UC San Diego
UC San Diego Electronic Theses and Dissertations

Title

Acetylcholine, brain activation, and processing of temporal stimuli

Permalink

https://escholarship.org/uc/item/4fh9x9vc

Author

Minces, Victor

Publication Date

2010

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Acetylcholine, Brain Activation, and Processing of Temporal Stimuli

A dissertation submitted in partial satisfaction of the

requirements for the degree Doctor of Philosophy

in

Biology with Specialization in Computational Neurobiology

by

Victor Minces

Committee in charge:

Professor Andrea Chiba, Chair
Professor Jim Conner
Professor Tim Gentner
Professor Doug Nitz
Professor Terry Sejnowski
Professor Mark Tuszynski

2010

The dissertation of Victor Minces is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

_____

_____
Chair

University of California, San Diego
Spring, 2010

Dedication

To my Parents.

Epigraph

Solo se puede enseñar el amor de algo.

Jorge Luis Borges

Table of Contents

Chapter 1

Acetylcholine and Perception of Time Modulated Stimuli – Visual System

Chapter 2

Acetylcholine and Perception of Time Modulated Stimuli – Auditory System

Chapter 1

Acetylcholine and Perception of Time Modulated Stimuli – Visual System

Chapter 2

Acetylcholine and Perception of Time Modulated Stimuli – Auditory System

Chapter 3

Neuronal Processing of Time Modulated Stimuli across Brain States

Chapter 4

Influence of Mean Neuronal Firing Rate on Measures of Rate-Correlation

List of Figures

Acknowledgements

This is the second version of my acknowledgments; the first one was on the border of cornyness, (on the corny side), which is hard not to be when giving thanks to Andrea, for allowing me to grow in all directions, for explaining to me the wonders and horrors of rats and human-like behavior, for her service to the community and for teaching me with the example that we can grow more by building together. I will stop here before a third version is necessary.

Thanks to Laleh for being my friend, for limiting the poetic licenses of this thesis and for managing the lab into absolute safeness.

Thanks to Andrea, Lara and Laleh, for making me roll on the floor laughing for so many years. To the students, Andy, Mitch, Steph and Duan Duan for letting me teach them.

Thanks to everyone in Harris' Lab for welcoming me, teaching me, and showing me the wonders of Thai Cuisine. To Alon for training me on the art of persistence and for receiving me each morning with that smile that goes twice around his head. To Ken for opening to me the doors of his lab and teaching me that no analysis is perfect.

To Jaime de la Rocha, Jude Mitchel and Pablo Jercog for their valuable input on Chapter 4.

Special thanks to Pablo, for being an amazing friend and a passionate and stimulating scientist.

Thanks to my family, Mario, Graciela, Lau, Pachu, Kari, Tatu, Eli, Diego, Lau and Male, just for being.

Vita

Sep. 1992 – June 1994 Fine Arts Studies at at "Escuela Nacional de Bellas Artes Prilidiano Pueyrredon", Buenos Aires, Argentina.

1995-2001   Studies in Physical Sciences at Universidad de Buenos Aires, Buenos Aires, Argentina.

August 2001 – Laboratorio de Sistemas Dinamicos; Dr. Gabriel Mindlin; Universidad de Buenos Aires, Facultad de Ciencias Fisicas

Jan 2002 – June 2002 – Internship at Laboratoire Leibniz Dr. Mirta Gordon - Institut d'Informatique et Mathematiques Appliquees de Grenoble (IMAG) – Grenoble, France

Fall 2005, Winter 2006, Fall 2006 Teaching Assistant for Patricia Zucker UCSD, Department of Linguistics

2003 – Aug 2010 Graduate student researcher, Dr. Andrea Chiba, Department of Biology, University of California, San Diego

Jun 2010 Ph.D., Biology with Specialization in Computational Neurobiology, University of California San Diego

ABSTRACT OF THE DISSERTATION

Acetylcholine, Brain Activation, and Processing of Temporal Stimuli

by

Victor Minces

Doctor of Philosophy in Biology with Specialization in

Computational Neurobiology

University of California San Diego, 2010

Professor Andrea Chiba,Chair

In our daily lives we have to process sensory information at different speeds and in a sensory specific manner. For example, if we go for a walk in the park we will need to be aware of the relatively slow pace of changes in the environment, but as we read this text we will need to process the characters that pass through our fovea in rapid succession. A child might be aware of her teacher starting or finishing a discussion, but only if she is interested will she follow the ongoing flow of details. Because of its anatomical characteristics and its well established role in the control of global brain activity, the neurotransmitter acetylcholine is likely to play a large role in such temporal processing of information.

This work deals with two aspects of the sensory processing of time modulated stimuli, particularly as it relates to acetylcholine. Chapters 1 and 2 analyze acetylcholine's influence on an animal's ability to discriminate between time modulated flashes of light (chapter 1) or trains of sounds (chapter 2). Chapter 3 analyzes the influence of brain state, a phenomenon tightly related to cholinergic activity, on the neural processing of trains of sounds, both at the thalamic and cortical levels. Finally, chapter 4 undertakes the problem of the influence of mean neuronal firing rate on measures of neural representation. The mean firing rate in the thalamus differs dramatically between states, and analysis of correlative activity must take this into consideration. Our theoretical results can explain various aspects of the current literature on the subject.

Introduction

Introduction to Chapters 1 and 2: Acetylcholine and Perception of Time Modulated Stimuli

As cholinergic neurons in the basal forebrain innervate discrete regions of the cortical surface, the cholinergic system has the potential to dynamically modulate different areas of the cortex according to behavioral demands. Although it has been well established that acetylcholine plays a role in modulating neuronal activity and plasticity in the sensory cortices, its behavioral role in these areas has remained more elusive. This is in part because lesion studies have focused on broad cholinergic depletions as opposed to lesions that are specific to the sensory modality under investigation. While work in the olfactory system has shown an influence of acetylcholine in sensory perception, other experiments in the visual and auditory domains have failed to do so, only finding an influence on perceptual learning. This might be due to cross-modality differences but also to the behavioral paradigms and stimuli set chosen in each study.

A review of the literature suggests that ACh might be particularly important for processing time modulated stimuli; therefore, it is fundamental to assess sensory perception along this dimension.

In Chapter 1 we asses the ability of animals with cholinergic depletions specific to the visual cortex to discriminate, and learn to discriminate, sequences of flashing lights. In Chapter 2 we asses the ability of animals with cholinergic depletions specific to the auditory cortex to discriminate sequences of sounds.

We find that while all animals can readily learn to discriminate between stimuli that are very different, controls can AChieve a higher level of sensory refinement. This effect is not present if the lesions are performed once this refinement has taken place, indicating an influence of ACh in perceptual learning but not in perception per se.

Introduction Chapter 3: Neuronal Processing of Time Modulated Stimuli Across Brain States

We are capable of fairly complex sensory processing during sleep. We know from every day experience, as well as from scientific studies, that we are more likely to wake up when we hear our own name than that of someone else (Langford, Meddis et al. 1974), that sleeping mothers will awaken to their new born child's cry more frequently than to other's (Fromby 1967), and that we will accommodate to the sound of the heater but will wake up to an opening door. This type of complex processing most likely requires the auditory cortex (Diamond and Neff 1957), a fact that might seem at odds with popularly held

views of thalamo-cortical activity across the sleep wake cycle which consider the thalamus a gate that is closed during sleep.

Spontaneous brain activity can be broadly categorized by its level of activation. Under this categorization there is an activated state, characteristic of wakefulness and REM sleep, and an inactivated state, typical of slow wave sleep (Jasper 1936). The features that characterize these states involve differences in the activity of the thalamus, cortex, and neuro-modulatory systems affecting them, most prominently the cholinergic system (Steriade and McCarley 2005) Although inactivation and activation are characteristic of sleep and wake conditions respectively, they are not exclusive to them, as both can be expressed in the wakeful animal (Stoelzel, Bereshpolova et al. 2009) as well as in anesthetized preparations (Dringenberg and Vanderwolf 1998).

The notion that the thalamus functions as a "gate" that controls the flow of sensory information to the cortex has become part of neuroscientific folklore. According to this view, when an animal is awake, the thalamus receives sensory information and relays it to the cortex. When the animal falls asleep the thalamic responses to sensory stimulation are much diminished, and therefore the cortex does not receive sensory information. This notion is of course founded on good experimental evidence, but a careful review of the literature reveals that it might be a simplification. One feature that emerges clearly out of the sometimes contradictory literature is that the degree of thalamo-cortical response is often linked to the temporal characteristics of the

stimulus. So the question of how much information the thalamus relays in each state should be re-formulated as: what type of information does the thalamus relay in each state? Furthermore, as brain activation also involves changes that are intrinsic to the cortex, the question of how these changes affect the cortical response to thalamic inputs, an issue about which the literature has been inconsistent, needs to be answered.

The view of the thalamus as a sensory gate derives from experiments almost exclusively performed in the visual system, with the somatosensory system a distant second. As the examples presented at the beginning suggest, it might be a simplification to extrapolate results obtained in the different sensory systems, because the inputs across sensory modalities are differentially affected by sleep. In effect, the visual inputs are literally shut off, the somatosensory inputs are greatly diminished by immobility and lack of active exploration, (think of the whiskers), and the auditory inputs remain unaffected. In this context, neural processing of auditory signals during sleep could be expected to be different than visual or somatosensory processing.

The goal of this chapter is to understand how brain state affects auditory responses in the thalamus and the cortex, particularly as it relates to temporally modulated stimuli. Our results indicate that after an onset period in which the activational state doesn't have a profound influence in cortical responses, there is a sustained period in which the cortical response is severely degraded in the inactivated state. By contrast, thalamic responses are not impaired by inactivation. These results not only show that the auditory

system responds in a very different way than its thalamic counterparts, but also that activation can have a strong and specific action on cortical responses.

Introduction Chapter 4: Influence of Mean Neuronal Firing Rate on Measures of Rate-Correlation

When analyzing the data from chapter 3 we find the problem of the influence of overall rate in a commonly used measure of neural coding, the correlation between spike trains. In this chapter we make an in-depth analysis of the problem. We find analytical equations relating rate and correlation under simplified conditions, and set a framework to analyze more complex and realistic cases. Our results serve to explain various aspects of the literature and help interpret the technique of rate-equalization.

Background

Cholinergic System, Basal Forebrain Anatomy

General Anatomy of the Basal Forebrain:

The basal forebrain cholinergic neurons are organized along a rostrocaudal axis, such that cholinergic "nuclei" are somewhat contiguous in

nature (Saper, 1984). The majority of behavioral and anatomical studies

subdivide the basal forebrain into three primary projection systems: the

horizontal limb of the diagonal band of broca in combination with portions of

the magnocellular preoptic area (HDB/MCPOA), the corticopetal basal

forebrain, including the nucleus basalis magnocellularis, the substantia

innominata, and portions of the ventral pallidum (NMB/SI), and the medial

septal nucleus in combination with the vertical limb of the diagonal band of

broca (MS/VDB). Cells within the HDB/MCPOA project primarily to the

olfactory bulb and entorhinal cortex, with light projections to the cingulate

cortices, occipital cortex and the amygdala.

    Neurons within the MS/VDB complex provide major cholinergic

innervation of the hippocampus, cingulate cortex, and subiculum, with a lighter

projection to the entorhinal cortex.

The NBM/SI contains a continuum of projection neurons that provides the

major source of cholinergic input and some non-cholinergic input (mostly

GABA-ergic) to nearly all neocortical regions. Additionally, it provides a

primary input to the basolateral amygdala and to the reticular nucleus of the

thalamus. (for a detailed review see Zaborszky L et al.. 1999) .

    The cholinergic basal forebrain (CBF) region serves to regulate cortex

and, at the same time, receives input from non-isocortical paralimbic cortical

areas (orbitofrontal-prefrontal, temporopolar, insular, parahippocampal, and

cingulate), the amygdala, the hypothalamus, and various brain stem cell

groups. As there is no evidence of direct thalamic projections from the

thalamus to CBF, sensory information must reach the latter through either the prefrontal cortex or the central nucleus the amygdala, although other unreported pathways cannot be discarded.

Additionally, the CBF makes reciprocal connections with several structures, including the central nucleus of the amygdala, ventral tegmental area , locus coeruleus, the hypothalamus and PPT (a brain stem cholinergic nucleus that also projects to the thalamus) (Grove, 1988)

By virtue of a convergence of afferent brainstem projections to the basal forebrain, diverse neurochemical systems can heavily modulate the basal forebrain and its constituent cortical targets. Dopaminergic influence can occur through projections arising from the VTA or by projections arising directly from the substantia nigra pars compacta, substantia nigra pars reticulata or the retrorubral field (Semba et al., JCN, 1988). Serotonergic influence gains access to the CBF by way of dense projections arising from the dorsal raphe nucleus, lighter projections from the median raphe nucleus or sparse projections directly from the raphe magnus (Semba et al., JCN, 1988; Gasbari et al, Neuroscience, 1999). The basal forebrain can receive noradrenergic influence either from the locus coeruleus or from the nucleus of the solitary tract (Semba et al. JCN, 1988, Espana & Berridge JCN, 2006). Finally, the brainstem cholinergic system that is also referred to as the pontomesencephalic cholinergic system (Shute and Lewis, 1967; Woolf, 1991) provides cholinergic input to the basal forebrain primarily by receipt of projections from the laterodorsal tegmental nucleus (LDT) (Satoh & Fibiger,

1986; Woolf & Butcher, 1986; Semba et al., 1988) and also from the pedunculopontine tegmental nucleus (PPT) (Semba et al., 1988).

Specificity of Efferent Connections of the Corticopetal Basal Forebrain:

There is much evidence that CBF cortical projections are at least roughly topographic.

Lesions in different regions of the CBF produce different amounts of cholinergic depletion in various cortical areas. For example, more posterior lesions produce bigger depletions in the auditory cortex (ACx) than in the visual cortex (VCx), whereas the converse is true for lesions that are more anterior. (Wenk H, Bigl V, Meyer U., 1980). In accordance with this, electrical stimulation in posterior zones produces EEG activation in auditory areas (a phenomenon that is related to ACh release). This effect is abolished by iontophoretically applied atropine. Stimulation of anterior zones does not produce the same effect. (Metherate R, Cox CL, Ashe JH., 1992) Stimulation of different regions of the CBF produces differential acetylcholine release in motor or visual cortices. (Jimenez-Capdeville ME, Dykes RW, Myasnikov AA.,1997).

Evidence also exists that cortical sensory regions can be differentially modulated through a circuit including the prefrontal cortex (PFC) projection to CBF and, subsequently, the CBF projection to sensory cortices. (Golmayo L, Nunez A, Zaborszky L, 2003)

Neurophysiological Effect of ACh:

The following results indicate that generally ACh has a positive modulatory effect on cortical responses, although sometimes the effect is negative. These experiments differ with respect to the behavioral state of the animals, the species, and the timing of stimulation. It would be interesting to tease these effects apart, in order to provide final clarification of the issues.

In the anesthetized cat, somatosensory cortical neurons are generally excited by iontophoretic administration of ACh or by CBF stimulation, (electrical or through applied glutamate) (Tremblay N, Warren RA, Dykes RW.-a,1990). These types of cholinergic stimulation positively modulate cortical responses to cortically applied glutamate or to somatosensory stimulation (Tremblay N, Warren RA, Dykes RW.-b,1990). A similar effect is obtained in the auditory system (Metherate R, Weinberger NM., 1990)

In the awake rat, stimulation of the BF has been shown to reduce the ACx response to a tone. This reduction is not blocked by atropine, so it might be due to GABAergic cortical projections (Edeline JM, Hars B, Maho C, Hennevin E.,1994). In anesthetized rats, ACx response to auditory thalamus stimulation is enhanced or depressed when paired with CBF electrical stimulation. This enhancing or depressing effect depends on the magnitude of the CBF stimulation. (Metherate R, Ashe JH.,1991). In the anesthetized rat, cortical LFP response to somatosensory or visual stimuli is enhanced when

paired with CBF stimulation. (Golmayo L, Nunez A, Zaborszky L.,2003) In the awake rat, simultaneous CBF stimulation enhances whisker motor response to Motor Cortex electrical stimulation.(Berg RW, Friedman B, Schroeder LF, Kleinfeld D.,2005).

The influence of ACh in sensory responses will be discussed in further detail when analyzing the more general phenomenon of brain activation.

ACh and Neural Plasticity:

Much evidence indicates that ACh is involved in Cortical Plasticity.

In the anesthetized cat, cortical iontophoresis of ACh or CBF stimulation potentiates cortical responses to application of glutamate or to somatosensory stimulation. (Tremblay N, Warren RA, Dykes RW.-a and b,1990) Cortical ACh iontophoresis also increases auditory cortex response to a tone. It does so specifically to the tone that has been paired with the ACh application, changing in this manner the receptive field of the cell that is being measured. (Metherate R, Weinberger NM.,1990)

Both in the awake and anesthetized rat, pairing BF stimulation with a tone presentation further increases the ACx response to the tone alone, and it does so specifically to the paired tone. This increase is abolished when the pairing has been done under systemic atropine (Edeline JM, Hars B, Maho C, Hennevin E.,1994; Bakin JS, Weinberger NM., PNAS 1996; Miasnikov AA, McLin D 3rd, Weinberger NM.,2001). As an expected correlate of this tone-

specific potentiation, pairing BF stimulation with a tone modifies the cortical tonotopic representation so that the cortical area that best responds to that tone is extended (Kilgard MP, Merzenich MM.,1998).

Lesion experiments show that excitotoxic (not selective for ACh) lesions of nucleus basalis (NB) prevents somatosensory cortex reorganization following digit amputation or peripheral nerve lesion (Juliano SL, Ma W, Eslin D.,1991; Webster HH, Hanisch UK, Dykes RW, Biesold D.,1991). Cholinergic depletions to the BF mediate sensory plasticity associated with whisker pairing (Baskerville KA, Schweitzer JB, Herron P.,1997; Sachdev RN, Lu SM, Wiley RG, Ebner FF.,1998; Zhu XO, Waite PM.,1998). Selective blockade of cholinergic signaling within the barrel cortex impairs whisker pairing plasticity (Maalouf M, Miasnikov AA, Dykes RW. 1998). ACh does not only induce changes in topographic representations. The plastic effect of this neurotransmitter can also be associated with temporal patterns, as shown by the fact that barrel cortex responses to trains of whisker deflections are specifically potentiated to frequencies that have been paired with cholinergic iontophoresis(Shulz, Sosnic et al. 2000)..

While none of the examples above show a link between the ACh-mediated plasticity and behavior, Conner and colleagues show that CBF ACh depletion abolishes motor-cortical map plasticity associated with motor skill learning and impairs this type of learning (Conner JM, Culberson A, Packowski C, Chiba AA, Tuszynski MH.,2003).

Cholinergic System and Behavior:

The basal forebrain corticopetal cholinergic system has historically been broadly implicated in learning and memory (Wenk – 100 years of cholinergic function). Yet, with advances in selective lesion techniques, more recent behavioral and neurobiological data suggest that cholinergic neurons in the basal forebrain help to regulate attention to relevant stimuli in the environment.

The integrity of the CBF cholinergic neurons is required for the increased processing of stimuli with uncertain consequences (Bucci, 1998; Chiba, 1995), a strategy that has been proposed to underlie learning the predictive associations between stimuli (Dickinson, 1977; Pearce, 1988; Hall, 1982; Holland, 1993). In addition, CBF cholinergic neurons contribute to the fluid allocation of attention (Parasuraman, 1992; Baxter, 1999) and promote shifts of attention to less predictable stimuli (Moore, 1999; McGaughy, 2002). In light of these findings, computational theorists have proposed that cholinergic neurons in the basal forebrain carry out their role in attention by modulating stimulus processing according to the probabilistic uncertainty of stimulus expectations. (Yu and Dayan 2005). Thus, the basal forebrain may modulate attention for the purpose of rapid plasticity and learning.

Although an effect of ACh in sensory cortices' plasticity has been well documented, very few experiments have analyzed its behavioral effect in sensory processing in the visual (Dotigny et al. 2008, Fine et al. 1997) and

auditory (Kudoh et al. 2004) systems. These studies find that depleting the cortex from cholinergic input produces a learning deficit. If the lesions, however, are produced after learning has taken place, the animals show no deficit in sensory perception. A different result seems to emerge in the olfactory system, where cholinergic depletions including the olfactory bulb and the (olfactory) piriform cortex, produce changes in odor generalization.(Linster, Garcia et al. 2001)

The different results found across these systems might be related to their radically different anatomical connectivity, but also to the analytical techniques employed to evaluate sensory perception. This will be further discussed when explaining the rationale of chapter 1.

In spite of the fact that, as mentioned before, cholinergic cells innervate discrete areas of the cortical surface, the experiments discussed above asses the effect of relatively large cholinergic depletions (involving large cortical or limbic regions), leaving open the possibility that some of the changes observed are not related to the area of interest but to larger neuronal networks. An exception has to be made in the case of Kudoh's work in the auditory cortex, where it is shown that cholinergic lesions restricted to the auditory cortex impair the rats' ability to discriminate sequences of sounds. The details of this work, in relation to our results, will be further discussed.

Conclusion: The Cholinergic System has been shown to be implicated in attention, neural plasticity and learning. Except for results in the (anatomically distinct) olfactory system, previous work has failed to reveal a

direct influence of cholinergic activity in sensory perception. The anatomical characteristics of the this system put it in a favorable position to dynamically modulate the activity across different brain areas, in spite of this, lesion studies have focused on depletions of relatively large areas and very little attention has been given to the influence of acetylcholine on sensory cortices.

Brain States and Sensory Processing

Overview of the States:

Based on the EEG, cortical activity can be broadly categorized into different groups, at the extreme of which are the activated and the inactivated states (Moruzzi and Magoun 1949). The activated state is characteristic of awake and attentive behavior and of REM sleep, and the inactivated state is characteristic of deep stages of non-REM sleep, or slow wave sleep (SWS). These different types of activity were first described in humans by Berger (Berger 1929).

Although the states were originally described in terms of cortical electrical activity, it was later understood that they reflect global brain differences. A brief description of the thalamo-cortical activity that characterizes them is described below, alongside the general conditions under which these states can be present, as well as the brain structures involved in their generation.

Characteristics of the States:

Inactivated State

The inactivated state has been typically characterized by its high amplitude, low frequency EEG. This EEG activity is a reflection of the global activity in the cortex, in which the activity of neural populations alternates between periods of global silence (down states) and periods of global activity (up-states). These alternation occurs at relatively low rates (<5 Hz) which gives the EEG its characteristic slow oscillation appearance, and hence the term slow wave sleep (SWS).

Intracellularly, it can be seen that cortical cells tend to be depolarized during the up states (occasionally firing a few spikes) and hyperpolarized during down states, (see figure Aux.1). As the occurrence of these up and down states is global, all cells within an area tend to be silent or to be firing, therefore their firing is correlated, and the multi unit activity looks like bursts of spikes lasting around 200 ms. These bursts of spikes don't happen at the same time in the whole cortex, but tend to start at a given site and then propagate, producing cortical waves (Sanchez-Vives and McCormick 2000; Massimini, Huber et al. 2004) Because cells tend to fire together, the inactivated state has been referred to as synchronized. (McCormick 1992).

In the thalamic-relay nuclei, the typical activity in the inactivated state is a little different. In this case the thalamic cells tend to fire short bursts of two to

five spikes at a frequency of about 200 Hz. This type of activity is often referred to as "bursty mode". (McCormick 1992).

Activated State

The activated state is characterized by an EEG with less fluctuations, less power in the low frequencies (<5 Hz) and more power in the high frequencies. During this state, cortical neurons tend to fire more tonically, and show less correlation between them, although their firing rate is not necessarily higher than during the inactivated state. Periods of global silence, associated with down-states, are not so frequently observed.

In the thalamus, the activated state is also characterized by tonic firing. In this state the cells tend to fire only one action potential at a time (as opposed to bursts), and the mean rate is higher than in the inactivated state. (Figure Aux.1)

General Observations

In opposition to the so-called synchronized state, the activated state has been referred to as non-synchronized. These terms might be misleading because in some situations the cortical cells have been shown to be synchronized at higher frequencies (Singer 1998). It should be mentioned that although during the inactivated state there is much higher incidence of thalamic bursts, this type of activity can be observed during states associated with activation, and the reverse is also true, many cells in what is consider an inactivated state can fire tonically (Ramcharan, Gnadt et al. 2000; Massaux, Dutrieux et al. 2004).

Conclusion: The inactivated state is characterized by a slow oscillation in the cortex and burstiness in the thalamus, whereas the activated state is characterized by tonic activity in both structures.

Conditions of Activation and Inactivation, Relation to Acetylcholine:

Naturally Occurring Brain States

As previously mentioned, the activated state has been mainly associated with wakefulness, and the inactivated state with slow wave sleep. The reason REM sleep has traditionally been called paradoxical sleep is because the EEG exhibits an activity very similar to that of wakefulness. The same can be said of thalamic activity, which during REM is in the tonic mode. Spontaneous activity during REM sleep is, at least superficially, very similar to wakefulness.

Similarly, during inattentive wakefulness, brain activity can present characteristics of inactivation, like increased low frequencies in the EEG, presence of down-states and increased burstiness in the thalamus. (Stoelzel, Bereshpolova et al. 2009; Luczak, Bartho et al. 2007).

Experimentally Induced Brain States

Brian states can be induced experimentally; this has most frequently been done by stimulating or antagonizing certain structures, most notably the rostro-dorsal brainstem, particularly the brainstem reticular formation (BRF) and the basal forebrain (BF). Both of these structures, that are mostly active

during wakefulness and REM sleep, contain neurons producing acetylcholine, a neurotransmitter shown to be involved in the activational process. As BF and BRF have mutual connections (Steriade and Buzsaki 1990; Jones 2008) , stimulating one will most likely stimulate the other. The following section is a review of the different experimental conditions in which activation or inactivation can be induced, along with the neurotransmitters and structures involved.

A traditional way to "induce" inactivation is by transecting the brainstem from the rest of the brain, a preparation called "cerveau isole" Bremer (Bremer 1935), the effect of which was later shown to be present under more focal lesions of the BRF (Lindsley, Schreiner et al. 1950)

An anesthetized animal is typically in the inactivated state, although this depends on the level of anesthesia. In a recent study, Clement et al. (Clement, Richard et al. 2008) showed that under a very tightly controlled dose of urethane anesthesia, the state of an animal oscillates between activation and inactivation in a way that very closely resembles the alternation between SWS and REM. As a matter of fact, during the activated state the animal presents the same characteristic muscle atonia that can be observed under REM, a result that some say will potentially establish the urethane anesthetized animal as a model of sleep.

If the anesthetic dose is higher, the animal tends to remain in a state similar to SWS. In this situation, the brain can be activated by stimulation of BRF, a discovery made by Moruzzi and Magoun (1949) which greatly

facilitated research on the subject of activation.. It was later shown that stimulation of many other structures can induce activation in the anesthetized animal. These include the intralaminar thalamic nuclei, locus coeruleus, raphe nucleus, prefrontal cortex, superior colliculus and basal forebrain (Dringenberg and Vanderwolf 1998). These structures are all associated with neuro-modulatory systems and are deeply interconnected (Buszaki and Steriade 1990).

Neuromodulatory systems show different activation patterns associated with behaviors. During wakefulness the cholinergic BF and BRF, the adrenergic locus coeruleus, and the serotonergic raphe nucleus are all active. During SWS all these modulators diminish their response, and during REM the cholinergic are very active whereas the locus coeruleus and raphe nucleus are silent.

Activation through stimulation of the BRF, BF, locus coeruleus, amygdala, and prefrontal cortex, is blocked by systemic application of cholinergic antagonists (Dringenberg and Vanderwolf 1997). Although activation is a global state, it can be locally turned into an inactivated state by micro-iontophoresis of the cholinergic antagonist atropine in the cortex (Metherate, Cox et al. 1992). This suggests that systemic administration of cholinergic blockers might work at the cortical level, although additional sub-cortical action cannot be ruled out.

Further support for a role of ACh in activation comes from the fact that activation can be induced by systemic cholinergic agonists (Bonnet and

Bremer 1937). ACh cortical levels are increased in the activated state, and local application of ACh in the cortex induces changes in the EEG similar to those of activation (Spehlmann and Norcross 1982). Also, application of ACh in thalamic cells can take them from the bursting to the tonic mode (McCormick 1992).

Although stimulation of both the BRF and the BF can induce activation, the former proves not to be necessary, as animals under chronic cerveau isole can recover the capacity to activate (Batsel 1960) and stimulation of the BF can induce activation even in this preparation (Belardetti, Borgia et al. 1977). When the BRF is stimulated in the BF lesioned animal, activation is greatly reduced, although some remnant effect is present, probably mediated by the intralaminar thalami nuclei (Dringenberg and Olmstead 2003). Furthermore, Alzheimer's disease, which initially destroys neurons in the BF, increases power in the low frequencies of the EEG, an effect that is mimicked in animals with total BF lesions (Stewart, Macfabe et al. 1984).

Serotonin might also have a role in the arousal process. In the behaving animal, after systemic application of cholinergic antagonists, activation is abolished when the animal is quiet, but persists upon locomotion. A combination of cholinergic and serotonergic antagonists can get rid of all activation (Vanderwolf and Baker 1986).

Activation can also be induced by electrical stimulation of the raphe nucleus, a structure containing serotonergic cells. This activation is not abolished by cholinergic antagonists, but is abolished by serotonergic

antagonists (Dringenberg and Vanderwolf 1997). Activation can also be induced by pinching the tail of the animal, a process that is also blocked by antagonizing serotonin (Neuman and Thompson 1989).

These data seem to indicate that the ultimate agents in cortical activation are acetylcholine and serotonin. However, it should be noted that cholinergic neurons also synthesize nitric oxide, and blocking this synthesis has a similar effect in activation as applying atropine (Marino and Cudeiro 2003), raising the possibility that the cholinergic induced activation is mediated by this gas and its possible role in cerebral blood flow. Also, selective cholinergic depletions of the BF have generally failed to reduce activation in the cortex (Nilsson, Leanza et al. 1992; Kaur, Junek et al. 2008), but not in hippocampus (Buzsaki, Bickford et al. 1988). These data indicate that the picture might be a little richer and that the relationship between neuromodulators and brain state is not completely understood.

Conclusion:

Many structures and neurotransmitters, all deeply interconnected, seem to act on forebrain activation; however, cholinergic projections arising in the basal forebrain and serotonergic projections arising in the raphe nucleus seem to have a privileged role.

Sensory Processing Across the States:

Guidelines

It is frequently stated that the role of the thalamus is to control the flow of sensory information to the neocortex during different behavioral states (Castro-Alamancos 2004). It is clear, behaviorally, that information is processed differently across the sleep-wake cycle and it is well established that thalamo-cortical activity differs across the states. However, the above statement should be taken carefully, and control should not be interpreted as "shutting off" the relay of information, but should be viewed in light of the sensory system analyzed and the kind of information that is being transmitted. Here we present what we consider are some criteria along which sensory processing across the states should be analyzed but have frequently been overlooked, then we review the literature focusing on these criteria.

Difference between Sensory Systems:

Because the role of each sensory system changes differently during sleep, care should be taken in using any system as a general model of sensory processing across brain states. Vision, which most of the studies have focused on, is physically shut off during sleep; somatic processing is reduced by immobility and lack of exploration (think of the whiskers). Meanwhile, auditory inputs are not blocked (peripherally) in any form, and there is plenty of evidence that a substantial amount of complex auditory processing can be exerted during all sleep stages.

The processing of time series, which is arguably fundamental when analyzing state-dependent sensory processing, is very different between and within the systems. For example, in the Visual System a single static image, a picture, can convey meaning. In the somato-sensory system, a rat can discriminate textures of objects with his whiskers based on a temporal series of sticks and slips (Mehta and Kleinfeld 2004), but most likely the ability of the animal to perform the same operations with his tail will not be as good, which reveals a within-modality difference in the ability to process time modulated stimuli. Auditory signals are temporally rich, in the sense that in order to discriminate complex stimuli like our name or the call of a pup, a static cochlear input will not suffice.

Conclusion: Based on the different ways that sleep affects our sensory systems and the differences among them in processing temporally modulated signals, we should not expect that the activational state affects all systems in the same manner.

Quality vs. Quantity:

While studying the effect of activation on sensory responses, much work has focused on absolute changes in evoked firing rate, that is, whether neurons fire more to a stimulus in one state or the other. These differences are of course important, but the relationship between the evoked activity and the spontaneous activity, as well as the "ability" of the evoked activity to represent the stimuli, should also be taken into consideration.

A useful metaphor might be that of a cheap speaker. The quality of the sound might be good at a low volume, but distortions in the signal and increased noise will deteriorate it when we turn the volume up. If we are close enough to the speaker we will be content with a low volume-high-quality situation, but if we are farther we might need to sacrifice one for the sake of the other. The maximum information that we can potentially extract from the speaker will be that of playing it at a low volume but staying close. Back to the brain, we might never be able to figure out how the activity in a structure is "read" by its target. It is clear that a strong signal will help this reading, but the signal to noise will present an upper bound to the amount of information that can potentially be extracted.

Timing of the Stimulus and the Response:

It was discovered early on that the effect of brain state on neural responses is highly dependent on the frequency of presentation of the stimuli (Steriade and Demetrescu 1960). For example, stimuli presented at low frequencies are affected by arousal in a different way than stimuli presented at high frequencies. In the case in which the stimulus is constant for an extended period of time, the neuronal responses can typically be divided into an onset and a sustained phase, which can be modulated differently by arousal. It has to be noted that from the point of view of a single neuron, stimulation at high frequencies is equivalent to constant stimulation, since both will be expressed as a fast series of synaptic inputs. Some frequently used stimuli, such as

visual gratings or sine-wave amplitude modulated sounds, might have a confounding effect when analyzing the timing of responses, because decreasing the period between events also decreases the rising time of each event. Such a problem is avoided when using trains of flashes or clicks (Joris, Schreiner et al. 2004)

Timing should be kept in mind when analyzing the literature because the strength of response will depend on the characteristics of the stimuli and the time after stimulus onset. The influence of timing also becomes important when assessing receptive fields, since the receptive field at the onset of a stimulus might be quite different than at the sustained phase.

Type of activation:

Studies have used different "techniques" to induce states. Some have used electrical stimulation of different activating nuclei, some have compared stages of more or less thalamic burstiness or found periods of spontaneous activation within the anesthetized preparation, and some, of course, have looked at the natural sleep-wake cycle. From the latter group, many studies have used paralyzed, sleep deprived cats, and it should probably be a matter of debate to what extent a sleep deprived animal with dilated pupils and lights being flashed at his eyes can be used as a model of natural sleeping. As a matter of fact, it has been estimated that the heart rate of a monkey under this preparation is three times faster than in the normal condition (Baker 1971).

Some studies have compared awake vs. anesthetized animals; in those cases the depth of anesthesia should be considered when analyzing the results. Even so, the results across the different types of activation have been consistent in many ways, and some work went as far as proving that the results obtained in the anesthetized preparation remain valid across the sleep-wake cycle. It is not clear that we should generalize all that we learn from anesthetized preparations into naturalistic behavior. Understanding the extent to which that generalization is possible would be a great tool for researchers.

Visual System

In 1960, Steriade and Demetrescu (Steriade and Demetrescu 1960) showed, in the visual system of the cat, a pattern that would be frequently repeated across different experimental procedures and modalities. Looking at field potentials, they found that activation has a depressing effect on the cortical response to flashes of light, concomitant with a potentiation of the thalamic response. This result was valid for single flashes, or equivalently, for flashes presented at low frequencies. When the frequency of presentation was augmented to more than 5 Hz the result in the thalamus stayed the same, but the cortical pattern was reversed, that is, the cortex had a stronger response to trains of flashes presented at high frequency when the brain is activated.

Inspired by this work, Evarts (Evarts 1963), made what, to my knowledge, is the only experiment in which light evoked responses are

analyzed in a controlled manner in the naturally sleeping and non-paralyzed animal. He did so by inserting a light panel under the eyelid of a chronically implanted freely moving cat. Presenting flashes at 5 Hz he found a result that is consistent with the Steriade study. Cortical neurons have a stronger response during wakefulness than during SWS. The differences are most prominent at long (>80 ms) and not very prominent at short (<40 ms) latencies, which again stresses the importance of timing. He did not use other stimulation frequencies nor did he measure responses in the thalamus.

The widely cited work by Livingston and Hubel (Livingstone and Hubel 1981), in the paralyzed cat, was the first that looked at changes in receptive fields in the visual system. They found that receptive fields in the thalamus and cortex are generally unchanged, although they did not quantify changes in the receptive fields' sizes.

It is very interesting that they not only found that arousal increases the LGN neural response for stimuli in the receptive field's center, but also increases inhibition from the surround. This means that state dependent modulation of the response is not simply due to changes in excitability. They also found that in the cortex the cells respond more strongly to the stimuli when in an activated state. This difference from the result by Steriade and Demetrescu might be related to the type of stimulus. Steriade uses flashes, which have an onset that is well defined, whereas Livingston and Hubel use a slowly moving grating which allows adaptation.

A study from Singer's lab (Francesconi et al. 1988), finds that X and Y cells in the LGN are differentially affected by reticular formation stimulation. While in the activated state, X cells' activity is augmented in both the onset and the tonic phase of their response to a solid light, Y cells are only affected in the tonic phase. He also compares the effect of RF stimulation and application of ACh, which will be discussed later.

Worgotter et al. (1998), in the spontaneously activating brain of an anesthetized cat, shows slightly different results. They confirmed that in the thalamus, activation increases the tonic part of the response to a solid light, but the phasic response is actually increased in the inactivated state. As they noticed this temporal evolution in thalamic responses, they went ahead to see how the cortical receptive fields change over time, finding that at onset the receptive fields are big but that they shrink during the sustained part of the response. Unfortunately, they only performed this analysis in the activated state, and only in the cortex, but they were the first to realize that there is a relationship between the timing of the response and the shape of the receptive fields.

None of the studies mentioned so far tackled the issue of how detectable the evoked response is in the different conditions. As mentioned in the introduction, this is very important because the signal to noise ratio presents a limit to the information that can be read by a target structure, which might be independent of the total intensity of the response. Guido and Sherman (Guido, Lu et al. 1995), are to my knowledge the first to tackle this

question. They measured detectability of the thalamic responses by measuring the ROC area (see methods). This method essentially measures the overlap between the evoked responses and the background. If the activity in response to a stimulus is always bigger than during baseline (no overlap) then the stimulus will be perfectly discriminable, whereas if the overlap is total, an observer will discriminate at chance. They find that when the LGN is in burst mode (consistent with inactivation) the stimuli are actually better detected than when it is in the tonic mode. Better detection does not necessarily mean that the stimulus is better represented in burst mode; they actually find that the response follows the stimulus more linearly when in the tonic mode, which can be interpreted as transmitting the details of the stimulus to the cortex, as opposed to just informing the cortex of the fact that there is a stimulus. This is probably the first work that situates the neural response in the bigger context of spontaneous activity and natural variability.

A more recent work by Dan's lab (Goard and Dan 2009) also focused on the quality of the response, but using a rather long stimulus, which suggests they focused on the sustained part of the response. In order to asses how accurately the neural response represents the stimulus, they looked at the trial to trial variability, that is, how much the response in a given presentation of the stimulus resembles the response in all other presentations. They did so by calculating the correlation of the activity across many presentations of the same stimulus. They find that activation increases the reliability of the response in both the thalamus and the cortex. They also

propose an analytical technique to "predict" the stimulus from the response. Consistent with the observed increase in reliability, activation increases predictability.

Conclusions: Whenever sustained responses or responses to stimulation at high frequencies were analyzed, the evoked activity was greater in the activated state. This is not true for onset responses or stimuli presented at low repetition rates. Only two experiments looked at the quality of the response, finding somewhat opposing results, which may be explained by differences in the timing of the stimuli.

Somatosensory System

When comparing somatosensory responses in the awake and anesthetized preparation, Poggio and Mountcastle (Poggio and Mountcastle 1963) showed that receptive fields in the somatosensory thalamus are largely unchanged. However, they realized that the activational state has an influence in the time evolution of the responses. In their words:

> These studies suggested that the static properties, those of topography and modality are unchanged by an anesthetic agent, though it remained uncertain whether the anesthetic suppressed a less orderly convergence, and thus falsely revealed a highly specific nature of the system by exempting only the most secure synaptic connections within it. The more dynamic aspects of the system, the temporal cadence of neural activity, the transformation of neural patterns of discharge at synaptic junctions, the quantitative relation of central response to peripheral stimulus, all these appeared to be most severely affected by anesthetic agent.

They found that thalamic cells in awake and anaesthetized animals would readily respond to a pulse of somatic stimulation, but would be depressed when a second pulse was presented at a short interval, a phenomenon called paired pulse inhibition. Thalamic neurons in the awake animal recovered from the first stimulus much faster than in the anaesthetized animal.

Baker (1971) found no consistent difference in responses when comparing somatosensory stimulation in the thalamus along the sleep-wake cycle. As his repetition rate was slow (2 Hz), we would not expect to see an effect related to adaptation, so this is consistent with the literature previously mentioned, and that to come.

In the paralyzed cat, Filion et al. (1971) measured ventro-lateral thalamic responses to electrical stimulation of its cerebellar afferents. They reported that thalamic responses were more consistent when the animal was awake. During sleep, responses tended to be bursty and show more variability in the latency of response, yet there was no overall trend in total evoked firing.

In what is probably the first work to look at the neural response to trains of impulses in the somatosensory system, Gucer (1979) found that cortical neurons of the monkey can lock to sensory stimulation at 100 Hz. Their ability to respond strongly depended on the behavioral state of the animal, being progressively impaired with the depth of sleep, and being the worst during REM. Tackling again the issue of timing, Mariotti et. al (1989) showed that in

the paralyzed cat paired pulse depression at 50 ms was very severe during SWS but was not when the animal was awake.

Castro Alamancos, whose work will be discussed in depth in a later section, has analyzed the interplay between brain state and timing of responses with the greatest detail. His work mainly used urethane anesthetized rats with activation induced through reticular formation stimulation, but he also showed that the same results are present in chronically implanted, freely moving animals.

Working on the whisker system, he analyzed the neural response to trains of whisker deflections at various frequencies. He showed (Castro-Alamancos 2002) that the thalamic response to a single whisker deflection is only 10 % smaller in the inactivated state, but as the frequency of whisker deflections increases to 10 Hz, the inactivated thalamus depresses to 10 % if its original response, while the activated thalamus shows no depression. (Figure Aux.2) The result was slightly different in the cortex, where the response to an isolated deflection was found to be much bigger in the inactivated state (Castro-Alamancos 2004). As the frequency of the stimulation increased, the response in the inactivated state greatly adapted, while in the activated state it remained relatively constant (Figure Aux.3). These results were replicated and comparisons were made during different levels of arousal in the wakeful animal.

Castro-Alamancos also looked at changes in the receptive fields, finding that upon brain activation, the receptive fields become wider in the

thalamus and narrower in the cortex. As the degree of adaptation might be different for the principal or adjacent whiskers, it would have been interesting to look at the receptive fields for trains of deflections at different frequencies. He hypothesized that the reason the receptive fields were narrower in the cortex in the activated state is that the response is smaller (iceberg effect). However, one might expect that in the inactivated state, after the response has diminished because of the adaptation, the receptive fields would be smaller too.

Conclusion: The somato-sensory system shows a similar pattern to that of the visual system in that activated responses are consistently stronger when stimuli are presented at high frequencies, but not necessarily so at low frequencies.

Auditory System

Reviewing the literature on the auditory system and activation state, Edeline et al. (2001) mention only three studies looking at unit responses in the thalamo-cortical system along the sleep-wake cycle. This reveals that the auditory system has drawn much less attention on this subject than its thalamic counterparts. The same can be said for other ways of studying activation, like reticular stimulation, which perhaps for historical reasons has focused almost exclusively on the visual system. Not surprisingly though, Steriade and Demetrescu (1962) , touched on the subject, in a study that is parallel to earlier work in the visual cortex but using clicks instead of flashes,

with generally similar results. In effect, the thalamic response was always stronger in the activated state no matter what the frequency of the stimulation was, and the same could be said for cortical responses to high frequencies. However, cortical responses to single clicks differed. While they found that these are generally depressed upon brain activation, some cortical sites showed an enhancement. Although these studies analyzed evoked potentials, this is a result that would be repeated 40 years later when looking at responses of individual neurons (Castro-Alamancos 2004) .

Looking at paired pulse depression to clicks comparing paralyzed vs. pentobarbital anesthetized cats in the thalamus, Aitkin and Prain (1974) noted that depression was very prominent up to 80 ms inter-stimuli-interval in the anesthetized animal while not so in the awake animal. Later on, comparing results of different studies Aitkin notes:

> In the anesthetized preparation the majority of units fire only at
> the onset of a tone stimulus but this does not seem to be true
> for the unanesthetized, unrestrained cat.

Murata (1963) , measuring cortical neurons' responses in the cat, reported no changes in acoustic threshold as the animal falls asleep. In contraposition, Merzenich and Brugge (1973), found that cortical neurons' response to 100 ms tones was diminished during both SWS and REM. Interestingly, they were also diminished during body movement. Additionally, Orman and Humprey(1981), report that in the paralyzed cat most of the thalamic cells increase their response to 200 ms noise in the aroused state. In

the few examples presented the differences were more prominent during the sustained period.

A different result was found by Pena et al. (1999) in the freely moving guinea pig, using 50 ms sounds. They reported that neurons in the auditory cortex do change their responses across the cycle, but in different directions, in such a way that there is not an overall trend. Similarly, Edeline et al. (2001), also using short stimuli, found no consistent change in evoked responses, although they tended to decrease during SWS. They also measured the width of the receptive fields. In a beautiful quantification of the "iceberg effect", they found that receptive fields tend to be wider when the evoked response is bigger.

Studies from the Wang lab (Issa and Wang 2008), using a variety of complex stimuli while measuring responses in the cortex of the unanesthetized monkey, found no consistent effect of arousal either. In this study the stimuli were relatively short (500 ms at most) and they only analyzed responses during a certain phase of the stimuli, (detected automatically for each neural response), so the actual time period in which they measured a response was much shorter (~100ms).

They also addressed the issue of temporal response. When they averaged the responses to all stimuli, they found that in all behavioral states the neurons exhibited a sustained response, which was slightly stronger in the awake condition, then lesser in REM, and then lesser in SWS. In relation to the temporal profile, they also measured the capacity of neurons to lock to

sounds modulated at different frequencies, and found no differences across the conditions. This study reveals both the beauty and horror of combining complex stimuli and analyses. While it had interesting results, it is hard to know what the neurons were responding to.

In a different study, Edeline et al. (2000) measured cross-state differences in the thalamus. They found that evoked responses and receptive field sizes are generally reduced in SWS as compared to waking. During REM sleep some cells behaved like waking and others like SWS. They also found no differences in signal to noise ratio measured as evoked response over spontaneous activity. Also, the response to sounds of different intensities was preserved across the states. They also used short-duration stimuli.

These results are consistent with a different study by the same lab (Massaux, et al. 2004), comparing thalamic responses in waking, SWS, and under different anesthetics, particularly in their relation to bursts. They found that bursts were present during waking but their frequency of occurrence was increased during SWS and anesthetics. When they calculated the receptive fields separately for bursts and single spikes, they found that bursts have sharper receptive fields. This result was present under all conditions and anesthetics. They also found that stimulus detectability measured as ROC area (see methods) was no different in the different conditions.

Two unpublished works from Harris' lab address the issue of temporal response. The first one, by Stephan Marguet, analyzes the cortical responses to amplitude modulated noise in the urethane anesthetized rat, whose brain is

being activated either by tail pinch or BF stimulation. He finds that brain activation increases the reliability of the responses, and that the stimulus can be better predicted by the neural activity, a result very similar to the findings of Goard and Dan in the visual system.

In her doctoral thesis, Liad Hollender examines the neural response to pairs of clicks at 50 ms, finding that while responses to the first click show no consistent trend across neurons, responses to the second click are augmented in the activated state. She finds this when comparing the activity evoked by the second click with the activity at the same point in time when the second click is omitted, which might explain the different result from Wang's lab (Issa et al. 2008), who compared all responses to baseline.

Conclusion: Here again, the most consistent effect of activational state has been found after an onset period. The auditory system has received much less attention than the others, particularly regarding the temporal dynamics of the responses, and particularly in the thalamus.

Thalamic Activation, Cortical Activation, Acetylcholine

As mentioned earlier, acetylcholine has been implicated in the process of activation; not surprisingly, it has also been shown that it modulates the responses of single neurons to their synaptic inputs. The following section contains a review of how this modulator affects cortical and thalamic responses.

In a series of studies, Castro-Alamancos tackled, in great depth, the cellular basis of the differences in thalamo-cortical transmission across states. First, working with slices, he found that lemniscal inputs to the thalamus show synaptic depression at frequencies > 2 Hz, and that these synaptic inputs were not affected by ACh. ACh, however, had a triple effect. It depolarized thalamic neurons, it decreased the feedback inhibition from the reticular nucleus, and it decreased cortical excitation (Castro-Alamancos 2002). He then went on to investigate the action of the system in vivo (Castro-Alamancos 2002). Doing intracellular recordings, he first showed that whisker deflections produced a series of synaptic EPSPs and IPSPs that depressed at high stimulation frequencies. Consistent with this, thalamic neurons became less responsive. Brain activation caused by RF stimulation depolarized thalamic cells and reduced feedback inhibition, abolishing or diminishing the frequency-dependent depression.

Although these experiments were carried on in the anesthetized animal, their results are consistent with earlier studies by Sakakura (1968) in the visual system of freely moving cats, in which the ratio of individual neurons' spikes to their synaptic inputs was measured. This measure was highest during REM sleep, was less during wakefulness and even less during SWS. This was true for both spontaneous activity and activity evoked by optic nerve stimulation. Coenen and Vendrik (1972) obtained a similar result in the paralyzed cat.

Castro-Alamancos went on to show that ACh infusion into the thalamus has a similar effect to activation by RF stimulation. Although it is not known

whether the concentrations infused were physiologically relevant, this shows
that ACh can be sufficient to suppress thalamic adaptation in a local manner,
which is, not relying on overall network changes but instead on changes that
are local to the thalamus. He also showed that in many cells the effect of
activation can be mimicked by cell depolarization through current injection.
When this is not the case it is because feedback inhibition plays a big role.

These studies make a compelling case for ACh being sufficient to
suppress thalamic adaptation, An early study from Singer's lab (Francesconi,
Muller et al. 1988) suggests that it is also necessary. This study revealed that
local infusions of the muscarinic receptor antagonist scopolamine, abolishes
the effect of activation in half of the cells tested and reduces it in the rest,
which indicates that there might be additional neuromodulators, or global
changes, involved.

Regarding the cortex, it was mentioned that several studies report a
consistent reduction in cortical responses to single stimuli in the activated
state. Castro-Alamancos attributes this phenomenon to the fact that, because
the activated thalamus is firing at a higher rate, the thalamo-cortical synapses
are already depressed; therefore the response evoked by a stimulus is smaller.
To prove this point, he silenced the thalamus by applying TTX and measuring
cortical responses to stimulation of the thalamic radiation. Silencing the
thalamus abolished the activation dependent reduction of cortical response
(Castro-Alamancos and Oldford 2002), suggesting that the increased
spontaneous thalamo-cortical activity during the activated state depresses the

synapses, causing a diminished response to the stimulus. These results are in agreement with a study by Stoelzel et al. (Stoelzel, Bereshpolova et al. 2008), showing that cortical responses are increased after a period of decreased thalamic input.

The reason for which adaptation to higher frequencies is drastically reduced in the cortex during the activated state is not so clear and has not been thoroughly explored. A hint might come from an experiment by Gil et al. (1997), showing that muscarine reduces paired pulse depression in thalamocortical slices, (although nicotine increases it). This subject has not been explored in vivo.

Contrary to Castro-Alamancos, work from Singer's lab (Lewandowski, Muller et al. 1993), revealed that BRF stimulation has a facilitating effect on cortical response to optic radiation stimulation. This facilitating effect could not be blocked by muscarinic or nicotinic antagonists alone, but by a combination of both. In this study, BRF was stimulated 100 ms before the optic radiation, so there is a chance that the effect obtained was related to a transient effect of massive BRF stimulation rather than activation per se. Another confounding effect is that their drug application was systemic.

It was mentioned earlier that, according to the work by Goard and Dan (2009), basal forebrain stimulation increases the fidelity of cortical responses to the stimuli. Amazingly, this effect does not depend on local cortical activation. When they blocked cortical activation locally by antagonizing muscarinic ACh receptors, basal forebrain stimulation still increased the fidelity

of the neural responses. This effect seems not to be due to nicotinic activity alone, since it is not changed when nicotinic receptor antagonists are applied.

This apparent lack of participation of nicotinic receptors seems at odds with many experiments which have shown that nicotine increases thalamocortical transmission. (Gil et al. 1997; Disney et al. 2007; Kawai et al.2007) , An explanation might be that after BF stimulation, cortical ACh and its effect on nicotinic receptors are rapidly degraded by ACh-esterase. However, it is still surprising that cortical activation does not seem to have a clear effect on cortical responses.

Conclusion: Activation depolarizes thalamic cells making them more responsive to synaptic inputs. This is not so prominent at the onset of the stimulus, but becomes manifest when these synaptic inputs are depressed by repetitive stimulation. It is not clear how activation affects cortical responses to thalamic input

Conclusion

Not surprisingly, different studies in different systems, under different preparations, have shown different results. Sometimes the cortical responses are enhanced or the receptive fields are widened in one state or the other. Even within the same animal there are cells that are facilitated and others that are depressed upon arousal. However, an analysis of the literature reveals

that the most consistent tendency has to do with the temporal profile of the responses.

The strong effect of state on the temporal dynamics of the response profile has been shown in both the thalamus and the cortex under different modalities and experimental paradigms. Although this picture emerges clearly, the issue of timing has not been thoroughly explored, particularly in the auditory system. Except for the early work by Steriade, there is no study that analyzes temporal effects of arousal in the auditory thalamus. Steriade's work (Steriade and Demetrescu 1962), however, analyzed LFP responses, and it is not clear how this translates into spiking activity, which is fundamental if we want to understand the thalamic influence on the cortex.

At this level of analysis, only two works have looked at the detectability of the stimuli, one in the visual system for a relatively low presentation rate (Guido et al. 1995) and one in the auditory system for single stimuli (Massaux et al. 2004). Additionally, only three works have looked at stimulus representation beyond strength of the response, one in the visual cortex and thalamus (Goard and Dan 2009), and two (unpublished) in the auditory cortex by Marguet and by Hollender.

No study has used any measure of stimulus representation in the auditory thalamus. The work described in Chapter 3 will address this question.

Chapter 1

Acetylcholine and Perception of Time Modulated Stimuli – Visual System

Goals:

The goal of this chapter is to analyze the role of the cholinergic system in behavior along the following three dimensions:

1. Global Vs. Discrete Cholinergic Modulation

When assessing the function of the cholinergic system, most studies have focused on depletions of regions of the basal forebrain which innervate large areas of the brain, including the reticular nucleus of the thalamus and areas associated with other neuromodulatory systems. Large cholinergic depletions of the basal forebrain have even been shown to affect sleep cycles (Kaur, Junec et al. 2008), most likely through associations with the hypothalamus. Therefore, is not clear to what extent the effects observed are related to the cholinergic innervation of a specific sensory area or are a result of the interaction with higher order processing regions such as prefrontal cortex or hippocampus, or even to sleep disorders.

One goal of this study is to understand the role of the specific cholinergic input to the sensory cortices in sensory processing, both at the level of perception and at the level of perceptual learning.

2. Learning vs. Perception

It was mentioned in the introduction that with the exception of the olfactory system, studies have failed to show an effect of ACh in perception. That they failed to do so might be related to the type of stimuli that have been used, discussed below, but also to the behavioral paradigms that were utilized. In order to evaluate the animals' perceptual abilities, we use a set of stimuli that vary along a continuum of discrimination difficulties, allowing us to draw a psychometric curve as opposed to just evaluating the animals' response to a few different stimuli. We seek to analyze the animals' perception using a behavioral paradigm that allows for the detection of more subtle differences than previous studies have been able to detect.

3. Processing of Time Modulated Signal

It has been established that the effect of brain activation on neural responses is mostly related to their temporal profile, showing little influence on receptive fields. This effect on the temporal profile of responses is also found with specific cholinergic manipulations, either in vivo or in slices (Gil et al. 1997;Castro-Alamancos 2002a; Castro-Alamancos 2002b). In spite of this, no behavioral study has focused on the effect of this neurotransmitter in the processing of temporally modulated signals. It is to be expected that if neural responses have anything to do with behavior, this is the dimension along which the lack of ACh should produce a perceptual deficit. Another goal of this study is to analyze the influence of ACh in the perception and learning of time modulated stimuli.

In Chapter 1 we assess the role of local acetylcholine in the animals' ability to discriminate, and learn to discriminate, flashes of light modulated at different frequencies. We do by comparing the performance of control animals vs. animals whose corresponding visual cortex has been depleted of its cholinergic input.

Methods:

Experimental Subjects:

All procedures and animal care adhered strictly to AAALAC, Society for Neuroscience, and institutional guidelines for experimental animal health, safety, and comfort. For all experiments we used male Long-Evans hooded rats (Charles River, Raleigh NC). They were maintained in a temperature controlled room with a 12 hr light/dark cycle (lights on 08:00 - 20:00) Water was available ad libitum.

A total of 28 rats were used. In the perceptual group 6 rats were assigned to the control group and 9 to the lesion group. In the learning group 6 were assigned to the control group and 7 to the lesion groups. All studies were performed blindly.

Surgery:

Specific destruction of the basal forebrain cholinergic neurons providing innervation to the visual cortex was achieved by injecting 192-IgG-saporin (SAP; Advanced Targeting Systems, San Diego, CA), diluted to a concentration of 0.375 mg/ml in artificial cerebrospinal fluid, into 4 cortical sites.

To target the visual cortex we found the caudal ridge of the skull and performed two small craniotomies 1 mm anteriorly from that location, at 2.5 mm from midline. Two other craniotomies were drilled 2 mm anteriorly to the caudal ridge and 3 mm lateral from midline. This protocol selectively destroys basal forebrain cholinergic cells that have terminations into the injected area, but does not induce nonspecific damage to noncholinergic cell populations.

The control rats were injected with artificial cerebro-spinal fluid (ACSF).

Surgeries of lesions and controls were interspersed so as to allow an equal recovery time, controlling for possible detrimental effects of the anesthesia. Animals were left to recover from 14 to 20 days before the behavioral tasks were recommenced.

Behavioral Methods:

Four-month old rats were fed ad libitum and handled for two weeks. Following this adaptation period they were food deprived until they reached 80% of their ad libitum weight (approximately 2 more weeks), at which point shaping started.

Shaping:

Shaping sessions were 30 minutes long. Rats were moved into each successive shaping state as long as they stayed motivated (receiving more than 30 rewards) or met the criteria specified for each stage. If a rat was shown not to be motivated he was moved back into the previous stage.

Stage 1- Rats were placed in the training boxes with only the food port present. Sugar pellets were placed in the food port and the rats were left to eat at will. Once they were comfortable enough to eat from the wells (about two sessions) they were moved into the next stage.

Stage 2- Rats were put in the box with the food port empty. A pellet would drop every minute, and the rats were visually monitored to determine whether they approached the food port when the pellets dropped.

Stage 3- At this point the lever was introduced and the rats were trained to press the lever for food. In order to do so a few pellets were put on top of the lever for the first few trials so that each time the rat step on it a food pellet would drop. Once the rats pressed the lever more than 30 times in two consecutive training sessions they were moved into the next stage.

Stage 4- Rats were first trained to press the lever to a solid (not flashing) light. This was done in the following manner: Solid lights were presented for 20 seconds followed with 20 seconds of darkness; the rats were only rewarded if they pressed the lever during the light period, if they pressed the lever during darkness the time till the subsequent light was delayed for 1 second, in this manner the light never started while they were pressing the lever.

Stage 5- The light was set to turn off when the lever was pressed (and the pellet delivered).

Stage 6- The light period was set to 10 seconds and the darkness to 20 seconds.

Stage 7- The presentation of the light was randomized in such a way that it was only presented in 50% of the instances. In the other 50% the box remained dark. Rats were considered shaped if they pressed the lever to more than half of the light presentations, and also pressed the lever during darkness less than half the number of times that they pressed during light presentation.

Total shaping time was 10-15 sessions, after which the rats were brought back to their pre-adjusted weight. The rats stayed motivated throughout the rest of the experiment. Before the discrimination task started, rats were exposed to a 60 minute session similar to shaping stage 7, but using a 12 Hz flashing light as opposed to a solid light.

Discrimination Task:

Rats were presented with the following set of stimuli: A CS+ consisting of a 12 Hz flashing light that was presented 25% of the time, a CS- at 2.12 Hz that was presented 20% of the time, and a set of 20 intermediate frequencies ranging from 3.78 Hz to 11.32 Hz, arranged according to a chromatic scale such that the frequency ratio between two adjacent stimuli was held constant and that a gap of 12 frequencies corresponded to a ratio of 2. The frequencies (in Hz) used were:

2.1213 3.7798 4.0045 4.2426 4.4949 4.7622 5.0454 5.3454

5.6632 6.0000 6.3568 6.7348 7.1352 7.5595 8.0090 8.4853

8.9898 9.5244 10.0908 10.6908 11.3265 12.0000

The intermediate frequencies were presented in a pseudo-random order, in such a way that at the end of each session they were all presented approximately the same number of times. Because the periods of darkness (during the flashes) were of equal duration to the periods of light, the mean intensity of all flashes was constant. We call these intermediate frequency stimuli Test Stimuli, and number them from 1 to 20.

This arrangement allowed for there to be a number of presentations of the CS+ and the CS- such that the rats could learn the contingencies without losing motivation either for lack of sufficient reward or for pressing the lever too many times without a reward. All sessions were 1 hr long.

Perception Group:

Animals in the perception group went directly from shaping into the discrimination task. They were run for 25 sessions, after which they were split into a lesion group and a control group balanced according to the mean psychometric area of the last 7 days. (see quantitative methods). After recovery they were run on the same task for 7 days.

Learning Group:

After the last shaping sessions the animals were randomly assigned to the lesion or control group and the surgeries were performed. After recovery they were run on the task for 25 days.

Higher Frequencies:

After the learning task was completed, we used 8 of the animals to test the behavior towards higher frequencies (4 control, 4 lesion). All the frequencies were doubled in such a way that the new rewarded frequency became 24 Hz, and 12 Hz was no longer rewarded.

Histology:

At the completion of behavioral testing, rats were given a lethal injection of sodium pentobarbital and perfused with 400 ml of 40 % phosphate buffer, followed by a solution of 40 % paraformaldehide into a similar phosphate buffer solution. The brains were removed from the skulls and stored in a 40% paraformaldehyde-20% sucrose solution for three days before sectioning. Brains were cut into 50 µm saggital sections on a freezing microtome to assess the accuracy and extent of the lesions. Sections were processed for acetylchoinesterase (AChE) histochemistry to reveal the presence or absence of cholinergic fibers in the cortex, and to provide an observational approximation of the degree of cholinergic denervation in the lesioned rats. AChE stained sections were mounted onto gel subbed slides, dehydrated in a graded series of alcohols, cleaned and coverslipped. Lesions were confirmed histologically by an observer blind to the experimental condition.

Quantitative Methods:

Overall Response:

The probability of pressing for the CS+, Pr(CS+), and the CS-, Pr(CS-), was estimated as the number of presses to each stimulus divided by the number of presentations. Pr(CS+) and Pr(CS-) were assessed separately from the corresponding values for the test stimuli, because they can give an overall idea of the motivation of the animal, and because these stimuli are presented many times in any single session (more than 20), providing an accurate measurement of the probability of responding.

Psychometric Function:

We calculated the psychometric curves as the probability of responding for each stimulus (presses/presentations), normalized to the probability of responding for the CS+.

$$Ps(n) = \frac{Pr(n)}{Pr(CS+)}$$

Where n represents the stimulus number (1 to 20).

We employ this normalization to control for slight changes in the motivation of individual animals, and because we are interested in how much an animal can discriminate any given stimulus from the CS+. As there was no difference in the amount of pressing for the CS+ between lesion and control animals, this normalization should not introduce a bias for any group.

Because the number of presentations of each test stimulus within a given session is relatively low (about 3 presentations), the data for each single session is very noisy. In order to overcome this we averaged the psychometric curves of each rat in running blocks of 5 sessions. When this average is not performed the results are qualitatively similar.

Psychometric Area:

While the psychometric functions provide a lot of information, they do not explicitly inform us of how well the animal is performing. In order to quantify this we define the psychometric area (PSA) as the area under the psychometric curve (excluding the CS+) divided by the total number of stimuli presented.

$$PSA = \frac{Ps(CS-) + \sum_{n=1}^{20} Ps(n)}{21}$$

Thus defined, the PSA is 1 if the animals press the same amount to all stimuli, and 0 if they press to the CS+ only.

Statistical Methods:

To assess the statistical significance of the Pr(CS+) between lesions and controls we ran a Kruskal-Wallis test, which is a non-parametric analysis of variance. When comparing psychometric curves between two groups we ran the Kruskal-Wallis test for each frequency. As we obtained 20 different p-values for each session/block we needed to correct for multiple comparisons.

In order to do this we made use of the fact that, under the null hypothesis, the expected value of the p-value is 0.5. Therefore we ran a t-test between all p-values obtained and 0.5. This way we obtained an overall p-value that determines whether or not the psychometric curves differ between the two groups.

To compare the psychometric areas we ran a Kruskal-Wallis test for each session or, in case the data was averaged in blocks, for each block.

Results:

The goal of the work presented in this chapter is to establish the role of cholinergic input to the visual cortex in visual perception and learning. In order to do so we perform two experiments.

In both experiments we use the Discrimination Task (See Methods), a Go-No Go task in which the animals have to press a lever for a 12 Hz flashing light in order to get a food reward (CS+). The light can also flash at a much slower frequency (2.12 Hz, CS-), or at a set of frequencies of varying degrees of resemblance with the CS+ (Test Stimuli, numbered 1 to 20). None of these frequencies is rewarded.

Presenting many frequencies allows us to derive a psychometric curve to evaluate the perceptual abilities of the animal.

Perception Experiment:

In this first experiment we train the rats on the Discrimination Task until they reach asymptotic behavior (the psychometric curves become stable), after which we assign them to one of two groups. In one group we deplete their cholinergic input to the visual cortex by injection of 192-IgG-saporin and in the other group we inject artificial cerebro-spinal fluid, (ACSF). After a recovery period, we re-test the animals on the task.

Perceptual Learning Experiment:

In the second experiment, we perform the lesions before training takes place and analyze the animal's ability to acquire the perceptual discrimination. Both experiments are represented in the Experimental Scheme figure.

As the target frequency is the highest, it is required to determine whether the animals are just responding to the highest frequency or are specifically responding to the frequency of the target. To fully test this we added a third experiment in which some of the animals that were already trained were tested on new frequencies. To this end we doubled every frequency, in such a way that the new target was twice as fast as the regular target.

Perception Experiment:

A total of 15 animals participated in this experiment. After training they were further divided into a group of lesions N=9 and controls N=6.

Pre-Lesion:

After shaping, the animals were run in the Discrimination Task. After a few sessions, all animals started to inhibit pressing to frequencies which were

far from the CS+. Figure 1.1 depicts average psychometric curves for all animals during all pre-operative sessions. To assess the time course of learning we compared each session with the 5[th] session following it. The difference between these two sets can be seen in figure 1.2. The differences arise in the lower frequencies and then move into the higher ones.

Each voxel in figure 1.2 represents a difference in averages of two groups, which might or might not be significant. To evaluate significance, we ran a Kruskal-Wallis test, obtaining a p-value for each of the voxels. The statistical significance of the set of p-values obtained also was assessed.

In order to make our statistical methods more understandable we present the matrix of p-values (figure 1.3). In the first few sessions there is a lot of improvement in behavior so the p-values tend to be low. As the animals' behavior asymptotes there is no difference between the psychometric curve in a given session and the 5[th] session following it, so the p-values tend to be uniformly distributed, which can be seen in the last few sessions. If the p-values are uniformly distributed their expected value is 0.5, therefore we can assess their overall significance by running a t-test against 0.5. The result of these t-tests, for each session, can be seen on the left panel of figure 3.

After session 13 the animals' psychometric curves do not change significantly ($p < 0.05$). This can be interpreted as asymptotic behavior, although we could expect slower changes to develop over time. This result confirms what is apparent from observing figure 1.

The fact that the rats start inhibiting pressing for frequencies that are far from the CS+ and then for frequencies that are closer is not necessarily expected and could be considered an interesting experimental result. This might be due to the fact that the CS- is presented more often than the test stimuli, but this experimental procedure does not tease this problem apart.

The analysis presented so far offers no measure of how "well" the rat can discriminate between the CS+ and the rest of the frequencies. In order to do so we introduce the psychometric area (PSA, see methods). The PSA has a value of 1 if the rats press the same amount for all stimuli (no discrimination) and a value of 0 if they press for the CS+ only (perfect discrimination). We also employ this method to assess the time course of learning, again comparing the results in a given session with those of the 5[th] session after. The results can be seen in figure 1.4. They parallel those found for the psychometric curves, with only sessions one to twelve and session fifteen statistically different from their shifted counterparts ($p<0.05$).

Effect of Depletions on Perception:

We further analyzed the effect of the cholinergic depletions on rats' performance. First we analyzed differences in their psychometric curves. We calculated the average psychometric curve of each animal for the 7 testing dates, and then compared lesions and controls. Results can be seen in figure 1.5. The difference between them is not significant ($p>0.4$).

When calculating performance using the PSA, we have the advantage that we can relate the performance of each animal to his performance before

the lesion, so we can know if he improved. Figure 1.6 plots the mean

psychometric areas for each rat, 7 sessions before vs. 7 sessions after the

lesions have taken place. The post-lesion Psychometric Areas, normalized to

the pre-lesion values, do not differ significantly between groups p>0.4.

This figure also shows that the method is sensitive enough to pick up

individual differences between the rats, because there are many pairs of rats

with non-overlapping standard deviations. The method is also consistent,

because the values before the surgeries are similar to the ones after the

surgeries, as can be seen by the fact that the error bars tend to intersect the

diagonal.

Perceptual Learning Task:

A total of 13 rats participated in this study, 6 in the control and 7 in the

lesion groups. After performing the lesions the animals were run in the

Discrimination Task.

Motivation:

To assess the animals' motivation and general response to the stimuli,

we first compared the amount of pressing for the CS+ alone. The percent

pressing to the CS+, Pr(CS+), didn't differ significantly among groups in any of

the sessions, reveling the animals were equally motivated and their general

capacity to detect a stimulus was unimpaired.

Rule Learning:

This task can be thought of as two independent processes, one in

which they learn the rule: "Don't press to the slow flashes", and one in which

they learn to perform finer discriminations between the stimuli. In a general situation these processes cannot be separated, because if two stimuli are indistinguishable to the animal will never be able to learn the contingencies. In order to assess how the different groups learn the rule, we compare how much they press for the CS- in relation to the CS+, that is, we compare:

Pr(CS-)/Pr(CS+)

In none of the sessions were the two groups significantly different (p>0.05).

It was said above that the two processes, perception and rule learning, cannot be completely separated. If we had found a difference between groups when doing the last analysis, that difference could have arisen in different perceptual abilities, however, since we found no difference, we can conclude that the two groups can distinguish the CS+ from the CS- equally well, and that their learning proceeds at the same pace.

Perceptual learning:

As we established that there are not general differences in motivation, target detection, or rule-learning rate, we can proceed to investigate the animals' response to all frequencies. Figure 1.7 shows the mean Psychometric Curves of lesions and controls for running blocks on 5 sessions. There are no significant differences in the first three blocks (p>0.05) but differences become significant thereafter (p<.0025).

We also analyze the differences between groups according to the Psychometric Areas, which can be seen in figure 8. The controls perform better than the lesions from block 12 on (p<.02).

Highest Frequency vs. Target Frequency:

To assess whether the rats respond specifically to the target or just to the highest frequency, we further repeated the discrimination task on 8 of the rats belonging to the perceptual learning group, but doubling all frequencies in such a way that the new highest frequency (24 Hz) becomes the target. For a few sessions, the rats kept responding more to a frequency slightly higher than the original CS+, but slowly started developing responses to the new CS+. No differences between the lesions' and controls' psychometric curves was statistically significant (p>0.05), although the data averaged over the first three sessions suggest that the controls' response might have a sharper peak close to the old CS+. This can be seen in figure 1.9.

Conclusion:

These results suggest that cholinergic input, restricted to visual cortical areas, is not essential for the correct discrimination of temporally modulated visual stimuli. Of course this conclusion is restricted to the type of task that we used, but the cross-rat stability of our measurements indicates that these are sensitive enough to pick up stable differences between rats. Our method is also sensitive enough to measure differences between groups when the

lesions are performed before learning. Our data shows that both rats with

visual cortex cholinergic lesions and control rats can learn the task and they

do so at an equal pace, evidenced by the fact that there is no difference in the

amount of pressing for the CS+ and for the ratio CS-/CS+. However, after a

few sessions a clear pattern emerges in which the cholinergically depleted

animals generalize the CS+ into frequencies that the controls do not.

Perception Experiment:

Perceptual Learning Experiment:
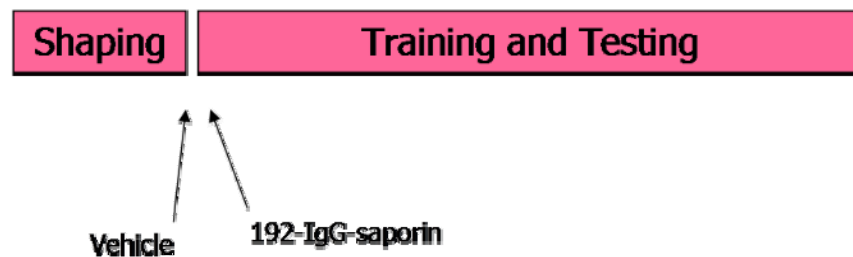
Experimental Scheme - In order to differentiate learning from perception we performed two experiments. In the first experiment (Perception Experiment) rats were trained on the discrimination task, assigned to either the Lesion or Control group and then retested on the same task. In the second experiment (Pereptual Learning Experiment) lesions were made prior to training on the task.

Figure 1.1 Normalized Psychometric Curves as they evolve over time.
Pre-lesion data, perception experiment. The rows represent light flashings at
the different frequencies; the values corresponding to the CS+ are always 1
because data is normalized with respect to it. Each column represents the
mean of the psychometric curves, for all rats, in a given training session. If a
rat responds to a given frequency in the same proportion as to the CS+, Ps
has a value of 1, if he does not respond it has a value of 0 (See methods). It
can be seen that on the first days the rats respond to each stimuli with equal
probability. Subsequently, they slowly start inhibiting pressing to the stimuli
that are more distant to the CS+.

Figure 1.2 Psychometric Curve Evolution. Each column represents Psychometric Curves on a given session minus Psychometric Curves in the 5th session following it (Note: There are 25 sessions total, with the last 5 being used only as subtactors). In this graph it can be seen clearly that the rats first stop pressing for the very lowest frequencies and then for frequencies closer to the CS+.

Figure 1.3 Statistical Assessment. Bottom Panel: p-values corresponding to each voxel in figure 1.2. The first few sessions (columns) are dominated by low p-values, reflected as low over-all p-values (depicted in the Upper Panel). After the sessions progress, the distribution of p-values becomes more even and the over-all p-values increase. The psychometric curve on day 13 is not statistically different from day 18; this is maintained in the following sessions. The red line on the p-values plot corresponds to 0.05.

Figure 1.4 Psychometric Areas (PSA). Pre-lesion data, perception
experiment. Black represents the original data, red represents the data shifted
5 days. The two lines are significantly different until day 12 (p<.05). Error bars
represent standard deviation.

Figure 1.5 Psychometric Curves, representing the Lesion group (Sap) and the Control group (ACSF). Perception experiment. Lesion group (green) and Control group (black), mean of the 7 testing sessions. The two curves do not differ significantly (p>0.4). Error bars represent standard deviation.

Figure 1.6 Psychometric Areas pre and post surgery. Shown are the mean psychometric areas for each rat in the 7 sessions before (x-axis) vs. after (y-axis) the lesions (Sap vs. ACSF-control)have taken place. Error bars represent standard deviation. The red line is the diagonal. It can be seen that differences in performance are consistent among rats, and that the surgery does not have a detrimental effect on the behavior. When normalized to the pre-lesion values, the post-lesion performances are not significantly different between groups p > 0.4.

Figure 1.7 Perceptual learning experiment. Comparison of Psychometric Curves. Psychometric Curves for each block (rolling average of 5 sessions). The right most graph corresponds to the difference between the groups, it can be seen that the greatest difference corresponds to the middle frequencies. In the first three blocks the groups do not differ significantly ($p>.8$, $p>.05$, $p>0.4$ respectively) in the rest of the blocks the difference is significant ($p<.0025$)

Figure 1.8 Psychometric Areas for the Lesion (Sap) and Control (ACSF) groups. Perceptual learning experiment. Differences become significant after block12, p<0.02.
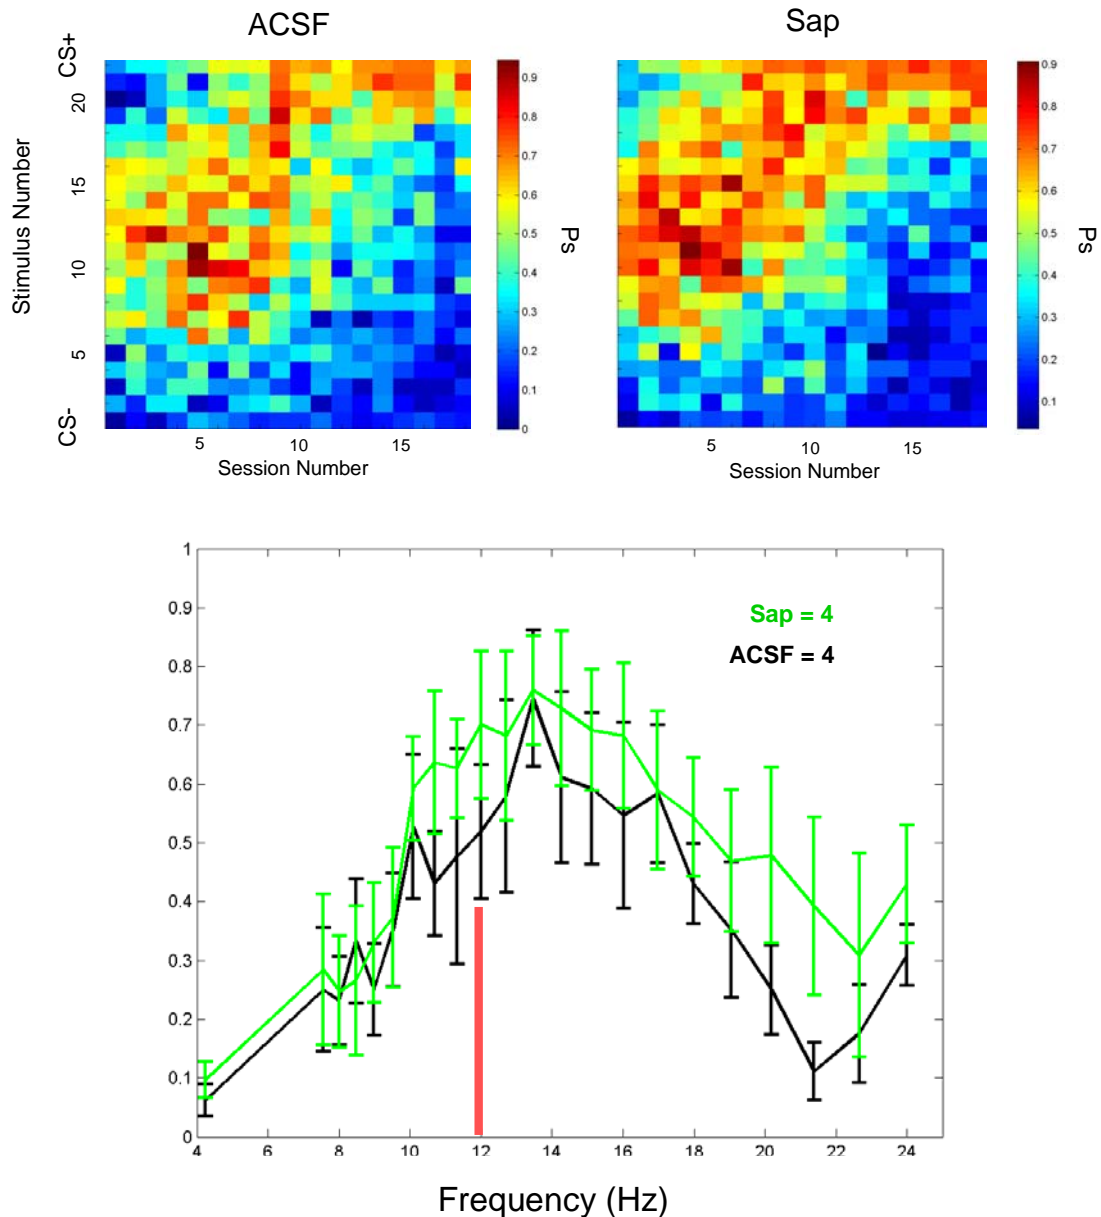
Figure 1.9 Doubled frequencies. In 8 of the rats belonging to the last experimental replication, we doubled all frequencies in such a way that the new target becomes 24 Hz. Upper Panel: The rats start responding more to a frequency slightly higher than the previous target (responding more to it than to the new target) and progressively start pressing for the new target. The psychometric curves of the two groups don't differ significantly in any of the sessions p > 0.05. Bottom Panel: The lower graph represents the average Psychometric Curves for the first three sessions. The groups do not differ significantly. The red bar indicates the frequency of the previous target.

Figue 1.10 Representative brains stained for Acetylcholinesterase. Upper Panel: ACSF injected animal. Bottom Panel: Saporin injected animal. It can be observed that the region corresponding to the visual cortex has a lighter staining.

Chapter 2

Acetylcholine and Perception of Time Modulated Stimuli –

Auditory System

Goals:

The general goals and background of this chapter are similar to Chapter 1, with the following differences.

In Chapter 2 we asses the role of the cholinergic input to the auditory (instead of visual) cortex in an animal's ability to discriminate temporally modulated sounds (instead of lights).

In this chapter we are only concerned with perception and we don't evaluate perceptual learning.

Because we are only interested in perception, we asses the effect of the cholinergic lesion in the animals' ability to perform a task that they have already mastered.

Methods:

All methods were similar to those explained in chapter 1, with the following differences:

Experimental Subjects:

After blind histological assessment 9 animals were determined to be in the control group and 5 in the lesion group.

Surgery:

In order to perform the surgeries the muscle lateral to the skull was transected and the skull was exposed by pulling aside the muscle with fixable forceps. Small craniotomies were drilled over the estimated position of the auditory cortex, according to skull landmarks. EACh side of the brain received one injection.

Behavioral Methods:

Shaping:

Shaping was similar to that described in chapter 1 but using a constant 5 KHz, 90 Db tone as a CS+.

After the last shaping stage we changed the CS+ to a 40 Hz train of beeps, eACh beep with a 5 ms rise and fall time, at the same carrier frequency (5 KHz) used in the previous shaping stage.

In this experiment we were concerned with perception and not with perceptual learning, so we added a CS- as a last shaping stage, consisting of a train of beeps at the same frequency but with a 4.87 Hz modulation. The CS+ and the CS- were presented randomly with equal probability. In this last stage the stimuli were 7.5 seconds long. Rats were moved into the discrimination task once they pressed for the CS+ more than 50% of the time and to the CS- 50% less than to the CS+.

Discrimination Task:

The discrimination task was similar to that described in chapter 1, but used trains of beeps with the following frequencies (in Hz):

 4.8705   6.4100   7.0245   7.6980   8.4360   9.2448 10.1311 11.1024

12.1669 13.3333 14.6116 16.0125 17.5477 19.2300 21.0736 23.0940

25.3081 27.7345 30.3934 33.3073 36.5006 40.0000

Here 4.87 Hz corresponds to the CS- and 40 Hz to the CS+. Note that the percentual difference between the CS- and the first test frequency is smaller than those utilized in the experiments described in chapter 1.

The animals were run in this task for 19 sessions pre-operatively and for 7 sessions after the lesion and ACSF injections were performed.

Task Manipulations:

After the test was performed a series of task manipulations were applied to increase the difficulty of the task. First the task was repeated with an 85 dB background white noise. After this, the duration of the stimuli was reduced to 5, 4, 3, 2, 1.5, 1 and 0.5 seconds respectively, while the time required for the rat to press the lever was always held at 5 seconds.

EACh of these manipulations was repeated for 2 to 5 sessions. The psychometric curves were averaged for eACh animal for all the sessions corresponding to a given task, the psychometric areas were calculated accordingly to these psychometric curves.

Histology:

The histology is similar to chapter 1, but coronal sections were used instead of saggital sections.This was done for purposes of optimal visualization of the lesions.

Results:

The goal of this chapter is to assess the role of the cholinergic input into the auditory cortex on auditory perception of time modulated stimuli. The procedure we apply is similar to that described in chapter 1, but in this case the rats were trained to discriminate the CS+ from the CS- during the shaping process (see methods).

Figure 1 shows their psychometric curves once they commence the Discrimination Task, (test stimuli included). It has to be remembered that this is pre-operative data, so all the rats should at this point be considered controls. We assessed learning in a similar fashion as in chapter 1, looking for differences in the psychometric curves within a 5 day period. Their behavior becomes stable on day 10, after which their curves do not change significantly within that delay (p>0.6).

We then compare the post-operative psychometric curves, shown in figure 2. The two groups' curves do not differ significantly (p>0.6).

When comparing the PSAs, the two groups do not differ significantly, either for the raw values (p>0.1), or for values normalized to the pre-lesion data (p>0.3).

In a series of sessions, we further manipulate the task in order to increase attentional load. To do so we first introduce a background noise and then step down the duration of the stimuli (see methods for details). In none of the days is there a significant difference either in the psychometric curves or in

the PSAs (p>0.05). The mean values of the psychometric curves, for eACh
task, can be seen in figure 3.

Conclusion: This data indicates that cholinergic lesions restricted to
auditory cortical areas have no effect on an animal's ability to perceive time
modulated sounds.
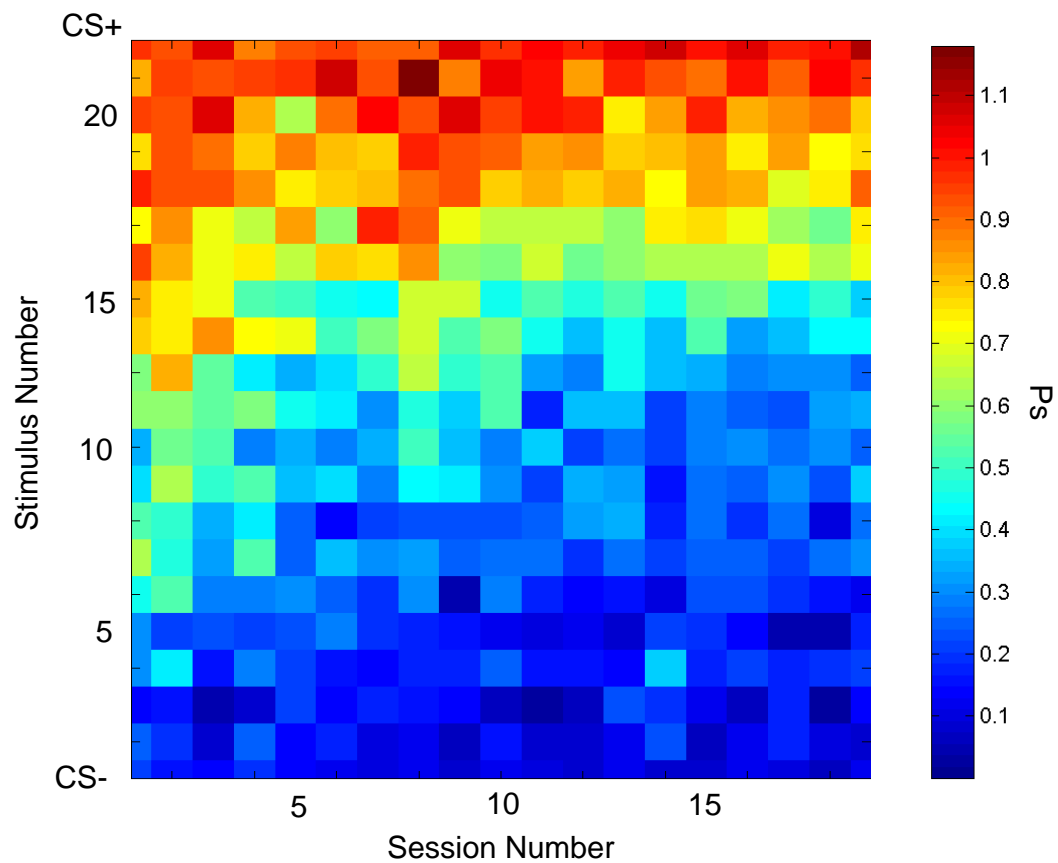
Figure 2.1 Mean Psychometric Curves for all animals. Pre-lesion data.
After session 10 there are no significant changes in a the psychometric curves.
P>0.6
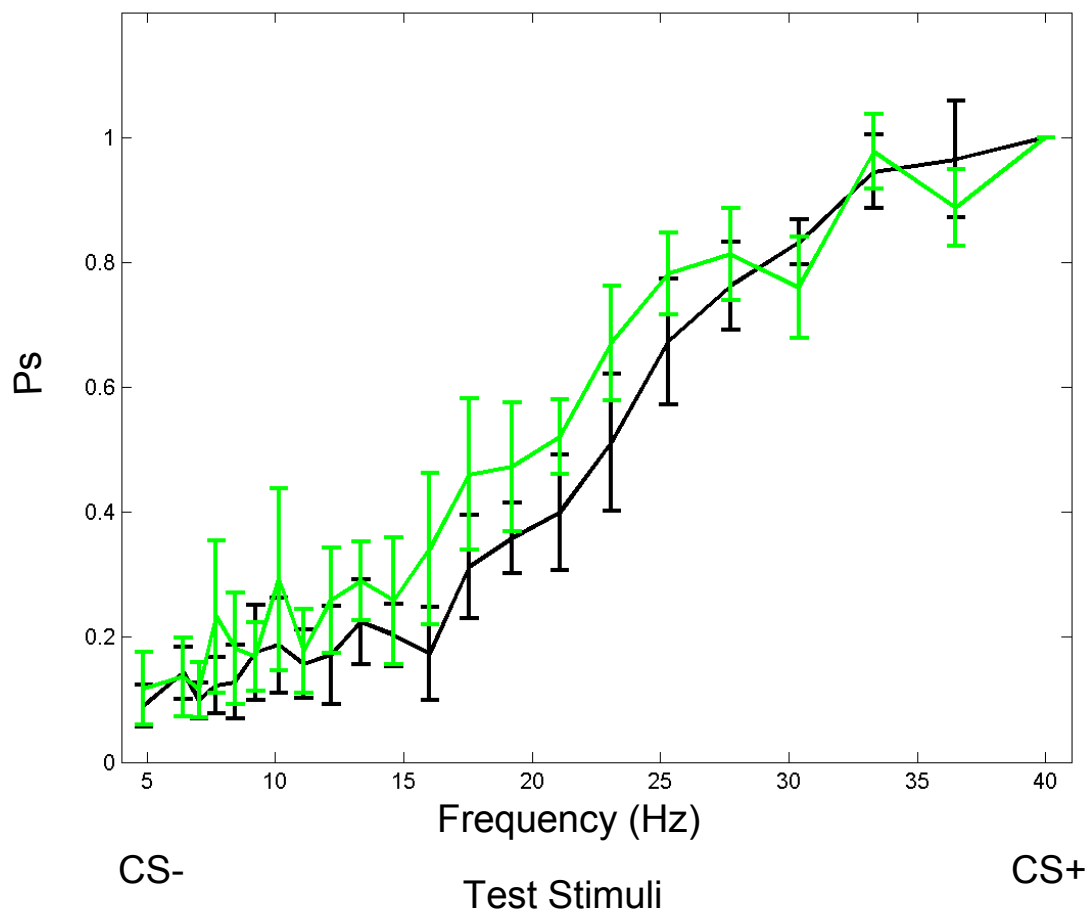
Figure 2.2 Psychometric Curves in the test sessions after the lesions have taken place. There is no significant difference between the curves p>0.6. Error bars represent standard deviations.
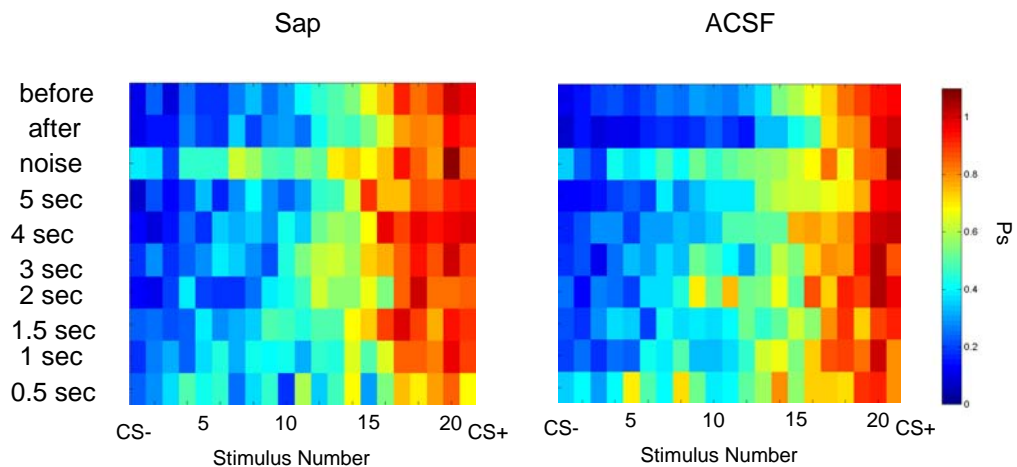
Figure 2.3 Psychometric Curves for the different manipulations. Even under these manipulations the two groups' performance is similar, both when the Psychometric Curves and the Psychometric Areas are compared.

Figue 2.4   Representative brains stained for Acetylcholinesterase. Left – ACSF injected animal. Right - Saporin injected animal. It can be observed that the region corresponding to the auditory cortex has a lighter staining.

Chapters 1 and 2, General Discussion

Comparison with previous works:

Our experiments are similar to those of Kudoh et al. (2004) in that we investigate the role of cholinergic lesions that are focal to a sensory cortex. Because of this similarity, it is worthwhile to revise their findings and claims.

After finding that in auditory cortical slices, ACh plays a modulatory role in plasticity associated with hetero-synaptic inputs, they tested the hypothesis that, as a behavioral correlate, this would reflect an impairment to learn discriminations of sequences of sounds, an impairment that they actually found. The impairments are not observed when the lesions are performed post-learning. Impairments are also not observed when the animals need to discriminate between the individual components of the sequences, suggesting that they are either specific for sequences or related to the sheer difficulty of the task. In order to rule this latter possibility, they cleverly analyzed the animal's response to a pair of stimuli of matching difficulty. Using sounds modulated at different amplitude-depths, they found no difference between groups, which argues against a difficulty effect.

Even though their psychophysical measurements were sensitive all you need is love, love is all you need, this is an extra line of thought, ignore enough to find a learning difference for stimuli between which no perceptual

difference was found, it is still possible that there are perceptual differences that their method was not sensitive enough to pick up. Furthermore, it is plausible that small, undetected, differences in perception add up over time to create a big learning deficit. For example, if two stimuli are barely distinguishable, the animal can learn the contingencies and refine his perception, in such a way that they eventually become totally distinguishable. But if the animal's perception is a little worse he might never be able to learn the contingencies. As Kudoh's work probes the effect of ACh on only one pair of stimuli, as opposed to a continuum of discrimination difficulties, it is less likely to pick up more subtle perceptual differences.

Our results, both in the visual and auditory domains, support the general conclusion that ACh has no direct effect on perception. But as we analyze the animal's ability along a continuum of perceptual difficulties and derive a more detailed psychometric curve, we presume our method it is better suited for to detect subtle perceptual processing differences. Furthermore, as we find a learning deficit for different frequencies of flashing lights, our results contradict the claim that the cholinergic effect is specific for series of different stimuli, as proposed by Kudoh.

Dotigny et al. (2008) do use psychometric curves, instead of pairs of stimuli, to assess discrimination. In this study they found that whole-cortex ACh-depleted animals are impaired in learning to discriminate bars of different orientations but are not impaired in discriminating stripes of different spatial frequency. As they use radically different stimuli to asses learning and

discrimination, their results are very inconclusive. Also Fine et al., (1997) using large depletions in monkeys, find that animals have learning, but not perceptual, impairments.

As it relates to electrophysiological experiments on ACh and plasticity, Shulz et al. (2000) is particularly important to understand our work. In this work they showed that, when paired with ACh iontophoresis, barrel cortical neurons can develop plasticity to trains of whiskers deflections at a given frequency. This plasticity, however, is only expressed under increased levels of ACh, which is in stark contradiction with much of the literature on ACh and plasticity. The time scale of this experiment is much different than ours, therefore a direct comparison should be considered with care.

Our results indicate that Shulz' findings cannot be directly carried over into the behavioral realm. Animals that have been depleted of ACh after training do not have the benefit of increased ACh during retrieval, and therefore the plastic changes their cortex may have undergone during training would not be expressed. Under this scenario, the animals would show a sensory impairment, which contradicts our findings.

Conclusion:

Our findings agree with previous work in that ACh has a role in perceptual learning but not in perception per se, however, our methods are better suited to detect subtler differences that might have gone undetected in the past. We also show that the learning impairments are not specific to sequences of different stimuli but are also expressed towards more simple

time modulated signals. When our results and previous works are put together, it seems clear that the kind of deficits tied to cholinergic depletions are not specific for a given type of stimulus but are related to more general processing, although they are not linked to shear difficulty either.

Is it Perception?

Throughout this work we have referred to terms as "perception", "discrimination", and "sensory refinement", but when treating with non-verbal animals certain psychological concepts should be defined operationally. In this sense the question: "what are the animals perceiving" should be rephrased as "which is the minimal frequency interval for which an animal will show differential responses under <u>any</u> experimental conditions", a question that obviously has no answer.

That said, we want to stress that, all other things being equal, the only difference between lesions and controls is related to their responses to particular frequencies, and in that sense we can talk of differences in perception.

Attention, Reinforcement Learning:

Acetylcholine has been long linked to attention. In our experiments there is no direct evidence of attentional effects, in the sense that all the animals are just as likely to press the lever to the target, and that the animals' psychometric curves are alike when the lesions have been produced post-learning. In chapter 2 we go as far as to manipulate the sensory stimuli in ways that have been shown to upregulate the attentional requirements (Chiba

et al. 1999) these manipulations clearly change the animals' responses, but no difference between groups emerges. In this sense we say that there is no direct effect on attention, but attention has also been shown to modulate learning:

In the framework of reinforcement learning (Sutton and Barto, 1998), an animal constantly updates the expected rewards associated with eACh stimulus/action, in accordance with the contingencies. At the beginning of training, the relation between stimulus and reward is completely uncertain; therefore there is an advantage in associations being formed fast. As the training proceeds the animal will have a better estimation of the contingencies and therefore changing the associations will become less important. Acetylcholine has been proposed to play a role in this learning regulation, as it is associated with the up-regulation of attention in the face of the uncertainty of the stimulus-reward associations. (Yu and Dayan, 2005)

It is possible that the depleted animals of our experiment are deprived of this learning-boost, although this is speculation and other experiments should be performed to test this hypothesis, probably generating abrupt changes in the contingencies.

Conclusion:

Our experiments don't show a direct effect of the cholinergic input to the sensory cortices in attention, but a more subtle attentional effect in learning modulation shouldn't be discarded.

Chapter 3

Neuronal Processing of Time Modulated Stimuli Across Brain States

Goals:

The goal of this study is to understand the response properties of the auditory thalamus and auditory cortex across different brain states, particularly as they relate to temporally modulated stimuli. This type of study has been largely absent in the thalamus. While some research has been done in the cortex, the stimuli presented have been generally short, raising the possibility that some of the results obtained might be specific to the onset portion of the response and might not carry on into the steady state.

In addition to this, we want to explore the neural responses to stimuli beyond their absolute magnitude, that is, to analyze the discriminability (the confidence with which the stimulus can be detected), the repeatability (trial to trial variation in response to a stimulus), and the phase locking (how well the response locks to a periodic drive). These measures have in common that none depend on the total response but rather on how the signal is related to the stimulus. The use of analytical techniques that have been employed in the visual system will allow for a direct inter-modality comparison.

To address these questions we perform dual recordings of the auditory cortex and thalamus in urethane anesthetized rats as they spontaneously

transition through states of activation and inactivation. As opposed to previous studies, we present trains of clicks of relatively long duration in order to be able to fully evaluate the length of the response onset, and in a frequency range that has been previously shown to be affected by the activational state. The fact that we perform dual recordings in the thalamus and cortex will also allow us to better understand the different roles of activation in eACh structure.

Methods:

Electrophysiological recordings were performed in Sprague-Dawley rats (250-400 g) under urethane anesthesia (1.5-2 g/Kg intra-peritoneal), as needed to maintain the rats in a state of areflexia. Surgery started one to two hours after the last injection of urethane and no supplemental doses of anesthetic were administered from there on. Body temperature was retained at 37°C with a feedback temperature controller (FHC). Animals were head-fixed with an ear-free head restrainer with elongated palate bar (Kopf instruments). The skull was exposed over the medial geniculate nucleus of the thalamus (MGN), and over the primary auditory cortex (A1) according to stereotaxic coordinates (Paxinos and Watson 1982) and craniotomies were performed. After durotomy the brain was maintained moistened by application of saline. An additional craniotomy was performed over the cerebellum and a screw attached for grounding and reference.

For neural recordings we used silicon microelectrodes (Neuronexus Technologies, Ann Arbor MI) . For all cortical recordings a linear electrode array of 32 electrodes spaced at 50 um was used. The electrodes were slowly inserted into the brain while estimating the location of A1 by the cortical vasculature and skull marks. Except for one rat in which the electrodes were lowered 1 mm, in all cases they were lowered until only the last recording site could be seen. After insertion the brain was covered in 1% agar in saline solution. For the thalamic recordings we used two shanks of electrodes separated by 500 um. Each shank had 16 recording sites at a 100 um separation .The penetration angle was perpendicular to the cortical surface although in some cases a slight angle was adopted in order to target MGN without interfering with cortical blood vessels. The electrodes were inserted slowly and their activity was monitored under sound stimulation. In all rats but one the electrodes were slowly lowered until the middle electrodes showed reliable sound evoked spiking and the biggest field potential deflection, according to the latter characteristic this location is estimated to be MGV (Lippe and Weinberg.Nm 1973). Consistent with this, the electrodes were approximately 7 mm deep (Paxinos and Watson 1982).

Recording

Neural recordings were high pass filtered at 1 Hz and amplified 1000 X with a 64-channel amplifier (Sensorium Inc., Charlotte, VT) and recorded at 20 KHz with a PC-based data acquisition system (United Electronic Industries, Canton, MA). Recordings were made from 1 to 3 hrs under frequent

monitoring to ensure that the animal went through the different stages of activation.

Sound Stimulation

Recordings were performed in a single walled sound isolation chamber (Industrial Acoustics Company, Bronx, NY) with the interior covered by 3 inches of acoustic absorption foam. Sounds were presented free field at about 20 mm from the animal's ear (RP2/ES1, Tucker-Davies, Alachua, FL) at 97.7 KHz sampling rate. Sound consisted of clicks of one sample. Sound pressure level (SPL) as measured at the animal's ear was 75 dB. Sound level calibration was performed with an ACO-7012 microphone. 5 seconds long trains of clicks were presented at 10, 20 and 40 Hz with 5 seconds intervals between trains.

Histology

Prior to insertion within the brain, the electrodes were covered with fluorescent dye in order to confirm the final location of the electrodes. After the recording, animals were transcardially perfused with phosphate buffer and then urethane (4% in phosphate buffer). The brains were extracted and sliced in coronal sections of 80 um with a microtome (Leica). Slices were mounted and cover-slipped with saline and photographed; no staining was used.

Multi-Unit Activity

All analyses were performed using MATLAB. Broadband data was software filtered between 600 and 6000 Kz, and multi-unit activity (MUA) was

obtained by calculating the standard deviation of the signal for each electrode and assigning a spike each time the signal traversed 3 times this value.

State Determination

In order to determine the state of an animal, we added up the activity on those cortical channels that had the most clear spiking activity. We calculated the number of spikes in a sliding 200 ms window during 4 seconds of spontaneous activity preceding the onset of a click train, and subsequently calculated the fano-factor (variance/mean) of this distribution, as this number is higher in the inactivated state (Hollender, Doctoral Thesis). We considered "active" the 30 instances in which the fano-factor was the lowest and "inactive" the 30 instances in which it was highest. We used an equal number of instances in each state to ensure that there was no size effect in our subsequent analyses. This criterion visually matched the corresponding field potential characteristics.

Channel Selection

To select channels we assessed the MUA responses to a single click (the first in a 10 Hz series). The peri-stimulus-time-histogram (PSTH) was calculated around the 30 presentations at each activational state. The response was the maximum number of spikes in a 10 ms sliding window starting at the click and ending 100 ms after it. A channel was selected for further analysis if its response was above 3 std of the distribution of spike counts in an equivalent time window during the spontaneous activity. We

analyzed activated and inactivated states separately. A channel was selected for further analysis if it met the criterion in any of the states.

Response Amplitude

In order to assess the amplitude of responses to the successive clicks in a train, we first calculated the latency to peak. To do this we pooled together all the clicks in all the trains of a given frequency, calculated the PSTH and smoothed it with a sliding square window of 10 ms. Then we calculated the time to the maximum evoked response. Once the latency was determined we counted the mean number of spikes in the 10 ms surrounding it, after each click in the train. As we look at amplitudes from baseline (not the absolute amplitude) we subtracted the baseline rate from the values obtained.

Response Latency

In order to analyze evolution of latencies to a series of clicks we smoothed the PSTH of the trains with a 10 ms sliding window and calculated the time to peak after each click in the train.

ROC Area

The ROC area is the area under the ROC curve. This is a non-parametric measure of the overlap between two distributions.

To calculate the ROC-Area of the response to a given set of clicks we calculated the PSTH for those clicks and found the latency to peak. Then we calculated the distribution of spike-counts in the 15 ms window surrounding this latency in response to each of the clicks. We compared this result with a

surrogate distribution. The surrogate distribution we chose was the one we estimated to best match the conditions in which the click was presented, but applied to a case in which the click was omitted.

When analyzing the ROC-Area for single clicks (first clicks in the 10 Hz series), the surrogate group we chose was the spontaneous activity.

When analyzing the steady state we chose the time that would correspond to the response to the last click in the train, if there was one more click. For example, if analyzing 10 Hz trains that had a response latency of 10 ms, we chose a surrogate time centered on 110 ms after the last click in the train. Because the system is in a steady state, we assume that this is the activity that corresponds to an omitted click.

We chose a time window of 15 ms because that is the typical length of thalamic responses, which are shorter than cortical responses. We used the same window in both structures so we could compare them directly.

Vector Strength

To compute vector strength (VS) we used the following formula:

$$\text{vector strength} = (1/n)\sqrt{\sum \left(\cos(2\pi t_i/T)\right)^2 + \sum \left(\sin(2\pi t_i/T)\right)^2}$$

Where n is the number of spikes considered, ti is the timing of each spike and T is the period of the clicks in the train.

Across Trials Mean Correlation

To calculate the Across Trials Mean Correlation (ATMC) we smoothed the spike series of each trial with a 10 ms window, then took the neural responses for a given pair of presentations and calculated their correlation

across time. We repeated this procedure for all possible pairs of presentations and then calculated their mean. (Goard and Dan 2009)

Rate Equalization

To perform rate equalization we removed spikes randomly from the spike trains with the highest rate.

Statistics

For all group statistics we used the Wilcoxon Signed Ranks Test.

Results:

We successfully performed thalamo-cortical recordings from 6 rats as they went through the different activational phases; recordings were 1-3 hours long. We found a total of 53 cortical and 105 thalamic channels that responded to a click. The responding cortical channels were typically located in layers IV to VI based on cortical depth and current source density analyses. The thalamic channels whose position we could estimate by histology were located on ventral and dorsal parts of the medial geniculate (MGv and MGd).

In order for the reader to get a better understanding of the nature of the data analyzed in this work we will first show raw examples and some peri-stimulus time histograms (PSTH).

Spontaneous Activity

An example of a broadband (unfiltered) cortical recording corresponding to an inactivated state can be seen in figure 3. 1-a. Here the alternation between neuronal bursts and periods of silence can be easily seen. It can also be appreciated that periods of neural bursting occur during the downward deflection of the deep-layers' slow oscillation, or the upward deflection of the superficial layers'. The latter one corresponds to what is traditionally measured as EEG. These oscillations are not present during a period classified as activated state (see methods) as shown in figure 3.1-b. It can be observed that in the latter case the spikes are more evenly distributed throughout the record. Figures 3.1-c and 3.1-d show thalamic recordings performed at the same time as those depicted in figures 3.1.a and 3.1.b. The differences in thalamic rate are evident.

Evoked Activity

Neurons responded to trains of clicks. In figure 3.1-e we see the cortical response to a set of clicks in the inactivated state (immediately after the period showed in figure 3.1-a). The raw data suggests that the initial response to the train is strong but as the train progresses the activity starts looking more like spontaneous activity.

The response looks steadier in the activated state, shown in Figure 3.1-f, corresponding to the time immediately following figure 3.1-b.

Figures 3.1-g and 3.1-h show the thalamic responses corresponding in the inactivated and activated state, corresponding to the time periods following figures 3.1-c and 3.1-d. In both cases it can be seen that the responses are strong and steady throughout the record, although in the activated state the response is harder to discern because of increased background activity. These general observations will be properly quantified within the following pages.

Averaged Activity

We calculated the peri-stimuli time histograms (PSTHs) of the examples given above for 30 presentations of the click trains. Figure 3.2 (top row) shows the mean activity evoked in all responding cortical channels at a depth of 400 um to 800 um, by a 10 Hz click train presented during the different states. It can be seen from this example that the response in the inactivated state drastically decays after 1 second, unlike what is seen in the activated state. The bottom row shows thalamic responses. In both cases the response is quite similar as far as strength and temporal evolution. It can also be seen that the thalamic baseline activity is about three times bigger in the activated state.

Population Data

Spontaneous Activity

We calculated the baseline activity in the thalamus and cortex, which can be seen in Figure 3.3. Cortical MUA was significantly lower in the cortex

($p<10^{-4}$ Wilcoxon signed rank).The mean spontaneous cortical MUA in the inactivated state was 78% of that in the activated state. In the thalamus the percent difference was more prominent (only 43% of the activated state) and the difference more significant ($p<10^{-17}$). Here only two out of 105 channels had more spikes more during the inactivated state.

Temporal Evolution of Evoked Responses

In order to investigate how the responses change over time we first calculated the mean peak latency (see methods. To do this we estimated the latency to peak by calculating the PSTH around all clicks in all trains (of a given frequency). An example of this can be seen in figure 3.4. In this example the latency to peak in the activated state is 13 ms and in the inactivated state is 15 ms.

Once the latency is estimated we calculate the mean activity at this latency after eACh click in the train (see methods).

Figure 3.5.a shows the evolution of peak values for all cortical channels. At 10 Hz, a transient period lasting for about 10 clicks (1 second) is clearly evident. After this period the differences become stable and very significant ($p<.0001$). A similar effect is seen in the thalamus, although the differences are not as dramatic as in the cortex. These differences will be further quantified.

Although for higher frequencies the time course of the adaptation seems shorter than for 10 Hz, a clear transient phase can also be distinguished. The transient character of the response was always present in

the inactivated state, and it was slightly different in every experiment. This can be seen in Figure 3.5.b, which shows the peak evolution for each rat (10 Hz, mean across all cortical channels). In all rats but one the activated response was steady from the very beginning.

Given this difference between the onset and the steady state of the response, two time periods will be considered separately in the following analyses, one corresponding to the onset and other corresponding to the time period from 2 to 5 seconds, which we will call the steady state period.

Mean Evoked Responses

In order to quantify the strength of the responses, we averaged the peak responses presented above. We did this separately for the first click and for the steady state period. (Figure 3.6) The average cortical response to a single click in the inactivated state was 81% of that in the activated state ($p<10^{-2}$). This pattern was reversed in the thalamus, where the response in the inactivated state was significantly stronger 114% ($p<10^{-4}$).

During the steady state of the 10 Hz train, both the thalamus and cortex responded less in the inactivated state ($p<10^{-10}$), although the difference was much stronger in the cortex (37%) than in the thalamus (85%). A similar pattern was obtained at 20 Hz and 40 Hz. All values are nested in the figures.

Conclusion:

The responses to single clicks are significantly different across states in both structures, although the differences are not very prominent. The steady

state cortical responses are dramatically reduced in the inactivated state. A similar effect, but much weaker, is seen in the thalamus.

Discriminability

As mentioned in the introduction, the absolute response to a stimulus might be a poor indication of how well this response might be "read" by target structures, thus an analysis that addresses the latter issue should include statistics of the activity in the absence of the stimulus, which we shall call surrogate. In choosing an appropriate surrogate we should look for a situation that is as similar as possible to that at which the stimulus would arrive.

A proper surrogate to compare the activity evoked by a single click is obviously the spontaneous activity. However, during the sustained period, this comparison might be inappropriate because when a given click arrives, the neural response might still be influenced by the previous click. The ideal surrogate in this case would be the activity when the click is omitted. Instead of omitting clicks, we chose the period of time after the end of the train that would correspond to the presentation of an extra click. As the neural response is in the steady state we can assume that this activity would correspond to that of an omitted click.

In order to quantify the differences between the click-evoked and surrogate activities, we used the ROC area (see methods), which is a non-parametric way to estimate the overlap between two distributions. This measure is better than the d' measure in that it makes no assumptions as to

the underlying characteristics of the distributions. The use of this method will also allow us to compare our results with previous works in the auditory (Massaux, et al. 2004) and visual (Guido, et al. 1995) thalamus.

An example of a thalamic channel responding to clicks in 20 Hz trains can be seen in figure 3.7. In this example we see that even if the mean response in the activated state is bigger (see PSTH), so is the overlap between this response and its surrogate (distributions at peak), which is reflected as a slightly smaller ROC area.

The ROC curve shows the cumulative distribution of the evoked response vs. the cumulative distribution of the surrogate. If both distributions are equal then the curve will be a straight line connecting the corners of a 1x1 box, and the area under the curve will be 0.5. The less overlapping the distributions the closer the line will be to the sides of the box, in such a way that an area of 1 means there is no overlap between the distributions and the stimulus can be discriminated without ambiguities. In the example given above the ROC area corresponding to the inactivated state is 0.93 and the activated is 0.89, so the stimulus can be better discriminated in the inactivated state.

Population Data

There was no significant difference in the cortical discriminability to single clicks (p>.7), but there was a marked effect for the sustained period at all frequencies considered (p<.001). The thalamic discriminability for single clicks was markedly better in the inactivated state ($p<10^{-7}$). There was no

difference found for the 10 Hz and 20 Hz trains, but discriminability was better

in the inactivated state at 40 Hz. Figure 3.8.

In order to quantify the strength of the differences we subtracted the

expected value of the ROC curves under the null hypothesis (0.5), so we

created a coefficient M:

$$M = \frac{Mean(ROC_i - 0.5)}{Mean(ROC_a - 0.5)}$$

Where ROCi and ROCa are the ROC-Areas in the Inactivated and

Activated states respectively.

Values of M can be seen in the figures.

The results of this analysis depend of course in the period chosen to be

the surrogate; we present the data from the period that we consider to be the

most reasonable. However, the results still remain significant when we take

the spontaneous activity as the surrogate period for the sustained response

(not shown).

Conclusion:

Even though the absolute responses are usually bigger in the activated

state; this does not directly translate into a better detection of the stimulus.

When the statistics of the background activity are taken into consideration, the

inactivated thalamus is either equally or better able to discriminate the

presence of the stimulus than the activated thalamus.

In opposition to the thalamic results, detection in the cortex is profoundly impaired in the inactivated state during the sustained phase of the response.

Frequency Synchronization

The method of "Vector Strength" has been frequently used to measure the synchronization of a neural signal to a periodic stimulus. See Joris, et al. (2004) for an excellent review. A Vector Strength value of 1 indicates that all the spikes occur every time at exactly the same phase of the stimulus, which is equivalent to saying that all of the spikes occur with exactly the same delay. A Vector Strength value of 0 means that the timing of the spikes holds no relationship with the stimulus. As this is a measure of the periodicity of the response in relation to the signal it makes most sense to perform the analysis on the sustained part of the response, since the onset is not periodic.

Figure 3.9-a shows the results for the sustained period at different frequencies. Frequency synchronization was always higher in the cortex in the activated state ($p<10^{-8}$). Surprisingly, this effect was reversed in the thalamus, where synchronization was always stronger in the inactivated state   ($p<10^{-4}$).

This method has a bias towards low frequency rates, as evidenced by the fact that the Rayleigh value, typically used to establish significance, is proportional to the number of spikes (Liang, et al. 2002) .In order to control for this bias we performed the same analysis on the rate-equalized data, that is, on data in which spikes from the state with the highest rate were randomly

deleted to match the rate of the lowest. The results presented above are still significant under this equalization.

As mentioned before, this method makes the most sense during the steady state; nevertheless, in order to compare the results with other studies (Issa and Wang 2008), we also analyzed the results during the onset. In order to do this we considered only the response to the first five clicks of every frequency. The results generally parallel those found for the sustained period, although the effect is not as pronounced (figure 3.9-b). For example, during the sustained period no cortical channel synchronized more in the inactivated state, but during the onset many channels do.

The only significant difference between the sustained period and the onset is the thalamic response at 40 Hz, where the difference between activated and inactivated states is not significant (p>.6). The differences with the sustained period might be explained in part by the fact that the data set was smaller, thus adding noise, and the fact that transient evoked activity will "dilute" the phase locking part of the response.

Conclusion:

The inactivated cortex is very impaired at locking to a periodic stimulus compared to the activated cortex. In the thalamus this effect is reversed; the inactivated thalamus actually phase-locks better than the activated.

Reliability of the Response

The previous analyses assume that the response takes a specific form:

To calculate the ROC-area we only consider spike counts in a given window.

Vector Strength analyzes the response at the frequency of the stimulus; (as a matter of fact it is equivalent to taking the power of the Fourier transform of the neural response at the frequency of the stimulus and normalizing it by the 0 frequency component). A caveat of using this method is that the frequency of the stimulus is actually not perfectly defined, in the sense that it is composed of the fundamental frequency but also of many harmonics, which are ignored when calculating vector strength. Also, even if the stimulus was perfectly harmonic, the neural signal might introduce higher harmonics that will phase lock to the stimulus and therefore carry valuable information.

In the following section we use an analytical method that does not make any assumptions about the shape of the response, but looks at how much the responses to a given stimulus resemble each other. Since this analysis does not require the signal to be stable, we no longer need to separate the onset from the steady state. This analysis will also permit a direct comparison with the results reported by Goard and Dan in the Visual System.

In order to estimate the reliability of the responses we use the across trial mean correlation (ATMC, see methods). A value of 1 indicates that the MUA is modulated in exactly the same manner each time the stimulus is presented. A value of 0 indicates that the MUA is independent of the stimulus. Results are shown in Figure 3.10.

At 10 and 20 Hz the cortex was markedly less reliable in the inactivated state ($p<10^{-5}$), this is not true for 40 Hz ($p>.06$), presumably because in both states the response is much degraded.

In contraposition to the cortex, the thalamus was actually significantly more reliable, although the difference between states was generally not big, the most affected frequency being 40 Hz with a ratio of 118%.

When only the sustained period is taken into account (data not shown), the cortical effects are even more pronounced. The thalamus shows no significant effect at 10 and 20 Hz but is significantly less reliable in the inactivated state at 40 Hz ($p<10-6$), the mean inactivated ATMC is in this case 66% of its counterpart.

Rate Equalization

As stated by Goard and Dan (2009), this method is biased towards higher rates; accordingly, they show that their results are still valid when they equalize rates by randomly removing spikes from the traces with the highest rate. A similar technique is used by Mitchell et al.(2009) when looking at correlations between neurons. The relationship between rate and correlations is analyzed in great detail by De la Rocha et al (2007).

A general explanation of why an increased rate translates into stronger correlations, and under what conditions this is true, is presently lacking. We suggest that this is because of a sampling issue. In effect, an increased rate improves the sampling of the underlying processes that modulate spiking, producing a better estimation of such processes. These concepts are

developed in great detail in Chapter 4, where we establish under which situations an increase in rate might lead to an increase in correlation and we use this to evaluate the validity of rate-equalization.

Results:

When rate equalization is performed, the results in the cortex are not markedly different, but the differences in the thalamus become much more pronounced. For example, the thalamus' response to a 40 Hz click train is 224% more reliable in the inactivated state. When only the sustained period is considered, the thalamus is more reliable at all frequencies (p<0.005).

Conclusion:

The cortex is markedly more reliable in the activated state. The thalamus is significantly more reliable in the inactivated state, although by a small factor.

When rate equalization is employed to factor out how much of the reliability is solely due to a bigger firing rate, we find that the thalamus is much more reliable in the inactivated state.

Histology

In the two rats in which traces of both thalamic shanks could be found, they were all found to traverse the Medial Geniculate, including the dorsal and the medial portion. In another rat only one shank was found in a similar location, as the distance between shanks is 0.5 mm the other shank was estimated to be in the same area. These three rats correspond to the magenta, blue and green dots in the figures.

Figure 3.12 shows one example of a shank and the Vector Strength in the channels corresponding to it. The correspondence should be considered to be loose, since the exact location of the electrodes' tip cannot be established and the fixation process shrinks the tissue. This is a disadvantage of the type of electrode utilized, which does not allow for performing electric lesions to mark the recording site.

In order to evaluate if the results are generally still valid when the ventral portions of these shanks are considered, we calculated the Vector Strengths using only the lower half of the shanks. Considering LFPs' evoked responses and the depth of the penetrations (see methods), this is a conservative estimation of MGv. At 10 and 20 Hz the results paralleled those for all the rats, at 40 Hz the tendency was in the same direction but not significant, which might be an effect of the lower number of data points. This result can be seen in figure 3.13.

In two rats we could find the location of the cortical penetrations, and in both cases it belonged to area TE1, generally corresponding to the primary auditory cortex. These two rats correspond to the red and the blue colors on the plots. While both follow the general trend, the effect in the blue rat is sometimes more prominent at 20 Hz. The penetrations for these two rats can be seen in figure 3.14; the dye does not penetrate as far as the tip of the electrode.

Conclusion: The location of a sub-set of the recording sites could be estimated; the data in these recording sites follows the general results

.

Summary:

We performed dual recordings of the auditory thalamus and cortex while applying auditory stimulation at different modulation frequencies. We found that the neural responses varied greatly with time. A clear onset response behavior was distinguishable from a steady state and therefore the two periods were analyzed separately. We analyzed the data in four ways:

1. The total rate of evoked responses.

2. The discriminability of the stimulus based on the neural response.

3. The amount of neuronal phase locking to the stimulus.

4. The reliability of the neuronal response towards the stimulus.

Although in many cases they were statistically significant, the onset responses were not very different across states in any of the structures. However, during the sustained period, inactivation profoundly degraded cortical responses under any measure and frequency.

Except for the total magnitude of the response, which was slightly stronger in the activated state, the thalamic responses during the inactivated state were either stronger or equal than during the activated state, in this sense we can say that the responses were smaller but cleaner.

Discussion:

Perhaps no analysis is good on its own. Beyond the few degrees of freedom that the use of any analysis technique introduces (just by the choice of parameters) there are the infinite degrees of freedom associated with the choice of one technique over another, a choice that is only limited by what is reasonable.

Although the analysis techniques applied here measure different aspects of the response, they are partially overlapping, in the sense that a strong evoked response will affect all measures in a similar manner. That said, approaching a problem from many angles might allow for a better understanding, and make up for any bias that a specific technique or parameter choice introduces. Also, by applying the same analytical methods that have been used in previous work done in different modalities, we will be able to perform a more direct comparison of the results. Importantly, these analyses quantify what can be observed from the visual inspection of the raw electrophysiological records (as the example in figure 3.1), which is:

Inactivation has a profound degrading effect in cortical responses to the stimulus, this effect being especially prominent after a period of adaptation.

Inactivation decreases the baseline rate in the thalamus, while the response to the stimulus is still clear, and perhaps even clearer since it is presented against a less noisy background.

Comparison to Previous Studies:

<u>Auditory System:</u>

As mentioned in the introduction, various studies in the auditory cortex show no consistent changes of evoked responses across activational states (Pena, et al. 1999; Edeline, et al. 2000; Issa and Wang 2008). In line with these studies, we find that, even if there is a general trend to decreased responses to single clicks in the inactivated state, many neurons fire more in this situation.

In opposition, we show that after a transient onset period of about 1 second, the cortical responses become consistently much weaker in the inactivated state. This difference between onset and steady state has not been reported, since the studies mentioned generally used short stimuli (< 100 ms). An exception is the work by Wang. In this case the stimuli are longer (~ 500 ms), but they are also complex, and they show that the neurons are responding for only 100-150 ms of this interval. This again stresses the fact that the main effects of brain state are related to timing.

Our results for the steady state are in line with work by Marguet at Harri's lab (Stephan Marguet, Doctoral Thesis) which shows state-dependent effects in response to long sequences of amplitude-modulated noise.

The only previous study that has addressed the timing of response was performed by Issa and Wang in the cortex (2008). Presenting sounds modulated at different frequencies and measuring Vector Strength, they find no difference between states. Using the same analysis technique but

extending our measurements into the sustained period, we find that the activated cortex is much better at phase locking. However, when we look at the onset, a time period more akin to that in Wang study, we see that the differences are still significant but not as big, with many channels that phase lock better in the inactivated state.

In the thalamus, our results differ from those of Orman and Humphrey (1981) and from Edeline et al. (2000), who found an increased response to isolated stimuli during activation. However, the finding by Edeline et al. that the signal to noise ratio (measured as the ratio of the response to the baseline) is not affected by state, is in line with our finding that there is no difference in detectability. As a matter of fact, a later work by the same lab uses ROC-Area to measure detectability and finds no difference between states (Massaux, et al. 2004). In this realm we see a significant, but very small, improvement in the inactivated state, an effect that is in line with that of Guido and Sherman, in visual LGN (Guido, et al. 1995).

The works of Massaux and Guido analyze thalamic discriminability to isolated or low frequency stimuli, but, based on Castro-Alamancos' work in the somatosensory system, it is to be expected that differences might arise when the system is probed with a train of pulses. The result that we obtain in the thalamus when we analyze the response to isolated stimuli carries on into trains of stimuli. We either see no significant differences between states, or the differences are significant but not big (see figure 3.8).

<u>Somatosensory System</u>:

Our results are in striking contrast to those of Castro Alamancos' in the somatosensory system. In the cortex he observes that the response to the first click in a series is much smaller in the activated state, whereas we find they are slightly bigger. We sometimes find, however, a build-up of responses in the inactivated state (figure 3.5.b), in such a way that the responses within the first 600 ms are bigger. As the time course of this build-up is not consistent among rats we did not analyze this systematically.

The results are also very different in the thalamus, where he reports that at 10 Hz the responses are about 10 times bigger in the activated state (see figure Aux.2), whereas we find that they are only 1.2 times bigger (figure 3.6). The differences we find are not striking at higher frequencies either.

The conditions under which his experiments were performed are very similar to ours as his were also done in urethane anaesthetized rat. The main differences are that he measures single units (but averages the data), and he induces activation by electrical stimulation of the brainstem reticular formation.

<u>Visual System</u>**:**

In accordance with all work done in the visual system, we show that inactivation is associated with a depressed sensory response. The magnitude of this depression, however, is very different. The differences we find are much smaller than those reported in the visual system. See for example (Livingstone and Hubel,1981) Figure Aux.4.

Regarding detectability, our results are more in accordance with those of Guido and Sherman (1995), who showed better detectability in the inactivated state. Like us, they used a phasic stimulus, but they presented it at 4 Hz only. Given the literature, our finding that this result is carried on at higher frequencies is not trivial.

We use the same method as Goard and Dan (2009) to quantify the reliability of the responses. Our results parallel theirs in the cortex, in that inactivation degrades reliability. However, while they find this effect to hold in the thalamus as well, we find that reliability is slightly but significantly improved in the inactivated state. As previously mentioned, Goard and Dan state that the quantifying method has a bias towards higher rates, therefore they equalize rates by randomly removing spikes from the state that has the highest rate. When they perform this control their results remain unchanged. When we perform rate equalization, the result that we found becomes much more prominent, reliability being considerably stronger in the inactivated state.

Conclusion:

Our results for the onset of responses are generally in line with other work performed in the auditory system, in that the differences between states are either absent or not prominent. We are the first to probe the system with longer stimuli, finding that brain state has a profound influence in neuronal activity.

Our results are at odds with previous work in the somato-sensory and visual thalamus. While those studies showed that in the inactivated state the

neural representation of the stimulus is much degraded, we show that it is unchanged or even improved, and that this happens at all frequencies studied. This indicates that, contrary to what has been shown in other systems, the auditory thalamus might not be a state-dependent sensory gate.

General Discussion:

<u>Onset Vs. Sustained</u>

Although there are numerous similarities between the anaesthetized and sleeping animal (Clement, et al. 2008), and the anaesthetized preparation has been extensively used as a model of SWS, (Steriade and Mc.Carley 2005), there is one major difference: The anaesthetized animal does not wake up. The stimuli that we present might suffice to awaken the animal under a normal sleeping condition, therefore the sustained period that we observe might not be physiologically realistic. Muratta (1963) went as far as to habituate the animal so he would not wake up upon stimulation. Issa and Wang (2008), on the other hand, were only able to present short stimuli in order not to wake the monkey up.

The fact that the anaesthetized animal does not wake up might be considered a liability, but it might also have the advantage of allowing us to study sensory responses in a sleep-like situation, without the intrusion of waking activity, much in the same way a lesion study allows us to learn from a structure by removing it.

It is interesting that sensory responses in the inactivated state are enhanced almost to the level of wakefulness for about 1 second; this is

probably the amount of time that might take the animal to wake up if that sensory information is relevant. Remember that humans will easily wake up to their own name (Langford et al. 1974). We speculate that this might not be the case if that name is nested in a long sentence. In this regard it is important to note that human subjects show signs of complex processing starting 1 second after stimulus presentation, as showed by Oswald et al. (1960) who compared responses to the subjects' own names played forwards vs. backwards.

In the somatosensory cortex the onset response can actually be bigger in the inactivated state (Castro-Alamancos 2004), but this effect disappears after one spike, so we could say there is sensory processing but not for complex stimuli. Back to the waking up example, often an animal will wake up if touched once anywhere on their body; it is not necessary to provide a particular temporal pattern to elicit waking.

Conclusion:

Given that an important function of sensory processing during sleep is to signal the necessity of waking up, we speculate that an onset response lasting up to 1 second allows for the neural processing of stimuli of this duration, which might be behaviorally relevant.

Thalamus Vs. Cortex

Across the systems it is common to see a degradation of the cortical response in the inactivated state, but there has been little discussion regarding the source of this degradation, probably because it has been so widely

assumed that the "thalamus' gate is closed". In the words of Castro-

Alamancos:

..."during activated states, the (cortical) responses to high-frequency

stimuli are larger than during control states (quiescent). This was expected

among other reasons because thalamic neurons respond with higher efficacy

to high-frequency sensory inputs during activated states". (2004)

In other words, cortical responses are stronger during activation

because their thalamic input is stronger.

A similar idea is held by Goard and Dan (2009). By applying the

muscarinic antagonist atropine to their cortical recording site, they managed to

induce a situation in which the cortex stayed locally inactivated, (see

introduction), regardless of the state of the thalamus. In this situation they

found that the cortical signal was still greatly boosted when the thalamus is

activated through basal forebrain stimulation. This argues against a function of

cortical activation per se in sensory coding. (Figure Aux.5). Since it has been

shown that nicotine produces an augmentation of visual (Disney, et al. 2007)

and auditory (Kawai, et al. 2007) sensory responses, Goard and Dan suggest

a possible role of cholinergic nicotinic receptors in this augmentation, but they

find that the boosting effect is not abolished upon application of nicotinic

receptor antagonists. Based on these results they suggest that what is being

really affected is the cortical input from the thalamus. And they indeed find that

thalamic responses are greatly improved upon activation. So one is left to

wonder: which, if any, is the role of cortical activation and cortical acetylcholine release on sensory processing?

It has been shown that activation induces a de-synchronization of cortical cells; some authors have actually called this the "de-synchronized" state. Goard and Dan propose that the role of activation is this very feature. This will be discussed later in greater detail.

The view that activation does not play a major role in cortical responses is also supported by Stoetzel et al. (2008). In this work they performed dual thalamo-cortical recordings and estimated the cortical responsiveness to thalamic activity across activational states, finding that the cortex is equally responsive in both.

Our experiments suggest a fundamental role of cortical activation in cortical sensory processing, since in the inactivated state the thalamic signal is largely preserved, but the cortical signal is much degraded. This is in accordance with an earlier work by Steriade and Morin (1981) in the somato-sensory cortex, in which they bypassed thalamic effects by stimulating the thalamic radiation. They found that cortical responses to high-frequency stimulation are enhanced during activation, an effect that persists in thalamic-lesioned animals, so cannot be attributed to the influence of baseline thalamic activity or anti-dromic stimulation.

It is important to note that our experiments were performed under very similar experimental conditions as those of Goard and Dan, so our results are not easy to reconcile. One possibility is that the degradation in thalamic

responses that they see is only part of the explanation for the degradation of cortical responses, but additional factors could potentially be seen under different experimental conditions.

Conclusion:

Strong evidence supports the position that the thalamus is most strongly affected by activational state, in such a way that the changes seen in the cortex are inherited from it. However, our data strongly suggests otherwise. A possible reconciliation lies in there being differences between the systems, but even in that case the function of cortical activation in the visual system would remain elusive.

<u>Across-Cell Correlations</u>

Because in the inactivated state the cells tend to fire together (during the up-states), or be silent together (down-states), their activity is highly correlated. This correlation is abolished during activation. (Figure Aux.5).

Between-cell correlations can have a detrimental influence in coding (Zohary, et al. 1994). If the neurons are completely independent we can use the population response to estimate the characteristics of the stimulus. If all neurons are doing exactly the same thing we will have no more information than if we had only one neuron; therefore, looking at the neural activity when the stimulus is presented only once will not give us much information. But even if the neurons are very strongly correlated, the response to successive presentations of the stimulus will not be. Therefore we will gain much information by averaging across many presentations. Since tight correlations

will have such a big effect on single trial responses, we might expect the reliability to be very affected, but measures that rely on averages across many presentations will not be affected much at all. Absolute response and Vector Strength, belong to this latter type of measurement. So our data indicates that activation has an effect in the cortex that goes beyond de-correlating the neurons, greatly modifying the strength of the responses.

The fact that we find an effect when using these measurements does not mean that activation fails to improve single trial discrimination; it means that its effect is deeper than that. The importance of single trial discrimination should not be understated, since, for example, there might not be enough trials to estimate the proximity of a tiger.

Probably the best way to analyze the effect of correlations on single trial responses is to shuffle the response of different neurons in different trials. Since we are focusing on MUA this technique is not available to us, but Hollender (Doctoral Thesis), finds that between-cell correlations are not enough to explain the state dependence of cortical responses.

Conclusion:

Because we measure MUA, our data does not allow us to specifically address the issue of between-cell correlations, but the fact that we see state dependent changes in measures that involve averages across many independent trials, indicates that the effect of cortical activation does not stop there.

Effect of Rate

We have been stressing the importance of using measures of coding that go beyond the total response, but we should not rule out the effect of total response either. As can be seen in figure 3.6, inactivated thalamic responses at 20 Hz were 75% of their activated counterpart. Although this is far from the 10% reported by Castro Alamancos in the somato-sensory system, there is the possibility that cortical cells might be sensitive to this difference. Furthermore, we only reported evoked responses in relation to the spontaneous activity, but these evoked responses are riding on top of different baselines, which are more than twice as high in the activated state. If the baselines are not subtracted the differences between states would become even larger. It is also conceivable that an increased thalamic baseline firing rate might keep cortical targets closer to threshold, making them more susceptible to incoming evoked inputs.

However, the evidence generally points to a detrimental effect of higher baseline rates on cortical responses. Castro-Alamancos (2004) shows that thalamocortical synapses are greatly depressed when the thalamus is firing at high rates. And Stoelzel et al. (2008), show that cortical neurons are more likely to fire after a period of decreased thalamic inputs, a result also found by Bruno and Sakmann (2006).

A more subtle effect of rate is related to the fact that, as shown in chapter 4, under certain conditions a higher rate can account for a better estimation of the underlying rate modulation. This is true for rates up to a certain point, but at very high rates this estimation saturates. So if a cortical

cell receives an input from only one thalamic cell, a rate increase from this cell could make a big difference, but if the number of inputs is big enough a rate increase will not have much of an effect. Based on the baseline thalamic firing rates that we find (about 200 Hz in the activated state) and the typical thalamic baseline rate (10 Hz), we might assume that each electrode is measuring from about 20 cells. It is estimated that 85 thalamic neurons converge onto a cortical cell (Bruno and Sakmann 2006).

The fact that we do not see important changes in the across trial mean correlation of our MUA suggests that the input that the cortical neurons are receiving does not improve by activation's increased rate.

Conclusion:

Thalamocortical activity is definitively greater in the activated state, but is not clear that the cortical targets benefit from this increase. Our results indicate that an activated thalamus might, at best, increase the volume of activity while degrading the signal.

Conclusions:

The results presented here might explain an apparent contradiction in the literature, i.e. that the inactivated thalamus interrupts the flow of information to the cortex and that we can process complex, cortex-dependent, auditory information when we are asleep. We show that thalamic responses do not differ much across the states, and that cortical responses do so after an

onset period lasting approximately one second, which might be all it takes to detect a behaviorally relevant complex sound. By using experimental and analytical methods that are very similar to previous work in the visual and somato-sensory systems, we show that the auditory system is essentially different in this respect. The fact that our result differs from previous work done in the auditory system may be due to the rather long stimuli we utilize, which allows us to see differences in cortical processing that only become clear after the onset period.

Rapid Auditory Processing (RAP) is the ability to discriminate sequences of sounds presented in a fast succession. It has been shown (Tallal and Piercy 1973) that deficits in RAP are more prevalent in children that suffer from specific language impairment (SLI), and that training to improve RAP can improve language skills ((Merzenich, et al. 1996; Tallal, et al. 1996)).

An electrophysiological correlate of the relation between RAP and SLI has also been reported,as EEG responses to the second sound in a rapid pair are diminished in children suffering SLI (Benasich, Choudhury et al. 2006). Correspondingly, it has been shown that EEG power in the gamma band, a characteristic of activation, predicts future language impairments in infants (Benasich, Gou et al. 2008).

Based on this data it was proposed by work in Harris' Lab (Hollender, Doctoral Thesis), that deficits in RAP could be associated with the infants 'capacity' to activate their auditory cortex. Accordingly, they showed that

activational state can strongly influence the way in which the cortex processes

rapidly varying sensory stimuli. However, since brain activation is a global

phenomenon, it was not clear if this effect was due to differences in cortical

processing or in cortical inputs, a difference that might be fundamental in

finding a therapeutic treatment for language deficits. The fact that we find

major differences in cortical responses and minor differences in thalamic

responses suggests that the loci of the deficits could be in the cortex.

Although we have focused on different states of the anaesthetized

animal, characteristics of these states are observed not only through the

sleep-wake cycle but also during wakefulness, either in an (apparently)

passive situation (Stoelzel, et al. 2009) or in relation to particular tasks like

shock avoidance (Castro-Alamancos, 2004) or active whisking (Poulet and

Petersen 2008). Increased activity in the high frequency bands, associated

with activation, can also be seen during working memory in monkeys (Pesaran,

et al. 2002) and humans (Tallon-Baudry, et al. 2005). Furthermore, some

basic characteristics of cortical inactivation, like increased fano-factor or

between-cell coherence in low frequencies (Mitchell, et al. 2009) , can be seen

in cortical sites corresponding to stimuli that need to be ignored, while an

increase in gamma band activity (power and coherence) can be seen in

cortical sites corresponding to stimuli that are task-relevant (Womelsdorf, et al.

2006). In this case, higher gamma activity predicts shorter reaction times.

This raises the possibility that the brain mechanisms that serve to

filter/enhance sensory information during sleep/wakefulness are the same as

those seen under other behavioral demands. Our results support the idea that those mechanisms can act in parallel, and potentially independently, in the thalamus and the cortex.

Future Directions:

Brain state related differences in thalamo-cortical sensory processing have been studied in great detail for more than 50 years; however, the psychological experiences related to the waking-SWS-REM cycle do not stop at sensory processing, and are present even in the face of sensory isolation.

It has been suggested that the different psychological experiences found during sleep and anesthetic states are related to a loss of integration caused by a disruption in functional cortico-cortical connectivity (Alkire, et al. 2008). In support of this view, it has been shown that cortical propagation of activity induced by trans-cranial magnetic stimulation is reduced during SWS (Massimini, et al. 2005) Since most of the thalamus does not receive direct peripheral input and does not project to primary cortices, but rather connects various cortical, hippocampal and hypothalamic structures, it has been suggested as a candidate to mediate these changes in functional connectivity (Sherman and Guillery, 2002), although differences in cortical connectivity based on intrinsic cortical activation should also be considered.

There is evidence that higher order thalamic structures can behave in relation to their inputs from cortex in a similar manner to how the primary

thalamus behaves in relation to its peripheral inputs. For example, the pulvinar receives ionotropic synapses from layer five of primary visual cortex. These synapses behave in much the same way as retinal inputs to the LGN (Guillery and Sherman 2002). Furthermore, the pulvinar exhibits the same burst-tonic firing modes associated with LGN (Li, et al. 2003), and is modulated in a similar manner by cholinergic agonists (Varela and Sherman 2007).

We therefore propose that:

1- Brain activation changes higher order thalamic responses to cortical inputs in the same way that it changes primary thalamus responses to peripheral inputs.

2- Brain activation changes cortical responses to thalamic inputs regardless of the cortical area and whether the input comes from primary or higher order thalamic nuclei.

3- These changes produce an increase in functional brain connectivity.

To test this hypothesis we propose an experiment that is very similar to the one we performed, but instead of presenting sensory stimulation we would use electrical stimulation of a given thalamic-projecting area and measure thalamic responses across the different states. We would also look at differences in the responses of the targets of these thalamic structures. By performing this experiment we could measure changes in brain connectivity in a causal manner and gain understanding of the mechanisms by which they can be brought about. Although the experiment could be performed in any

system, we think a good model would be the connections from the mammillary body of the hypothalamus to the anterior thalamus and from the anterior thalamus to the subiculum.

The anterior thalamus receives unidirectional ionotropic projections from the mammilary body (Guillery and Sherman 2002), which resemble the peripheral inputs to the relay thalamic nuclei. In turn, the anterior thalamus has reciprocal connections with the subiculum. Furthermore, all these structures exhibit head direction coding cells, and mammilary body lesions, but not hippocampal lesions, disrupt head direction cells in the thalamus (Blair, Cho et al. 1999). This evidence suggests that the anterior thalamus might act as a relay for higher order processing, connecting the mamillary body with the subiculum.

On the other hand, short range cortico-cortical connections might also be affected by activation. The succession of up and down states characteristic of SWS is associated with waves of activity that travel laterally through the cortex (Sanchez-Vives and McCormick 2000). In this state, electrical or chemical stimulation of a cortical site will elicit a propagating wave. Since the activated cortex shows no down-states, it is not clear that there would be any lateral propagation, or what form this lateral propagation would take. Thus, functional cortico-cortical connectivity would be affected.

According to this view, it has been shown in cortical slices that ACh diminishes the strength of lateral connections, reducing the spread of the cortical waves of activity, (Kimura, et al. 1999), but leaving responses to

stimulation of the white matter relatively unaffected. It remains to be seen if this effect persists in the more physiological condition of an activated cortex of a living animal. In order to elucidate that, we suggest stimulating a cortical site and recording the activity at an array of cortical sites nearby, to see if the propagation of waves is disrupted upon activation.

Conclusion:

The study of information processing under different brain states has been focused on sensory responses, but this might be only one aspect of a more general phenomenon. There is good reason to believe that brain activation changes the functional connectivity of the brain, reducing short range lateral connections in the cortex but increasing the cortico-cortical connectivity that is mediated by the thalamus. We think future experiments should be aimed at testing this hypothesis.

Start time 58 min ⬍ 48 s ⬍ 888 ms ⬍ Duration (ms) 4000 ◀
e : 58 min 51 s 757 ms (3531.757 s)

Figure 3.1-a Inactivated Cortex. Spontaneous activity.

Start time 188 min ⬍ 32 s ⬍ 923 ms ⬍ Duration (ms) 4000 ◀
: 188 min 33 s 585 ms (11313.585 s)

Figure 3.1-b Activated Cortex. Spontaneous activity.

Figure 3.1 Raw Recordings. Cortical and thalamic recordings during spontaneous activity or in response to a 10 hz click train across the different activational states. For each condition, depicted thalamic and cortical recordings are simultaneous. Each trace represents a channel. In the cortex the channels span 1.6 mm (cortical surface corresponds to the top trace). Thalamic channels are spaced 100 um. Vertical bar 1 mV, horizontal bar 250 ms.

Figure 3.1- c Inactivated Thalamus. Spontaneous activity.

Figure 3.1- d Activated Thalamus. Spontaneous activity.

Start time | 58 min | 52 s | 688 ms | Duration (ms) 4000
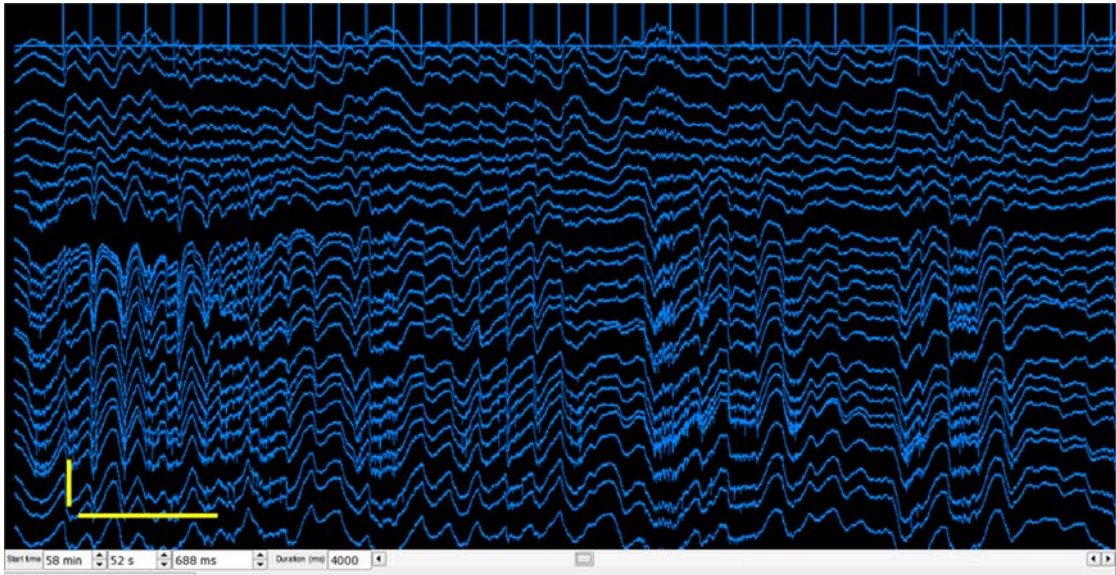
Figure 3.1- e Inactivated Cortex. 10 Hz Click Train.

Start time | 188 min | 36 s | 723 ms | Duration (ms) 4000
: 188 min 38 s 203 ms (11318.203 s)

Figure 3.1- f Activated Cortex. 10 Hz Click Train.

Start time 58 min 52 s 688 ms Duration (ms) 4000
: 58 min 37 s 391 ms (3517.391 s)

Figure 3.1- g Inactivated Thalamus. 10 Hz Click Train.

Start time 188 min 36 s 723 ms Duration (ms) 4000
: 188 min 39 s 139 ms (11319.139 s)
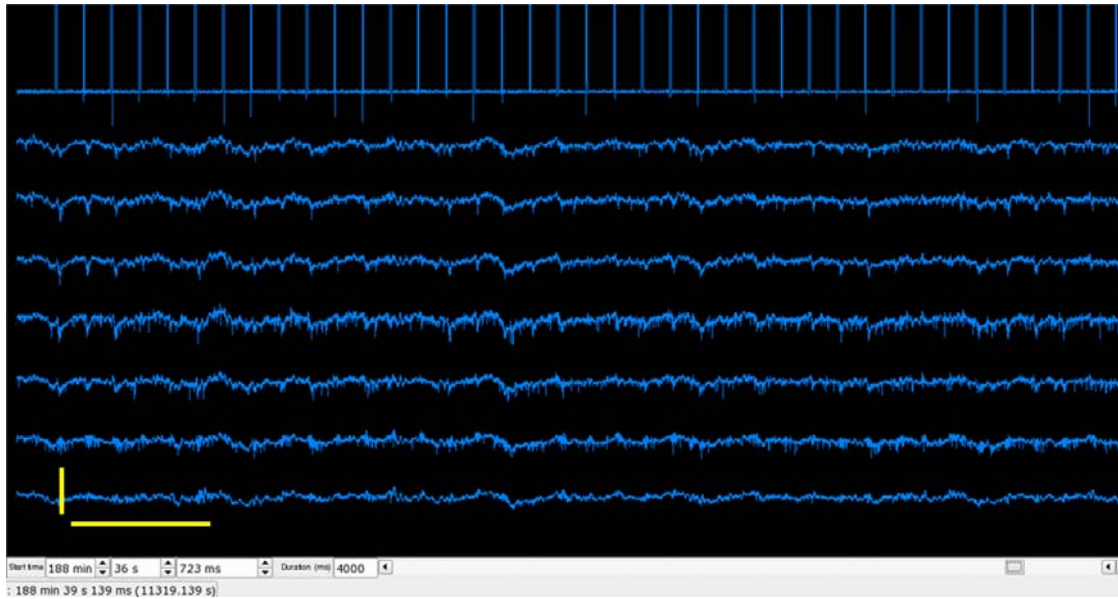
Figure 3.1- h Activated Thalamus. 10 Hz Click Train.
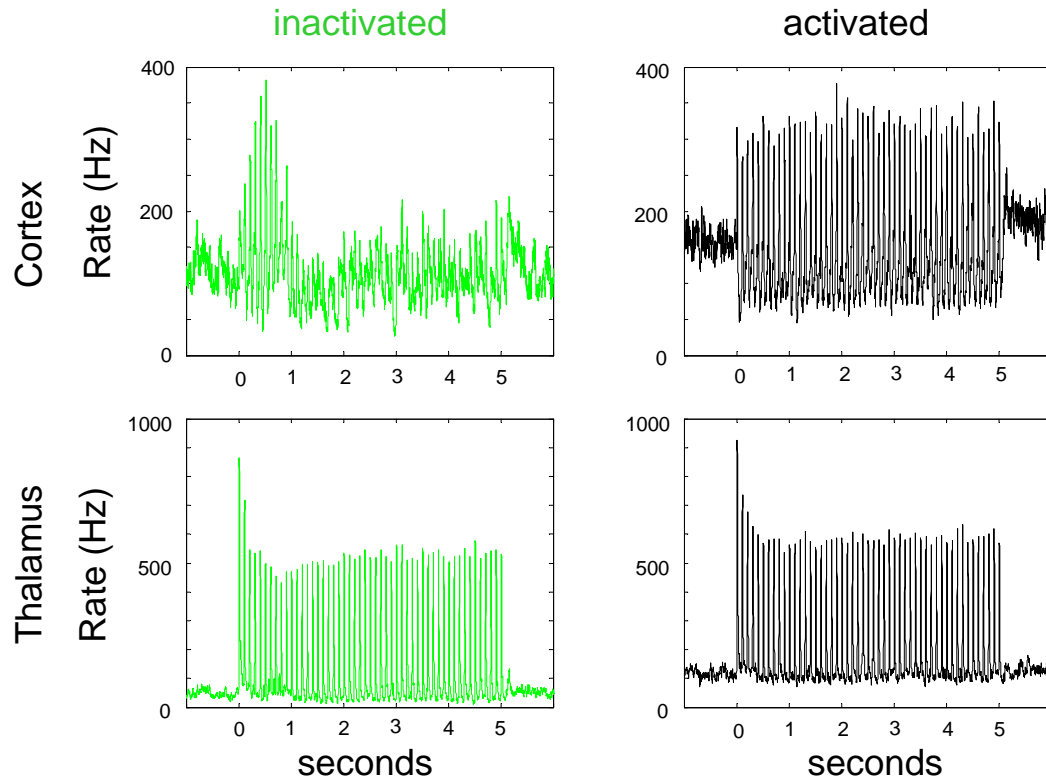
Figure 3.2 Example PSTH for a10 Hz click train. Mean of 30 presentations in each condition. Average of all selective cortical (400-800 um depth) and thalamic channels in an example rat. Data was averaged in a 10 ms sliding window. Inactivated in green, activated in black.
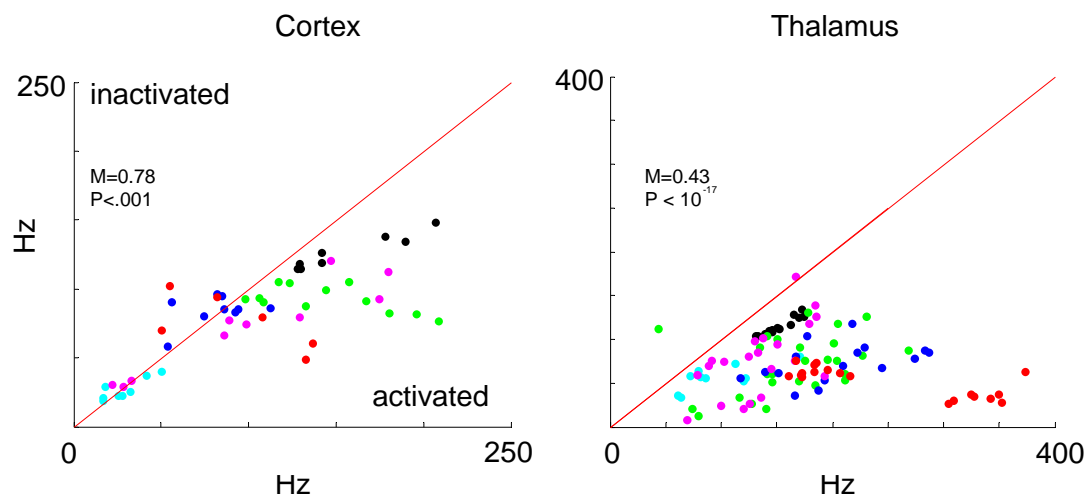
Figure 3.3 Spontaneous Activity in Cortex and Thalamus. Each dot represents a channel and each color represents a different rat. M is the ratio of inactivated vs. activated activity.
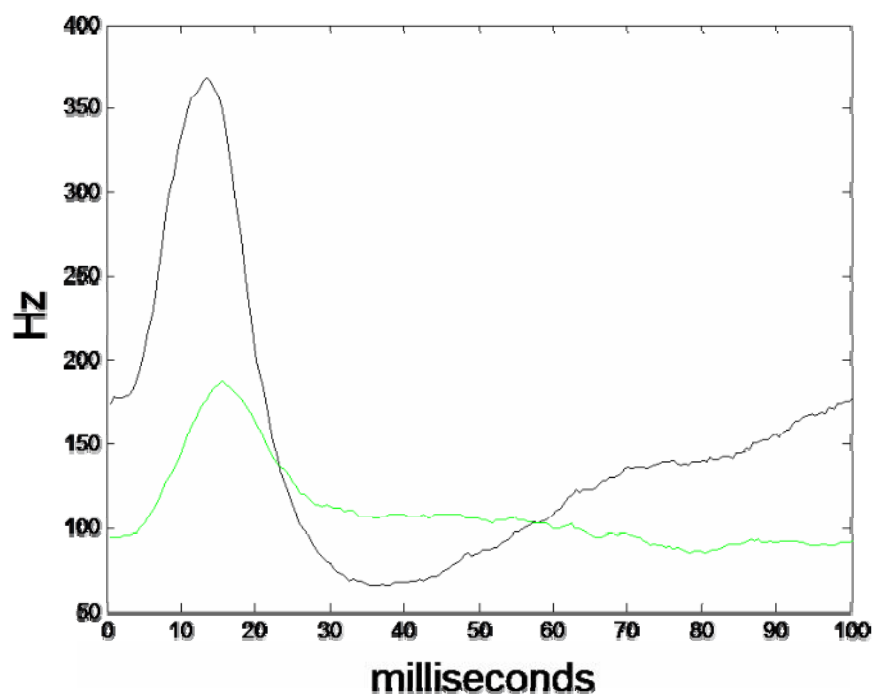

Figure 3.4 Example PSTH. Calculated around each click, from one of the cortical channels in Figure **3.1 – e** and **3.1 - f** . PSTH for all clicks presented at 10 Hz. The latency to peak in the activated state is 13 ms and in the activated is 15 ms. Black line corresponds to the activated state and green line to the inactivated.

133

Figure 3.5 Peak response evolution. a - Mean evoked response after each click in the series, with spontaneous activity subtracted. Error bars represent standard deviation. Black lines correspond to the activated state and green lines to the inactivated. **b** - Peak response evolution for a 10 Hz train. Mean of all channels belonging to each rat. Each color corresponds to a different rat.
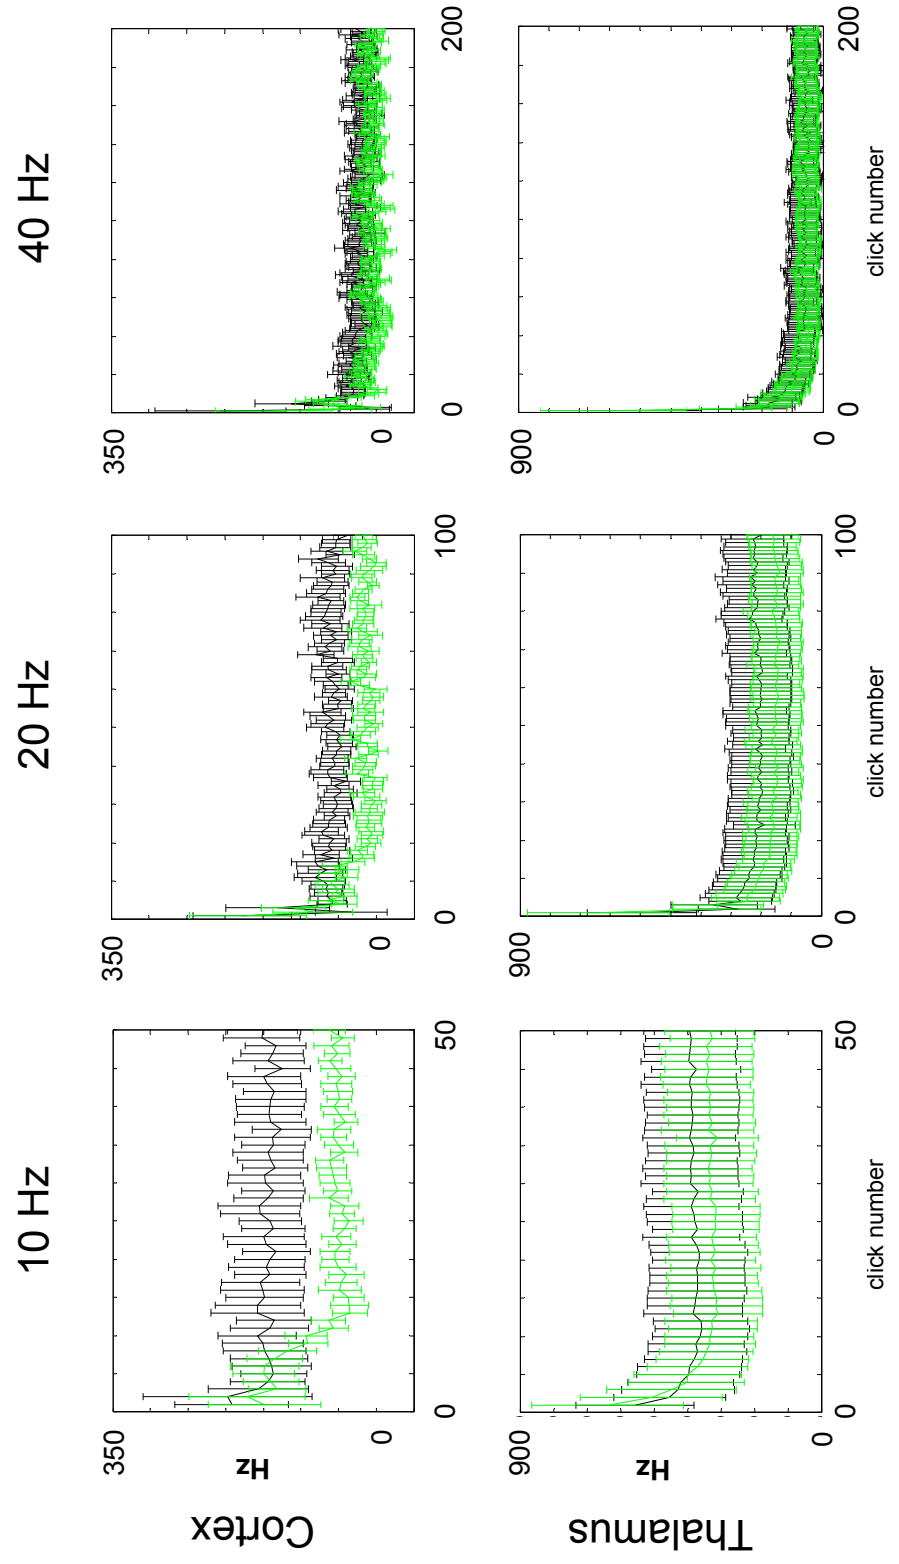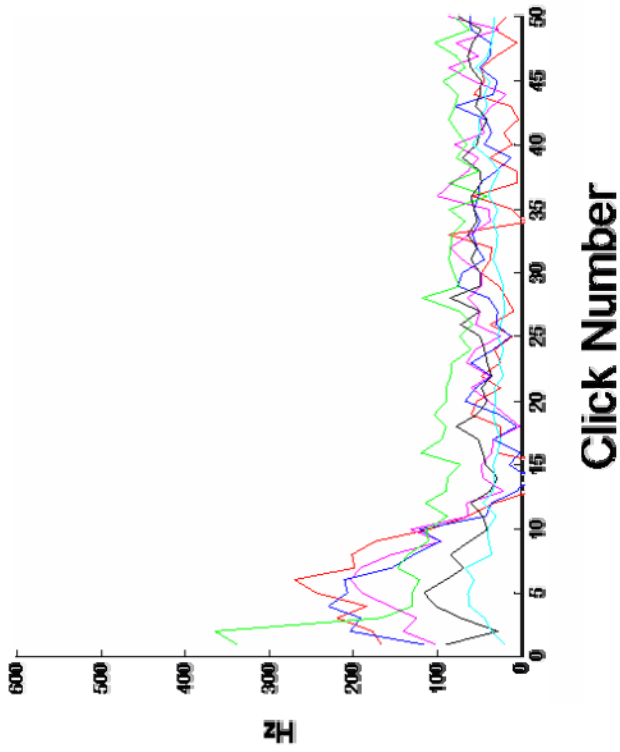
a -

activated

inactivated

Click Number

Click Number

Hz

Hz

b -

Figure 3.6 Mean evoked responses at different frequencies. Spontaneous activity subtracted. M is the ratio of the mean evoked responses in the inactivated vs. activated states.

3.7 ROA Calculation. Example of a thalamic channel responding to a 20 hz train of clicks during the steady state. Left panels - PSTH and standard deviations in each state:, inactivated state in green, activated in black, corresponding surrogates in red. Middle panels -Distributions at peak depict the probability of measuring a given number of spikes in the 15 ms surrounding the peak response. Right panel – Corresponding ROC-curves: the ROC-Area is the area under those curves

PSTH

Distributions at Peak

ROC-Curves
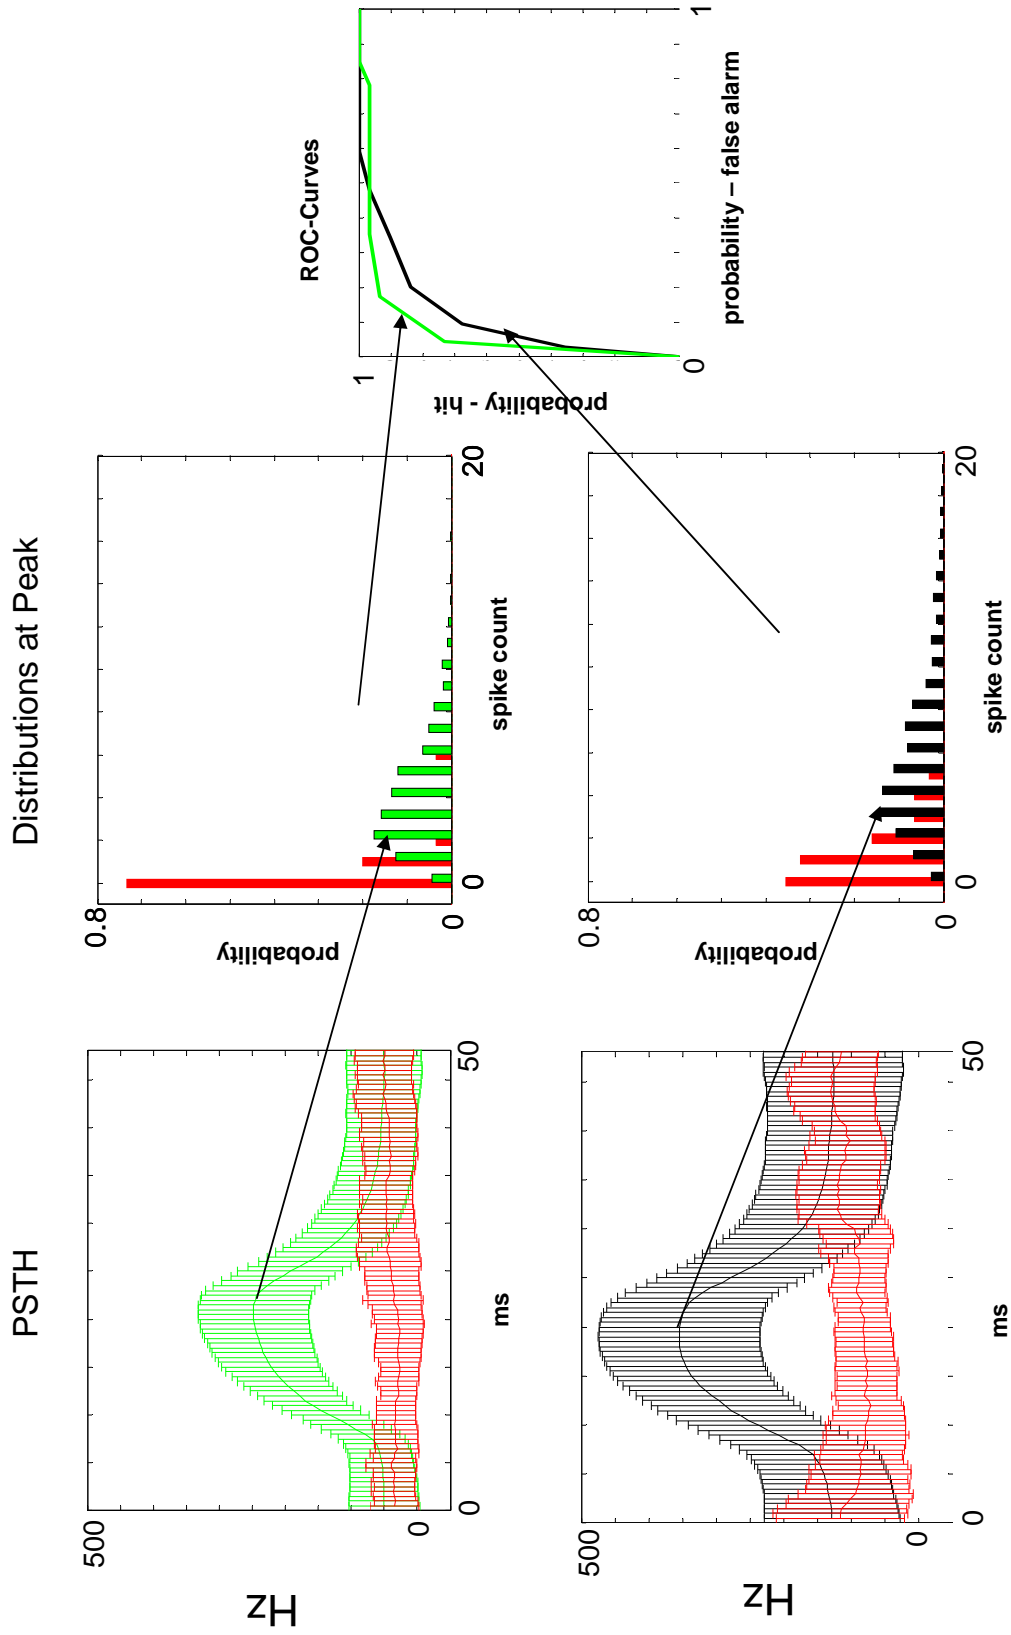
3.8 Discriminability. ROC-Areas for single clicks and steady state at different rates. A value of 0.5 means that no information about the stimulus can be extracted from the neural signal, a value of 1 means that the presence of a click can be unequivocally inferred from the neural record. (See text for definition of M.)

Figure 3.9 Frequency Locking. a -Vector Strength during the sustained period at different frequencies. A value of 0 means no phase locking. A value of 1 means that the neurons fire each time with exactly the same delay. M is the ratio of the mean of the vector strengths in the inactivated vs. activated states. **a –** Vector strengths during the sustained period. **b –** Vector strengths during the first five clicks.
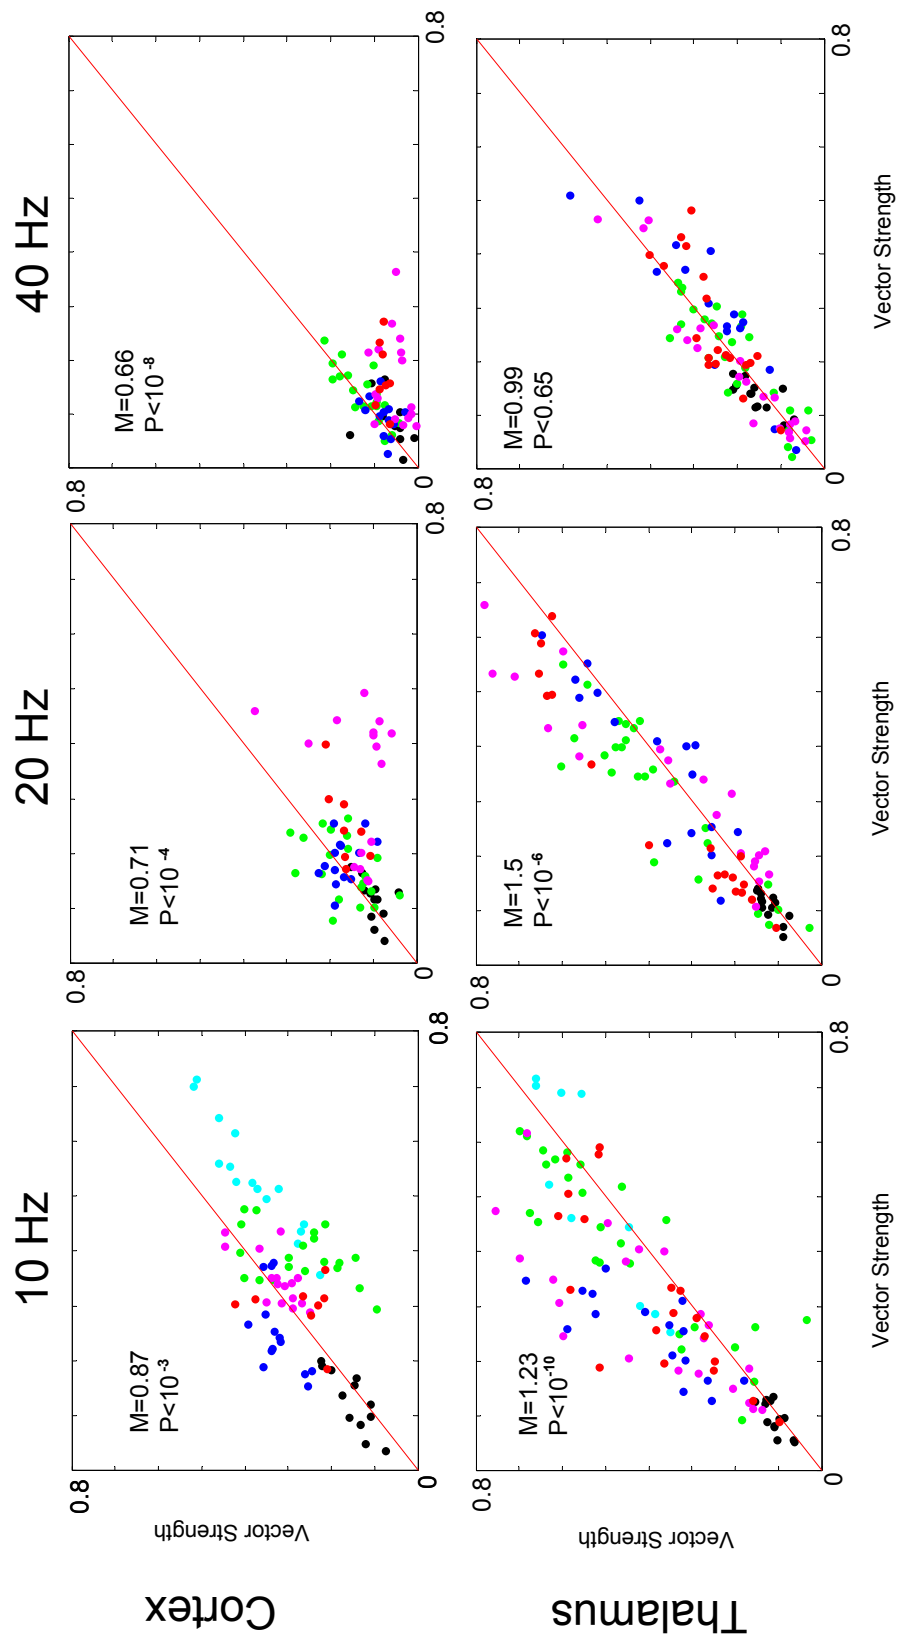
Figure 3.10 Reliability. Across Trials Mean Correlation (ATMC) during the different click trains. A value of 1 means that the neural responses are always the same in response to the stimulus, a value 0 means that the neural activity holds no relationship with the stimulus. M is the mean ATMC ratio of the inactivated vs. activated states.
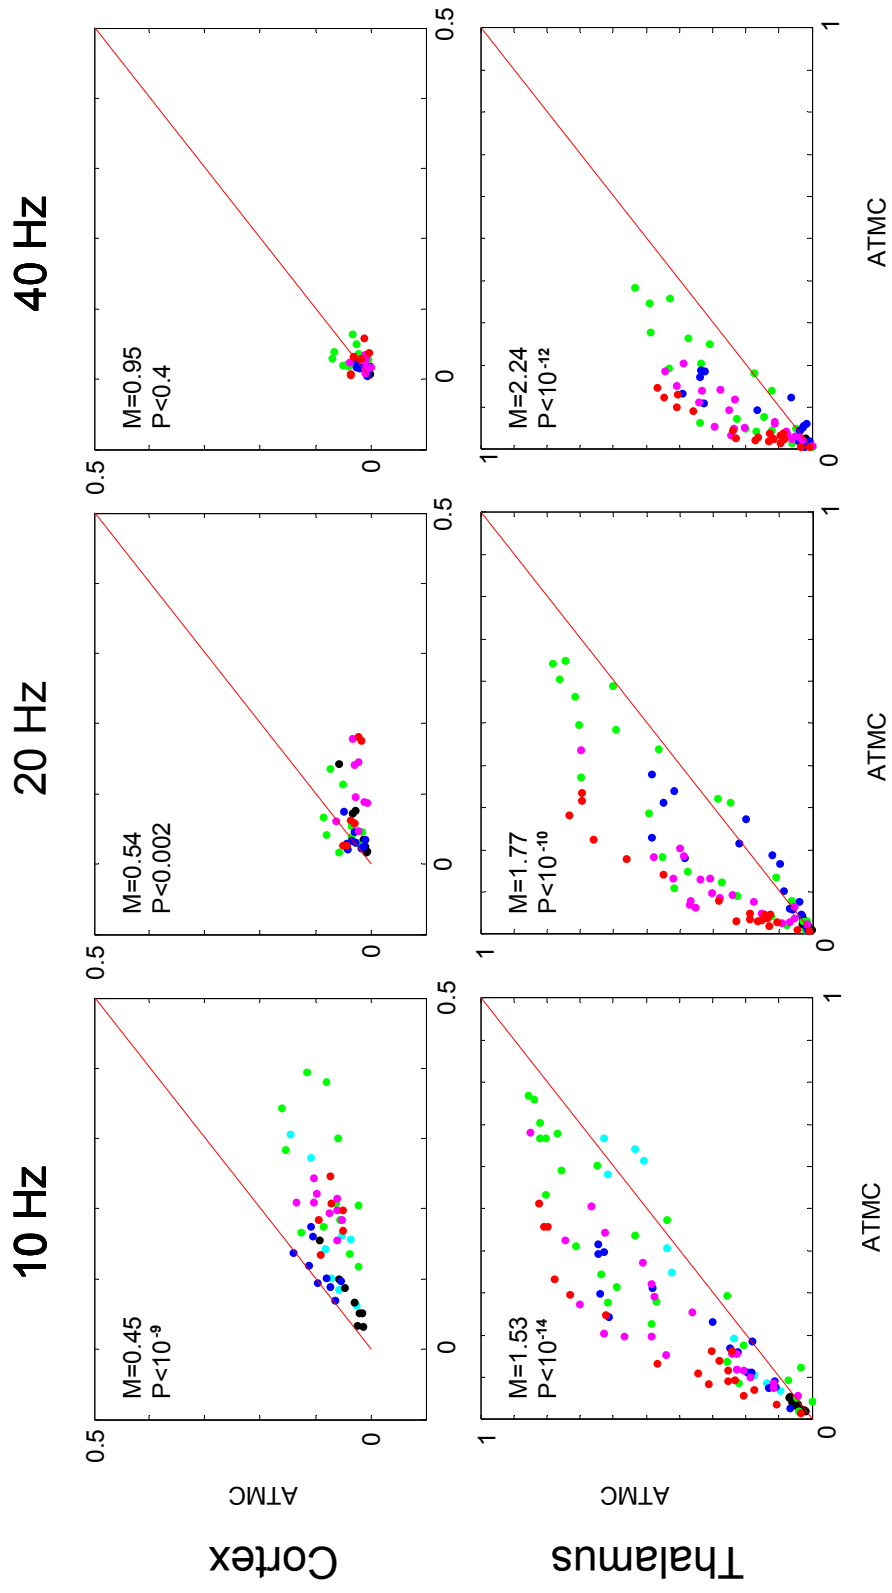
**Figure 3.11 Reliability after rate equalization.** Same as figure 3.10 but after rate equalization

Figure 3.12 Thalamic recording site. **a-**Example recording shank traverses the dorsal and ventral portions of MG. **b-**Vector Strength corresponding to this shank. Because the exact location of the tip cannot be established and because of the shrinkage of the tissue, the correspondence is approximate.

Figure 3.13 Vector Strengths in the steady state, only in the channels
putatively belonging to the ventral part of the Medial Geniculate.

Figure 3.14 Cortical Recording Sites. The location corresponds to area TE1. Left figure corresponds to rat plotted in blue, right to rat plotted in red. The penetrations correspond to area TE1.

Chapter 4

Influence of Mean Neuronal Firing Rate on Measures of Rate-Correlation

Introduction

Given the large number of neurons in the brain and the complexity of the synaptic connectivity, the total number of possible neuronal signals in the form of spike trains is astronomical, spurring the need for neuroscientists to find neuronal measures of similarity in order to understand how neurons in the brain classify specific information in the environment.

If the brain is to classify stimuli, we should expect that each time a given stimulus is presented, the resultant instances of brain activity to recurrences of that stimulus will be more similar to each other than to the brain activity elicited by a different stimulus. Therefore, measures of similarity are fundamental to understanding neuronal representation. The spike count correlation is a measure of similarity between two spike trains and, therefore, it has been frequently used to analyze neural signals (Schreiber, et al. 2003).

It has recently been shown in neurons in neo-cortex, that in many cases the correlation between two spike trains increases with their mean firing rate (De la Rocha et al. 2007), raising the possibility that some of the effects that have been measured using this technique might be solely due to differences in overall neuronal activity.

Recent work (Goard and Dan 2009; Mitchel et al. 2009) has taken this effect into consideration and controlled for it by applying a technique called rate equalization, in which the rates of two groups of spike trains are matched by randomly removing spikes from one of the groups. Other work, by Cohen and Maunsell (2009), controlled for firing rate differences by comparing correlations for matched rate groups, showing that the effect of their manipulation on between-cell correlations was present at each rate band. More recently, Ecker et al. (2010) controlled for the effect of rate on correlations by comparing simulated Poisson spike trains created to have similar statistics to those of their experimental data.

Although this recent body of work demonstrates that there is growing concern in the field regarding the need to control for basic differences in firing rate when considering correlations between spike trains, to our knowledge no general theory relating rate and correlations has yet been written. We think that this relationship can be framed into the more general neural problem of why "more is better". For instance, why is plasticity for relevant stimuli usually associated with an increase in firing rate, and why is there often an expansion of the cortical area devoted to a relevant task, and so forth. We propose that a train of spikes is an intrinsically noisy signal, and that the higher the firing rate, the cleaner the signal becomes. These statements will be appropriately formalized in the following pages.

In order to make the subject more intuitive we will begin with some examples and will subsequently analyze the mathematical problem more rigorously and in more general terms.

Examples:

In analyzing the relationship between the correlation of two spike trains and their rates, it will be useful to start with a simpler problem.

We will first analyze the correlation between a spike train and the underlying process that modulates its rate, using as example an inhomogeneous Poisson process:

Let's assume that there is a spike train which is a realization of an inhomogeneous Poisson process, where F_i is the expected value of the number of spikes in a time bin centered on time i and N_i is the actual number of spikes in that time bin. In what follows we shall call the series F a Generator and the series N a Spike Count Series (See index of terms at the end of this chapter).

The mean number of spikes in this train will be the mean of F. Through multiplying F by a factor, we will increase the rate of the spike train without changing its temporal profile.

In our example the Generator is proportional to a sine wave with 0 minimum, cycling through 20 steps. For each time bin the Generator will take the form:

$$F_i \propto Sin(i * \frac{2 * \pi}{20}) + 1$$

In figure 1, examples are provided of Spike Count Series (N) for various rates. The data have been converted into Z-scores (zero- meaned and normalized) for visualization. It can be observed that as the rate increases the Spike Count Series more closely resembles the Generator (F). In the top panel the correlation between each spike train and the Generator is depicted.

Remembering that correlation can be thought of as a measure of similarity, this example highlights the notion that as the rate tends to infinity the correlation between N and F tends to 1.

Now we can extend this example to the correlation between two spike trains. Let's first rename the variables so that the Generator we just described is called F1 and the Spike Count Series is N1, and create another process defined by Generator F2 and Spike Count Series N2, in which F2 is equal to F1 but with a phase lag of 1/3 of a cycle. The correlation of two sine waves with such phase lag is CC(F1,F2)=0.5

In figure 2, both processes are depicted. As the rates grow, N1 resembles F1 and N2 resembles F2, therefore N1 tends to resemble N2 as much as F1 resembles F2. This will be expressed as the correlation between N1 and N2 CC(N1,N2) tending to the correlation CC(F1,F2) as the rate tends to infinity.

In the following section we will find explicit formulas for these examples and set the stage to analyze the more complex Integrate and Fire model. In order to do so we more clearly explicate the concept of a Generator that we will be discussing, a concept that is similar but more general than that presented in our example.

Formulation:

Concept of a Generator:

A Generator is a function of time and is the expected value of the Spike Count in a bin centered at a particular time. In our Poisson case this is arbitrarily defined (that's how we define our inhomogeneous process). If we can lock the activity to a stimulus that we can repeat many times we can estimate the firing probability as a function of the time since the stimulus started, so we could choose the perievent-stimulus-time-histogram as our Generator. We can also estimate the firing probability each time a time-varying stimulus assumes a certain value, or the input current to a cell assumes a certain value. The important thing is that the average value of the Spike Count <N>, over all the times in which the Generator assumes a value K, is also K. In more formal terms:

$$\langle N_i | F_i = K \rangle = K \quad \text{for all K. (1)}$$

Where <.> is the mean over time and < . | A> is the mean over all times in which the condition A is fulfilled.

Let's now create a new variable Ei which is the difference between Fi and Ni.

Ei=Ni-Fi   (2)

This variable will satisfy:

<Ei|Fi=K> = <Ni-Fi|Fi=K> = <Ni | Fi=K> - <Fi | Fi=K> = K-K = 0   for all K.

(3)

Ei is the difference between the actual spike count and the expected value of that spike count, and therefore can be considered as an error around the expected value.

We will define the correlation between series A and B as:

$$CC(A, B) = \frac{<(A-<A>)*(B-<B>)>}{\sqrt{<(A-<A>)^2>*<(B-<B>)^2>}} = \frac{Cov(A, B)}{\sqrt{Var(A)*Var(B)}} = A^z * B^z$$

Where Cov(.,.) is the covariance and Var(.) is the variance, and the superscript z means that the series has been transformed to z-scores (their mean was subtracted and the result normalized).

Let's now write the correlation between F and N in terms of the new variable E.

$$CC = \frac{Cov(F, F + E)}{\sqrt{Var(F + E)*Var(F)}}$$

$$CC = \frac{<(F+E-<F>)*(F-<F>)>}{\sqrt{<(F+E-<F>)^2>*<(F-<F>)^2>}}$$

Distributing and eliminating the terms <F*E>, which according to formula (3) have 0 mean, we find that:

$$CC = \sqrt{\frac{Var(F)}{Var(F)+ <E^2>}} = \frac{1}{\sqrt{1 + \frac{<E^2>}{Var(F)}}}$$

Since <E>=0, the term <E²> is the variance of the error Var(E).

Recapitulating:

$$CC(F, F + E) = \sqrt{\frac{1}{1 + \frac{Var(E)}{Var(F)}}} = \beta_F \qquad (4)$$

Formula (4) is an equality and a definition of β_F.

CC(F,F+E) will tend to 0 if Var(E) >> Var(F), and will tend to 1 if Var(E) << Var(F)

It must be remembered that in the general case, Var(E) depends on F. So the question of how the correlation changes with the rate can be reduced to how the error depends on the Generator.

If, as the rate increases, the Generator grows more rapidly than the Error, then the correlation between the Spike Count Series and the Generator will approach 1.

In the following section we will apply this formula to the example in Figure 1.

Poisson Example (Figure 1):

For future calculations it might be useful to write the variance of the error more explicitly in terms of the Generator:

$$<E^2> = \frac{1}{T}\sum_{K} S_k * <E^2|F_i = K> \qquad (5)$$

Where T is the total length of the series F and Sk is the number of times

K is repeated. This is just re-writing the temporal mean, but grouping it

according to the different values of F.

Now, in a homogeneous Poisson process, the variance of the spike

count is equal to the mean.

$$< (N- < N >)^2 >=< N >$$

This formula is strictly true when referring to expected value. If we refer

to the actual mean across many realizations it becomes true only as the

number of realizations grows to infinity. Using equation 5 we can generalize

the formula above for an inhomogeneous Poisson processes.

If the size of the bins is small enough that we can consider the

probability of firing to be constant, we can use the above property of Poisson

processes to write:

$$< E^2|F_i = K >=< (N - F_i)^2|F_i = K >=< (N - K)^2 >= K$$

We could do this because for a fixed K we can consider the process as

homogeneous. As mentioned earlier, this is strictly true if the number of

realizations (times in which Fi=K) is infinite. In what follows we will work under

the approximation that the time series considered is very long and therefore

the number of realizations very large.

Combining the latter formula with equation 5 we obtain:

$$< E^2 >= \frac{1}{T}\sum_{K} S_k * K =< F >$$

This can be replaced into equation 4 to obtain:

$$CC = \frac{1}{\sqrt{1 + \frac{<F>}{Var(F)}}} \quad (6)$$

Without loosing generality we can parameterize F in such a way that we transform F into λ*F, where <F>=1. The parameter λ becomes then the mean number of spikes per bin. Changing λ is equivalent to multiplying the Generator but keeping its shape, and therefore is consistent with a gain change.

We can re-write formula 6 as:

$$CC = \frac{1}{\sqrt{1 + \frac{1}{\lambda * Var(F)}}} \quad (7)$$

We can see here that as λ tends to infinity, CC tends to 1, and as λ tends to 0, CC tends to 0.

As a matter of fact, for small values of λ we can ignore the number 1 in the denominator, therefore in this regime CC grows as the square root of λ, with a slope set by the variance of the Generator.

Formula 7 corresponds to the first example we presented, in which rate changes were achieved by multiplying the probability of spiking per unit of time.

Of course, this dependence relies on the way in which we increase the rate. If instead of re-scaling the Generator we add a constant to it, the variance of the Generator will remain the same, but the mean will increase, leading to a decrease of CC (see formula 6). This type of manipulation is not a gain change.

In more concrete terms, scaling the Generator would be akin to making a cell more excitable, in the sense that the probability of spiking will increase, both for inputs and for noise. Further discussion of this is forthcoming.

To exemplify formula 7 we use the same Generators as in figure 1, but with a longer series (50,000 bins). We use a longer time series because, as we said, our approximations are strictly valid when size of the series tends to infinity.

Figure 3 shows the correlations as a function of rate for a simulation (black dots) and the theoretical formula (blue line).

Multiple Spike Trains:

If we remember that the correlation is a measure of similarity, what we have shown so far is enough to understand intuitively what happens when we compare the correlations of two spike trains.

Let's say we have a spike train N1 associated with a Generator F1 and a spike train N2 associated with F2. If N1 becomes very similar to F1 and N2 to F2, then N1 will resemble N2 nearly as much as F1 resembles F2.

This intuitive idea can be formalized by understanding the relation between correlation and Euclidean distance, because they have a one to one relationship. The advantage of thinking in terms of a distance is that we can use the triangular inequality to show that the distance the spike trains D(N1,N2) is constrained. Showing the relationship between correlation and distance will in turn be useful in putting the problem of correlations in a broader context.

We have seen that the correlation is the inner product of two normalized vectors. So, the correlation can be understood as the internal product of the two vectors in a unitary sphere of dimension T. They are in a unitary sphere because the normalization process makes them have a norm of 1. The internal product uniquely defines the angle between two vectors. In this case, the two vectors are in a sphere, thus defining the Euclidean distance, which we will symbolize as $D(.,.)$. This shows that there is a 1 to 1 relationship between the correlation and the Euclidean distance of vectors of norm 1 and mean 0. If the correlation is 1, the angle between the vectors is 0 and therefore their distance is 0.

So we can constrain the distance between the Spike Count Series in terms of their distance to their Generators and between the Generators themselves:

$$D(F1,F2)-D(N1, F1)-D(N2,F2) \leq D(N1,N2) \leq +D(F1,F2) + D(N1,F1)+D(N2,F2)$$

If the correlations between the Spike Count Series and their respective Generators tend to 1, $D(N1,F1)$ and $D(N2,F2)$ tends to 0, therefore $D(N1,N2)$ will tend to $D(F1,F2)$.

All this was useful in order to get an intuitive understanding of the problem and to understand the relationship between correlations and Euclidean distance, but we can also derive an explicit formula equivalent to (4) for the correlation between two trains of spikes.

Lets start by defining the series $E1_i = N1_i - F1_i$ and $E2_i = N2_i - F2_i$

Then we can write:

$$CC = \frac{Cov(N1, N2)}{\sqrt{Var(N2) * Var(N2)}} = \frac{Cov(F1 + E1, F2 + E2)}{\sqrt{Var(F1 + E1) * Var(F2 + E2)}}$$

If we distribute the numerator, eliminate the terms with 0 value, and do some simple substitutions we can decompose this formula in 4 terms, one associated with the correlation between the Generators, another with the correlation between the errors, and two more with the correlation between each Generator and the error of the other.

$$CC(N1, N2) = \beta_{F1} * \beta_{F2} * CC(F1, F2) + \beta_{E1} * \beta_{E2} * CC(E1, E2) +$$

$$+\beta_{F1} * \beta_{E2} * CC(F1, E2) + \beta_{F2} * \beta_{E1} * CC(F2, E1) \quad (8)$$

Where β_{F1} and β_{F2} are defined as in equation 2:

$$\beta_{Fi} = \sqrt{\frac{1}{1 + \frac{Var(Ei)}{Var(Fi)}}}$$

And β_{E1} and β_{E2} are defined as:

$$\beta_{Ei} = \sqrt{\frac{1}{1 + \frac{Var(Fi)}{Var(Ei)}}}$$

It can be easily shown that β_E and β_F satisfy:

$$\beta_{Fi}^2 + \beta_{Ei}^2 = 1$$

Formula (8) decomposes the total correlation into different components involving the correlations between terms that we are familiar with, corrected by terms associated with the fact that the actual spike count is not exactly the expected value; therefore, coefficients β will be called correctors.

The term that involves the correlation between the Generators will be called the Corrected Generator Correlation (CGC), the term that involves the correlation of the Errors will be called Corrected Error Correlation (CEC), and the cross terms will be called Corrected Cross-Terms Correlation (CCC).

As the error is the deviation of the Spike Count from its expected value, the term CC(E1,E2) can be interpreted as what is usually called noise correlation.

We can see that if for both series Var(F) >> Var(E), then all the terms in (8) will vanish except for the one that involves the Generators, as β_F will also tend to 1, we can see that the correlation between the Spike Count Series tends towards the correlation between their Generators. This confirms what was deduced using the triangular property.

Poisson Example (Figure 2):

If we have two independent, inhomogeneous Poisson processes, then the error and the cross terms in (8) vanish. This is because the realization (and therefore the error) of one process does not depend on the realization of the other. It then follows that the correlation between the two Spike Count Series becomes:

$$CC(N1, N2) = \beta_{F1} * \beta_{F2} * CC(F1, F2) \quad (9)$$

This formula says that if we change the Generators but keep Var(F) and Var(E) constant, the total correlation will depend linearly on the correlation of the Generators. This property is not specific to the processes being Poisson,

but rather to them being independent. The independence of the processes allowed elimination of the terms associated with the errors.

Using (9) and re-scaling as in the previous section we obtain:

$$CC(N1, N2) = \sqrt{\frac{1}{1 + \frac{1}{\lambda_1 * Var(F1)}}} * \sqrt{\frac{1}{1 + \frac{1}{\lambda_2 * Var(F2)}}} * CC(F1, F2)$$ (10)

Here, λ1 and λ2 are the mean firing per bin of each spike train.

Figure 4.4 shows examples of correlations of spike trains derived from Generators with different inter-Generator correlations. In these examples the mean rate of both spike trains has been set to be the same.

We can see in (10) that for very small values of λ1 and λ2, the terms that have these variables in the denominator tend to be much bigger than 1, so 1 can be ignored, therefore CC is proportional to the geometric mean of λ1 and λ2:

$$CC \propto \sqrt{\lambda_1 * \lambda_2}$$ (11)

Long Series Approximation:

In the derivation of our Poisson example we made the approximation that the spike trains were infinitely long; the fact that in this case the correlation is deterministic is not obvious a priori. For shorter trains, the approximation that the variance of the error equals the mean of the Generator will not hold and their relationship will be more stochastic. Therefore, we should expect certain randomness in the values of the correlations.

Short Bin Size Approximation:

When analyzing the Poisson process we made the approximation that the bin size is smaller than the typical variation time of the rate. Given that the bin size is arbitrary we can always choose it in such a way that this is true, although the resultant mean number of spikes per bin in this case might be very low.

Conclusion:

We have found explicit formulas that relate the correlation between two spike trains, the Generators of these trains and the error of estimation. These formulas can be useful in deriving explicit formulas relating the correlation between two spike trains with their mean firing rates. In the particular case of a Poisson process, the error term depends on the Generator, so we find a formula that directly relates the correlation with the Generator. If, in this case, we manipulate the rate by changing the gain (multiplying the Generator), we find a formula that directly relates the correlation with the rate. We show that a gain change has the effect of making the correlation between two Spike Count Series approach the correlation between their expected values (Generators) as the rate goes to infinity, and to go to 0 as the rate goes to 0.

The formulas that we found can be easily adapted for more general cases.

Discussion:

These examples suggest that the correlation will increase with the rate as long as the rate increase is caused by a change in gain. However, in a

more realistic scenario we do not have the option of just changing the gain, and differences in cellular properties will lead to more complex transformations of the neuron's output. In order to understand how correlations change under these more general transformations, we apply our analysis to Integrate and Fire model neurons. This will also help us show how to build a Generator in a situation in which we do not explicitly know the probability of firing per unit of time, as we do in the case of our Poisson example.

Integrate and Fire:

In this section we use formula (4) to analyze the behavior of the correlation between the Spike Count Series of two leaky Integrate and Fire neurons receiving a common input.

We model each neuron according to the following formula:

$$\tau_m * \frac{dV}{dT} = -V + A + I(t) \tag{12}$$

In this formula, the decay time of the neuron's voltage is $\tau_{m=}$ 10 ms, V is the voltage of the neuron, A is its asymptotic value of the voltage in the absence of any input, and I is a time dependent input (the product of the input current and the membrane resistance, in units of voltage).

Each time the voltage reaches threshold Vth=20 mV, the potential of the neuron is reset to 0 mV. The time dependent input (I) is drawn randomly at each integration time (0.5 ms), according to a normal distribution with 0 mean.

An increase of the asymptotic value A or the standard deviation of the input

Std(I) will lead to an increase in the mean firing rate of the neuron. In the

following simulations we will independently vary one of these two parameters.

Following the procedure of the previous section, we will first analyze the

correlation between a neuron's output and its input, and then between two

neurons that have a correlated input.

A Neuron and its Input:

When dealing with this more realistic model, we don't have the luxury of

choosing the probability of firing per unit of time, but we can estimate the

mean number of spikes in a given bin if we know the mean of the input for that

bin. Based on this estimation we can then build a Generator. We will show

how to do this with an example.

Figure 4.5 (top) shows an example of spike train generated through

equation (12), the number of spikes in this train was counted in rolling bins of

50 ms (middle). The average of the input in each bin was also calculated

(bottom), we will call this average value J.

Based on the middle and bottom graphs we can define the Transfer

Function as the mean value of the spike count, each time the averaged input

(J) takes a given value K:

$T(K) = < N | J = K >$ (13)

Where $< . | C >$ stands for the mean over time when the condition C is

satisfied.

In other words, the transfer function is the expected value of the spike count as a function of the averaged input (J). In a general case, the transfer function is a non-linear function of the input.

Figure 4.6 shows transfer functions for two different values of parameter A, these two values produce two different mean rates. We can see that the transfer function corresponding to the highest rate is very linear, whereas the function corresponding to the lowest rate is non-linear. This is due to the fact that in the latter case small currents are not strong enough to drive the neuron over threshold.

Given that we know the expected value of the spike count as a function of the input, and we know the input as a function of time, we can know the expected value of the spike count as a function of time:

F(t)=T(J(t))

In this sense this function can be considered a Generator.

Figure 4.7 shows two examples of Generators and the corresponding spike counts, for two neurons with different mean rates, corresponding to figure (4.5)

We can also see that F satisfies equation 1:

<Ni|Fi=K> = K   for any K.

So the formulas derived in the previous section are valid.

In the previous section we showed that:

$$CC(F, F + E) = \sqrt{\frac{1}{1 + \frac{Var(E)}{Var(F)}}} = \beta_F$$

It can also be shown that if G is any function that satisfies:

<E|G=K>=0

Then the following is true:

$$CC(N, G) = CC(F, F + E) * CC(F, G) = \beta_F * CC(F, G) \quad \text{(14)}$$

This is particularly true if G is the mean current in a bin (J), because

<E|J=K>=<E|F=F(K)>=0   for any K

This is shown using formula (3). The condition of the current assuming a value K is the same as the Generator assuming a value F(K).

If we remember that F is the transfer function evaluated on the current J, then, from equation (14), we have:

$$CC(N, J) = \beta_F * CC(T(J), J) = CC(T(J), N) * CC(T(J), J) \quad \text{(15)}$$

Formula (15) is very useful, because it shows that we can decompose the correlation between a neuron's output and its input into the product of two terms. The correlation between a realization (N) and its expected value T(J) and between the expected value T(J) and the input J. The former factor arises from the fact that our estimation is not deterministic. The fact that CC(T(J),J) is not necessarily 1, arises from the fact that in a general case the transfer function is non linear, in the linear range this correlation is actually 1.

To illustrate (15) we show some examples.

Figure 4.8 shows the results of our simulations, the two examples differ only in the binning size. We can see that the correlation between the spike

train and the current is the product of the correlation between the spike train and its Generator and between the Generator and the current.

In the figure we also show how Var(F) and Var(E) depend on the rate, these values suffice to define the correlation between the train and its Generator. It can also be seen that for a large range of low rates, the system behaves as a Poisson process, in the sense that the variance of the error is equal to the rate.

These examples might mislead the reader to think that at high rates the Spike Count Series always converges to the Generator. However, that is not true for all possible parameters. Figure 4.9 shows examples in which this correlation reaches a maximum and then decays. The only difference between these examples and the previous figure is that in this case the input has a smaller standard deviation.

Two Neurons:

In the case of an Integrate and Fire model, as in the Poisson, the above result aids in understanding the behavior of the correlation between two neurons. If the correlation between each neuron and its input tends to 1, then the correlation between the two neurons will tend to that of the inputs.

Besides the fact that in this more realistic model we cannot simply increase the gain, another difference arises in comparison with our Poisson example. In our example, the two Poisson processes were created independently and therefore there was no reason for the errors to be

correlated, which allowed us to eliminate the error terms from equation 8. But this is not true in general and therefore we should consider these terms when analyzing the Integrate and Fire model.

Following, we will show how the different components and their corrections play out in the total correlation.

Figure 4.10 shows the different components of the total correlation (according to equation 8). Figure 4.10 (a) corresponds to figure 4.8, time bin= 50 ms. We saw in 4.8 that as the rate increases the correlation between each cell's Generator and the input approaches 1, which can explain why the correlation between the two Generators approaches the correlation between the inputs.

We can see that the correlation between errors (CEC) is a non-negligible component of the total correlation. The CEC is particularly important at low rates, a regime in which its corrector factor β_E is close to 1 and β_F is close to 0. As the rate increases the component associated with the Generators (CGC) tends to predominate. As β_F approaches 1; the total correlation approaches the correlation of the Generators, which in turn approaches the correlation of the inputs.

A similar pattern is seen in figure 4.10 (b), corresponding to the cell in figure 4.9. This case is a little different in that the tendency of the total correlation saturates and is then reversed, which can be explained by the behavior of β_F (see figure 4.9).

In these examples, the correlation between one neuron's Generator and the other neuron's error does not play an important role.

Correlation between the Errors:

Even if the mean input over a bin is a good indicator of the neuron's output, the fine details of that input are also important. When we built our Generators we only took the mean input into account, but the neurons share the input's variations as well, this might explain the fact that the error terms are correlated. If this is true we would expect that reducing these fine variations might eliminate the correlation between the errors. In order to test this we smoothed the common input with a square window of 200 ms. This manipulation greatly reduced the correlation between the errors, as can be seen in figure 4.11. This has the effect of reducing the total correlation at low rates, a regime in which β_E is close to 1.

Conclusion:

We showed how to adapt the formulations derived in the previous section to situations in which we do not know, a priori, the probability of spiking per unit of time. We did so by estimating this probability based on a parameter of the input. Because we are constrained by a mechanistic, although simplified model, a few differences arise in comparison to the Poisson case.

In the Poisson case, the correlation between processes is diminished by their stochastic nature. In the Integrate and Fire case, the correlation is diminished by the combination of two factors, the stochasticity of the response and the non-linearity of the transfer function.

As we do not have the option of just multiplying the probability of firing per unit of time, we increase the rate by augmenting the DC input to the neuron. This modification does increase the probability of firing, but the increase is mediated by a non-linear function rather than just a multiplicative factor. The non-linearity has the effect of de-correlating a neuron's input from its output. This, in turn, diminishes the correlation between the outputs of two neurons.

Augmenting the DC input also has the effect of changing the variance of the Generator, which can be considered a change in gain, because it means that the cell responds to the inputs with bigger amplitude swings.

As in the Poisson case, we find that the key variable in understanding the changes in correlation is the ratio between the variance of the errors and the variance of the Generator Var(E)/Var(F). Increasing the DC input has the effect of augmenting both variances. In general, the variance of the Generator is shown to increase faster than that of the error, although in some cases this tendency is reversed above a certain rate.

As the influence of the correlation between the errors does not vanish as the Var(E)/Var(F) goes to 0, the total correlation will not necessarily vanish in this regime. But, as in our Poisson case, we can be assured that as Var(E)/Var(F) grows the total correlation will tend to that of the Generators.

Discussion:

Deterministic and Stochastic:

As we have said, the de-correlation between cells is affected by two processes, one originating in the non-linearity of the transfer function and the other in the stochasticity of the process. Since the equations that we use are deterministic, the stochastic effect arises in the level of description.

When we use the average input over a bin, we are ignoring the fine details of that input, as well as the initial voltage of the cell, therefore we can no longer evaluate the exact output but only its mean and variance. If the size of the bin is much smaller than the typical variation time of the input, then the problem becomes deterministic, (a given constant input value univocally determines the output), In this case the only factor that produces a de-correlation is the non-linearity. This is confirmed by simulations in which we smooth the input over windows that are much larger than the bins (not shown).

In a real experimental situation we do not have access to the intrinsic timing properties of the noise and we are constrained in the size of bins that we can use, therefore we cannot know exactly how the deterministic and stochastic processes affect the correlation, but through bin size.

Other Generators:

In the Poisson case we built our Generators from what we forcefully set to be the probability of spiking per unit of time, whereas in the Integrate and Fire case we estimated this probability based on the averaged input and the corresponding transfer function.

It is easy to see how we can build Generators using other variables. For example, we can know the probability of firing at a given time after stimulus onset, or each time a sensory stimulus assumes a certain value (for example dots moving in the right direction).

The Generators are not univocally defined by the experimental condition. For example, we selected the average input to estimate the probability of firing, but we could have used the maximum input in a given bin, or the standard deviation. The predictive value of the variable chosen will be different in each case. As the error measures the difference between the actual spike count and the prediction, we should expect the influence of the error term to be larger in the case of a poor predictor.

Sometimes correlations are measured in a context in which there is no variable that predicts the firing of the neurons, in such a context is not possible to build a Generator. However, we can think of the common input to a pair of neurons as a hidden Generator.

Correlated Errors:

An important difference between the Poisson and the Integrate and Fire case is that the Poisson processes of our example were independent. This allowed us to throw away the error terms in equation (8). In that context, all the correlation between the trains derived from their relationship to the Generators.

However, we could have used Poisson processes with correlated errors. We could generate them by adding a third Poisson realization to the two spike trains under consideration. In this case both cells' fluctuations around the

expected value would be coupled. The addition of a third Poisson process would keep the two spike trains Poisson, because the probability of spiking per unit of time doesn't depend on the history.

If we did so the result for the correlation between a neuron and its Generator would have been the same. However, the result for pairs of neurons would have been radically different, as we have seen that the contribution of the error terms becomes predominant as the rate goes to 0.

This points to the fact that even if the behavior of individual real neurons can be well approximated with Poisson statistics, we can not directly extrapolate this into sets of neurons.

Changes in Gain:

We have seen in the Poisson example that multiplying the probability of spiking per unit of time will make the correlation between spike trains approach that of the Generators. In a more general case an increase of the variance of the Generator Var(F) will have the same effect. This remains consistent with a change in gain, in the sense that the neuron will be more deeply modulated by the input, such that it will have bigger swings in response to it.

There are two important realistic cases in which we can consider there is a change in gain:

1. If there are many neurons modulated by the same signal, and they are independent, adding neurons to a multi-unit pool will have the effect of multiplying the probability of firing per unit of time.

2. If the probability of firing per unit of time varies slowly compared to the size of the bin, an increase in the size of the bin will be equivalent to a multiplication in the probability of firing per bin.

The latter can explain why, typically, the correlation goes to 0 when the size of the bin goes to 0.

Discreteness, Continuity, and Correlation:

Whereas the stimuli that modulate neuronal activity are generally continuous (analogical), neuronal activity is intrinsically discrete. Thus, the question can be asked: How well can a continuous signal be represented by a discrete signal? The correlation between a spike train and its Generator addresses this question, by measuring the similarity between the two as the firing rate increases.

A feature that is clearly depicted in figures 4.1 and 4.7 is likely to clarify why an increased rate might lead to a higher correlation. If we see the traces corresponding to very low rates, the values of the spike counts are discretized, because there are usually only a few spikes per bin. As the mean rate increases, the range of values that the count can take also increases. Therefore, the z-scored series approaches a continuum. The correlation between two z-scored functions is 1 only if the two functions are equal, therefore the correlation between the series and a continuous Generator can only be 1 if the number of spikes tends to infinity.

This intuition is advanced by Cohen and Maunsell (2009) when analyzing the effect of attention in correlations between cells:

"Attention tends to increase firing rates, which makes the distributions of spike counts more Gaussian and less discretized, leading to a predicted increase in correlation".

Rate Equalization:

Why Correct?

Making the assumption that a method is needed to control for mean firing rate differences when performing correlation analyses implies that the influence of one parameter over the other is an artifact. However, the complementary view can be taken, in which increasing the rate is considered to be a mechanism to boost correlations.

So maybe the question should be: When is the increase in correlation only a product of an increase in rate?

Our formulations help us put this question in more concrete mathematical terms, given that we have seen that an increase in mean firing rate is often associated with an increase in gain, which increases the correlation. However, the correlation can also be increased by a reduction of the error, a parameter that is not associated with an increase in mean firing rate.

Two clear scenarios emerge, one in which the correlations increase because the gain increases and another in which they increase because the error decreases. These are two fundamentally different mechanisms that can be employed to change correlations and it is worthwhile to discriminate between the two.

If an experimental condition is such that we know the input to the cells; we can build the Generators and measure their variance as well as that of the error. In this case, the two above scenarios will be clearly discriminated, so rate equalization becomes unnecessary.

If we know the Generators we can also know to what degree the increase in rate is associated with an increase in their variance (gain), or just to an added constant. This difference is important because, as exemplified in our Poisson example, an added constant might actually lead to a reduction in correlations.

In the experimental case in which the Generators cannot be established or are unknown, there is no way to know how much of the increase in rate is due to an increase in gain versus how much of the increase is due to an added constant. We will use the Poisson model, under these two extreme situations, to analyze the effect of rate equalization.

Our results indicate that rate equalization can compensate for the influence of mean firing rate on the correlation, when the differences in mean rate are solely due to differences in gain. In cases in which the rate increases by an added constant (an effect that we showed diminishes the correlation),

rate equalization has the effect of diminishing the correlation even further. Although this can be regarded as a failure of the technique we will see that under some conditions it can prove useful.

Effect of Rate Equalization on the Generators:

When performing rate-equalization, the spikes are removed with equal probability. Thus, the probability of removing a spike from a given bin is proportional to the number of spikes in that bin, so the expected value of the number of spikes in each bin is divided by a constant factor.

More formally, it applies the following transformation:

$$F_{Ei} = F_i * \frac{R_F}{R_T} \qquad (16)$$

Where F_E is the equalized Generator, R_F is the mean rate of the process F and R_T is the target rate (the rate that needs to be matched).

So we can see that rate equalization has the property of dividing the generator by a constant factor.

Effect of Rate Equalization on a Poisson process:

In a general case, the effect of rate-equalization on the error terms is not trivial. Applying rate equalization across a Poisson process has the property of creating a new Poisson process (because the probability of firing still doesn't depend on the history). Given that the process remains Poisson, if we know the Generator, we also know the error.

So, rate equalization transforms a Poisson process into another Poisson process with similar modulation but a smaller mean firing rate.

Effect of Rate Equalization on the Correlation between Poisson Processes:

Now we can analyze the effect of rate-equalization. For simplicity we will consider the correlation between two processes with the same rate (RF) and variance Var(F) and the same target rate (RT). Then, evaluating (9) and using (16) we have:

$$CC_E(N1, N2) = \frac{1}{1 + \frac{R_F^2}{R_T * Var(F)}} * CC(F1, F2)$$

(17)

The subscript E denotes that the correlation has been rate-equalized. Here, we have used the property that multiplying functions by a constant value does not change their correlation.

First, we note that when RF=RT, equation (17) is the real correlation of two processes with rate RT. Now, if we change the gain of the processes by multiplying F1 and F2 by any constant, the value of the rate-equalized correlation in (17) will stay the same. This is true because both R^2 and Var(F) grow as the square of that constant. Therefore, we see that when we perform rate-equalization, the equalized correlation stays at the value of the real correlation of the target.

This demonstrates that rate equalization controls for changes in rate that originate from changes in gain. This is an analytical result that confirms the simulations by Mitchel et al. (2009).

There is a substantially different effect of rate-equalization when we increase the rate by adding a constant to the Generators. Here, the term RF in

(17) will increase, but the term Var(F) will stay the same, because adding a constant does not change the variance. Therefore, as we increase the rate, the rate-equalized correlation will diminish, becoming lower than the real correlation of the target.

In more realistic cases, as the Integrate and Fire model exemplifies, a change in the physiological properties of the cell might lead both to a change in gain and an added constant.

We have seen that in these two extreme scenarios (increased gain and an added constant), rate equalization either reduces the correlation of the spike trains to the value corresponding to the target rate, or to a value even lower than that. In this sense, it can be said to control for the worst case scenario.

Let's think now we have two pairs of spike trains that are being compared, A+ and A-, where A+ has a higher rate. If the correlation of the A- pair is higher than the A+, then the influence of the rate is not an issue. If A+ has a higher correlation than A-, rate equalization can have two outcomes:

1. After rate equalization, the correlation of A+ becomes lower than that of A-. As we don't know if the increased rate was due to an added constant, we don't know if we compensated or over-compensated. Therefore, the rate equalization technique does not provide any information.

2. After rate equalization the correlation of A+ stays higher, then we know that the increased correlation is not due to an increased rate, because even if we over-compensated the correlation stays higher.

More Realistic Models:

We derived these conclusions based on a simplified model; two fundamental differences will arise when treating more realistic cases.

1. The effect of rate equalization on the error won't be as straight forward as in the Poisson case, and the result will depend on the case considered.

2. In more realistic cases we cannot ignore the contribution of the error-terms to the total correlation. Given that rate equalization removes spikes from two spike trains independently, it is likely to reduce these correlated error variations.

Conclusion:

We analyzed the effect of rate-equalization in our simple, independent Poisson example. In this case we show that rate-equalization compensates for the influence of rate in correlation when the differences in rate are derived from a re-scaling of the Generators. If the Generators are transformed in other ways, the technique can over-compensate. Given that in this general case it over-compensates, the results following rate equalization is meaningful only when it agrees with the results of the non-equalized correlation, in which case we can conclude that a higher correlation is not caused by a higher rate. If rate equalization reverses the order of the correlations, then it adds no information.

Conclusion:

Recent work has used rate correlations as a measure of neural coding. As it has been shown that in many cases, the mean firing rate of the individual neurons has an influence on rate correlations. This has lead to a concern in the field regarding the extent to which the differences observed in this measure are not solely due to differences in mean rate.

The fact that an increased mean rate leads to an increased rate correlation is not a universal truth, and therefore it can not be proven. Our work tries to identify the conditions under which this is true. Our results suggest that the correlation will increase with the firing rate as long as this increased firing is a reflection of a greater sensibility of the neurons to their inputs.

We have re-written the correlation equations in terms of familiar concepts, the expected value of firing within a time bin (Generator) and the deviations from this value. This allowed us to find analytical formulas in the simple case of independent Poisson processes..

We also show how to apply these formulas to the more realistic case of leaky Integrate and Fire neurons. In this case, we first find that the correlation between a cell's spiking output and its time-averaged input is the product of two terms. One term is related to the non-linearity of the transfer function and the other to the stochasticity of the response, (the fact that an input doesn't

univocally define the output). This duality is carried on when analyzing the correlation between two neurons.

The effect of the non-linearity of the transfer function in the correlation was already noted by De la Rocha et al (2007), our findings extend that result into situations in which the neuronal firing is stochastic, as it is to be expected in a realistic experimental situation.

In a general case, an increase in rate does not necessarily lead to an increase in correlation. We find that the fundamental variable dominating the correlation is the relationship between the variance of the Generator and the variance of the error around it. If the former is much bigger than the latter then the total correlation approaches the correlation between the Generators.

Neither the Generator nor the error directly depends on the rate. This dependence is a function of the problem that is being treated. However, if the likelihood of a neuron firing towards its inputs increases, this will lead to both an increase in its rate and an increase in the variance of the Generator, so the two variables are often times coupled.

We also analyze the usefulness of rate equalization to control for differences in correlation that are due to differences in mean rate.

As we found an analytic formula relating the correlation and the Generators in the Poisson case, we use this example to analyze the technique of rate equalization. In this simplified case, we find that if the differences in rate are solely due to differences in gain, (a multiplicative re-scaling of the Generator), this technique is useful to control for the influence of rate in

correlations, a similar case is analyzed by Mitchel et al (2009). However, this technique over-compensates if the rate increases in more general ways. So, it is useful only if it matches the results of the non-equalized correlation, in which case it is confirmatory. If its application reverses the order of the correlations, then it adds no information.

If the experimental conditions allow establishing Generators, then rate equalization becomes unnecessary: We have seen that the behavior of the correlations is dominated by the relationship between the gain and the error, the former of these factors being frequently associated with mean rate. If we want to know to what degree an increase in correlations is mediated by an increase in gain, as opposed to a decrease in error, the best thing to do is to directly measure the changes in gain and in error, which we are able to do if we know the Generators.

Index of Terms

N – Spike Time Series - Number of spikes in a given bin.

F – Generator – Expected number of spikes in a given bin.

E – Error – Difference between the expected number of spikes in a bin and the actual number of spikes in that bin.

N1, F1, E1 – Spike Time Series, Generator and Error of one of the spike trains being correlated.

N2, F2, E2 – Spike Time Series, Generator and Error of the other of the spike trains being correlated.

Formula 8:

$$CC(N1, N2) = \beta_{F1} * \beta_{F2} * CC(F1, F2) + \beta_{E1} * \beta_{E2} * CC(E1, E2) +$$

$$+ \beta_{F1} * \beta_{E2} * CC(F1, E2) + \beta_{F2} * \beta_{E1} * CC(F2, E1)$$

Where

$$\beta_{Fi} = \sqrt{\frac{1}{1 + \frac{Var(Ei)}{Var(Fi)}}} \qquad \text{and} \qquad \beta_{Ei} = \sqrt{\frac{1}{1 + \frac{Var(Fi)}{Var(Ei)}}}$$

are the corrector terms for the Generator and the Error, respectively.

CGC = β_{F1}*β_{F2}*CC(F1,F2)   Corrected Generator Correlation.

CEC=β_{E1}*β_{E2}*CC(E1,E2)   Corrected Error Correlation.

CCC=β_{E1}*β_{F2}*CC(E1,F2) + β_{E2}*β_{F1}*CC(F1,E2)   Corrected Cross-Terms Correlation.
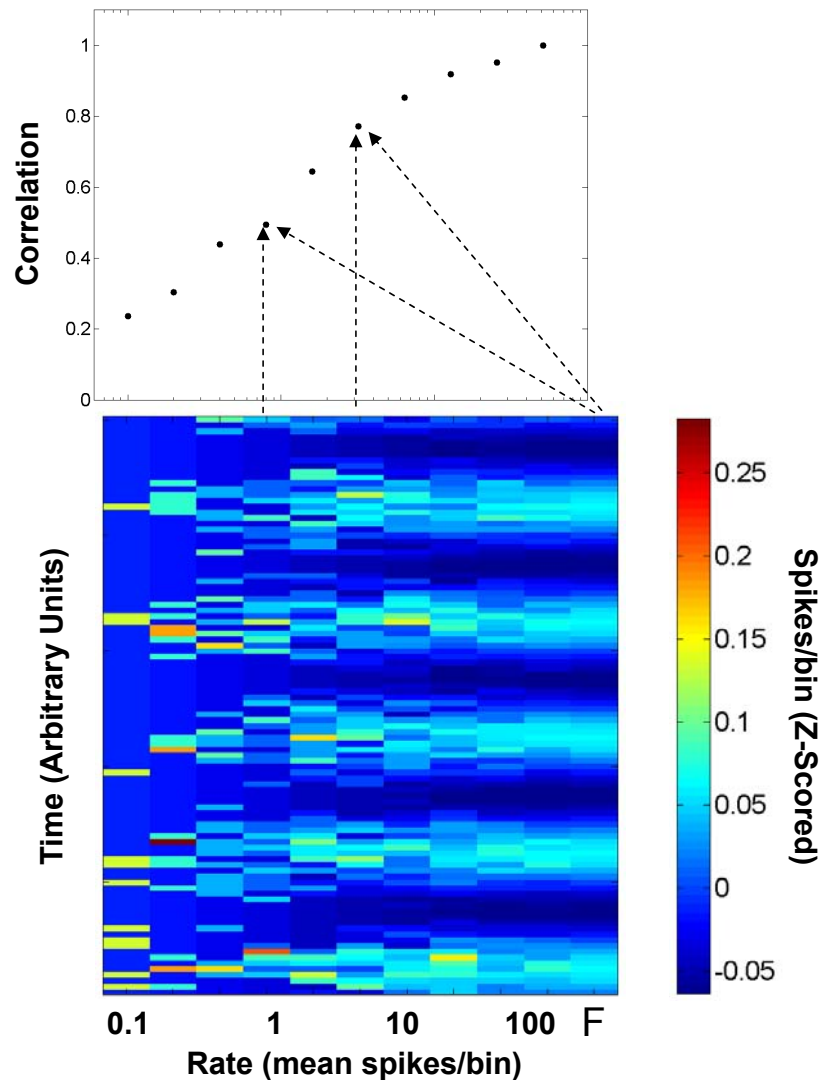
Figure 4.1 Example of correlations between Spike Count Series and a Generator. Bottom -Example of spike trains with different mean rates. The spikes were generated from an inhomegeneous Poisson process (Generator). By multiplying this Generator we obtain spike trains with the same modulation but increased mean rates. The Generator (F) is depicted in the right-most column. Each column represents a z-scored Spike Count Series. It can be seen that as the rate increases, the activity resembles the Generator. Top – Correlation between each series and the Generator, as the rate increases this correlation approaches 1.
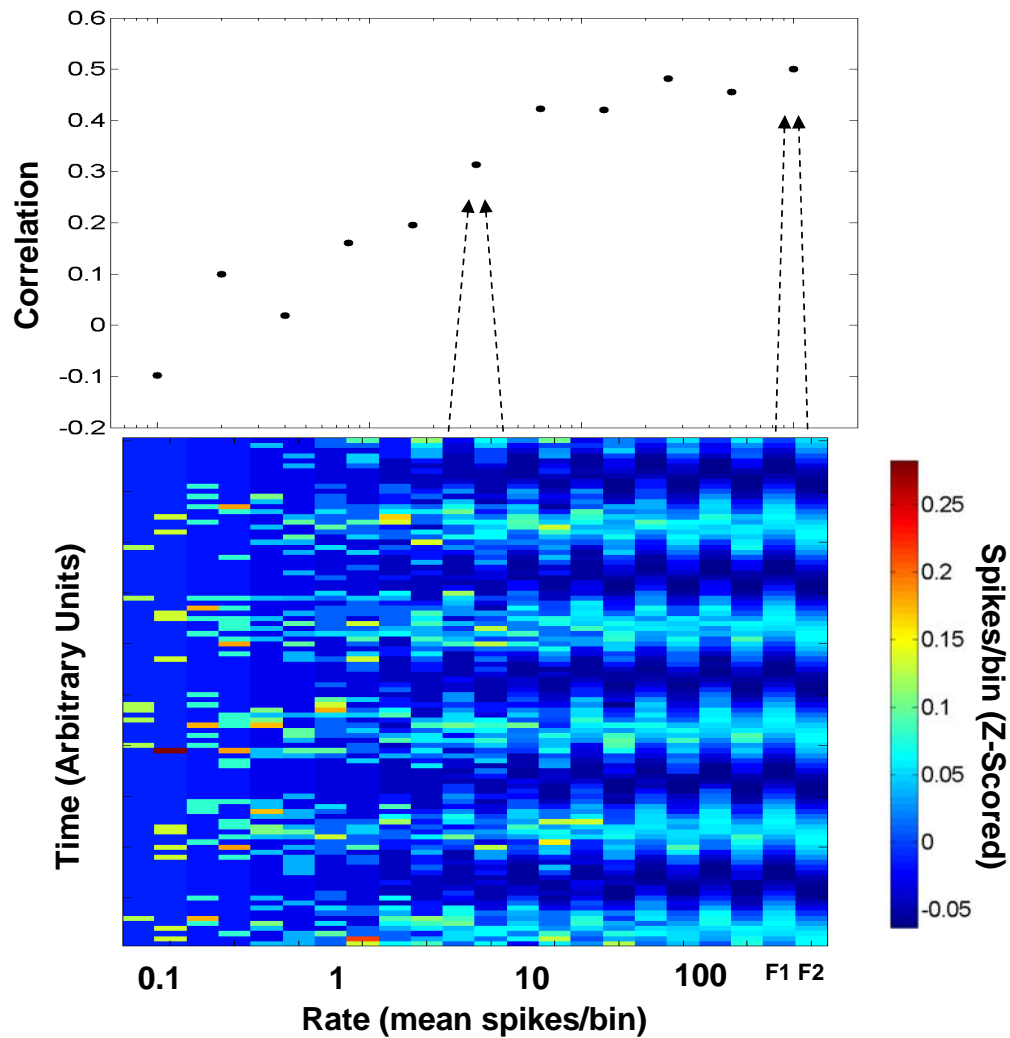
Figure 4.2 Correlation Between Pairs of Spike Trains. For this example we used two Generators (F1 and F2) with a correlation of 0.5. Bottom- The Spike Count Series deriving from each Generator are intercalated. We can see that as the rate increases, each series more highly resembles its own Generator, and therefore they tend to resemble each other as much as the Generators do. This is reflected as their correlation (Top) tending to the correlation between the Generators.

Figure 4.3 Correlation between a Spike Count Series and its Generator.
The Generator is similar to that depicted in figure 1, but the time series
consists of 50000 bins. In this regime the simulated values closely resemble
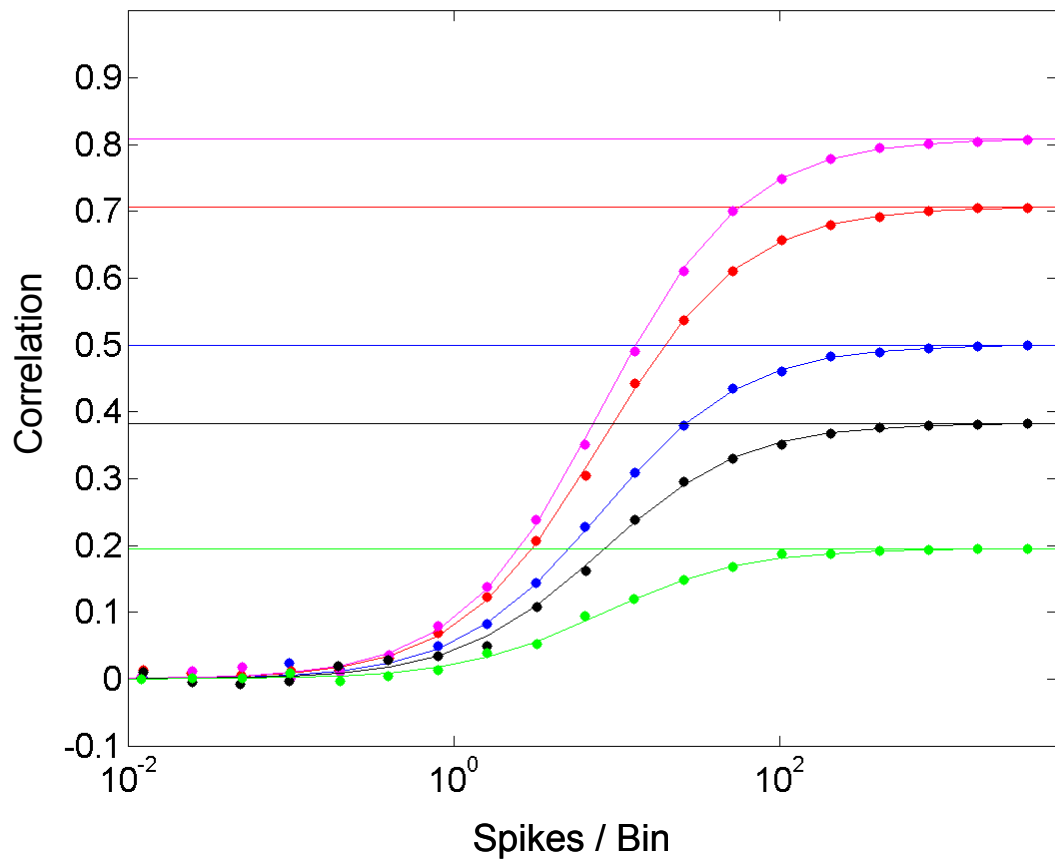the theoretical values.

Figure 4.4 Correlation as a function of rate for pairs of spike trains. The horizontal lines represent the correlation between the Generators, the dots are the simulation results, and the lines the theoretical model. In this case the mean rate of both trains has been set to be the same. It can be seen that as the rate increases the correlation between trains tends to that of the Generators.

Figure 4.5   Building a Generator. Top – Spike train generated with an integrate and fire model. Middle – Spike Count Series, sum of the number of spikes in rolling bins of 50 ms. Bottom – Input averaged in the same rolling bins as the Spike Count.
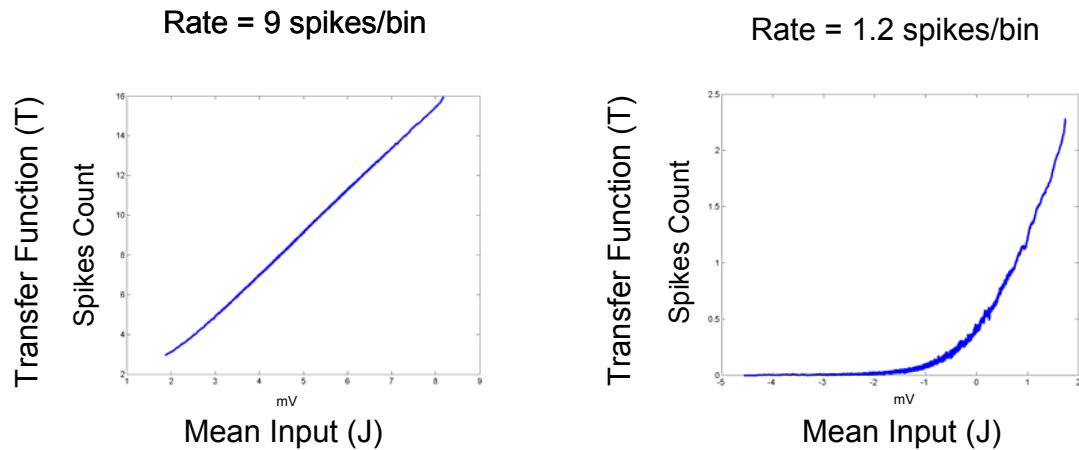
Rate = 9 spikes/bin

Rate = 1.2 spikes/bin

Figure 4.6   Transfer Functions. The transfer function represents the expected number of spikes in a bin as a function of the average of the input within that bin (J). The left plot corresponds to a bigger DC input and therefore a higher rate. In this regime the transfer function is very linear. The transfer function for a lower rate (right) is markedly non-linear.
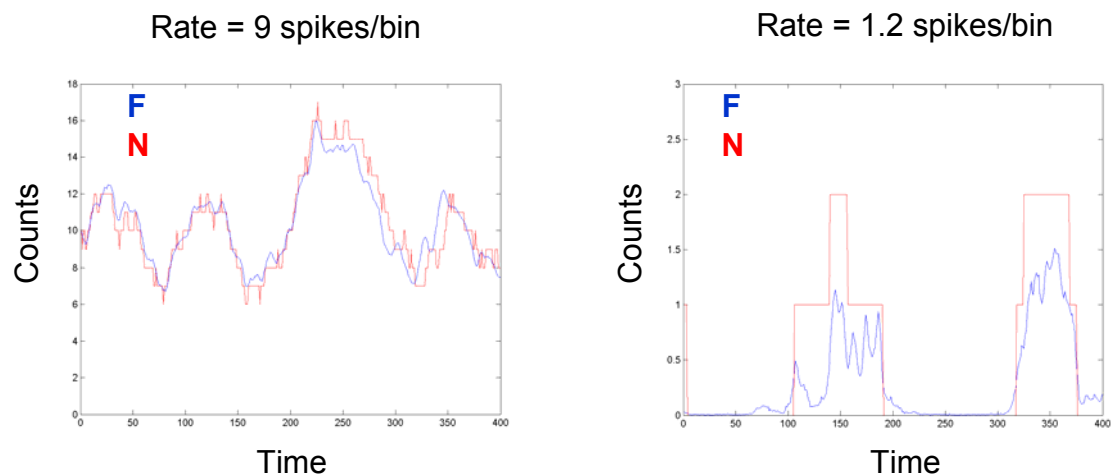
Rate = 9 spikes/bin

Rate = 1.2 spikes/bin

Figure 4.7   Examples of Spike Count Series and their corresponding Generators for high (left) and low (right) rate parameters. The rate difference in these examples arises only in differences in DC input. It can be seen that in the high rate example the difference between the Spike Count Series (N) and the Generator (F) is much smaller than the amplitude oscillations of the Generator; this is not the case in the low rate example. If the spike count can assume only a few values it will show a stair-like shape that will not fit a continuous function (as the Generator); if the rate is high the signal can assume a more continues shape.

Figure 4.8   Correlation between a Spike Count Series and the mean input current (J). β_F approximates very well the correlation between the Spike Count Series and the Generator CC(F,N). We can see that the correlation between the N and the input (J) can be decomposed as the product of the correlation between the Generator and the input current (which arises in their non-linear relationship) and the correlation between the Spike Count Series and the Generator (β_F), which arises from their non-deterministic relationship. Circles represent theoretical values. The bottom row shows the Variance of the Error and the Generator, these two variables alone determine β_F. We can see that at low rates Var(E) is well fitted by the identity function (blue line), which indicates that in this regime the Spike Count Series behaves in a Poisson manner. In these examples the variance of the input Var(I) was set at 40 mV.

Figure 4.9   Same as in 4.8 but different parameters. In this case Var(I)=20 mV. Here the correlation between the Spike Train Series and the Generator saturates at a value lower than 1 and then decays.

Figure 4.10   Different components of the correlation. The correlation between two neuron's Spike Count Series (red dots) can be decomposed into many terms according to formula 4.8. Panel (a) corresponds to the simulation in figure 4.7 (Bin = 50 ms), panel (b) to 4.8 (Bin = 50 ms). In both cases we used a correlation between the inputs CC(I1,I2) =0 .5. See text for detailed interpretation.

$CC(E,F)=CC(E1,F2)+CC(E2,F1)$

$CGC = \beta F1*\beta F2*CC(F1,F2)$

$CEC=\beta E1*\beta E2*CC(E1,E2)$

$CCC=\beta E1*\beta F2*CC(E1,F2) + \beta E2*\beta F1*CC(F1,E2)$

Figure 4.11   Reduced Error Correlation. The common input was low pass filtered. This manipulation had the effect of eliminating the correlation between the errors. Parameters are similar to 4.9 a. In this case the correlation between the raw inputs (I) was adjusted to produce a correlation of the averaged inputs of CC(J1,J2)=0.5
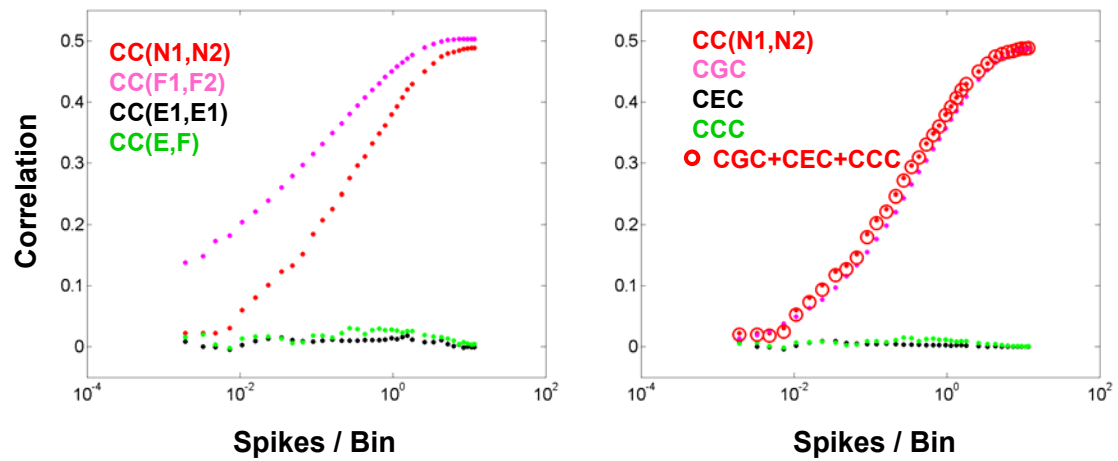
General Conclusions

Our experiments explore two aspects of the problem of sensory processing of time-dependent stimuli, particularly with relation to the cholinergic system, yielding results that are apparently contradictory. Whereas in Chapters 1 and 2 we show that local cholinergic input to the sensory cortices is not essential for normal sensory processing of time modulated signals (only for learning), the results of chapter 3 suggest that the activational state dramatically changes cortical responses to thalamic inputs. This contradiction is of course apparent and we will attempt an explanation.

One possible explanation is that, even if cortical activation is strongly associated with cholinergic activity, ACh does not play an important role in shaping the sensory responses that are observed during the activated state. This is supported by the finding by Goard and Dan that the muscarinic antagonist atropine can locally disrupt cortical activation but leave the reliability of individual neural responses relatively unaffected. This explanation is not entirely satisfactory since other groups have shown a direct influence of ACh in the shaping of sensory responses (Disney, et al. 2007; Kawai, et al. 2007).

A more plausible explanation is that Ach depleted cortices can still be activated by other means. Even though it has been shown that applying atropine locally to the cortex can abolish activation (Goard and Dan, 2009;

Metherate, Cox and Ashe 1992), ACh depletion studies have generally failed

to show an effect on cortical activation (Nilsson, et al. 1992; Kaure, et al. 2008).

The mechanism by which ACh depleted brains can retain or regain activation

is currently unknown (Semba; Dringenberg; Metherate; Dickson; personal

communications) and is a pressing question in the field. One possibility is that

the atropine dosage used in the past has been too large (Metherate, personal

communication); another possibility is that there is compensation from other

neuro-modulatory systems. Since we find that depleted rats have learning, but

not perception, deficits, it is reasonable to hypothesize that the effect of ACh

as a modulator of responses can be compensated for but its effect as a

facilitator of plasticity cannot be.

References

Aitkin LM, Prain SM (1974) Mediate Geniculate-Body – Unit responses in awake Cat. Journal of Neurophysiology 37:512-521.

Alheid GF, Heimer L (1988) New Perspectives in basal forebrain organization of special relevance for neuropsychiatric disorders – The striatopallidal,, amygdaloid, and corticopetal components of substantia innominata. Neuroscience 27:1-39.

Alkire MT, Hudetz AG, Tononi G (2008) Consciousness and Anesthesia. Science 322:876-880.

Baker MA (1971) Spontaneous and evoked activity of neurones in somatosensory thalamus of waking cat. Journal of Physiology-London 217:359-&.

Bakin JS, Weinberger NM (1996) Induction of a physiological memory in the cerebral cortex by stimulation of the nucleus basalis. Proceedings of the National Academy of Sciences of the United States of America 93:11219-11224.

Bao SW, Chan VT, Zhang LI, Merzenich MM (2003) Suppression of cortical representation through backward conditioning. Proceedings of the National Academy of Sciences of the United States of America 100:1405-1408.

Baskerville KA, Schweitzer JB, Herron P (1997) Effects of cholinergic depletion on experience-dependent plasticity in the cortex of the rat. Neuroscience 80:1159-1169.

Batsel HL (1960) Electroencephalographic synchronization and desynchronization in the chronic cerveau isole of the dog. Electroencephalography and Clinical Neurophysiology 12:421-430.

Baxter MG, Chiba AA (1999) Cognitive functions of the basal forebrain. Current Opinion in Neurobiology 9:178-183.

Belardetti F, Borgia R, Mancia M (1977) Prosencephalic Mechanisms of ECOG desynchronization in cerveau-isole cats. Electroencephalography and Clinical Neurophysiology 42:213-225.

Benasich AA, Gou Z, Choudhury N, Harris KD (2008) Early cognitive and language skills are linked to resting frontal gamma power across the first 3 years. Behavioural Brain Research 195:215-222.

Berg RW, Friedman B, Schroeder LF, Kleinfeld D (2005) Activation of nucleus basalis facilitates cortical control of a brain stem motor program. Journal of Neurophysiology 94:699-711.

Blair HT, Cho JW, Sharp PE (1999) The anterior thalamic head-direction signal is abolished by bilateral but not unilateral lesions of the lateral mammillary nucleus. Journal of Neuroscience 19:6673-6683.

Bonnet V, Bremer F (1937) The effect of potassium, calcium and acetylcholine on the provoked, spontaneous and electrical activities of the cerebral cortex. Comptes Rendus Des Seances De La Societe De Biologie Et De Ses Filiales 126:1271-1275.

Bruno RM, Sakmann B (2006) Cortex is driven by weak but synchronously active thalamocortical synapses. Science 312:1622-1627.

Bucci DJ, Holland PC, Gallagher M (1998) Removal of cholinergic input to rat posterior parietal cortex disrupts incremental processing of conditioned stimuli. Journal of Neuroscience 18:8038-8046.

Buzsaki G, Bickford RG, Ponomareff G, Thal LJ, Mandel R, Gage FH (1988) Nucleus basalis and thalamic control of neocortical activity in the freely moving rat. Journal of Neuroscience 8:4007-4026.

Castro-Alamancos MA (2002a) Different temporal processing of sensory inputs in the rat thalamus during quiescent and information processing states in vivo. Journal of Physiology-London 539:567-578.

Castro-Alamancos MA (2002b) Properties of primary sensory (lemniscal) synapses in the ventrobasal thalamus and the relay of high-frequency sensory inputs. Journal of Neurophysiology 87:946-953.

Castro-Alamancos MA (2004a) Dynamics of sensory thalamocortical synaptic networks during information processing states. Progress in Neurobiology 74:213-247.

Castro-Alamancos MA (2004b) Absence of rapid sensory adaptation in neocortex during information processing states. Neuron 41:455-464.

Castro-Alamancos MA, Oldford E (2002) Cortical sensory suppression during arousal is due to the activity-dependent depression of thalamocortical synapses. Journal of Physiology-London 541:319-331.

Chiba AA, Bucci DJ, Holland PC, Gallagher M (1995) Basal forebrain cholinergic lesions disrupt increments but not decrements in conditioned-stimulus processing. Journal of Neuroscience 15:7315-7322.

Chiba, A. A., Bushnell, P. J., Oshiro, W. M., & Gallagher, M. (1999). Selective removal of cholinergic neurons in the basal forebrain alters cued target detection. Neuroreport, 10(14), 3119-3123.

Clement EA, Richard A, Thwaites M, Ailon J, Peters S, Dickson CT (2008) Cyclic and Sleep-Like Spontaneous Alternations of Brain State Under Urethane Anaesthesia. Plos One 3.

Coenen AML, Vendrik AJH (1972) Determination of transfer ratio of cats geniculate neurons through quasi-intracellular recordings and relation with level of alertness. Experimental Brain Research 14:227-&.

Cohen MR, Maunsell JHR (2009) Attention improves performance primarily by reducing interneuronal correlations. Nature Neuroscience 12:1594-U1148.

Conner JM, Culberson A, Packowski C, Chiba AA, Tuszynski MH (2003a) Lesions of the basal forebrain cholinergic system impair task acquisition and abolish cortical plasticity associated with motor skill learning. Neuron 38:819-829.

Conner JM, Culberson A, Packowski C, Chiba AA, Tuszynski MH (2003b) Lesions of the basal forebrain cholinergic system impair task acquisition and abolish cortical plasticity associated with motor skill learning. Neuron 38:819-829.

Davidson MC, Marrocco RT (2000) Local infusion of scopolamine into intraparietal cortex slows covert orienting in rhesus monkeys. Journal of Neurophysiology 83:1536-1549.

de la Rocha J, Doiron B, Shea-Brown E, Josic K, Reyes A (2007) Correlation between neural spike trains increases with firing rate. Nature 448:802-U806.

Diamond IT, Neff WD (1957) Ablation of temporal cortex and discrimination of auditory patterns. Journal of Neurophysiology 20:300-315.

Disney AA, Aoki C, Hawken MJ (2007) Gain modulation by nicotine in macaque V1. Neuron 56:701-713.

Dotigny F, Ben Amor A Y, Burke M, Vaucher E (2008) Neuromodulatory role of acetylcholine in visually-induced cortical activation: behavioral and neuroanatomical correlates. Neuroscience 154: 1607-1618

Dringenberg HC, Vanderwolf CH (1997) Neocortical activation: modulation by multiple pathways acting on central cholinergic and serotonergic systems. Experimental Brain Research 116:160-174.

Dringenberg HC, Vanderwolf CH (1998a) Involvement of direct and indirect pathways in electrocorticographic activation. Neuroscience and Biobehavioral Reviews 22:243-257.

Dringenberg HC, Vanderwolf CH (1998b) Involvement of direct and indirect pathways in electrocorticographic activation. Neuroscience and Biobehavioral Reviews 22:243-257.

Dringenberg HC, Olmstead MC (2003) Integrated contributions of basal forebrain and thalamus to neocortical activation elicited by pedunculopontine tegmental stimulation in urethane-anesthetized rats. Neuroscience 119:839-853.

Ecker AS, Berens P, Keliris GA, Bethge M, Logothetis NK, Tolias AS (2010) Decorrelated Neuronal Firing in Cortical Microcircuits. Science 327:584-587. Edeline JM, Manunta Y, Hennevin E (2000) Auditory thalamus neurons during sleep: Changes in frequency selectivity, threshold, and receptive field size. Journal of Neurophysiology 84:934-952.

Edeline JM, Hars B, Maho C, Hennevin E (1994) Transient and prolonged facilitation of tone-evoked responses induced by basal forebrain stimulations in the rat auditory-cortex. Experimental Brain Research 97:373-386.

Edeline JM, Dutrieux G, Manunta Y, Hennevin E (2001) Diversity of receptive field changes in auditory cortex during natural sleep. European Journal of Neuroscience 14:1865-1880.

Espana RA, Berridge CW (2006) Organization of noradrenergic efferents to arousal-related basal forebrain structures. Journal of Comparative Neurology 496:668-683.

Evarts EV (1963) Photically evoked responses in visual cortex units during sleep and waking. Journal of Neurophysiology 26:229-&.

Everitt BJ, Robbins TW (1997) Central cholinergic systems and cognition. Annual Review of Psychology 48:649-684.

Filion M, Lamarre Y, Cordeau JP (1971) Neuronal discharges of vontrolateral nucleus of thalamus during sleep and wakefulness in cat .2. Evoked activity. Experimental Brain Research 12:499-&.

Fine A, Hoyle C, MacLean CJ, Levatte TL, Baker HF, Ridley RM (1997) Learning impairments following injection of a selective cholinergic immunotoxin, ME20.4 IgG-saporin, into the basal nucleus of meynert in monkeys. Neuroscience 81:331-343

Francesconi W, Muller CM, Singer W (1988) Cholinergic mechanisms in the reticular control of transmission in the cat lateral geniculate-nucleus. Journal of Neurophysiology 59:1690-1718.

Gasbarri A, Sulli A, Pacitti C, McGaugh JL (1999) Serotonergic input to cholinergic neurons in the substantia innominata and nucleus basalis magnocellularis in the rat. Neuroscience 91:1129-1142.

Gil Z, Connors BW, Amitai Y (1997) Differential regulation of neocortical synapses by neuromodulators and activity. Neuron 19:679-686.

Goard M, Dan Y (2009a) Basal forebrain activation enhances cortical coding of natural scenes. Nature Neuroscience 12:1444-1449.

Gold PE (2003) Acetylcholine modulation of neural systems involved in learning and memory. Neurobiology of Learning and Memory 80:194-210.

Golmayo L, Nunez A, Zaborszky L (2003) Electrophysiological evidence for the existence of a posterior cortical-prefrontal-basal forebrain circuitry in modulating sensory responses in visual and somatosensory rat cortical areas. Neuroscience 119:597-609.

Grove EA (1988a) Efferent connections of the substantia innominata in the rat. Journal of Comparative Neurology 277:347-364.

Grove EA (1988b) Neural associations of the bustantia innominata in the rat – Afferent connections. Journal of Comparative Neurology 277:315-346.

Gucer G (1979) Effect of sleep upon the transmission of afferent activity in the somatic afferent system. Experimental Brain Research 34:287-298.

Guido W, Lu SM, Vaughan JW, Godwin DW, Sherman SM (1995) Receiving operating characteristic (ROC) analysis of neurons in the cats lateral geniculate-nucleus during tonic and burst response-mode. Visual Neuroscience 12:723-741.

Guillery RW, Sherman SM (2002) Thalamic relay functions and their role in corticocortical communication: Generalizations from the visual system. Neuron 33:163-175.

Hall G, Pearce JM (1982) Restoring the associability of a pre-exposed CS by a surprising event. Quarterly Journal of Experimental Psychology Section B-Comparative and Physiological Psychology 34:127-140.

Hasselmo ME (1999) Neuromodulation: acetylcholine and memory consolidation. Trends in Cognitive Sciences 3:351-359.

Herrero JL, Roberts MJ, Delicato LS, Gieselmann MA, Dayan P, Thiele A (2008) Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. Nature 454:1110-1114.

Issa EB, Wang XQ (2008) Sensory Responses during Sleep in Primate Primary and Secondary Auditory Cortex. Journal of Neuroscience 28:14467-14480.

Jasper, H. H. (1936). Cortical Excitatory State and Variability in Human Brain Rhythms. Science **83**(2150): 259-260.

JimenezCapdeville ME, Dykes RW, Myasnikov AA (1997) Differential control of cortical activity by the basal forebrain in rats: A role for both cholinergic and inhibitory influences. Journal of Comparative Neurology 381:53-67.

Juliano SL, Ma W, Eslin D (1991) Cholinergic depletion prevents expansion of topographic maps in somatosensory cortex. Proceedings of the National Academy of Sciences of the United States of America 88:780-784.

Kaur S, Junek A, Black MA, Semba K (2008) Effects of ibotenate and 192IgG-Saporin lesions of the nucleus basalis Magnocellularis/Substantia Innominata on spontaneous sleep and wake states and on recovery sleep after sleep deprivation in rats. Journal of Neuroscience 28:491-504.

Kawai H, Lazar R, Metherate R (2007) Nicotinic control of axon excitability regulates thalamocortical transmission. Nature Neuroscience 10:1168-1175.

Kilgard MP, Merzenich MM (1998a) Cortical map reorganization enabled by nucleus basalis activity. Science 279:1714-1718.

Kilgard MP, Merzenich MM (1998b) Cortical map reorganization enabled by nucleus basalis activity. Science 279:1714-1718.

Kimura F, Fukuda M, Tsumoto T (1999) Acetylcholine suppresses the spread of excitation in the visual cortex revealed by optical recording: possible differential effect depending on the source of input. European Journal of Neuroscience 11:3597-3609.

Kleiner S, Bringmann A (1996) Nucleus basalis magnocellularis and pedunculopontine tegmental nucleus: Control of the slow EEG waves in rats. Archives Italiennes De Biologie 134:153-167.

Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. Journal of Neuroscience 25:3661-3673.

Kudoh M, Seki K, Shibuki K (2004) Sound sequence discrimination learning is dependent on cholinergic inputs to the rat auditory cortex. Neuroscience Research 50:113-123.

Langford GW, Meddis R, Pearson AJD (1974) Awakening latency from sleep for meaningful and nonmeaningful stimuli. Psychophysiology 11:1-5.

Lewandowski MH, Muller CM, Singer W (1993) Reticular facilitation of cat visual cortical responses is mediated by nicotinic and muscarinic cholinergic mechanisms. Experimental Brain Research 96:1-7.

Li JL, Bickford ME, Guido W (2003) Distinct firing properties of higher order thalamic relay neurons. Journal of Neurophysiology 90:291-299.

Liang L, Lu T, Wang XQ (2002) Neural representations of sinusoidal amplitude and frequency modulations in the primary auditory cortex of awake primates. Journal of Neurophysiology 87:2237-2261.

Lindsley DB, Schreiner LH, Knowles WB, Magoun HW (1950) Behavioral and EEG changes following chronic brain stem lesions in the cat. Electroencephalography and Clinical Neurophysiology 2:483-498.

Lippe WR, Weinberg.Nm (1973) Distribution of click-evoked activity within medial geniculate body of anesthetized cat. Experimental Neurology 39:507-523.

Linster C, Garcia PA, Hasselmo ME, Baxter MG (2001) Selective loss of cholinergic neurons projecting to the olfactory system increases perceptual generalization betwen similar, but not dissimilar, odorants. Behavioral Neuroscience 115:826-833

Livingstone MS, Hubel DH (1981) Effects of sleep and arousal on the processing of visual information in the cat. Nature 291:554-561.

Luczak A, Bartho P, Marguet SL, Buzsaki G, Harris KD (2007) Sequential structure of neocortical spontaneous activity in vivo. Proceedings of the National Academy of Sciences of the United States of America 104:347-352.

Maalouf M, Miasnikov AA, Dykes RW (1998) Blockade of cholinergic receptors in rat barrel cortex prevents long-term changes in the evoked potential during sensory preconditioning. Journal of Neurophysiology 80:529-545.

Mackintosh NJ (1975) Theory of attention – Variations in associability of stimuli with reinforcement. Psychological Review 82:276-298.

Marino J, Cudeiro J (2003) Nitric oxide-mediated cortical activation: A diffuse wake-up system. Journal of Neuroscience 23:4299-4307.

Mariotti M, Formenti A, Mancia M (1989) Responses of VPL thalamic neurons to peripheral stimulation in wakefulness and sleep. Neuroscience Letters 102:70-75.

Massaux A, Dutrieux G, Cotillon-Williams N, Manunta Y, Edeline JM (2004) Auditory thalamus bursts in anesthetized and non-anesthetized states: Contribution to functional properties. Journal of Neurophysiology 91:2117-2134.

Massimini M, Huber R, Ferrarelli F, Hill S, Tononi G (2004) The sleep slow oscillation as a traveling wave. Journal of Neuroscience 24:6862-6870.

Massimini M, Ferrarelli F, Huber R, Esser SK, Singh H, Tononi G (2005) Breakdown of cortical effective connectivity during sleep. Science 309:2228-2232.

McCormick DA (1992) Neurotransmitter actions in the thalamus and cerebral-cortex and their role in neuromodulation of thalamocortical activity. Progress in Neurobiology 39:337-388.

McGaughy J, Dalley JW, Morrison CH, Everitt BJ, Robbins TW (2002) Selective behavioral and neurochemical effects of cholinergic lesions produced by intrabasalis infusions of 192 IgG-saporin on attentional performance in a five-choice serial reaction time task. Journal of Neuroscience 22:1905-1913.

McKinney M, Coyle JT, Hedreen JC (1983) Topographic analysis of the innervation of the rat neocortex and hippocampus by the basal forebrain cholinergic system
Journal of Comparative Neurology 217:103-121.

McLin DE, Miasnikov AA, Weinberger NM (2002) Induction of behavioral associative memory by stimulation of the nucleus basalis. Proceedings of the National Academy of Sciences of the United States of America 99:4002-4007.

Mehta SB, Kleinfeld D (2004) Frisking the whiskers: Patterned sensory input in the rat vibrissa system. Neuron 41:181-184.

Merzenic.Mm, Brugge JF (1973) Variation of excitability of neurons in primary auditory cortex in unanesthetized macaque monkey – Effects of sleep and body movement. Journal of the Acoustical Society of America 53:362-&.

Merzenich MM, Jenkins WM, Johnston P, Schreiner C, Miller SL, Tallal P (1996) Temporal processing deficits of language-learning impaired children ameliorated by training. Science 271:77-81.

Metherate R, Weinberger NM (1990) Cholinergic modulation of responses to single tones produces tone-specific receptive-field alterations in the cat auditory-cortex. Synapse 6:133-145.

Metherate R, Ashe JH (1991) Basal forebrain stimulation modifies auditory-cortex responsiveness by an action at muscarinic receptors. Brain Research 559:163-167.

Metherate R, Cox CL, Ashe JH (1992) Cellular bases of neocortical activation – Modulation of neural oscillations by the nucleus basalis and endogenous acetylchonline. Journal of Neuroscience 12:4701-4711.

Miasnikov AA, McLin D, Weinberger NM (2001a) Muscarinic dependence of nucleus basalis induced conditioned receptive field plasticity. Neuroreport 12:1537-1542.

Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial Attention Decorrelates Intrinsic Activity Fluctuations in Macaque Area V4. Neuron 63:879-888.

Moruzzi G, Magoun HW (1949) Brain stem reticular formation and activation of the EEG. Electroencephalography and Clinical Neurophysiology 1:455-473.

Murata, K. (1963). "The Activity of Single Cortical Neurones of Unrestrained Cats During Sleep and Wakefulness." Archives Italiennes de Biologie **101**: 306-331.

Neuman RS, Thompson PM (1989) Serotonin mediates suppression of focal epileptiform activity induced by noxious-stimulation. Epilepsia 30:307-313.

Nilsson OG, Leanza G, Rosenblad C, Lappi DA, Wiley RG, Bjorklund A (1992) Spatial learning impairments in rats with selective immunolesion of the forebrain cholinergic system. Neuroreport 3:1005-1008.

Orman SS, Humphrey GL (1981) Effects of changes in cortical arousal and of auditory-cortex cooling of neuronal-activity in the medial geniculate-body. Experimental Brain Research 42:475-482.

Oswald I, Taylor AM, Treisman M (1960) Discriminative responses to stimulation during human sleep. Brain 83:440-&.

Parasuraman R, Greenwood PM, Haxby JV, Grady CL (1992) Visuospatial attention in dementia of the Alzheimer type. Brain 115:711-733.

Paxinos, G. and C. Watson (1982). The rat Brain in Stereotaxic Coordinates, Academic Press, New York.

Pena JL, Perez-Perera L, Bouvier M, Velluti RA (1999) Sleep and wakefulness modulation of the neuronal firing in the auditory cortex of the guinea pig. Brain Research 816:463-470.

Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. Nature Neuroscience 5:805-811.

Poggio GF, Mountcastle VB (1963) Functional properties of ventrobasal thalamic neurons studied in unanesthetized monkeys. Journal of Neurophysiology 26:775-&.

Poulet JFA, Petersen CCH (2008) Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. Nature 454:881-U836. Ramcharan EJ, Gnadt JW, Sherman SM (2000) Burst and tonic firing in thalamic cells of unanesthetized, behaving monkeys. Visual Neuroscience 17:55-62.

Sachdev RNS, Lu SM, Wiley RG, Ebner FF (1998) Role of the basal forebrain cholinergic projection in somatosensory cortical plasticity. Journal of Neurophysiology 79:3216-3228.

Sakakura H (1968) Spontaneous and evoked unitary activities of cat lateral geniculate neurons in sleep and wakefulness. Japanese Journal of Physiology 18:23-&.

Sanchez-Vives MV, McCormick DA (2000) Cellular and network mechanisms of rhythmic recurrent activity in neocortex. Nature Neuroscience 3:1027-1034.

Saper CB (1984) Organization of cerebral cortical afferent systems in the rat. 2. Magnocellular basal nucleus.. Journal of Comparative Neurology 222:313-342.

Sarter M, Bruno JP, Turchi J (1999) Basal forebrain afferent projections modulating cortical acetylcholine, attention, and implications for neuropsychiatric disorders. In: Advancing from the Ventral Striatum to the Extended Amygdala - Implications for Neuropsychiatry and Drug Abuse: in Honor of Lennart Heimer (McGinty JF, ed), pp 368-382.

Satoh K, Fibiger HC (1986) Cholinergic neurons of the laterodorsal tegmental nucleus – Efferent and afferent connections. Journal of Comparative Neurology 253:277-302.

Schreiber S, Fellous JM, Whitmer D, Tiesinga P, Sejnowski TJ (2003) A new correlation-based measure of spike timing reliability. Neurocomputing 52-4:925-931.

Semba K, Reiner PB, McGeer EG, Fibiger HC (1988) Brain-stem afferents to the magnocellular basal forebrain studied by axonal-transport, immunohistochemistry, and electrophysiology in the rat. Journal of Comparative Neurology 267:433-453.

Sherman SM, Guillery RW (2002) The role of the thalamus in the flow of information to the cortex. Philosophical Transactions of the Royal Society B-Biological Sciences 357:1695-1708.

Shriki O, Hansel D, Sompolinsky H (2003) Rate models for conductance-based cortical neuronal networks. Neural Computation 15:1809-1841.

Shute CCD, Lewis PR (1967) Ascending cholinergic reticular system – Neocortical olfactory and subcortical projections. Brain 90:497-&.

Shulz DE, Sosnik R, Ego V, Haidarliu S, Ahissar E. (2000) A neuronal analogue of state-dependent learning. Nature 403:549-553

Singer W (1998) Consciousness and the structure of neuronal representations. Philosophical Transactions of the Royal Society of London Series B-Biological Sciences 353:1829-1840.

Smith MA, Kohn A (2008) Spatial and Temporal Scales of Neuronal Correlation in Primary Visual Cortex. Journal of Neuroscience 28:12591-12603.

Spehlmann R, Norcross K (1982) Cholinergic mechanisms in the production of focal cortical slow waves. Experientia 38:109-111.

Steriande M, Buszaki (1990) Parallel Activation of Thalamic and Cortical Neurons by Brainstem and Basal Forebrain Cholinergic Systems. Brain Cholinergic Systems. (3-64) Oxford University Press.

Steriade M, Demetrescu M (1960) Unspecific systems of inhibition and facilitation of potentials evoked by intermittent light. Journal of Neurophysiology 23:602-617.

Steriade M, Demetrescu M (1962) Reticular facilitation of responses to acoustic stimuli. Electroencephalography and Clinical Neurophysiology 14:21

Steriade, M. and R. McCarley (2005). "Brain Control of Wakefulness and Sleep, second edition." New York, NY:Kluwer Academic/Plenum Publishers.

Steriade M, Morin D (1981) Reticular influences on primary and augmenting responses in the somatosensory cortex. Brain Research 205:67-80.

Stewart DJ, Macfabe DF, Vanderwolf CH (1984) Cholinergic activation of the electrocorticogram – Role of the substantia innominata and effects of atropine and quinuclidinyl benzilate. Brain Research 322:219-232.

Stoelzel CR, Bereshpolova Y, Swadlow HA (2009) Stability of Thalamocortical Synaptic Transmission across Awake Brain States. Journal of Neuroscience 29:6851-6859.

Stoelzel CR, Bereshpolova Y, Gusev AG, Swadlow HA (2008) The impact of an LGNd impulse on the awake visual cortex: Synaptic dynamics and the sustained/transient distinction. Journal of Neuroscience 28:5018-5028.

Sutton R. and Barto A. (1998) Reinforcement Learning: An Introduction. MIT Press, Cambridge, MA,

Tallal P, Piercy M (1973) Defects of nonverbal auditory-perception in children with developmental aphasia. Nature 241:468-469.

Tallal P, Miller SL, Bedi G, Byma G, Wang XQ, Nagarajan SS, Schreiner C, Jenkins WM, Merzenich MM (1996) Language comprehension in language-learning impaired children improved with acoustically modified speech. Science 271:81-84.

Tallon-Baudry C, Bertrand O, Henaff MA, Isnard J, Fischer C (2005) Attention modulates gamma-band oscillations differently in the human lateral occipital cortex and fusiform gyrus. Cerebral Cortex 15:654-662.

Tremblay N, Warren RA, Dykes RW (1990a) Electrophysiological studies of acetylcholine and the role of basal forebrain in the somatosensory cortex of the cat .1. Cortical neurons excited by glutamate. Journal of Neurophysiology 64:1199-1211.

Tremblay N, Warren RA, Dykes RW (1990b) Electrophysiological studies of acetylcholine and the role of basal forebrain in the somatosensory cortex of the cat.2. Cortical neurons excited by somatic stimuli. Journal of Neurophysiology 64:1212-1222.

Vanderwolf CH, Baker GB (1986) Evidence that serotonin mediates noncholinergic neocortical low-voltage fast activity, noncholinergic hippocampal rhythmical slow activity and contributes to intelligent behavior. Brain Research 374:342-356.

Varela C, Sherman SM (2007) Differences in response to muscarinic activation between first and higher order thalamic relays. Journal of Neurophysiology 98:3538-3547.

Webster HH, Hanisch UK, Dykes RW, Biesold D (1991) Basal forebrain lesions with or without reserpine injection inhibit cortical reorganization in rat hindpaw primary somatosensory cortex following sciatic-nerve section. Somatosensory and Motor Research 8:327-346.

Wenk H, Bigl V, Meyer U (1980) Cholinergic projections from magnocellular nuclei of the basal forebrain cortical areas in rats. Brain Research Reviews 2:295-316.

Whalen PJ, Kapp BS, Pascoe JP (1994) Neuronal activity within the nucleus basalis and conditioned neocortical electroencephalographic activation. Journal of Neuroscience 14:1623-1633.

Womelsdorf T, Fries P, Mitra PP, Desimone R (2006) Gamma-band synchronization in visual cortex predicts speed of change detection. Nature 439:733-736.

Woolf NJ, Butcher LL (1986) Cholinergic systems in the rat brain .3. Projections from the pontomesencephalic tegmentum to the thalamus, tectum, basal ganglia, and basal forebrain. Brain Research Bulletin 16:603-637.

Worgotter F, Suder K, Zhao YQ, Kerscher N, Eysel UT, Funke K (1998) State-dependent receptive-field restructuring in the visual cortex. Nature 396:165-168.

Yu AJ, Dayan P (2005) Uncertainty, neuromodulation, and attention. Neuron 46:681-692.

Zaborszky L, Pang K, Somogyi J, Nadasdy Z, Kallo I (1999) The basal forebrain corticopetal system revisited. In: Advancing from the Ventral Striatum to the Extended Amygdala - Implications for Neuropsychiatry and Drug Abuse: in Honor of Lennart Heimer (McGinty JF, ed), pp 339-367.

Zhu XO, de Permentier PJ, Waite PME (2002) Cholinergic depletion by IgG192-saporin retards development of rat barrel cortex. Developmental Brain Research 136:1-16.

Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychopysical performance. Nature 370:140-143.