michaelmas 2019

For Hans-Werner Kahn,

and Eri and Claudia Kahn

I wish you could be here to celebrate

&

For Christine Kahn and Sarah Kahn

I cannot say how glad I am to have you

Abstract

Sleep is a fundamental drive akin to hunger and thirst. Yet, despite decades of progress, we cannot say why the body needs sleep. One promising possibility is that sleep plays a critical role in the plasticity of neuronal circuits and in the consolidation of memories. This is underscored by the wealth of behavioural evidence linking sleep to enhanced memory recall. However, there are diverging theories of how sleep benefits memory and how it influences neuronal networks. In my thesis, I am exploring a mechanism by which sleep can influence synaptic strength. Furthermore, I am using this mechanism to contribute towards testing a hotly debated hypothesis on the function of sleep, the so-called synaptic homeostasis hypothesis of sleep (SHY).

My first chapter is an *ex vivo* characterisation of a mechanism by which the cortical slow oscillation (a dominant pattern of activity during sleep) affects neuronal plasticity. Specifically, the slow oscillation consists of alternating periods of high (UP state) and low (DOWN state) firing rates. Using single cell recordings, we find that the UP state induces weakening of subthreshold synaptic inputs, whereas the DOWN state has no such effect.

In my second experimental chapter, I analyse *in vivo* electrophysiological recordings from the somatosensory cortex of anaesthetised mice. I confirm that sensory responses become weaker when paired with UP states but not DOWN states. One prediction of SHY is that local changes in synaptic strength are associated with local changes in the slow oscillation. I find that local reductions in sensory responses are tightly associated with local reductions in the slow oscillation. My analysis thus supports SHY's prediction, even though it also suggests that further experiments are needed to confirm this.

In my third experimental chapter, I use freely moving mice to investigate whether UP states can facilitate synaptic weakening during natural sleep. I record field potentials and multi-unit responses in primary motor cortex in response to transcallosal electrical stimulation in the contralateral hemisphere. To be able to pair the electrical stimulation with the UP and DOWN state, I developed a closed-loop detection system for UP and DOWN states based on spike rates during natural sleep. I find that evoked responses in the LFP are decreased after UP state pairings but not DOWN state pairings or stimulation during waking.

Together, the data presented in this thesis suggest that cortical UP states constitute periods of enhanced potential for synaptic weakening in the absence of other sleep oscillations.

Acknowledgments

Above all, I would like to express my deepest gratitude to Ed Mann and Vlad Vyazovskiy. Thank you for giving me a scientific home, for supervising me, for being inspiring scientists, for the countless enlightening discussions about how the brain might work (and what the heck sleep might be there for), and for being genuinely warm, supportive and simply amazing people. Ed, thank you for believing in me and taking me into your lab. Vlad, thank you for coming aboard my DPhil train and letting me explore what brain slices can do when they have legs attached. I'm incredibly grateful to have done my DPhil with both of you and hope we will have many more discussions, collaborations and pints in the future.

Yves Weissenberger and Michael Lohse, what made my DPhil experience truly unique was having two such brilliant scientists around me pretty much every day. It was an honour and just plainly lot's fun. I shall never again flinch when I'm told that I'm full of sh… But most of all, thanks for being the amazing friends you are.

Julian Bartram and Antonia Langfelder, it has been endless fun to work with you and I am grateful for having been able to continue what you started. My thesis stands on your shoulders. Julian, it has been a privilege to patch away at unholy times of the day and to introduce the mouse brain[1] to all the colours of heavy metal and German swearing. I also owe a big thank you to the rest of Ed Mann's lab, especially Pantelis, Sonia, Sofia, Matt, Simon and Kat. May the god of patching ever keep you in his favour.

[1] Ex vivo, all experiments in full agreement with animal welfare standards. No swearing was directed at live animals

The lab members of the Vyazovskiy lab: thank you for all the help, collaboration and discussions in the last years. Laura and Tomoko, thank you for introducing me to the world of sleep recordings and analysis. Cristina and Lukas, it has been an honour to collaborate with you and scientifically grow up next to you. A special huge thanks to both of you for reading parts of this thesis. Cristina, I am incredibly grateful for how supportive (and tolerant) you were of me in the long last stretch of my thesis. Finishing the thesis often felt like traveling a long grey winter road (riding a blind donkey). Thank you for making sure there were some unicorns, butterflies, lots of sunshine and even the odd thundercloud.

I would also like to thank three people who so warmly opened the door to science to me all those years ago: Isabelle Mansuy, Li-Huei Tsai, and especially Johannes Gräff.

Last but not least, my deepest gratitude to my close friends. Thank you for being there despite my months-long social disappearances and, in some cases, despite the many airmiles between us. I suspect the years to come will not ease on the time-restrictions but know this: It's a privilege to have such deep friendships in life and maintaining them is among the most important things to me in life.

Table of Contents

CHAPTER 1: INTRODUCTION 12

1.1 SLEEP, SLOW WAVES AND SYNAPTIC PLASTICITY 12

1.1.1 Sleep and it's substages across the animal kingdom 12

1.1.2 The slow oscillation 15

1.1.3 Slow wave activity as a global and local marker for sleep pressure 18

1.1.4 The coordination of slow waves with other sleep oscillations 19

1.1.5 The memory function of sleep 21

1.2 THE SYNAPTIC HOMEOSTASIS HYPOTHESIS OF SLEEP 25

1.2.1 A definition of SHY and its fundamental premises 25

1.2.2 Where to be SHY 27

1.2.3 How to be SHY - what is synaptic strength? 29

1.2.4 SHY's first premise: waking causes net synaptic strengthening 31

1.2.4.1 Does net synaptic strength depend on sleep history? 31

1.2.4.2 Does learning cause net strengthening of synapses? 33

1.2.5 How could sleep support net weakening? 36

1.2.6 Is synaptic weakening during sleep physiologically relevant? 37

1.3 THESIS OBJECTIVES 38

CHAPTER 2: MATERIALS AND METHODS 40

2.1 ANIMALS 40

2.2 BRAIN SLICE EXPERIMENTS (CHAPTER 3) 40

2.2.1 Preparation of brain slices 40

2.2.2 Whole-cell patch clamp recordings in layer 3 of medial entorhinal cortex 41

2.2.3 Ex vivo electrical stimulation and plasticity protocol 42

2.2.4 Ex vivo pharmacology 43

2.2.5 Data analysis and statistics for patch clamp experiments 43

2.3 *IN VIVO* ANAESTHESIA EXPERIMENTS (CHAPTER 4) 44

2.3.1 Surgery (A.Langfelder, P.Antonoudiou) 44

2.3.2 Recordings, whisker stimulation and experimental design (A.Langfelder) 44

2.3.3 Analysis of LFP and MUA (MCK) 46

2.3.4 Analysis of SWA and detection of individual slow waves (MCK) 48

2.4 FREELY BEHAVING ANIMALS (CHAPTER 5) 49

2.4.1 Implants and surgical procedure 49

2.4.2 Electrophysiological recordings 51

2.4.3 Experimental Design 54

2.4.4 Closed loop electrical stimulation 56

2.4.5 Histology and identification of recording sites (MCK + Lukas Krone) *57*

2.4.6 Scoring of vigilance states *58*

2.4.7 Pre-processing of local field potentials and multiunit activity *59*

2.4.8 Quantification of evoked responses and plasticity thereof *62*

2.4.9 Statistics *64*

CHAPTER 3: CORTICAL UP STATES INDUCE SELECTIVE WEAKENING OF SUBTHRESHOLD INPUTS EX VIVO 67

3.1 INTRODUCTION 67

3.2. RESULTS 68

3.2.1 Cortical UP states facilitate weakening of sub- but not suprathreshold inputs *68*

3.2.2 UP state pairings are ineffective following potentiation but weakening is reversible *71*

3.2.3 Up state induced weakening involves a presynaptic expression mechanism but is induced postsynaptically *73*

3.2.4 UP state induced weakening depends on GSK-3 beta but not caspase-3 ..... *76*

3.3 DISCUSSION 78

CHAPTER 4 – NEURONAL PLASTICITY IS ASSOCIATED WITH CHANGES IN CORTICAL SLOW WAVES *IN VIVO* UNDER ANAESTHESIA 84

4.1 INTRODUCTION 84

4.2 RESULTS 85

4.2.1 Whisker-evoked responses and UP/DOWN state pairings in vivo *85*

4.2.2 Whisker evoked responses are weakened following UP but not DOWN state pairings *88*

4.2.3 Plasticity is associated with a dose-dependent decrease in slow wave activity *92*

4.2.4 Changes in slow waves following plasticity are comparable to changes following natural sleep homeostasis *94*

4.3 DISCUSSION 96

4.3.1 UP state induced plasticity *96*

4.3.2 Changes in SWA accompany changes in evoked responses *99*

4.3.3 Changes in the slow oscillation following UP state pairings *100*

4.3.4 The effects of anaesthesia *101*

CHAPTER 5: INFLUENCE OF SLOW WAVES ON EVOKED RESPONSES DURING NATURAL SLEEP 104

5.1 INTRODUCTION 104

5.2 RESULTS 105

5.2.1 Laminar responses to contralateral electrical stimulation *105*

5.2.2 Closed-loop stimulation during ON and OFF states in freely moving mice *110*

5.2.3 Effects of ON and OFF state pairings on evoked responses during waking *115*

5.2.4 Effects of ON/OFF state pairings on evoked responses during NREM sleep *121*

5.2.5 Effects of electrical stimulation on sleep architecture and slow wave activity *125*

5.3 DISCUSSION 130

5.3.2 Online detection of ON/OFF states *131*

5.3.3 Plasticity *132*

5.3.4 Effects of stimulation on sleep *135*

6. GENERAL DISCUSSION **138**

6.1 UP state induced plasticity during sleep *138*

6.2 Consequences for SHY *141*

6.3. Next steps *142*

7. BIBLIOGRAPHY **143**

Chapter 1: Introduction

1.1 Sleep, slow waves and synaptic plasticity

1.1.1 Sleep and it's substages across the animal kingdom

«sleep is for whimps». This sentiment by Margaret Thatcher unfortunately summarises an attitude towards sleep that is shared by many in modern society. While there is now a solid body of evidence suggesting that sleep is crucial for normal brain function, we are still not able to pinpoint what exactly makes sleep so important (Krueger et al. 2016; Wulff et al. 2010). Indeed, it seems most likely that sleep serves more than one function, even if its original adaptive value may have had a single cause. Identifying the functions of sleep could have far-reaching consequences for humanity: We may be able to enhance specific functions of sleep while suppressing others, we might one day be able to speed up sleep's work and gain up to 8h per day, and we can simply adjust our behaviour to maximise the beneficial effects of sleep

Sleep is commonly defined through behavioural and/or physiological criteria. Behavioural criteria include reduced motor output, increased sensory response thresholds, species-specific postures and sleep sites and a homeostatic need for sleep (Campbell & Tobler 1984; Tobler 1995). Physiological criteria for sleep in mammals and birds rely on the electromyogram (EMG) and the electroencephalogram (EEG), which also allow to further subdivide sleep into stages: Rapid eye movement sleep (REM) displays wake-like EEG patterns activity and complete loss of muscle tone, except for the characteristic rapid eye movements and phasic twitches in limbs and face muscles. In contrast, non-REM (NREM) sleep is defined by the occurrence of large amplitude slow waves and sleep spindles in the EEG. REM sleep further distinguishes

itself by decreased thermoregulation and irregular heart and breathing rates. In humans but not in other animals NREM sleep is further subdivided into 3 stages.

Applying the above criteria, sleep-like states have been proposed to occur in almost all species studied thus far (Vorster & Born 2015; Lesku et al. 2009; Tobler 1995) including reptiles, amphibians, fish (Zhdanova et al. 2001), insects such as fruit flies (Shaw et al. 2000), bees (Zwaka et al. 2015) and cockroaches (Tobler & Neuner-Jehle 1992). Even the roundworm *C.elegans* has a sleep-like state during development and in response to stressors (Raizen et al. 2008). Nevertheless, it is worth noting that the absence of something is always more difficult to prove than its presence, and that there are studies finding no evidence of sleep (see (Lesku et al. 2009)). Indeed, at least one species of Mexican cavefish may have minimized its need for sleep (Duboué et al. 2011).

Moreover, one can ask whether these purely behavioural definitions for fish and invertebrates are sufficiently specific. For instance, if one would look at human behavioural data from which all sleep was deleted, it would still be possible to find periods that fulfil the behavioural criteria for sleep. In each of our daily lives there are periods where we move less, perhaps rest at a specific type of place (e.g. a bench in the park) and are slower to respond than at other times (e.g. sitting on a bench vs driving on the highway). It is also not difficult to imagine that making us engage in mental or physical exercise for several hours might increase the need for having such a period of rest. Indeed, I would argue that the main reason for knowing that a human is not asleep when resting on a park bench is that we have a clear reference (i.e. we can put ourselves in the position of the observed person and we know how human sleep looks like). Once we do not know whether a species does sleep, it becomes difficult to tell what different

levels of arousal and locomotion mean. One possibility to help this issue is to examine neuronal and molecular pathways that drive behaviour. For instance, the sleep of many invertebrate species is affected by caffeine or hypnotics (Bjorness et al. 2009) and the zebrafish hypocretin system is involved in regulating states that fulfil the behavioural criteria of sleep (Prober et al. 2006). However, to come back to the park bench example, caffeine would likely also modulate how long we would want to sit still. Therefore, it is hard to differentiate between sleep and the presence of different levels of alertness without a clear neuronal marker such as the EEG. On the other hand, it is important to acknowledge that neuronal markers of sleep are based on correlational observations, not on an understanding of sleep. In other words, we do not fully understand the purpose of the neuronal activity patterns we use to detect sleep or to what extent other activity patterns and circuit architectures could fulfil the same functions. In summary, sleep-like behaviours are a fundamental trait across all animals, and it seems very possible that fish and invertebrate models can reveal key functions of sleep. However, mammals and birds have the clear advantage of displaying brain activity that strongly resembles human sleep.

Rodents have become a popular model organism for sleep, and I used the mouse for the entirety of my thesis. The popularity of the rodent is undoubtedly justified, given the presence of key aspects of human sleep such as slow wave sleep, sleep spindles, and REM sleep, and given the accessibility of the brain to neuronal recordings and genetic manipulations. However, it is crucial to be aware that sleep in mice displays some quite prominent differences to its human counterpart. For instance, while laboratory mice may spend roughly 50% of an entire day asleep, they do so in a highly fragmented manner, switching between states over a hundred times a day (Franken et al. 1998).To exemplify this, a mouse can easily have over 30 individual bouts of REM sleep within

24h (Franken et al. 1998) while a human would be expected to have around four. Similarly, the average NREM episode in mice lasts around 5-10 minutes while it lasts 2-3 hours in humans (Vorster 2018). An additional difference is that mice are opportunistic heterotherms, i.e. they can enter torpor. Last, and probably also least, mice are mostly active during the night in the laboratory. It is difficult to predict what consequences these differences have for the function of sleep in mice, but they should certainly be kept in mind.

1.1.2 The slow oscillation

It was already recognised by the pioneers of the EEG that the brain can exhibit high amplitude, slow waveforms in the range of 0.5 – 4 Hz (Berger 1933). Berger considered these waves clinical anomalies, and slow waves were indeed heavily studied as markers for tumours (Walter 1936). However, it was also quickly recognised that slow waves occur during natural sleep (Loomis et al. 1935). While EEG slow waves were central to many discoveries in the emerging sleep field, the first clear insights into their cellular origin were made through intracellular recordings in anaesthetised cats (Wilson & Groves 1981; Steriade et al. 1993). Steriade and colleagues simultaneously recorded the membrane potential of single cells and the EEG in cats. They established that EEG slow waves are highly correlated with slow oscillations (< 1Hz) in single neurons, which oscillate between a hyperpolarised, low firing rate DOWN state and a depolarised, high firing rate UP state. They (and others) observed the slow oscillation not only in neocortical areas but also in lower and higher order thalamic nuclei (e.g. (Sheroziya & Timofeev 2014)) and the reticular nucleus of the thalamus (Steriade et al. 1993). Moreover, slow oscillations in the LFP and multi-unit activity have been reported in the granular cells of the cerebellum (Ros et al. 2009; Canto et al. 2017), the amygdala

(Collins et al. 2001) and the hippocampus (Wolansky et al. 2006). Steriade and colleagues also noticed that a large fraction of excitatory and inhibitory neurons (ca. 90%) in cortex shared such slow oscillations in the membrane potential (Steriade, Nuñez, et al. 1993a). Importantly, this high degree of participation refers to subthreshold slow waves, which mainly reflect the input a cell receives. When one examines the probability that a given cell spikes in an UP state, participation is much sparser, ranging from 40% in layer 5 and roughly 15 % in layers 2/3 and 6 of cortex (Sakata & Harris 2009).

The mechanisms for generating slow waves have proven surprisingly intricate and are still incompletely understood (Neske 2015). Moreover, the vast majority of what we know comes from anaesthetised animals or even brain slices. Nevertheless, it is clear that cortex alone is sufficient to generate a version of the slow oscillation because slow waves persist after thalamic lesions (Steriade, Nuñez, et al. 1993b), in a (sufficiently large) isolated gyrus *in vivo* (Timofeev et al. 2000) and in cortical brain slices (Sanchez-Vives & McCormick 2000). Specifically, layer 5 is a likely origin for many cortically-generated slow waves. UP-state firing in this layer tends to precede other layers *in vivo* and *in vitro* (Sanchez-Vives & McCormick 2000; Sakata & Harris 2009) and transection of laminar connections *in vitro* abolishes superficial but not infragranular slow waves (Sanchez-Vives & McCormick 2000). In addition, optogenetic activation of neuronal subpopulations in layer 5 can synchronise ongoing slow waves and drive OFF-ON transitions (Beltramo et al. 2013).

While cortex is thus necessary and sufficient for slow waves, such cortex-only slow waves appear very different from sleep slow waves. Indeed, under natural conditions the thalamus (and other structures) may well play equally important role in the

generation of slow waves. In accordance with this, pharmacological inactivation of thalamic nuclei reduces the incidence of slow waves and delta power (David et al. 2013; Lemieux et al. 2014; Doi et al. 2007) and bidirectional TRN modulation has a bidirectional effect on delta power during sleep (Lewis et al. 2015). Moreover, thalamic stimulation is effective at initiating and entraining UP states in various paradigms including sensory stimulation (Ngo et al. 2013; Petersen et al. 2003), optogenetic stimulation of thalamic fibers in brain slices (Favero & Castro-Alamancos 2013) and during natural sleep (David et al. 2013; Gent et al. 2018). Moreover, the previously mentioned study by (Lewis et al. 2015) demonstrated that localised or global activation of the TRN during waking can induce local or widespread slow waves, respectively. The final argument for a strong thalamic role is that thalamic activity often precedes activity in directly connected cortical areas under anaesthesia (Steriade & Contreras 1995; Slézia et al. 2011) and during natural sleep (Gent et al. 2018). Finally, it has been proposed that thalamic slices, like cortical slices, can intrinsically generate synchronised slow waves under the right circumstances (Hughes et al. 2004; Crunelli & Hughes 2010).

In summary, I would argue that slow oscillations in the EEG of a naturally sleeping animal reflect a mixture of several circuit motifs that can impose UP and DOWN states on cortex. From the perspective of a single cortical column, at least 3 different types of slow waves should be possible: Locally generated SOs (e.g. from layer 5), thalamically generated SOs, and inherited slow waves, which were generated somewhere else and have spread via horizontal or thalamic transmission (Amzica & Steriade 1995; Massimini et al. 2004a) . Indeed, this latter type may be the most common occurrence in the majority of cortex as no single region appears as a major source of slow waves, even

though roughly half of all slow waves initiate in frontal regions and travel caudally (Massimini et al. 2004b).

1.1.3 Slow wave activity as a global and local marker for sleep pressure

The motivation to sleep, like other essential needs such as hunger or thirst, is regulated by a combination of acute need and anticipation. In other words, sleepiness and hunger grow stronger if we have not slept or eaten in a long time but even after a night of complete sleep deprivation, we feel more sleepy at 3 AM compared to 5 hours later, when we would normally wake up. The interaction between such homeostatic and circadian contributions to sleepiness has famously been formalised by (Daan et al. 1984; Borbély 1982) who suggested that circadian and homeostatic sleep pressure sum up linearly. The former, a homeostatic pressure for sleep, is familiar to us all and it is thus relatively unsurprising that prolonged wakefulness increases subjective tiredness, shortens sleep latency and leads to rebound sleep (Fröberg et al. 1975; Borbély 1982). What may be surprising, however, is that rebound sleep duration does not fully compensate for the lost sleep time, the most dramatic example for which is 264 h of waking which was only followed by 14 h of sleep (Gulevich 1966). Thus, it was proposed that a second compensatory mechanism may be an increased intensity of sleep, and slow-wave activity may be a physiological marker for the intensity of sleep (Daan et al. 1984; Borbély 1982). This hypothesis was based on a solid body of data showing that SWA is increased after sleep deprivation, decays throughout sleep (Blake & Gerard 1937), is decreased when sleep is preceded by a nap (Dijk et al. 1987), and correlates well with arousal thresholds (Daan et al. 1984; Borbély 1982). However, the authors did not only propose SWA to reflect sleep intensity, they also proposed that it is largely independent of circadian factors, thus making it a specific measure for

homeostatic sleep need. Key evidence for this view came from forced desynchrony experiments, where subjects slept at intervals longer than the circadian period (e.g. 28h), thus constantly shifting sleep time relative to circadian time. In these experiments, slow wave activity was predictable solely by sleep debt but not circadian time while sleepiness was a function of both (Dijk & Czeisler 1995)

While EEG slow wave activity is now widely accepted as a marker for global homeostatic sleep pressure, more recent data suggest that certain aspects of sleep can occur in a highly localised manner. For instance, localised somatosensory stimulation or deprivation before sleep lead to localised increases (Kattler et al. 1994) and decreases (Huber et al. 2006) in EEG power, respectively. Furthermore, regional brain activity during waking can show sleep-like activity in a sleep-debt dependent fashion (Rector et al. 2005; Vyazovskiy et al. 2011). Finally, there is a strong anteroposterior gradient in slow wave activity across cortex. Together, these data suggest that slow wave activity can be regulated on a local level and that localised variations and changes in SWA may reflect localised aspects of sleep pressure that exist in conjunction with global sleep need.

1.1.4 The coordination of slow waves with other sleep oscillations

Other patterns of synchronised network activity occur during sleep and interact with the slow oscillation. Sleep spindles are relatively fast (10 - 15 Hz in mice) bursts of high amplitude oscillatory activity in the EEG that can last up to three seconds (Loomis et al. 1935; Steriade et al. 1993). They get their name from the impression that their amplitude ramps up and down, giving them a spindle-shaped envelope. Spindles are most likely initiated in the reticular nucleus of thalamus, where hyperpolarisation-activated inward currents interact with low-threshold calcium channels (Steriade et al.

1987; Fuentealba & Steriade 2005). Akin to slow waves, spindles can reflect sleep pressure (though becoming less prominent) and have been extensively linked to neuronal plasticity and memory(Rosanova & Ulrich 2005). For instance, spindle incidence correlates with prior learning, correlates with recall performance and their manipulation can affect memory (Gais et al. 2002; Latchoumane et al. 2017a).

Another key oscillation for learning and memory is the hippocampal sharp-wave ripple (SWR). SWRs are high amplitude deflections in the LFP (the sharp wave), which are crested by a high frequency component (the ripple, 150-250 Hz). These oscillations arise mostly from intrahippocampal dynamics and can also occur during wakefulness. Like spindles and slow waves, SWR incidence and amplitude depends on sleep pressure and correlates with learning and memory recall. Moreover, interrupting SWRs after learning impairs memory recall and cell assembly stabilisation (Girardeau et al. 2009a; van de Ven et al. 2016).

While SOs, SWRs and spindles can be observed independently in a brain slice (Schlingloff et al. 2014; von Krosigk et al. 1993; Sanchez-Vives & McCormick 2000), the intact forebrain is extensively interconnected, providing ample opportunity for the oscillations to interact. Indeed, the LFP and unit activity in the SO-range and spindle-range in cortex is correlated with hippocampal SWR activity (Sirota et al. 2003; Siapas & Wilson 1998; Staresina et al. 2015). Similarly, spindle and SO activity is correlated within cortex and the spindle-triggered and SO-triggered average EEG and LFP displays clear SO and spindle activity, respectively (Mölle et al. 2002; Steriade 1999; Klinzing et al. 2016). However, this coupling is not absolute: Around 60% of slow spindles are associated with slow waves and around 15% of slow waves are associated with a spindle (Muehlroth et al. 2019). Nevertheless, there is substantial evidence that

these three oscillations functionally interact. For instance, spindles recorded in the thalamus become less synchronous upon removal of cortex (Contreras et al. 1997), and optogenetic stimulation in the reticular nucleus can induce simultaneous spindles and SWR only if delivered into the UP state of the SO (Latchoumane et al. 2017a). Moreover, the co-occurrence of a spindle and slow wave can alter dendritic calcium dynamics (Niethard et al. 2018).

In summary, NREM sleep is the scene of an intricate interplay between SOs, SWR and spindles, which can appear in isolation, pairs, and even triplets. Each of these at least seven combinations can vary regarding key characteristics (i.e. frequency, amplitude, locality) and has different spatiotemporal dynamics (i.e. hippocampus connects directly to prefrontal but not somatosensory cortex). In the light of this diversity, a degree of functional heterogeneity seems likely. Nevertheless, the coupling of these oscillations is a very elegant and likely mechanism or systems consolidation during sleep (Klinzing et al. 2019).

1.1.5 The memory function of sleep

If there were a measure that ranked biological problems according to their importance and our lack of understanding of their function, sleep may well receive the highest score. It is very likely that no animal can forgo sleep, yet we cannot say why that is. Nevertheless, researchers have produced a wealth of insight into sleep's purpose. Indeed, this research suggests that the difficulty with elucidating sleep's function may be its involvement in many parallel functions, which may depend on phylogeny, developmental stage and even immediate environment. The most commonly proposed functions of sleep include an immune function, thermoregulation, (extra)cellular housekeeping, energy homeostasis, and memory (Krueger et al. 2016). For the

remainder of my thesis, I will only focus on functions under the umbrella of "neuronal plasticity and memory". However, it is important to keep in mind that the abovementioned parallel nature of sleep's functions may not be correct. In that scenario, many observations about "functions" of sleep could be consequences of one key function. For instance, memory could suffer after sleep deprivation because the brain's immune system is dysfunctional and thus circuitries and/or molecular consolidation processes become dysfunctional.

The idea that sleep plays a role in learning and memory is old (Jenkins & Dallenbach 1924), and there is unequivocal evidence that sleep is beneficial for memory in several circumstances (Rasch & Born 2013; Jeffrey M Ellenbogen et al. 2006). Thus the question is not whether sleep can benefit memory but whether it does so actively or passively, whether it is permissive (no sleep = no consolidation) or modulatory (sleep allows different quality or quantity of consolidation), and which types of memory are sleep-dependent (Vertes 2004; Jeffrey M Ellenbogen et al. 2006; Vyazovskiy et al. 2017).

From a behavioural perspective, sleep has been shown to benefit procedural (Walker et al. 2002) as well as declarative memory tasks (Rolls et al. 2011; Edwin M. Robertson et al. 2004) and critical period plasticity (Frank et al. 2001). However, not all tasks benefit from sleep (Edwin M Robertson et al. 2004; Donchin et al. 2002; Dadvand et al. 2006; Sawangjit et al. 2018a), thus arguing against a universal need for sleep in all memory processes. From this purely behavioural point of view, I see two challenges for concluding on an active and permissive role of sleep in memory consolidation. An active role of sleep is difficult to prove with purely behavioural means because memory consolidation can be impaired by providing interfering experiences (Underwood &

Postman 1960; Squire & Alvarez 1995; Edwin M. Robertson et al. 2004). It thus seems difficult to behaviourally exclude the possibility that sleep simply provides a period of low levels of interference (Vertes 2004), although some authors have attempted experiments to account for such interference effects and might thus disagree with me here (Jeffrey M. Ellenbogen et al. 2006; Jeffrey M Ellenbogen et al. 2006). An interesting problem for a permissive role of sleep in memory is that many studies find sleep deprivation only effective at interfering with memory within the first few hours after acquisition (Sawangjit et al. 2018b). Hence, if a subject is sleep deprived for a few hours after learning, the memory does not consolidate as well or not at all. The timeframe of such "need for consolidation sleep" seems very short from an intuitive and evolutionary perspective. If an animal had an obligatory need to sleep within 2-4 hours after learning anything, it would severely restrict behaviour and it feels that our society would be more acutely aware of this.

An important argument going beyond pure behaviour is that we observe various physiological processes during sleep that are highly suggestive for a role in memory consolidation. For instance, neuronal sequences that occurred during learning are replayed during sleep in hippocampus and cortex (Wilson & McNaughton 1994; Ji & Wilson 2007) and several forms of synaptic restructuring appear sleep-dependent (de Vivo et al. 2017a; Li et al. 2017). Moreover, these processes are often modulated by experience exactly like one would expect from a substrate of memory consolidation (Klinzing et al. 2019). For instance, representations of emotionally salient locations are preferentially replayed during NREM sharp wave-ripples (Girardeau et al. 2009b) and the amount of sleep spindles increases with learning and correlates with recall performance (Gais et al. 2002). Importantly, these processes can lead us to a mechanistic accounts of how sleep might interact with memory (Rasch & Born 2013).

Of course, these observations prompt two questions: Firstly, are these processes exclusive to sleep (sleep being permissive) or perhaps just more salient or common? Indeed, replay has been observed during wakefulness (Karlsson & Frank 2009) and such awake replay may be sufficient to reshape the representation of learned information (Liu et al. 2019). However, awake replay differs from replay during sleep and may serve a different function (Tang et al. 2017). The second question is whether any of these processes are necessary for memory consolidation.

Perhaps the most striking argument for a role of sleep in memory is the manipulation of consolidation by manipulating specific processes during sleep without disrupting sleep itself. Beautiful examples of this are the enhancement or memory recall by enhancing or interfering with slow waves (Marshall et al. 2006; Ngo et al. 2013; Fattinger et al. 2017), the disruption of REM theta rhythms disrupting consolidation in mice (Rolls et al. 2011), the decrement in critical period plasticity and consolidation after complete silencing of relevant brain regions during sleep (Jha et al. 2005; Sawangjit et al. 2018b), the decrement in memory performance after suppression of sharp wave ripples (Girardeau et al. 2009a) and the bidirectional modulation response of memory performance by corresponding manipulations of sleep spindles (Latchoumane et al. 2017b). These studies are clearly tantalizing arguments for an active role of sleep, yet our limited overall understanding of memory and other roles of sleep leaves open alternative explanations for these findings. Firstly, because we do not have a complete mechanistic understanding of memory consolidation and how new learning interferes with this, we cannot exclude that (i) we may increase/decrease the amount of interfering learning even during sleep, (ii) we modulate processes that are not sleep-dependent but are perhaps just easier to detect and thus specifically manipulate during sleep. Secondly, it is possible that memory is affected indirectly as a consequence of manipulating

another function of sleep. For instance, enhancing slow waves could primarily support an immune function or the cellular maintenance function of sleep (Vyazovskiy & Harris 2013; Besedovsky et al. 2017).

Overall, it is clear that sleep benefits memory in many settings and that neuronal activity during sleep is relevant for this. However, it is not trivial to assess the relative importance of passive and active contributions of sleep and the relative importance of other functions of sleep towards memory consolidation. In addition, more experiments are needed to address the distinction between sleep being permissive to consolidation or simply beneficial. Clearly, a fruitful avenue will be to keep developing and refining mechanistic models of how sleep can benefit memory consolidation. The more precisely we can modulate sleep's presumed role in memory, the more unequivocally we can conclude on the above questions. The two currently most prominent ideas on how sleep can directly and actively benefit memory are the synaptic homeostasis hypothesis of sleep (Tononi & Cirelli 2003) and what I will refer to as the systems consolidation hypothesis of sleep (Klinzing et al. 2019). Of note, these ideas are not mutually exclusive (Born & Feld 2012). Nevertheless, they often make separable predictions that can be individually tested. My thesis ended up being most strongly relevant to SHY and I will thus focus on exploring this hypothesis for the rest of this introduction. However, many of the elements relevant to systems consolidation have been introduced above.

1.2 The synaptic homeostasis hypothesis of sleep

1.2.1 A definition of SHY and its fundamental premises

The fundamental idea behind SHY is that sleep counterbalances a net synaptic strengthening that occurs due to learning during waking (Tononi & Cirelli 2003). This

focus on net synaptic strength is what clearly distinguishes SHY from other ideas about a key role of sleep in memory and synaptic plasticity.

SHY is motivated by several theoretical considerations regarding brain function. Firstly, it fields the uncontentious argument that synaptic transmission costs energy and causes cellular stress, thus making synaptic strength a likely target for tight biological oversight. Secondly, SHY argues that neuronal coding is sparse, which should ultimately require low synaptic strength. Thirdly, Tononi and Cirelli argue that learning is associated with net synaptic strengthening and should mainly happen during wake, while renormalisation of the network should happen during sleep, when there is no bias caused by incoming information.

When discussing and testing SHY, it is crucial to distinguish it from related theories and research fields. Indeed, the general idea that sleep benefits memory and neuronal wiring in some form has been intensely researched well before and in parallel with SHY (Frank 2016). Similarly, sparse coding and homeostatic plasticity are entire fields of their own. So, what are the central pillars that distinguish SHY and make it a scientifically testable hypothesis? What is the set of criteria that must not be false for SHY to remain plausible? First, learning (or at least waking) has to be associated with a net increase in synaptic strength (modified quote from C.Cirelli, oral presentation, London 2018). Secondly, synaptic renormalisation should overwhelmingly happen during sleep (modified quote C.Cirelli). If renormalisation happens equally in all states, it is unlikely a key function of sleep. This also implies that there must be a mechanistic explanation for why sleep is particularly effective at synaptic renormalisation or why it exclusively happens during sleep. Thirdly, assuming synaptic renormalisation during sleep occurs, it must be physiologically relevant. In other words, it is not enough that an effect exists,

it also must be strong enough to be relevant. Indeed, the key function of sleep should be able to account for key characteristics of sleep and lack thereof, such as sleep homeostasis, memory consolidation and reduced cognitive performance following sleep deprivation.

Many of the above points are vague, and thus the theory is in danger of becoming untestable. For instance, the criteria for SHY might be met by specific synapses with specific circuits under specific circumstances (e.g. recurrent excitatory synapses in early development of visual circuits) but not others. This leaves us with two options: We may test the exact wording of SHY, which stipulates that if we average across all synapses in all neurons, then learning will increase this number. The second option is to allow SHY some limited room to be only partially true. For instance, it could be accurate for neocortex and/or hippocampus but not the cerebellum. Similarly, it could be true during development but not adulthood. In my view, the latter option should be favoured, as it would still be a significant advancement for our understanding of sleep.

Below, I will briefly review evidence on SHY. I will focus on the evidence pertaining to the influences of sleep history on net synaptic strength. I will then move on to reviewing evidence on the physiological relevance of synaptic weakening. But before delving into assessing SHY, I will explore the notion of "net synaptic strength" and the technical challenges of measuring it, which are both crucial for evaluating the findings in my thesis.

1.2.2 Where to be SHY

The notion of average synaptic strength requires thinking about where the majority of synapses in the brain are located. As a first approximation, it seems clear that the focus

in humans can be limited to the cortex and cerebellum, as they together make up 90% of human brain mass and include 99% of all neurons (Azevedo et al. 2009). Furthermore, while the cerebral cortex only accounts for 20% of neurons, it constitutes 80 % of human brain mass. Hence, cerebral cortex likely accounts for the vast majority of human synapses. This discrepancy between number of neurons and mass fits with the notion that the number of synapses on cortical excitatory neurons ranges between 5000 and 20000 (DeFelipe et al. 2002) while the exceptionally numerous cerebellar granule cells receive and form well below 1000 synapses. In the mouse, the two structures account for 80% of neurons and only just above half the mass, with cortex alone accounting for 45% of brain mass (Herculano-Houzel et al. 2006). This "decreased cortical dominance" is exemplary for all non-primate mammals (Herculano-Houzel 2010) and suggests that in many mammals a relevant proportion of synapses are subcortical. Therefore, testing SHY on mouse cortex alone may not be sufficient to validate SHY but it is likely necessary. Mouse cortex accounts for roughly half the brain's synapses and it seems unlikely that sleep has a fundamentally different function in the cortex of mice and humans.

Another issue to consider is that the majority of cortical neurons are made of excitatory and inhibitory cells that form synapses onto each other, leading to at least four functionally different types of synapses. About 80-90% of these synapses are excitatory (DeFelipe et al. 2002), thus leaving SHY mainly with excitatory synapses onto excitatory neurons and onto interneurons. While interneurons only constitute a fifth of cortical neurons electron microscope-based tracing found that nearly half the local synapses of a pyramidal cell in mouse visual cortex are made onto interneurons (Bock et al. 2011; Lee et al. 2016). This is likely an overestimation of excitatory-to-inhibitory connectivity because most synapses of excitatory neurons are not local (Stepanyants et

al. 2009). Nevertheless, it seems to me that the population of cortical excitatory synapses onto interneurons is likely large enough to be relevant to SHY. Thus, again we can state that net weakening of excitatory-excitatory connectivity during sleep is necessary for SHY to be true, yet it may not be sufficient. Of note, the argument may even be slightly more complicated from an energy-perspective, as many inhibitory neurons fire at higher rates than their excitatory counterparts. Hence, the energy-expenditure due to inhibitory synapses may be much larger than their number suggests.

In summary, cortical excitatory synapses onto excitatory neurons are a good starting point for testing SHY. If they do not conform to the hypothesis, it is likely wrong. But if they do, SHY must be tested on other populations of synapses.

1.2.3 How to be SHY - what is synaptic strength?

Synaptic strength, like sleep, is an intuitively simple concept that can become surprisingly slippery under scrutiny. It is commonly defined by the overall effect of presynaptic activity on the postsynaptic neuron. This definition includes all aspects of neuronal communication, from synaptic transmission (including glia) over dendritic and somatic integration to action potential initiation and propagation. Notably, this definition makes synaptic strength a fast-changing parameter, as it is readily and strongly modulated by neuromodulators and short-term plasticity (on the order of seconds to minutes). However, neuromodulation and short-term plasticity are unlikely to underlie long-term memory and should thus mostly be seen as a confound for testing SHY. Long-term changes in synaptic strengths are usually (but not exclusively) attributed to changes in postsynaptic receptors or the presynaptic release machinery. Therefore, I will henceforth use the term synaptic strength to refer to long-lasting aspects of synaptic transmission strength. Before completely side-lining

neuromodulation, however, it is worth considering that from the perspective that synaptic transmission costs energy, a neuromodulatory interpretation of SHY could be fruitful. In other words, if the function of sleep is to reduce synaptic transmission and provide neuronal rest (Vyayovskiy and Harris 2004), the vastly reduced amounts of neuromodulation during sleep could be crucial.

Synaptic strength is commonly measured using structural or physiological measures. The former includes imaging of dendritic spines, dendritic glutamate receptors, or biochemical quantification of glutamate receptors of brain lysates or biochemically isolated synapses. Physiological measures usually quantify responses to either electrical, optogenetic, or sensory stimulation and range from single cell subthreshold recordings and local field potentials, to extracellular or EEG/EMG responses. Structural measures can suffer from the fact that they may not accurately reflect relevant parameters. For instance, biochemical quantification of AMPA receptors in brain tissue conflates synaptic and non-synaptic proteins. On the other hand, spine size correlates relatively well with physiological measures of synaptic strength in well-controlled conditions. However, sleep is a systemic event affecting parameters that could significantly affect spines, such as brain temperature, pH and interstitial volume. Furthermore, not all spines form functional synapses and we do not yet understand the purpose of many dendritic dynamics, which may be particularly relevant for studies on development. Nevertheless, for testing SHY structural measures are likely more suitable because they are more unbiased and generalized than physiological measures, which are always focussed on a single circuitry. Moreover, physiological measures are heavily influenced by behavioural state, such as locomotion (Niell & Stryker 2010) and by other physiological variables, such as temperature (Moser et al. 1993). Lastly, the most practical measures for studying sleep, the LFP and spiking activity, have a complicated

relationship with synaptic strength even in constant conditions. The LFP reflects a population wide convolution of excitation, inhibition and the synchrony between these events (Haider et al. 2016; Buzsáki et al. 2012). Spiking, on the other hand, reflects the strength of inputs as well as their spatiotemporal integration across a neurons' dendrites. Thus, while physiological measures are clearly useful for probing aspects of SHY, structural measures seem better suited specifically to test the first premise of SHY.

1.2.4 SHY's first premise: waking causes net synaptic strengthening

If waking does not cause net potentiation, SHY should be considered wrong. Furthermore, the authors of SHY argue that net potentiation should be due to learning, which means that if learning does not, as a rule, lead to net potentiation (in cortex), we should likely also consider SHY to be false. The latter point seems less obligatory as it may be possible that a net increase occurs without being related to learning in the computational sense (i.e. see the "neuromodulatory argument" in the following section).

This leaves us with two routes to assess the first premise of SHY: Firstly, the prediction of a net increase in synaptic strength during waking can be directly tested without further theoretical considerations. Secondly, we can draw on the vast field of learning and memory to assess the notion that learning is associated with net potentiation.

1.2.4.1 Does net synaptic strength depend on sleep history?

In recent years, several studies employed structural and physiological measures to test the interaction between sleep history and net synaptic strength. Using electron microscopy, (de Vivo et al. 2017b) surveyed layer II neocortical synapses and observed

a roughly 20% reduction in spine size in mice culled after being predominantly asleep for 7h compared to mice that were predominantly awake. Importantly, this was the case for both spontaneous waking at a different circadian time and for enforced waking at the same circadian time. One limitation of this study is that only synapses between excitatory neurons were studied (i.e. only one out of four possible types of ionotropic synapses) and that the mice were only 4 weeks old, which means they should be considered adolescent. Indeed, an earlier study using repeated *in vivo* spine imaging found net negative spine turnover after sleep in developing but not adult mice (Maret et al. 2011). This suggests that sleeps influence on net synaptic strength may depend on the stage of development. In line with this, ultrastructural studies in rats quite consistently report a net increase in spines during development (Koss et al. 2014; Juraska 1982). This could suggest that the data from (de Vivo et al. 2017b), may not necessarily hold in adult animals. Partially addressing this developmental concern is a study from (Diering et al. 2017) which repeatedly imaged surface AMPA receptors on dendritic spines of adult mice and reported a mean decrease in quantity of receptors 4h into the light cycle compared to 4h into the dark cycle. Unfortunately, the design of these experiments does not allow a discrimination between an effect of circadian time and an effect of sleep. The same "circadian concern" can be raised for key results in of the first studies to assess SHY in mice (Vyazovskiy et al. 2008), which observed similar net decreases of synaptic proteins in adult rats culled after a period of sleep. However, the authors also found increased evoked cortical local field potentials following sleep deprivation. This contradicts a circadian effect, even though evoked potentials are not trivial to interpret in this context (see page 29). Indeed, a study using evoked responses in cats directly conflicts with these results, finding enhanced responses after sleep (Chauvette et al. 2012), even though circadian factors were also not controlled for. In

summary, there is considerable support for a net negative effect of sleep on “unbiased” structural measures of synaptic strength. Nevertheless, it would be beneficial for the field to independently replicate the report (de Vivo et al. 2017b) and to conduct a similarly well controlled study in adult mice.

1.2.4.2 Does learning cause net strengthening of synapses?

It is beyond doubt that learning involves changes in synaptic strength. Furthermore, a clear intuition is that learning something new should be accompanied by an increase in excitatory connectivity. After all, this is exactly what Hebb had in mind when he famously suggested that learning may involve “some growth process or metabolic change” in the synapse between neuron A and B, if neuron A repeatedly takes part in firing neuron B (Hebb 1949). Indeed, the discovery of long-term potentiation (LTP) (Bliss & Lomo 1973) provided an elegant mechanism for Hebb’s postulate. However, today Hebbian plasticity usually refers to both LTP and long-term depression (LTD) (Abbott & Nelson 2000), exemplifying the notion that excitatory plasticity is bidirectional. Indeed, I would argue that the isomorphism of “new memory, hence new connection” is appealing and intuitive but there is no reason for it to be true in a complex learning system such as the cortex. To illustrate this, imagine a network of 100 neurons fully connected to each other, each synapse having a possible weight between zero and one. If you set most initial weights close to zero, learning almost certainly involves net potentiation. Conversely, if all weights are initiated almost at saturation, learning will likely cause weights to decrease on average. Consistent with this view, current reviews on the mechanisms of information storage focus to similar degree on

strengthening and weakening of synapses and on the formation and elimination of synapses (Holtmaat & Caroni 2016; Chklovskii et al. 2004).

On first glance, experimental evidence seems to support the notion that cortical and hippocampal learning temporarily involve increased numbers of synapses. For instance, the exposure to an enriched environment as well as spatial learning leads to transient increase in dendritic spines per dendritic branch (Moser & Trommald 1994; Rampon et al. 2000; O'Malley et al. 2000) in the hippocampus. Similarly, auditory fear conditioning, (Moczulska et al. 2013) and motor learning, (Xu et al. 2009) lead to a transient increase in spines and spine turnover in cortex. Moreover, monocular deprivation causes a lasting increase in spine density and selective whisker trimming induced equal gains and losses of spines in the barrel cortex (Holtmaat et al. 2006). While these observations are all correlational, there is evidence that shrinking newly potentiated spines after learning indeed interferes with memory (Hayashi-Takagi et al. 2015), suggesting new spines may be necessary for memory.

In the light of SHY, the problem with all these studies is that net changes in spine numbers are impossible to link to net changes in synaptic weights. Indeed, 24h after learning only 30% of new spines actually form synapses (Cane et al. 2014). More importantly, changes in the size of persistent spines likely also play a major role in learning. Hence, for all we know at this juncture, it could well be that learning transiently increases net connectivity (or even merely spine numbers) but has no consistent effect on synaptic weights. Intriguingly, this is precisely what is suggested by a study examining synapse number and size after learning (Eyre et al. 2003). The authors confirm a transient increase in spine number but find a concurrent decrease in synapse size. Hence, while learning clearly causes an increase in spine turnover and can

cause a net increase in synapse number, there is currently no evidence that learning causes a transient or permanent net increase in synaptic weights.

While I would thus argue that the learning field does not support or refute the idea of requiring net synaptic increases, it does clearly acknowledge the need for homeostatic checks on synaptic plasticity (Abbott & Nelson 2000). There are various proposals of mechanisms that could homeostatically control different parameters, such as firing rates or synaptic weights. Intriguingly, one form of homeostatic plasticity, synaptic scaling (Turrigiano 2012), has recently been shown to be unilaterally dependent on sleep. Homeostatic scaling refers to the adjustment of synaptic strength in order to keep neurons within a range of firing rates. (Hengen et al. 2016) found that homeostatic upregulation of firing could only occur during wakefulness. While spontaneous firing rates at best indirectly reflect synaptic weights, this line of investigation could ultimately establish whether homeostatic adjustments to synaptic weights depend on sleep or not.

To conclude, SHY's first premise has received considerable support, especially in developing mice. However, there are alternative explanations for the findings and the learning field cannot provide conclusive theoretical support that SHY is truly needed for learning. The findings of (Diering 2007) provide a good example why the observations of SHY could well be a side effect of a more relevant computational goal of sleep: The reported net decrease in spine size during sleep was driven by 58% of spines decreasing while 42% of spines increased. This at the very least argues that SHY acts in concert with other mnemonic functions of sleep (Born & Feld 2012). Moreover, it underscores the need to test other aspects of SHY, such as how weakening may work and whether it is physiologically relevant.

1.2.5 How could sleep support net weakening?

A common criticism of SHY is that it is unclear how sleep may facilitate net synaptic weakening (Frank 2012). Certainly, the ability to specifically manipulate such a mechanism will ultimately be necessary to proof the physiological importance of net weakening during sleep. So how could sleep bring about a net weakening? To me, there are two non-exclusive albeit separable sets of possibilities: Firstly, sleep's characteristic activity patterns (slow waves, sharp waves, spindles) could be sufficient to cause net weakening. Secondly, alterations in neuromodulatory tone or gene expression could alter the laws of neuronal plasticity, biasing them towards weakening.

There is some evidence that sleep's oscillations can promote plasticity. In mouse brain slices stimulation mimicking sharp wave ripples in the hippocampus can promote LTP (Sadowski et al. 2016), and spindle-like stimulation can induce LTP in the somatosensory cortex (Rosanova & Ulrich 2005). Moreover, one study aimed to mimic slow oscillations by injecting a SO-like pattern of hyperpolarising currents into cat pyramidal cells (Chauvette et al. 2012). When they paired this sleep-like stimulation with synaptic inputs, they observed LTP. It should be noted, however, that the most depolarised state these cells experienced was -65 mV, at which they were held in between negative current injections. While this adequately approximates the mean membrane potential of cells during UP states, it fails to account for the strong variation around this mean (after all, UP states are regularly associated with crossing the spike threshold at -45 mV) (Steriade et al. 2001) and the fact that UP states are generated by a barrage of excitatory and inhibitory synaptic activity (Shu et al. 2003), and not intrinsic currents. Indeed, it has been suggested previously that periods of modest depolarisation can be sufficient to generate LTD (Sjöström et al. 2004).

The interaction of neuromodulation and synaptic plasticity has been studied extensively. Several key modulatory systems drastically decrease their activity during NREM sleep, including the locus coeruleus (Aston-Jones & Bloom 1981), the basal forebrain (Xu et al. 2015), the raphe nucleus (McGinty & Harper 1976) and the tuberomammillary nucleus (Takahashi et al. 2006). Strikingly, all these neuromodulators have been implicated in neuronal plasticity (Pawlak et al. 2010; Takeuchi et al. 2016; Letzkus et al. 2011; Seol et al. 2007), even though it can be difficult to disentangle a role in memory from a role in other cognitive processes, such as sensory processing and attention (Sara 2009). Nevertheless, it is clear that NREM sleep provides a profoundly altered neuromodulatory environment compared to waking, which almost certainly has an influence on the rules of neuronal plasticity.

1.2.6 Is synaptic weakening during sleep physiologically relevant?

A second premise of SHY is that net weakening during sleep is an important function of sleep. In other words, experimentally altering net synaptic weights should (partially) mirror the effects of altering sleep history. This premise of SHY has so far received less attention than the first, which may be partially owed to the difficulty of manipulating synaptic weights.

Several studies have addressed whether learning and or altered synaptic weights can affect slow wave activity, thus testing if SHY may be able to account for slow wave activity homeostasis. Firstly, a computational study has shown that alterations in synaptic weights of a thalamocortical model could indeed alter SWA (Olcese et al. 2010). Secondly, several studies have modified waking experience and observed

changes in SWA. For instance, (Vyazovskiy et al. 2000) unilaterally trimmed the mouse whisker pad and observed reduced SWA activity in the contralateral hemisphere. Similarly, depriving mice of vision during development leads to reduced SWA in adulthood (Miyamoto et al. 2003). In humans, blindfolding during the day leads to reduced SWA during sleep and learning leads to localised increases in SWA (Huber et al. 2004). Furthermore, preventing practice leads to performance decrements and decreases in SWA (Huber et al. 2006). These studies show clearly that SWA is influenced by daytime activity and learning. However, it is difficult to dissociate use-dependent processes (energy stores, reactive oxygen species etc.) from the effects of synaptic plasticity.

The hitherto most direct of the interaction between synaptic potentiation and SWA came from (Huber et al. 2007), who applied a TMS potentiation protocol in humans and found increased SWA. This indeed strongly suggests that induction of plastic changes is associated with changes in SWA. However, whether rTMS protocols cause plasticity in the synapse (e.g. as opposed to excitability) is not yet clear (Cooke & Bliss 2006).

In summary, there is little doubt that local SWA is influenced by learning and local daytime dynamics. There is also direct evidence that induction of potentiating plasticity causes potentiation of SWA. However, it would still be valuable to gain direct evidence that local synaptic plasticity can locally modulate SWA and that interfering with synaptic plasticity during sleep inhibits SWA homeostasis.

1.3 Thesis objectives

The goal of my thesis is to shed light on the mechanisms by which slow waves could contribute to synaptic plasticity during sleep and to ask whether synaptic plasticity can

affect the slow oscillation in turn. With respect to SHY, my work aims to go beyond correlating synaptic weights with sleep history and to contribute towards establishing the mechanisms by which sleep can affect synaptic strength. Moreover, I aim to contribute towards the question whether synaptic weakening during sleep has physiological relevance. I approach this question by asking whether synaptic weakening can locally affect the slow oscillation.

The specific goals were as follows:

Chapter 3: To test the impact of cortical UP states on the rules of neuronal plasticity in an *ex vivo* preparation. Specifically, to examine the effects of pairing subthreshold and suprathreshold synaptic inputs with the cortical UP state. This was done using patch clamp recordings in brain slices of the mouse medial entorhinal cortex.

Chapter 4: To induce synaptic plasticity in the anaesthetised barrel cortex and test whether this results in local changes in the slow oscillation.

Chapter 5: To test the impact of cortical UP states on neuronal plasticity in naturally sleeping animals. Specifically, to test whether pairing transcallosal electrical stimulation with cortical UP or DOWN states during NREM sleep differentially affects the evoked response in the subsequent waking period.

Chapter 2: Materials and Methods

2.1 Animals

All experiments were carried out in accordance with the UK Animals (Scientific Procedures) Act of 1986. All animals used in this thesis were C57BL/6 purchased from Harlan Laboratories and kept on a regular (non-inversed) 12h light/dark cycle.

2.2 Brain slice experiments (chapter 3)

2.2.1 Preparation of brain slices

For the *ex vivo* experiments in chapter 3, animals aged 14 – 21 days were deeply anaesthetised with isoflurane (4%) and decapitated. The brain was rapidly removed into ice-cold artificial cerebrospinal fluid (aCSF) containing (in mM): 126 NaCl, 3.5 KCl, 1.25 NaH_2PO_4, 1 $MgSO_4$, 2 $CaCl_2$, 24 $NaHCO_3$, and 10 glucose, at pH 7.2-7.4 when supplied with carbogen gas (95 % O_2, 5 % CO_2). The concentrations of magnesium and calcium used here are optimised to minimise spontaneous UP state generation while maintaining the ability to electrically induce them. Lower calcium concentrations lead to spontaneous UP state generation (e.g.(Sanchez-Vives & McCormick 2000), and personal observations) while higher magnesium concentrations would impair the ability to electrically trigger them. The reasons for the former are likely a passive (charge screening) or active effect (Lu 2010) on neuronal excitability and for the latter the dependency of UP states on NMDA receptors (Castro-Alamancos 2015, personal observations). The brain was then transferred to a vibratome to obtain 350 µM thick horizontal sections containing the medial entorhinal cortex. Slices were allowed to recover in an interface chamber at room temperature for at least one hour before

experiments were started. For recordings, slices were transferred to a submerged recording chamber maintained at approximately 33 ^{o}C and fixated using a harp.

2.2.2 Whole-cell patch clamp recordings in layer 3 of medial entorhinal cortex

Borosilicate glass micropipettes were pulled to a resistance of 6-8 MΩ and filled with intracellular solution containing (in mM): 110 potassium-gluconate, 40 HEPES, 2 ATP-Mg, 0.3 GTP, 4 NaCl (pH 7.2; 270–290 mOsmol l^{-1}) and Biocytin (Sigma-Aldrich) at a concentration of 4 mg ml^{-1}. Recordings were made using an Axon Multiclamp 700 B amplifier (Molecular Devices), low-pass filtered at 2 kHz and digitised at 10 kHz with an ITC-18 A/D board (Instrutech). Custom-made (Ed Mann & Julian Bartram, minor modifications by M.Kahn) scripts in Igor Pro (WaveMetrics) were used for to acquire the data in consecutive 15 s long sweeps and for triggering electrical and intracellular stimulation.

Whole-cell current clamp recordings were obtained from layer 3 principal cells in the medial entorhinal cortex (mEC), which were identified based on localisation, morphology and electrophysiological properties in response to a step protocol. The input resistance was measured throughout the experiment based on a 20 pA hyperpolarising current step and was on average 207.1 ± 12.7 (n=21) MΩ , which is consistent with previous studies. Cells were excluded from the analysis if the input resistance or membrane potential in the last 10 minutes of a plasticity experiment was different from baseline by more than 30% and 8 mV respectively.

2.2.3 Ex vivo electrical stimulation and plasticity protocol

2.2.3 *Ex vivo* electrical stimulation and plasticity protocol

Afferent inputs to recorded cells were stimulated using current pulses (0.1 ms, 50-100 μA) delivered through monopolar tungsten electrodes placed in layer 5 of mEC. The currents were generated by stimulus isolation boxes (DS3 Isolated Constant Current Stimulator, Digitimer). The precise location of the stimulation electrode and the stimulation strength were adjusted to reliably trigger an excitatory post synaptic potential (EPSP) while minimising clear inhibitory components, disynaptic EPSPs or direct dendritic activation.

To trigger UP states in mEC, a second tungsten electrode was lowered into layer 2/3 of mEC at a horizontal distance of at least 200 μM from the recorded cell. As described previously (Mann et al. 2009) brief stimulation (0.1 ms, 50-200 μA) reliably brought the network into an UP state. As the ability of a slice to produce UP states likely reflects relevant parameters for plasticity, the presence of inducible UP states was verified irrespective of the experimental group.

During baseline and post pairing stimulations, afferent stimulation was triggered every 15 s. During subthreshold UP state pairings, an UP state was triggered every 15 s, followed 100 ms later by afferent stimulation at 5 Hz. This pairing was repeated 10 times with 5 afferent stimulations per UP state, resulting in a total of 50 pairings. Note that, while balanced across experimental groups, in some experiments the total of 50 pairings was achieved by 5 UP states and 10 stimulations per UP state. DOWN state pairings were identical to UP state pairings except that no UP states were triggered. During suprathreshold UP state pairings (Figure 8 a), conditions were identical to subthreshold UP state pairings, but each afferent stimulation was followed by a current injection through the patch pipette in order to evoke one or rarely two action potentials.

The current step necessary for this was calibrated before the beginning of the experiment (10 ms long steps starting 5 ms post stimulation, usually less than 1 nA). For suprathreshold burst pairings the current injections were longer and started earlier (20 ms step starting 2 ms after afferent stimulation).

2.2.4 *Ex vivo* pharmacology

In experiments where a drug was added into the patch pipette or the bath, at least 30 minutes were allowed for the drug to take effect before baseline recordings were started. The NMDA receptor antagonist MK-801 (200 μM, Sigma-Aldrich) was used at a slightly lower concentration than is common (e.g. 500 μM) because higher concentrations caused changes in the input resistance in our hands. While MK-801 was dissolved directly into the internal solution, SB415286 (Abcam) was dissolved in 0.01 % DMSO. As a control for this, we added 0.01 % DMSO to the internal solution. Importantly, neither drug interfered with the occurrence of UP states. For inhibition of potentially presynaptic NR2B containing NMDA receptors, Ro 25-6981 (0.5 μM, Sigma-Aldrich), was directly added to the bath.

2.2.5 Data analysis and statistics for patch clamp experiments

Data were analysed using semi-automated, custom written procedures in Igor Pro (Wavemetrics; E.Mann, J.Bartram, M.Kahn). To analyse the slope of the EPSP for an experiment, two points on the slope of the EPSP were set manually. The slope was then calculated for each evoked EPSP, grouped into 1-minute bins and normalised to mean of the baseline period. The peak response was also measured and generally reflected the slope. To statistically quantify the extent of plasticity, the mean normalised EPSP slope

in the last 10 minutes of the recording was used. Statistical testing was performed using SPSS (IBM).

2.3 *In vivo* anaesthesia experiments (chapter 4)

2.3.1 Surgery (A.Langfelder, P.Antonoudiou)

For acute experiments in the barrel cortex under anaesthesia in vivo, adult mice (6 - 8 weeks) were deeply anaesthetised with isoflurane (2 %) and placed on a stereotaxic apparatus and connected to a heating pad and rectal probe to maintain core temperature at 37 ^{o}C throughout the experiment. A local anaesthetic (marcaine) was applied under the scalp before the skull was exposed and a square craniotomy (2.5 x 2.5 mm) was drilled with the approximate coordinates of the D1 barrel (1.5 mm posterior and 3 mm lateral to bregma) at its centre. Once the craniotomy was complete a one-time bolus (1 mg/kg) of medetomidine was injected intraperitoneally (IP), and tubing was inserted into the peritoneal space and connected to a syringe pump (World Precision Instruments) to supply a constant dose of medetomidine (2 mg/kg/h) for the rest of the experiment. The pump was started 5 minutes after the bolus injection at which point the isoflurane level was lowered to 0.5 %.

2.3.2 Recordings, whisker stimulation and experimental design (A.Langfelder)

To obtain extracellular recordings in layer 2/3, an 8x8 extracellular electrode array (200 µM distance between electrodes; Microprobes) was inserted at a depth of 0.2 mm into the barrel cortex. The signal was grounded through a wire inserted into the scalp. The signal was acquired at 30 kHz using the Intan RHD 2000 interface software, an Intan RHD2164 64-channel amplifier board and a RHD2000 amplifier (Intan Technologies).

2.3.2 Recordings, whisker stimulation and experimental design (A.Langfelder)

Whisker deflection was achieved by inserting individual whiskers (D1 or B2) into a glass capillary attached to a piezoelectric wafer controlled by a piezo amplifier (Physikalische Instrumente). To deflect a whisker, a sine wave (12V, 40 ms) was sent to the amplifier through an Arduino board. The Arduino board was used to generate stimulation patterns (i.e. baseline vs pairings). To achieve closed-loop stimulation during UP or DOWN states, one channel with clear slow-waves in the LFP was selected in the recording software and a manual UP state threshold was set for this channel. During UP state pairings, each threshold crossing triggered five whisker deflections at 10 Hz. During DOWN state pairings, the identical burst of whisker deflection was triggered when the threshold was not crossed for 200 ms.

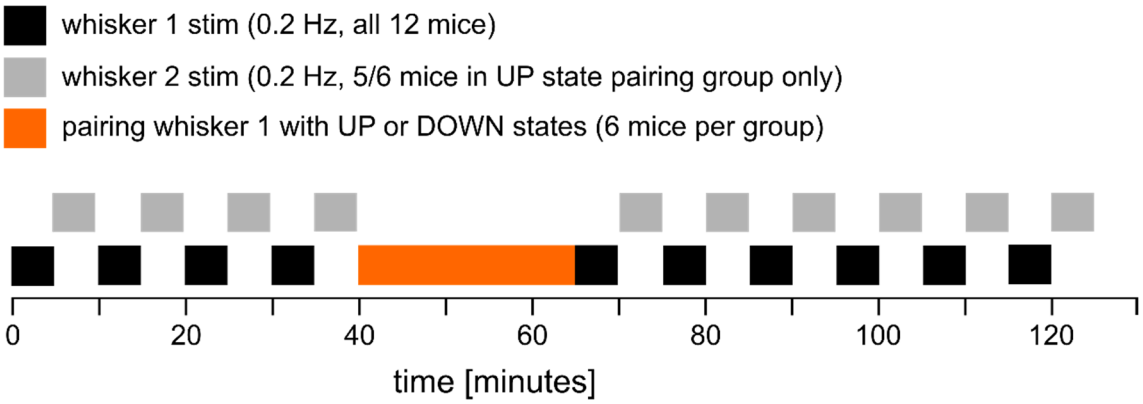


Figure 1 Experimental design
Boxes represent 5-minute blocks of whisker deflection at 0.2 Hz and concurrent LFP/MUA recordings.

Recordings were structured as follows (Figure 1): First, a 5-minute period of sham pairings were recorded for each mouse to allow assessment of UP state detection. Subsequently, baseline whisker deflections occurred in blocks of 5 minutes, during which a single whisker was deflected once every 5 s. After each block, there followed a 5-minute block where nothing was recorded (n=7) or where a second whisker was deflected (n=5). In total, the baseline recording for each single whisker consisted of four 5-minute blocks of recordings spaced over 40 minutes. During pairings, the previously

described rules for UP or DOWN state stimulation were applied to a single whisker. Following the pairing, six blocks of 5 minutes were recorded, thus the experiment lasted for another 60 minutes after the pairing. As can be seen in Figure 1, this experimental design introduced the unbalanced confounding factor of a second whisker deflection. This is discussed in the results and discussion section of chapter 4.

2.3.3 Analysis of LFP and MUA (MCK)

For analysing the LFP, the raw signal was sampled down from 30 kHz to 500 Hz using the "resample" function in Matlab. This function automatically applies an anti-aliasing lowpass filter (FIR). Before further analysis, each trial and channel (i.e. one whisker stimulation ± 4s) was manually inspected for noise to minimize the chances that measurements of slow wave activity or evoked responses were affected by obvious intrusions of noise. Consistently noisy channels and trials in which transient noise occurred were excluded from analysis. When a trial was excluded, this was done for all channels even though noise was not always global. This procedure did not affect the overall results presented in chapter 4. Moreover, I could not find any correlation between the amount of 50 Hz noise and any of the findings presented in this chapter (data not shown).

For analysing whisker evoked responses in the LFP, the DC offset was removed for each trial (i.e. one whisker stimulation ± 4s) by subtracting the voltage value immediately before each stimulation. Subsequently, all trials in each 5-minute recording block were averaged and the peak was defined as the minimum value within 15 ms of stimulation. Next, all channels were excluded that had a mean peak amplitude below 50 µV (across all recording sweeps). Finally, for analysis looking at changes in the evoked response over time across mice, the mean across all channels was calculated for each

mouse (i.e. resulting in one data point per mouse and 5-min time bin). Of note, the influence of virtually every single analysis decision described here was tested against multiple other choices and none would change the main results described in the results chapter. Examples include other methods for DC correction, selection of responsive channels (e.g. based on t-tests), averaging across channels before or after normalising to baseline, using mean vs median averaging at various steps.

For spike sorting, recording blocks were concatenated and reformatted into binary files using custom written Matlab code. Spike sorting was completed using the Kilosort software (Pachitariu et al. 2016), and the resulting clusters were manually curated using the Phy software (https://github.com/cortex-lab/phy). Kilosort was usually run using the default parameters (changing parameters did not usually result in strikingly different numbers of spikes or clusters following manual curation). In Phy, clusters were mainly assessed based on the shape of their average waveform and their autocorrelogram. Due to the 200 µM spacing of electrodes in the array, the same cell practically never appeared on more than one channel, making it difficult or even impossible to isolate single units (Chung et al. 2017). Thus, while a few clusters per recording appeared well isolated, all clusters detected on one channel were merged into one multi-unit cluster. Hence, the key advantage of spike sorting in this case was to remove artefactual threshold crossings.

To analyse multi-unit activity, spike timing information from Kilosort was imported into Matlab and all putative single and multi-unit cluster on one channel were merged and binned into 1 ms bins (i.e. one spike time histogram per channel). Similar to the LFP, evoked MUA in each channel was averaged into 5-minute bins. A channel was considered responsive, if the number of spikes within 50 ms after stimulation was in the

2.3.4 Analysis of SWA and detection of individual slow waves (MCK)

95th percentile of a Poisson distribution based on the number of spikes 50 ms before stimulation. Channels were only included if the average evoked response fulfilled this criterion in every 5-min time bin. Note that the latter criterion led to a small increase in exclusion of channels that lost their responsiveness over time, which could be due to plasticity or loss of recording quality. This was done to err on the side of biasing the analysis against our hypothesis. However, like in the LFP, the specific criterion for including channels made no difference to the main conclusions.

To quantify the evoked response in the MUA, the sum of spikes in the 50 ms before stimulation was subtracted from the 50 ms after stimulation. This measure has been called response magnitude (Armstrong-James et al. 1992) and can help to isolate changes in evoked spiking from changes in baseline firing rate.

2.3.4 Analysis of SWA and detection of individual slow waves (MCK)

To analyse slow wave activity, a 3 s period starting 1.9 s after each stimulation was multiplied with a Hanning window to limit edge artefacts and transformed into the frequency domain using the fast Fourier transform algorithm (FFT) in Matlab. Note that the exact size of this window did not seem to strongly affect the result (2 – 4 s in duration), which is consistent with other results suggesting that the stimulation does not strongly bias the measure of SWA in our case (see results). Slow wave activity was then defined as the average power in the frequency bands between 0.5 and 4 Hz. Like for the LFP and MUA, the time course of SWA across all mice was calculated by calculating the mean SWA in 5 min bins across all channels for each animal.

Correlations between the change in SWA and the change in evoked responses were computed by first calculating the relative change in both measures for each responsive

channel (i.e. % change = mean $_{\text{post-pairing}}$/mean $_{\text{pre-pairing}}$). Subsequently, all channels across all experiments were pooled and outliers were removed. They were defined as measures beyond 1.5 inter-quartile ranges from the median. Note that including outliers did not change the quality of result.

Individual slow waves were detected in a manner similar to previous studies (Vyazovskiy 2007, Riedner 2007). The LFP was bandpass filtered between 0.5 and 4 Hz using a Type II Chebyshev filter (3rd order, stopband 0.1 – 8 Hz, stopband attenuation 20 dB, passband ripple 3 dB) in each stimulation trial and channel separately. A slow wave was then defined as a pair of zero crossings (i.e. downwards followed by upwards) which was separated by at least 100ms.

2.4 Freely behaving animals (chapter 5)

2.4.1 Implants and surgical procedure

All implants were prepared manually the day before the surgery. For the frontal and occipital EEG recordings, silver wires were wrapped around blunted skull screws and soldered to a 90-Degree connector (Pinnacle Technology Inc. Lawrence). For the EMG, the end of a silver wire was bent into a U-shape and then twisted, to avoid sharp edges. This was done on two separate wires that were soldered to the above-described EEG head stage. The laminar probe (Neuronexus) has a ground and a reference wire, which was each soldered to a male connector pin, which could then be connected during surgery. The laminar probe was stained with the dye Dil (DiIC18(3), Invitrogen) before the surgery to aid the localisation of the electrode tract.

2.4.1 Implants and surgical procedure

Adult animals aged 10-20 weeks were used for all experiments. To induce anaesthesia, the mouse was exposed to a prefilled chamber with 4% isoflurane in medical oxygen and, once the breathing rate approached 80 min^{-1}, the mouse was transferred to a heating pad and the isoflurane dose was lowered to approximately 2%. After the scalp was shaved and sterilised using iodine and ethanol, the animal was transferred to a stereotaxic frame. At this point, Metacam (1-2 mg kg^{-1}), Vetergesic (0.08 mg kg^{-1}) and Dexamethasone (0.2 mg kg^{-1}) were injected subcutaneously and artificial tears were applied. Once the head was fixed, a rectal probe was inserted to maintain core temperature around 37 ^{o}C. From this point forward, all procedures were conducted under aseptic conditions, which often required assistance from a second person for parts of the procedure.

The scalp was opened, and the straightness of the skull was verified by levelling bregma and lambda and the points 1 mm lateral to Bregma. To minimise the loss of implants, the skull's surface was roughened using the scalpel and edging gel and then the coordinates for implantation were marked as shown in (Figure 2). Usually, the holes for the reference (cerebellum), ground (cerebellum or left occipital) and the two EEG screws (frontal and occipital) were drilled first and the screws were then immediately inserted using a screwdriver. Subsequently, the hole for the stimulation electrode was drilled and the electrode was carefully and slowly inserted. All screws were then fixed with dental cement before a craniotomy was made for the laminar electrode. Once the bone was removed, the dura was carefully rolled back with a syringe tip and the laminar

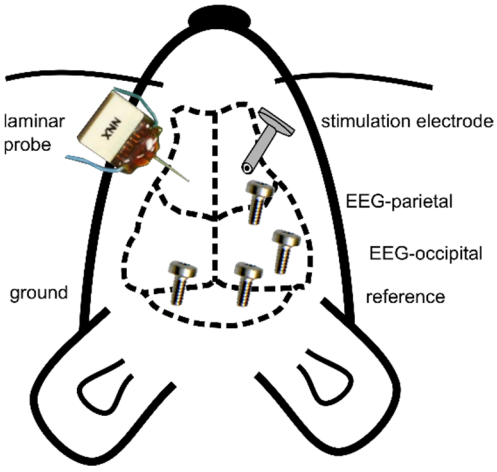


Figure 2 Schematic of implants. EMG wires in the neck muscles are not shown

probe was immediately inserted until the last of its 16 contacts was below the cortical surface. The craniotomy was immediately sealed with a silicone gel (KwikSil; World Precision Instruments, FL, USA). The entire head stage was then cemented and the EMG wires inserted into the neck before the skin was sutured, if necessary. Animals were given subcutaneous saline injections following the surgery. Following the surgery animals were carefully monitored at least once a day for seven days and analgesics were administered orally or subcutaneously, if necessary.

2.4.2 Electrophysiological recordings

Animals were moved to recording chambers at least 3 days before the start of any recording. At least 1 day into this habituation phase was allowed before the EEG head stage was connected to a cable bridging the animal and the pre-amplifier and another day before the laminar probe was connected to the pre-amplifier. An example of sleep recordings during data acquisition is shown in (Figure 3).

EEG and EMG signals were routed via an S-box (Tucker Davies Technologies, (TDT)) to a PZ-5 pre-amplifier (TDT), where they were differentially digitised (relative to the cerebellar screw or the contralateral EMG wire, respectively) at 25 kHz. The signal was then sent to a RZ-2 signal processing system (TDT), which interfaced with the Synapse recording software (TDT). The RZ-2 sampled the signal down to 700 Hz (applying an adequate low-pass filter at 45 % of the final sampling frequency) and stored it at 305 Hz. Whenever possible, the signal was stored in this “raw” form in addition to versions with high pass filters more suitable for online monitoring (i.e. 0.5 and 10 Hz for EEG and EMG, respectively).

2.4.2 Electrophysiological recordings

The signal from the laminar probe was routed directly to the PZ-5 and sampled at 50 kHz. To obtain continuous LFP data (and limit data size) one version of the signal was down sampled to 305 Hz identical to the EEG signal. For stimulation-evoked activity, a snippet of the LFP was stored at 3.5 kHz whenever the recording software triggered a stimulation. Specifically, the snippet started 500 ms before the stimulation and was 1.5 s long. An important consequence of this procedure is that there may be a small (< 1ms) delay between the time when the software sends out the trigger and when the current is applied by the stimulation box.

To record spiking activity, the laminar signal (at 25 kHz) was filtered between 500 and 3000 Hz and a manual threshold was set for each channel individually. Whenever the signal on a given channel crossed the threshold, the timestamp of threshold crossing and a 1.56 ms long snippet of the signal was stored at 12 kHz. This procedure has the advantage that it strongly reduces the considerable data load of recording 16 channels for days at 25 kHz. On the other hand, it irrevocably discards data. In other words, some spikes are too small to trigger the threshold while some noisy events do trigger it and thus create a 1.56 ms long refractory period where spikes will be lost.

2.4.2 Electrophysiological recordings

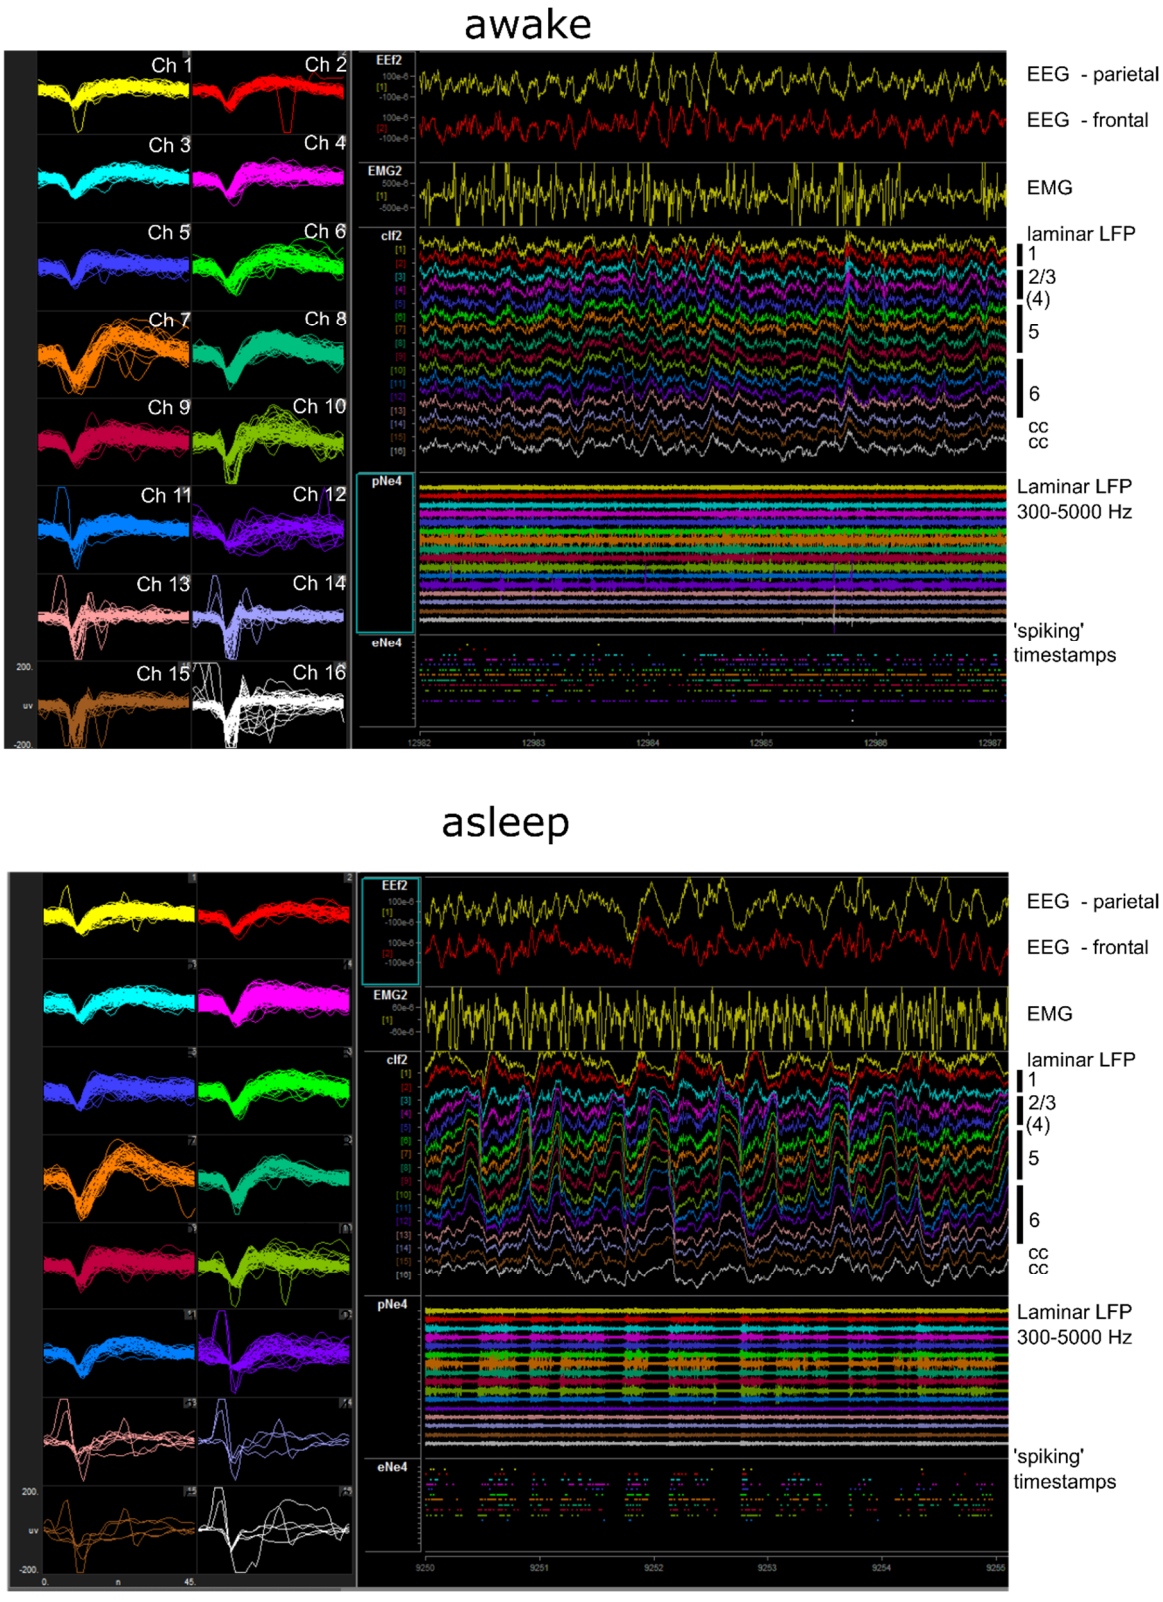


Figure 3 Example recordings of freely moving animals
Left columns: Snippets of spike-waveforms recorded by thresholding the bandpass filtered LFP (300-5000 Hz) for each channel
Right column, from top to bottom: electroencephalogram (EEG), electromyogram (EMG), LFP filtered between 0.5 and 100 Hz, LFP filtered between 300-5000 Hz, timestamps of threshold crossings of the high pass-filtered LFP.

2.4.3 Experimental Design

Every mouse was subjected to at least 4 basic experimental conditions on separate days. As the project evolved, some mice were subjected to further (up to 3) control conditions. Each condition began with eighty stimulations at 0.2 or 0.1 Hz approximately at ZT 1 (one hour after lights-on). During this "baseline waking" period, mice were kept awake by providing novel objects. Following this baseline stimulation, animals were exposed to the different experimental conditions (described below) for approximately 2.5 h. I will henceforth refer to this as the "pairing" period, because electrical stimulation was typically paired with a specific state (even though no stimulation may occur in some cases). After this pairing period, a post-pairing wake stimulation followed in all conditions. This post-pairing wake stimulation was always identical to the pre-pairing wake stimulation on all days for the same mouse (very subtle differences in baseline stimulation parameters occur in a few mice, but all variance is between mice, never within mice).

As shown in (Figure 23, page 116), the four basic conditions were (1) sleep-mock: stimulation was targeted alternatingly at ON/OFF states but the stimulation box was turned off. (2) Sleep-ON and (3) sleep OFF where stimulation was targeted selectively at ON and OFF states, respectively. (4) wake-stim: the same number of stimuli were delivered as during (2) and (3) but the animal was kept awake with novel objects. The inter-stimulus interval was similarly constrained as during (2) and (3) but was adjusted such that the same number of stimuli were delivered in approximately the same amount of time. In addition to the four basic conditions, most mice were subjected to an LTD (5) and LTP (6) protocol. The LTP "pairing period" consisted of 60 burst stimulations (7 stimuli at 300 Hz per burst), delivered every 10s. The LTD pairing period consisted

of 900 stimulations at 1 Hz. In both conditions the animal was kept awake with novel objects for the approximately 10 minutes of pairing. After delivery of the standard post-pairing wake stimulation, the animal was kept awake for an additional 50 minutes before an additional post-pairing waking stimulation was delivered (to assess stability of long-term changes). The final condition several (3/7) animals were exposed to was (7) wake-mock, which was identical to (4), wake-stim but the stimulation box was turned off. This condition is designed to control for the possibility that waking causes an increase in the evoked response independent of stimulation and that stimulation causes a decrease in the evoked response independent of behavioural state. Due to the low n-number in this condition, it is not included in the current analysis.

The number of stimulations delivered during these pairing protocols were determined by the first experimental day in each animal. The animal was allowed to sleep for up to 2.5 hours and the only constraint on stimulation numbers were this the minimum inter stimulus interval (10 s), the amount of NREM sleep and the number of ON/OFF detections. In all subsequent experimental days, the same number of stimulations were delivered (except for mock stimulation days). Therefore, the total duration of the experiment was allowed to vary slightly (i.e. ~ 20 mins) but was kept constant whenever possible.

To avoid a systematic effect of repeated stimulation, these conditions were randomly counterbalanced with two limitations: Firstly, wake-stim (4) was never done as the first condition, because the number of stimuli delivered during the pairing period was constrained most strongly by the ON/OFF detection algorithm. Secondly, LTP and LTD protocols were mostly done last to avoid potential long-lasting changes to the network.

2.4.4 Closed loop electrical stimulation

Electrical stimulation was delivered to the animal through a bipolar concentric stimulation electrode (Plastics One) attached to a PSI6x stimulus isolation unit. The stimulus isolation unit was coupled to a stimulation box (Grass Instruments), on which the stimulation parameters could be set manually. Once the parameters were set, the stimulation box could be triggered by means of a TTL pulse from the RZ-2 system (TDT), which was controlled by the recording software. One day before experiments started, an input-output curve was obtained.

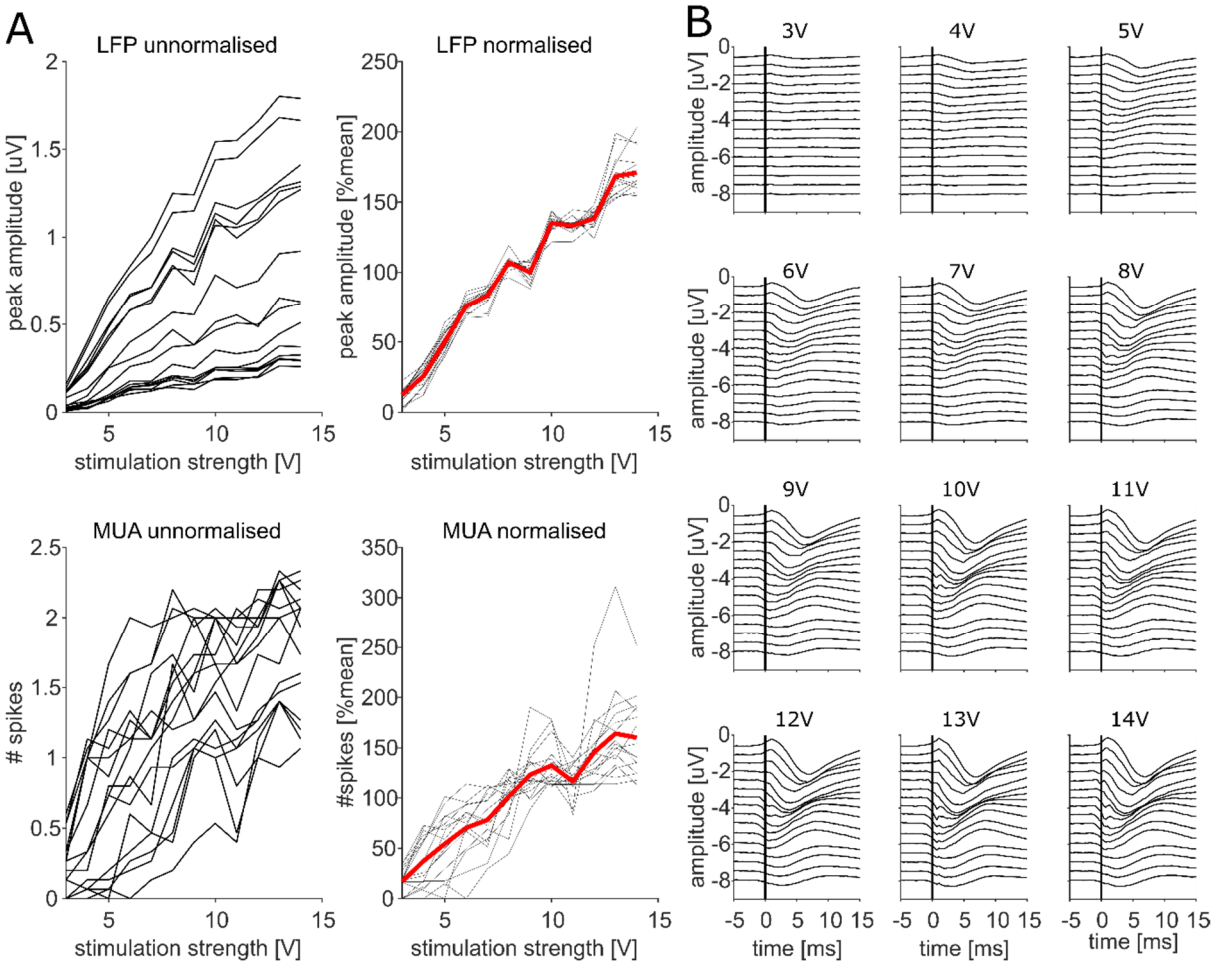


Figure 4 Example input-output curve showing all channels of one mouse
Left columns show unnormalized measures of peak amplitude an multi-unit activity (without spike-sorting)
Right columns show the same measures normalised to the mean of each channel. Grey traces show individual channels and red trace show mean across channels

2.4.5 Histology and identification of recording sites (MCK + Lukas Krone)

Procedures for online data processing and closed-loop stimulation were custom written in the proprietary object-oriented programming environment supplied by TDT. In brief, the incoming spikes were summed across all channels over a predefined time window (usually 50-150 ms). Whenever this running sum went below a predefined threshold (usually 0-1 spike), a potential DOWN state was registered. To avoid stimulating during waking a running root mean squared of the EMG signal was used and a manual threshold was set for it. The reasoning for parameter selection and the success of this procedures are detailed in the corresponding results section (page 105).

2.4.5 Histology and identification of recording sites (MCK + Lukas Krone)

After experiments were completed, animals were deeply anaesthetised with an intraperitoneal injection of Pentobarbital (Euthanal). Once the animal reached deep anaesthesia (as verified by loss of righting, pedal and corneal reflexes) microlesions were performed to aid laminar identification of recording sites. For microlesions, the laminar probe was connected to an impedance testing device (NanoZ, Plexon Inc), which was used to pass current (10 µA for 10 s) through four equally spaced channels of the laminar probe. The bottom channel was always lesioned first, as the first lesion can damage the other channels. Animals were then transcardially perfused with PBS and 4% paraformaldehyde and the head of the animal was then stored in PFA (i.e. the implant was not removed at this point, which improved the quality of histology) and moved into PBS-A (acidified PBS) after a few days. Brains were embedded in agarose and cut into 50 µM thick coronal sections. The sections were stained with DAPI and imaged using a fluorescence microscope. The sections containing the electrode tract were identified using the red Dil-fluoresence and were imaged at 1.6, 2.5 and 5x magnification.

2.4.6 Scoring of vigilance states

The recording location in the rostrocaudal and mediolateral dimension was identified using the (Watson & Paxinos 2012) mouse brain atlas. The cortical layer of each laminar contact was identified in the 5x-magnification images. First, the site(s) of the microlesions were identified in the DAPI, or background fluorescence (GFP filter) images. Second, the position of lesions, the Dil staining, and the length of the electrode were used to determine the position of each contact. Layer 1 was identified based on the low density of neurons compared to Layer 2/3. Similarly, the beginning of layer 5 was identified based on the lower cell density in Layer 5. While layer 4 is comparatively small in primary motor cortex, it exists and can be identified as a small increase in cell density right above Layer 5 (Skoglund et al. 1997; Yamawaki et al. 2014). Layer 6 was also identified based on the higher density of cells compared to Layer 5.

This protocol was brought to our lab by Dr Lukas Krone and I am grateful for his help with this procedure.

2.4.6 Scoring of vigilance states

Data were extracted from the raw data format of the recording software, resampled to 256 Hz using and bandpass filtered using custom Matlab scripts (V. Vyazovskiy, 0.5-100Hz for EEG/LFP and 10-50 Hz for EMG, 3rd order phase conserving type II Chebyshev filter). The signals were then converted to the ASCII format and from there converted into edf-files using the open source software Neurotraces. The edf-files were visualised in the software SleepSign (Kissei Comtec Co, Nagano, Japan). To score vigilance states, the LFP, EEG and EMG data were examined in 4-s epochs. If present, timing of electrical stimulation was also visualised. Waking was defined as a low voltage, high frequency EEG with a high variance in the EMG. In contrast, NREM sleep was defined as high amplitude EEG signals containing slow waves (and high delta

power) and exhibiting low EMG tone and variance. The EMG commonly displayed clear heartbeat artefacts during all sleep episodes. REM sleep was defined as wake-like activity with sleep-like EMG signal and usually high theta activity in the occipital derivation. When an animal displayed wake-like activity for less than 4 consecutive episodes (i.e. 16 s) within a NREM bout, this was scored as brief-awakening. Episodes of all four types (NREM, REM, waking, brief awakening) were flagged if they contained clear artefacts in any EEG or LFP channel. When sleep scoring was complete, the SleepSign software returned the vigilance states and the power spectra for each 4 s episode, the latter were calculated in 0.25 Hz frequency bins using a Hanning window.

Special consideration was given to 4 s episodes containing stimulation events. For stimulations aimed at waking periods, the episode was only scored as NREM if there was sleep like activity in the 2 s before or after stimulation. Vice versa, if stimulation was aimed at NREM episodes, activity was scored as REM or waking if the activity 2 s before or after the stimulation resembled the respective state. The same "over-sensitive" procedure was applied with regards to artefacts. In summary, episodes with stimulation events were treated such that false misses of erroneous stimulations were minimised.

2.4.7 Pre-processing of local field potentials and multiunit activity

The 1.5 s long LFP snippets (sampled at 3-6 kHz) surrounding each stimulation were imported into Matlab using the supplier's (TDT) Matlab software developing kit. Pre-processing of the snippets involved removing line noise and slow drift using a regression-based algorithm (http://chronux.org/, (Mitra & Bokil 2009)). Regression-based approaches were chosen to avoid introducing filtering artefacts. For analysis of the peak and slope of the evoked response, potential DC-offsets were accounted for by

subtracting the mean of the 5 ms preceding the stimulation from the entire snippet for each channel and trial separately.

Spiking activity was always processed with WaveClus 3 (Chaure et al. 2018). This software was chosen because it was designed explicitly to perform well on single channel recordings as well as multi-channel recordings. In contrast, most other spike-sorting algorithms are optimised for polytrode recordings (Pachitariu et al. 2016; Chung et al. 2017), where a single unit is recorded on >1 channel. Such cases are likely rare in the present recordings, given the relatively large distance between channels (100 µM). The suitability of WaveClus for such single channel recordings is further supported by a direct comparison based on simulated single channel recordings (Wild et al. 2012). Of note, WaveClus is designed to work on raw, continuous voltage signals, to which the program applies a thresholding step. In the present analysis, thresholding was already applied during the recording and only putative spike times and waveforms were recorded. Thus, the thresholding step in Waveclus was omitted. Following clustering by WaveClus, clusters identified as noise were excluded. In spite of this, inspection of stimulation-evoked waveforms revealed significant occurrences of clearly noise-like waveforms following stimulation. As exemplified in (Figure 5), there were several

characteristics that coincided with noise-like spike waveforms: (1) complete lack of baseline firing (2) complete lack of variability in spike timing following stimulation (3) spike half-widths longer than 0.5 ms.

For both LFP and MUA analysis, only responses that occurred during epochs scored as artefact-free waking (see page 58) were included in the analysis. For an average experiment (i.e. one mouse and stimulation type), out of 80 stimulations before the "pairing", 67.4 (SD 12.7) were included in the analysis. Similarly, 64.8 trials (SD 14.1) were included from the 80 stimulations following the pairing. To explore whether trial-exclusion may be biased towards a specific condition, pre/post pairing periods or mice, a linear fixed-effects model was fit predicting the number of excluded trials based on subject-ID, condition, and pre-post pairing period. This model showed clearly that some

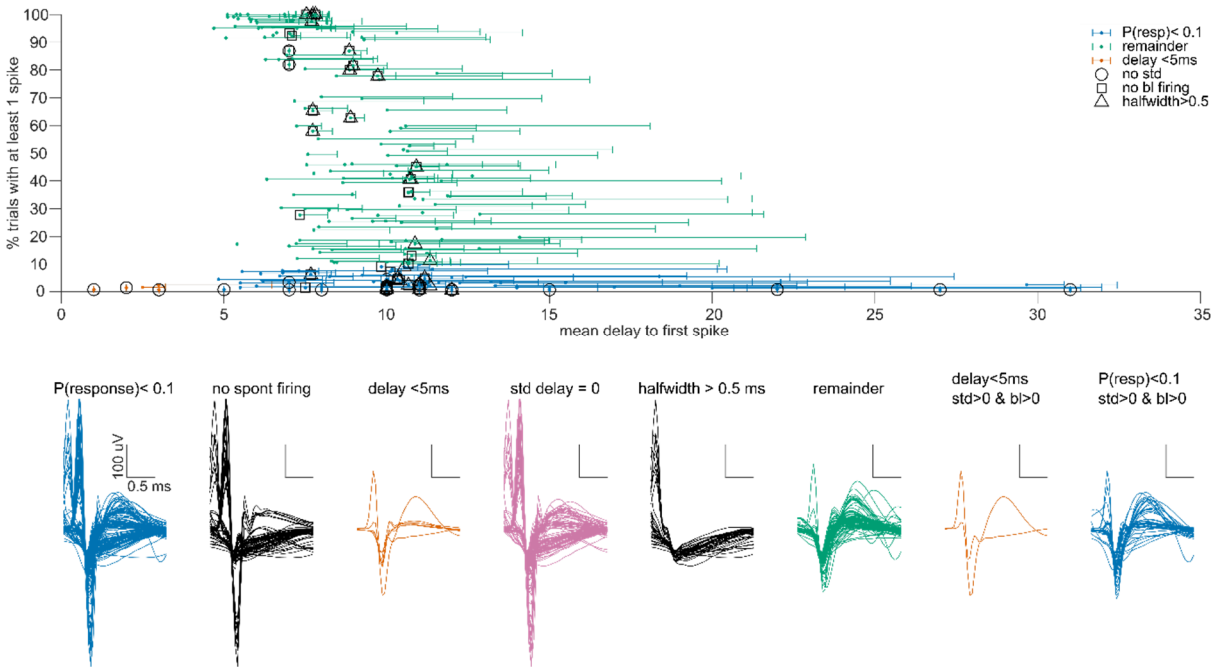


Figure 5 Characteristics of putative spike waveforms and artefacts isolated using the WaveClus algorithm
Top panel: Stimulation-locked characteristics of all clusters isolated from one mouse across 7 experimental days. P(resp)<0.1: clusters responding to stimulation with at least one spike only in 10% of all trials. Remainder: clusters without any of the other flagged characteristics. Delay < 5 ms : clusters with an average delay to first spike < 5ms. No std: Clusters which always respond at the exact same time after stimulation. No bl firing: clusters without activity in any trial in the 500 ms pre-stimulation period . halfwidth > 0.5: spike-halfwidth > 0.5 ms.
Bottom panel: Average waveforms of clusters shown in top panel. All scales are identical to the leftmost panel.

mice had more exclusions than others but suggested no effect of condition or pre/post pairing period. (model R^2: 0.694, ANOVA against the null hypothesis that regression weight=0; F=28.15, p<0.0001 for mouse-ID, F=2.49, p=0.11 for before/after pairing, for all other terms p>0.2). The same conclusion was reached from performing two separate ANOVAs testing the effect of condition and mouse-ID on the number of excluded trials.

2.4.8 Quantification of evoked responses and plasticity thereof

Quantification of responses was performed subsequent to all above-described pre-processing steps.

To quantify LFP responses, the peak was defined as the minimum value in the period 1 – 20 ms following stimulation. This period did not involve the stimulation artefact. Unless mentioned otherwise, this peak measurement was always performed on the averaged waveforms. To identify channels with a significant response, the peak was measured across all trials on a given experimental day and compared to a normal distribution based on the mean and standard deviation of the baseline period. A channel was considered responsive if p< (0.05 / 16); 16 being the number of channels.

To quantify the MUA, several approached were considered. One common approach is to sum the number of spikes following stimulation across a given period. This approach suffers considerably in the present case because the responses are often a mixture of excitation and inhibition (such as brief oscillatory responses) and because the width of the response varies (i.e. some channels show excitation from 5-7 ms following stimulation while other respond at 10-15 ms). The latter problem means summing across a predefined period would be quite insensitive to some responses. Consistent

with these concerns, defining responsive channels based on the sum of spikes over a predefined period had a very high false negative rate based on visual inspection. Therefore, we divided spike counts into 1 ms bins and aimed to establish which of the bins significantly responded to stimulation. To this end, a Poisson regression model (a generalized linear model, GLM) was fitted to all trials of a given channel. The GLM was fit using a lasso regularization method via the command *lassoglm* in MATLAB (threefold cross-validation). In essence, this method tries to predict spike rates in each post-stimulation bin based on the baseline spike rate and a specific regression weight for each time bin. Hence, if the spike rate in a given time-bin can be explained simply by the baseline rate, the regression weight for this time bin should be zero. The results from this procedure appeared to accurately capture the evoked response based on visual inspection (Figure 6). In order to reduce the computational effort for this procedure, channels without any responsive time-bins were excluded beforehand. This was done by comparing the spike rate in the most responsive time bin to the baseline level via a permutation test (5000 permutations).

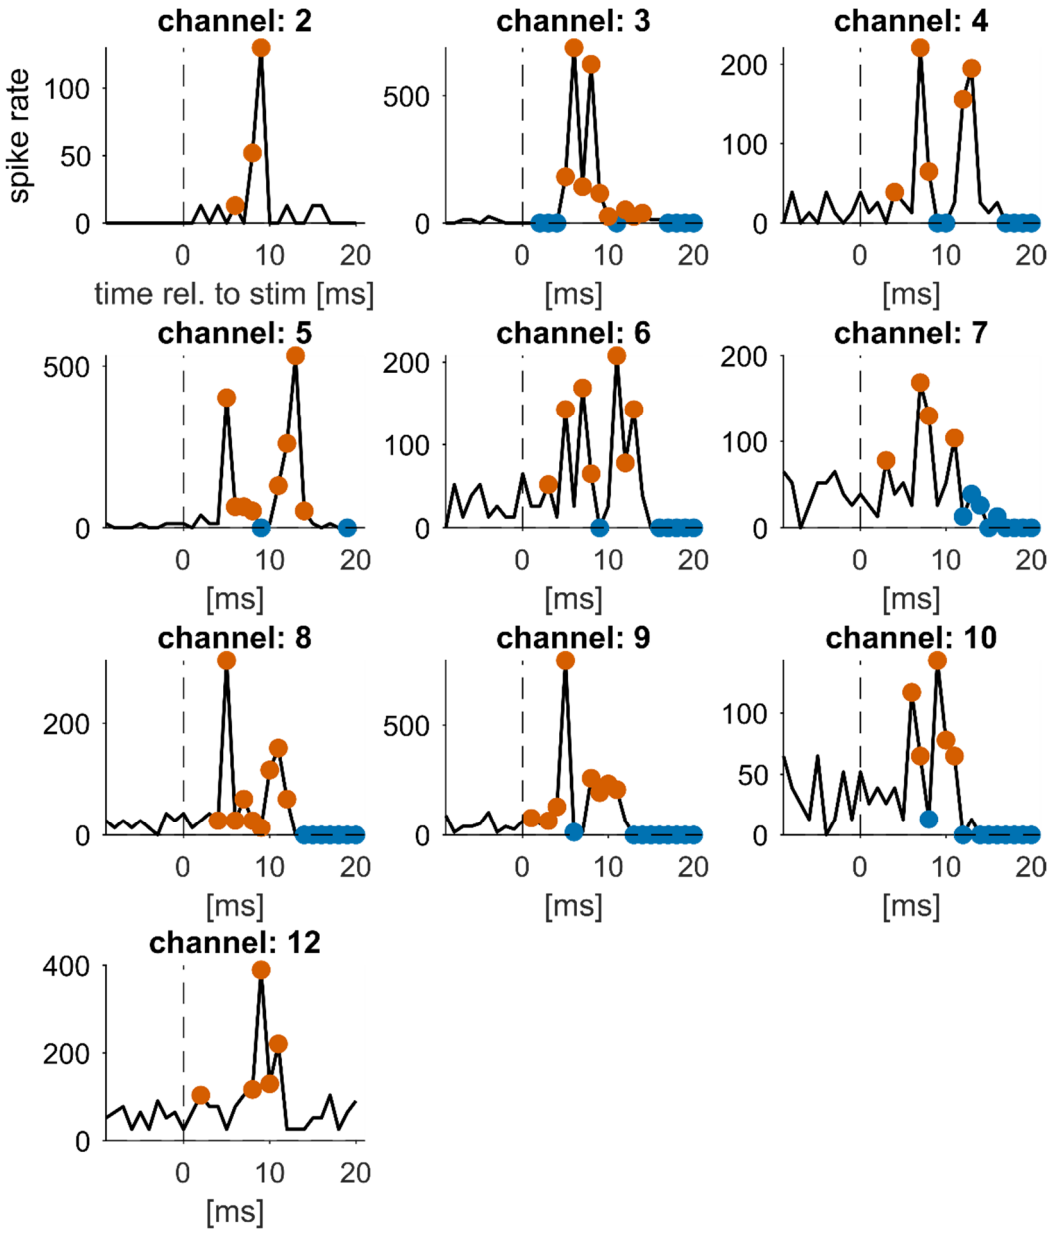


Figure 6 Stimulation-evoked spike rates and GLM-based responsive bins for responsive channels of one mouse
Black traces are mean spike rates across all included trials and dashed line indicates timing of stimulation. Orange and blue circles indicate time-bins identified as time bins positively and negatively modulated by stimulation, respectively. The GLM was only run on channels identified as responsive, based on a permutation test on the time-bin with the maximal response.

2.4.9 Statistics

The experimental design posed several statistical challenges. Most notably, each mouse experienced several treatments and observations were often nested (sometimes also referred to as hierarchical). For instance, each mouse experienced several treatments

and we recorded spiking activity from 16 channels. Hence, the channels are nested within mice. Furthermore, in many cases there were missing observations (e.g. for a few mice it was not possible to apply all treatments, and some channels were not always responsive to stimulation).

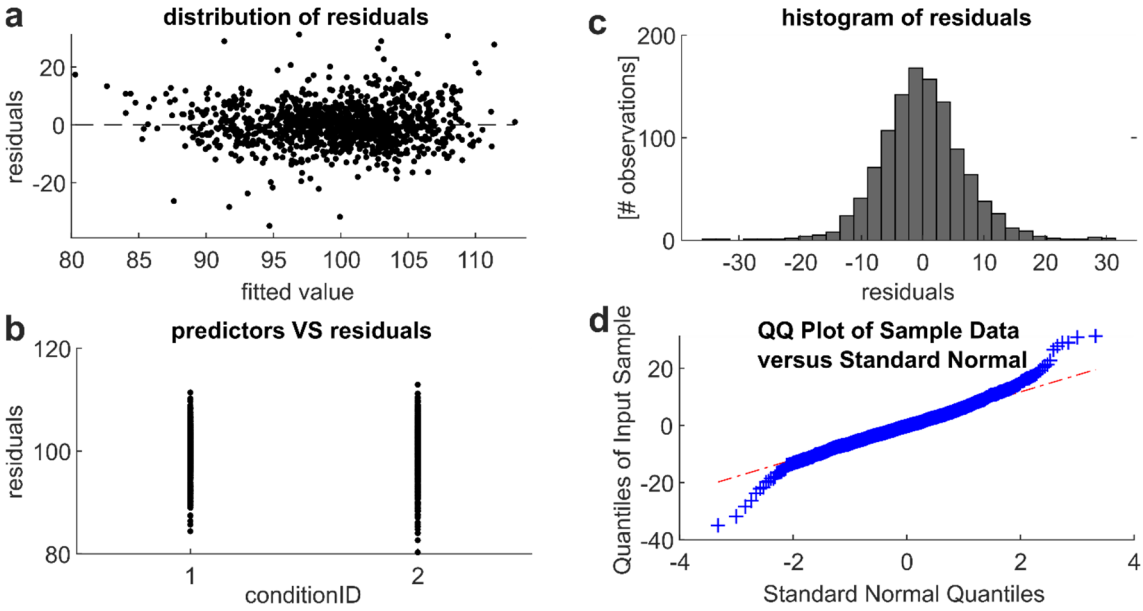


Figure 7 Testing the assumptions of linear mixed effects models
A) Homogeneity of variance requires the variance of the fitting error (residuals) to be independent of the value of the fitted value. This plot also shows deviations from the assumption of linearity.
B) The equality of variance is also assessed separately for each predictor variable
C) Normality of error requires the residuals to be normally distributed. A histogram allows visual examination of this property.
D) A quantile-quantile plot of the residuals offers another way to examine normality. If the residuals are perfectly normally distributed, all points would lie on the red line.

To address these challenges, we used linear mixed models (LME) (Harrison et al. 2018). This method has several advantages, most notably it can account for the above-mentioned nested nature of experiments and it can readily handle missing datapoints (e.g. a noisy or unresponsive channel on one day). LMEs have 4 key assumptions: independence, homogeneity of variance, normality of error, and linearity. These assumptions were always visually inspected using plots like the one shown in (Figure 7).

2.4.9 Statistics

To test the effect of condition (ON, OFF, mock, etc.) on sleep stages (time in NREM, REM wake), repeated measures MANOVA is an appropriate statistical test. However, when the data violate the assumptions of this test, its power and suitability suffer. Therefore, whenever a repeated-measures MANOVA was theoretically warranted, the data were tested regarding its assumptions. First, boxplots were used to identify univariate outliers and the Shapiro-Wilk or Shapiro-Francia test for univariate normality were performed using Matlab. Furthermore a separate correlation matrix of dependent variables (e.g. time in NREM vs REM etc.) was computed for each condition (mock, ON, OFF etc.). When there were too many outliers or cases of collinearity (Pearson's R > 0.9), the test was not used and multiple repeated measures ANOVAs were performed instead. As a consequence, the result will be more readily interpretable at the cost of an increased chance of type-I errors.

Chapter 3: Cortical UP states induce selective weakening of subthreshold inputs ex vivo

3.1 Introduction

The cortical slow oscillation during NREM sleep has long been hypothesised to play a role in sleep's beneficial effect on cognition and particularly memory. Several recent studies have provided elegant backing for the role of slow waves in memory consolidation. For instance, enhancement (Marshall et al. 2006; Ngo et al. 2013; Miyamoto et al. 2016) and interruption (Fattinger et al. 2017) of slow waves have opposing effects on memory consolidation. Strikingly, even imposing slow waves optogenetically on a waking brain led to significant improvements of memory performance (Miyamoto et al. 2016). While it is thus likely that slow waves are beneficial for at least some instances of memory consolidation, it is less clear by what mechanism the slow oscillation exerts this effect.

The ambiguity on the effect of slow waves on connectivity also has consequences for leading theories on the function of sleep for neuronal networks. For instance, the synaptic homeostasis hypothesis (Tononi & Cirelli 2003) proposes that a key function for sleep is to facilitate a net weakening of synapses, whereas the systems consolidation hypothesis (Rasch & Born 2013) proposes that memories are transferred from short-term to long-term storage during sleep. While not mutually exclusive, the two theories certainly pose a challenge to plasticity mechanisms during sleep.

Slow waves in the EEG and LFP are generated by synchronous, slow oscillations in the local neuronal network between hyperpolarised, low firing rate DOWN states and

depolarised, high firing rate UP states (Steriade et al. 2001). While the slow oscillation clearly bidirectionally interacts with various subcortical areas (Lewis et al. 2015), it can persist in an isolated slab of cortex (Sanchez-Vives & McCormick 2000) and likely in part originates from cortical layer 5 (Lőrincz et al. 2015; Beltramo et al. 2013; Neske 2016).

Here, we used cortical brain slices of adolescent mice to investigate the plasticity rules during cortical UP and DOWN states. We focussed on the medial entorhinal cortex because it is a crucial interface between hippocampus and neocortex and because its strong recurrent connectivity readily produces UP states. We used patch clamp recordings to monitor evoked potentials as a proxy of synaptic strength and make use of the precise control of UP states afforded by the *ex vivo* preparation. We find that pairing subthreshold synaptic inputs with cortical UP states induces synaptic weakening, whereas we do not observe any weakening of suprathreshold pairings. Synaptic weakening was dependent on postsynaptic NMDA receptors even though it appears to display signs of pre and postsynaptic expression. Finally, synaptic weakening depended on GSK-3beta, a key regulator of plasticity that is differentially regulated by sleep (Peineau et al. 2009).

3.2. Results

3.2.1 Cortical UP states facilitate weakening of sub- but not suprathreshold inputs

To investigate the impact of cortical UP states on synaptic plasticity we obtained brain slices of the medial entorhinal cortex (mEC) from adolescent (P14-21) mice. We adjusted the extracellular calcium and magnesium concentrations such that UP states

did not occur spontaneously but were reliably evoked by electrical stimulation in layer 3 ((Mann et al. 2009), Figure 8 a & b). Such evoked UP states were network events as verified by concurrent LFP recordings (Figure 8 c), they were indistinguishable from spontaneously occurring UP states and, as reported by (Digby et al. 2017) they were longer than neocortical UP states thus mirroring the properties of entorhinal UP states *in vivo* (Hahn et al. 2012).

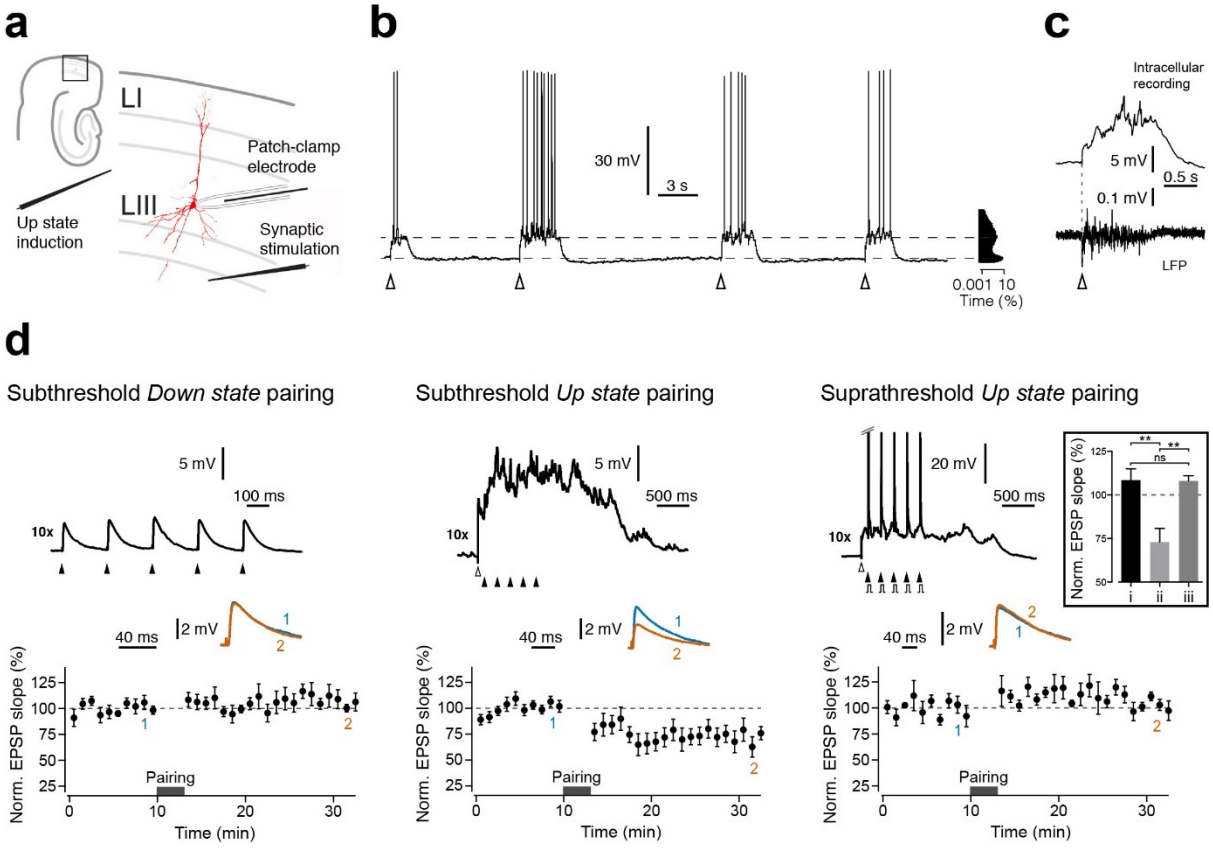


Figure 8 UP states induce weakening of subthreshold but not suprathreshold inputs
*a) Experimental setup in brain slices of mEC; (b) Examples of electrically evoked UP states in intracellular recordings. Arrows indicate UP state induction. Inset in (b) depicts bimodal membrane potential histogram. (c) Simultaneous intracellular and extracellular recordings, showing that intracellular UP states are associated with network activity. (d) Effects of UP/DOWN state pairing on synaptic strength. Upper row: representative intracellular Vm traces during a single pairing; Middle row: representative average evoked EPSP before (blue) and after (orange) pairing. Bottom row: Average time course across experiments of normalised EPSP slope. Inset in the rightmost column: average (+/- SEM) slope in last 10 minutes of recording. **: p<0.01, one-way ANOVA followed by Tukey's test.*

We obtained current clamp recordings of pyramidal cells in layer 3 of mEC and recorded excitatory postsynaptic potentials (EPSPs) triggered by electrical stimulation

in layer 5 of mEC (15 s inter-stimulus interval). As feedforward activity might obscure measurements of the EPSP peak, we quantified the slope and not the peak of the evoked EPSP. Once a recording displayed a stable baseline of the EPSP slope, the cell was subjected to one of three pairing conditions, all of which involved a total of 50 stimulations at 5 Hz (Figure 8d). These three conditions resulted in significantly different values of the EPSP slope measured in the period from 10 to 20 min post pairing ($F(2, 18) = 9.63$, $P < 0.01$, $\eta 2 = 0.52$; One-way analysis of variance (ANOVA)). First, as a control, L5 stimulations were delivered at 5 Hz every 15 s without triggering an UP state (Figure 8 d, left). This DOWN state pairing resulted in an EPSP slope of 108.4 +/- 6.5 % (n=6) compared to baseline. In contrast, when an UP state was triggered every 15 s immediately before delivering the L5 stimulation, the slope decreased significantly to 72.9 +/- 7.9 % (Figure 8 d, middle, $n = 9$; $P < 0.01$ cf Down state pairing, One-way ANOVA followed by Tukey's post hoc test; Cohen's $d = 1.95$). While the spontaneous firing of the postsynaptic cell was not controlled during these UP state pairings, the postsynaptic spike rate did not correlate with weakening of the EPSP ($R^2 = 0.08$, $P = 0.46$, $n = 9$) and weakening also occurred in cells that did not spontaneously spike during UP states at all. To test how precisely timed postsynaptic spiking interacts with cortical UP states, the third experimental condition was identical to the UP state pairings with the addition of a suprathreshold current injection into the postsynaptic cell, which triggered a spike within 15 ms of each L5 stimulation. The result of this experiment, where each EPSP during an UP state was reliably followed by a spike, was that the EPSP slope changed to 108.0 +/- 3 % ($n = 6$; n.s. cf Down state pairing, Cohen's $d = 0.02$; $P < 0.01$ cf subthreshold Up state pairing, Cohen's $d = 1.92$; One-way ANOVA followed by Tukey's post hoc test). Note that we used 5 Hz stimulation during pairings based on *in vivo* firing rates of the mEC during slow waves (Hahn

2012), however, subthreshold UP state pairings also led to significant weakening of the EPSP slope compared to baseline with pairing frequencies of 2.5 Hz (72.2 ± 5.2%, n = 6, $t(5) = 5.35$, $P < 0.01$, one-sample t-test) and 7.5 Hz (72.0 ± 6.0%, n = 7, $t(6) = 4.63$, $P < 0.01$, one-sample t-test). These experiments suggest that UP states may induce synaptic weakening of subthreshold inputs, while during suprathreshold pairings the weakening is either prevented or simply additively interacts with spike-timing dependent potentiation.

Note that a concern with this experiment is that EPSPs were triggered with electrical stimulation, which recruits excitatory, inhibitory and perhaps even neuromodulatory inputs in a relatively artificial way. Julian Bartram (Bartram/Kahn *et al.,* 2017) addressed this concern by essentially reproducing the above findings using a paired-patching configuration, where two connected excitatory cells were directly monitored (data no shown).

3.2.2 UP state pairings are ineffective following potentiation but weakening is reversible

Memory consolidation during sleep arguably involves synapses which have recently undergone plasticity. We therefore investigated how the above-described UP state induced plasticity in mEC interacts with recently potentiated synapses. Thus, we repeated the UP state pairing condition from above, and additionally injected sufficient current through the patch pipette to trigger a burst of postsynaptic action potentials following each EPSP (total of 100 pairings at 5 Hz, Figure 9, top). Postsynaptic bursting has previously been reported to facilitate potentiation (Pike et al. 1999)and we indeed consistently observed potentiation of the EPSP slope following burst pairings, which

did not seem to reach baseline levels within 50 minutes after pairing (Figure 9). To test whether subthreshold UP state pairings can affect previously potentiated synapses, one group of cells was exposed to a subthreshold UP state pairing 20 minutes after the UP state burst pairing (n=6) while the other group was left undisturbed (n=7). A mixed ANOVA analysis on the EPSP slope after the burst pairing revealed a significant effect of time ($F(1,11) = 8.57$, $P < 0.05$; $\eta^2 = 0.48$) but no interaction between time and group ($P = 0.90$; $\eta^2 = 0.001$), suggesting that the additional subthreshold UP state pairing did not modify the timecourse of plasticity. Following this analysis, we pooled both groups and confirmed that there was significant potentiation compared to baseline in the period 10-20 minutes (158.5 ± 10.5%, n = 13, $t(12) = 5.47$, $P < 0.001$) and 30 – 40 min post

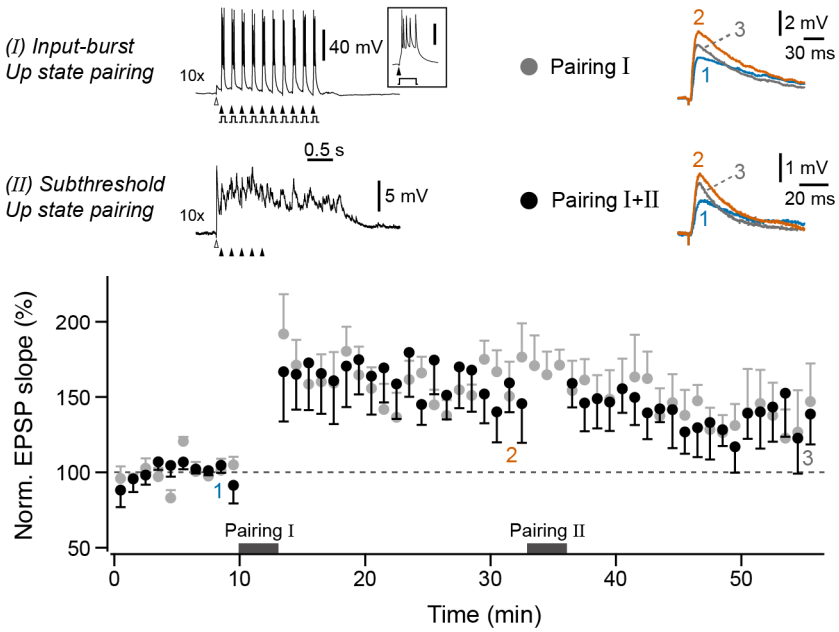


Figure 9 UP state pairings have no detectable effect on previously potentiated responses
(i) Representative intracellular trace during the first pairing, which was designed to induce synaptic potentiation in both groups. Following each stimulation, current was injected to make the cell fire a burst (inset shows single burst). (ii) Representative recording of the second pairing, which was administered to one group (N=6) but not the other (N=6), and followed 20 minutes after the first pairing. This subthreshold pairing is identical to the ones that previously yielded synaptic weakening. Coloured EPSPs are averages of responses at baseline (blue), after the burst-pairing (orange), and following the 2nd pairing (grey), or simply passage of time, respectively. Bottom panel: average responses (normalised to baseline mean) across experiments (+/- SEM).

input-burst pairing (133.5 ± 12.8%, n = 13, t(12) = 2.62, P < 0.05; one sample t-tests, with Bonferroni correction for multiple comparisons; Figure 9). Even though other interpretations are possible (see discussion), this experiment is consistent with the interpretation that UP state induced weakening cannot act on synapses that have previously been potentiated.

Though without my personal involvement, it is worth mentioning here that our group also explored what happens if the burst pairing occurs *after* the induction of synaptic weakening. In this case, the burst pairing still leads to potentiation (data not shown), suggesting that weakening cannot occlude strengthening.

3.2.3 Up state induced weakening involves a presynaptic expression mechanism but is induced postsynaptically

Synaptic plasticity can be expressed through changes in presynaptic release and/or through changes to the postsynaptic currents evoked by a constant amount of neurotransmitter. Similarly, the locus of induction can be mainly situated at the presynaptic terminal or at the postsynaptic density. A classic way of investigating the locus of expression is a coefficient of variation analysis (CV, (Faber & Korn 1991) which in the case of subthreshold UP state pairings indicated a presynaptic component of expression (data not shown). However, the CV analysis can be confounded by postsynaptic changes, such as changes in the number of silent synapses (Kerchner & Nicoll 2008). Another measure sensitive to changes in presynaptic release is the paired pulse ratio - the change in the magnitude between two closely spaced evoked responses. Therefore, we conducted an additional series of subthreshold UP state pairings experiments, in which each baseline and postpairing stimulation was followed by

another stimulation within 50 ms (Figure 10 a). We defined the paired-pulse ratio as the relationship between the slope of the first and second EPSP before and after the subthreshold UP state pairing, and found that it significantly increased from 1.35 ± 0.11 during baseline to 1.69 ± 0.19 following pairing; (n = 7, t(6) = 2.80, P < 0.05, paired t-test; Figure 10 b). Furthermore, there was a strong correlation between the magnitude of synaptic weakening and the change in the paired pulse ratio (R^2 = 0.73, n = 7, P < 0.05; Pearson linear correlation test). Together, these experiments suggest that subthreshold UP state pairings weaken synaptic transmission at least in part by means of changing presynaptic release.

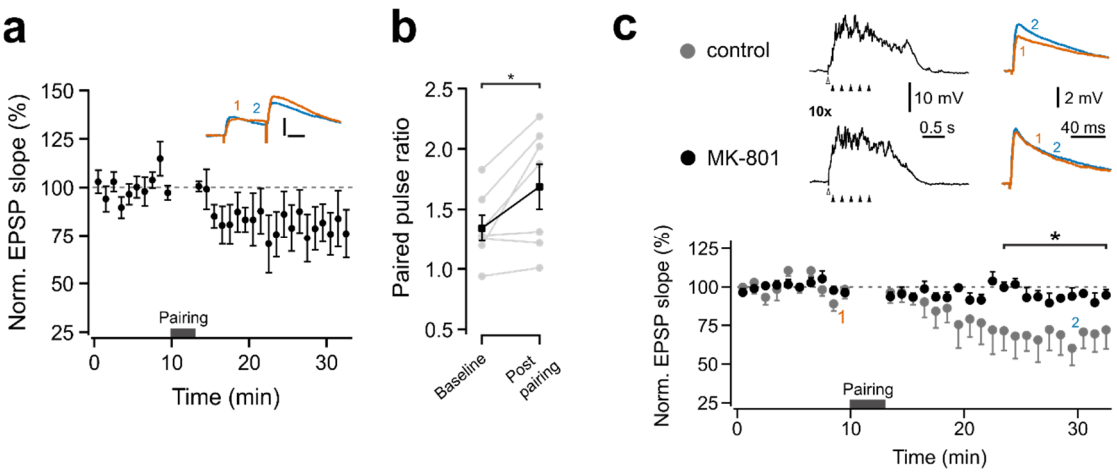


Figure 10 Expression of UP state induced synaptic weakening has a presynaptic component but depends on postsynaptic NMDA receptors. A) Subthreshold UP state pairings significantly decrease the paired pulse ratio.
*B) Quantification of change in paired-pulse ratio. C) Pharmacological blockade of postsynaptic NMDA receptors through intracellular delivery of MK-801 inhibits UP state induced weakening. All traces are mean ± SEM. *: p<0.05.*

Next, we wanted to investigate the synaptic locus for the induction of synaptic plasticity. To this end, we took advantage of the patch clamp technique, which allows to deliver a drug specifically into the recorded postssynaptic cell. We thus included either an inhibitor of NMDA receptors (MK-801, 200 µM) or just the standard internal solution in the patch pipette, allowed at least 30 minutes for diffusion into the cell before starting a baseline recording, and then repeated the subthreshold UP state

pairings. As shown in (Figure 10 c) inhibition of postsynaptic NMDA receptors led to significantly less synaptic weakening (94.42 ± 3.3% of baseline slope (n = 6)) compared to the control group (68.8 ± 11.7% (n = 5), $t(9) = 2.30$, $P < 0.05$, unpaired t-test).

While the above method should be reasonably selective for postsynaptic NMDA receptors, the open channel blocker MK-801 may in small quantities escape through postsynaptic receptors and bind presynaptic NMDA receptors. In addition, the positive pressure applied to cells during patching may lead to small doses of the inhibitor reaching the presynaptic terminal. Given that it is uncommon for postsynaptic NMDA receptors to induce presynaptically expressed plasticity, we wanted to further assess whether presynaptic NMDA receptors are implicated in the induction of synaptic weakening. To this end, we took advantage of the predominantly presynaptic location of NR2B-subunit containing NMDA receptors in the developing entorhinal cortex (Woodhall et al. 2001) and bath applied Ro 25-6981 (0.5 μM), which at this concentration is reasonably specific for NR2B-subunit containing NMDA receptors (Figure 11). As presynaptic NMDA receptors are thought to tonically modulate presynaptic release and mediate short-term plasticity, we first assessed the impact of 0.5 μM Ro 25-6981 on baseline synaptic transmission and paired pulse facilitation. As perhaps expected, addition of the NR2B antagonist to the bath reduced the slope of the EPSP (64.3±9.6% of baseline, n = 6, $t(5) = 3.72$, $P < 0.05$, one sample t-test, Figure 11 a, top). Surprisingly, we were unable to detect an effect of the drug on paired-pulse facilitation (n=6, P=0.45, paired t test, Figure 11 a, bottom), which could be due to statistical (i.e. power) or biological reasons (see discussion). When Ro 25-6981 was bath applied 30 minutes before the beginning of subthreshold UP state pairings, the EPSP slope was still significantly decreased (78.2 ± 4.9% (n = 4, $t(3) = 4.48$, $P < 0.05$; one sample t-test Figure 11 b), to an extent that was not different to interleaved control

experiments (81.8 ± 4.7%, n = 4; t(6) = 0.54, P = 0.69; unpaired t-test). While subunit specific pharmacology can be difficult to interpret, this experiment is consistent with our previous finding that UP state induced weakening is dependent on post- and not pre-synaptic NMDA receptors.

3.2.4 UP state induced weakening depends on GSK-3 beta but not caspase-3

Ultimately, it would be ideal to have a specific molecular target with which plasticity during sleep can be manipulated. Therefore, and to further understand how UP state associated plasticity functions, we tested the involvement of two interesting candidate molecules. Firstly, we tested GSK-3 beta, which has been proposed as a regulator of the balance between long term potentiation and depression, with downregulated GSK-3

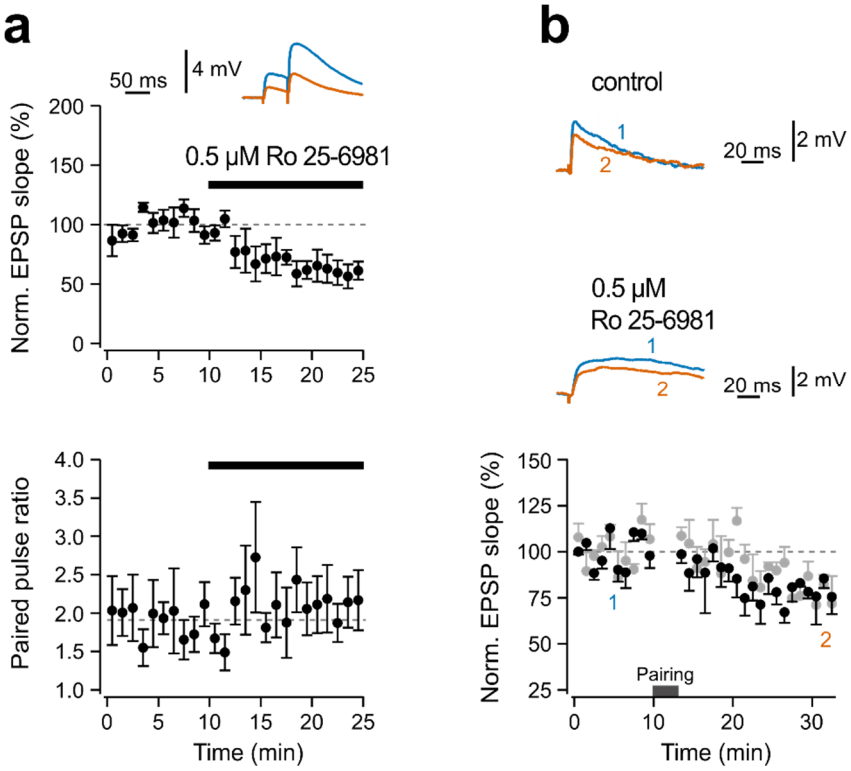


Figure 11 NR2B-containing NMDA receptors are not necessary for plasticity
a) Pharmacological inhibition of NR2B-containing, mostly presynaptic, NMDA receptors decreases the slope of the evoked response (top)but has no detectable influence on the paired-pulse ratio (bottom).
b) Inhibition of NR2B-containing NMDA receptors does not prevent the induction of synaptic weakening induced by subthreshold UP state pairings. All traces are mean ± SEM.

levels after LTP being responsible for reduced capabilities for LTD (Peineau et al. 2007). Furthermore, (Vyazovskiy *et al.,* 2008) showed that GSK-3 beta is hyperphosphorylated during sleep. To test the involvement of GSK-3 beta in UP state induced plasticity, we included an inhibitor of GSK-3 (SB415286) or a vehicle control (0.01% DMSO) in the patch pipette and conducted subthreshold UP state pairings (Figure 12 a). Interestingly, blocking GSK-3 beta seemed to block synaptic weakening (post pairing EPSP slope = 106.2 ± 5.2% baseline (n = 6), compared to 80.16 ± 4.1% (n = 6) in the control group ($t(10) = 3.91$, $P < 0.01$, unpaired t-test).

Finally, we wanted to investigate the involvement of the caspase-3 dependent plasticity pathway. It has been shown that caspase-3 can be involved in LTD that is associated with spine pruning (Li et al. 2010). This would be an interesting trajectory, as during development dendritic spine density follows a similar inverted U-curve as does sleep slow wave activity. Sleep has thus been often hypothesised to be linked to developmental pruning, but such an association has never been shown. To test the involvement of Caspase-3, slices were incubated in an inhibitor of caspase 3 (2 μM DEVD-FMK) for 30 minutes before the beginning of an experiment. To control for this incubation period, control slices were treated identically but incubated in ACSF alone. As shown in (Figure 12 b), the mean EPSP slope was similarly decreased in the group exposed to Caspase-3 inhibition (N=5, 80.66 ± 12.2 SEM) compared to controls (77.26 ± 9.02 SEM, $p = 0.503$, two-tailed independent samples T-Test). Importantly, there was still significant weakening across both groups ($p < 0.05$, two-tailed one-sample t-test, N=10).

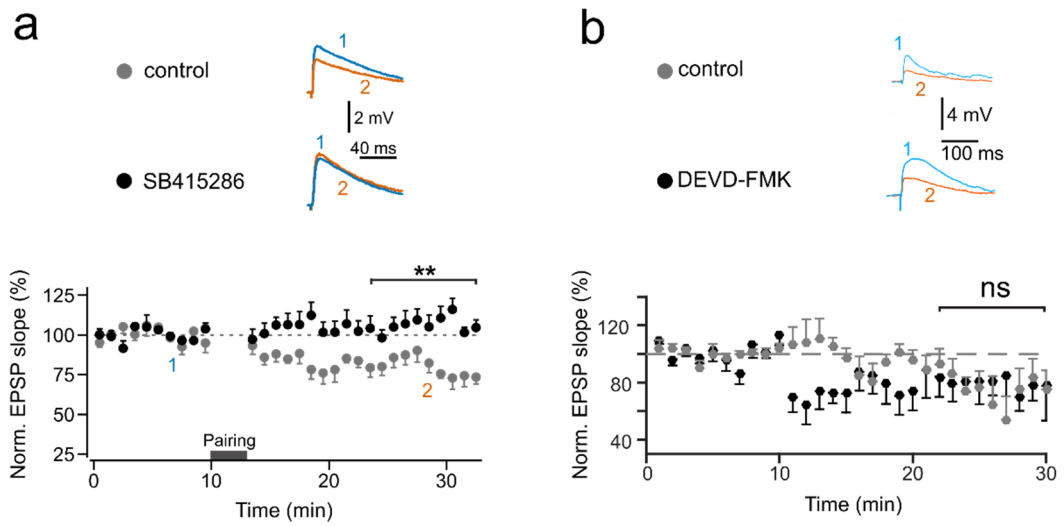


Figure 12 The mechanism of UP date induced weakening involves GSK-3 but not Caspase-3
*a) Including an inhibitor of GSK-3β in the patch pipette (10 μM SB415286) impairs UP state induced weakening (**: p<0.01)*
b) Incubation of slices in an inhibitor of caspase-3 (2 μM DEVD-FMK) did not lead to significantly different changes in EPSP slope in the last 10 minutes of recording. All traces are mean ± SEM.

3.3 Discussion

In this chapter, I described my contributions to our investigation into the interaction between cortical UP states and plasticity of excitatory inputs to LIII pyramidal cells in slices of the medial entorhinal cortex of adolescent mice (Bartram/Kahn *et al.,* 2017). We first established that pairing *sub*threshold inputs with UP states consistently leads to synaptic weakening while pairing with DOWN states does not. Furthermore, when we paired *supra*threshold inputs with UP states, plasticity was also not detectable. I have consciously avoided the term long-term depression to describe this phenomenon, as, owing to the difficulty of long-lasting stable recordings, we routinely only recorded for 20 minutes following the pairing. However, in such instances when we did recordfor longer, the weakening never recovered. Furthermore, the paired-patching experiments conducted by J.Bartram suggest that UP state induced plasticity in input specific, thus it seems to group with Hebbian rather than homeostatic types of plasticity. In contrast to

what may be expected from a classic spike-timing dependent plasticity protocol (Markram et al. 1997), the presence of synaptic weakening seemed independent of postsynaptic spiking. However, our observations fit well with other studies showing that pairing inputs with slight depolarisation of the cell can lead to LTD in neocortex (Sjöström et al. 2004). Interestingly, Sjöström and colleagues found that depolarisation induced LTD and classic, timing dependent LTD occluded each other, indicating perhaps also that plasticity during UP states is not mechanistically distinct from classic LTD pathways. As such, I personally see our description of UP state-induced weakening less as a novel form of plasticity and more as a demonstration that a known form of LTD may be at play during a physiologically highly relevant form of activity. This is not a trivial finding because, while the UP state is indeed a period of somatic depolarisation, it is also a period of intense excitatory and inhibitory synaptic activity engulfing large parts of the dendritic tree (Waters & Helmchen 2004).

We subsequently demonstrated that pairing inputs with postsynaptic bursts during UP states led to lasting (more than 50 mins) potentiation and that subsequent subthreshold UP state pairings were ineffective at changing the time course of this potentiation. This experiment should be interpreted with caution, because the EPSP decreased by itself after potentiation (i.e. there was a main effect of time but no interaction between time and group). Unfortunately, the already low success rate of these experiments made it impractical to wait for the EPSP to stabilise after burst pairing. In consequence, we may be examining a period in which acute and long-term maintenance of LTP are fading into each other. Nevertheless, it seems that whichever mechanism is responsible for mediating the time course of potentiation is overlapping or incompatible with UP state induced weakening. This is interesting in the light of our demonstration of an

involvement of GSK-3 beta, which has been shown to be downregulated after LTP, and thereby inhibit subsequent LTD (Peineau et al. 2007).

Nevertheless, it might be slightly surprising that subthreshold UP state pairings are not able to de-potentiate the EPSP. One further experiment could thus be to test whether classical 1 Hz LTD can de-potentiate burst LTP (Fujii et al. 1991).This would clarify, whether this entorhinal burst LTP simply cannot be de-potentiated or whether UP state induced weakening cannot act as de-potentiating stimulus. Lastly, note that the burst pairing was delivered during an UP state, making it possible that UP state induced weakening and burst pairing elicited two independent and opposite plasticity effects with different time courses.

We next sought to localise the expression of plasticity using two physiological measures, the CV and paired pulse facilitation. While both of these methods have separate caveats, and indeed silencing of synapses might confound both methods (Costa et al. 2017), they both point towards a presynaptic contribution in expressing plasticity. One way to shine more definite light on this issue would be to use glutamate uncaging, thus circumventing presynaptic release. Nevertheless, our findings are once more consistent with (Sjöström et al. 2004), who also concluded on a presynaptic expression of plasticity based on a CV analysis.

Induction of plasticity was clearly dependent on postsynaptic mechanisms. This is shown by the impaired plasticity upon specific pharmacological manipulation of the postsynaptic cell with an NMDA receptor antagonist and an inhibitor of GSK-3 beta. This is in broad agreement with timing-dependent LTD (Bi & Poo 1998; Sjöström et al. 2001) and other forms of LTD (Collingridge et al. 2010). In contrast to (Sjöström et al. 2004), blocking presumed presynaptic NR2B receptors had no effect on plasticity. This

could mean a mechanistic difference between our and Sjöström's findings. However, in my opinion it is more likely that there is one or more flaws in the interpretation of pharmacological experiments targeting pre- or post-synaptic NMDA receptors. This includes, of course, our own experiments. Firstly, most of the commonly used "subunit specific" antagonists, are mostly tested in comparison to the most common NMDA receptors containing NR2A-subunits , and not on other isoforms containing NR2C, NR2D, or the more recently identified NR3 subunits. For instance, (Sjöström et al. 2004; Sjöström et al. 2003) used ifenprodil, which has been tested in frog oocytes to be specific for NR2B over NR2A containing receptors (Williams 1993).This leaves open the possibility of binding other subunits or even having other targets (e.g.(Brittain et al. 2012)). Secondly, the idea that some NMDA subunits are selectively expressed at specific cellular location is oversimplified. For instance, it was recently found that around 50 % of all naturally occurring NMDA receptors in cortex are heterotrimeric and significantly affected by ifenprodil (Hansen et al. 2014).To come back to our results, it could thus simply be that a combination of differences in age, brain region, and pharmacological agent accounted for the different result. Interestingly, Sjöström's presynaptic LTD depended on endocannabinoid signalling, which is an established retrograde signalling molecule. It would thus be interesting to see whether UP state induced plasticity also depends on endocannabinoids.

Our finding that GSK-3 is involved in plasticity provides an interesting link to molecular changes during sleep, as (Vyazovskiy et al. 2008) showed that GSK-3 beta phosphorylation is increased in cortex and hippocampus during sleep. It could be that increased GSK-3 activity is facilitating synaptic weakening, but it could of course simply be a consequence of less LTP. Either way, it would be interesting to see the effects of GSK-3 beta manipulation on synaptic weights during natural sleep. Of note,

the inhibitor we used (SB415286) inhibits two isoforms of GSK-3 (Bain et al. 2007), however GSK3β is the much more likely driver of the effects on our case as it is highly expressed brain tissues (Lau et al. 1999). More generally, the involvement of GSK-3 in this form of plasticity could be investigated with additional methods, such as viral vectors.

Overall, our *ex vivo* results indicate that UP states may provide a period in which *sub*threshold but not *supra*threshold inputs are weakened. In the central nervous system, no single cell can reliably drive another cell by itself. It thus follows that UP states may indeed promote the survival of well-coordinated firing patterns while weakening more noisy activity. This could provide an interesting new angle into the debate about how sleep can cause net synaptic weakening (e.g. (de Vivo et al. 2017b)) *and* be crucial for consolidating memory. A tantalizing (and frequently asked) question, however, is whether we would expect this type of plasticity to happen during waking. After all, UP states closely resemble waking activity and have been proposed to be "fragments of waking" (Destexhe et al. 2007). A provocative perspective would be, why not? I am not aware of a study who has tested the effects of prolonged subthreshold pairings in the wake brain. The closest study is (Feldman 2000) who found that random inputs to barrel cortex neurons tend to be depressed based on a bigger window for spike timing depression compared to potentiation. Indeed, one main criticism of the synaptic homeostasis hypothesis of sleep is, that there is no overwhelming reason why synaptic weakening cannot continuously happen during waking. Another, less provocative, perspective would be that the neuromodulatory tone (e.g. acetylcholine, noradrenaline) is practically absent during slow-wave sleep (and in brain slices), which likely has a crucial impact on plasticity rules. Nevertheless, ultimately this conundrum is best solved with experiments in naturally behaving animals (see chapter 5).

This larger issue of generalisability is, in my view, the strongest limitation of all our experiments. The parameter space of potentially relevant experimental factors is tremendous. Examples for such factors are the brain region, age of the animal, ion concentration, recording temperature, cell type, neuronal pathway, and other factors that differ in the intact brain. It is thus encouraging that we (A.Langfelder, unpublished) and others (Gonzales-Rueda 2018), have confirmed the basic principle of UP state induced weakening, as I will touch upon in the next chapter.

Chapter 4 – Neuronal plasticity is associated with changes in cortical slow waves *in vivo* under anaesthesia

4.1 Introduction

The synaptic homeostasis hypothesis of sleep (SHY, e.g. (Tononi & Cirelli 2003)) proposes that a key function of sleep is to facilitate a net weakening of synapses, which may counteract a net strengthening of synapses during waking. As synaptic strengthening and weakening occur in local cortical circuits, this hypothesis could elegantly account for local changes in sleep intensity after learning (Huber et al. 2004). Several studies in recent years were consistent with the notion of net weakening during sleep (Vyazovskiy et al. 2008; de Vivo et al. 2017b; Hengen et al. 2016), even though some contrasting studies (Chauvette et al. 2012) suggest that it may not hold true under all circumstances.

While this central prediction of SHY has thus been tested to some degree, it remains unclear whether synaptic weakening could, in principle, account for key physiological variables affected by sleep, such as sleep pressure. Perhaps the most commonly used measure of sleep pressure in both clinical and animal work is slow wave activity (SWA) during non-rapid eye movement (NREM) sleep, quantified as the spectral power in the band between approximately 0.5 and 4 Hz. It is broadly accepted that SWA reflects the homeostatic sleep drive on a global cortical level (Borbely 1982) and that it can also act on a local level (Kattler et al. 1994; Vyazovskiy et al. 2000; Huber et al. 2007; Huber et al. 2004). In the light of this, the authors of SHY speculate that changes in synaptic strength could in some instances account for changes in SWA that occur during sleep. Indeed, a computational simulation of a corticothalamic network supported the idea that

synaptic weakening can reduce the power of slow waves (Olcese et al. 2010). However, one difficulty with answering this question has been that it is unclear how synaptic weakening may occur during sleep.

We (Bartram/Kahn et al. 2017) and others (Gonzales-Rueda et al. 2018) have recently found that the active phase of the slow oscillation (UP state) may facilitate synaptic weakening of cortical excitatory connections in brain slices and under anaesthesia. Here we used this UP state induced synaptic weakening protocol to ask whether the resulting plasticity in turn influences SWA. We record evoked responses to whisker deflections and SWA in the primary somatosensory cortex and find that the induction of plasticity (UP state pairings) but not control stimulation (DOWN state pairings) is associated with a reduction in slow wave activity. This reduction is local and highly correlated with the local extent of change in the evoked response.

I would like to note again (see methods) here that all credit for setting up and conducting these recordings goes to Antonia Langfelder, who wanted to test whether UP state induced plasticity exists in this setup. After Antonia left the lab, I spike-sorted and re-analysed these data from the above-outlined perspective. While some of the below analyses have been conducted in a similar manner by Antonia and Ed Mann, all the analysis I present here is my own.

4.2 Results

4.2.1 Whisker-evoked responses and UP/DOWN state pairings *in vivo*

We recorded the local field potential (LFP) and multi-unit activity (MUA) in the mouse barrel cortex using 8x8 channel electrode arrays with 200 μm spacing, thus sampling

several barrel columns (Figure 13 a). On average, 32.3 (with SD of 16.6) channels had clear negative responses in the LFP. Spontaneous activity under medetomidine-isoflurane anaesthesia was dominated by the slow oscillation, as is evident from the presence of alternating periods of high and low spike rates, the former coinciding with negative deflection in the LFP of superficial layers (Figure 13 b). As expected from previous studies, single whisker deflection caused a transient negative deflection in the superficial LFP together with a brief elevation of spike rates in the MUA. Moreover, the responses were strongest in one "principal channel" and decayed significantly with distance from this channel (Figure 13 c&d).

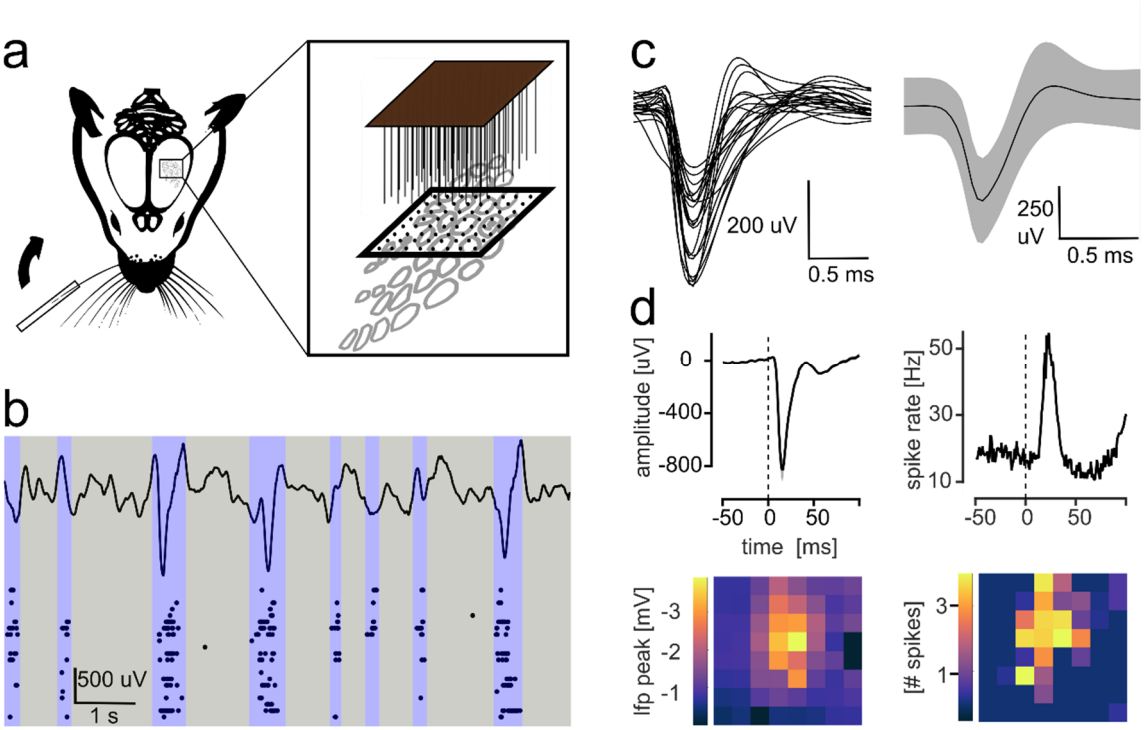


Figure 13 Experimental setup, spontaneous activity and evoked responses
a) 64-channel multi-electrode arrays were inserted into the barrel cortex. Electrode spacing is 200 µM. b) Representative LFP traces and spike times of spontaneous activity under medetomidine-isuflorane anaesthesia. LFP trace is from one channel and each line of the raster plot corresponds to one channel. Blue and grey shading represents UP and DOWN states, respectively. c) left: mean MUA waveforms across the entire experiment of the channels shown in b. right: largest amplitude waveform with standard deviation in grey. d) top: evoked LFP and MUA response for one representative channel; bottom: distribution of LFP peak amplitude and response magnitude across all 64 channels in one representative experiment.

In order to test the differential impact of UP and DOWN state pairings on evoked responses, whisker deflection occurred once every 5 s during a baseline period, after which whisker deflection was paired with either UP or DOWN states (Figure 14a). During pairings, online detection of an UP or DOWN state triggered 5 whisker deflections at 10 Hz. To estimate the success rate of the online UP/DOWN state detection in each experiment, a mock-pairing was recorded before the start of each recording session (Figure 14 b). To quantify state detection with minimal assumptions about what defines an UP state, the number of spikes in the 50 ms following each state detection was compared to the mean number of spikes for each recording (Figure 14c, left). As expected from successful state detection, the distribution of spike rates at pairing onset are almost completely nonoverlapping ($p<0.001$, independent samples Kolmogorov-Smirnov test)

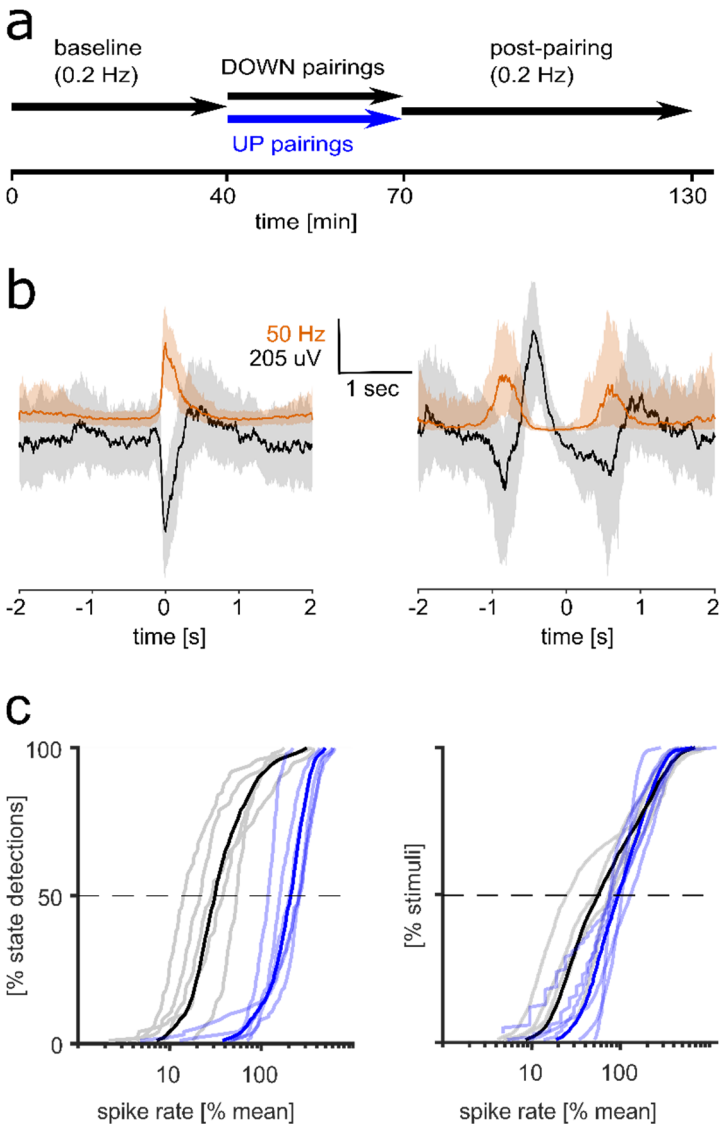


Figure 14 Experimental design and UP/DOWN state pairings
a) Experimental design: before and after pairing, whisker deflection occurred once every 5 seconds. During pairings, each state detection triggered a 10 Hz burst of 5 stimuli
b) median ± 25th percentiles of spike rate (orange) and LFP (black) upon mock UP (left) or DOWN (right) state detection
c) Quantification of pairing success rate. Left: cumulative distribution of spike rates within 50 ms after mock UP state (blue) or DOWN state (black) detection. Right: identical analysis for all stimuli, which would have been delivered after a state detection.

between groups and the UP state pairings are associated with significantly higher median spike rates (210.88 ± 24.8 % of mean spike rates for UP state pairings compared to 31.53 ± 6.89 % in DOWN state pairings, p=0.0043; two-sided Wilcoxon rank sum test, n=6 and 5, respectively). While this suggests that state detection is adequate, it does not guarantee that all 5 stimuli (of a 10 Hz burst triggered upon state detection) fall within the correct state. Indeed, the same analysis for all mock stimuli suggests that significant amounts of stimuli fall into incorrect states (Figure 14c, right). While the two distributions are still significantly different ($p < 0.001$ independent samples Kolmogorov-Smirnov test, note that this treats individual percentiles as independent samples), the difference is clearly smaller. This is perhaps best illustrated by the mean evoked spike rate, which is still significantly higher during mock UP state stimulation than in mock DOWN state stimulations (UP pairings 97.16 ± 8.41 % (22.6 ± 7.98 spikes), of mean spike rate vs DOWN pairings: 57.36 ± 9.743 % (14.02 ± 4.76 spikes), during DOWN pairings, p=0.017, Wilcoxon rank-sum test; n=6 and 5, respectively). In conclusion, in mock pairings UP and DOWN states are detected with high accuracy and lead to significant enrichment of UP or DOWN state pairings in the respective group. However, the amounts of UP state pairings that occur in the DOWN state pairing group may be sufficient to be physiologically relevant.

4.2.2 Whisker evoked responses are weakened following UP but not DOWN state pairings

To test the differential impact of UP and DOWN state pairing on sensory-evoked responses, a baseline period of evoked responses was obtained before pairing whisker deflections with either UP states or DOWN states for 30 minutes, using closed-loop online detection based on the LFP (Figure 14 a & methods). Following pairing, whisker

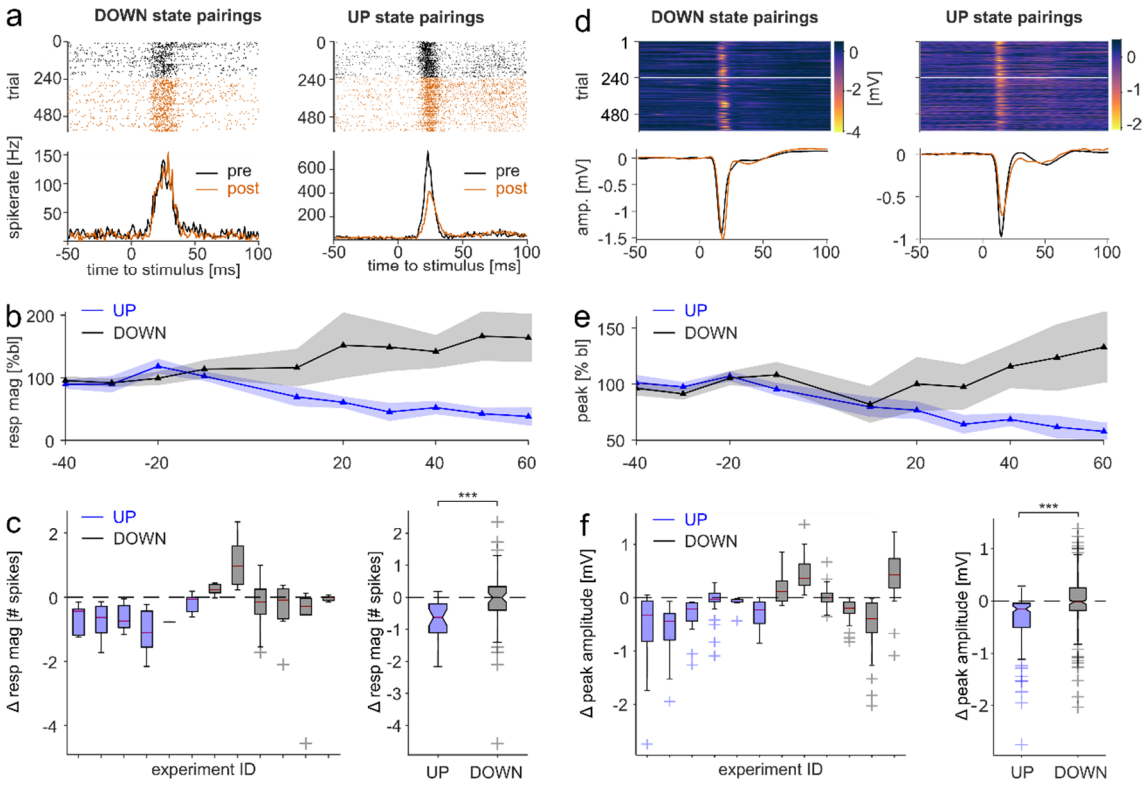


Figure 15 Evoked responses are decreased following UP state but not DOWN state pairings
a) Example raster plots (top panels) and average PSTH (bottom panels) for DOWN (left) and UP (right) state pairings. Black traces are pre-pairing responses and orange traces are post-pairing responses
b) Average time course of response magnitude across mice in UP (blue) and DOWN (black) state pairings.
c) Left: change in response magnitude across channels for each individual animal in UP and DOWN state pairings, right: change in response magnitude across all channels and animals
d) Top panels: example trial-by-trial LFP traces from same animal as in a). bottom panels: average LFP across all trials (black: pre-pairing, orange: post-pairing)
e) Average time course of LFP peak amplitude across all mice
f) left: absolute change in LFP peak amplitude across all channels for each individual animal in UP and DOWN state pairings, right: change in LFP peak amplitude across all channels and animals

deflection continued for 60 minutes identical to the baseline stimulation. Note that in some experiments instead of interleaving 5-minute blocks of recording and pauses, deflection of two separate whiskers was interleaved (see methods). As expected from previous studies, the response magnitude in the MUA (Figure 15 a&b) and the peak of the evoked LFP (Figure 15 d&e) decreased following UP state but not DOWN state pairings ($P < 0.05$ for an interaction between group and time, mixed ANOVA; n=6 for each group but see methods). However, as shown in (Figure 15 c & f), there was considerable variability in plasticity across experiments and even across channels within single experiments. Hence, the mean change across all channels may be driven by

individual channels with large changes. Indeed, across all channels of one condition (UP or DOWN pairings), the changes in LFP and MUA were not normally distributed ($P<0.001$, Shapiro-Wilk). To address this, the changes in evoked responses were pooled across all channels of a given condition. As shown in (Figure 15 c&f), the resulting distributions are still significantly different ($p<0.001$, for LFP peak [n=214, 174; DOWN, UP] and MUA response magnitude [$p< 0.001$ n= 61,30 DOWN, UP] independent samples Wilcoxon rank-sum test), thus confirming the effect of UP state pairings on evoked responses.

One important methodological difference between this and previous studies is the lack of control over postsynaptic spiking in this setup. While one might thus expect a correlation between response strength and the degree of depression, no significant relationship was detectable (see discussion). It is known that under anaesthesia, UP and DOWN states have differential effects of neuronal gain (Reig 2015). Therefore, it is possible that changes in the average evoked response are due to changes in the frequency of UP and DOWN states. To address this concern, individual responses were separated according to the spiking activity in the 100 ms preceding whisker deflection. As expected from previous work, evoked LFP and spiking responses during UP states (i.e. 1 or more spikes) were smaller than responses during DOWN states (only pre-pairing responses were used for this analysis, Figure 16).

Importantly, a mixed-ANOVA analysis on the normalised responses (Figure 16 c&f) showed a main effect of time but no interaction between time and group suggesting that different frequencies of UP and DOWN states cannot explain plasticity (LFP: time*state: df=9, F=1.286, p=0.304; main effect of time: df=9, F=12.431, p=0.000195; no main effect of group: df=1, F= 1.307,p=0.286; Greenhouse-Geisser correction

applied for both; MUA: time*state: df=9, F=0.833, p=0.474; main effect of time: df=9, F= 9.496653, p=0. 00066; no main effect of group: df=1, F=3.617; p=0.094, Greenhouse-Geisser correction applied for both). Overall, our results are consistent with previous intracellular studies and thus the most parsimonious interpretation is that UP state pairings induced changes in synaptic weights.

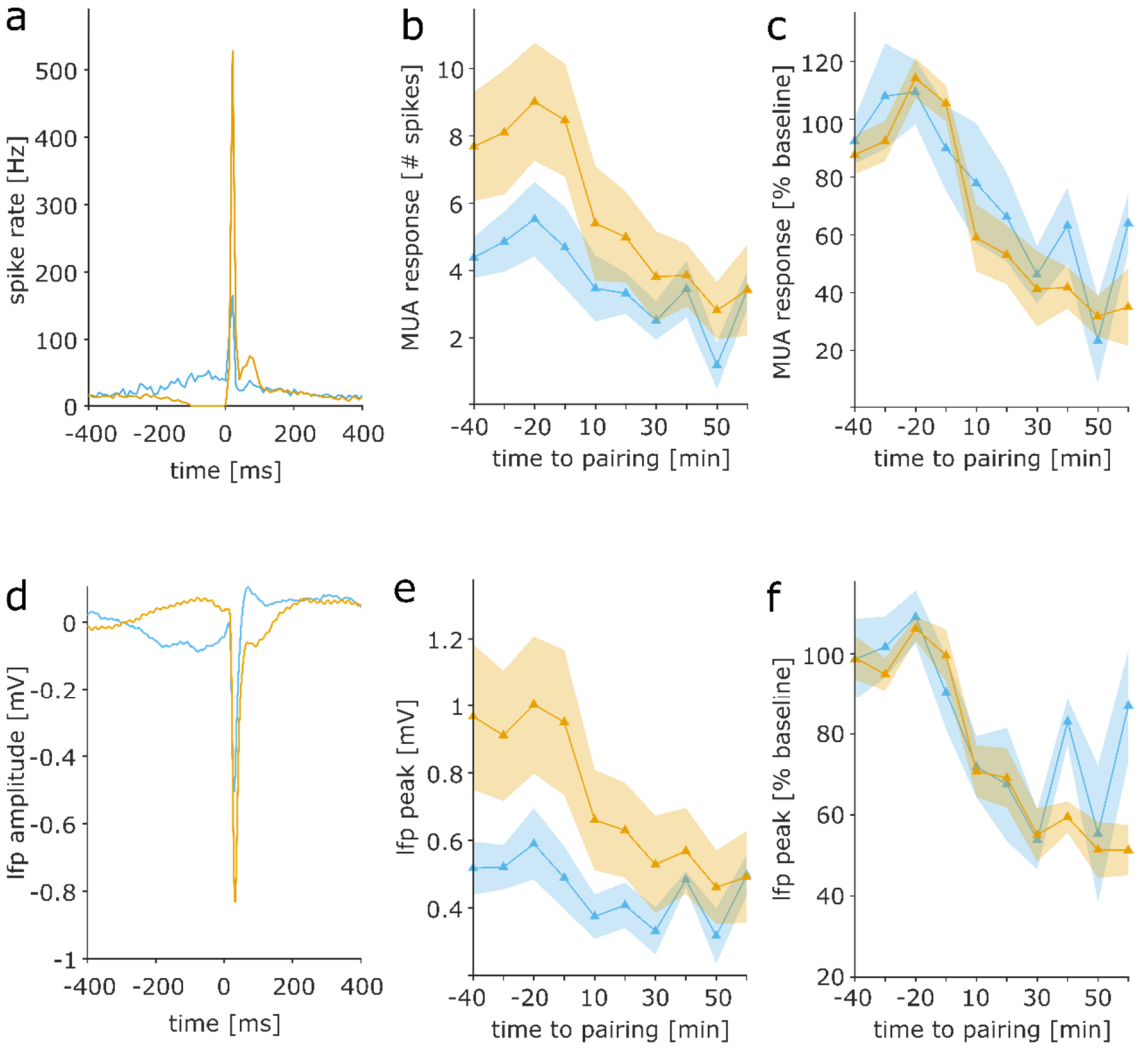


Figure 16 Plasticity in the "UP state pairing" condition is not caused by changes in the slow oscillation
a) Example PSTH of pre-pairing whisker-evoked spiking responses falling into DOWN states (orange) or UP states (blue).
b) & c) Mean time course (across mice) of whisker-evoked spiking responses in the UP-pairing condition sepparated according to state at the time of stimulation. b) absolute values illustrating smaller UP state responses c) relative change (compared to pre-pairing), illustrating similar trajectory once the offset is corrected
d-f) Same as a-c but for whisker-evoked LFP peak amplitude

4.2.3 Plasticity is associated with a dose-dependent decrease in slow wave activity

Having replicated UP state induced weakening of evoked responses in this setup, we next wanted to know whether plasticity of neuronal responses is associated with overall changes in SWA. We thus quantified the average SWA across all responsive channels for each animal in a three second epoch beginning two seconds after each whisker deflection. We observed that, following UP state pairings SWA was reduced to 89 % (± 3.2 % SEM) of its baseline value whereas it was not on average reduced following DOWN state pairings (97 % ± 2.15, $P < 0.05$ for an interaction between group and time, mixed ANOVA; n=6 per group, Figure 17 a & b). This relationship between response plasticity and SWA extended to individual recording sites, as there was a significant correlation between the change in peak LFP response and SWA following UP state pairings (Figure 17 c&d, Pearson's r: 0.57, $p<0.001$, n=310, this correlation was computed following outlier exclusion to avoid distortions, when all data points were included Pearson's r was 0.61, $p<0.001$, n=384). One concern with this analysis is that slow waves can be influenced by sensory stimuli (Ngo et al. 2013), and it is thus possible that decreased SWA is simply a consequence of weakened whisker input. To address this concern, we made use of a subset of experiments where we recorded from the same sites but alternated between blocks of stimulating two different whiskers. To isolate the effects of altered sensory input, we selectively analysed the baseline period prior to the induction of plasticity. For most recording sites, stimulating a different whisker had a strong impact on the evoked response, comparable to the effects of plasticity. We thus asked whether the changes in evoked response due to different whisker deflection (during the baseline period) correlate similarly with changes in

SWA. As shown in (Figure 17 c), the correlation is just short of being significant and, importantly, much smaller than for the plasticity (following outlier exclusion: Pearson's r: 0.13, p=0.052, n=210, including all data points r:0.31, p=0.31, n=320), suggesting that altered sensory drive cannot explain altered SWA in this instance. Taken together, these results are consistent with the idea that local changes in neuronal connectivity can affect the local expression of slow waves.

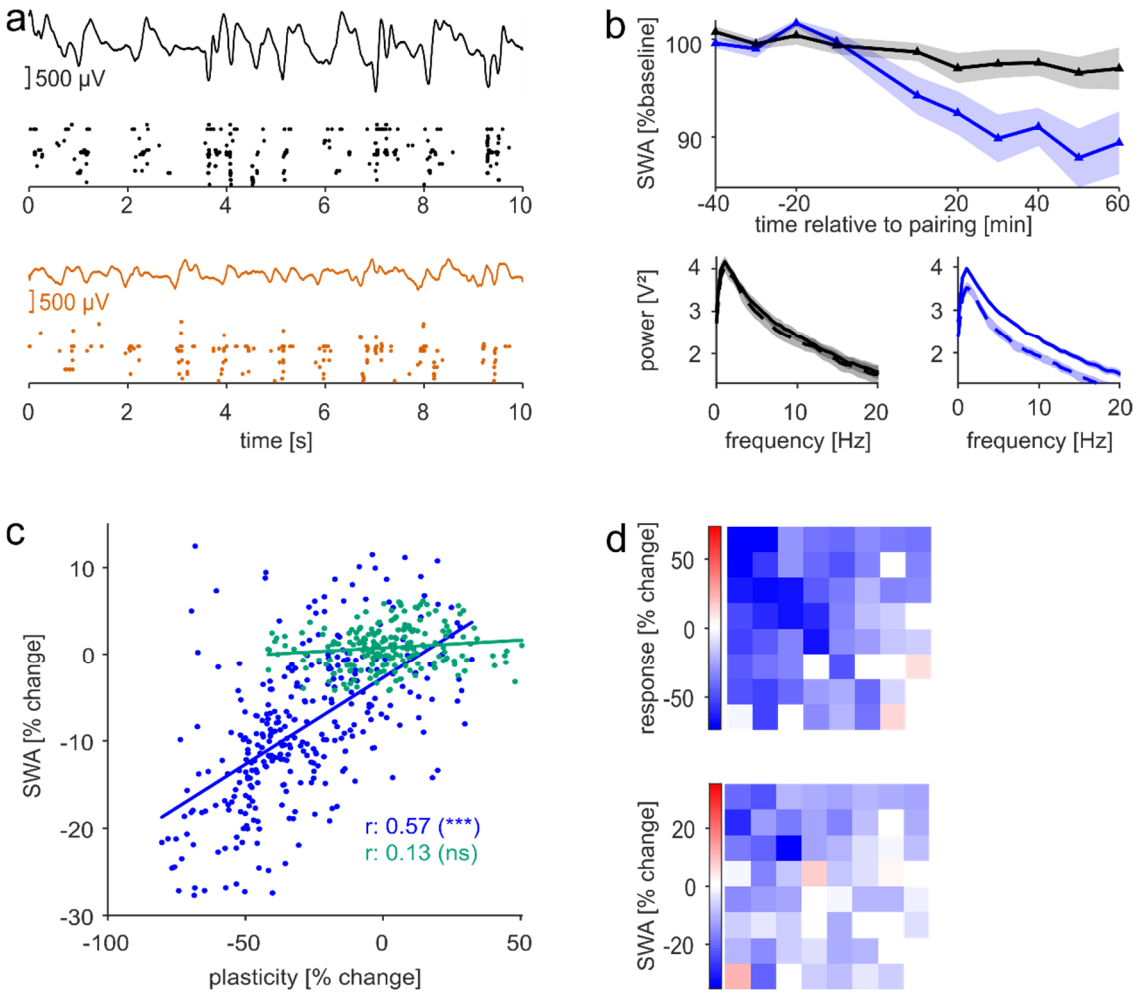


Figure 17 Slow wave activity decreases dose-dependently following UP state but not DOWN state pairings

a) Example LFP and MUA activity before (black traces) and after (orange traces) UP state pairings

b) Top: mean timecourse (± SEM) of SWA activity across experiments in UP state pairing (n=6) and DOWN state pairing experiments (n=6). Bottom: Mean power spectral density (±SEM) before (solid lines) and after pairing (dashed lines) in DOWN (black, left) and UP (blue, right) state pairings.

*c) Correlation between change in evoked response and change in SWA, each dot represents one channel. Blue dots: Change from pre- to post-pairing in UP state pairing experiments. Green dots: control analysis comparing changes from stimulating a different whisker in the pre-pairing period of UP state pairing experiments (i.e. the "plasticity" measure reflects different sensory input rather than synaptic plasticity). r = Pearson correlation coefficient. ***: p<0.001*

d) Example of spatial pattern of changes in evoked responses (top) and SWA (bottom) following UP state pairings across 8x8 channel array.

4.2.4 Changes in slow waves following plasticity are comparable to changes following natural sleep homeostasis

A change in the power spectrum of an oscillation provides little information about what specific aspect of the oscillation has changed (e.g. decreased power in in SWA could result from a reduction in amplitude or incidence of individual slow waves, or a mixture thereof). Previous studies in freely moving mice associated the decrease in SWA from early to late sleep with decreased slope and amplitudes of individual slow waves, while the total number of slow waves remained constant (Vyazovskiy et al. 2009; Vyazovskiy et al. 2007). We thus wanted to investigate whether the decrease in SWA activity that we observe under anaesthesia is associated with a comparable change in individual slow waves.

We therefore detected individual slow waves in a two second window beginning 3.5 seconds after each stimulation using criteria similar to (Vyazovskiy et al. 2009). As expected, the detected slow waves had a negative peak associated with the UP state and a positive peak associated with a subsequent DOWN state. Furthermore, larger amplitudes were associated with larger modulation of spiking activity (Figure 18a). We quantified several aspects of the average slow wave in each channel, including the amplitude and slope of both peaks and the duration of the UP state. Focussing only on UP state pairing experiments, we found that all amplitudes and slopes were strongly reduced while the duration and number of slow waves seemed largely unaffected (Figure 18 b, Bonferroni corrected two-sided t-test test across experiments (n=6): a1:0.0055, a2: 0.0046, s1: 0.0036, s2: 0.0021, duration: 0.5075, number of events: 0.4411. Across channels: $p<0.001$ for a1, a2, s1, s2, and p=0. 0058 and p=0. 0504, two-sided signed-rank test).

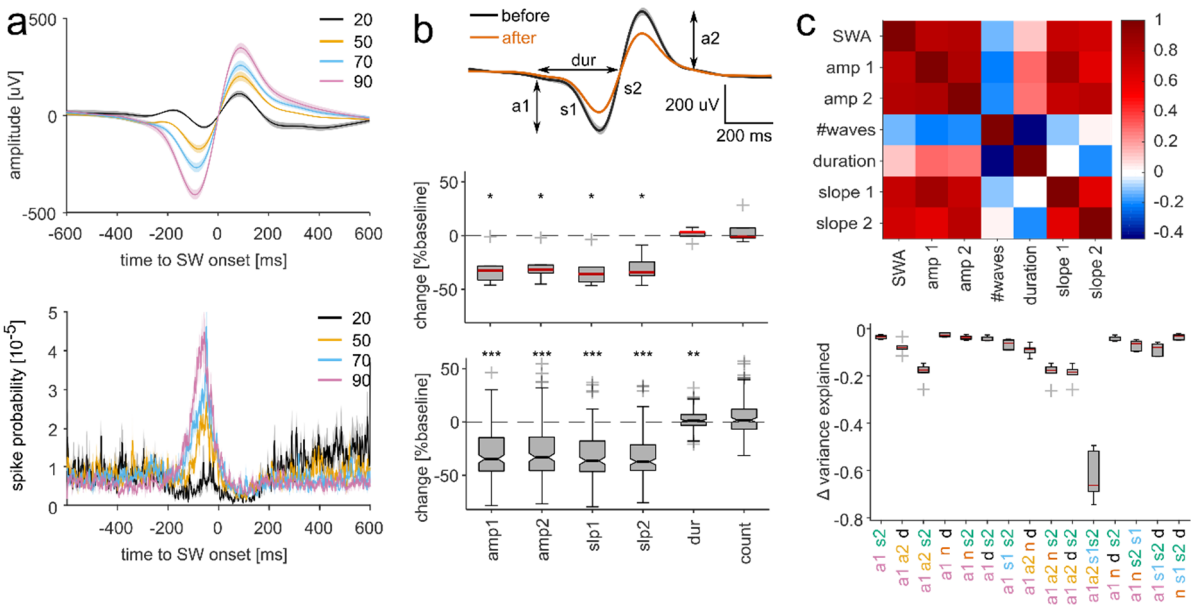


Figure 18 Changes in slow waves following UP state pairings
a) Examples of individual slow waves detected in the LFP of one mouse (top panel), and the MUA associated with the detected slow waves (bottom panel). Colours depict different percentiles in the LFP peak amplitude of detected slow waves.
*b) Changes in the shape of slow waves following UP state pairings. Top panel shows average slow wave in one channel of one mouse and the various measures that can be computed. A1: first amplitude, A2: second amplitude, s1: first slope, s2: second slope, dur: duration. Middle and bottom panels show the change in various measures across mice (middle) and across channels (bottom). *p<0.05, ***p<0.001*
c) Top: correlation matrix between the measures of slow waves shown in b), as well as the number of detected slow waves (#waves) and SWA. bottom: variance in SWA explained by a regression model including all measures except the ones shown on the x-axis.

While there is no physiological imperative that individual aspects of a slow wave should be tightly correlated, this was clearly the case in our recordings (Figure 18 c). Specifically, we examined these correlations using each 3.5 s window between stimulation events as one observation. The amplitudes of the first and second peak (UP state followed by DOWN state) were strongly positively correlated with each other and their respective slopes. In contrast, the duration of slow waves was only weakly positively correlated with these measures and the number of slow waves was negatively correlated with all measures. One interpretation of these results is that the slow oscillation is simply scaled after UP state pairings, thus leaving the duration and incidence of individual waves largely unaffected. Because the overall number of slow waves did not change, the negative correlation between the number of slow waves and

their amplitude cannot explain our results but is still surprising. One possibility is that smaller slow waves have shorter refractory periods and/or that some faster but smaller oscillatory events have been detected as slow waves.

4.3 Discussion

In this chapter I reported my analysis of slow wave induced plasticity *in vivo* under anaesthesia, which aimed to ask whether weakening of connectivity is associated with weakening of local slow waves in the primary somatosensory cortex.

4.3.1 UP state induced plasticity

As a first step, I spike-sorted the recordings and ensured that the effects of recording noise are minimal. I then showed, as Antonia Langfelder has before me, that pairing whisker deflections with UP states but not DOWN states leads to a net weakening of responses. This effect persisted when whisker deflections were grouped according to stimulation during UP or DOWN state, which suggests that changes in the frequency of UP and DOWN states cannot account for plasticity in the response. A key problem with the interpretation of this dataset is, that 5/6 animals in the UP state pairing group received a slightly different pattern of stimulation. While the one animal that underwent the "regular" stimulation pattern was one of the clearest cases of weakening, this difference should be addressed by further experiments. Unfortunately, this was difficult as departmental construction essentially incapacitated the room for these experiments for almost a year. The reason I still consider our experiments valuable is because they are highly consistent with what would be expected based on our previous work in slices (Bartram/Kahn et al. 2017) and the work from Ole Paulsen's group in the barrel cortex of mice *in vivo* (Gonzales-Rueda et al. 2018).

While our plasticity protocol shares many key features with (González-Rueda et al. 2018), we use extracellular methods to record sensory responses across multiple sites in barrel cortex while Gonzales-Rueda and colleagues optogenetically stimulate L4 cells and record intracellularly in pyramidal cells in layer 2/3. Thus, given the data alone it is difficult to establish where plasticity happens in our case. Whisker evoked signals travel via the brainstem over one of several possible thalamic nuclei before entering the cortex mainly through layer 4 (Fox 2008). The whiskers are highly compartmentalised throughout these stations, so that one whisker primarily innervates the layer 4 of one column, whereas other columns are mostly innervated through layer 2/3 cortico-cortical projections. While plasticity could theoretically occur at any synapse in this pathway, the mice we used (6-8 weeks old) where past the critical period of thalamocortical plasticity, thus they were at an age when layer 2/3 synapses have been shown to be much more plastic (Fox 1992; Glazewski & Fox 1996). Moreover, it is worth noting that while we activate a larger circuit than Gonzales-Rueda et al., activation of L4 to 2/3 synapses is common to both studies. This leaves two plausible possibilities: First, plasticity may mainly have occurred at L4-L2/3 synapses of the principal barrel. Significant reduction of evoked responses at other locations would thus reflect reduced feedforward excitation. Second, plasticity occurred also at L2/3-2/3 synapses. While the data cannot give further insight, the second option seems much more likely to me because lateral layer 2/3 circuits would have experienced significant amounts of subthreshold pairings (see two paragraphs below for more detailed statistics). One would thus have to assume L2/3 but not L4-L2/3 connections are immune to subthreshold UP state pairings. Furthermore, changing L4-2/3 connectivity in one barrel would be very unlikely to change SWA across several barrels (see below). Overall, there is a degree of uncertainty to the localisation of plasticity but most likely, it

occurred at L4-2/3 synapses in the principal barrel and also in L2/3-L2/3 cortico-cortical synapses across neighbouring barrels.

Another pressing matter is whether the plasticity that we observe is due to changes in synaptic strength. Other explanations could be the depth of anaesthesia, neuromodulatory tone, changes in intrinsic excitability or tissue damage. The latter can be accurately quantified as being possible with a probability of less than 5%, as it would have occurred in the controls as well. Depth of anaesthesia and (thus) neuromodulatory tone could both change due to UP state pairings, yet it would be surprising for such changes to last for an hour. Lastly, changes in excitability clearly do occur, as evident by the altered spontaneous firing rate (data not shown). However, changes in excitability are well known to accompany LTP and LTD protocols beginning with their first description (Bliss & Lomo 1973). Indeed, based on the above-outlined likelihood of other possibilities, synaptic plasticity would seem the most likely cause for changes in excitability.

One last aspect worth discussing is that we had little control over correlated postsynaptic spiking. Thus, the question is what the distribution of spiking probability in response to single whisker deflection in the barrel cortex is. We cannot directly assess this in our dataset because our arrays are designed to record from multiple columns rather than sample single units. However, whisker responses in layer 2/3 of the barrel cortex are a well-researched together with other models for early cortical sensory processing in general. For instance, (Brecht & Sakmann 2002) recorded intracellular responses to deflection of the principal whisker in layer 4 cells of rats under urethane anaesthesia and reported low rates of suprathreshold responses (0.14 spikes per average cell and whisker deflection). In contrast to this, a large proportion of cells shows clear

subthreshold responses to whisker deflection(Moore & Nelson 1998). Indeed, the authors found that while 20 out of 24 neurons had subthreshold responses, only 7 neurons exhibited suprathreshold responses (in at least 20% of trials). Moreover, subthreshold responses were evoked by deflection of most whiskers (80%) whereas suprathreshold responses were highly limited (10% of whiskers). The situation is similar in the visual system (Carandini & Ferster 2000), exemplifying the notion that subthreshold tuning is considerably broader than suprathreshold tuning in early cortical processing, at least under anaesthesia (see (Barth & Poulet 2012)for an excellent review). In conclusion, the majority of responsive neurons in barrel cortex will likely have displayed subthreshold responses. However, especially in the principal barrel, a significant proportion of neurons will have displayed suprathreshold responses.

4.3.2 Changes in SWA accompany changes in evoked responses

Having confirmed the presence of UP state induced plasticity to the best of my abilities, I moved on to show that a reduction of slow wave activity (0.5 – 4 Hz) accompanies weakening of the evoked response in animals that underwent UP but not DOWN state pairings. Moreover, within single UP state pairing experiments, the extent of response weakening was highly variable and predicted the change in SWA.

This finding is a correlation and thus subject to its weaknesses. First, it could be that changes in SWA drive changes in evoked responses, which I discussed above and think is unlikely. Secondly, another factor could influence both measures without a direct causal connection between the two. Some of these factors have also been discussed above (neuromodulatory tone, tissue damage). A more likely scenario would have been

that the relatively short period during which we measured SWA was influenced by sensory stimulation. This is unlikely because changing the stimulated whisker had no comparable effect to UP state pairings.

A further challenge for interpreting our result is the localisation of the channels. Given that we use a flat array, the intuitive assumption is that differences between channels are due to horizontal distance. However, the brain is curved and inserting a rigid array will undoubtedly result in slight variations in the depth of individual channels. Hence, differences in plasticity and changes in slow wave activity could be driven by horizontal and/or vertical variation. Horizontal variation in plasticity and SWA would fit well with findings on local changes in SWA (Huber *et al*., 2004). Vertical variation in SWA is much less explored territory, as it requires chronic recordings with laminar probes. However, there is evidence of layer specific occurences of slow waves during REM sleep, suggesting that slow waves can be localised to specific layers. Ultimately, we cannot allocate the contribution of these two processes. Nevertheless, our results suggest that local changes in connectivity can locally alter slow wave activity.

4.3.3 Changes in the slow oscillation following UP state pairings

To support the notion that plasticity can drive SWA homeostasis, plasticity not only needs to affect SWA, it also needs to do so by changing the slow oscillation in a similar manner to sleep. In our experiment, weakening of evoked responses was associated with smaller slow waves with decreased slopes. However, the duration and incidence of slow waves was barely affected. This matches well with the findings from (Vyazovskiy et al. 2007), who did also not detect a change in the overall incidence or duration, but found increased incidence of smaller slow waves and an opposing trend for larger slow waves. In contrast to our study, (Vyazovskiy et al. 2007) found that sleep affected the

first slope of the slow wave more strongly than the second slope. In our study, all amplitudes and slopes of individual slow waves were equally affected and highly correlated. This could be due to a range of possibilities, including differences in the analysis procedure or recording site (superficial in our study and deep layers in (Vyazovskiy et al. 2007)). Interestingly, Vyazovskiy and colleagues find an increase in multipeak waves in late sleep. This could be one explanation of the slight increase in slow wave duration in our study.

It will be interesting to explore how the dynamics of ON and OFF states resemble the dynamics during natural sleep, as observed by (Vyazovskiy et al. 2009). In my opinion, sleep homeostasis, plasticity and perhaps even memory could be elegantly linked via the correlation structure of neuronal activity. At sleep onset, synaptic weights are high, and neurons are strongly correlated, thus producing synchronous ON/OFF states and therefore large slow waves. As plasticity desynchronises the network, ON/OFF states become less synchronised (Vyazovskiy et al. 2009) and slow waves become smaller. Intriguingly, strong correlations are a major obstacle for the storage capacity of a neuronal network.

4.3.4 The effects of anaesthesia

One key caveat of our study is the use of anaesthesia instead of natural sleep. Indeed, the publication of the results by (González-Rueda et al. 2018) were accompanied by a preview (Timofeev & Chauvette 2018), which suggested the results may be due to amnesic function of anaesthetics (Veselis et al. 1997). While this concern can ultimately only be addressed by repeating this experiment during natural sleep, there are some theoretical and practical limitations of the argument by (Timofeev & Chauvette 2018)

Firstly, anterograde amnesia during anaesthesia is not at all in disagreement with natural sleep. In fact, sleep is a profoundly amnesiac state. Intuitively, most everyone knows the frustrating feeling of forgetting a dream's content while one is trying to remember or retell it (Meier et al. 1968). Moreover, sleep walkers can complete goal directed actions and hold entire conversations with no or minimal recall (Zadra et al. 2013). Hence, anterograde anaesthesia is, if anything, a commonality between anaesthesia and sleep.

Secondly retrograde amnesia is not a at all a common by-product of all anaesthetics. Indeed, while Benzodiazepines and Propofol are favoured by surgeons for producing strong anterograde amnesia (Veselis et al. 1997), studies in mice find no retrograde effects of propofol (Pang et al. 1993) or halothane (Rosman et al. 1992). Indeed, Diazepam given to mice after training on a Y-maze task actually enhanced memory consolidation (Cahill et al. 1986). Similarly, propofol had no retrograde effect of memory in adult humans (Bulach et al. 2005) and a meta-analysis found no evidence for a retrograde amnesic effect of benzodiazepines or dexmedetomidine in young adults (Viana et al. 2017). Nevertheless, there is evidence that sedative doses of medetomidine (which we used in this study) can indeed affect memory consolidation (Sirviö et al. 1992). Moreover, I was not able to find studies on the retrograde effects of urethane, which was used by Gonzales-Rueda and colleagues but is a carcinogenic substance and is thus mostly used in acute recordings (Kumar & Clover 2015).

Thirdly, our study and the study by (González-Rueda et al. 2018) used different anaesthetics, which likely act via separate molecular pathways. While we used medetomidine and a small amount of isoflurane, (González-Rueda et al. 2018) used urethane. Medetomidine is an alpha-2 adrenergic agonist, while urethane likely acts on a broader range of targets (Hara & Harris 2002). Nevertheless, we cannot exclude that

the two anaesthetics had an overlapping and LTD-promoting effect. In summary, it the concern about anaesthesia is of course valid. On the other hand, the literature does not at all produce an expectation that anaesthetics would be expected to produce retrograde amnesia or even synaptic weakening.

Chapter 5: Influence of slow waves on evoked responses during natural sleep

5.1 Introduction

In the previous chapters, I explored the interactions between cortical UP and DOWN and states and neuronal plasticity rules in brain slices and under anaesthesia. The results from these chapters suggest that UP states may provide a window of enhanced synaptic weakening, and that localised synaptic weakening may be accompanied by localised decreases in slow wave activity. While anaesthesia and slices are valuable models, we do not have a complete understanding of what sets natural NREM sleep apart from these models. Indeed, it is possible that idiosyncrasies of these models are responsible for the findings in the previous chapters (Timofeev & Chauvette 2018). Therefore, this chapter is dedicated to testing the findings of the previous chapters in naturally sleeping mice.

While the previous chapters focussed on entorhinal and somatosensory cortex, we decided to use callosal connections between the primary motor cortices to study plasticity during natural sleep. This pathway has been shown to exhibit classical long-term potentiation and sleep-dependent plasticity in freely moving rats (Vyazovskiy et al. 2008).

In order to test the differential impact of UP and DOWN states on evoked responses, it is necessary to be able to detect them online while recording. This has been done by various previous studies in humans and mice using the EEG (Besedovsky et al. 2017; Ngo et al. 2013; Massimini et al. 2004a) However, ultimately slow waves are driven by local cortical activity, which is reflected in cortical ON and OFF states in extracellular

recordings (e.g.(Vyazovskiy et al. 2009)). If possible, it thus appears more appropriate to use local spiking activity to detect the phase of the slow oscillation rather than the more global and ambiguous EEG or LFP signal.

In this chapter, we record multi-unit activity and local field potential across the laminae of primary motor cortex of freely moving mice. We use the evoked response to transcallosal electrical stimulation as a proxy of synaptic strength. To test the effect of pairing the stimulation with ON and OFF states in accordance with previous chapters, we first develop a procedure to detect ON and OFF states online. To our knowledge, this is the first study using the spiking activity to achieve close-loop stimulation during NREM sleep.

5.2 Results

5.2.1 Laminar responses to contralateral electrical stimulation

We chronically implanted seven mice with frontal and occipital screws to monitor the EEG and with two wires in the neck muscles to measure the EMG. To establish an indirect readout of synaptic strength, we implanted a 16-channel laminar probe into the primary motor cortex (M1), contralateral to an electrical stimulation electrode (Figure 2, page 50). We chose this callosal M1-M1 pathway because it has been previously used in rats to investigate the interaction between sleep and synaptic plasticity (Vyazovskiy et al. 2008). For each mouse, we established an input-output curve for electrical stimulation at least one day prior to plasticity experiments (see Figure 4 and related methods, page 56). Because the results from the first results chapter suggest that UP states may specifically weaken sub-threshold inputs, we selected the weakest

stimulation level that evoked a response in both the MUA and the LFP (see Figure 3, Figure 4 and related methods).

We first examined the LFP and MUA response to contralateral stimulation across cortical layers and across the baseline period of all experimental conditions. The characteristics of this response can generate insights into the circuitry that the stimulation engages and are crucial to judge the quality of online ON/OFF state detections in the next section. Moreover, the comparison across the experimental conditions is necessary to flag potential systematic differences in the baseline response.

Significant spiking responses to electrical stimulation occurred with an average probability of 51 ± 18% (mean ± SD, n=7 mice with 16 channels each) across layers 1, 2/3, 5 and there was a significant effect of layer on response probability (Table 1). Post-hoc comparison revealed that layer 6 was less likely to respond compared to other layers (Table 1, Figure 19 d). The spiking response generally involved a period of increased firing, followed by a period where spike rates fell below the spontaneous rates. The increased firing rate began on average 3.66 ± 0.93 ms (mean ± SD; n=72 channels[2] from 7 mice, Figure 19 b & c) after stimulation, and started significantly later in layer 1 compared to L5 and L6 (Table 1, Figure 19). This is in line with the visual impression from 4/7 mice, that the response began in deeper layers and then moved upwards (see example in Figure 19 a). Notably, in every experiment there was at least one channel that significantly responded within 1 ms of stimulation (mean time to first responsive time bin in any channel across mice: 1.59 ± 0.69 ms (mean ± SD). This

[2] This is the number of unique channels across all mice that were "responsive" in at least 1/6 conditions. However, not all these channels were responsive (and thus used for analysis) in every condition

could be due to unaccounted stimulation-induced noise or antidromic activation. The spiking response peaked between 5 and 10 ms, and, on average lasted until 10.7 ± 2.13 ms (mean ± SD) after the stimulus. There was an effect of layer on the peak firing rate, with layer 6 responding less than L2/3 and 5.

Measure	Model	Condition effect	Layer effect	Post hoc comparisons						
				Cond.	L1			L2/3		L5
					L2/3	L5	L6	L5	L6	L6
MUA: probability channel responds to stimulation	$\log P(response) = layer + condition +$ (1|channel:mouse)+ (1|mouse) Model: binomial GLME	**$p<10^{-7}$** dF=5 X^2=40.9 test: LLR	**p=0.001** dF=3 X^2=15.3 test: LLR	LTD/LTP < all others	ns	ns	0.09	ns	**	*
MUA: first responsive 1 ms time bin	$t_{1stResp} = layer + condition +$ (1|mouse)+ (1|channel:mouse) Model: LME	p=0.22 dF=5 X^2=6.9 test: LLR	**p=0.011** dF=3 X^2=10.9 test: LLR	-	ns	**	*	ns	ns	ns
MUA time of peak	$t_{peak} = layer + condition +$ (1|mouse)+ (1|channel:mouse) Model: LME	p=0.45 dF=5 X^2=4.7 test: LLR	p=0.25 dF=3 X2=4.03 test: LLR	-	-	-	-	-	-	-
MUA peak firing rate	$Fr_{max} = layer + condition +$ (1|mouse)+ (1|channel:mouse) Model: LME	**p=0.02** dF: 5 X^2: 12.6 test: LLR	**$p<10^{-4}$** dF=3 X^2=23.9 test: LLR	LTP v. OFF	ns	ns	ns	ns	***	***
LFP: probability channel responds to stimulation	*see same for MUA*	**$p<10^{-16}$** dF: 5 X^2: 110.4 test: LLR	**$p<10^{-7}$** dF=3 X^2=37.9 test: LLR	LTD/LTP < all others	ns	ns	**	ns	***	**
LFP: delay to response onset	*see same for MUA*	**p=0.006** dF: 5 X^2: 16.2 test: LLR	**$p<10^{-7}$** dF=3 X^2=38.6 test: LLR	LTP v LTD	ns	***	***	***	***	***
LFP: time of peak	*see same for MUA*	**$p<10^{-5}$** dF=5 X^2=35 test: LLR	**$p<10^{-15}$** dF=3 X^2=95.6 test: LLR	Wake v all others LTP v mock	ns	***	***	***	***	***
LFP: peak amplitude	*see same for MUA*	**$p<10^{-5}$** dF=5 X^2=26.29 test: LLR	**$p<10^{-8}$** dF=3 X^2=44.8 test: LLR	LTD v. all others	**	***	***	***	***	*

*Table 1 Statistics of laminar responses to callosal electrical stimulation. LLR: log likelihood ratio test comparing models with and without the parameter in question; LME: linear mixed effect model. *p<0.05, **p<0.01, ***p<0.001.*

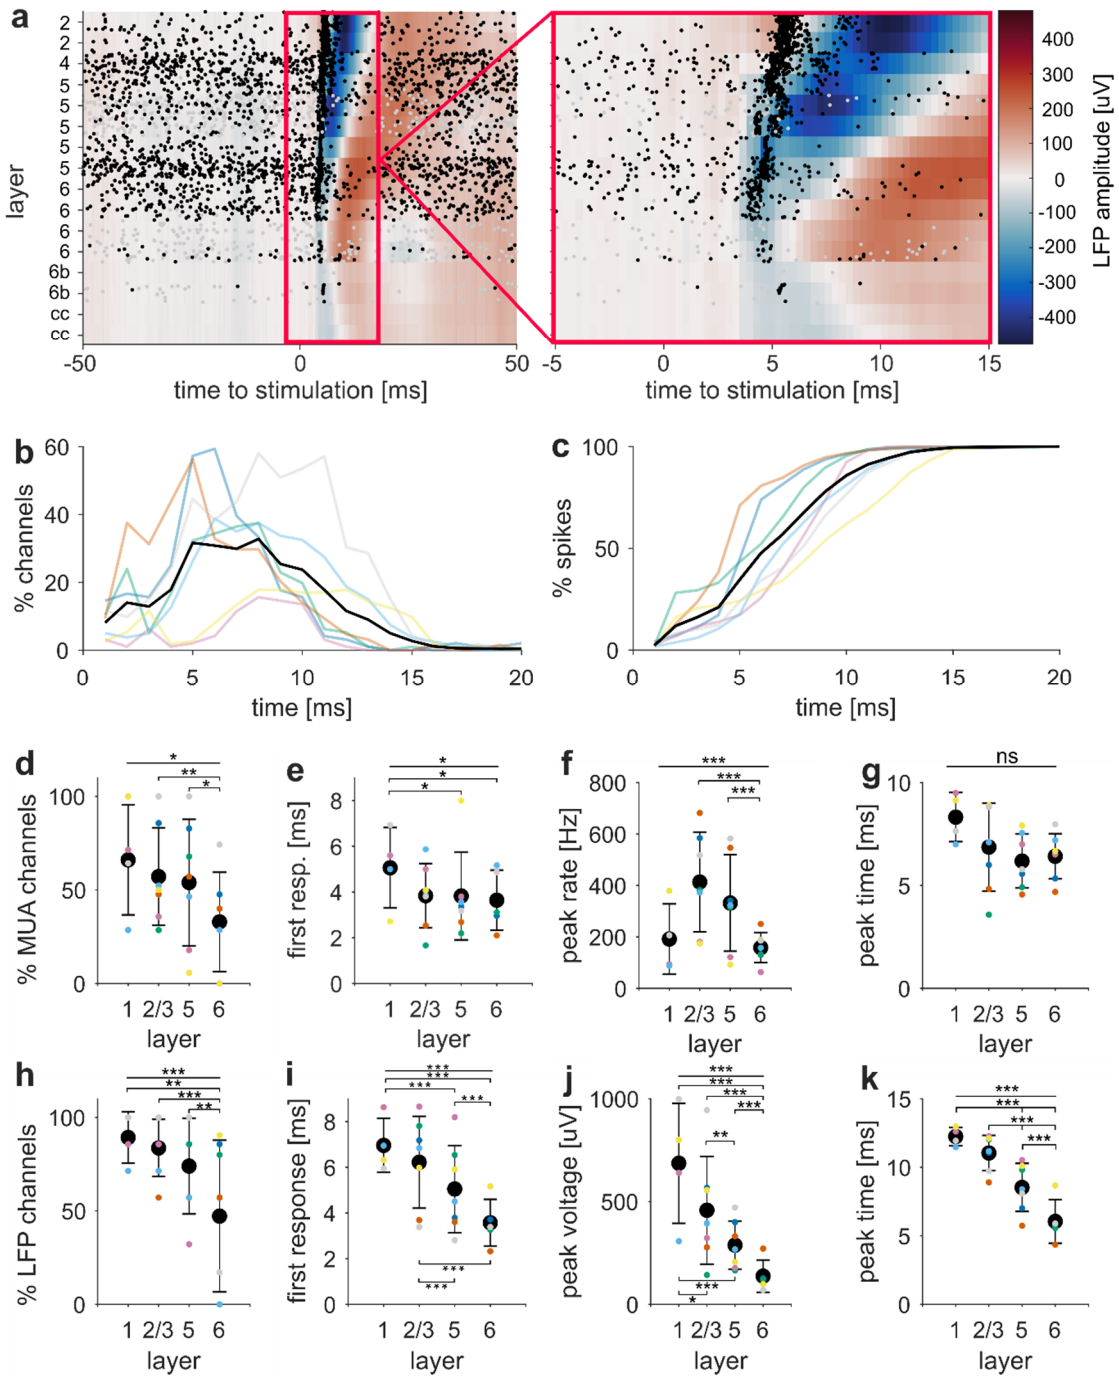


Figure 19 Laminar responses to transcallosal electrical stimulation in the LFP and MUA
a) Example LFP and MUA activity in M1 of one mouse. LFP is the average across all pre-pairing trials and MUA is from all 80 pre-pairing trials (for each channel, the spike times of all trials are concatenated vertically).
b) Distribution of responsive time bins across responsive channels for each mouse (colors) and on average (black)
c) Cumulative sum of spikes occurring during responsive time bins
(d-g): Laminar characteristics of the spiking response: (d) percentage of channels showing significant response to stimulation (e) first time bin with significant response, (f)peak firing rate and (g)time of peak firing rate
*(h-k): Laminar characteristics of the LFP response (d) percentage of channels showing significant response to stimulation, (e) first time bin with significant response, (f) peak voltage response and (g) time of peak response. *p<0.05, **p<0.01, ***p<0.001.*

The LFP response was closely related to the spiking response but appeared to be slightly delayed. Averaged across all responsive channels, the LFP had a negative peak of 314 μV ± 172 at 8.7 ± 1.85 ms after the stimulation (mean ± SD; n=7 mice with 16 channels each). The LFP response began on average 5 ±1.77 ms after the stimulation and lasted for 8.2 ± 1.9 ms (mean ± SD, n= 90 channels[3] from 7 mice), which is consistent with the durations of cortical postsynaptic potentials recorded in single neurons. Compared to the MUA, the LFP onset was on average delayed by 1.3 ±1.4 ms (mean ± SD), even though this could be due to an insufficient signal-to-noise ratio to accurately detect the response onset in the LFP. Similarly, the peak was delayed 5 ± 1.16 ms (mean ± SD) relative to the MUA response. Similarly, the LFP peak was delayed by 2.3 ± 0.83 ms (mean ± SD).

The anatomical layer had a similar but clearer influence on the LFP compared to the MUA. Similar to the MUA, responsiveness of channels declined with depth with layer 6 being significantly less likely to respond than all other layers (Table 1, Figure 19 h). Similarly, deep layers had smaller negative peaks, with almost all layer-wise comparisons confirming this pattern. Despite having small peaks, deeper layers tended to respond and peak considerably earlier than superficial layers.

Taken together, the layer profile of the evoked response in the LFP and MUA are consistent with a scenario where synaptic inputs reach deep layers first and then reach higher layers through cortico-cortical transmission. As can be seen in (Table 1), there was a difference in several characteristic of the baseline evoked response between the

[3] This is the number of unique channels across all mice that were "responsive" in at least 1/6 conditions. However, not all these channels were responsive (and thus used for analysis) in every condition

LTD/LTP experiments and most other experiments. This means that there may have been systematic differences between the LTP/LTD condition and the other experiments. One explanation could be that they were usually conducted last (while the other conditions were randomised). However, it is also reassuring that there is no reason to assume that there are systematic differences among the other experimental conditions.

5.2.2 Closed-loop stimulation during ON and OFF states in freely moving mice

To detect ON and OFF states online, spikes were summed across all channels of the laminar probe (Figure 20 a). OFF states were detected when the running sum of spikes was below a certain threshold (usually below 1 or 2 spikes) for a sufficient amount of time (50-125 ms). An ON state was defined as a period of high firing (10-30 Hz) for a prolonged period of time (same duration as DOWN state), following an OFF state. A challenge for this procedure is the trade-off between speed and accuracy and the trade-off between sensitivity and selectivity. Furthermore, the optimal parameters are not uniform across animals, in part due to different numbers of neurons recorded by each laminar probe. Therefore, we used a baseline recording of each mouse to simulate ON/OFF state detection with differing parameters. As expected, increasing the minimum duration of OFF/ON states leads to detection of larger amplitude slow waves in the LFP but also to fewer detections of ON/OFF states (data not shown)(Vyazovskiy et al. 2009). Furthermore, increasing the minimal duration of OFF/ON states led to an increased chance to detect a state towards its very end.

To validate our closed-loop stimulation approach, we conducted sham-stimulation experiments for each mouse. As shown in (Figure 20 b), OFF state detection was always preceded by a period of neuronal quiescence whereas ON state detections were preceded by increased spiking. Upon detection (i.e. when stimulation would occur), the

probability to transition out of the detected state began to increase exponentially (Figure 20 b-d). Therefore, the success rate of targeting stimulation into a specific state depends exponentially on the delay between stimulation onset and the onset and duration of the neuronal response in the target area. Based on the onset and duration of the LFP response described in the previous chapter, the success rate of targeting OFF states was 69 ± 2.5 % (mean ± SEM) and the success rate of targeting ON states was 89.7 ± 4.16% (mean ± SEM; Figure 20 f) .

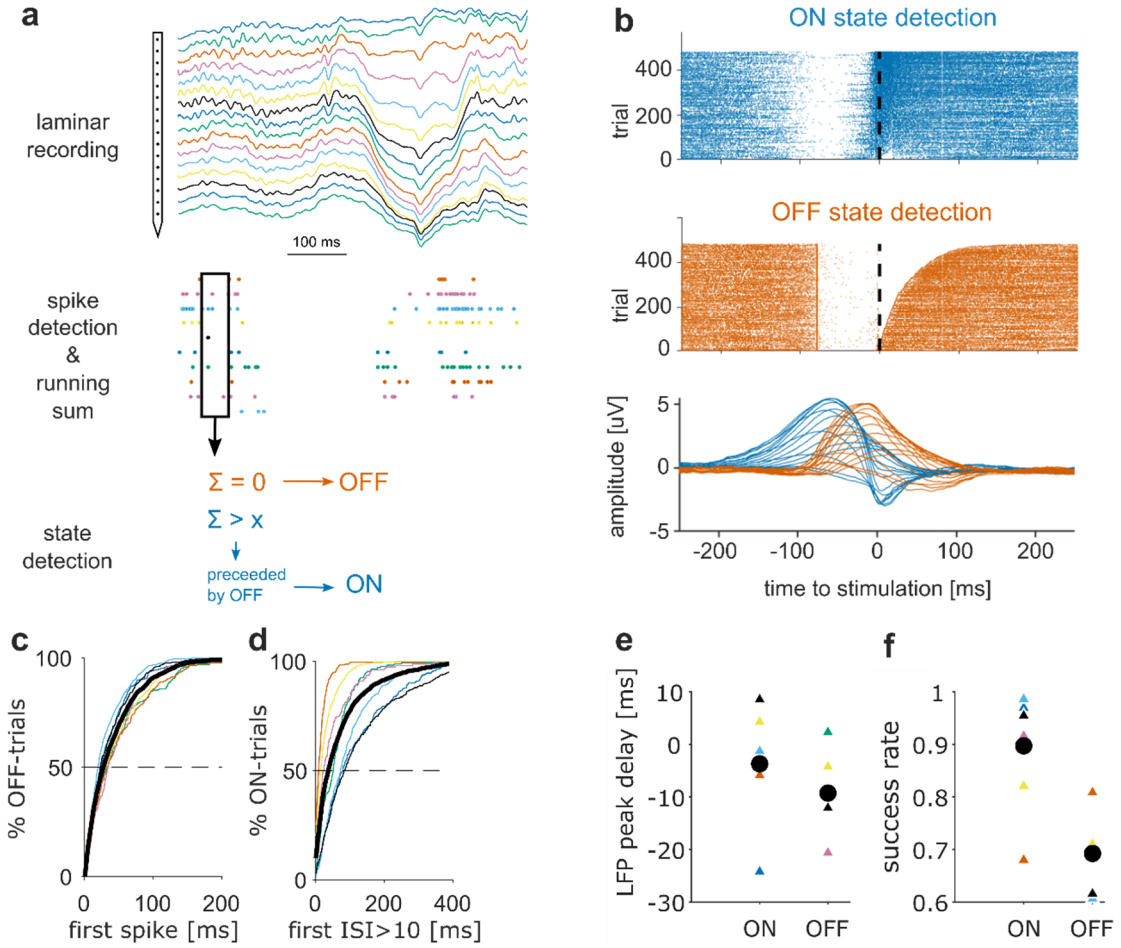


Figure 20 Closed-loop stimulation during ON and OFF states in freely behaving mice
a) Implementation of online ON/OFF state detection algorithm: The LFP is filtered in the spike range (300 – 5000 Hz), thresholded to isolate putative spikes, and summed across channels in a sliding window (50 – 100ms). In this example, no stimulation occurred (i.e. trace taken from the mock-condition).
b) Example mock stimulation traces from one animal. The approximately 400 detected ON states are sorted according to the spike rate in the period where the LFP response to electrical stimulation would occur. The OFF states are sorted according to the time to the first spike. Bottom panel shows average evoked LFP.
c&d) Cumulative sum of trials no longer conforming to the criteria for OFF/ON states relative to time of state detection (colors are individual mice, black line is mean) .C: % of OFF state detection trials having displayed at least one spike. D: % ON state detection trials with at least one inter-spike interval > 10ms.
e) Delay of the LFP peak relative to state detection. Black markers/lines represent mean and colours represent individual mice
f) Success rate of state detections based on response onset and duration from the previous section.

Another approach to validate ON/OFF state targeting is to examine the magnitude of the evoked response (Ngo et al. 2013). The prediction from several previous studies in anaesthetised animals and brain slices is that the magnitude of the response should be significantly modulated by ON/OFF states (Reig et al. 2015; Haider et al. 2007). In line with this, we find that the stimulation triggered increase in spiking (i.e. relative to pre-stimulation baseline) is larger during OFF state pairings compared to ON state pairings ($p= 0.0024$, paired t-test, n=7; Figure 21 a-g,i). However, when the baseline is not subtracted the opposite trend emerges, with responses during ON-states displaying a larger absolute peak ($p = 0.0676$, paired t-test, n=7; Figure 21 h&j). Exactly the same pattern emerges in the LFP response: when the baseline difference at stimulation onset is accounted for, the response is larger during OFF state stimulation. If it is not, there is no longer any evidence for a difference (Figure 22). Taken together, these results suggest that, on average, distinct states of the neuronal network are targeted in all mice.

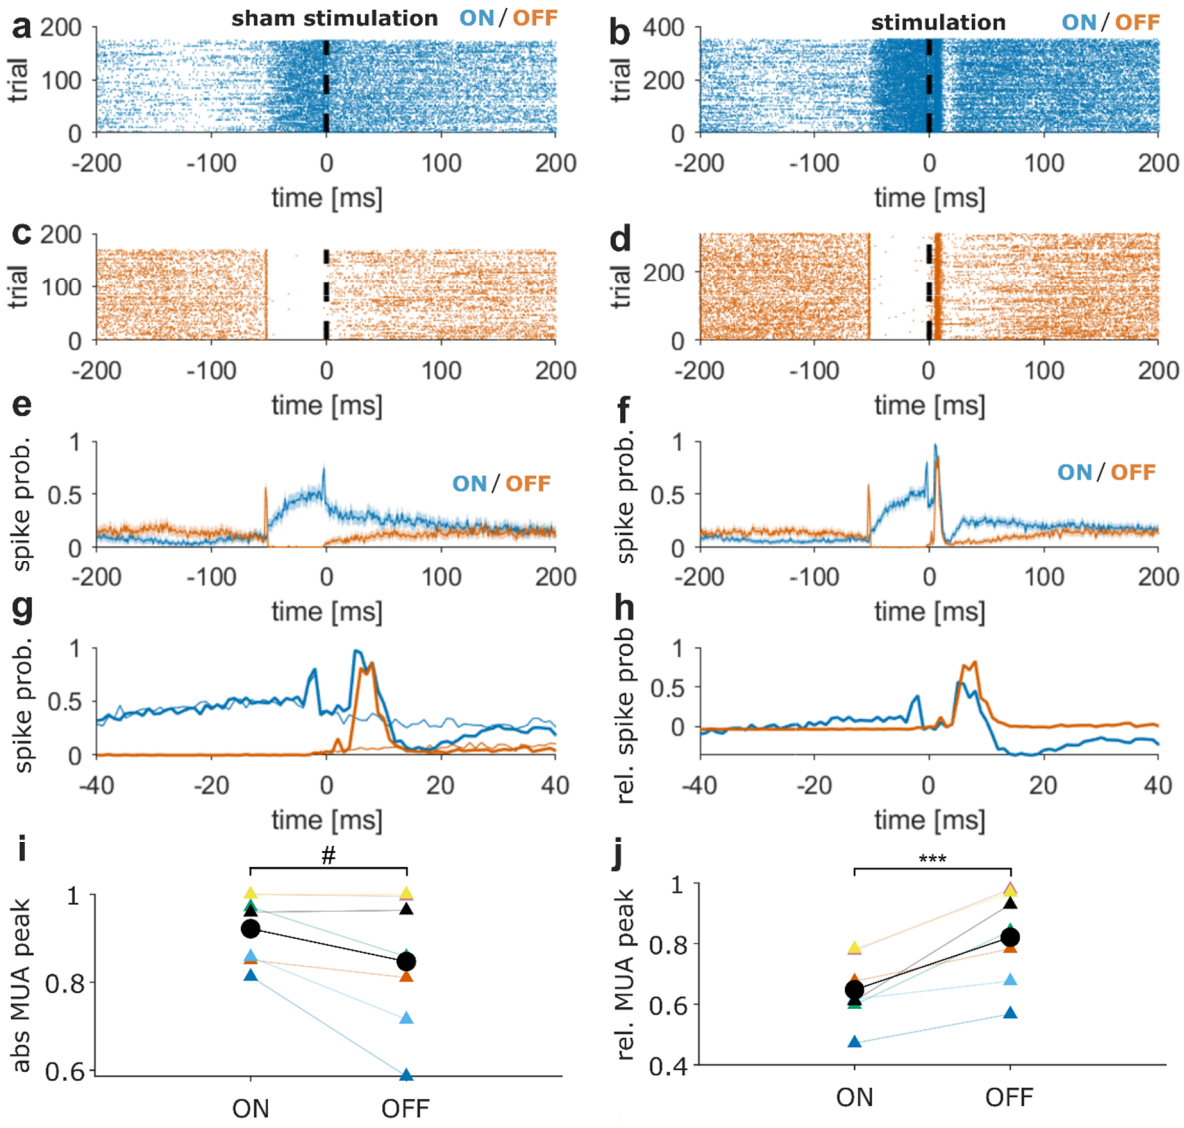


Figure 21 Spiking responses to ON/OFF stimulation compared to mock stimulation
a-f) Example responses to sham stimulation (a,c,e) and electrical stimulation (b,d,f) in one mouse. a-d: each line shows summed spikes across all channels in one trial. e&f show mean ± SD across trials. Note the slightly higher number of trials in the electrical stimulation condition (b,d,f) leads to denser visualisation.
g&h) Stimulation-evoked responses with different corrections for spontaneous firing rates: without normalisation (g), ON state stimulation reaches higher firing rates than OFF states stimulation whereas when the pre-stimulation baseline is subtracted (h), OFF state stimulation reaches higher firing rates. Thin lines in g represent mock stimulation spike rate.
*i&j) Quantification of unnormalized (i) and normalised (j) stimulation-evoked spike rates. #: 0.05 > p < 0.06, ***p<0.001, paired t-test.*

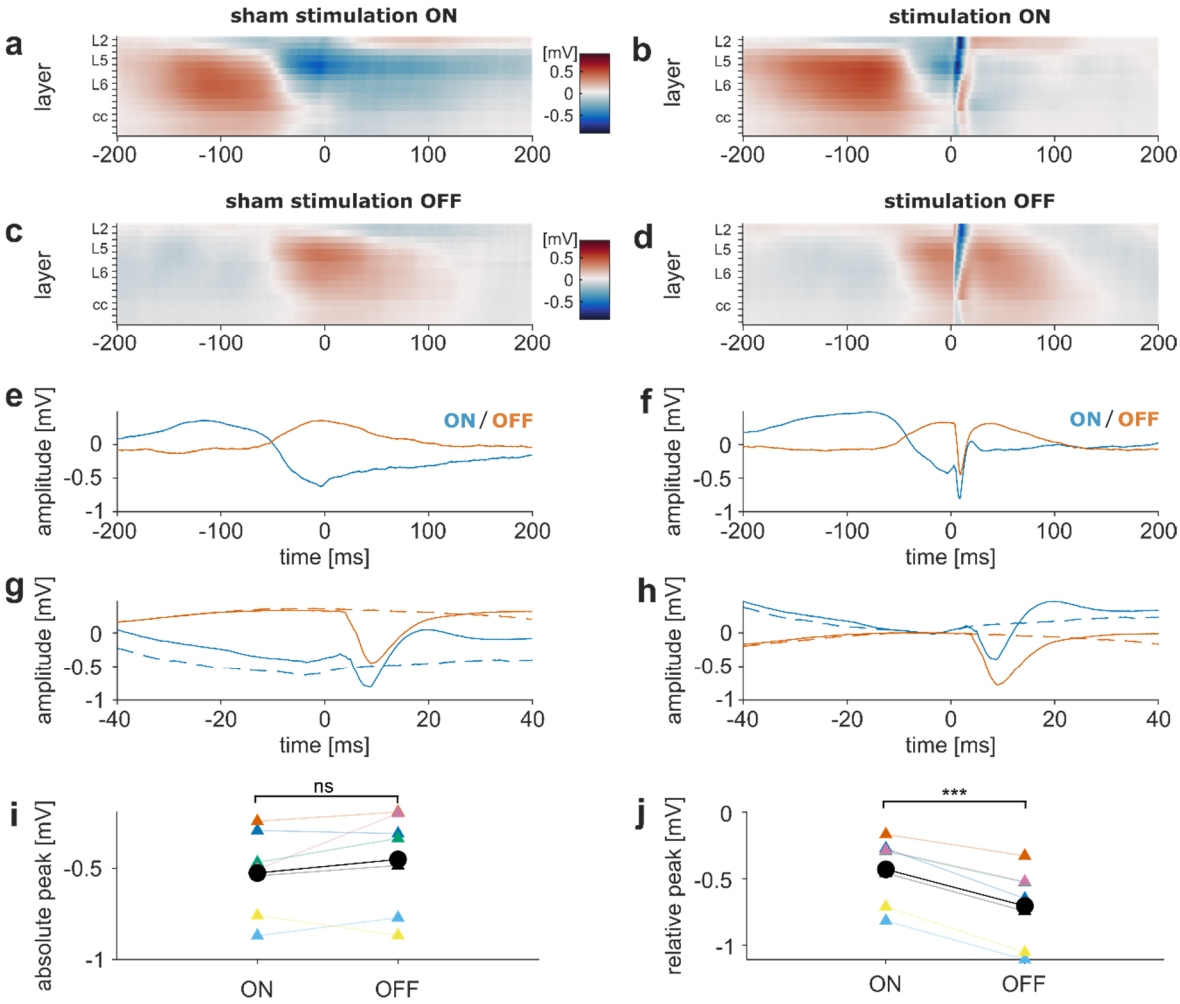


Figure 22 LFP responses to ON/OFF stimulation compared to mock stimulation
a-f) Example average responses to sham (a,c) and electrical (b,d) stimulation in one mouse. Each line is one channel and traces are aligned to stimulation onset. e&f show traces of topmost channel in mock and electrical stimulation example, respectively.
g&h) Effect of normalisation on stimulation-evoked LFP response from panel f: without normalisation (g) OFF/ON state stimulation occurs on a background of positive/negative LFP amplitudes. Subtracting this pre-stimulation offset (h) reveals clear differences in the amount evoked response.
*i&j) Quantification of unnormalized (i) and normalised (j) mean stimulation-evoked peak amplitudes across mice. *** $p < 0.001$*

5.2.3 Effects of ON and OFF state pairings on evoked responses during waking

We wanted to address the hypothesis that pairing an input to cortex with ON states weakens this input, whereas pairings with OFF states or waking activity has a weaker or opposing effect. Therefore, we recorded neuronal responses to contralateral electrical stimulation in awake mice exploring objects. We delivered 80 stimulations (0.01 Hz) before and after each of the following 7 different pairing-protocol shown in (Figure 23 a): Stimulation during ON states, stimulation during OFF states, sham stimulation (stimulation turned off), or during waking (novel objects were given to promote wakefulness when necessary). As a positive control, we also applied classical long-term potentiation/depression protocols.

The effect of the seven different pairing conditions (referred to as "condition") on the change in LFP peak amplitude from pre- to post pairing wakefulness (Figure 23 b) was assessed with linear mixed-effects models of the form[4]: $\Delta V = condition + V_{baseline} + condition * V_{baseline} + (1|\text{channel:mouse}) + (1|\text{mouse})$. The model supported a significant effect of condition on the change in LFP peak amplitude. As shown in (Figure 23 c & d), this was true for both the absolute[5] change (e.g. V_{post}/V_{pre}) and the relative change (e.g. $V_{post} - V_{pre}$). Moreover, the relative change in peak amplitude was not only affected by condition but also negatively associated with the baseline (pre-

[4] The terms in the brackets represent random effects, which account for the nested dependencies in the data. I.e. multiple channels per animal and repeated measurements per mouse.

[5] When predicting the change in the LFP in absolute terms, the pre-pairing response has to be included as a predictor, resulting in a model of this form: Δ(response) = pre-pairing + condition + condition*pre-pairing. Such interaction terms make the interpretation of results considerably more difficult and I thus focused on the relative changes. However, this model also had significant effects of condition (log-likelihood ratio test, $\chi^2(5)=37.53$, $p=10^{-7}$) and a significant interaction term (log-likelihood ratio test, $\chi^2(5)=57.53$, $p=10^{-11}$)

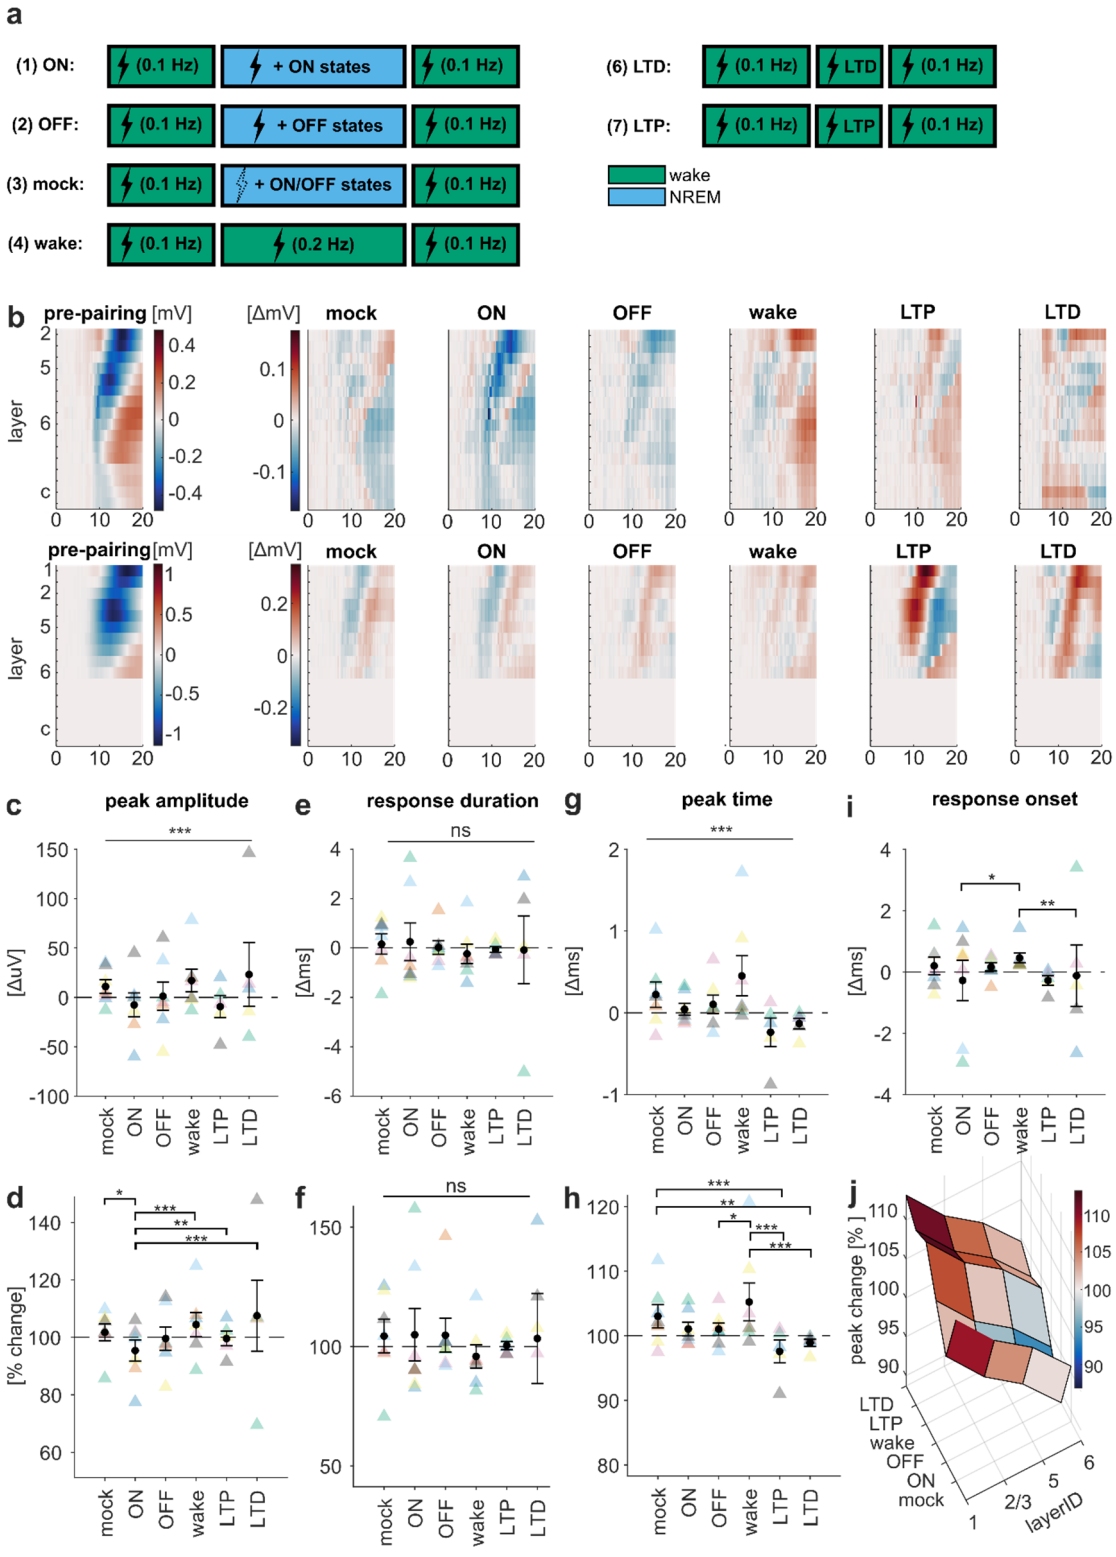


Figure 23 Effects of pairing protocols on evoked LFP responses
a) Experimental design: each mouse was subjected to conditions on separate day in random order. Baseline and post-pairing stimulation consisted of 80 stimulations at 0.1 Hz.
b) 2 examples of pre-pairing evoked responses and changes following the different pairing protocols.
c-f) Absolute changes in several measures of the LFP response across conditions.
d-h) Changes in LFP response measures normalised to the baseline measurement.
j) Estimated marginal means depicting the interaction of layer and condition based on a linear mixed model.
*Significance bars spanning all conditions indicate significant main effect of condition. Individual bars represent pairwise comparisons. *p<0.05, **p<0.01, ***p< 0.001*

pairing) amplitude (log-likelihood ratio test, $\chi^2(5)=25.78$, $p<10^{-4}$ for condition, and $\chi^2(1)=16.14$, $p<10^{-4}$ for baseline amplitude , $\beta_{(pre\text{-}pairing)} = -0.015 \pm 0.003$). There was no evidence for a significant interaction between baseline amplitude and condition. This suggests that the condition influences the change in evoked response and that larger responses tend to decrease more after pairing, independent of the condition. Pairwise post-hoc comparisons suggested that ON pairings negatively influence the change in amplitude compared to all other pairings save OFF state pairings (Tukey-adjusted contrasts for difference in β values ± SE : ON-mock: -7.2 % ± 2.4, p=0.03; ON-wake: -10.39 % ± 2.4, p<0.001; ON-LTP: -9.4 % ± 2.7, p:0.0073, ON-LTD: -10.92 % ± 2.71, p<0.001, ON-OFF: -4.71 % ± 2.4 , p=0.37). While wake pairings had a trend towards increasing the response, this was not significant. There was no effect of condition on the duration of the LFP response (Figure 23 e&f), but the peak time and the response onset were both affected (Figure 23 g-i).

To explore the influence of layers on the change in LFP amplitude (Figure 23 j), we included the histological layer assignments in the linear mixed model ($\Delta V = condition + V_{baseline} + layer + (1|\text{channel:mouse}) + (1|\text{mouse})$). This model supported an influence of layer (log-likelihood ratio test, $\chi^2(3)=10.39$, p=0.015) on the relative change in amplitude, but not the absolute change (log-likelihood ratio test, $\chi^2(3)=5.47$, p=0.14). Post hoc testing with the former model suggested that layer 6 showed a significantly more negative change in amplitude across conditions compared to layers 2/3 and 1 (Tukey-adjusted contrasts for difference in β values ± SE : L6-L2/3: -8% ± 2.8, p=0.019; L6-L1: -13.6 ± 4.6, p=0.010; L5-L1: -8 ± 3.9, p=0.08). This suggests that when the effect of condition is to increase the peak (e.g. wake-pairings), deeper layers show less of an increase (Figure 23 j). When the condition decreases the peak, like with the ON pairings, deeper layers show more of a decrease. However, this

result seems somewhat suspicious given the strong effect of layer on the pre-paring response (Figure 19) and the negative effect of pre-paring response on the change in amplitude. It seems possible that the model cannot fully account for the effect of pre-paring response (or that this effect is not entirely linear), and that the effect of layer is simply a remnant from its strong correlation with pre-paring amplitude.

In the spiking response, we first examined the difference in the mean number of spikes across all 1 ms time bins with a significant excitatory response to electrical stimulation (Figure 24 a&b, methods Figure 6, page 64). At first, there seemed to be a trend towards an interaction between condition and pre-pairing spiking (Figure 24 a&b, p=0.054, log likelihood ratio test, $\chi^2(5)=10.8$). However, plasticity has been reported to affect spontaneous spike rates (Kimura & Pavlides 2000), which would confound this measure. Indeed, a linear mixed model identified a clear effect of condition on the absolute (post-pre) and relative (post/pre) difference in spontaneous spike rates (Figure 24 c, $\chi^2(5)=11.371$, p=0.045, main effect of condition on relative difference; $\chi^2(5)=29.976$, $p<10^{-5}$ for the absolute difference). Pairwise post-hoc tests suggested that ON pairings increased spontaneous firing rates compared to mock and wake pairings.

To compensate for the changes in baseline spontaneous rates, we subtracted the mean spontaneous firing rate from future measures of mean evoked firing rates (Figure 24 d-f). Using this new measure, the absolute difference between pre- and post- pairing response was significantly influenced by an interaction of condition with pre-pairing response magnitude (Figure 24 d, $\chi^2(5) =16.792$, p=0.005). While none of the post-hoc comparisons reached significance, examination of the effect sizes suggests that the difference between LTP pairings and mock pairings, as well as the difference between LTP pairings and wake pairings were the main conditions that drove the significant

effect. However, even when every condition was tested separately against a zero-effect null hypothesis using a t-test, none of the conditions returned a significant effect. Moreover, like the mean number of spikes, there was no clear effect of condition on the peak firing rate or the timing of the peak firing rate (Figure 24 d&e).

In summary, these data support a weakening effect of ON state pairings on evoked responses during wakefulness in the LFP but not in the MUA. While I will discuss these results in more detail in the discussion section (page 132, section 5.5.3), one clear concern is that these data were gathered in freely moving mice, with limited control over behavioural state.

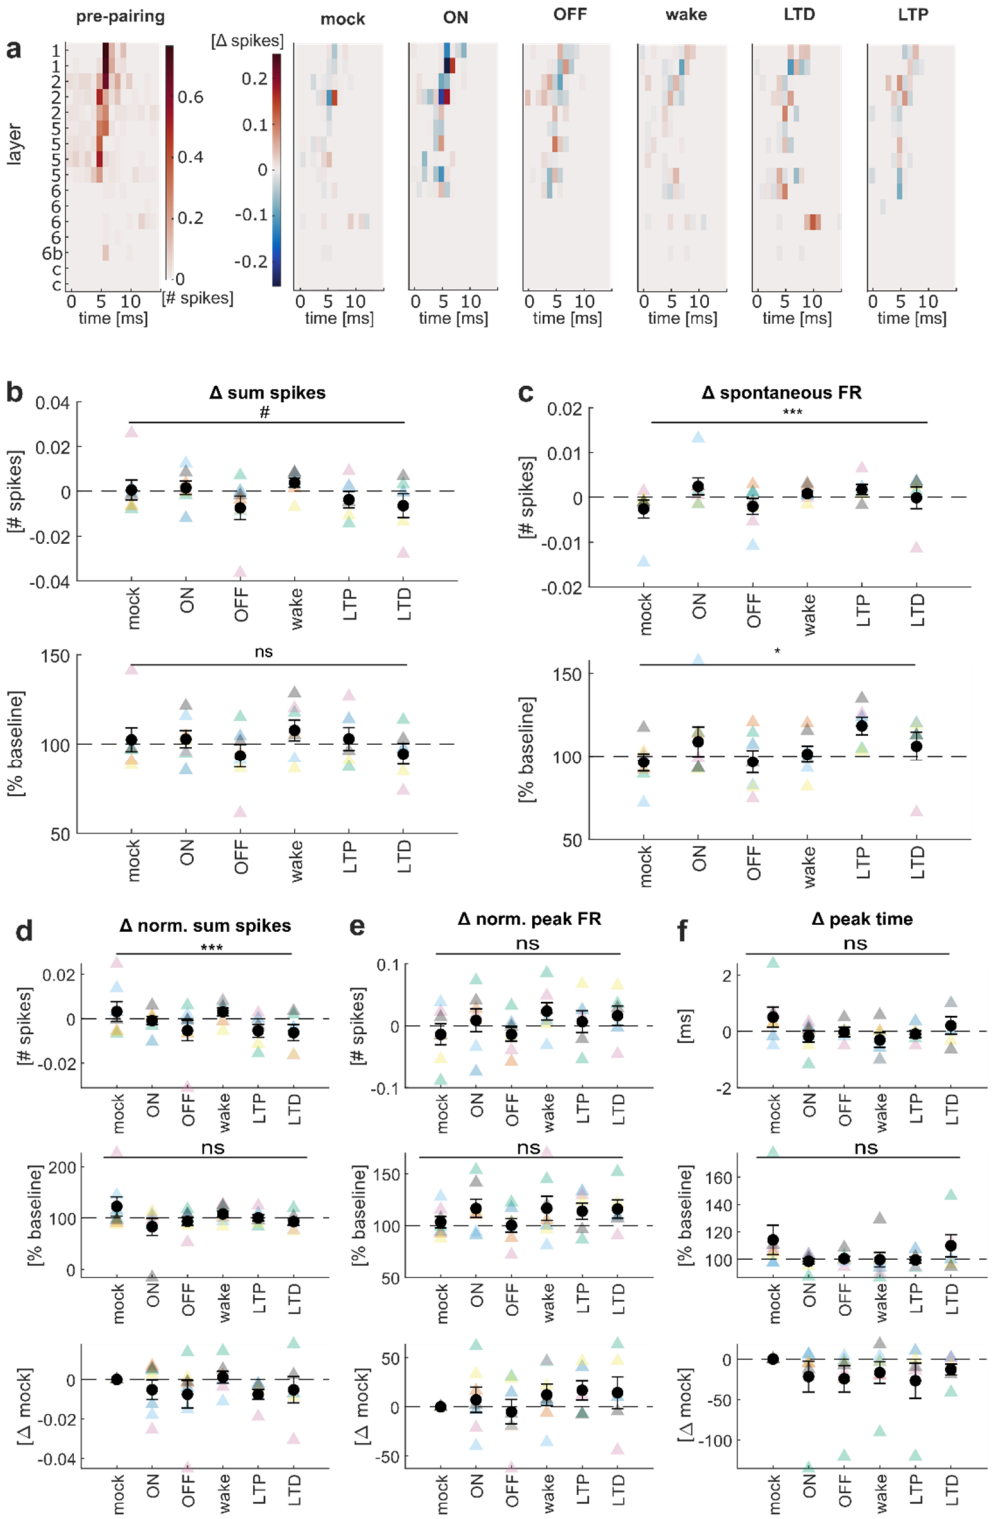


Figure 24 Effects of pairing protocols on evoked MUA responses
a) Example of mean pre-pairing evoked MUA responses and changes following the different pairing protocols. Spike counts are averages across trials per 1 ms bin. Only responsive time bins are shown.
b&c) Changes in evoked and spontaneous firing following the pairing protocols. Black markers represent mean ± SEM. Top row: absolute change (post/pre), bottom row: relative change (post/pre).
d-f) Changes in evoked firing corrected for spontaneous firing rates.
Top row: absolute difference in corrected evoked firing rate (post – pre), middle row: relative difference in corrected evoked firing (post/pre), bottom row: absolute change in each condition normalised to absolute change in the mock condition. (just for visualisation, makes no difference in statistics)
*Significance bars spanning all conditions indicate significant main effect of condition. *p<0.05, **p<0.01, ***p< 0.001; #: 0.05 > p < 0.06*

5.2.4 Effects of ON/OFF state pairings on evoked responses during NREM sleep

A major limitation of the above results is that they resulted from comparing stimulations during waking, with limited control over what the animal is doing when stimulation occurs. Therefore, it is likely that a significant portion of variance in the response strength is attributable to behavioural state rather than neuronal plasticity. For example, sensory responses during quiet waking have previously been reported to be larger than responses during active sensory exploration (Crochet & Petersen 2006). One readily available approach to shed light on this issue is to look at changes in evoked responses during NREM sleep. Therefore, we analysed the evoked responses across the pairing period in all conditions where the animal was allowed to sleep (i.e. ON, OFF and mock pairings). All responses to stimulation during artefact-free NREM sleep were binned into 10 equally sized bins across the pairing period (Figure 25 a). The mock condition was included because OFF and ON states should change from the beginning to the end of sleep even without stimulation (Vyazovskiy et al. 2009). It is thus important to ensure that changes in the evoked responses are not (trivially) confounded by these stimulation-independent trends.

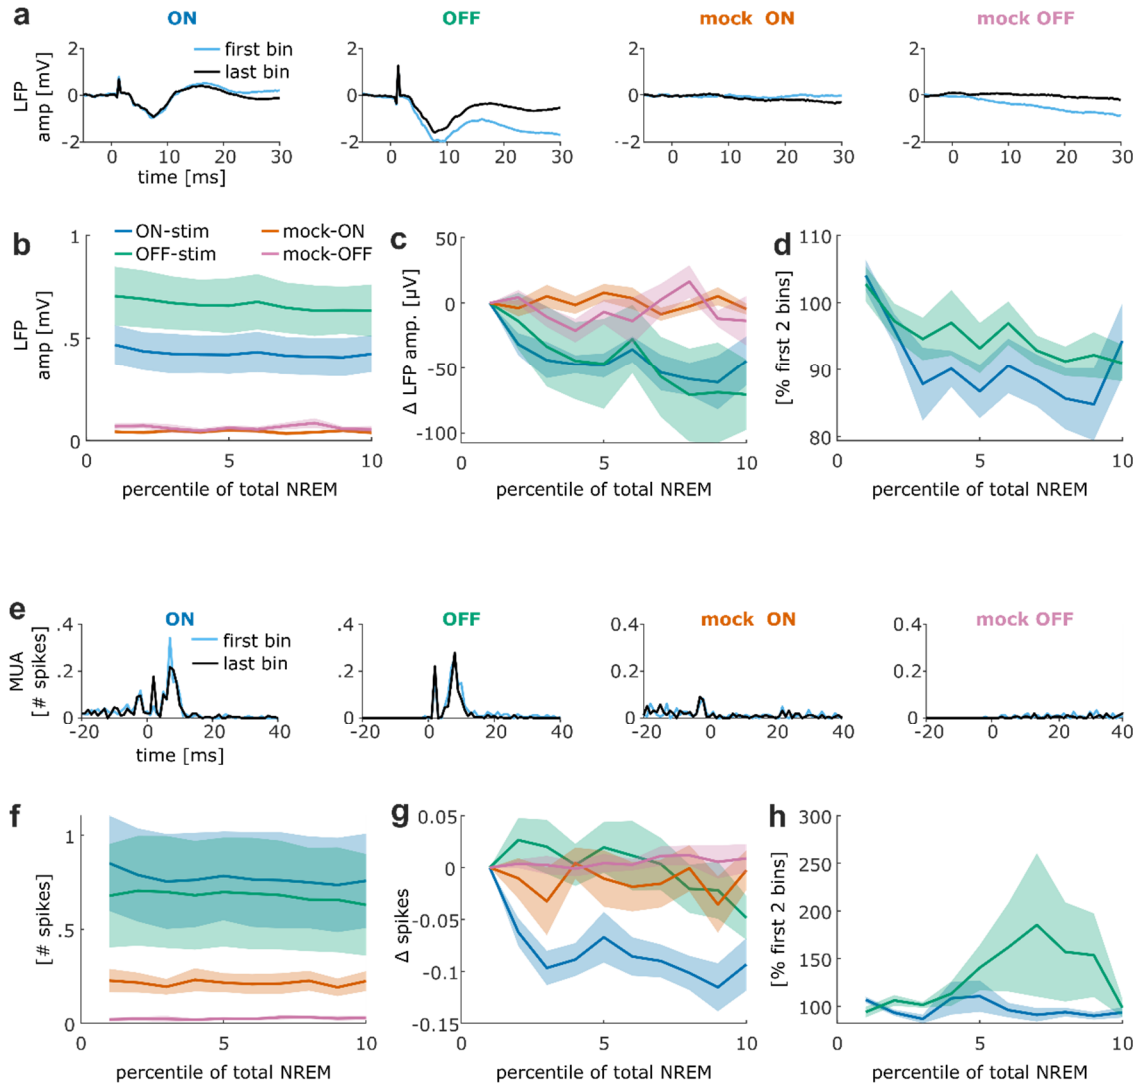


Figure 25 Plasticity in the LFP and MUA during NREM pairings
a) Example mean LFP traces of one animal. Blue trace: mean across trials in first 10% of NREM sleep time. Black traces: mean across last 10%. Mock ON/OFF are shown to demonstrate stimulation-independent effects. Note that mock-ON and mock-OFF are derived from the same recording.
b-d) Quantification of LFP peak across NREM sleep with different normalisations. b: unnormalised peaks displaying expected differences. c: absolute difference (subtractive) to first time bin (separately for each condition) d: ratio to first time bin (divisive). Mock conditions are not shown as they are very noisy (small baselines lead to noisy traces in divisive measures).
e-h) Same as a-d but for MUA (# spikes reflects sum of spikes within 20 ms of stim).
For statistics, see text and table below.

Measure	Model	Model assumptions	Time x condition	condition	time	Post-hoc: Tukey Contrasts (adjusted using Holm's method)
LFP: absolute amplitude	*Amp* = condition + time + condition*time + (1|mouseIDs)+(1|channel:mouse) model type: LME	OK	p>0.99 χ^2=6.2 ΔdF=27 test: LLR	p<10^{-16} χ^2=2374.9 ΔdF=3 test: LLR	p=0.83 χ^2=5.00 ΔdF=9 test: LLR	ON v. OFF: *** ON v. mON: *** ON v. mOFF: *** OFF v. mON: *** OFF v. mOFF: *** mOFF v. mON: 0.192
LFP: absolute change in amplitude rel. to first time bin	*Δ Amp* = condition + time + condition*time + (1|mouseIDs) + (1|channel:mouse) model type: LME note: t1=0 by definition; thus t1 is not included in model (otherwise homogeneity of variance violated)	OK	p<10^{-11} χ^2=106.95 ΔdF=24 test: LLR	p<10^{-16} F=234.56 dF=3 dFe=2975 test: type-3 ANOVA	p<10^{-16} F=17.01 dF=8 dFe=2947 test: type-3 ANOVA	Sig. contrasts for conditions: ON v. OFF: t10 ON v. mON: >t1 ON v. mOFF: >t1 OFF v. mON: >t1 OFF v. mOFF: >t1 mOFF v. mON: ns Sig. contrasts for time: ON: 2 v 8-9, 3 v 9, 6 v 9 OFF: 2 v 4-5, 7-10; t3 v 8-10; t6 v 8-10 testing t2 = 0 ON: p<0.01; OFF: p<0.01
LFP: relative change in amplitude rel. to first 2 time bins	*Δ Amp* = condition + time + condition*time + (1|mouseIDs) + (1|channel:mouse) model type: LME	OK	p<0.005 χ^2=23.1 ΔdF=9 test: LLR	p<10^{-16} F=1648.74 dF=2 dFe=22.15 test: type-3 ANOVA	p<10^{-16} F=26.85 dF=9 dFe=1611.4 test: type-3 ANOVA	Sig. contrasts by conditions: ON v. OFF: t 3,4,5,6,8,9 Sig. contrasts by time bin: ON: t1 v all; 2 v 3-9; 10 v 5-9 OFF: t1 v. v all
MUA: sum of spikes within 20 ms after stimulation	*(Σ spikes)* = condition + time + condition*time + (1|mouseIDs) + (1|channel:mouse) model type: LME	Not met	-	-	-	-
MUA: absolute change in the sum of spikes within 20 ms after stimulation	*Δ(Σ spikes)* = condition + time + condition*time + (1|mouseIDs) + (1|channel:mouse) model type: LME as above t1=0 by definition, thus excluded from model	OK	p=0.21 χ^2=29.1 ΔdF=24 test: LLR	p<10^{-16} χ^2=196.04 ΔdF=3 test: LLR	p=0.02 χ^2=17.99 ΔdF=8 test: LLR	Contrasts: conditions: ON vs all others: *** OFF v. mON: p=0.06 OFF v. mOFF: ns mOFF v. mON: *** sig contrasts for time bin 2 v 9, 5 v 8 all time bins are <0 with p<0.01
MUA: relative change in evoked spikes	*As above*	Not met	-	-	-	-

Table 2 Statistics for effects of stimulation on evoked responses during NREM sleep
*LLR: log likelihood ratio test comparing models with and without the parameter in question; LME: linear mixed effect model. ***: p<0.001*

We analysed the evoked response during NREM sleep using three different ways of normalisation: Without normalisation, the evoked responses in the MUA and LFP displayed the expected differences between conditions but no interaction between

condition and time or main effects of time (Figure 25 b&f, Table 2). The second normalisation we applied was to calculate the difference to the first time bin for each condition (Figure 25 c&g). This analysis revealed a significant interaction between condition and time in the LFP. Post-hoc contrasts between conditions suggested that at all time bins the peak amplitude in the ON and OFF state stimulation was significantly more decreased than in the two mock conditions (Table 2). However, there was no meaningful difference between ON and OFF state stimulation (Table 2). To further explore the temporal dynamics of the decline in evoked response in the ON and OFF state condition, we computed contrasts between the time bins in these conditions (Table 2). This analysis suggested that several time bins significantly differed in both conditions, which suggests that the change occurs gradually, rather than all at once. In the MUA, this "subtractive normalisation" revealed only simple effects of time and condition, without an interaction (Table 2). Interestingly, post-hoc tests suggested that the ON state condition had a stronger decline in evoked responses than the OFF state condition. However, mock-ON also decreased more than mock-OFF, even though on a much smaller scale. Nevertheless, this does raise the possibility that the observed changes in the ON state condition could be due to complex interactions between the slow oscillation, the detection algorithm and the stimulation. This concern might be substantiated by the visual impression that the majority of change in the ON condition occurs from the first to the second time bin.

Finally, we used divisive normalisation to the mean of the first two time bins (Figure 25 d & h). For this analysis we excluded the mock conditions because division by very small values leads to noisy results. Interestingly, in the LFP this analysis suggested that ON state pairings lead to a more reduced response than OFF state pairings in all but the last time bin (Table 2). In the MUA, visual inspection suggested a similar trend but the

small number of spikes in the OFF states led to a distribution of data points that would have required very advanced statistics. Given that subtractive normalisation for the MUA suggested that this analysis is confounded by changes in slow waves, we did not conduct these statistics.

In summary, the development of evoked responses across NREM sleep supports the idea that stimulation changes the evoked responses in the LFP but cannot really shed further light on the MUA. Moreover, divisive normalisation suggests that changes during ON state pairings may be larger than during OFF state pairing, while subtractive normalisation does not support this. Importantly, this analysis has its own caveats, most notably that ON and OFF state pairings, *by definition,* have a different baseline response size.

5.2.5 Effects of electrical stimulation on sleep architecture and slow wave activity

We next sought to investigate whether and how stimulation affects sleep. This is relevant from two perspectives: Firstly, if stimulation during ON and/or OFF states has an immediate effect on sleep (e.g. waking the animal up), then this would be a confound for interpreting the effect of stimulation on plasticity. Secondly, the previous chapter's results suggest that changes in evoked responses may be associated with changes in slow wave activity.

The most straightforward way by which any manipulation can affect sleep is by changing sleep architecture, i.e. the distribution of time the animal spends in each vigilance state. A shown in (Figure 26 a), there was no evidence for an effect of stimulation on the relative time spent in NREM sleep, REM sleep or awake (n = 7 mice,

3 separate repeated measures ANOVA[6], effect of pairing condition on % NREM with sphericity assumed: p=0.518, F(2,12)=0.695; pairing v wake: p=0.251, F(2,12)=1.555, for pairing v REM Friedman's test was used, p=0.180; RM-ANOVA was also non-significant but outliers forced non-parametric testing).

To test whether stimulation has immediate or delayed effects on the progression of SWA across sleep, we assessed the time course of SWA in the EEG and LFP using only epochs not directly containing a stimulation event (Figure 26 c &d). As expected, homeostatic decline of SWA resulted in a significant main effect of time on SWA in linear mixed models[7] run separately for the EEG and the average SWA in the LFP (LLR, $\chi^2(9)=112.9$ and 147.9 for EEG and LFP, $p<10^{-16}$ for both, n=7 mice). In addition, the parietal EEG displayed a significant effect of condition on SWA ($\chi^2(2)=21.62$, $p=10^{-5}$), whereas there was no such effect of condition on SWA in the LFP ($\chi^2(2)=2.7$, p=0.25). There was no evidence for an interaction between condition and time in the EEG/LFP (LLR, $\chi^2(18)=16/12$, p=0.59/0.8). Post-hoc comparisons in the EEG suggested that OFF state pairings were associated with significantly reduced SWA, compared to ON state pairings and mock pairings (Tukey contrasts, $P<10^{-5}$ for both comparisons in the EEG). This suggests that, independent of when the measurement was taken, the OFF state pairing condition always displays lower SWA. This contrasts slightly with the visual impression that the first- and last-time bin are not different and is likely due to insufficient power. Indeed, if the same analysis is repeated

[6] A repeated measures MANOVA might be more adequate but several of its assumptions (outliers, collinearity) were violated, thus strongly reducing its power. The test also confirms a non-significant effect (Wilk's lambda = 0.197). Note that conducting multiple RM-ANOVA's should increase false-positives, thus biasing against the reported results of no effect.

[7] $SWA_{\%\,mock} = condition * time + condition + time + (1|\text{mouse})$.

in the LFP without averaging across channels (i.e. a linear mixed model that accounts for the nested nature of this analysis), there is a significant interaction between condition and time (data not shown). When analysed like this, the significant difference between OFF states and the other conditions emerges in the second time bin but disappears for the last bin. All other conclusions remain identical. This suggests that OFF state pairings are associated with lower SWA activity and that this difference is not developing slowly over time but is either present at all times or emerges after the first 10 % of NREM sleep.

The slow oscillation is not the only network phenomenon during natural NREM sleep, and other events, such as spindles, have been associated with plasticity. To test the effect of stimulation on other frequencies we calculated the difference between the average power spectra (again including only 4 s episodes without stimulation) across conditions (Figure 26 e). As expected from the previous results, there was a significant interaction between condition and frequency in the LFP ($\chi^2(1)=36.9$, $p=10^{-9}$, log-likelihood ratio test) and in the frontal EEG ($\chi^2(79)=108.17$, $p=0.016$, log-likelihood ratio test). Post-hoc test on individual 0.25 Hz frequency bins suggested that EEG power in the frequencies between 0.25 and 2.5 Hz were smaller in OFF compared to ON state pairings and not significantly different in other bands (pairwise contrasts without correction for multiple comparison, bins[8] with $p <= 0.1$: 0 Hz, 0.75-2 Hz, remainder is $0.01>p<0.05$, estimated differences ranged from 7 ± 4.11 % (2.5 Hz) to 12 ± 4.11 % (1.25 Hz)). In the LFP (Figure 26 e, bottom) there was also evidence that ON and OFF pairing had differential effects on power spectra (interaction between

[8] Bins have a width of 0.25 Hz, the reported frequencies represent the lower limit. E.g. 0 Hz means 0-0.25 Hz.

condition and frequency: $\chi^2(79)=685.8$, $p=10^{-16}$, log-likelihood ratio test, LME with channels and mouse as nested random effects). Post-hoc tests for individual frequency bins indicated that sleep during ON state pairings had more power compared to sleep during OFF state pairings in low frequency bands (0-2 Hz, $p<0.001$ for all bins except 1.5 Hz with $p=0.005$), but it had lower power in several higher frequency bands (pairwise contrasts without correction for multiple comparison, 5.25 – 5.75 Hz, $p<0.01$ for all, 7 - 8.25 Hz, $p<0.05$, 9.25 – 9.75 Hz, $p<0.05$, 10.5 – 20 Hz, $p<0.01$ for most bands).

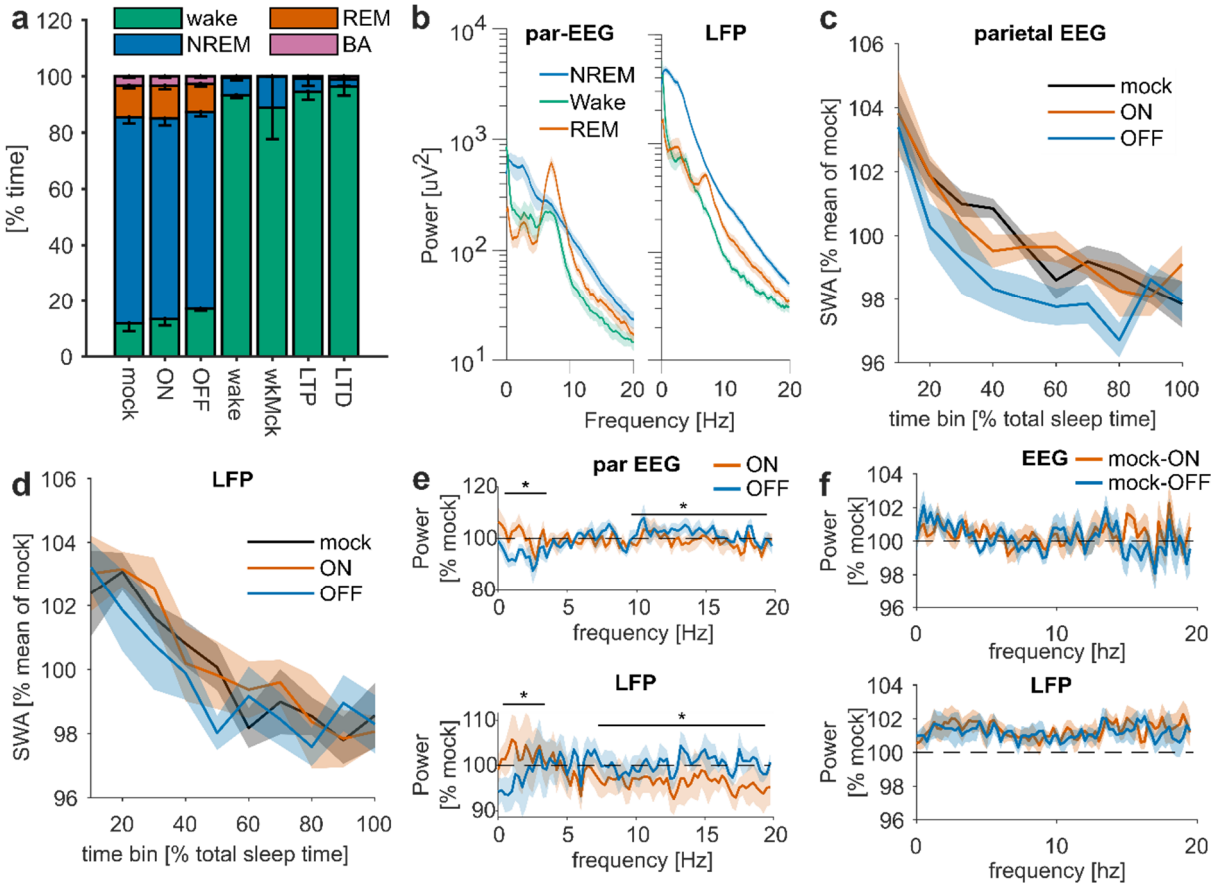


Figure 26 Effects of electrical stimulation on sleep
a) Impact of stimulation on the distribution of time spent in vigilance states (BA: brief awakenings, error bars represent SEM)
b) Mean EEG and LFP power spectra across all mice (n=7, shading: SEM)
c&d) Evolution of SWA across sleep in the frontal EEG and LFP (mean ± SEM across 7 mice). SWA is normalised to mean of the mock condition.
*e) Mean spectral difference between ON/OFF stimulation and mock stimulations in the parietal EEG and LFP * : $p<0.05$, pairwise comparisons*
f) Mean spectral differences between mock-ON/OFF episodes and remaining NREM episodes in mock-pairing experiments.

To examine which of the conditions differed in comparison to the mock condition, a separate LME was fit to relative change in power compared to the mock condition. For OFF state pairings, there was a significant effect of frequency ($\chi^2(79)=181.27$, $p=10^{-10}$, log-likelihood ratio test) whereas there was no effect for ON state pairings ($\chi^2(79)=78.7$, p=0.48 log-likelihood ratio test). For OFF state pairings, the model coefficients for individual frequency bins suggested a significant decrease in power in several low bins and an increase in 3 higher frequency bands (t-tests using Satterthwaite's method, *p<0.05, **p<0.01, ***p<0.001, significant decreases: 0.5 Hz*, 0.75-1.25 Hz**,2 Hz**,2.25 Hz*,2.2-2.75 Hz***, 3.5 Hz *, and significant increases: 8 Hz*,10.25 Hz* 10.5 Hz**; no correction for multiple comparison was applied, changes at 2.2-2.75 Hz would remain significant with corrections). In the LFP a separate model fit to the relative change in power during OFF state pairings supported a significant effect of frequency on the change in power compared to mock pairings. ($\chi^2(79)=685.8$, $p=10^{-16}$, likelihood ratio test). Again, post-hoc tests suggested a significant decrease in low frequency bands and an increase in one higher frequency band (p<0.05 for 0.5-1 Hz, and for 13.25 Hz). In slight contrast to the EEG, there was also a significant effect of frequency in the ON state pairing condition ($\chi^2(79)=685.8$, $p=10^{-16}$, likelihood ratio test), however no post-hoc comparisons reached significance. Together, these data indicate that ON and OFF state pairings have a differential effect on the power spectra of NREM episodes that do not contain a stimulation event. Interestingly this is the case for both the LFP and the EEG. Furthermore, this difference is likely driven by the OFF state pairings, which decrease several frequencies in the SWA range and increase (fewer) frequencies in the spindle range.

One caveat of analysing only epochs that do not contain stimulation is a possible effect of biased epoch selection, which could affect the spectrum independently of stimulation.

To quantify this bias, we focussed on the mock stimulation experiments and compared the average NREM spectra of mock-ON or mock-OFF stimulation epochs with the average NREM spectra of remaining episodes (Figure 26 f) . In the EEG, there was a significant main effect of frequency but no effect of condition and no interaction ($\chi^2(79)=0.96$, p=0.99 for condition*frequency and $\chi^2(1)=2.6$, p=0.1 for condition; $\chi^2(79)=110$, p=0.01 for frequency). Post-hoc tests suggested a significant increase in several low frequency bands (0.5 Hz**,1 Hz*,1.5 Hz**), and a decrease in several high frequency bands (13.5 Hz** 16.75 Hz*, 17 Hz***, 18 Hz***). In the LFP, there was no evidence for an interaction between frequency and condition but only a significant main effect of condition ($\chi^2(1)=0.96$, p=0.32 for condition*frequency and $\chi^2(1)=5.16$, p=0.018 for condition; $\chi^2(79)=77.6$, p=0.52 for frequency). While this suggests that mock detections of ON and OFF states do differ in their relative power (independent of the frequency), this difference was very small (0.17 % ± 0.05). More importantly, the intercept of the model was estimated at 101.2 % ± 0.36, which suggests that epochs with detected ON and OFF states tend to have slightly increased power across all spectra. Together, this confirms that episodes selected by the ON/OFF detection can bias the power spectrum, even without stimulation. However, we found no difference between the ON and OFF conditions, which makes this bias an unlikely explanation for the differences observed when stimulation is delivered.

5.3 Discussion

In this chapter I aimed to test the interaction between ON/OFF states and synaptic plasticity in freely behaving mice. Based on the previous two chapters, the hypothesis was that pairing synaptic inputs with ON states during NREM sleep would weaken these inputs more strongly than pairing them with OFF states or waking activity. To test

this hypothesis I chose a similar model to (Vyazovskiy et al. 2008), who showed an effect of sleep on plasticity in freely moving rats. Specifically, I implanted an electrical stimulation electrode into the primary motor cortex and recorded evoked responses using a 16-channel laminar electrode in the contralateral primary motor cortex.

The evoked responses in the LFP began and peaked in deeper layers before the superficial layers. In the MUA, this trend was much less clear and likely present in only a subset of animals. An early response in deeper layers would be consistent with a strong innervation of layer 5 by callosal projections (Petreanu et al. 2007). However, such a pattern is also conceivable via polysynaptic pathways (i.e. contralateral M1 – region X – M1) and not least via antidromic activation. However, even if the laminar probe were placed perfectly in the area most strongly innervated by the stimulated area, or if only antidromic activation would occur, most of the recorded responses (particularly in the LFP) would likely be the consequence of local synaptic connectivity.

5.3.2 Online detection of ON/OFF states

To test the above hypothesis, I developed a method for online detection of ON/OFF states. While closed-loop stimulation during slow waves is becoming increasingly more popular, to my knowledge all studies are based on the LFP or EEG (Gulati et al. 2017; Ngo et al. 2013; Garcia-Molina et al. 2018; Fattinger et al. 2017), which only indirectly reflects UP and DOWN states. I sought to detect slow waves directly based on the summed spiking activity.

One biologically effective means for closed loop stimulation relies on setting an (adaptive) negative threshold to detect presumed DOWN states and then targeting stimuli into the subsequent UP state by using the average delay between positive and

negative peak for each individual (Ngo et al. 2013). This method affects memory (Ngo et al. 2013), changes SWA, and influences the immune system (Besedovsky et al. 2017)

5.3.3 Plasticity

Having developed a functional state detection algorithm, I used it to test whether ON or OFF state pairings had a specific impact on the evoked response. I chose a within subject design, wherein each of seven mice received all pairing protocols: ON pairings, OFF pairings, mock, pairings, and wake pairings. As positive controls for neuronal plasticity, classical LTP and LTD protocols were also applied to several mice.

In the LFP, my analysis suggests that ON pairings significantly reduced the evoked LFP peaks compared to all other conditions except of OFF pairings, where there was only a trend. However, OFF pairings did not significantly reduce the peak amplitude. Hence, there is no evidence that OFF pairings reduced the LFP amplitude whereas ON pairings did so. The lack of significant difference between ON and OFF pairings could have to do with the incomplete separation between ON and OFF states by the algorithm. Indeed, in 2 mice the success rates were only around 70%, which may be enough for significant plasticity. Wake pairings were only significantly different compared to ON pairings but showed a clear trend to increase the evoked response compared to other conditions. Indeed, the time of peak was significantly delayed by wake pairings, which could support the notion that wake pairing has a significant effect on its own.

In the MUA, ON state pairings and wake pairings seemed to increase baseline firing rates. After correcting for changes in baseline firing, there was a significant effect of condition on the change in the mean number of evoked spikes. However, there was no evidence that ON state pairings led to a weakening. Indeed, only LTP pairings seemed

to decrease spiking, though none of the effects would withstand correction for multiple comparisons. Overall, plasticity in the MUA appeared to be very subtle and did not show any conclusive directionality.

Several technical and biological factors could contribute towards blurring the effects of plasticity in the MUA. More technical factors include the non-plastic contribution of antidromic spiking, the insufficient removal of noise in the spiking, insufficient samples size, and the effects of uncontrolled behavioural states during stimulation (discussed later). Antidromic spiking is unlikely to be a strong contributor because the expectation would be units/channels with high-fidelity to respond to stimulation, which rarely occurred (Fuller & Schlag 1976). Nevertheless, non-synaptic responses (noise or antidromic activity) still influenced evoked responses because most animals had at least one channel that responded within 1 ms of stimulation, which is inconsistent with synaptic delays and axonal conduction delays across the corpus callosum. Aside from these technical limitations, there is the possibility that plasticity in the MUA manifests in more complex measures, such as the synchrony in the spiking across neurons (Teleńczuk & Destexhe 2014). Note that this is not trivial to measure using the multi-unit activity in this study. Finally, it is possible that the observed result is an interesting combination of several factors: The relationship between evoked spiking and synaptic weights is complex (Carvalho & Buonomano 2009), and the majority of synaptic responses are sub-threshold and combine excitatory and inhibitory events. As a consequence, the required sample size to observe an effect would have to be much larger in the MUA than in the LFP, and would possibly require single unit resolution.

The LTP and LTD protocols were included as positive controls, and it is thus somewhat surprising that neither of them showed clear indications of altering the LFP or MUA

response. Indeed, close individual examination of the differences in the mean evoked responses from before to after the pairing revealed hardly any commonalities across mice (data not shown) and measuring the slope instead of the amplitude of the LFP (Vyazovskiy et al. 2008) had no effect on this outcome. LTP is a very robust phenomenon in brain slices of various brain regions and has repeatedly been observed in the hippocampus *in vivo* (Bliss & Lomo 1973; Whitlock et al. 2006). On the other hand, *in vivo* LTP in cortex of adult mice has not frequently been reported and when it has, it is not as reliable or only works for very specific pathways but not others (Glazewski et al. 1998). There are many plausible reasons why the same LTP protocol worked with a similar number of animals in the hands of (Vyazovskiy et al. 2008) but not in mine. However, whatever the reasons are, this finding suggests that my experiment is not a perfect replica of (Vyazovskiy et al. 2008) but it does not at all preclude sleep-dependent forms of plasticity.

It is crucial to stress that all the above results were derived from averaging across stimulations that occurred while mice were awake and engaged in variable behaviours. It is well known that various behavioural states can strongly influence spontaneous and evoked activity (Niell & Stryker 2010; Petersen et al. 2003; Vyazovskiy et al. 2008). I was not able to identify a strong correlation between the evoked LFP/MUA response and the EMG variance, EEG/LFP delta power or EEG/LFP theta power (data not shown). However, this could simply reflect the fact that the diversity of mouse behaviour is not easily captured by such measures. Another approach is to manually annotate the videos of the mice at the time of each stimulation. This effort is currently underway.

Another way around the confounding factor of behavioural states is to measure evoked responses during sleep. This analysis supported the notion that ON and OFF state pairings induced changes in the LFP evoked response. Using subtractive normalisation, there was no difference between ON and OFF stimulation but using divisive normalisation ON states showed a stronger decrease. This result is interesting because it matches the results from the evoked responses during wakefulness. As shown in (Figure 25), the reduction in LFP amplitude after ON state pairings was much clearer with divisive rather than subtractive normalisation. Of course, this analysis has its own significant limitations. As shown in (Figure 22, page 114), the state of the neuronal network modulates the evoked response. It is very likely that this modulation is not simply a discrete switch from ON to OFF state but that it has a significant continuous component (i.e. the bigger the ON states, the weaker the response). Given that slow waves change from the beginning to the end of sleep, this could strongly confound this analysis. Additionally, the detection algorithm may also be biased towards detecting slightly different types of ON/OFF states as sleep progresses. Thus, while this analysis cannot convince in isolation, in complements well the analysis of responses during wakefulness. It confirms the notion that the LFP and MUA display different trends in response to stimulation and it slightly soothes the concern that the decrease in evoked responses after ON state pairings is due to differences in behavioural state.

5.3.4 Effects of stimulation on sleep

If and how state-specific stimulation modulates sleep is of two-fold importance: Firstly, it is a confound because stimulation could have an arousing (Segundo et al. 1955) or sleep-promoting (Akert et al. 1951) effect and could thus affect plasticity via any of the proposed mechanisms by which sleep may affect plasticity (see introduction). Secondly,

a change in SWA that correlates with changes in the evoked responses would be expected if the results from the previous chapter apply to this experiment. So how could we tell the confound from the interesting result? Firstly, a strongly arousing or sleep promoting effect should change not only sleep depth (measured as SWA), but also the time spent asleep. I did not find any evidence for this. Secondly, plasticity usually develops gradually (e.g. Figure 15 on page 89) and the changes in SWA would be expected to follow such a time course. On the other hand, if stimulation directly affected sleep depth, there should be little or no delay.

To assess this, power spectral density was assessed in all 4 s epochs that did not include a stimulation. This analysis revealed that OFF state pairings significantly decreased SWA activity compared to mock and ON state pairings, in a manner not linearly dependent on time in both the LFP and EEG. Furthermore, frequencies above approximately 11 Hz had increased power in the OFF compared to ON state pairings. This strongly indicates a direct effect of OFF state pairings on sleep for several reasons: Firstly, the rate of change in SWA does not seem to differ between OFF pairings and ON pairings, suggesting that stimulation simply decreases SWA to a similar degree at most time points. Secondly, the same effects were observed in the LFP and the EEG. The EEG is located on the same hemisphere as the stimulation electrode, which makes UP state induced plasticity a less likely (though plausible) driver for this effect. Thirdly, the shift from lower to higher frequencies seems more consistent with arousal than with local changes in SWA.

Our findings do not fully agree with previous studies using closed-loop stimulation. (Ngo et al. 2013) calculated spectra across all 4 s epochs during the pairing period and found an increase in EEG - SWA when auditory stimulation was targeted to the UP

state (their fig S1) and a decrease when the OFF state was targeted (their figure S4). However, when they excluded epochs including stimulation from the latter analysis the effect disappeared. It is thus possible that the effects of ON state stimulations are comparable. However, in our hands, OFF states seem to reduce SWA when stimulation periods are not included. Furthermore, Ngo et al. find a trend towards increasing spindle power across all episodes with UP state pairings and very weak changes with DOWN state pairings. We observe the opposite, ON state pairings have decreased spindle power. This could again be explained by the fact that we excluded stimulation epochs from our analysis.

6. General Discussion

In this thesis I investigated the bidirectional interactions between cortical slow waves and synaptic plasticity. I began in brain slices, where cortical UP states weakened subthreshold inputs but not suprathreshold inputs. I then moved to an anesthetised preparation, where I observed that localised weakening of sensory responses is associated with a localised weakening of slow waves. Finally, I conducted experiments during natural sleep and attempted to test whether cortical ON states can weaken electrically evoked responses. While the evoked responses in the LFP weakened in response to ON state pairings but not OFF state pairings, I did not observe any trend in the spiking response.

I see two major technical limitations in this work. Firstly, the dataset on the effect of plasticity on SWA contained unbalanced differences in the stimulation patterns beyond the separation into the UP and DOWN state groups. While the data do not support an obvious influence of this difference on the findings, it should ideally be experimentally addressed. Secondly, the sample size for the naturally sleeping mice is relatively small and effects could be accounted for by differences in behaviour at the time of stimulation.

While I completely acknowledge the limitations of the data presented in this thesis, I will below discuss the major questions and consequences that our work would raise if the limitations can be addressed.

6.1 UP state induced plasticity during sleep

Together, the data from (González-Rueda et al. 2018) and our study (Bartram/Kahn et al. 2017) suggest that cortical UP states may bias synapses towards weakening, even

though this does not happen indiscriminately. Indeed, both studies agree that suprathreshold inputs are preserved (but not strengthened), which could reflect an elegant mechanism to increase the signal to noise ratio of the network (Tononi & Cirelli 2019) and enhance the relative strength of coherent inputs.

To me, a key question is, whether the critical factor in all these experiments are the sleep-specific activity patterns or, alternatively, the sleep-specific cellular environment (e.g. gene expression, neuromodulation, cellular metabolites etc.). In other words, UP states could be important for plasticity and memory simply because they contain most of the spiking and/or most of the synchronised spiking across brain regions. What really matters for plasticity is the cellular environment in which this spiking occurs. In an alternative scenario, something unique in the physiology of the slow oscillation could influence the plasticity rules applied to spiking patters.

We observed UP state induced weakening in brain slices as well as during anaesthesia and potentially during natural sleep. This could suggest a degree of invariance of UP state induced plasticity to genetic and neuromodulatory environments. However, all these conditions share comparably low levels of several key modulators, such as acetylcholine or noradrenaline. It is thus intriguing that (Miyamoto et al. 2016) were able to enhance memory consolidation by imposing a slow oscillation like pattern during a period of sleep deprivation. This could further support the idea that UP states provide a unique cellular environment for plasticity. Moreover, the difference between UP states and wakefulness is also partially controlled for in our experiments in freely moving mice. As shown in (Figure 23), delivering electrical stimulation during ON states did weaken the response while stimulation during wakefulness had no such effect.

If UP states provide a unique environment for plasticity, how would they do so? After all, as mentioned in chapter 3 (discussion), intracellular recordings from UP states strongly resemble recordings during wakefulness (Destexhe et al. 2007). In the light of this, one straightforward possibility is that the DOWN state is what distinguishes the UP state: It seems likely that near constant depolarisation (wakefulness) creates a very different intracellular environment than intermittent depolarisation. Indeed, it is well known that many ion channels crucial for dendritic processing inactivate upon depolarisation and only de-inactivate upon hyperpolarisation. For instance, T-type calcium channels are almost completely inactivated when the membrane potential is consistently above -60 mV but become de-inactivated upon hyperpolarisation (Leresche & Lambert 2017). Moreover, these channels have strong links to both LTP and LTD in cortex (Birtoli & Ulrich 2004; Leresche & Lambert 2017). Our *ex vivo* data also support a role of calcium dynamics in UP state induced weakening (Bartram/Kahn *et al.*, 2017). Using two-photon calcium imaging of dendritic spines, we found that incoming subthreshold inputs produce a stronger calcium influx during UP states compared to DOWN states (even though the evoked potential is actually smaller).

Calcium dynamics provide a direct link to the question why UP states would induce synaptic weakening and not strengthening. Considerable evidence suggests that very low calcium influx may have no consequences, while high calcium concentrations induce LTP and intermediate concentrations induce LTD (Lüscher & Malenka 2012). It is tempting to speculate that subthreshold inputs during DOWN states would induce calcium influx below the plasticity threshold and that UP states would push these dynamics towards LTD. Correlated presynaptic activity – as may occur during spindles – could then lead to LTP.

In summary, our data could support the following scenario: The UP state – enabled by the preceding DOWN state – boosts calcium dynamics such that subthreshold inputs are depressed. Moreover, coordinated activity brought about by replay and/or coordinated spindles and sharp-wave ripples leads to suprathreshold pairings and thus strengthening (or at least preserved synaptic strength on a background of weakening).

6.2 Consequences for SHY

As mentioned in my introduction, testing SHY involves (i) proving that sleep history influences net synaptic weights in cortex. (ii) Finding a mechanism for how this happens and manipulating this mechanism to (iii) prove that net weakening is functionally relevant. Our work is relevant for (ii) and (iii). It suggests that UP states may bias the network towards synaptic weakening. Admittedly, this mechanism is not easy to manipulate. However, one starting point may be GSK-3β, which regulates the balance between LTP and LTD, is upregulated during sleep (Vyazovskiy et al. 2008) and was crucial for UP state induced weakening in our hands.

We used UP state induced weakening to test whether it will in turn affect the slow oscillation under anaesthesia. Local and global changes in SWA are the most widely used marker for sleep pressure and are thus very likely reflecting a functionally important function of sleep. We found that UP state induced weakening is indeed associated with localised changes in SWA. This would lend further support to the idea that (at least local) SWA homeostasis can be driven by synaptic weakening and that SHY has functional relevance.

6.3. Next steps

In the short term, both *in vivo* data sets could be significantly enhanced by additional samples. If confirmed, there are several questions directly relating to our findings:

1) UP states vs. wake-like activity *ex vivo*: It would be interesting to explore whether our plasticity regimen truly depends on a hyperpolarisation-activated component. One could repeat our experiment in a pharmacological environment that produces more constant depolarisation and then optogenetically or electrically induce DOWN states. Alternatively, one could directly block or knock out specific channels.
2) Neuronal plasticity and SWA *in vivo*: It would be interesting to further explore whether local changes in excitatory synaptic strength can affect SWA. One interesting tool for this could be optical control of excitatory synaptic strength: These recently described methods may permit manipulating postsynaptic AMPA receptor concentrations using light (Sinnen et al. 2017; Kakegawa et al. 2018). If SHY is correct, one would expect to be able to increase SWA by increasing AMPA receptor contents. Moreover, inhibiting AMPA receptor internalization should reduce the decline of SWA.
3) Contrasting effects of spindles and slow waves on evoked responses:
   It is plausible that UP states in isolation promote net weakening while spindles are involved in strengthening. If a suitable online-spindle detector could be used, contrasting ON state pairings with spindle pairings could be an intriguing experiment.

7. Bibliography

Abbott, L.F. & Nelson, S.B., 2000. Synaptic plasticity: taming the beast. *Nature Neuroscience*, 3(Supp), pp.1178–1183.

Akert, K., Koella, W.P. & Hess, R., 1951. Sleep Produced by Electrical Stimulation of the Thalamus. *American Journal of Physiology-Legacy Content*, 168(1), pp.260–267.

Amzica, F. & Steriade, M., 1995. Short- and long-range neuronal synchronization of the slow (< 1 Hz) cortical oscillation. *Journal of Neurophysiology*, 73(1), pp.20–38.

Armstrong-James, M., Fox, K. & Das-Gupta, A., 1992. Flow of excitation within rat barrel cortex on striking a single vibrissa. *Journal of neurophysiology*, 68(4), pp.1345–58.

Aston-Jones, G. & Bloom, F.E., 1981. Activity of norepinephrine-containing locus coeruleus neurons in behaving rats anticipates fluctuations in the sleep-waking cycle. *Journal of Neuroscience*, 1(8), pp.876–886.

Azevedo, F.A.C. et al., 2009. Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *The Journal of Comparative Neurology*, 513(5), pp.532–541.

Bain, J. et al., 2007. The selectivity of protein kinase inhibitors: a further update. *Biochemical Journal*, 408(3), pp.297–315.

Barth, A.L. & Poulet, J.F.A., 2012. Experimental evidence for sparse firing in the neocortex. *Trends in neurosciences*, 35(6), pp.345–55.

Bartram, J. et al., 2017. Cortical Up states induce the selective weakening of subthreshold synaptic inputs. *Nature communications*, 8(1), p.665.

Beltramo, R. et al., 2013. Layer-specific excitatory circuits differentially control recurrent network dynamics in the neocortex. *Nature Neuroscience*, 16(2), pp.227–234.

Berger, H., 1933. Über das Elektrenkephalogramm des Menschen. *Archiv für Psychiatrie und Nervenkrankheiten*, 98(1), pp.231–254.

Besedovsky, L. et al., 2017. Auditory closed-loop stimulation of EEG slow oscillations strengthens sleep and signs of its immune-supportive function. *Nature Communications*, 8(1), p.1984.

Bi, G.Q. & Poo, M.M., 1998. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 18(24), pp.10464–72.

Birtoli, B. & Ulrich, D., 2004. Firing mode-dependent synaptic plasticity in rat neocortical pyramidal neurons. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 24(21), pp.4935–40.

Bjorness, T.E. et al., 2009. Control and function of the homeostatic sleep response by adenosine A 1 receptors. *Journal of Neuroscience*.

Blake, H. & Gerard, R.W., 1937. Brain Potentials During Sleep. *American Journal of Physiology-Legacy Content*, 119(4), pp.692–703.

Bliss, T. V & Lomo, T., 1973. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *The Journal of physiology*, 232(2), pp.331–56.

Bock, D.D. et al., 2011. Network anatomy and in vivo physiology of visual cortical neurons. *Nature*, 471(7337), pp.177–182.

Borbély, A.A., 1982. A two process model of sleep regulation. *Human neurobiology*, 1(3), pp.195–204.

Born, J. & Feld, G.B., 2012. Sleep to Upscale, Sleep to Downscale: Balancing Homeostasis and Plasticity. *Neuron*, 75(6), pp.933–935.

Brecht, M. & Sakmann, B., 2002. -Dynamic representation of whisker deflection by synaptic potentials in spiny stellate and pyramidal cells in the barrels and septa of layer 4 rat somatosensory cortex. *The Journal of Physiology*, 543(1), pp.49–70.

Brittain, M.K. et al., 2012. Ifenprodil, a NR2B-selective antagonist of NMDA receptor, inhibits reverse Na+/Ca2+ exchanger in neurons. *Neuropharmacology*, 63(6), pp.974–982.

Bulach, R., Myles, P.S. & Russnak, M., 2005. Double-blind randomized controlled trial to determine extent of amnesia with midazolam given immediately before general anaesthesia. *British Journal of Anaesthesia*, 94(3), pp.300–305.

Buzsáki, G., Anastassiou, C.A. & Koch, C., 2012. The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13(6), pp.407–420.

Cahill, L., Brioni, J. & Izquierdo, I., 1986. Retrograde memory enhancement by diazepam: its relation to anterograde amnesia, and some clinical implications. *Psychopharmacology*, 90(4).

Campbell, S.S. & Tobler, I., 1984. Animal sleep: A review of sleep duration across phylogeny. *Neuroscience & Biobehavioral Reviews*, 8(3), pp.269–300.

Cane, M. et al., 2014. The relationship between PSD-95 clustering and spine stability in vivo. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 34(6), pp.2075–86.

Canto, C.B. et al., 2017. The Sleeping Cerebellum. *Trends in neurosciences*, 40(5), pp.309–323.

Carandini, M. & Ferster, D., 2000. Membrane potential and firing rate in cat primary visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 20(1), pp.470–84.

Carvalho, T.P. & Buonomano, D. V., 2009. Differential Effects of Excitatory and Inhibitory Plasticity on Synaptically Driven Neuronal Input-Output Functions. *Neuron*, 61(5), pp.774–785.

Chaure, F.J., Rey, H.G. & Quian Quiroga, R., 2018. A novel and fully automatic spike-sorting implementation with variable number of features. *Journal of Neurophysiology*, 120(4), pp.1859–1871.

Chauvette, S., Seigneur, J. & Timofeev, I., 2012. Sleep oscillations in the thalamocortical system induce long-term neuronal plasticity. *Neuron*, 75(6), pp.1105–13.

Chklovskii, D.B., Mel, B.W. & Svoboda, K., 2004. Cortical rewiring and information storage. *Nature*, 431(7010), pp.782–788.

Chung, J.E. et al., 2017. A Fully Automated Approach to Spike Sorting. *Neuron*, 95(6), p.1381–1394.e6.

Collingridge, G.L. et al., 2010. Long-term depression in the CNS. *Nature Reviews Neuroscience*, 11(7), pp.459–473.

Collins, D.R., Pelletier, J.G. & Paré, D., 2001. Slow and Fast (Gamma) Neuronal Oscillations in the Perirhinal Cortex and Lateral Amygdala. *Journal of Neurophysiology*, 85(4), pp.1661–1672.

Contreras, D. et al., 1997. Spatiotemporal patterns of spindle oscillations in cortex and thalamus. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 17(3), pp.1179–96.

Cooke, S.F. & Bliss, T.V.P., 2006. Plasticity in the human central nervous system. *Brain*, 129(7), pp.1659–1673.

Costa, R.P. et al., 2017. Functional consequences of pre- and postsynaptic expression of synaptic plasticity. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 372(1715).

Crochet, S. & Petersen, C.C.H., 2006. Correlating whisker behavior with membrane potential in barrel cortex of awake mice. *Nature Neuroscience*, 9(5), pp.608–610.

Crunelli, V. & Hughes, S.W., 2010. The slow (. *Nature neuroscience*, 13(1), pp.9–17.

Daan, S., Beersma, D.G. & Borbely, A.A., 1984. Timing of human sleep: recovery

process gated by a circadian pacemaker. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 246(2), pp.R161–R183.

Dadvand, L., Stryker, M.P. & Frank, M.G., 2006. Sleep does not enhance the recovery of deprived eye responses in developing visual cortex. *Neuroscience*, 143(3), pp.815–826.

David, F. et al., 2013. Essential Thalamic Contribution to Slow Waves of Natural Sleep. *Journal of Neuroscience*, 33(50), pp.19599–19610.

DeFelipe, J., Alonso-Nanclares, L. & Arellano, J.I., 2002. Microstructure of the neocortex: Comparative aspects. *Journal of Neurocytology*, 31(3–5 SPEC. ISS.), pp.299–316.

Destexhe, A. et al., 2007. Are corticothalamic "up" states fragments of wakefulness? *Trends in Neurosciences*, 30(7), pp.334–342.

Diering, G.H. et al., 2017. Homer1a drives homeostatic scaling-down of excitatory synapses during sleep. *Science*, 355(6324), pp.511–515.

Digby, R.J. et al., 2017. Distinct mechanisms of Up state maintenance in the medial entorhinal cortex and neocortex. *Neuropharmacology*, 113, pp.543–555.

Dijk, D.J., Beersma, D.G.M. & Daan, S., 1987. EEG Power Density during Nap Sleep: Reflection of an Hourglass Measuring the Duration of Prior Wakefulness. *Journal of Biological Rhythms*, 2(3), pp.207–219.

Dijk, D.J. & Czeisler, C.A., 1995. Contribution of the circadian pacemaker and the sleep homeostat to sleep propensity, sleep structure, electroencephalographic slow

waves, and sleep spindle activity in humans. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 15(5 Pt 1), pp.3526–38.

Doi, A. et al., 2007. Slow oscillation of membrane currents mediated by glutamatergic inputs of rat somatosensory cortical neurons: in vivo patch-clamp analysis. *European Journal of Neuroscience*, 26(9), pp.2565–2575.

Donchin, O. et al., 2002. Mechanisms Influencing Acquisition and Recall of Motor Memories. *Journal of Neurophysiology*, 88(4), pp.2114–2123.

Duboué, E.R., Keene, A.C. & Borowsky, R.L., 2011. Evolutionary convergence on sleep loss in cavefish populations. *Current Biology*.

Ellenbogen, J.M. et al., 2006. Interfering with Theories of Sleep and Memory: Sleep, Declarative Memory, and Associative Interference. *Current Biology*, 16(13), pp.1290–1294.

Ellenbogen, J.M., Payne, J.D. & Stickgold, R., 2006. The role of sleep in declarative memory consolidation: passive, permissive, active or none? *Current Opinion in Neurobiology*, 16(6), pp.716–722.

Eyre, M.D. et al., 2003. Morphological changes in hippocampal dentate gyrus synapses following spatial learning in rats are transient. *European Journal of Neuroscience*, 17(9), pp.1973–1980.

Faber, D.S. & Korn, H., 1991. Applicability of the coefficient of variation method for analyzing synaptic plasticity. *Biophysical journal*, 60(5), pp.1288–94.

Fattinger, S. et al., 2017. Deep sleep maintains learning efficiency of the human brain.

Nature Communications, 8(1), p.15405.

Favero, M. & Castro-Alamancos, M.A., 2013. Synaptic Cooperativity Regulates Persistent Network Activity in Neocortex. *Journal of Neuroscience*, 33(7), pp.3151–3163.

Feldman, D.E., 2000. Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron*, 27(1), pp.45–56.

Fox, K., 1992. A critical period for experience-dependent synaptic plasticity in rat barrel cortex. *Journal of Neuroscience*, 12(5).

Frank, M.G., 2016. Circadian Regulation of Synaptic Plasticity. *Biology*, 5(3).

Frank, M.G., 2012. Erasing Synapses in Sleep: Is It Time to Be SHY? *Neural Plasticity*, 2012, pp.1–15.

Frank, M.G., Issa, N.P. & Stryker, M.P., 2001. Sleep enhances plasticity in the developing visual cortex. *Neuron*, 30(1), pp.275–87.

Franken, P., Malafosse, A. & Tafti, M., 1998. Genetic variation in EEG activity during sleep in inbred mice. *American Journal of Physiology - Regulatory Integrative and Comparative Physiology*.

Fröberg, J.E. et al., 1975. Circadian rhythms of catecholamine excretion, shooting range performance and self-ratings of fatigue during sleep deprivation. *Biological Psychology*, 2(3), pp.175–188.

Fuentealba, P. & Steriade, M., 2005. The reticular nucleus revisited: Intrinsic and network properties of a thalamic pacemaker. *Progress in Neurobiology*, 75(2),

pp.125–141.

Fujii, S. et al., 1991. Reversal of long-term potentiation (depotentiation) induced by tetanus stimulation of the input to CA1 neurons of guinea pig hippocampal slices. *Brain Research*, 555(1), pp.112–122.

Fuller, J.H. & Schlag, J.D., 1976. Determination of antidromic excitation by the collision test: Problems of interpretation. *Brain Research*, 112(2), pp.283–298.

Gais, S. et al., 2002. Learning-dependent increases in sleep spindle density. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 22(15), pp.6830–4.

Garcia-Molina, G. et al., 2018. Closed-loop system to enhance slow-wave activity. *Journal of Neural Engineering*, 15(6), p.66018.

Gent, T.C. et al., 2018. Thalamic dual control of sleep and wakefulness. *Nature Neuroscience*, 21(7), pp.974–984.

Girardeau, G. et al., 2009a. Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, 12(10), pp.1222–1223.

Girardeau, G. et al., 2009b. Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, 12(10), pp.1222–1223.

Glazewski, S. et al., 1998. Long-term potentiation in vivo in layers II/III of rat barrel cortex. In *Neuropharmacology*. pp. 581–592.

Glazewski, S. & Fox, K., 1996. Time course of experience-dependent synaptic potentiation and depression in barrel cortex of adolescent rats. *Journal of*

Neurophysiology, 75(4), pp.1714–1729.

González-Rueda, A. et al., 2018. Activity-Dependent Downscaling of Subthreshold Synaptic Inputs during Slow-Wave-Sleep-like Activity In Vivo. *Neuron*, 97(6), p.1244–1252.e5.

Gulati, T. et al., 2017. Neural reactivations during sleep determine network credit assignment. *Nature Neuroscience*, 20(9), pp.1277–1284.

Gulevich, G., 1966. Psychiatric and EEG Observations on a Case of Prolonged (264 Hours) Wakefulness. *Archives of General Psychiatry*, 15(1), p.29.

Hahn, T.T.G. et al., 2012. Spontaneous persistent activity in entorhinal cortex modulates cortico-hippocampal interaction in vivo. *Nature Neuroscience*, 15(11), pp.1531–1538.

Haider, B. et al., 2007. Enhancement of Visual Responsiveness by Spontaneous Local Network Activity In Vivo. *Journal of Neurophysiology*, 97(6), pp.4186–4202.

Haider, B. et al., 2016. Millisecond Coupling of Local Field Potentials to Synaptic Currents in the Awake Visual Cortex. *Neuron*, 90(1), pp.35–42.

Hansen, K.B. et al., 2014. Distinct Functional and Pharmacological Properties of Triheteromeric GluN1/GluN2A/GluN2B NMDA Receptors. *Neuron*, 81(5), pp.1084–1096.

Hara, K. & Harris, R.A., 2002. The Anesthetic Mechanism of Urethane: The Effects on Neurotransmitter-Gated Ion Channels. *Anesthesia & Analgesia*, 94(2), pp.313–318.

Harrison, X.A. et al., 2018. A brief introduction to mixed effects modelling and multi-

model inference in ecology. *PeerJ*, 6, p.e4794.

Hayashi-Takagi, A. et al., 2015. Labelling and optical erasure of synaptic memory traces in the motor cortex. *Nature*, 525(7569), pp.333–338.

Hebb, D.O., 1949. *The Organization of Behavior; A Neuropsychological Theory*,

Hengen, K.B. et al., 2016. Neuronal Firing Rate Homeostasis Is Inhibited by Sleep and Promoted by Wake. *Cell*, 165(1), pp.180–191.

Herculano-Houzel, S., 2010. Coordinated scaling of cortical and cerebellar numbers of neurons. *Frontiers in neuroanatomy*, 4, p.12.

Herculano-Houzel, S., Mota, B. & Lent, R., 2006. *Cellular scaling rules for rodent brains*,

Holtmaat, A. et al., 2006. Experience-dependent and cell-type-specific spine growth in the neocortex. *Nature*, 441(7096), pp.979–983.

Holtmaat, A. & Caroni, P., 2016. Functional and structural underpinnings of neuronal assembly formation in learning. *Nature Neuroscience*, 19(12), pp.1553–1562.

Huber, R. et al., 2006. Arm immobilization causes cortical plastic changes and locally decreases sleep slow wave activity. *Nature Neuroscience*, 9(9), pp.1169–1176.

Huber, R. et al., 2004. Local sleep and learning. *Nature*, 430(6995), pp.78–81.

Huber, R. et al., 2007. TMS-induced cortical potentiation during wakefulness locally increases slow wave activity during sleep. *PLoS ONE*, 2(3).

Hughes, S.W. et al., 2004. Synchronized Oscillations at α and θ Frequencies in the

Lateral Geniculate Nucleus. *Neuron*, 42(2), pp.253–268.

Jenkins, J.G. & Dallenbach, K.M., 1924. Obliviscence during Sleep and Waking. *The American Journal of Psychology*, 35(4), p.605.

Jha, S.K. et al., 2005. Sleep-dependent plasticity requires cortical activity. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 25(40), pp.9266–74.

Ji, D. & Wilson, M.A., 2007. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience*, 10(1), pp.100–107.

Juraska, J.M., 1982. The development of pyramidal neurons after eye opening in the visual cortex of hooded rats: A quantitative study. *The Journal of Comparative Neurology*, 212(2), pp.208–213.

Kakegawa, W. et al., 2018. Optogenetic Control of Synaptic AMPA Receptor Endocytosis Reveals Roles of LTD in Motor Learning. *Neuron*, 99(5), p.985–998.e6.

Karlsson, M.P. & Frank, L.M., 2009. Awake replay of remote experiences in the hippocampus. *Nature Neuroscience*, 12(7), pp.913–918.

Kattler, Dijk & Borbély, 1994. Effect of unilateral somatosensory stimulation prior to sleep on the sleep EEG in humans. *Journal of sleep research*, 3(3), pp.159–164.

Kerchner, G.A. & Nicoll, R.A., 2008. Silent synapses and the emergence of a postsynaptic mechanism for LTP. *Nature reviews. Neuroscience*, 9(11), pp.813–25.

Kimura, A. & Pavlides, C., 2000. Long-Term Potentiation/Depotentiation Are Accompanied by Complex Changes in Spontaneous Unit Activity in the Hippocampus. *Journal of Neurophysiology*, 84(4), pp.1894–1906.

Klinzing, J.G. et al., 2016. Spindle activity phase-locked to sleep slow oscillations. *NeuroImage*, 134, pp.607–616.

Klinzing, J.G., Niethard, N. & Born, J., 2019. Mechanisms of systems memory consolidation during sleep. *Nature Neuroscience*.

Koss, W.A. et al., 2014. Dendritic remodeling in the adolescent medial prefrontal cortex and the basolateral amygdala of male and female rats. *Synapse*, 68(2), pp.61–72.

von Krosigk, M., Bal, T. & McCormick, D., 1993. Cellular mechanisms of a synchronized oscillation in the thalamus. *Science*, 261(5119), pp.361–364.

Krueger, J.M. et al., 2016. Sleep function: Toward elucidating an enigma. *Sleep Medicine Reviews*, 28, pp.42–50.

Kumar, A.H.S. & Clover, A.J.P., 2015. Intraperitoneal co-administration of low dose urethane with xylazine and ketamine for extended duration of surgical anesthesia in rats. *Laboratory Animal Research*, 31(4), p.174.

Latchoumane, C.-F. V. et al., 2017a. Thalamic Spindles Promote Memory Formation during Sleep through Triple Phase-Locking of Cortical, Thalamic, and Hippocampal Rhythms. *Neuron*, 95(2), p.424–435.e6.

Latchoumane, C.-F. V. et al., 2017b. Thalamic Spindles Promote Memory Formation during Sleep through Triple Phase-Locking of Cortical, Thalamic, and

Hippocampal Rhythms. *Neuron*, 95(2), p.424–435.e6.

Lau, K.F. et al., 1999. Expression analysis of glycogen synthase kinase-3 in human tissues. *The journal of peptide research : official journal of the American Peptide Society*, 54(1), pp.85–91.

Lee, W.-C.A. et al., 2016. Anatomy and function of an excitatory network in the visual cortex. *Nature*, 532(7599), pp.370–374.

Lemieux, M. et al., 2014. The impact of cortical deafferentation on the neocortical slow oscillation. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 34(16), pp.5689–703.

Leresche, N. & Lambert, R.C., 2017. T-type calcium channels in synaptic plasticity. *Channels (Austin, Tex.)*, 11(2), pp.121–139.

Lesku, J.A., Martinez-Gonzalez, D. & Rattenborg, N.C., 2009. Sleep and sleep states: Phylogeny and ontogeny. In *Encyclopedia of Neuroscience*.

Letzkus, J.J. et al., 2011. A disinhibitory microcircuit for associative fear learning in the auditory cortex. *Nature*, 480(7377), pp.331–335.

Lewis, L.D. et al., 2015. Thalamic reticular nucleus induces fast and local modulation of arousal state. *eLife*, 4, p.e08760.

Li, W. et al., 2017. REM sleep selectively prunes and maintains new synapses in development and learning. *Nature Neuroscience*, 20(3), pp.427–437.

Li, Z. et al., 2010. Caspase-3 activation via mitochondria is required for long-term depression and AMPA receptor internalization. *Cell*, 141(5), pp.859–71.

Liu, Y. et al., 2019. Human Replay Spontaneously Reorganizes Experience. *Cell*, 178(3), p.640–652.e14.

Loomis, A.L., Harvey, E.N. & Hobart, G., 1935. Potential Rhythms of the cerebral cortex during sleep. *Science*, 81(2111), pp.597–598.

Lőrincz, M.L. et al., 2015. A distinct class of slow (~0.2–2 Hz) intrinsically bursting layer 5 pyramidal neurons determines UP/DOWN state dynamics in the neocortex. *Journal of Neuroscience*, 35(14), pp.5442–5458.

Lüscher, C. & Malenka, R.C., 2012. NMDA receptor-dependent long-term potentiation and long-term depression (LTP/LTD). *Cold Spring Harbor perspectives in biology*, 4(6).

Mann, E.O., Kohl, M.M. & Paulsen, O., 2009. Distinct roles of GABA(A) and GABA(B) receptors in balancing and terminating persistent cortical activity. *J Neurosci*, 29(23), pp.7513–7518.

Maret, S. et al., 2011. Sleep and waking modulate spine turnover in the adolescent mouse cortex. *Nature Neuroscience*, 14(11), pp.1418–1420.

Markram, H. et al., 1997. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275(5297), pp.213–215.

Marshall, L. et al., 2006. Boosting slow oscillations during sleep potentiates memory. *Nature*, 444(7119), pp.610–613.

Massimini, M. et al., 2004a. The sleep slow oscillation as a traveling wave. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 24(31),

pp.6862–70.

Massimini, M. et al., 2004b. The sleep slow oscillation as a traveling wave. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 24(31), pp.6862–70.

McConnell, A.R. et al., 2011. Friends with benefits: on the positive consequences of pet ownership. *Journal of personality and social psychology*, 101(6), pp.1239–52.

McGinty, D.J. & Harper, R.M., 1976. Dorsal raphe neurons: depression of firing during sleep in cats. *Brain Research*, 101(3), pp.569–575.

Meier, C.A. et al., 1968. Forgetting of Dreams in the Laboratory. *Perceptual and Motor Skills*, 26(2), pp.551–557.

Mitra, P. & Bokil, H., 2009. *Observed Brain Dynamics*,

Miyamoto, D. et al., 2016. Top-down cortical input during NREM sleep consolidates perceptual memory. *Science (New York, N.Y.)*, 352(6291), pp.1315–8.

Miyamoto, H., Katagiri, H. & Hensch, T., 2003. Experience-dependent slow-wave sleep development. *Nature neuroscience*, 6(6), pp.553–4.

Moczulska, K.E. et al., 2013. Dynamics of dendritic spines in the mouse auditory cortex during memory formation and memory recall. *Proceedings of the National Academy of Sciences of the United States of America*, 110(45), pp.18315–20.

Mölle, M. et al., 2002. Grouping of spindle activity during slow oscillations in human non-rapid eye movement sleep. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 22(24), pp.10941–7.

Moore, C.I. & Nelson, S.B., 1998. Spatio-temporal subthreshold receptive fields in the vibrissa representation of rat primary somatosensory cortex. *Journal of Neurophysiology*.

Moser, E., Mathiesen, I. & Andersen, P., 1993. Association between brain temperature and dentate field potentials in exploring and swimming rats. *Science*, 259(5099), pp.1324–1326.

Moser, M.-B. & Trommald, M., 1994. *An increase in dendritic spine density on hippocampal CAl pyramidal cells following spatial learning in adult rats suggests the formation of new synapses (hippocampus/plasticity/confocal microscopy/complex environment)*,

Muehlroth, B.E. et al., 2019. Precise Slow Oscillation–Spindle Coupling Promotes Memory Consolidation in Younger and Older Adults. *Scientific Reports*, 9(1), p.1940.

Neske, G.T., 2016. The Slow Oscillation in Cortical and Thalamic Networks: Mechanisms and Functions. *Frontiers in Neural Circuits*, 9, p.88.

Neske, G.T., 2015. The Slow Oscillation in Cortical and Thalamic Networks: Mechanisms and Functions. *Frontiers in neural circuits*, 9(January), p.88.

Ngo, H.-V.V. et al., 2013. Auditory Closed-Loop Stimulation of the Sleep Slow Oscillation Enhances Memory. *Neuron*, 78(3), pp.545–553.

Niell, C.M. & Stryker, M.P., 2010. Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex. *Neuron*, 65(4), pp.472–479.

Niethard, N. et al., 2018. Cortical circuit activity underlying sleep slow oscillations and spindles. *Proceedings of the National Academy of Sciences of the United States of America*, 115(39), pp.E9220–E9229.

O'Malley, A. et al., 2000. Transient spine density increases in the mid-molecular layer of hippocampal dentate gyrus accompany consolidation of a spatial learning task in the rodent. *Neuroscience*, 99(2), pp.229–232.

Olcese, U., Esser, S.K. & Tononi, G., 2010. Sleep and Synaptic Renormalization: A Computational Study. *Journal of Neurophysiology*, 104(6), pp.3476–3493.

Pachitariu, M. et al., 2016. Fast and accurate spike sorting of high-channel count probes with KiloSort. , pp.4448–4456.

Pang, R. et al., 1993. Effect of propofol on memory in mice. *Pharmacology Biochemistry and Behavior*, 44(1), pp.145–151.

Pawlak, V. et al., 2010. Timing is not everything: neuromodulation opens the STDP gate. *Frontiers in Synaptic Neuroscience*, 2, p.146.

Peineau, S. et al., 2007. LTP inhibits LTD in the hippocampus via regulation of GSK3beta. *Neuron*, 53(5), pp.703–17.

Peineau, S. et al., 2009. The role of GSK-3 in synaptic plasticity. *British Journal of Pharmacology*, 153(S1), pp.S428–S437.

Petersen, C.C.H. et al., 2003. Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 100(23), pp.13638–13643.

Petreanu, L. et al., 2007. Channelrhodopsin-2–assisted circuit mapping of long-range callosal projections. *Nature Neuroscience*, 10(5), pp.663–668.

Pike, F.G. et al., 1999. Rapid report: postsynaptic bursting is essential for "Hebbian" induction of associative long-term potentiation at excitatory synapses in rat hippocampus. *The Journal of physiology*, 518 (Pt 2)(Pt 2), pp.571–6.

Prober, D.A. et al., 2006. Hypocretin/orexin overexpression induces an insomnia-like phenotype in zebrafish. *Journal of Neuroscience*.

Raizen, D.M. et al., 2008. Lethargus is a Caenorhabditis elegans sleep-like state. *Nature*.

Rampon, C. et al., 2000. Enrichment induces structural changes and recovery from nonspatial memory deficits in CA1 NMDAR1-knockout mice. *Nature Neuroscience*, 3(3), pp.238–244.

Rasch, B. & Born, J., 2013. About sleep's role in memory. *Physiological reviews*, 93(2), pp.681–766.

Rector, D.M. et al., 2005. Local functional state differences between rat cortical columns. *Brain Research*, 1047(1), pp.45–55.

Reig, R. et al., 2015. Gain modulation of synaptic inputs by network state in auditory cortex in vivo. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 35(6), pp.2689–702.

Robertson, E.M., Pascual-Leone, A. & Miall, R.C., 2004. Current concepts in procedural consolidation. *Nature Reviews Neuroscience*, 5(7), pp.576–582.

Robertson, E.M., Pascual-Leone, A. & Press, D.Z., 2004. Awareness modifies the skill-learning benefits of sleep. *Current biology : CB*, 14(3), pp.208–12.

Rolls, A. et al., 2011. Optogenetic disruption of sleep continuity impairs memory consolidation. *Proceedings of the National Academy of Sciences*, 108(32), pp.13305–13310.

Ros, H. et al., 2009. Neocortical networks entrain neuronal circuits in cerebellar cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 29(33), pp.10309–20.

Rosanova, M. & Ulrich, D., 2005. Pattern-specific associative long-term potentiation induced by a sleep spindle-related spike train. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 25(41), pp.9398–405.

Rosman, E. et al., 1992. Halothane anesthesia causes state-dependent retrieval failure in mice. *Physiology & Behavior*, 52(3), pp.449–453.

Sadowski, J.H.L.P., Jones, M.W. & Mellor, J.R., 2016. Sharp-Wave Ripples Orchestrate the Induction of Synaptic Plasticity during Reactivation of Place Cell Firing Patterns in the Hippocampus. *Cell Reports*, 14(8), pp.1916–1929.

Sakata, S. & Harris, K.D., 2009. Laminar Structure of Spontaneous and Sensory-Evoked Population Activity in Auditory Cortex. *Neuron*, 64(3), pp.404–418.

Sanchez-Vives, M. V & McCormick, D.A., 2000. Cellular and network mechanisms of rhythmic recurrent activity in neocortex. *Nature Neuroscience*, 3(10), pp.1027–1034.

Sara, S.J., 2009. The locus coeruleus and noradrenergic modulation of cognition. *Nature Reviews Neuroscience*, 10(3), pp.211–223.

Sawangjit, A. et al., 2018a. The hippocampus is crucial for forming non-hippocampal long-term memory during sleep. *Nature*, 564(7734), pp.109–113.

Sawangjit, A. et al., 2018b. The hippocampus is crucial for forming non-hippocampal long-term memory during sleep. *Nature*, 564(7734), pp.109–113.

Schlingloff, D. et al., 2014. Mechanisms of sharp wave initiation and ripple generation. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 34(34), pp.11385–98.

Segundo, J.P., Naquet, R. & Buser, P., 1955. effects of cortical stimulation on electrocortical activity in monkeys. *Journal of Neurophysiology*, 18(3), pp.236–245.

Seol, G.H. et al., 2007. Neuromodulators Control the Polarity of Spike-Timing-Dependent Synaptic Plasticity. *Neuron*, 55(6), pp.919–929.

Shaw, P. et al., 2000. Correlates of Sleep and Waking. *Science*.

Sheroziya, M. & Timofeev, I., 2014. Global Intracellular Slow-Wave Dynamics of the Thalamocortical System. *Journal of Neuroscience*, 34(26), pp.8875–8893.

Shu, Y., Hasenstaub, A. & McCormick, D.A., 2003. Turning on and off recurrent balanced cortical activity. *Nature*, 423(May), pp.288–293.

Siapas, A.G. & Wilson, M.A., 1998. Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron*, 21(5), pp.1123–8.

Sinnen, B.L. et al., 2017. Optogenetic Control of Synaptic Composition and Function. *Neuron*, 93(3), p.646–660.e5.

Sirota, A. et al., 2003. Communication between neocortex and hippocampus during sleep in rodents. *Proceedings of the National Academy of Sciences of the United States of America*, 100(4), pp.2065–9.

Sirviö, J. et al., 1992. The effects of dexmedetomidine, an alpha2 agonist, on learning and memory, assessed using passive avoidance and water maze tasks in rats. *Neuropharmacology*, 31(2), pp.163–168.

Sjöström, J., Turrigiano, G.G. & Nelson, S.B., 2004. Endocannabinoid-Dependent Neocortical Layer-5 LTD in the Absence of Postsynaptic Spiking. *J Neurophysiol*, 92, pp.3338–3343.

Sjöström, P.J., Turrigiano, G.G. & Nelson, S.B., 2003. Neocortical LTD via coincident activation of presynaptic NMDA and cannabinoid receptors. *Neuron*, 39(4), pp.641–54.

Sjöström, P.J., Turrigiano, G.G. & Nelson, S.B., 2001. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32(6), pp.1149–64.

Skoglund, T.S., Pascher, R. & Berthold, C.H., 1997. The existence of a layer IV in the rat motor cortex. *Cerebral Cortex*, 7(2), pp.178–180.

Slézia, A. et al., 2011. Phase advancement and nucleus-specific timing of thalamocortical activity during slow cortical oscillation. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 31(2), pp.607–17.

Squire, L.R. & Alvarez, P., 1995. Retrograde amnesia and memory consolidation: a neurobiological perspective. *Current opinion in neurobiology*, 5(2), pp.169–77.

Staresina, B.P. et al., 2015. Hierarchical nesting of slow oscillations, spindles and ripples in the human hippocampus during sleep. *Nature Neuroscience*, 18(11), pp.1679–1686.

Stepanyants, A. et al., 2009. The fractions of short- and long-range connections in the visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 106(9), pp.3555–60.

Steriade, M., 1999. Coherent oscillations and short-term plasticity in corticothalamic networks. *Trends in neurosciences*, 22(8), pp.337–45.

Steriade, M. et al., 1987. The deafferented reticular thalamic nucleus generates spindle rhythmicity. *Journal of Neurophysiology*, 57(1), pp.260–273.

Steriade, M. et al., 1993. The slow (< 1 Hz) oscillation in reticular thalamic and thalamocortical neurons: scenario of sleep rhythm generation in interacting thalamic and neocortical networks. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 13(8), pp.3284–99.

Steriade, M. & Contreras, D., 1995. Relations between cortical and thalamic cellular events during transition from sleep patterns to paroxysmal activity. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 15(1 Pt 2), pp.623–42.

Steriade, M., Mccormick, D.A. & Sejnowski, T.J., 1993. Thalamocortical Oscillations in the Sleeping and Aroused Brain. *Science*, 262(5134), pp.679–685.

Steriade, M., Nuñez, A. & Amzica, F., 1993a. A novel slow (< 1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 13(8), pp.3252–65.

Steriade, M., Nuñez, A. & Amzica, F., 1993b. Intracellular analysis of relations between the slow (< 1 Hz) neocortical oscillation and other sleep rhythms of the electroencephalogram. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 13(8), pp.3266–83.

Steriade, M., Timofeev, I. & Grenier, F., 2001. Natural Waking and Sleep States: A View From Inside Neocortical Neurons. *Journal of Neurophysiology*, 85(5), pp.1969–1985.

Takahashi, K., Lin, J.-S. & Sakai, K., 2006. Neuronal Activity of Histaminergic Tuberomammillary Neurons During Wake–Sleep States in the Mouse. *Journal of Neuroscience*, 26(40), pp.10292–10298.

Takeuchi, T. et al., 2016. Locus coeruleus and dopaminergic consolidation of everyday memory. *Nature*, 537(7620), pp.357–362.

Tang, W. et al., 2017. Hippocampal-Prefrontal Reactivation during Learning Is Stronger in Awake Compared with Sleep States. *The Journal of Neuroscience*, 37(49), pp.11789–11805.

Teleńczuk, B. & Destexhe, A., 2014. Local Field Potential, Relationship to Unit Activity. In *Encyclopedia of Computational Neuroscience*. New York, NY: Springer New York, pp. 1–6.

Timofeev, I. et al., 2000. Origin of Slow Cortical Oscillations in Deafferented Cortical Slabs. *Cerebral Cortex*, 10(12), pp.1185–1199.

Timofeev, I. & Chauvette, S., 2018. Sleep, Anesthesia, and Plasticity. *Neuron*, 97(6), pp.1200–1202.

Tobler, I., 1995. Is sleep fundamentally different between mammalian species? *Behavioural Brain Research*.

Tobler & Neuner-Jehle, 1992. 24-h variation of vigilance in the cockroach Blaberus giganteus. *Journal of sleep research*, 1(4), pp.231–239.

Tononi, G. & Cirelli, C., 2019. Sleep and synaptic down-selection. *European Journal of Neuroscience*.

Tononi, G. & Cirelli, C., 2003. Sleep and synaptic homeostasis: a hypothesis. *Brain Research Bulletin*, 62(2), pp.143–150.

Turrigiano, G., 2012. Homeostatic synaptic plasticity: Local and global mechanisms for stabilizing neuronal function. *Cold Spring Harbor Perspectives in Biology*, 4(1).

Underwood, B.J. & Postman, L., 1960. Extraexperimental sources of interference in forgetting. *Psychological Review*, 67(2), pp.73–95.

van de Ven, G.M. et al., 2016. Hippocampal Offline Reactivation Consolidates Recently Formed Cell Assembly Patterns during Sharp Wave-Ripples. *Neuron*, 92(5), pp.968–974.

Vertes, R.P., 2004. Memory Consolidation in Sleep: Dream or Reality. *Neuron*, 44(1), pp.135–148.

Veselis, R.A. et al., 1997. The comparative amnestic effects of midazolam, propofol, thiopental, and fentanyl at equisedative concentrations. *Anesthesiology*, 87(4), pp.749–764.

Viana, K.A. et al., 2017. What is the level of evidence for the amnestic effects of sedatives in pediatric patients? A systematic review and meta-analyses R. A. Veselis, ed. *PLOS ONE*, 12(7), p.e0180248.

de Vivo, L. et al., 2017a. Ultrastructural evidence for synaptic scaling across the wake/sleep cycle. *Science*, 355(6324), pp.507–510.

de Vivo, L. et al., 2017b. Ultrastructural evidence for synaptic scaling across the wake/sleep cycle. *Science*, 355(6324), pp.507–510.

Vorster, A.P. & Born, J., 2015. Sleep and memory in mammals, birds and invertebrates. *Neuroscience and Biobehavioral Reviews*.

Vyazovskiy, V., Borbély, A.A. & Tobler, I., 2000. Unilateral vibrissae stimulation during waking induces interhemispheric EEG asymmetry during subsequent sleep in the rat. *Journal of sleep research*, 9(4), pp.367–71.

Vyazovskiy, V. V et al., 2008. Molecular and electrophysiological evidence for net synaptic potentiation in wake and depression in sleep. *Nature neuroscience*, 11(2), pp.200–8.

Vyazovskiy, V. V et al., 2017. Sleep homeostasis, habits and habituation. *Current Opinion in Neurobiology*, 44, pp.202–211.

Vyazovskiy, V. V. et al., 2009. Cortical Firing and Sleep Homeostasis. *Neuron*, 63(6),

pp.865–878.

Vyazovskiy, V. V. et al., 2011. Local sleep in awake rats. *Nature*, 472(7344), pp.443–447.

Vyazovskiy, V. V. et al., 2007. Sleep Homeostasis and Cortical Synchronization: II. A Local Field Potential Study of Sleep Slow Waves in the Rat. *Sleep*, 30(12), pp.1631–1642.

Vyazovskiy, V. V. & Harris, K.D., 2013. Sleep and the single neuron: the role of global slow oscillations in individual cell rest. *Nature Reviews Neuroscience*, 14(6), pp.443–451.

Walker, M.P. et al., 2002. Practice with Sleep Makes Perfect: Sleep-Dependent Motor Skill Learning. *Neuron*, 35(1), pp.205–211.

Walter, G., 1936. The location of cerebral tumours by electroencephalography. *Lancet*, 2, pp.305 – 308.

Waters, J. & Helmchen, F., 2004. Boosting of Action Potential Backpropagation by Neocortical Network Activity In Vivo. *Journal of Neuroscience*, 24(49), pp.11127–11136.

Watson, C. & Paxinos, G., 2012. *Paxinos and Franklin's the Mouse Brain in Stereotaxic Coordinates, Fourth Edition*,

Whitlock, J.R. et al., 2006. Learning induces long-term potentiation in the hippocampus. *Science (New York, N.Y.)*, 313(5790), pp.1093–7.

Wild, J. et al., 2012. Performance comparison of extracellular spike sorting algorithms

for single-channel recordings. *Journal of Neuroscience Methods*, 203(2), pp.369–376.

Williams, K., 1993. Ifenprodil discriminates subtypes of the N-methyl-D-aspartate receptor: selectivity and mechanisms at recombinant heteromeric receptors. *Molecular pharmacology*, 44(4), pp.851–9.

Wilson, C.J. & Groves, P.M., 1981. (Accepted January 29th, 1981). , 220, pp.67–80.

Wilson, M.A. & McNaughton, B.L., 1994. Reactivation of hippocampal ensemble memories during sleep. *Science (New York, N.Y.)*, 265(5172), pp.676–9.

Wolansky, T. et al., 2006. Hippocampal Slow Oscillation: A Novel EEG State and Its Coordination with Ongoing Neocortical Activity. *Journal of Neuroscience*, 26(23), pp.6213–6229.

Woodhall, G. et al., 2001. NR2B-Containing NMDA Autoreceptors at Synapses on Entorhinal Cortical Neurons. *Journal of Neurophysiology*, 86(4), pp.1644–1651.

Wulff, K. et al., 2010. Sleep and circadian rhythm disruption in psychiatric and neurodegenerative disease. *Nature Reviews Neuroscience*, 11(8), pp.589–599.

Xu, M. et al., 2015. Basal forebrain circuit for sleep-wake control. *Nature Neuroscience*, 18(11), pp.1641–1647.

Xu, T. et al., 2009. Rapid formation and selective stabilization of synapses for enduring motor memories. *Nature*, 462(7275), pp.915–9.

Yamawaki, N. et al., 2014. A genuine layer 4 in motor cortex with prototypical synaptic circuit connectivity. *eLife*, 3.

Zadra, A. et al., 2013. Somnambulism: clinical aspects and pathophysiological hypotheses. *The Lancet. Neurology*, 12(3), pp.285–94.

Zhdanova, I. V. et al., 2001. Melatonin promotes sleep-like state in zebrafish. *Brain Research*.

Zwaka, H. et al., 2015. Context Odor Presentation during Sleep Enhances Memory in Honeybees. *Current Biology*, 25(21), pp.2869–2874.