Statistical analysis of neural correlates in Decision-making

Marina Martinez-Garcia

TESI DOCTORAL UPF / 2014

Director de la tesi

Prof. Dr. Gustavo Deco

Department of Information and Communication Technologies

The court's PhD was appointed by the rector of the Universitat Pompeu Fabra on ....., 2014.

Chairman

Member

Member

Member

Secretary

The doctoral defense was held on ....., 2014, at the Universitat Pompeu Fabra and scored as .....

PRESIDENT MEMBERS

SECRETARY

A la memòria de dues grans dones, Puri i Balbina.

Nuestras horas son minutos
cuando esperamos saber,
y siglos cuando sabemos
lo que se puede aprender.

Antonio Machado
Proverbios y cantares.

Abstract

We investigated the neuronal processes which occur during a decision-making task based on a perceptual classification judgment.

For this purpose we have analysed three different experimental paradigms (somatosensory, sisual, and suditory) in two different species (monkey and rat), with the common goal of shedding light into the information carried by neurons.

In particular, we focused on how the information content is preserved in the underlying neuronal activity over time. Furthermore we considered how the decision, the stimuli, and the confidence are encoded in memory and, when the experimental paradigm allowed it, how the attention modulates these features. Finally, we went one step further, and we investigated the interactions between brain areas that arise during the process of decision-making.

Resum

Durant aquesta tesi hem investigat els processos neuronals que es produeixen durant tasques de presa de decisions, tasques basades en un judici lògic de classificació perceptual. Per a aquest propòsit hem analitzat tres paradigmes experimentals diferents (somatosensorial, visual i auditiu) en dues espècies diferents (micos i rates), amb l'objectiu d'il·lustrar com les neurones codifiquen información referents a les tàsques.

En particular, ens hem centrat en com certes informacions estan codificades en l'activitat neuronal al llarg del temps. Concretament, com la informació sobre: la decisió comportamental, els factors externs, i la confiança en la resposta, bé codificada en la memòria. A més a més, quan el paradigma experimental ens ho va permetre, com l'atenció modula aquests aspectes. Finalment, hem anat un pas més enllà, i hem analitzat la comunicació entre les diferents àrees corticals, mentre els subjectes resolien una tasca de presa de decisions.

Acknowledgments

La tesi que esteu llegint ha sigut el fruit de la col.laboració entre 4 conjunts de persones, és per tant fruit de l'esforç comú. A tots els meus col.laboradors moltes gràcies per acompanyar-me i aprendre amb mi durant aquests anys:
Juan Abolafia, por tu paciencia y comprensión, you are the milk!!
Mario Pannunzi, por todos los "Buenos dias!" con un buonissimo caffè en la mano.
Andrea Insabato, por tu infinita paciencia con este pingüino patoso.
Jose L. Pardo-Vázquez, te agradezco todas las discusiones, las echaré de monos!
Adrià Tauste Campo per ensenyar-me i ajudarm-me a fer el treball ben fet.

Gracias a mi supervisor de tesis, Gustavo, por introducirme en el mundo de la Neurociencia, y por la comfianza depositada en mi. Telegráficamente gracias!

Dar un sincero agradecimiento a Edmun Rolls, Mavi Sanchez-Vives, Carlos Acuña Castroviejo y Ranulfo Romo, sin vuestra colaboración, no hubiese podido hacer este trabajo.

También agradecer a todos los miembros de mi tribunal: Adrián Ponce-Alvarez, Toni Guillamon, Carlos Acuña Castroviejo, Mavi Sanchez-Vives y Laura Dempere-Marco.

Als meus companys de despatx, per tos els dies bons, i els menys bons compartits! Ruggero, Nano, Manu, Alba, Judith, Bàrbara, Paola, Carlos. I a tots els meus companys de grup durant aquests anys, per tantes vivències conjuntes.

A la meua familia barcelonina: Gemma, Rosa, Tomeu, Pau, Noe, Lluis, Vicent, Montse, i molts més. Als *matemàtics* que l'esperit fme-ero perdure fins l'infinit!, gràcies per estar en molts moments. A Màxim per ser el meu amic.

Durant aquests anys m'he trobat a varios *Dory* pel cami, persones que sense un peque m'han ajudat: als huevones, piltrafilles, socràtic@s, informatico-informaticos... gràcies, continueu amb mi!

A tots els ulls extra, que m'han ajudat a redactar/ corregir/ estructurar la tesi, gràcies!, heu sigut tots molt amables i de molta ajuda.

A tota la gent de la terreta, ara si xampany i pastes!

Als meus pares i al meu germà, gràcies per ensenyar-me a estimar les matemàtiques i per fer-me menjar tants musclos. Mare per ser pilar segur on sostindre'm sempre. A mon pare per sempre creure en mi. A Manel, per descobrir el món amb mi, i finalment, a Jordi, per ser per estar.

Contents

List of Figures

List of Tables

Abbreviations

Neurophysiology		
	S1	Primary **S**omatosensory Cortex
	S2	Secondary **S**omatosensory Cortex
	A1	Primary **A**uditory Cortex
	MPC	**M**edial **P**remotor **C**ortex
	VPC	**V**entral **P**refrontal **C**ortex
	PMv	**V**entral **P**remotor Cortex
	DPC	**D**orsal **P**remotor **C**ortex
	PFC	**P**refrontal **C**ortex
	M1	**P**rimary **M**otor Cortex
	OFC	**O**rbito-**f**rontal **C**ortex
	LIP	**L**ateral **I**ntraparietal
Experimental		
	2AFC	Two-**A**lternative **Fo**rced-**C**hoice task
	ISI	**I**nter stimuli **I**ntervals
	ISI	**I**nter-**S**pike **I**ntervals
	PSTH	**P**eristimulus **H**istogram
	TRO	**T**est bar's **R**elative **O**rientation
Data Analysis		
	LRM	**L**inear **R**egression **Mo**del
	MI	**M**utual **I**nformation
	DI	**D**irected **I**nformation
	Ff	**F**ano **F**actor
	HMM	**H**idden **M**arkov **M**odel
	BIC	**B**ayesian **I**nformation **C**riterion
Neuronal Models		
	SF	**S**ynaptic **F**acilitation Model
	DDM	**D**rift **D**iffusion **M**odels

Chapter 1

Introduction

Decision-making can be defined as the cognitive process of selecting one out of several possible outcomes. Of all factors that intervene in the process, only those which are observable from an outsider's point of view can be quantified. Thus, only outcomes that result in a defined, clear action can be the object of our study. Furthermore, subjective factors introduce unknown variables, making it more difficult to find which underlying rules apply to the decision-making process. Although most of our day-to-day decisions do not meet any of these requirements, it is our hope that knowledge of the more controllable cases will shed some light on the less controllable cases.

While this subject can be approached from many points of view, including philosophy, economics or marketing, our approach stems from the field of neuroscience. The problem of decision-making has become the center of interest of many neuroscientists aiming to understand the neural basis of intelligent behaviour by linking perception and action.

Behavioural aspects of decision-making have been deeply studied for over a century [1]. However, only as the technology of measuring devices (such as electrodes) evolves we can delve deeper into the neurophysiology (Hodgkin and Huxley [2]). In the late 50's this experimental technology allowed scientists to record neuronal activity from animals, and go in depth into neural codification. The most celebrated work in this direction was done by Hubel and Wiesel [3, 4], who related visual stimulus to the neuronal activity from the cortex of cats.

Our interest focuses on the neural correlates of decision-making process. For instance, we want to understand the role of neural activity in perception, memory, confidence or attention, which are key components that lead to the final choice.

There has been remarkable progress in the context of *perceptual discrimination tasks*, in which the principal factors involved in the design are controlled. For example, we can establish a logical rule by which one of the above *possible outcomes* is chosen. Then by recording the subject's brain activity, and more specifically the neuronal activity, we can relate these factors with neuron's dynamics. We are aware that this kind of tasks are a simplification of the decision-making process, and they are far away from the day-to-day decisions humans make. Despite their limitations, these *perceptual discrimination tasks* offer a wealth of information, since they involve several cognitive processes of interest such as perception, working memory, comparison or confidence.

Two-alternative forced-choice tasks (2AFC) are a subclass of *perceptual discrimination tasks* in which the choice is obtained as an objective relationship between two different sensory inputs $(S1, S2)$. At the end of the task the subject must report the choice via a motor response. 2AFC are the focus of this thesis. The reason is twofold: on one hand, most cognitive factors of interest make manifest appearances in them, so there is no loss of generality in restricting thus our attention. On the other hand, 2AFC are the source of all experimental data made available to us.

The aim of this thesis is to shed light into the neuronal processes which occur during a decision-making task based on a perceptual judgment. We ask ourselves about the mechanism the brain (through its neurons) uses to solve these classification tasks. We analyse several internal and external factors involved in the decision-making process:

$$Inputs\,(S1, S2) \rightarrow \left\{ \begin{array}{c} \textbf{Black Box/ Brain:} \\ Perception \\ Working\,memory \\ Comparison \\ Confidence \\ Attention \\ \dots \end{array} \right\} \rightarrow Output\,(Response)$$

In particular, we focused on how the information is preserved in the underlying neuronal substrate in time. Furthermore we considered how the decision, the stimuli, and confidence are encoded in memory, and when appropriate, how attention modulates these pieces of information. Finally, we went one step further, and investigated the interactions between areas that arise during the process of decision making.

To solve the previously mentioned, as well as other related questions, we analyse data from two different species (monkeys, rats) in different sensory modalities (somatosensory, auditory and visual). In what follows we present a description of the tasks conducted. Each task enables us to study neural correlates for a specific parameter of interest. Through linear and non-linear, when appropriate, mathematical techniques, we analyse the relationship between the choice of the tasks and the information of electrophysiological origin at our hands (neuron voltage action potencial). Additional factors, such as the variability of the state of the brain, subject confidence, inter-neuron relationships, and stimuli are introduced into the analysis when it is appropriate to do so.

1.1 Perceptual Discrimination Tasks

Perception is related to the mental process associated with the interpretation of the sensory information. It is a mental representation which ultimately allows the subject to perform the task [5].

In a perceptual discrimination task the sensory information can be obtained through one of the five traditional senses: sight, hearing, taste, smell, and touch [6–10]. In this thesis we focus on three modalities: somatosensory, visual, and auditory, and through these we analyse the information carried by the neurons while the subjects perform the tasks.

Somatosensory discrimination task

One of most studied tasks in the somatosensory perception domain is the *vibrotactile discrimination task*, which was first performed by Mountcastle [11] and then by Romo [8, 12]. In this task the stimuli $(S1, S2)$ are stimulation frequencies (f_1, f_2). At the outset of the task, two consecutive mechanical vibrations with respective frequencies f_1, f_2 are delivered to the subject's fingertip. The monkey must compare the two frequencies and then report the highest frequency by pushing one of two buttons $(f_1 < f_2; f_1 > f_2)$ (see Fig. 1.1). The two frequencies are delivered with a time lag of 3 seconds, forcing the brain to store in memory the traces of the first frequency, to perceive separately the second frequency and then to compare both to make a decision. Hence, the processes of perception, memory, comparison and decision-making can be independently studied.

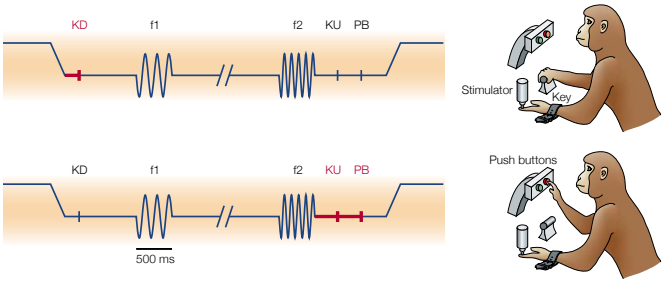

FIGURE 1.1: Vibrotactile discrimination task. The stimulator lowers. The monkey places its free hand on an immovable key (key down, KD). The monkey receives the first vibrotactile stimulation on its fingertip (f_1). After a delay of 3 s, a second stimulus is delivered (f_2). When the second stimulus ends the monkey releases the key (key up, KU) and then pushes one of the two buttons (PB) to indicate its decision ($f_1 < f_2; f_1 > f_2$). Figure adapted from [18].

Romo and colleagues recorded neurons from seven cortical areas: primary somatosensory cortex (S1) [13], secondary somatosensory cortex (S2) [14], medial premotor cortex (MPC) [15], ventral prefrontal cortex (VPC) [16], dorsal premotor cortex (DPC), prefrontal cortex (PFC) [17], and primary motor cortex (M1). This task allowed Romo and colleagues to study the neural encoding of the neurons. They showed that a coding mechanism based on the firing rate predicted strong correlation between the stimulus frequencies and the subject's response. We will come back to their findings in the next section.

Visual discrimination task

In this modality there is a wide range of decision-making tasks being studied. In this thesis we worked with the *bar orientation discrimination task*, designed by Acuña and colleagues [7, 19]. In this task the monkeys were trained to compare the orientation of a reference bar, presented during an initial interval of time with that of a test bar, presented during the second interval. Then, they had to decide whether the test bar was tilted right or left as compared to the reference bar (see Fig. 1.2).

The authors recorded activity from single neurons from the ventral premotor cortex (PMv). They showed that PMv neurons carry information about the first stimulus during the delay and comparison periods, and that neurons also reflect the comparison between stimuli, including the strength of the evidence. Their findings suggest that PMv is involved in shaping future behavior and in learning.

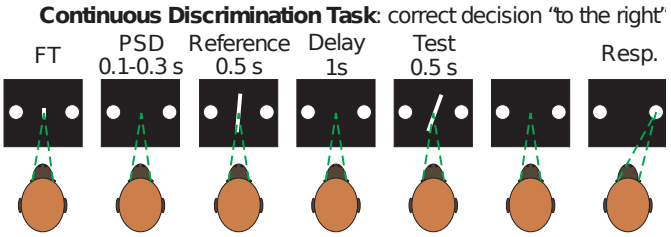

FIGURE 1.2: Visual discrimination task. Two circles appear at the center and at both sides of the screen (fixation time, FT). After a variable delay (100–300 ms), the first bar appears, followed by a delay of 1 s, and then the second stimulus appears. Once the second stimulus is over, the monkey makes a saccade to indicate the relative tilt of the bars. Correct discrimination is rewarded. Figure adapted from [19].

Auditory discrimination task

For this modality we analysed single-unit activity from the auditory cortex (A1) of rats while they performed an *interval-discrimination task*. The animals had to decide whether two auditory stimuli were separated by either 150 or 300 ms and nose-poke to the left or to the right accordingly (see Fig. 1.3). The special feature of this task is that it incoporates different brain states: engaged and idle. These allowed us to study the *Input-Output* relationship under different enviroment conditions.

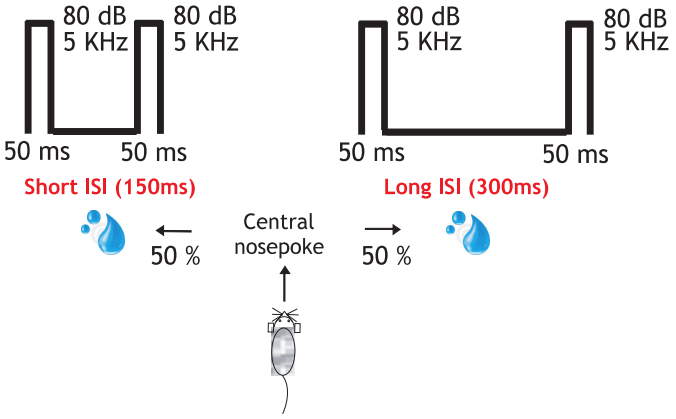

FIGURE 1.3: Auditory interval-discrimination task. The rats entered the central socket and two identical stimuli (50 ms; 80 dB; 5322 Hz) were presented through earphones. About 150 or 300 ms interstimulus interval indicated left or right reward delivery, respectively. Figure adapted from [10].

1.2 Neural code

In our context, the neural code can be defined as the set of rules and mechanisms that neurons would use to represent information about the task. These rules are encoded through neuronal activity, e.g. spikes or silence periods, and allow the system to solve the task.

In this domain, one of the natural questions is: *How can we look into the neuronal activities and predict what is going on in the outside world?*

Luna et al.[13] studied five possible neuronal codes used for S1 neurons during the vibrotactile discrimination task. The five possible codes analysed were: time intervals between spikes, average spiking rate during each stimulus, absolute number of spikes elicited by each stimulus, average rate of bursts of spikes and absolute number of spike bursts elicited by each stimulus. They showed that the code based on the spike count is the most likely code for that task. We have used this result to justify our analysis based on spike count throughout this thesis.

In this section we introduce two different methods which can be applied to spike count coding. In a more formal way, we want to characterize the relationship between the response and task parameters. From a mathematical point of view, the relationship between variables can be roughly divided into the following two categories: the linear relationship and the non linear relationship.

1.2.1 The linear regression models approach

To show the state of the art of the linear relationship between neuronal activity and task parameters we come back to the vibrotactile discrimination task explained above (see Fig. 1.1). As an example, we will summarise now the work done by Hernández et al.[8] and Lemus et al.[12]. The authors analysed single unit activity across seven different cortical areas (see Fig. 1.4), while the monkeys performed a variant of the previously described vibrotactile discrimination task.

In both studies the authors introduced a second delay period between the second stimulus and the report of the response. We call it *the postponed response delay period*. This unique operation requires the subject to store in working memory the decision during the second delay period.

They used a multivariate lineal regression model (LRM) to relate the neuron's firing-rates, $r(t)$, to the stimulus frequencies f_1 and f_2. To compute the firing-rate they counted the number of spikes in a certain time window t:

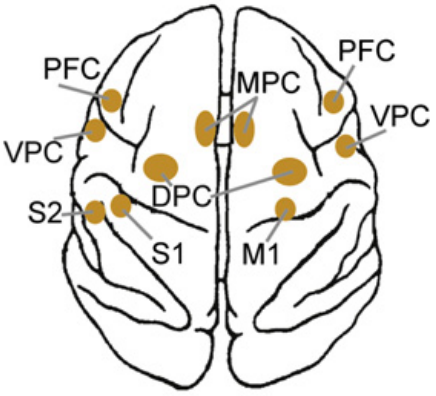

FIGURE 1.4: Top view of macaque brain hihlighting recorded cortical areas: Primary somatosensory cortex (S1), secondary somatosensory cortex (S2), medial premotor cortex (MPC), prefrontal cortex (PFC), ventral prefrontal cortex (VPC), dorsal premotor cortex (DPC), and primary motor cortex (M1). Figure adapted from [8].

$$r(t) = a_0(t) + a_1(t)f_1 + a_2(t)f_2$$

They analysed the significance of the linear regression coefficients (a_1, a_2) as a function of time to characterize the information carried by neurons about the frequencies f_1, f_2 and the decision. Specifically, they classified the function of each neuron as follows:

- If a_1 is significant and a_2 is not significant, the neuron encodes f_1 (green traces in Fig. 1.5).
- If a_2 is significant and a_1 is not significant, the neuron encodes f_2 (red traces in Fig. 1.5).
- If both a_1 and a_2 are significant, of opposite sign but different magnitude, the neuron combines also information about the response with a sensory component (black traces in Fig. 1.5).
- If both a_1 and a_2 are significant and $a_1 = -a_2$, the neuron encodes the decision motor report (blue traces in Fig. 1.5).

If the coefficients are positive the encoding process is called *positive*, otherwise it is called *negative*.

Using the above classification, the authors described the role played by each cortical area during perceptual discrimination. First, the primary somatosensory cortex, S1, was reported to be essentially sensory as its neurons positively encoded

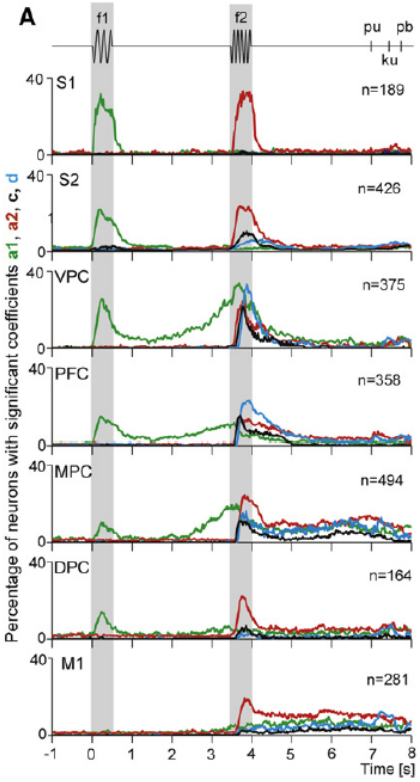

FIGURE 1.5: Percentage of neurons with significant coefficients as a function of time. Green and red traces correspond to neurons with significant a_1, a_2 coefficients, respectively. Black traces indicate neurons in which a_1, a_2 are significant, with opposite signs and different magnitude. Blue traces indicate neurons in which a_1, a_2 are significant, with opposite signs and equal magnitude. Figure adapted from [8].

information about f_1 and f_2 in each stimulation period and remained independent of these variables for the rest of the trial. Areas S2, PFC, VPC, MPC, M1 and DPC also generated neural representation of each stimulus during the whole task, but in contrast with S1, their neurons encoded f_1 and f_2 in a positive as well as a negative monotonic manner.

During the delay between each stimulus, the above mentioned areas, except M1, encoded traces of f_1 in a time window known as the working memory period. After the presentation of the second stimulus, they encoded information about f_2 and f_1 as well as information about the comparison between f_2 and f_1. Hence, these cortical areas encoded at various strengths and times the information of both past and current sensory information on which the perceptual decision report is based.

Even during the *postponed response delay period* the authors found neurons in MPC whose firing rates were correlated with the frequency of the f_1 stimuli and/or the f_2 stimuli (Fig. 2B in [12]). Also in this period, the information carried by

S2 to M1 about the comparison evolved into information that was consistent with the animal's choice. Although the percentage of each neuron type greatly varies from S2 to M1 (see Fig. 1.5), the above findings suggest that sensory, memory and comparison information is gradually conveyed to the frontal lobe circuits that in turn drive the motor circuits for movement execution. The authors also suggested that the neurons that reflect the sensory stimuli in the *postponed response delay period* (Fig. 2B in [12]) enable the subject to revise the report and thus allow for a change in the initial decision.

1.2.2 Non-linear relationships: Information Theory approach

During this section we update and reformulate the first question (*How can we look into the neuronal activities and predict what is ocurring in the outside world?*) from an information-theoretic Theory point of view. So we ask ourselves the following questions:

- *What is the capacity of a neuron to express information?*

- *How much of this "information" is due to a certain experimentally related parameter S (i.e. a sensory stimulus)?*

In this context the capacity of being informative is closely related to the variability of the neuron response (R): the more variable a neuron is, the more informative it can be. Our goal is to quantify the variability (information) due to a given parameter, for example a stimulus S. One way to solve this problem is to calculate the total variability of the neuron and then remove the variability which is not caused by the stimulus S.

Variability due to S= Total amount of variability - Variability NOT due to S

Although this last formula will be next rewritten next in precise terms using information-theoretic tools, we have choosen to include it because of its intuitive value.

Measuring the total amount of variability

The first step is to quantify the total information/variability of a neuron. The classic way to do so is through the *entropy*. In his celebrated paper [20] from 1948, Shannon introduced a quantitative measure of the information conveyed by the

occurrence of an event r, with probability $p(r)$: he defined $I(r) = \log_2 \frac{1}{p(r)}$. From here it was natural to define entropy as the averaged information content of events r of a given variable R:

$$H(R) = -\sum_{r \in R} p(r) \cdot \log_2 p(r) \tag{1.1}$$

Where $p(r)$ is the probability of observing the response r across all trials[1]. Using logarithms in base two, the entropy is expressed in *bits*. As can be seen from 1.1, entropy is the statistical average of the information stored in variable R, and as Shannon proved in [20] gives the minimum number of bits on average required to describe a random variable R .

When applied to neural responses, the entropy gives a measure of statistical variability. In principle, this variability may be due to different factors. To unravel the variability for a specific factor we must use a refinement of this measure.

Measuring the variability which is not due to S

An easy way to tackle the problem is to calculate the variability when the stimulus S is fixed. By setting the stimulus and computing again the entropy we can quantify the variability not caused by it (because it is fixed, and known). To do so, we first calculate : *Conditional entropy*:

$$H(R|S) = -\sum_{r \in R} \sum_{s \in S} p(s, r) \cdot \log_2 p(r|s). \tag{1.2}$$

Where $p(r|s)$ is the probability of observing response r given the stimulus s (i.e. conditioned on s) and $p(s, r)$ is the probability distribution of R and S.

Measuring the variability due to S

Finally, the difference between the unconditional and conditional entropy gives: A

$$I(S; R) = H(R) - H(R|S) = \sum_{s \in S, r \in R} p(s, r) \cdot \log_2 \left(\frac{p(s, r)}{p(s)p(r)} \right). \tag{1.3}$$

This quantity is known as the *mutual information* between stimuli and response. It is the reduction in uncertaninty of R due to the variable S. It measures the strength of association between two variables, and it can be regarded as a non-linear correlation. In particular, MI takes the value zero if the two variables are independent. A graphical representation of the relationship between entropy,

[1]The measure is well defined for $p(r) = 0$, because we used the convention that $0 \cdot \log = 0$, which is justified by continuity since $\lim_{r \to 0} r \cdot \log r = 0$.

conditional entropy and MI is shown in Fig. 1.6, in which we can see that the MI is the common area between $H(Y)$ and $H(X)$. Some of the most relevant properties of this quantity are: it is always positive and symmetrical ($I(S;R) = I(R;S)$).

In this thesis we use this concept as a nonlinear measure of the relationship between two different series. As with the LRM methods this measure can be applied to spike count (firing-rate, $r(t)$) series, but also to the rawest neuronal information: spike trains.

Next example is used to illustrate the presented Information Theory concepts. For more information on these concepts and their properties see [21].

$$I(X;X) = H(X)\,.$$

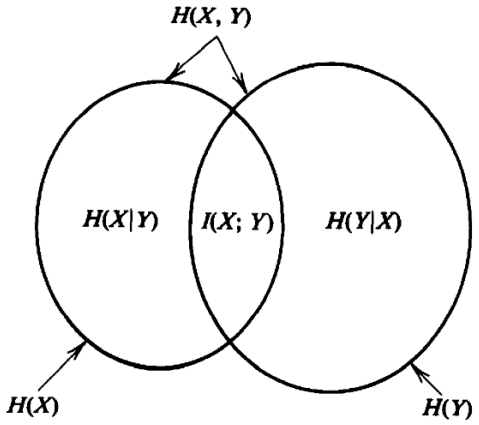

FIGURE 1.6: Relationship between entropy and mutual information. In our case we used R, S instead of X, Y. Adapted from [21].

Example of the Information Theory concepts

FIGURE 1.7: 8-sided dice.

Imagine that we have an 8-sided die like in Fig. 1.7. If someone rolls it and we are given only the parity of the result (even/odd), how much information is that?

In other words, what is the mutual mnformation between the die and the parity? We can answer this question using two procedures: by doing all the calculations, or via intuition.

Applying the formulas

What is the entropy of a roll?

The die can take values from 1 to 8 ($X = \{1, 2, 3, 4, 5, 6, 7, 8\}$) with equal probability, (we assume a fair die). The probability for each side (x_i), or value, is $p(x_i) = 1/8$. Applying equation 1.1 yields its entropy:

$$H(die) = -\sum_{i=1:8} p(x_i) \log_2 p(x_i) = 8\Big(\frac{1}{8}\log_2(8)\Big) = 3 \; bits$$

The next step is to compute the conditional entropy. For that we need to know the conditioned distribution and the probability of a throw ending with and odd/even result. As we have the same number of odd and even rolls: $p(odd) = p(even) = 1/2$. For the conditioned distribution computation:

- In the *odd* case:
  - For $x_i = \{2, 4, 6, 8\}$ $\quad p(x_i|odd) = 0$
  - For $x_i = \{1, 3, 5, 7\}$ $\quad p(x_i|odd) = 1/4$
- In the *even* case:
  - For $x_i = \{2, 4, 6, 8\}$ $\quad p(x_i|even) = 1/4$
  - For $x_i = \{1, 3, 5, 7\}$ $\quad p(x_i|even) = 0$

Now, applying the formula 1.2, the conditional entropy is:

$$\begin{aligned} H(X|parity) &= p(odd) \sum_{i=1:8} p(x_i|odd) \log_2 p(x_i|odd) \\ &\quad + p(even) \sum_{i=1:8} p(x_i|even) \log_2 p(x_i|even) \\ &= \frac{1}{2} \sum_{i=1,3,5,7} p(x_i|odd) \log_2 p(x_i|odd) + \frac{1}{2} \sum_{i=2,4,6,8} p(x_i|even) \log_2 p(x_i|even) \\ &= \frac{1}{2}\Big(4\frac{1}{4}\log_2(1/4)\Big) + \frac{1}{2}\Big(4\frac{1}{4}\log_2(1/4)\Big) = 2 \; bits \end{aligned}$$

Finally, the calculation of their mutual information, following equation 1.3, is:

$$I(X; parity) = H(X) - H(X|parity) = 3 - 2 = 1 \; bit$$

From an Encoding point of view

Decimal	Bits
1	001
2	010
3	011
4	100
5	101
6	110
7	111
8	000

TABLE 1.1: Encoding of 8 numbers in triplets.

Another way to understand the entropy is by using Shannon's theorem (1946), which states that the shortest length of a code is bounded by the entropy of a system. In this particular example (8 equiprobable symbols), the minimal bound can easily be reached, as is shown in the table 1.2.2, though this is not in general the case.

Then if we apply Shannon's Theorem, the entropy of the die is 3 bits, which of course matches the previous calculation.

In order to answer the second question we note that the last symbol of all the odd numbers is a 1 while the last symbol of all even numbers is a 0. So what we know about X if we know its parity is the last bit. In other words the mutual information between X and the parity is equal to 1 bit, which again matches the above calculation.

1.3 Outline of the Thesis

In chapter 2 we study a variant of the vibrotactile discrimination task previously explained (see Fig. 1.1). In this version the authors [8, 12], introduce a second delay between the end of the second stimulus (f_2) and the motor action (PB in Fig. 1.1). In that way they force the subject to hold in memory its choice, allowing us to ask about the internal mechanisms that enable decisions to be postponed for a period after the evidence has been provided. We analyse the firing rate ($r(t)$) of MPC neurons, to relate them with the response. We focus on the dynamics of the whole population. We also tried to understand, by using two different mathematical models, the possible mechanism that generates these dynamics.

In chapter 3 we study single neurons discharge while rats were performing the above auditory discrimination task 1.1. This experiment allows us to ask the following questions: How do the dynamics of A1 neurons vary under different brain states? or, What is the information content about the response (right, left) of the A1 neurons? Does it vary under the different brain states? To awswer these questions we studied the relationship between the firing rate ($r(t)$) of A1 neurons and the behavioural response across engaged and idle brain states.

In chapter 4 we look past the behavioural response and the internal parameters, such a the *confidence* in the response, and how this internal information is encoded. We studied neurons from the ventral Premotor cortex while monkeys were performing the above visual discrimination task (see Fig. 1.2). To understand how the confidence is encoded in the brain we study the firing rate ($r(t)$) of PMv neurons. But as confidence can vary from trial to trial even though the difficulty is fixed we apply a technique which allows us to analyse trial by trial in a big time window without lost information. For this analysis we use the rawest information that we have, the single spike trains, instead of the spike count used in previous chapters.

In the last chapter we take a full turn with respect to our previous approaches. In previous chapters we related the neuronal activity with a certain parameter. In this chapter, our aim is to study the interaction between different cortical areas to characterize communication flows. In particular, we analyse to what extent these communications are dependent on the key stages of the discrimination task: sensory encoding, working memory, effective decision-making and motor action. For that purpose we analysed data simultaneously recorded while monkeys were performing the above vibrotactile discrimination task and also a passive version of the task. The data recorded corresponds to single neurons from five areas simultaneously (S1, S2, SMA, DPC, and M1, see Fig. 1.4). These simultaneous recordings allow us to explore to what extent feed forward and feedback interactions across the study cortical areas are task dependent. Our aim in this chapter is to study neuronal causal correlations:

- To characterize feed forward and feedback communications across the five cortical areas under study at the different stages of the task. We determine the task-specificity of these interactions using a passive task.
- To investigate whether feed forward and feedback communications encode synergistically or redundantly information about the decision report and to classify cortical circuits accordingly.

Chapter 2

Neural and computational mechanisms of postponed decision

The work presented in this chapter was published in:
Proceedings of the National Academy of Sciences (PNAS) in May 2011

Authors: Marina Martinez-Garcia, Edmund T. Rolls, Gustavo Deco, and Ranulfo Romo

2.1 Introduction

An important aspect of decision-making is that actions must often be delayed after the information for the decision has been provided. We examine the mechanisms that underlie this in a well-known paradigm because research in decision-making is the comparison of two vibrotactile stimuli (f_1, f_2) applied to the fingertips with a fixed delay period between them [8]. To perform this cognitive task, the subject needs to store in working memory the information about the first stimulus, f_1, and to compare it with the second stimulus, f_2, to make the decision of whether $f_1 < f_2$ or $f_1 > f_2$, and report it immediately after the f_2 stimulus is released. To extend the analysis of the mechanisms' underlying decisions, which sometimes cannot be made immediately after the evidence is provided, Romo et colleagues [8, 12] introduced an additional delay between the second stimulus and the subject's response. In these works they analyzed neuronal activity by using a linear regression model. In the *Introduction* (1.2.1) we have summaryzed their findings.

In this chapter, we use a different approach to the linear regression model (LRM see section 1.2.1 for more information) to analyze the nature of the activity in

the postponed response delay period. We measure the mutual information (MI) between the neuronal activity and the postponed response, which takes into account the variability of the firing [21–24], and apply this to understanding neuronal activity in the MPC recorded in the same paradigm as [12]. We wished to measure the information in the firing rates during the delay period. To do so, we computed the MI between the variables "firing rate" and "category of response" (i.e., $f_1 < f_2$ or $f_1 > f_2$) throughout the 3-s delay period in which the response is being postponed. For the population of neurons with significant information during the second stimulus (f_2), when the evidence required to make the decision is available, many of the neurons have low information during the delay period after f_2 before the postponed response can be made, but recover the information when the response must be made, as illustrated in Fig. 2.1B. We propose here that the information can be recovered by a nonspecific input applied at the time of the response, and demonstrate this with two models. In one model, synaptic facilitation (SF) [25–27] occurring in the postponed response delay period allows the memory to be maintained with little firing, and little information from the firing, in the delay period. In a second model, graded firing rates in the attractor decision-making network [28] in the delay period allow the faster-firing neurons to maintain sufficient firing and information in the delay period so that the nonspecific recall cue can activate these and the lowerfiring neurons in the same attractor network to their full information value when the behavioural response is required at the end of the delay period.

2.2 Results

2.2.1 Information in the Neuronal Activity of Medial Premotor Cortex Neurons

The activity of a single MPC neuron in the postponed response delay task is illustrated in Fig. 2.1 A and B. The monkeys (Macaca mulatta) were trained to discriminate in frequency between two consecutive mechanical vibrations (f_1 and f_2) delivered to one fingertip [8]. The monkeys were asked to report the results of the decision about which vibrotactile stimulus, f_1 or f_2, had a higher flutter frequency after a fixed delay period (3 s) between the end of f_2 and the cue that triggered the beginning of the motor report. Sequential decision tasks with a delayed response require information about f_1, temporally stored in working memory, to be compared with the current information from f_2 to form a decision

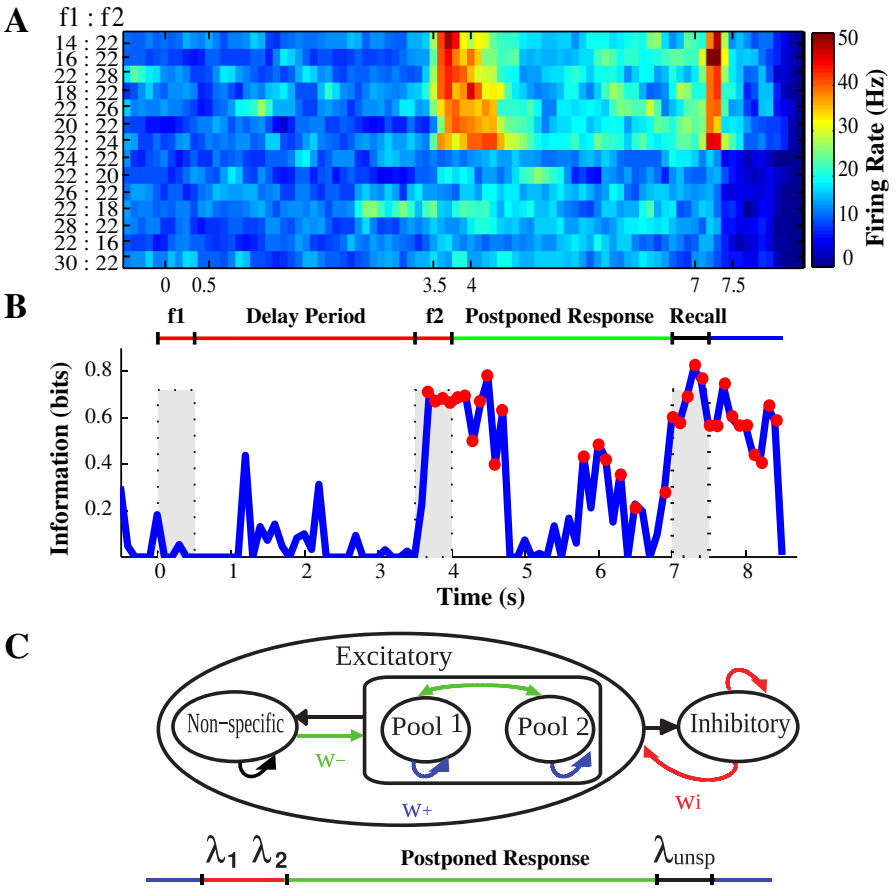

FIGURE 2.1: **Response of a single decision-making neuron recorded in the MPC. (A)** Color map of the firing rate (Hz) in bins of 200 ms moving in steps of 100 ms. Each row is the mean activity for a specific pair of stimuli $[f_1 : f_2]$. Labels indicate the pairs of vibrotactile frequencies $[f_1 : f_2]$ for each trial type. The trials are organized into two blocks: first with $f_1 < f_2$, and then $f_1 > f_2$; this makes the decision-making pattern of neuron firing clear: the neuron fires for decisions when $f_1 < f_2$, for a short time after f_2, and close to the time of the response. Only correct trials were included in the analysis. The time scale beneath the map shows the periods of stimulation and for reporting the decision. The f_1 period is 0.0–0.5s. The delay period between the stimuli is 0.5–3.5s. The f_2 period is 3.5–4.0s. The postponed response delay period is 4.0–7.0s. The behavioral response can be started at a signal given at the time of 7.0s. **(B)** The MI between the activity shown in A and the category of the response with a surrogate correction (200 surrogates) [29]. The red circles indicate the values that are significant at $p < 0.05$, tested with a first-order Monte Carlo method. The shaded rectangles show the stimulation periods $f_1$and f_2, and the report period. **(C)** Architecture of the spiking integrate-and-fire attractor network model of decision-making and of activity in the subsequent delay period. The model consists of two different neuronal populations: excitatory and inhibitory (interneurons). There are two types of excitatory population: selective (pool 1 and pool 2 for each of the two decisions) and nonselective. The recurrent arrows indicate recurrent connections between the different neurons in a pool, and the other arrows show the different connections between the groups. The selective pools first receive $\lambda_1 = 250$ Hz and $\lambda_2 = 150$ Hz during the red period. The inputs are then removed during a delay of 3 s in which the same unspecific input $\lambda = 204$ Hz (0.255 Hz for each of 800 synapses onto each neuron) is injected to both selective pools. For nonselective and inhibitory neurons, $\lambda_{unsp} = 0$.

of whether $f_2 > f_1$ or $f_2 < f_1$. After the discrimination, the subject must keep the decision or response in working memory for 3 s, and then report the outcome by pressing one of two push buttons (one for each option: $f_2 > f_1$, $f_2 < f_1$) after the postponed response delay (Fig. 2.1B). Fig.2.1A shows that the neuron fired during the f_2 period on trials in which $f_2 > f_1$, had a lower firing rate in the postponed response delay period (4–7 s), and then increased its rate when the signal was given that the behavioral response could be made (7.0–7.5 s). Fig. 2.1B, the MI analysis for the same neuron, shows that there was information in the firing rate about the response made at the end of each trial during and just after the f_2 period when the decision could be made, and that the information decreased during the postponed response delay period and increased again when the signal for the behavioral response was given.

The single-cell mutual information for the whole set of MPC neurons is shown in Fig. 2.2 *Top*. It is clear that some neurons do have low information in the postponed response delay period.These neurons also tend to have low information in the f_2 and response periods. Other neurons have higher amounts of information in the delay period, and these neurons tend to have higher information in the f_2 and response periods. The neurons are ordered according to how much information they have in the f_2 period. We confirmed that some of these neurons maintain their firing in the delay period even after a strict Holm–Bonferroni correction for multiple tests [30] described in section 2.4, Fig. 2.5 and Table 2.1.

More tellingly, we performed a multiple-cell information analysis, which tests how the information about the decision increases with the number of neurons in the sample [22, 31] (Figs. 2.5 and 2.7). We found that with 18 neurons taken at random from those with low single-cell information content in a 600-ms window in the delay period (5.2–5.8 s), the average information per neuron was 0.06 bits, and the total information provided was 0.51 bits, with a 90% correct prediction of the decision (Fig. 2.6). (These 18 neurons had low information even during f_2, on average 0.4 bits/ neuron, and we needed 14 such neurons selected at random to reach 1 bit of information). If we consider 16 randomly selected neurons from those with the higher information values shown in Fig. 2.3, then the multiple-cell information analysis showed that the average amount of information for each cell was 0.56 bits in the same 600-ms window in the delay period, and that with subsets of cells chosen at random from the 16 cells the information reached 1 bit and 100% correct with four to six cells (Fig. 2.7). The implication is that with just six MPC cells chosen at random from the set with higher information values in the delay period shown in Fig. 2.2, the animal could do the task perfectly, with 100% correct.

Overall, the single- and multiple-cell information analyses show that while some neurons do contain little information in the delay period about which a report will be made at the end of the postponed response delay period, other neurons do maintain significant information in the delay period, with four to six such MPC neurons being sufficient to account for the correct behavioral response at the end of the delay period.

We now consider two models that examine the basis of the recovery of the information at the end of the postponed response delay period when the behavioral response must be made. The models have in common the fact that a nonspecific external input applied at the time when the response can be made after the delay allows the information to be recovered in the neuronal firing, as shown in Fig. 2.2 *Top*.

2.2.2 Synaptic Facilitation Model

To explain the mechanism underlying the appearance of information about the decision during and just after f_2, then its disappearance for some neurons in the postponed response delay period, and finally its reappearance at the response time, we made an integrate-and-fire attractor network model [25, 26, 32] of the decision-making neurons (Fig. 2.1C) that is able to reproduce that pattern of activity by incorporating SF. There are two decision populations, or pools, of neurons, with pool 1 activated by stimulus f_1 via the λ_1 inputs and winning the competition if $f_1 > f_2$, and pool 2 activated by stimulus f_2 via the λ_2 inputs and winning the competition if $f_1 < f_2$. The global inhibition produced by the inhibitory neurons, together with the different λ_1 and λ_2 inputs to pools 1 and 2, provide the basis for the competition, which is influenced by the randomness of the spiking times of the neurons to produce probabilistic choice [22, 33]. We also implemented short-term SF (see section 2.4) in which the calcium-mediated SF makes the residual calcium level grow [34]. Each neuron that spikes increases the residual calcium level, u, in the presynaptic terminals, which in turn increases the release probability. The time constant for this process was 2 s. Details of the implementation and operation of the simulation are in section 2.4.

The results of the SF simulations are illustrated in Fig. 2.2 *Bottom* and Fig. 2.3 *Lower*. After a period of spontaneous activity before $t = 3.5$ s, the decision cues λ_1 and λ_2 are applied at $t = 3.5 - 4.0s$. If $\lambda_1 > \lambda_2$, pool 1 corresponding to a decision that $f_1 > f_2$ wins, and its firing rate and the MI between the firing and the behavioral response increases. In the delay period from 4 to 7 s, the decision

cues are no longer present, and the firing rate and the MI decrease to close to zero (Figs. 2.2 and 2.3). During this delay period, the synaptic facilitation between the neurons in pool 1 that occurred during f_2 in decision pool 1 remains, gradually decaying (Fig .2.8). When a nonspecific external input (λ_{unsp}) is applied at $t =$ 7.0–7.5 s to both pools 1 and 2 to reflect the moment when the subject receives the stimulus to give its behavioral response, then because of the altered synaptic calcium levels, the firing rate of one of the selective pools increases to the attractor activity level (Fig. 2.3), as does its information about the response to be made (Fig. 2.4), whereas the firing of the other selective pool remains with low activity, although a little higher than the spontaneous firing rate. In this way, the SF model recalls the selective firing for the correct response even though there was no firing in the delay period (Figs. 2.2 –2.4).

It has been shown that as the postponed response delay period increases, the performance of the subjects decreases [12]. We found that the performance of the SF model decreases in a similar way over periods of up to 3 s (Table 2.3). However, it is a prediction of the SF model that performance will decay as the shortterm memory period increases much beyond the time constant of the synaptic facilitation, 2 s. We therefore examine a firing rate model in the next section that can maintain the memory over much longer periods than this. Further, it is a property of the SF model that it can perform the delay task with no firing and no firing rate information in the delay period, because the memory is held in the facilitated synaptic weights (Figs. 2.2 –2.4); however, this was not found for the MPC neurons, the majority of which do have some firing during the delay period (Figs. 2.3 and 2.4), and retain some information in their firing rates in the delay period (Fig. 2.2). Moreover, the MPC neurons have a distribution of firing rates and of the single-cell information in the delay period, with some neurons maintaining their firing rates and firing rate-related information in the delay period well, and others less so (Figs. 2.2 –2.4). We therefore analyzed a different model of the delay-related firing, which has a graded distribution of firing rates.

2.2.3 Graded Firing Rate Attractor Network Model of the Activity in the Delay Period After the Decision.

An integrate-and-fire decision-making network was implemented in the way just described, but without any synaptic facilitation, and with graded firing rates. The gradation to the firing rates was implemented by replacing the equal and strong

synaptic weights (w_+) that connect the neurons within each specific decision-making pool with an exponentially graded set of synaptic weights. The graded synaptic weights had the same average value ($w_+ = 2.1$), but some were considerably stronger, and the distribution decreased exponentially to a value of 1, as described elsewhere [28] and in section 2.4. Graded firing rate distributions are not usually examined in these attractor decision-making networks because the mean field analysis used to set the network parameters requires the same value for the weights within a pool. However, because neurons in the brain typically have graded firing rate distributions, frequently close to exponential [22, 35], we have investigated the properties of decision-making networks with graded firing rate representations [28]. Here, we investigate the activity of a similar network in a short-term memory period after a decision, and analyze whether the firing rate distribution in the delay period is similar to that found for MPC neurons in that some neurons encode little information in their firing rates, and others more information. For these graded simulations, there was a predecision cue period of spontaneous firing; the decision cues were applied at 3.5–4.0 s, and in this period the decision was made; there was a 3-s delay period from 4 to 7 s, and then a nonspecific input was applied equally to the two decision pools (1 and 2) from 7.0 to 7.5 s, to investigate whether just pool 1, which had reached the decision during 3.5–4.0 s, could be restored to its high and selective activity with respect to pool 2 when the behavioral response was required.

The performance of the graded firing rate network for neurons in the winning pool (pool 1) is illustrated by rastergrams in Fig. 2.3. Some neurons continue firing in the delay period, whereas others decrease their rates considerably, to only slightly above the spontaneous rate shown in the spontaneous period. This pattern is qualitatively similar to that found for the MPC neurons (Fig. 2.3, *Upper*). Fig. 2.4 shows that the firing rates for these simulations are indeed graded in the delay period, and that, as is expected, the firing rates in the delay period are monotonically related to the firing rate in the decision period (3.5–4.0 s). Fig. 2.4 also shows that the MPC neurons have graded rates in the delay period, as well as in the f_2 decision period. Fig. 2.2 shows that the graded firing rate simulations have graded information conveyed by the different neurons, with some with relatively high firing rate information, and others with much lower firing rate information, in the delay period. This is qualitatively similar to the MPC neurons' information measures throughout the task, although as the MPC neurons are noisier from trial to trial than the graded simulation, the actual magnitude of the information is less for the MPC neurons than for the graded simulations.

Of particular interest and theoretical significance is that the graded simulations

recover the information when the external stimulus is applied nonselectively (i.e., equally) to pools 1 and 2 at $t = 7.0$–7.5 s. The concept here is that with low inputs during the delay period, the neurons overall have less activity than in the decision period when the decision cues are applied. However, the correct decision pools can maintain the identity of the decision by having just some neurons firing at a sufficient rate to keep the attractor active by the feedback of the firing rates through the graded synaptic weights. The neurons with the low weights in the graded distribution may have very little firing in the delay period, and indeed may be firing at a rate insufficient to maintain the attractor themselves in the delay period (as shown by further simulations). However, when the nonselective external signal is applied at $t = 7.0$–7.5 s equally to neurons in pools 1 and 2, the neurons in pool 1 that are already active are stimulated into higher firing, which has the effect of recruiting through the intrapool recurrent synaptic connections the other lower-firing neurons in pool 1, and also, by competition through the inhibitory interneurons, keeping the neurons in pool 2 at low activity. This results in the recovery of information during $t = 7.0$– 7.7 s, illustrated in Fig. 2.2, which models what is shown for the MPC neurons in Fig. 2.2. (The actual process at the time of the application of the external stimulus at $t = 7.0$–7.5 s can also be viewed as a decision-making process in which $\lambda_1 = \lambda_2$, but there is a bias to λ_1 from the greater activity left in pool 1 than in pool 2 at the end of the delay period).

In summary, the graded firing rate model of decision-making has properties that capture many of the properties of the MPC neurons (Figs. 2.2-2.4). These properties include the maintenance of low but significant information in the delay period, which is graded according to the graded firing rates, and the restoration of the information when it needs to be recalled, by a nonselective external input in the case of the network. This operation is different from that of the synaptic facilitation model, which can hold the memory in the synaptic facilitation with no firing rate or information evident in the delay period (Figs. 2.2-2.4), and which has only a limited short-term memory period.

2.3 Discussion

In this chapter we have shown using mutual information analyses that though some MPC neurons lose their information about the decision in a subsequent delay period before the behavioral response can be made, some neurons maintain that information, although at a lower level and with lower firing rates than during the

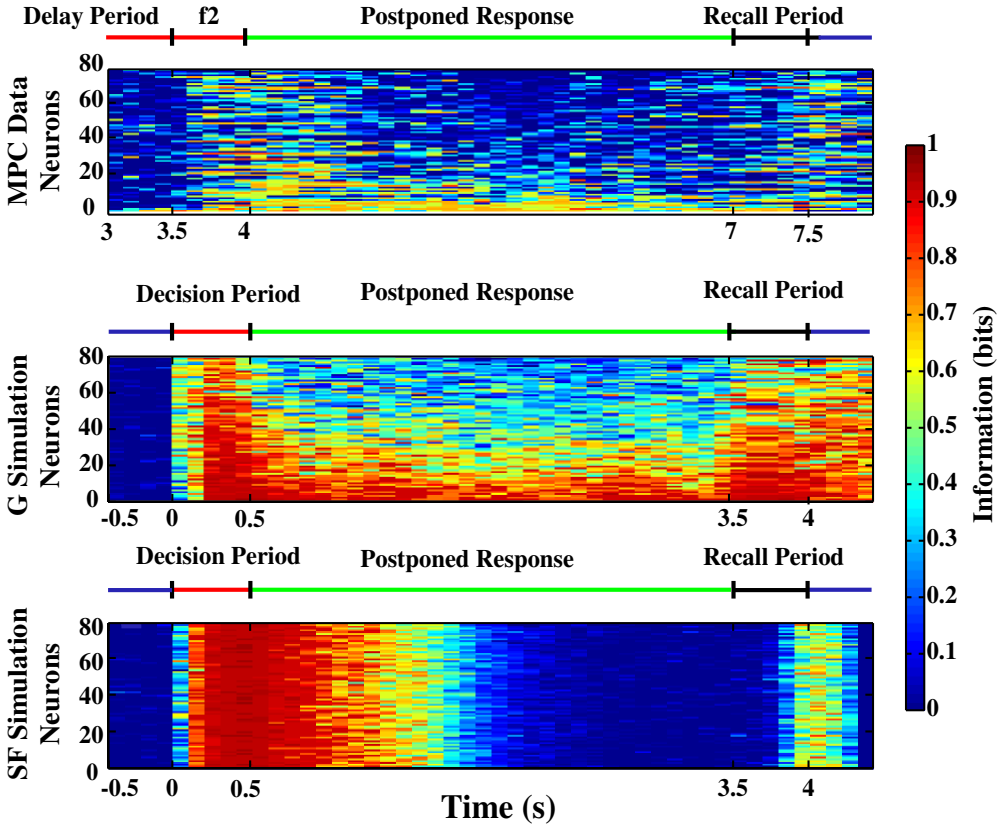

FIGURE 2.2: **Mutual information analyses as a function of time** (*Top*) Eighty neurons from the medial premotor cortex. Each row corresponds to a different neuron. All neurons in the dataset are included for which there were five or more trials for each condition. The times when the f_2 stimuli are applied (the delay period and the behavioral response period) are indicated (Fig. 2.1). In particular, the f_2 period is 3.5–4.0 s. The postponed response delay period is 4.0–7.0 s. The behavioral response can be started at signal given at 7.0 s. The mutual information shown is that between the firing rate in a 200-ms window (sliding every 100 ms) and the response made by the monkey. The calibration bar shows the information value for a single neuron. (*Middle*) Eighty neurons in pool 1 of the graded firing rate simulation. Each row is a single neuron. The rows are sorted by the amount of information during the decision period, 3.5–4.0 s, which corresponds to the f_2 period for the MPC neurons. The delay period is 4–7 s. The equal external inputs are applied at t = 7.0–7.5 s, labeled recall period. The mutual information shown is that between the firing rate in a 200-ms sliding window and the firing in the decision period. *(Bottom)* SF simulation (conventions as in *Middle*).

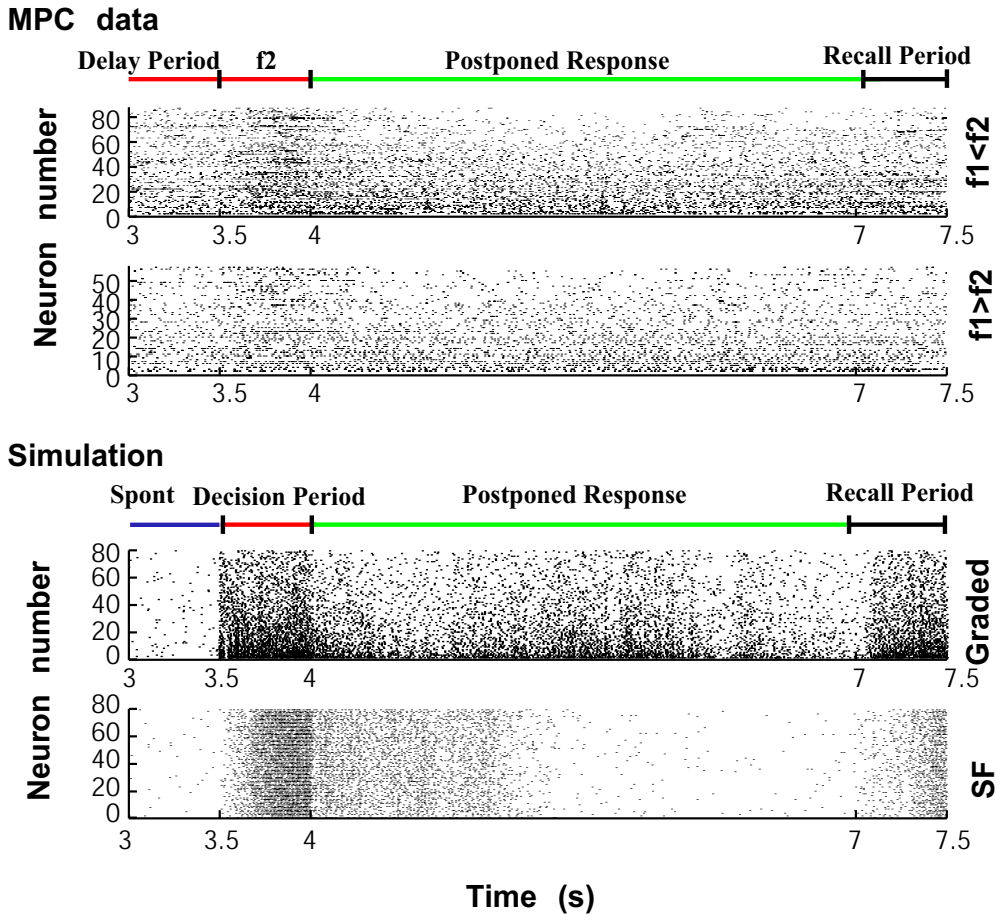

FIGURE 2.3: **Rastergrams to illustrate the firing for the MPC neurons, the graded simulations, and the SF simulations.** For the MPC neurons,*(Upper)* with firing to $f_1 < f_2$, and the *Lower* set of rasters is for neurons with firing to $f_1 > f_2$. One trial is shown for each neuron, and the trial selected is one in which that neuron by its highfiring rate encodes the decision. For the graded simulation, each row is the firing for a different neuron, and all of the data are from one simulation trial, to show how the rates for the different neurons remain graded throughout the trial, including when there is some fluctuation of average firing rate in the delay period. For the synaptic facilitation simulations, one trial is also shown. The times correspond to those in Fig. 2.1. The recall period for the MPC data was when the behavioral response could be initiated, and for the simulationswas when the unspecific input was applied to produce recall.

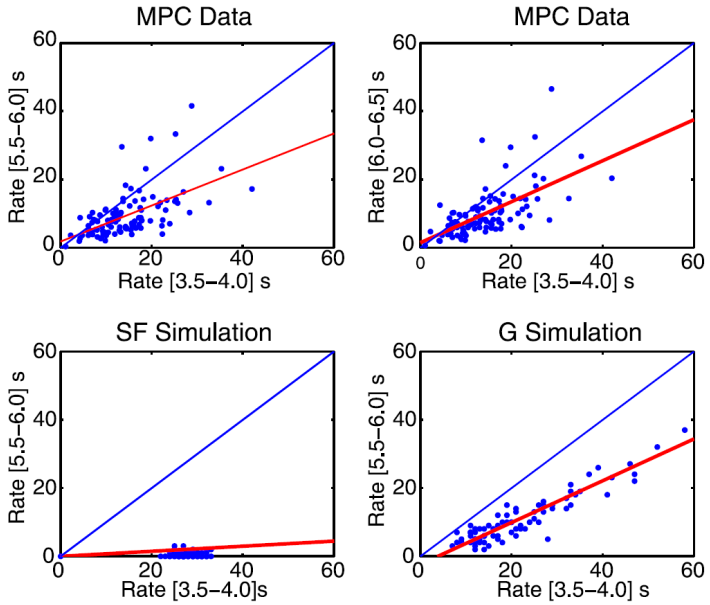

FIGURE 2.4: **The firing rate during the delay period (ordinate) vs. the firing rate in the decision period (abscissa)**. Each point represents one neuron *(Upper)* MPC neurons. Scatterplots show the rates for two different 0.5-s time windows in the delay period. *(Lower Left)* SF simulations. *(Lower Right)* Graded firing rate simulations. The blue line shows where an equal response would lie. Red lines show linear fits to the data.

decision period (f_2) (Figs. 2.2-2.4). We have shown previously that an integrate-and-fire decision-making attractor network can account for the decision-making itself in this task [26], and show here that the same network can also account for the memory of the decision in the subsequent delay period before the behavioral response can be made. This use of the same network to make the decision and to hold the decision in a short-term memory is a good and evolutionarily efficient property of this model of decision-making [22, 33].

The two mechanisms we model for the details of how these processes are implemented have different advantages. The SF mechanism is energy efficient, for it can, as we show here, maintain the evidence of a previous decision even with no neuronal firing (Figs. 2.2-2.4); however, it is limited by the maximum duration of its short-term memory, set by the time constant of the synaptic facilitation, which is $\sim 1 - 2s$ [25, 36]. The SF model can be made to have, as in the MPC (Figs. 2.2-2.4) and many other brain areas [12, 22, 37, 38], some continuing firing in the delay period by increasing the nonspecific inputs to the neurons in the delay period. In the MPC, in which some neurons have low single-cell information levels at some point in the delay, the multiple-cell information shows that a group of as few as seven neurons in the group with most single-cell loss of information never

fall below 0.42 bits, allowing 90% correct performance in the delay period (see section 2.4). The graded firing model of decision-making and subsequent short-term memory for the decision described here shows that low levels of firing in some neurons in an attractor population with graded firing rates can enable the attractor to be maintained, and later restored, even when some of the neurons in the graded representation have low activity in the delay period, and correspondingly low MI values. In the graded firing rate simulations analyzed, we took the parameters—in particular the inputs being applied equally to pools 1 and 2 in the delay period—down to levels that just enabled the attractor to be maintained. Because neurons in the MPC and in many other cortical areas [22] have graded representations, ours is an accurate model of the neurophysiological mechanism. And by using firing rates, the graded firing rate model can maintain its firing rates for potentially long periods, of tens of seconds.

In practice, it could well be that the cortex uses both mechanisms described here in the same network. That is, the level of firing that is required in a small proportion of neurons in the graded firing rate attractor population may be sufficient to produce some synaptic facilitation, and thus synaptic facilitation may be involved in the cortex in these types of networks. In this mixed scenario, however, the mechanism relies more on the graded nature of the firing rate representations, and for the attractors to be maintained by low firing rates of at least some neurons, especially when the memory must be for more than a few seconds.

Asynaptic facilitation mechanism has been proposed previously to help with the memory of f_1 during the delay between the stimuli f_1 and f_2 in a sequential decision-making task [26], but has not been suggested before for implementing the memory in a postponed decision task. We have shown here that synaptic facilitation provides a possible mechanism for remembering the decision during the delay period before the decision can be reported.

In this chapter we used MI to quantify the relationship between the firing of the MPC decision-making neurons during different parts of the task, and the decision taken when f_2 is applied. MI analysis is useful because it takes into account the variability of the firing from trial to trial, and at different times within a trial. The MI analysis shows that the information about the response is not significant until the end of the second stimulus (by which time f_1 and f_2 have been presented; Fig. 2.2). Immediately after f_2, the analysis shows an association between the firing rates and the later response of the monkey. Then for some neurons (Fig. 2.2), the information becomes low during the delay period only to come back again at the response time. At the report time the MI becomes significant as a consequence of the large difference in the firing rates of the two selective populations. The

multiple-cell information analysis is helpful in showing that if a population of neurons with low single-cell information values in the delay period are considered together, then because the information from the different neurons adds (Fig. 2.6), there is in fact some information provided even by these neurons in the delay period. The MI analysis also highlights that it is the high firing rate neurons that encode much information, as the information measure reflects the difference in the firing rates—that is, the number of spikes in a short time window between the two populations being considered [22, 31].

The relative contribution of the two mechanisms, synaptic facilitation vs. restoration of firing in the whole set from the firing in a subset with graded activity, could be investigated experimentally by using longer delay periods before the behavioral response can be made. The synaptic facilitation mechanism with its time constant of 2 s would predict very poor performance (by the neurons and behaviorally) if the delay period is increased to 5 s. In contrast, the mechanism that involves restoration of firing in the whole set of neurons from the firing remaining in a subset could in principle restore the firing in all of the neurons in the appropriate decision pool after much longer periods.

Finally, we note that in the first delay period, between f_1 and f_2, a sensory representation of f_1 must be stored. In the postponed response delay period, the result of the decision must be stored, and we show in this chapter that it is natural for the decision-making network to also store the results of the decision, for it is an attractor decision-making mechanism that we analyze. The networks that implement these memories must be different, and are shown to be different by the fact that different neuronal populations are engaged by these two processes [8].

2.4 Methods

2.4.1 Information Analysis

Single-cell information analysis

We are interested in knowing the information carried by neuronal responses, especially how the firing rate is related to the stimulus or to the category of the behavioral response. To estimate the information content carried by the neuron's firing rate, we performed mutual information (MI) analysis [21, 22], which quantified the average amount of common information contained in the variables R and

S —in our case, firing rate and category of behavioral response. In other words, the MI reflects the uncertainty removed from one by knowing the other. The MI was calculated following Eq. 1.3.

In the data, the calculation of firing rate was based on the mean number of spikes across trials, obtaining for each bin an $n-$dimensional vector of spike counts, where n is the number of frequency pairs. One frequency pair is the set of $[f_1 : f_2]$ stimuli presented on one specific trial. This computation was performed in sliding windows of 200 ms in steps of 100 ms for each pair of frequencies. The variable category (the behavioral response) is represented by ones with the sign of the difference between a specific frequency pair $[f_1 : f_2]$, i.e., $+1$ for $f_1 > f_2$, and -1 for $f_1 < f_2$. We only included "correct" trials in the analysis, i.e., only the trials on which the monkey solved the task correctly. Because the MI estimate is subject to statistical errors that can lead to an overestimate of the information, we corrected the information estimates using a first-order Monte Carlo method. In this correction procedure the mean information from many runs in which the stimuli and the responses are shuffled across trials is subtracted from the information estimate from the original unshuffled data [22, 39]. We applied this Monte Carlo method basing the correction on 50 shuffled runs. This method also leads to a test of the statistical significance of the corrected MI between firing rates and category [29].

The same method was used in the simulations to measure the MI between the neuronal firing in different 200- ms epochs through a trial, and the decision that had been taken by the network represented by which pool had the high firing rate in the decision period when the decision cues were applied (t =3.5–4.0 s).

We next show how we corrected for multiple comparisons in the single-cell information analysis using a Holm–Bonferroni multiple-test correction procedure [30]. First, we performed a statistical significance test of the single-neuron MI on 864 neurons from the MPC (medial premotor cortex), 323 from the pre- SMA (presupplementary motor cortex), and 252 from M1 (the primary motor cortex) [12] as follows. To calculate a p-value (for the null hypothesis of no information) for each single cell at a specific time window (in a 200-ms sliding window with 100-ms step), we applied a surrogate-based statistical test. The null hypothesis implemented by the surrogates corresponds to the absence of statistical dependencies between the firing rate of the neuron in a particular window and the behavioral response of the animal. To do this, we generated 200 surrogates by randomly reordering the responses and the firing rates on different trials for each cell. The p-value is calculated by comparing the estimated MI value of each single cell at a specific time window with the empirical distribution of the MI of the corresponding surrogates. The Holm–Bonferroni procedure considers the fact that we have many cells and

Area	Responsive	Significant MI in f_2 and recall period	H-B multiple test	
			k maximum	k minimun
MPC	867	180 (18.2 %)	21 (11.67 %)	2 (1.1 %)
pre-SMA	323	30 (9.0 %)	5 (53.33 %)	1 (3.33 %)
M1	252	34 (11.1 %)	18 (52.9 %)	1 (2.9%)

TABLE 2.1: The proportions of neurons in different areas that retain single cell information in the postponed response delay period when tested with a Holm–Bonferroni correction for multiple comparisons. Column 2 shows the number of neurons responsive in the task. Column 3 shows the number and proportion of neurons that show significant information in both f_2 and in the response period. Columns 4 and 5 show the maximum and the minimum across each of the 200-ms time windows in the delay period of the number of neurons k with significant information after the Holm–Bonferroni test has been applied. The number and proportion of neurons in an area is shown (see Fig. 2.5 for more details). f_2, second stimulus; MPC, medial premotor cortex; pre-SMA, pre-supplementary motor area; M1, primary motor cortex.

therefore multiple null hypotheses to test in a given time window. Let us assume that the significance α level is 0.05. The Holm–Bonferroni procedure consists of ordering the p-values and comparing the smallest p-value to α/k. The first null hypothesis is rejected if the p-value is less than α/k. After this, one tests the remaining $k-1$ null hypotheses by starting again with the same α, i.e., reordering the $k-1$ remaining p-values and comparing the smallest one to $\alpha/(k-1)$. This iteration is continued until the null hypothesis with the smallest p-value cannot be rejected. The result of the Holm–Bonferroni correction is accepting all null hypotheses that have not been rejected at previous steps.

We are interested in finding the percentage of neurons that maintain information about the behavioral response during the postdecision delay relative to those neurons that show significant information during the decision period of f_2 presentation and also during the final behavioral response period. Therefore, we selected the neurons with an MI > 0.26 bits during these periods: from f_2 to 50 ms later and at the response time (end of the postdecisional delay period) in at least three time windows. In that way we obtained 180 neurons (18.2%) from MPC; 30 neurons (9.0%) from pre-SMA; and 34 neurons (11.1%) from M1. The results after the Holm–Bonferroni correction are shown for each of the 200-ms time windows in Fig. 2.5. Table 2.1 shows that in the MPC between 21 and two (depending on the time window) of 180 MPC neurons are able to maintain the (single-cell) information in the delay period.

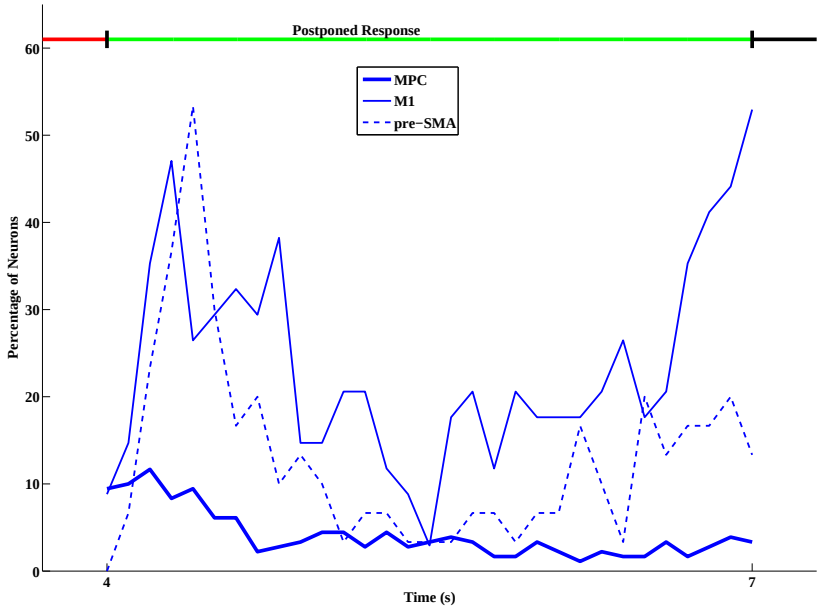

FIGURE 2.5: **Results of the Holm–Bonferroni correction**, for multiple tests for the single-cell information measure in the following areas: MPC, M1, and pre-SMA. The y axis represents the percentage of neurons for each 200-ms time window that have significant information after the correction has been applied during the postponed response delay period from 4 to 7 s. Time course of the task is shown above (Fig. 2.1).

Multiple-cell information analysis

When measuring the information about a set of stimuli S from the responses of many neurons, the response space becomes very large, as there are responses from every neuron to every stimulus. It becomes difficult to record a sufficiently large number of trials to sample this high dimensional space adequately. Rolls et al. [31] introduced a decoding procedure in which the stimulus s' (from the set S) shown on each trial is predicted from the neuronal responses. It is then possible in the low dimensional space of the number of stimuli in the set to compute the MI between actual stimuli s shown on a trial and the predicted stimuli s' based on the neuronal responses of the population of neurons. The MI between the stimuli and the predicted stimuli is then calculated as follows. Bayesian probability decoding using crossvalidation was used to generate a table of conjoint probabilities $P(s, s')$.The s' represents all possible stimuli, and hence belong to the same set S as each stimulus s. The MI value based on this probability decoding (I_p) was calculated as

$$< I_{\mathrm{p}} > = \sum_{s \in S} \sum_{s' \in S} \mathrm{P}(s, s') \log_2 \frac{\mathrm{P}(s, s')}{\mathrm{P}(s)\mathrm{P}(s')}. \tag{2.1}$$

A correction procedure for the sampling bias was used. The percentage correct was calculated using maximum-likelihood decoding in which the most likely stimulus that was shown on each trial is predicted from the neuronal response of all of the neurons on that trial. Examples and fuller descriptions of the use of these procedures are available [22, 31, 40–42].

We performed multiple-cell information analyses to measure how the information about the decision increases with the number of MPC neurons (6) in the sample. We found, taking 18 neurons at random from those with low single-cell information content in a 600-ms window in the delay period (5.2–5.8 s), that the average information per neuron was 0.06 bits, and that with 18 neurons the information provided was 0.51 bits, and 90% correct prediction of the decision (Fig. 2.7). (These 18 neurons had low information even during f_2, on average 0.4 bits/neuron, and it needed 14 such neurons selected at random to reach 1 bit of information during f_2). A further analysis showed that just seven such neurons provided 0.42 bits of information and 90% correct or better performance throughout the delay period.

If we consider 16 randomly selected neurons from those with the higher information values shown in Fig. 2.2, then the multiplecell information analysis showed that the average amount of information for each cell was 0.56 bits in the same 600-ms window in the delay period, and that with subsets of cells chosen at random from the 16 cells the information reached 1 bit and 100% correct with four to six cells. The implication is that with just six of the MPC cells chosen at random from the set with higher information values in the delay period shown in Fig. 2.2, the animal could do the task perfectly, 100% correctly.

2.4.2 Synaptic Facilitation Model

Implementation of the model

The decision-making network is illustrated in Fig. 2.1C, and operates according to the principles described elsewhere [32], with the addition of synaptic facilitation, as specified in the tabular material that follows. The model is biologically realistic and based on an attractor network [33]. The network has four neuronal populations or pools: one inhibitory pool (with NI = 200 neurons) and three excitatory pools or populations (with the total number of excitatory neurons NE = 800), of which one pool is nonselective and the other two are selective and specific for each decision. The selective pools are involved in a competition mediated by

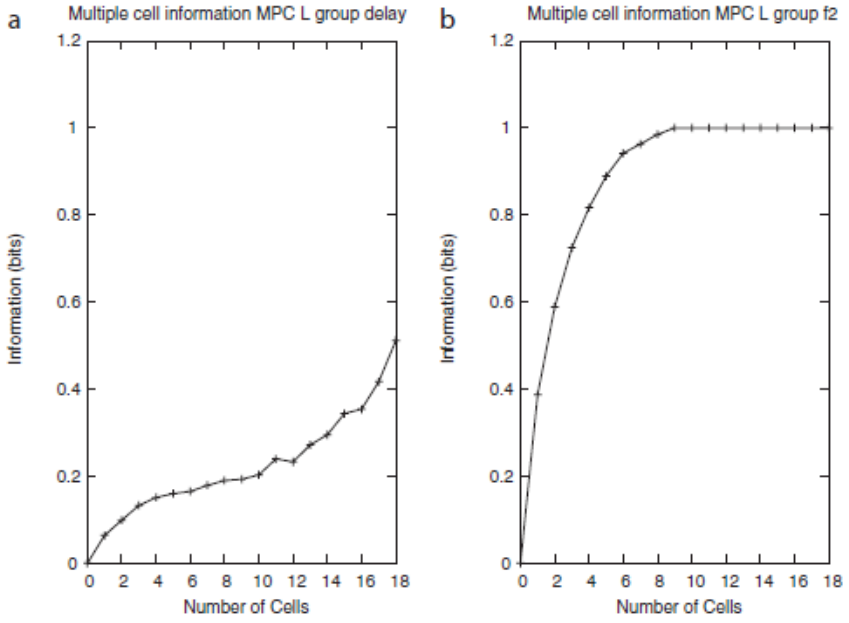

FIGURE 2.6: **Multiple-cell information for the group of MPC neurons with low single-cell information values. (A)** The values for the average information available in the responses of different numbers of MPC neurons on each trial taken in a 600-ms period toward the end of the delay period about which decision had been made. **(B)** The values for the average information available in the responses of different numbers of MPC neurons on each trial taken in a 600-ms period during f2 about which decision had been made. The decoding method was Bayesian probability estimation.

inhibition (inhibitory interneurons), in which only one pool wins (a winner-take-all model). The nonspecific group is connected to the selective pools; likewise, all three excitatory pools are connected to the inhibitory pool. All of the pools are self-connected (recurrent connections). We used integrate-and-fire neurons with three types of receptors mediating the synaptic currents: the excitatory recurrent postsynaptic currents are mediated by AMPA (fast) and NMDA (slow) receptors, and the inhibitory postsynaptic currents to both excitatory and inhibitory neurons are mediated by GABA receptors (see refs. [22, 26, 32] for more details). The external excitatory recurrent postsynaptic currents injected onto the network for λ_1, λ_2, and λ_{unsp} are driven only by AMPA receptors. The parameters used are shown in the following section.

In the simulations, first the network runs with a background external input of 3 Hz to each of the 800 synapses for external inputs onto every neuron, which remains on throughout the simulation. Then in a decision period corresponding to f_2 for t =3.5–4.0 s, each selective pool is driven by a different input, λ_1 and λ_2 respectively. This time symbolizes the first part of the vibrotactile decision-making task, i.e.,

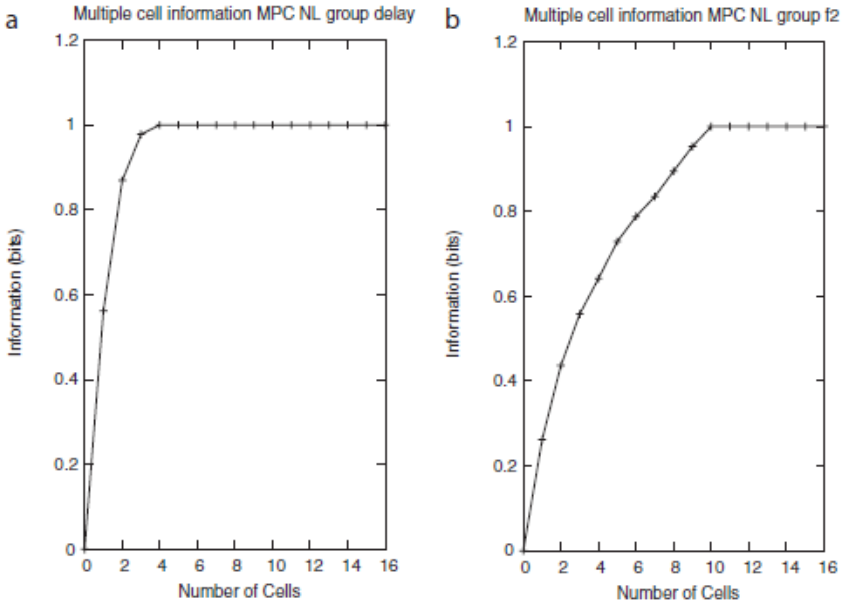

FIGURE 2.7: **Multiple-cell information for the group of MPC neurons with high single-cell information values. (A)** The values for the average information available in the responses of different numbers of MPC neurons on each trial taken in a 600-ms period toward the end of the delay period about which decision had been made. **(B)** The values for the average information available in the responses of different numbers of MPC neurons on each trial taken in a 600-ms period during f_2 about which decision had been made. The decoding method was Bayesian probability estimation.

from the beginning of f_1 to the end of f_2 (Fig. 2.1). The network responds to the external inputs (λ_1 and λ_2 applied in the decision period) by starting the competition between the two selective (decision) pools, the firing rates of which grow apart during this period as one pool moves to a high firing rate attractor level. The pool that reaches the high firing attractor reflects the decision that has been made, and the other selective pool remains firing at around the spontaneous firing rate level. As a result of the calcium-mediated SF, the residual calcium levels of the neurons in the winning selective pool have grown in this decision period, and the probability of spiking has increased [26], as illustrated in Fig. 2.8 *Left*. After this 0.5-s decision-making period, the decision cues are removed and the postponed response delay short-term memory period lasts from 4 to 7 s. In this delay period, because the decision cues have been removed, the firing rates of the two selective pools drop to a spontaneous level of activity (as shownby the rastergram in Fig. 2.3 and by the firing rate in the delay period shown in Fig. 2.4 *Lower Left*). Note, however, that the information in the delay period is still present in the network at the synaptic level but not in the firing rates, as reflected in the information analysis shown in Fig. 2.2 for the subsequent recall period.

Finally, at t = 7.0–7.5 s, both selective pools receive the same extra nonspecific external input (λ_{unsp}), to reflect the moment when the subject receives the stimulus to give its response. The extra nonspecific external rate was 0.255 Hz onto each of the 800 external input synapses of each neuron in a selective pool (with a set of Next = 800 external synapses for external inputs, this results in an additional Poisson background external input of 204 Hz to each neuron in both selective pools). When the external nonspecific input (λ_{unsp}) is injected into the selective pools the report period starts and, because of the altered synaptic calcium levels, the firing rate of one of the selective pools increases to the attractor activity level (Fig. 2.3), as does its information about the response to be made (Fig. 2.4), whereas the firing of the other selective pool remains with low activity, although a little higher than the spontaneous firing rate. The outcome of the competition in the postponed response period (t = 7.0–7.5 s in the simulations) is the report of the task.

Analyzing the computational activity of the SF model by the MI method (Fig. 2.2), we can see how the information becomes significant during the discrimination period, because the MI decodes the network responses to the injected inputs (λ_1 and λ_2 to pool 1 and pool 2, respectively), one higher than the other (in our simulations, $\lambda_1 > \lambda_2$). During the postponed response delay period, the MI in the firing rates about the decision becomes less significant as a result of the reset of the external input, with the variability of the firing rates as well as the low and nondifferential firing rates contributing to the low information at this time. Finally, the information available becomes high again during the report time owing to the λ_{unsp} injected, which acts to increase the rates, but selectively in the pool with the synaptic facilitation, as shown in Figs. 2.2 and 2.3. That is, in the SF model the rates are encoding the response only during two specific periods: after f_2, when they already have all of the sensory information, and when the response is demanded. With this criterion we confirm that the SF network remembers the discrimination and it is able to reproduce it when it is requested at the time of the postponed response without having to store the information about f_1, f_2 or the decision in short-term memory using high firing rates in an attractor. The simulation results are shown in Figs. 2.2–2.4.

In the simulations we consider a trial as correct if the activity of the more stimulated selective pool (with λ_1 = 250 Hz and λ_2= 150 Hz) is higher than the activity of the other selective pool, during all of the bins of the 100-ms period centered at the end of the λ_{unsp} period.

	Postponed response delay, %			
	1 s	1.5 s	2.5 s	3 s
Task performance	100	98.6	92.6	85.5
Model performance	100	99	92	83

TABLE 2.2: The performance of the task [12] versus the performance of the network with different postponed decision delays: 1, 1.5, 2.5 and 3 seconds.

Performance in the task by the subjects, and by the SF model, as a function of the postponed response delay period magnitude.

We tested the performance of the model as a function of the duration of the delay period, because synaptic facilitation is unlikely to be able to maintain a memory trace for long time periods, and it is a prediction of the model that performance will decay to zero as the short-term memory period increases much beyond the time constant of the synaptic facilitation, 2 s. We performed four simulations with different postponed response delays. In Table 2.2 we compare the recorded experimental data [12] with the performance of the SF model. Our model fits the data well. Analyzing the MI of the different postponed decision delays, we found that the maximum MI value at the end of the decision period in the network is 0.967 $\pm$ 0.001 bits, and it does not depend on the delay period, whereas the maximum MI value at the end of the nonspecific stimulation is not always the same, but depends on the delay period. The longer the delay period, the lower is the MI value: 0.95, 0.91, 0.63, and 0.37 bits for 1-, 1.5-, 2-, and 3-s postponed response delays, respectively.

Magnitude of the synaptic facilitation u in the four pools as a function of time for the synaptic facilitation model, and a mean-field analysis

For the SF model, the time evolution of u, the synaptic utilization, for the four different pools is shown in Fig. 2.8 *Left*.

The results of a mean field analysis [43, 44] for two different scenarios for the synaptic facilitation model are shown, one with synaptic facilitation (Fig. 2.8, *Lower Right*) and one without synaptic facilitation (Fig. 2.8, *Upper Right*). We plotted the value for the firing rate difference between pools 1 and 2 for a fixed value of λ_{unsp}, w_+ and U, for two values of U (one for each scenario, 0.15 and 0, respectively) to check that the synaptic facilitation network cannot recover the information after a delay period without synaptic facilitation. To compare both ranges of parameters λ_{unsp}, w_+, we performed the mean field analysis by

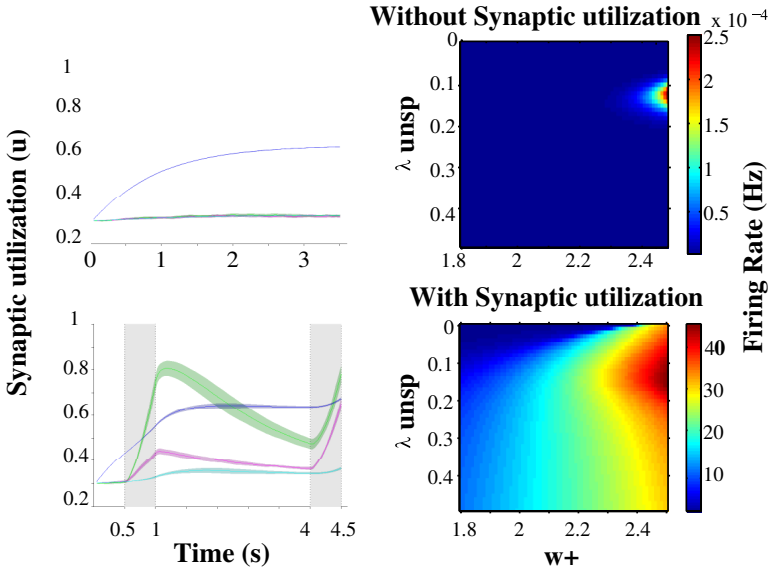

FIGURE 2.8: **The time evolution of the variable** *u(Left)*, the synaptic utilization, for the four different pools. *(Upper Left)* The network without stimulation. *(Lower Left)* the network with stimulation. The shaded gray rectangles show the period of stimulation; blue, the inhibitory pool; green, selective pool 1; pink, the selective pool 2; cyan, the nonselective pool. The plots are an average of >50 trials. The dark lines are the mean, and the colored shadows are 1 SD. *(Right)* Results of the mean field analysis [28, 43]. Each point represents the firing rate difference between pool 1 and pool 2 for a fixed value of (λ_{unsp}, w_+) and U. We used the values of the mean synaptic facilitation (one for each pool) at the time instant of 7,500 ms extracted from the simulation shown in the above plots.

multiplying the specific connectivity weight by the two specific sets of U, one for each scenario. To obtain the set that corresponds to a nonfacilitated network, we ran the network without any stimulation, and we took the U values after the network was stable. To obtain the set that corresponds to a facilitated network, we took the U values of the simulation at the time value of 7.5 s, and show the results in Fig. 2.8. We show that there is selective firing at the time of recall in this model with these parameters only when SF is present.

2.4.3 Graded Firing Rates in an Attractor Network Model

The attractor network used was similar to that described for the SF network, except that no synaptic facilitation was used, and the synaptic weights in the intrapool connections for the selective pools 1 and 2 were set to an exponential-like distribution to produce an exponential-like firing rate distribution, as described below and in more detail by [28].

Graded weight patterns

In an attractor network, the synaptic weights of the recurrent connections are set by an associative (or Hebbian) synaptic modification rule with the form

$$\delta w_{ij} = \alpha r_i r_j, \tag{2.2}$$

where δw_{ij} is the change of synaptic weight from presynaptic neuron j onto postsynaptic neuron i, α is a learning rate constant, r_j is the presynaptic firing rate, and r_i is the postsynaptic firing rate when a pattern is being trained [22, 45, 46]. To achieve this for the firing rate distributions investigated, we imposed binary and graded firing rates on the network by selecting the distribution of the recurrent synaptic weights in each of the two decision pools. To achieve a binary firing pattern, used for the SF simulations described, all of the synaptic weights between the excitatory neurons within a decision pool were set uniformly to the same value w_+.

Graded firing patterns were achieved by setting the synaptic weights of the recurrent connections within each of the decision pools to be in the form of a discrete exponential-like firing rate (r) distribution generated using methods taken from [31]:

$$P(r) = \begin{cases} \frac{4}{3} a\lambda e^{-2(r+r_0)} & \text{for } r > 0 \\ 1 - \sum\limits_{r_i \in r : i > 0} \frac{4}{3} a\lambda e^{-2(r_i+r_0)} & \text{for } r = 0 \end{cases} \tag{2.3}$$

where a is the sparseness of the pattern defined in Eq. 2.4, and r_0 isthe firing rate of the lowest discretized level. The population sparseness a of a binary representation is the proportion of neurons active to represent any one stimulus or decision in the set. The sparseness can be generalized to graded representations, as shown in Eq. 2.4;

$$a = \frac{(\sum\limits_{i}^{N_E} r_i)^2}{\sum\limits_{i}^{N_E} r_i^2}, \tag{2.4}$$

where r_i is the firing rate measured for neuron i in the population of NE excitatory neurons in the network [22]. We note that this is the sparseness of the representation measured for any one stimulus over the population of neurons [22, 35].

In the graded firing rate simulations, we use $a = 0.1$ to correspond to the fraction of excitatory neurons that are in a single decision pool. We chose 50 equal-spaced

discretized levels to evaluate the distribution $(0, \frac{1}{3} - r_0, \frac{2}{3} - r_0, ..., 3 - r_0)$. r_0 and λ are chosen so that first and second moments of the firing rate distribution are equal to the sparseness, i.e., $< r >=< r^2 >= a$. A weight matrix $W = \{w_{1,1}, \cdots, w_{1,fN_E}, w_{2,1}, \cdots, w_{fN_E,fN_E}\}$ was constructed by first sampling a firing rate for each neuron, r_i, using Eq. 2.3 and then setting w_{ij} based on the desired firing rates of each pair of neurons, as described in more detail by Hopfield [28]. The mean weight was set to a value close to 2.1, the maximum weight was ~ 5.0, and the minimal weight was 1.0. The final mean weight used for the simulations was 2.04, as this provided for satisfactory stability of the network in the spontaneous period, because stability is reduced by graded compared with binary firing rates [28].

Graded firing rate simulation protocol

Our focus was on the activity in the delay period after the decision had been taken, and on whether a low level of firing in the delay period could be restored to a high level, with a high information content, at the end of the delay period when an external input was applied equally to the two decision pools, 1 and 2. In a 0.5-s period of spontaneous firing from 3.0 to 3.5 s, the external rates λ_1 and λ_2 were 3.00 Hz applied to each of the 800 external synapses on each neuron. In the decision period from 3.5 to 4.0 s, λ_1 was 3.10 Hz and λ_2 was 2.98 Hz per external synapse, values that produced a decision on almost all trials of pool 1 winning. In the delay period from 4 to 7 s, λ_1 and λ_2 were set to the lowest value that enabled firing to be maintained reliably (although at a low level) by some neurons, 2.95 Hz per synapse. During the recall period from 7.0 to 7.5 s, λ_1 and λ_2 were set to the identical value of 3.05 Hz per synapse.

2.4.4 Model Summary

A	Model Summary
Populations	Two: excitatory, inhibitory
Topology	—
Connectivity	Full connected
Neuron model	Leaky integrate-and-fire, fixed threshold, fixed refractory period, MNDA
Channel models	—
Synapse model	Instantaneous jump and exponential decay for AMPA and GABA and exponential jump and decay for NMDA receptors
Plasticity	Synaptic facilitation
Input	Independent fixed-rate Poisson spike trains to each selective population
Measurements	Spike activity

B	Populations	
Total number of neurons	$N = 1000$	
Excitatory neurons	$N_E = 0.8 \cdot N$	Neurons in each selective pool: $N_{selective} = N_E \cdot sparseness$
Inhibitory neurons	$N_I = 0.2 \cdot N$	

C1	Neuron and Synapse Model
Type	Leaky integrate-and-fire, conductance-based synapses
Subthreshold dynamics	$$C_m \frac{dV(t)}{dt} = -g_m(V(t) - V_L) - I_{\text{syn}}(t)$$ $$I_{\text{syn}}(t) = I_{\text{AMPA,ext}}(t) + I_{\text{AMPA,rec}}(t) + I_{\text{NMDA}}(t) + I_{\text{GABA}}(t)$$
Spiking	If $V(t) > V_\theta \wedge t > t^* + \tau_{rp}$ 1. set $t^* = t$ 2. emit spike with time-stamp t^* 3. $V(t) = V_{reset}$
Synaptic currents	$$I_{\text{AMPA,ext}}(t) = g_{\text{AMPA,ext}}(V(t) - V_E) \sum_{j=1}^{N_{\text{ext}}} s_j^{\text{AMPA,ext}}(t)$$ $$I_{\text{AMPA,rec}}(t) = g_{\text{AMPA,rec}}(V(t) - V_E) \sum_{j=1}^{N_E} w_j s_j^{\text{AMPA,rec}}(t) u_j(t)$$ $$I_{\text{NMDA}}(t) = \frac{g_{\text{NMDA}}(V(t) - V_E)}{1 + \gamma \exp(-\beta\, V(t))} \sum_{j=1}^{N_E} w_j s_j^{\text{NMDA}}(t) u_j(t)$$ $$I_{\text{GABA}}(t) = g_{\text{GABA}}(V(t) - V_I) \sum_{j=1}^{N_I} s_j^{\text{GABA}}(t)$$

C2	Neuron and Synapse Model
Fraction of open channels	$$\frac{ds_j^{\text{AMPA,ext}}(t)}{dt} = -s_j^{\text{AMPA,ext}}(t)/\tau_{\text{AMPA}} + \sum_k \delta(t - t_j^k - \delta)$$ $$\frac{ds_j^{\text{AMPA,rec}}(t)}{dt} = -s_j^{\text{AMPA,rec}}(t)/\tau_{\text{AMPA}} + \sum_k \delta(t - t_j^k)$$ $$\frac{ds_j^{\text{NMDA}}(t)}{dt} = -s_j^{\text{NMDA}}(t)/\tau_{\text{NMDA,decay}} + \alpha\; x_j(t)(1 - s_j^{\text{NMDA}}(t))$$ $$\frac{dx_j(t)}{dt} = -x_j(t)/\tau_{\text{NMDA,rise}} + \sum_k \delta(t - t_j^k - \delta)$$ $$\frac{ds_j^{\text{GABA}}(t)}{dt} = -s_j^{\text{GABA}}(t)/\tau_{\text{GABA}} + \sum_k \delta(t - t_j^k - \delta)$$
Synaptic facilitation	$$\frac{u_j(t)}{dt} = \frac{U - u_j(t)}{\tau_F} + U(1 - u_j(t)) \sum_k \delta(t - t_j^k)$$

D	**Input**
Type	**Description**
Poisson generators	Fixed rate N_{ext} synapses per neuron, which each synapse driven by a Poisson process

E	**Measurements**
Spike activity	

Parameters

Integrate-and-fire			
C_{m} (E)	0.5 nF	C_{m} (I)	0.2 nF
g_{m} (E)	25 nS	g_{m} (I)	20 nS
V_{L}	-70 mV	V_{thr}	-50 mV
V_{reset}	-55 mV	V_{E}	0 mV
V_{I}	-70 mV		
$g_{\text{AMPA,ext}}$ (E)	2.08 nS	$g_{\text{AMPA,rec}}$ (E)	0.104 nS
g_{NMDA} (E)	0.327 nS	g_{GABA} (E)	1.25 nS
$g_{\text{AMPA,ext}}$ (I)	1.62 nS	$g_{\text{AMPA,rec}}$ (I)	0.081 nS
g_{NMDA} (I)	0.258 nS	g_{GABA} (I)	0.973 nS
$\tau_{\text{NMDA,decay}}$	100 ms	$\tau_{\text{NMDA,rise}}$	2 ms
τ_{AMPA}	2 ms	τ_{GABA}	10 ms
τ_{rp}(E)	2 ms	τ_{rp}(I)	1 ms
α	0.5 ms^{-1}	γ	$[\text{Mg}^{2+}]/(3.57\text{mM}) = 0.280$
β	$0.062\,\text{mV}^{-1}$	sparseness	0.10
N_{ext}	800		
U	0.15	τ_{F}	2000 ms
Connection			
w_+	2.17	w_i	0.97

TABLE 2.3: Parameters used in the integrate-and-fire simulations and the connection parameters. Capital letters: E corresponds to *excitatory*, and I to *inhibitory*.

Chapter 3

Decision-making under different brain states in A1

The work presented in this chapter was published in:
Frontiers in Integrative Neuroscience in October 2011
Journal of Neurophysiology in August 2013
In both cases, the authors are: Juan M Abolafia*, Marina Martinez-Garcia*, Gustavo Deco, and Maria V. Sanchez-Vives.

3.1 Introduction

In this chapter we present the study of a data-base of neurons from rat auditory cortex (A1) while they performed a decision-making task. The animals had to decide whether two auditory stimuli were separated by either 150 or 300 ms (see Fig 3.1). The special feature of this study is that the neurons were recorded under different brain states: during engaged (correct choices provided reward) and during idle states (performance not required).

In the first part of this chapter, we present a study on the firing patterns of the neurons across brain states. In particular, we wanted to characterise and interpret the modulations of the rate during the inter stimuli interval (ISI). The slow modulatory components could be locally generated or the result of a top-down influence originated in higher associative association areas. Such a neuronal discharge may be related to the computation of the interval time and contribute to the perception of the auditory stimulus.

In the second part of this chapter, we focus on the variability and the information content of the neurons. As a measure of the variability we used the Fano Factor $\left(Ff = \frac{\sigma^2}{\mu}\right)$. Which was markedly reduced, not only during stimulation, but also in between stimuli in engaged trials. Next, we explored if this decrease in variability was associated with an increased information encoding. Our information theory analysis revealed an increase of the information content in the auditory responses during engagement as compared to idle states, in particular in the responses to task-relevant stimuli. Altogether, we demonstrated that task-engagement significantly modulates the coding properties of auditory cortical neurons during an interval-discrimination task.

3.2 Task description

In this section we define with more details the interval-discrimination task described in the Introduction 1.1. The behavioural protocol consisted of four different recording stages with a total duration of 2.5 hours (Fig. 3.1 A). For the stages that required it, the animals were trained to poke their noses into the center socket (see 3.1 B), which immediately triggered the onset of two identical stimuli (80dB, 5322 Hz, 50 ms duration). The end of stimulus one ($S1$) was separated from the beginning of stimulus two ($S2$) by 150 or 300 ms. In a chronological order the different stages were:

- The "Initial-idle" stage (17 min) was done during idle listening of the animal. The auditory stimuli (50 ms, 80dB; 5322Hz), ISI, intertrial interval (2-3 s) and trial repetitions (180) were the same in each recording stage. During this recording the animals freely moved around the recording box (with occluded sockets) while listening to stimuli presentation, and they did not receive reward.

- The "Engagement-task" was the behaviourally relevant one. During the engagement stage the rat enters the central socket and two identical stimuli are presented through earphones. The animals had to remain in the center socket till the end of the stimuli presentation. They had to discriminate whether the two stimuli were separated by 150 or 300 ms. This required a left or right poke respectively, in order to get the water reward. In the behavioural task, false alarms (poking in the opposite side) or early withdrawals (withdrawal before stimuli termination) were punished with a 3 s time out and white noise (WAV-file,0.5 s, 80 dB SPL).

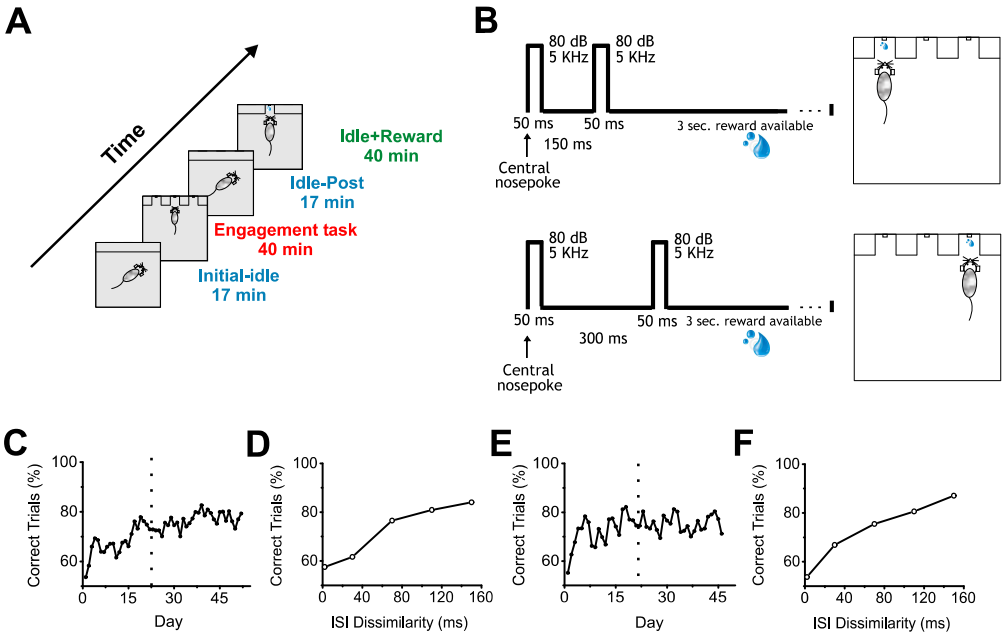

FIGURE 3.1: **Behavioural protocol and performance.** **A:** Sequence of recording stages in chronological order. The "initial-idle" stage was done during idle listening of the animal. The auditory stimuli (50 ms, 80dB; 5322Hz), ISI (150 and 300 ms), intertrial interval (2-3 s) and trial repetitions (180) were the same in each recording stage. The following stage was the behaviourally relevant one. Next, there was again a "idle-post" recording identical to the "initial-idle" one, and another idle recording but now followed by a reward after each pair of stimulus presentation. The total duration of the recording protocol was 2.5 hours. **B:** In the interval-discrimination task the rat entered in the central socket and two identical stimuli (50 ms; 80dB; 5322Hz) were presented through earphones. 150 or 300 ms ISI indicates left or right reward delivery, respectively. **C:** Animal 1 performance (correct trials (%)) against days of training. Dashed line indicates beginning of recorded sessions. **D:** Psychometric curve of performance for pairs of intervals differing between 2 and 150 ms. within the same session. Performance improved as the difference between both ISIs (ms) increased, while discrimination became more difficult for highly similar intervals.**E-F:** Same as C-D but for animal 2. Animal 1 and 2 were trained before implanting microdrives with tetrodes. Learning performance in animals 1 and 2 shown in C-F corresponds to short and long ISIs appearing at random trials.

- The "Idle-post" (17 min) recording was identical to the "initial-idle" one.
- The last stage was the "Idle+reward" recording. During this stage, the left and right sockets were occluded but not the central one, where the animal repeatedly entered and received a water drop 0.3 s after having listened to the same stimuli as in the engagement stage.

Note that there was movement in all phases of the experiment, either towards the sockets or around the cage. The aim of these recording sequences was to track the activity of a single unit, and compare response patterns between engaged versus idle brain states. The three idle stages (initial-idle, idle–post and idle+reward) had the same amount of trials (180-200 trials each), stimuli (50 ms; 80 dB; 5.3 kHz), ISIs (150 and 300 ms) and intertrial intervals (2 to 3 s) as the task-engagement stage. As soon as the animals reached 70% of behavioural performance (Fig. 3.1 C,E), they were implanted. After the last recording we obtained a psychometric curve to further evaluate the perceptual and behavioural effects of the short and long ISIs being presented to the animal during the behavioural task. During the psychometric curve the short ISI (150 ms) became longer, while the long ISI (300 ms) became shorter, allowing us to test the perceptual threshold of ISI discrimination (Fig. 3.1 D, F).

3.3 Study of the data-base: Slow Modulations

In the present study we recorded the activity of eighty-six neurons from the auditory cortex of the behaving rat. The aim of this section is to explore the slow modulation of neuronal activity in the intervals between stimuli while the rat was performing an interval-based decision-making task.

In principle, sensory areas such as A1 are primarily associated with stimulus encoding. However there is increasing evidence that early cortices, and in particular A1, are not only feature detectors. Multimodal responses [47, 48], attentional modulation [49–51], expectation [52], or reward-modulation [53] illustrate additional contextual aspects that modulate responses even in early sensory cortices.

While fast responses to auditory stimuli have been characterized in detail in auditory cortex, the slow modulation of neuronal firing to evoked and spontaneous activity has barely been studied. Slow modulation of sustained responses has been found to predict the behavioral decisions during auditory categorization tasks in

monkeys, including errors [54]. Sensory or behavioral events contingent on reinforcement can also result on slow modulation of firing rate or sustained firing as a consequence of a learning process [55]. Therefore, slow modulation of firing could constitute an anticipatory mechanism that associates events (stimulus–behavior–reinforcer) that are relevant or adaptive to the environment. These cognitive components associated to stimulus discrimination tasks have been more commonly associated to higher areas such as frontal areas [56].

3.3.1 Description of the singular patterns

Eighty-six single units from the rat auditory cortex were isolated and classified [57, 58] according to their phasic auditory responses as: onset (26%), onset + offset (13%), offset (2%), non-responsive (43%), suppressive (13%), and "other" (3%). The percentage of non-responsive neurons was similar to the one reported by means of cell-attached recordings in the head-fixed awake animal [58].

In this section we describe different patterns of neuronal discharges occurring during the intervals (ISI) in between stimuli and thus in the absence of auditory stimulation. We show that the neuronal firing occurring between auditory stimuli in the auditory cortex can be quite prominent and that is often modulated by attention. Finally, the possible functional role of this slow modulation of neuronal discharge is discussed.

We found that 17.4% (n = 14) of the recorded neurons showed a prominent neuronal discharge during the interstimulus interval, in the form of either an upward or downward ramp towards the second auditory stimulus. While in some of the neurons the activity during the interval ramped up toward the second stimuli (n = 6; Figs. 3.2 A, B, and 3.3 A, B, C), in others the activity ramped down following a sort of post-discharge (n = 6; Figs. 3.3 D, E and 3.4 A, B). In the two remaining neurons the activity during the interval remained rather in a plateau (Fig. 3.3 F). The neurons shown here further illustrate the large heterogeneity of neuronal responses that have been described in auditory cortex.

The Fig. 3.2 illustrates the PSTHs from two different neurons while the rat was performing the task. During the passive sound stimulation, the neuron in Fig. 3.2 A had a weak offset response to the first auditory stimulus and a subsequent decrease in the firing during the interval, that progressively increased towards the second stimulus. During attention these responses became more prominent (Fig. 3.2 A; *top* PSTHs). The offset response was increased, and neuronal activity

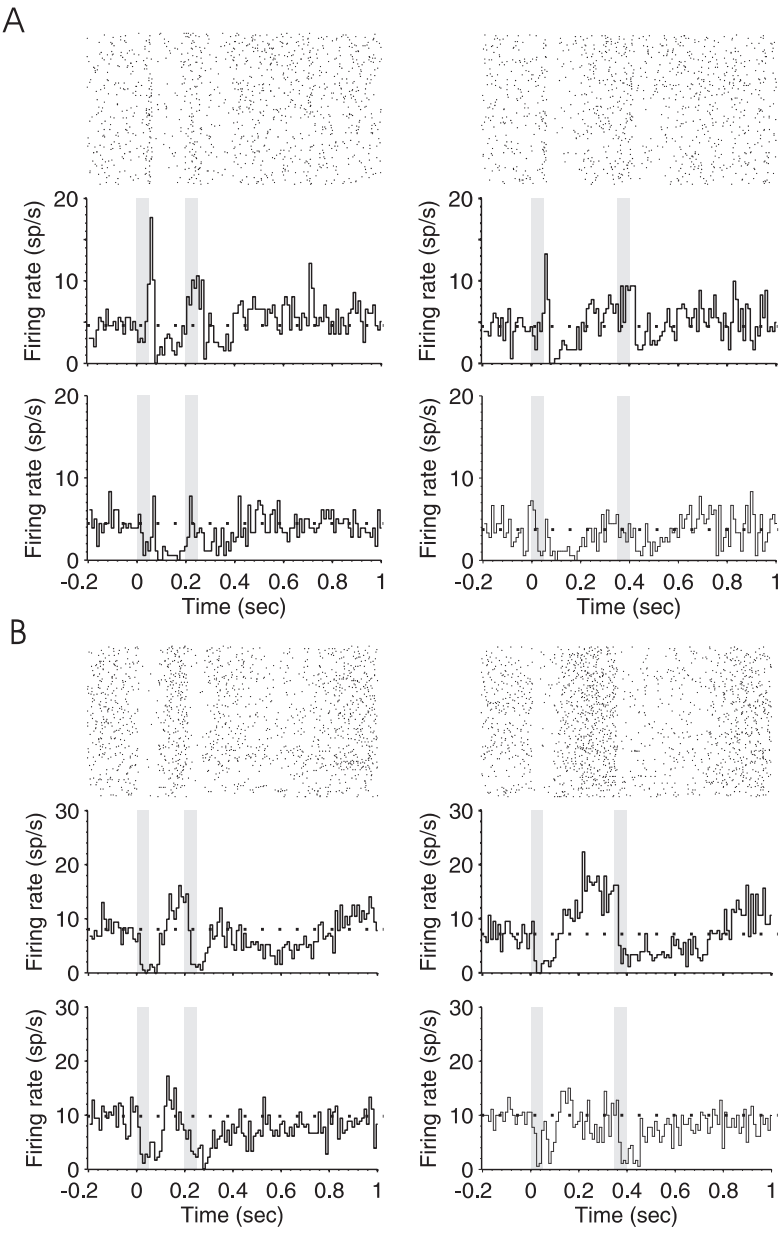

FIGURE 3.2: **Ramping-up activity during the interstimulus interval. A:** Upper part. PSTH (180 trials) of the response (*bottom*) and raster plot (*top*) of a single neuron during the attentive task shows the response pattern to two identical stimuli (50 ms; 80 dB; 5322 Hz) separated by 150 ms (*left*) and 300 ms (*right*). The gray boxes correspond to the periods of auditory stimulation. The PSTH shows an offset response to stimulus 1 and onset response during stimulus 2 while in the passive brain state (*lower part*) there is an overall decrease of excitability and no monotonic increase of spontaneous activity toward $S2$ presentation. **B:** Same as in (A), but this other neuron is silenced by the auditory stimulation followed by an increase of firing until $S2$ occurrence (*upper part*), while in the passive brain state (*lower part*) the neuronal firing is markedly reduced. In (A, B) the mean spontaneous activity is represented with a dotted line.

ramped up toward the second stimulus well above the spontaneous activity preceding the first stimulus. Interestingly, the response to the second stimulus was not an offset response but a sustained one. This is the case for both the short and the long interval trials, which were randomly given. The neuron illustrated in Fig. 3.2 B is of a different type, a "suppressive" response [58], since its discharge was silenced by auditory stimulation. This is well appreciated in the raster plots that correspond to the attentional trials (Fig. 3.2 B, top PSTHs). Still, even when the neuron was silenced by the auditory stimulation, its activity ramped up toward the second stimulus, more prominently in the attentional trials than in the passive ones. The second auditory stimulus again decreased its firing rate, which remained decreased for 200 ms following stimulation.

Out of the six neurons with increasing activity toward the second stimulus, all of them had an up-regulation of this activity during attentive trials. When the average firing rate during the first half of the interval was compared against that during the second, the activity increased in a 17% in passive trials and 246% in attentive ones for the short (150 ms) intervals. For the long (300 ms) intervals these values were 58 and 192% respectively.

In some cases, the activity occurring in between auditory stimuli was not ramping-up toward the second stimulus, as the one illustrated above, but rather appeared as a prominent post-discharge following the auditory stimulation (Figs. 3.3 D, E and 3.4 A) In the neuron illustrated in Fig. 3.4 A, each auditory response was followed by a post-discharge lasting around 200 ms. In this neuron, not only the auditory responses but also the auditory post-discharge was significantly increased by attention. A total of five neurons showed a similar modulation by attention, the post-discharge increasing an average of 45% (short ISIs) and 53% (long ISIs) in attentive versus passive trials. In one neuron, the post-discharge was decreased in a 40% by attention. In the case of the neuron illustrated in Fig. 3.4 A, the firing rate during the 200-ms preceding the first auditory stimulus was also significantly increased by attention. This is the period of time that takes place when the animal is heading to the central nose poke that triggers stimulus presentation.

The firing rate during the period preceding auditory stimulation was also significantly increased during attentive trials in the neuron displayed in Fig. 3.4 B, which on the other hand had a rather different auditory response. This neuron had a weak spontaneous discharge preceding the auditory stimulation, and no response to the auditory stimulus. However, a very large post-discharge followed each auditory stimulus. This unusual pattern of response took place during non-attentional trials. During attentional trials, those prominent post-discharges disappeared, and

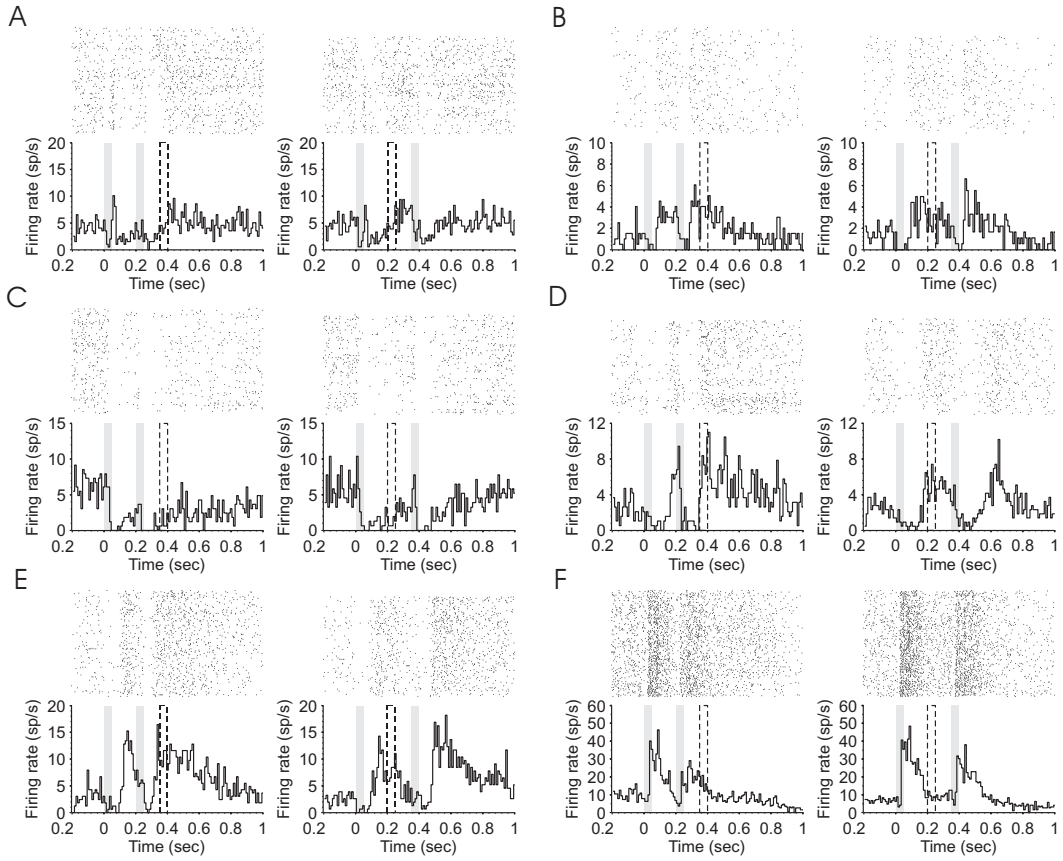

FIGURE 3.3: **Modulation of neuronal firing between stimuli presentation. A–C:** PSTH (180 trials) of two equal stimuli (50 ms; 80 dB; 5322 Hz) separated by either 150 ms (*left*) or 300 ms (*right*). Three example neurons during the attentive task show suppressive activity during stimuli presentation followed by an increased firing until $S2$ is presented. **D, E:** PSTHs of two different neurons showing firing activity during the interval that decreases toward the second stimulus. **F:** PSTH of an example neuron showing a late onset response accompanied by a sustained activity after stimuli termination.

instead, the discharge preceding the first stimulus was increased, as did the example in Fig. 3.4 A. An enhanced firing rate preceding the occurrence of the first stimulus could be related to stimulus expectancy [52] or to prediction of reward [53], both described in early sensory cortices.

3.3.2 An interpretation of the patterns

In the previous section we report about 14 particular neurons that showed prominent responses during the intervals between stimuli, with firing rates that either increased or decreased toward the second stimulus. These neuronal discharges could be referred to as spontaneous activity, since they occurred while there was no auditory stimulation. However the term "spontaneous activity" has been avoided

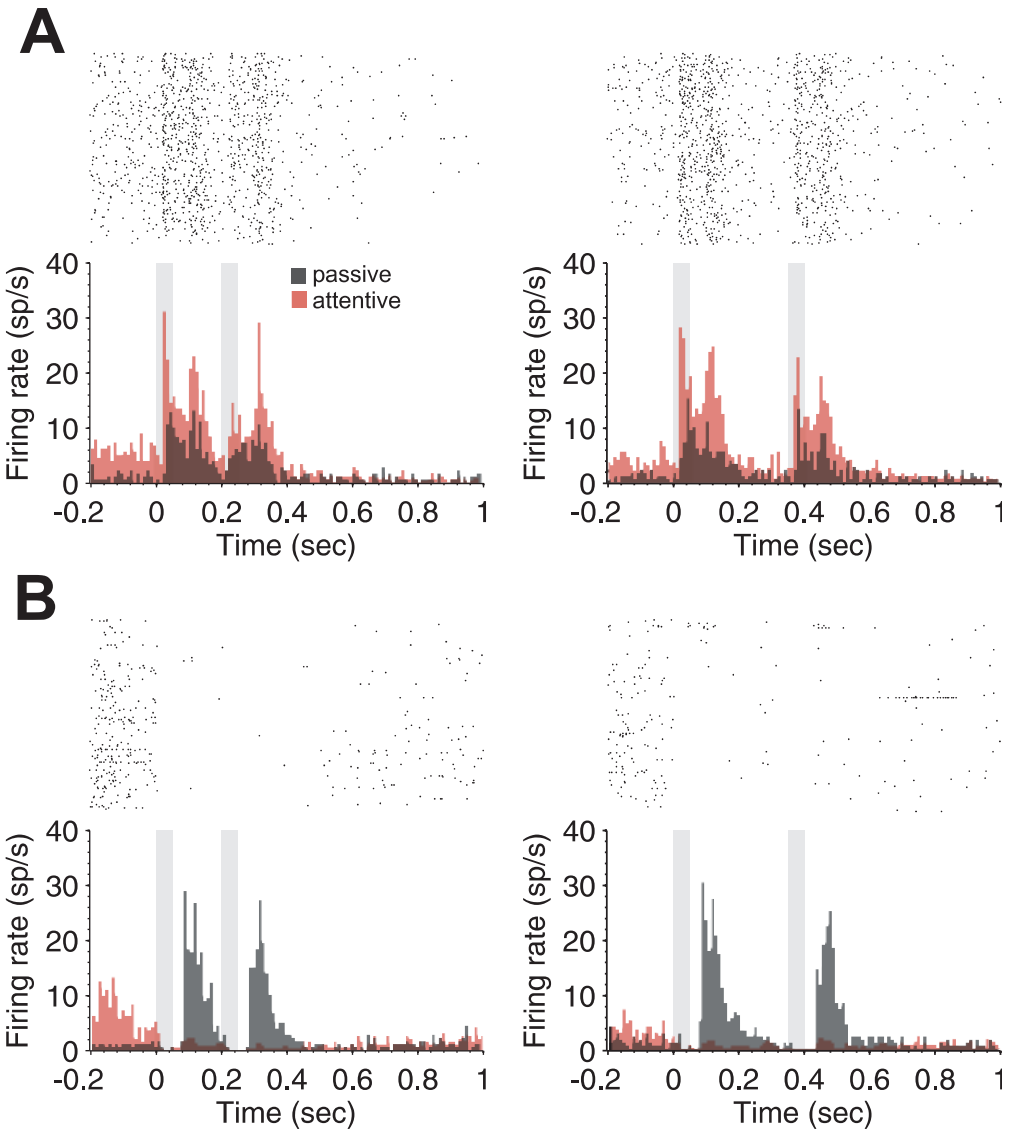

FIGURE 3.4: **Attention modulated firing rate in the periods in between auditory stimuli. A:** PSTH (180 trials) of a single neuron during the passive (gray) and attentive (red) task where two identical stimuli (50 ms; 80 dB; 5322 Hz) separated by 150 ms (*left*) and 300 ms (*right*) were delivered. Notice that not only the response but also prominent responses in the absence of stimuli (preceding and between stimuli) are up regulated by attention. **B:** Example neuron that had no auditory response, but a prominent and slow response after the termination of auditory stimuli during passive trials. This response was silenced during the attentive trials. However, notice that the activity preceding the first auditory stimulus was increased during attention (in red). In both cases (A, B), the raster plots correspond to the attentive trials. In all cases, 150 and 300 ms intervals were given randomly during the same sessions.

since these neuronal discharges were often associated to the preceding auditory responses, even if with a slow time course of over 150 ms. The neuronal discharges occurring in the absence of auditory stimulation were enhanced by attention in 12 (out of 14) cases, while they were decreased in the remaining two.

Most of the studies of the effects of attention on auditory responses have focused on how phasic responses modulate their response properties according to the brain state [51, 59–65]. Slow modulation of firing rate in the auditory cortex of the behaving monkey has been previously found to be related to the processing of stimuli, motor decision, or even reward [54, 55, 64], as it has been in primary visual cortex [53]. During behavioral experiments, this slow or sustained (up to several seconds) part of the response is related to event sequences during a task and provides a neuronal mechanism for anticipation and association of events related to hearing and relevant to behavior [55]. Altogether, slow modulation of firing could complement the representation of the timing of auditory stimuli as well as the codification of stimuli by means of phasic responses. A similar pattern was reported by [50, 64, 66].

Some studies have found no changes in spontaneous activity under attentional demands [51, 60, 61]. On the other hand, an increase in spontaneous firing rate at the end of the trial under attention with respect to the passive state has been reported, enhancement that could be reflecting motor-related aspects [67]. Single units from auditory cortex have also been shown to have enhanced sustained responses preceding a target stimulus [68]. Here we have shown that the spontaneous discharge is increased by attention in the period preceding the first stimulus in two neurons (Figs. 3.4 A, B).

The mechanisms for these slow modulations of firing rate are not known. One possibility would be that they reflect top-down modulation. Not only cortical, but also subcortical areas present modulation of spontaneous activity within tasks. Late trial neuronal activity in the monkey inferior colliculus has been described to be modulated by context, like a "reward expectation" signal [69]. Reward-modulation of the late activity after the end of the auditory stimulus has also been described in the rat auditory thalamus [70]. A difference of these ramping activities with respect to the ones we have illustrated (Fig. 3.2 is that the ramping-up here was preceding the second stimulus, and not the reward [64, 71]. The reward in our protocols occurred after the second stimulus, whenever the animal poked his nose in the correct side and thus triggered its delivery. It did not occur at a fixed time (usually after 1 s in the illustrated PSTHs). The ramping activity illustrated in Fig. 3.2 B between 0.6 and 1 s could be interpreted as such or associated to motor activity. We can speculate that the ramping-up activity in between stimuli

(Figs. 3.2 A, B) could be rather associated to stimulus expectation or to interval computation. In this respect, a recent study [52] showed that neurons from the primary auditory cortex increased their firing rate as the target approached. This firing rate modulation reflected a temporal expectation which improved sound processing, therefore increasing the probability of obtaining a reward.

In all, neuronal firing in early auditory cortex in the absence of auditory stimulation could provide a neuronal mechanism for anticipation and memory, reflecting a learning process where consecutive sensory and behavioral events are associated with reinforcement. The slow modulation of ongoing firing during the interval between stimuli and the post-stimulus period could act as a mechanism to track and integrate time between stimuli presen-tations and be part of the neuronal basis of interval-categorization by means of tonic firing, particularly in attentive stages.

3.4 Variability and information content in A1

In this section our main objective was to study the effects of task-engagement on firing variability not only on evoked responses during correct trials but also during the interstimulus interval (ISI).

3.4.1 Engagement diminishes variability during and after stimulation in A1

Both enhanced [72–74] and reduced responses [51] to stimulation have been observed while processing behaviourally relevant auditory stimuli. We first explored the effect of engagement on the firing rate of 33 neurons that comprised the onset and onset-offset ones. In most cases engagement significantly increased the spike firing (see section 3.6) during stimuli-evoked responses (n=22; Fig. 3.5, while the opposite trend was less common (n=11). The response to the second stimulus was typically decreased as a result of auditory adaptation processes [51, 75, 76]. Substantial adaptation was observed both in onset responses (Fig. 3.5 A) as well as offset ones (Fig. 3.5 B) (n=13). There was a trend for the average adaptation to the second stimulus to be lower in the engaged than in the idle state (n=10), although the difference was not statistically significant (p-value=0.23 and p-value= 0.37 for short and long ISIs respectively; Wilcoxon; Fig. 3.5 C). In the remaining three cases, adaptation increased during engagement.

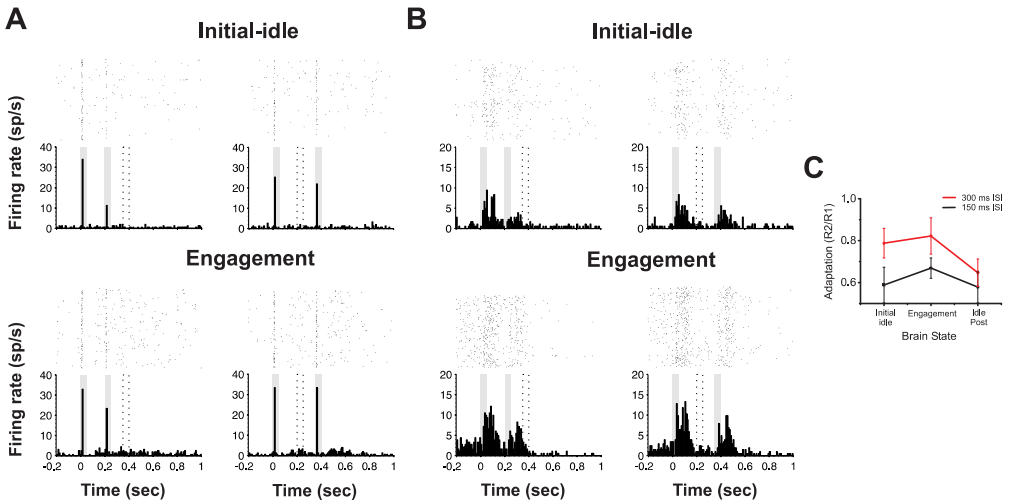

FIGURE 3.5: **Firing rate during the interval-discrimination task in different brain states. A:** Raster and persistimulus histogram (PSTH) of an example neuron for different brain states. The frequency response histogram (*bottom*) using 10-ms bins and the spikes raster plot (*top*) for 180 trials. A response peak can be observed at the onset of stimuli (50 ms; 80 dB; 5322 Hz) for 150 ms ISI (*left*) and 300 ms ISI (*right*). Grey bars indicate stimuli presentation, and dotted bars indicate stimulus presentation if the opposite ISI would have occurred. **B:** Same as in A for a different neuron. A specific increase of spontaneous activity previous to stimuli presentation is observed only during task-engagement. **C:** Average adaptation of thirteen neurons for short and long ISIs for different brain states. Error bars are S.D.

Next we studied whether engagement altered neuronal response variability in the auditory cortex of the behaving animal. We calculated the Fano factor (Ff) (spike-count variance divided by spike-count mean) in order to test how neuronal variability changes during and after stimulus-evoked responses as a function of the behavioural state of the animal. The Ff was calculated on a trial by trial basis. Since the evoked auditory responses were mostly phasic we found that ms bins showed the best time resolution and reflected the most accurately the changes in variability.

Fig. 3.6 shows, for two single units (A, B and C, D), the average Ff variation for short (A, C) and long (B, D) ISIs, and for the engaged state (red) versus idle states (blue). A significant reduction (see section 3.6) in Ff during task-engagement can be observed during stimulus presentation ($S1$ and $S2$) for short (Fig. 3.6 A) and long (Fig. 3.6 B) ISIs with respect to the idle ones. A second neuron with a prominent offset response is illustrated in (Fig. 3.6 C, D). This neuron had a reduction in Ff during the response onset and offset during engagement with respect to the idle state.

The variability in responses to stimuli was decreased not only during engagement,

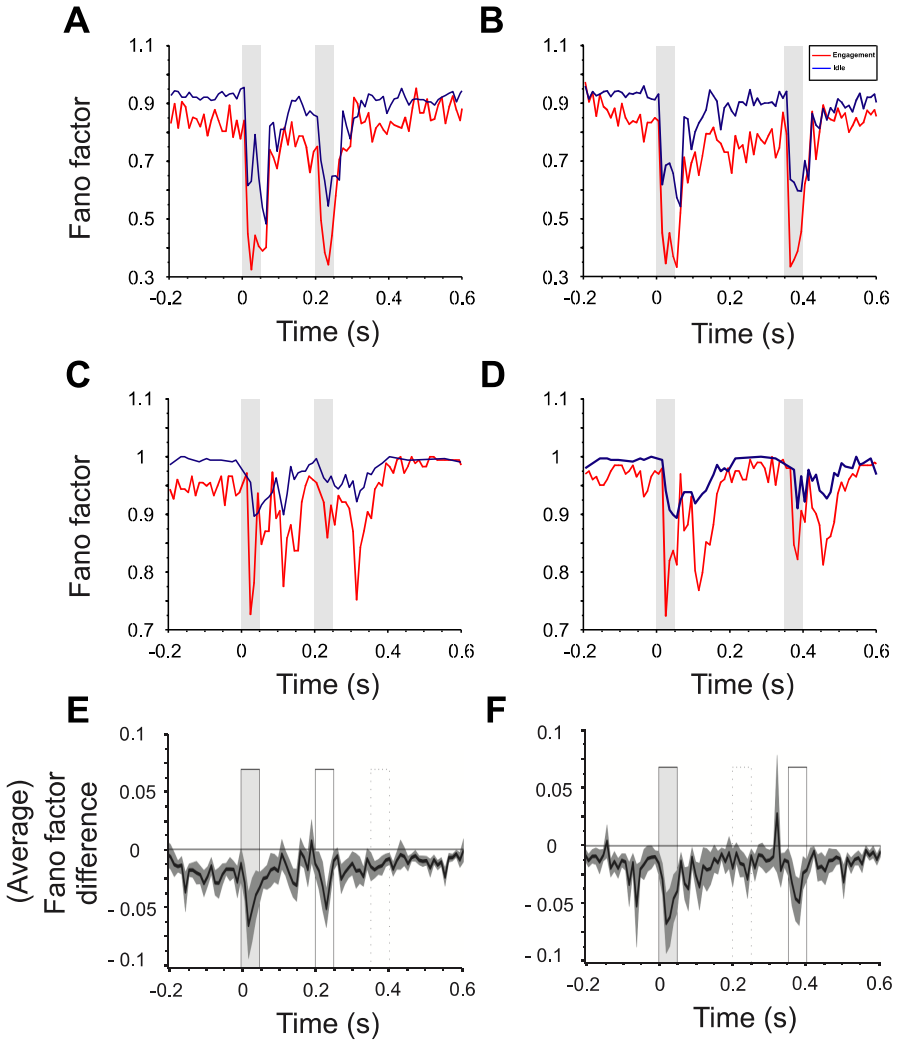

FIGURE 3.6: **Response variability is reduced during engaged brain states. A-B:**. Fano factor (Ff) of an example neuron during short (A) and long (B) ISI and for the engaged (red) versus idle stages (blue). Grey bars indicate the presentation of $S1$ and $S2$. A reduction in Ff during the engaged compared to the idle state can be observed during trials of both short (A) and long (B) ISI. Not only during stimulation but also during the ISI, and during the spontaneous activity period preceding stimulation (-0.2-0 s), variability was decreased in the engaged with respect to the idle state of the animal. **C-D:**. Same as in A, B for another example neuron with an onset-offset response pattern. In the engaged state, there was a reduction in Ff during the onset and offset compared to the idle state. Similarly, during the spontaneous activity period preceding stimulation (-0.2-0 s), there was a decreased variability in the engaged (C; D) compared to the idle state of the animal. **E-F:**. The difference in the Ff value for each bin of the trial was obtained between the engaged state and the average idle recordings. The obtained differences were averaged and the SEM errors are displayed with the grey shadow. Most of the values are negative, indicating lower variability in the engaged condition with respect to the idle one. Enhanced differences in variability are observed during stimuli presentation.

but also during the ISI. Accordingly, neuronal activity during the ISI (Fig. 3.6 A, B) showed a reduced variability during task-engagement for short and long ISIs when compared to the idle state. Furthermore, we found that during the spontaneous activity period preceding a stimulus presentation (200 ms), there was a marked decreased variability in the engaged compared to the idle state of the animal (Fig. 3.6 A-D). In most cases, a significant reduction of neuronal activity variability during engagement was observed (n=14; p=0.01 (Wilcoxon))(Fig. 3.6 E-F), while the opposite was only observed in one neuron. Statistical comparisons (Wilcoxon) showed a significant difference between the variability during $S2$ in the engaged state vs that in initial idle ($p < 0.008$) or idle-post (0.008). However, there was no significance ($p < 0.7$) when the two idle states were compared.

Out of these fourteen neurons with significantly modulated Ff during engagement, nine neurons showed a reduction of variability in the 200 ms preceding a stimulus presentation, while eight neurons showed a reduction of variability during the ISI. In all, we observed a significant reduction of variability along the trial duration for all studied neurons during engagement as compared to the idle brain state, and this was enhanced during stimuli presentation.

It has been shown by other authors that Ff is not contingent on the firing rate [77–80]. We also tested this and, for that purpose, we selected bins with similar firing rate (< 5% difference) from engaged and idle trials. We plotted for each selected bin the Ff value in the idle versus in the engaged state for each neuron (n=14; see Fig. 3.7 A-B). Bins were matched according to the same time location of the trial in the different brain states (Fig. 3.7 A) and also to different time location (Fig. 3.7 B). Fig. 3.7 shows that most of the values remain above the x/y main diagonal, indicating that Ff values are larger in the idle than in the engaged state than what would be expected by a change in firing rate. We also computed the mean distance of the values with respect to the x/y main diagonal, which reflects the difference between the Ff-idle and Ff-engaged. We found that the positive values of the difference was 0.14 (std: 0.05) and 0.09 (std: 0.07) (A and B, respectively) while for negative values was 0.08 (std: 0.06) and 0.05 (std: 0.03) (A and B, respectively). Thus the Ff is nearly times larger in the idle compared to the engaged state, and therefore a decrease in Ff during engagement is not a mere artefact of an increase in firing rate. Additionally, the number of values above the diagonal are 64% and 67% (A and B,respectively) while the ones below are 36% and 32% (A and B, respectively). Finally, we compared the statistical significance (Wilcoxon) between the values above and below the x/y main diagonal (i.e., engagement vs idle). We found no statistical significance for Fig. 3.7 A ($p < 0.3$) while the opposite was found in Fig. 3.7 B ($p < 0.00$)

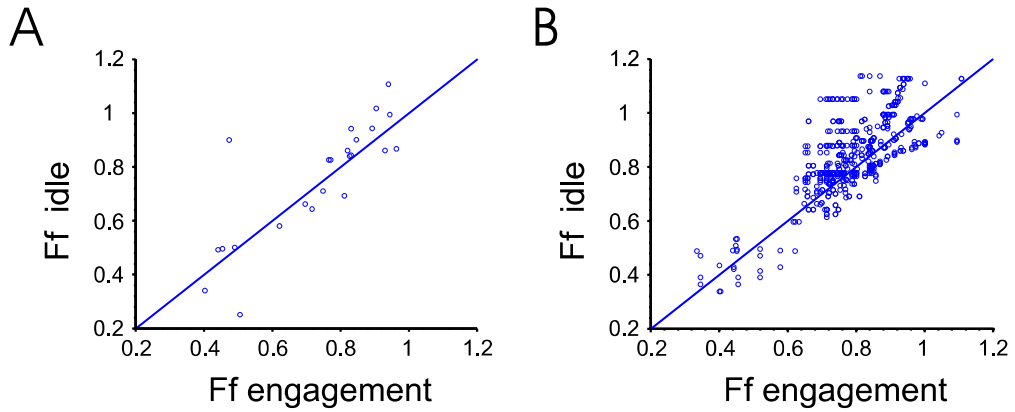

FIGURE 3.7: **The Fano factor is not dependent on the firing rate. A:** The Ff value of idle recordings (Y axis) is plotted against the Ff value of engaged ones (X axis). The Ff value was calculated for each pair of bins with equal spike count ($< 5\%$) between the idle and engaged states. Each pair of bins with equal firing rate had the same position in the trial in the different brain states. **B:** Each pair of bins with equal firing rate had different position in the trial. Only neurons with significant Ff were included (n= 14) for A and B. A significantly larger amount of the values remained above the x/y main diagonal for A and B.

possibly due to the increased number of values in the later. Thus from this section we conclude that during engagement there is a reduction in variability. In order to evaluate if this decrease in Ff is associated with an increased encoding capability we proceeded to use information theory to estimate mutual information (MI).

3.4.2 Mutual information is increased during engagement

MI analysis has been previously used to estimate the information content present in spike trains generated by neurons from the auditory cortex in both anesthetized [81, 82] and awake animals [83]. Here, we performed the MI analysis to find out whether single units in auditory cortex of the awake behaving animal encode information related with interval-discrimination of auditory stimuli. In our interval-discrimination task, the animals had to decide whether two identical stimuli were separated by 150 or 300 ms. In that task, the key stimulus that determines if the ISI category is "short" or "long" is the second one. MI between the variable "spike count" and the variable "ISI category" (150 or 300 ms), was calculated. Hence, we compared the MI value in the response to the first stimulus versus that to the second stimulus, in both idle and engaged states.

In order to compute the MI value we used here the bias corrected method of [84] as described in section 3.6. Furthermore, in order to evaluate the statistical

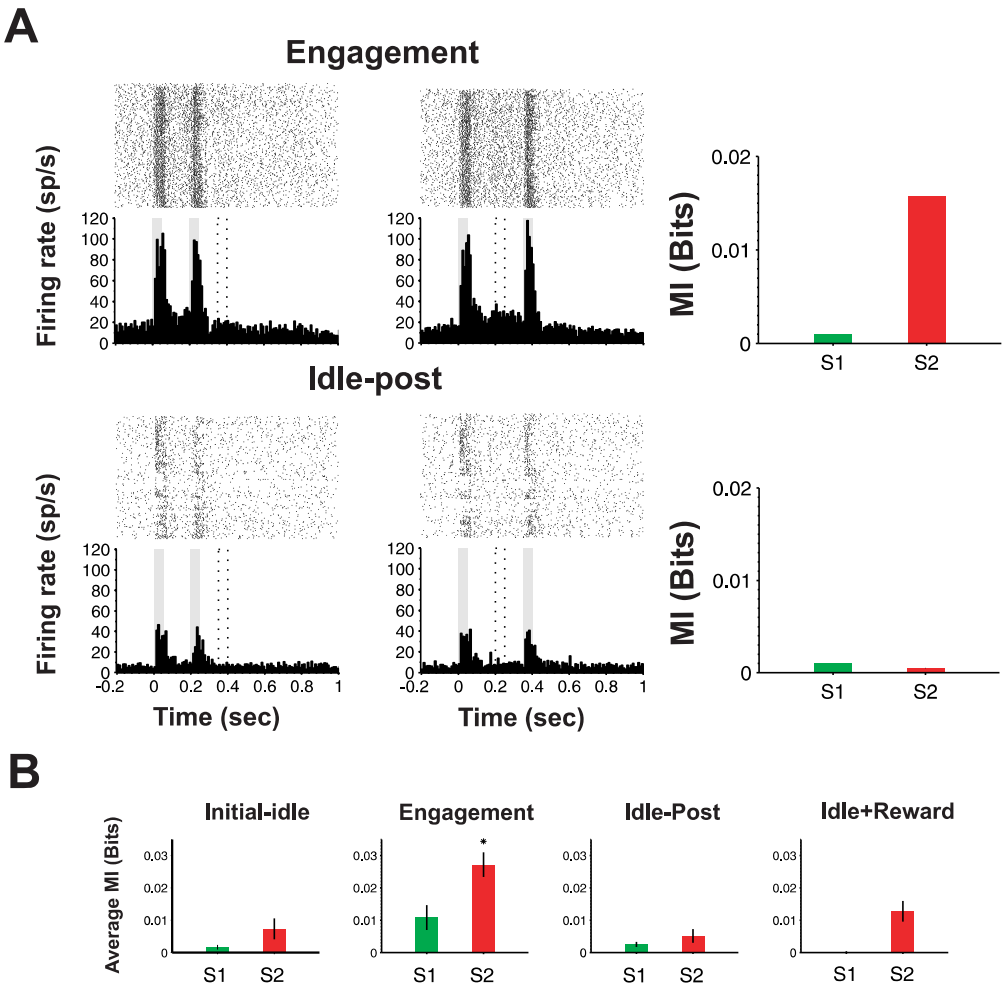

FIGURE 3.8: **Information content in response to auditory stimuli. A:** Raster plots and PSTHs of the neuronal discharge in different brain states. Identical auditory stimuli occurred in the grey boxes, with 150 ms interstimuli interval (*left*) and 300 ms interstimuli interval (*right*). In the *right* column, mutual information in the responses for both stimuli is represented in bits. Significance of surrogate test of mutual information is shown only during engagement in $S2$. **B:** The same pattern as in A (*right*) is shown for the average of ten neurons that passed the surrogate test of MI. The significance of surrogate test of mutual information is shown only during engagement in $S2$. Bars indicate SEM. Asterisks indicate $p < 0.05$.

significance of these MI values we used the surrogates method (see also 3.6). Given the sparse activity and the requirements for the calculation of MI [24] we found necessary to use 50 ms bins which comprised the stimulus duration ($S1/S2$). MI was calculated in twenty-one neurons with an average firing rate during the 50 ms duration of the auditory stimulus ($S1$) that was significant as defined in section 3.6. Of these twenty-one neurons 10 successfully passed the surrogates test of the spike count during $S2$ in the engaged brain state. One example neuron of these is shown in Fig. 3.8 A, showing a raster plot and peristimulus histogram (PSTH), for interstimulus intervals of 150 ms and 300 ms. In this case, MI values were higher during the response to $S2$ than to $S1$ in the engaged state ($S2$:0.016; $S1$:0.0), while it was not in the initial-idle stage ($S2$:0.0; $S1$:0.0), idle-post ($S2$:0.0; $S1$:0.0) or idle+reward ($S2$:0.0004; $S1$:0.0). Furthermore, we tested the surrogates significance of these MI values and we found that in the engaged estate the MI value during $S2$ significantly (0.007) passed the surrogates test.

Mean MI values during responses to $S1$ vs $S2$ stimuli (N=10) in the initial-idle ($S2$:0.007), idle-post ($S2$:0.005) and idle+reward ($S2$:0.010) were lower than in the engaged state ($S2$:0.028) (see Tab. 3.1). Even though these MI values could be interpreted a rather low, similar MI values have been found in the auditory cortex [85]. The surrogate test was only passed in the engaged brain state and these results suggest that the engaged state carries more information than the first one.

Statistical comparisons (Wilcoxon) of the MI were also performed between $S1$ and $S2$ for each brain state (see Tab. 3.1). We observed significant differences between MI values during $S1$ and $S2$ in engagement while not during the idle brain states (see Tab. 3.1). Additionally, $S2$ values of MI were compared (Wilcoxon) among brain states (see Tab. 3.1) and we also observed a significant difference of MI in $S2$ during engagement as compared with the idle states. Therefore, information content of spike trains evoked by auditory responses is augmented during the engagement in an interval-discrimination task.

A response profile of another example neuron showing an "onset-offset" pattern is illustrated in Fig. 3.9 A and B. In this case, spontaneous activity is increased during the time interval preceding stimulus presentation during task-engagement. The MI value in the response to $S2$ was 0.0465, with a surrogate significance of 0.002, during task-engagement, while MI value was lower during the idle ones (initial-idle: 0.0001; idle-post: 0.0065; idle+reward: 0.0001, see Fig. 3.9 C,left). This neuron further shows that the engaged state of animals has an effect on the information content in spike trains evoked by behaviourally-relevant stimuli.

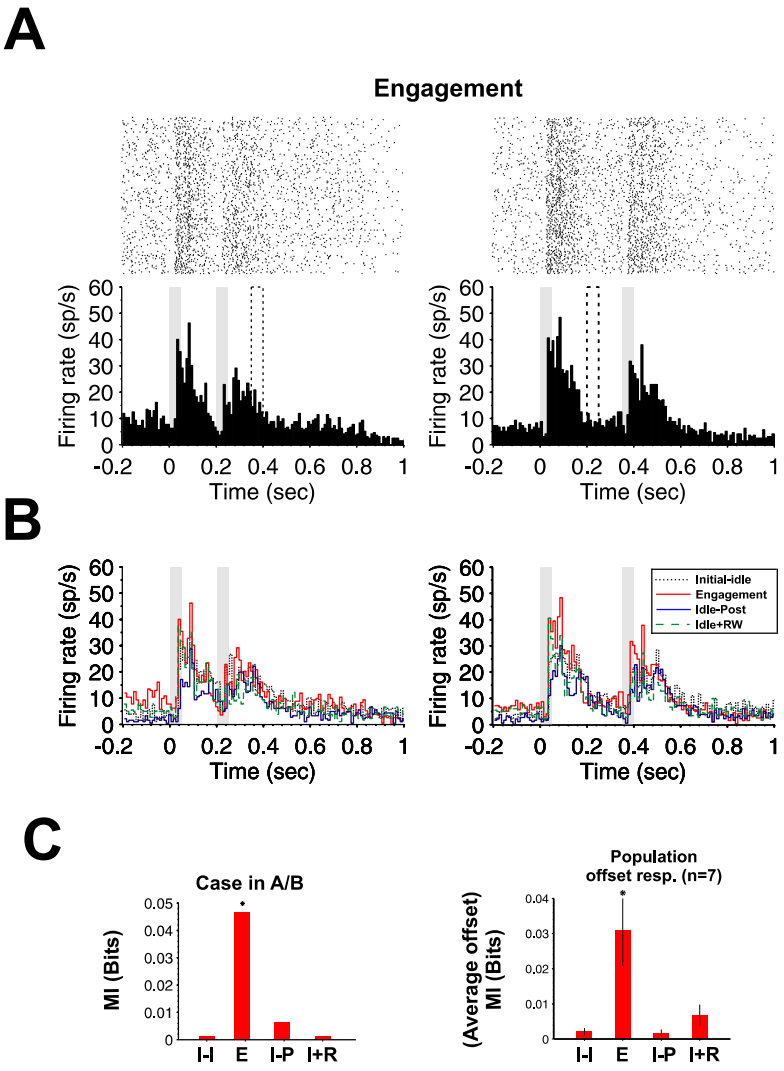

FIGURE 3.9: **Information content in response to a second auditory stimulus ($S2$). A:** Raster plot and PSTH (180 trials) of the firing rate of an example neuron to identical stimuli presentation to 150 (*left*) and 300 (*right*)-ms interstimulus intervals. Notice the prominent offset responses of this neuron. **B:** Overlaid PSTHs illustrate the response of the neuron in A while in different brain states. An increase in spontaneous activity preceding a stimulus presentation is observed only in the engaged state. **C:** Mutual information in the response to stimulus 2 ($S2$) in the different stages of the experiment (II (initial-idle); E (engagement); IP (idle-post); I+R (idle +reward). The graph on the left corresponds to the case illustrated in A and B only for $S2$. The significance of surrogate test of mutual information is shown only during engagement in $S2$. The graph on the *right* corresponds to MI contained in offset responses (n=7). The significance of surrogate test of mutual information is shown only during engagement in $S2$ offset. Bars indicate SEM. Asterisks indicate $p < 0.05$.

We also explored the information contained in the offset responses evoked once the stimulus was terminated. Seven neurons showed offset responses to auditory stimulation, while four of them showed additionally onset responses (e.g. Fig. 3.9 A, B). MI during the offset response component was calculated after $S2$ termination in those neurons that were classified as "onset-offset" or "offset" (n=7). We analyzed MI during a time window of the same duration as the one used to calculate MI during stimuli presentation(50 ms). The 50 ms window was taken around the peak of the offset response (25 ms before and after the peak) of $S2$. MI values of the population mean of offset neurons (Fig. 3.9 C; *right*) were significantly higher in the engaged state (0.0310) as compared with the idle (initial idle:0.0021 ; idle-post:0.0017; idle+reward: 0.0067) as evidenced with the surrogate test (see Tab. 3.1). In order to test the significance of that MI values statistical comparisons (Wilcoxon) of the MI were again performed between $S1$ and $S2$ for each brain state (see Tab. 3.1). As in the other case we observed significant differences between MI values during $S1$ and $S2$ in engagement (0.0468) while not during the idle brain states (see Tab. 3.1). When we compared $S2$ values of MI (Wilcoxon) among brain states we again observed a significant difference of MI in $S2$ during engagement as compared with the idle states (see Tab. 3.1). These results altogether suggest that offset neuronal response after $S2$ termination not only carries information, but carries a similar amount of information about the category of the ISI carried by the one of the responses of onset neurons during $S2$ presentation.

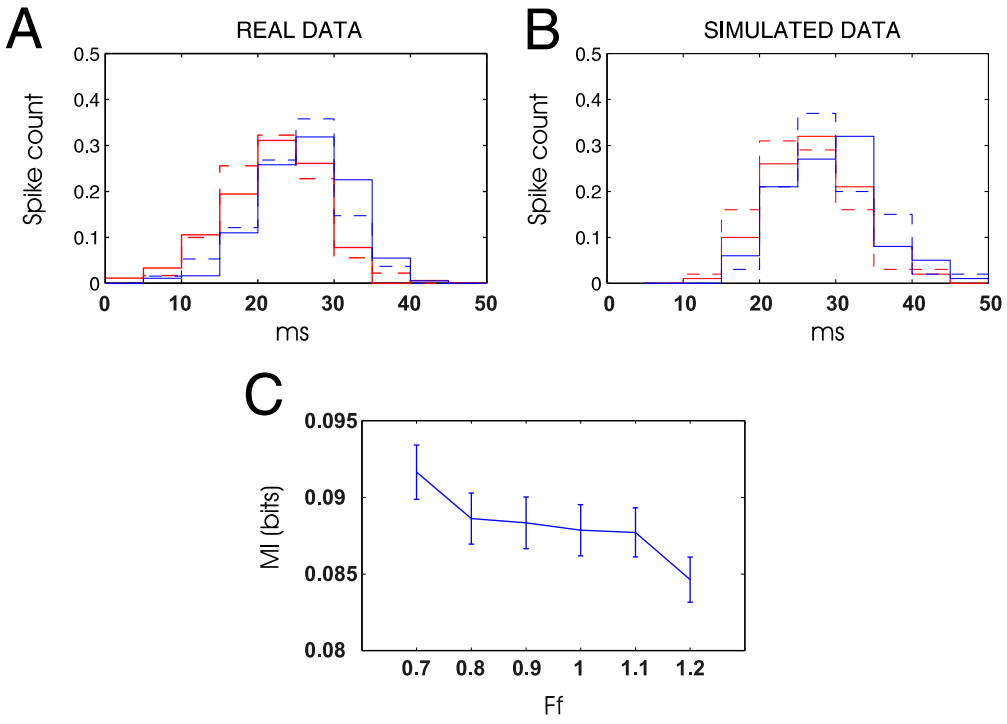

FIGURE 3.10: **Relationship between Fano Factor (Ff) and Mutual Information (MI). A:** Spike count distribution of an example neuron during the second stimulus presentation (50 ms) during the engaged state (red) and idle (blue) and for the left (dashed) versus right (solid) responses. **B:** Spike count distribution of simulated neuronal data with a gamma distribution. The same amount of trials (400) and parameter values (mean, and Ff) were used as in A in the four different conditions. **C:** The MI was obtained out of the simulated distribution of the engaged state shown in B. As the Ff value was increased the MI value decreased. Bars indicate S. D.

3.4.3 Relationship between Fano Factor and Mutual Information

In order to study the relationship between Ff and MI more systematically we developed a theoretical toy model that parametrised the experimental data. The toy model was defined by generating artificial spike train datasets whose inter-spike intervals follow a gamma distribution with a given mean firing rate (Fig. 3.10 B) as in the real data (Fig. 3.10 A) in order to test the relationship between Ff and MI (Fig. 3.10 C). Artificial datasets were modelled by gamma point processes, where Ff and the firing rate can be controlled. This model has been successively used to model spiking data [86, 87]. In this model, inter-spike intervals are independently drawn from a gamma distribution (see equation 3.1) that has two parameters: a scale parameter, r, that controls the intensity of the process (firing rate) and a shape parameter, α, that controls the variance of the distribution. Indeed, Nawrot et al[88] showed that $Ff = 1/\alpha$. For a given pair of r and α, we generated a spike

Summary Statistics State	MI in $S1$	Surrogate in $S1$	MI in $S2$	Surrogate in $S2$	P-val Between $S1$ and $S2$	P-val Between $S2$ Engaged vs. Idle	MI value in $S2$ Offset	Surrogate in $S2$ Offset	P-val Between $S2$ Offset Engaged vs. Idle
II	0.002	0.56	0.007	0.53	0.313	0.027*	0.0021	0.677	0.0082*
E	0.011	0.28	0.028	0.04*	0.010*		0.0310	0.0468*	
IP	0.002	0.50	0.005	0.51	0.562	0.039*	0.0017	0.7124	0.0047*
I+R	0.0002	0.70	0.013	0.60	0.250	0.031*	0.0067	0.2630	0.0156*

TABLE 3.1: Initial Idle (II); engagement (E); Idle-Post (IP); Idle+reward (I+R) MI in $S1$: mean MI value in $S1$ (bits).Surrogate in $S1$: mean value of the surrogate significance test in $S1$. MI in $S2$: mean MI value in $S2$ (bits).Surrogate in $S2$: mean value of the surrogate significance test in $S2$. p-value engaged vs idle: p-value between $S2$ in engagement and $S2$ idle. MI value in $S2$ offset: mean MI value in $S2$ offset (bits) Surrogate in $S2$ offset: mean value of the surrogate significance test in $S2$ offset p-value between $S2$ offset eng vs idle: p-value between $S2$ offset in engagement and $S2$ offset idle.$*p - val < 0.05$

train composed of 1000 consecutive inter-spike intervals. The length of the spike train (T) was divided into short non-overlapping time bins of 50 ms (equal to the stimulus period in the experiments) and the spike count (N) was calculated in each time bin. To avoid border effects we left aside the 10 first time bins and we stored the spike counts of the following 200 bins. As a result we obtained a spike count distribution for a given set of r and α, noted $f_{r,\alpha}(N)$.

$$p_\alpha(\tau) = \frac{(\alpha r)^\alpha \tau^{\alpha-1}}{\Gamma(\alpha)} \exp(-\alpha r \tau) \tag{3.1}$$

Using this procedure, we generated, for a given α and for two fixed values of the rate parameter, denoted r_1 and r_2, two spike count distributions, $f_{r_1,\alpha}(N)$, and $f_{r_2,\alpha}(N)$ computed the MI between the parameter r and the N. In our case r_1 corresponds to the mean firing rate during $S2$ in short ISI trials while r_2 corresponds to the mean firing rate during $S2$ in long ISI trials for a certain neuron. According to $Ff = 1/\alpha$, we varied α between 0.3, 1.1 while keeping r_1 and r_2 fixed. Then we calculated the MI for the pair (r_1, r_2), for all the α, between the stimulus and the response. This procedure was repeated 1000 times in order to estimate the error. As shown in Fig. 3.10 C, the MI of the simulated distribution (Fig. 3.10 B) increased for decreasing values of Ff. We found then a negative correlation between these parameters (corr:-0.77; $p < 0.05$) such that for lower Fano factor, mutual information increased. In conclusion, this toy model demonstrated that a parametric decrease of the Ff systematically increased the MI, generalizing therefore the experimental observations.

3.5 Discussion

We studied neuronal responses in rat auditory cortex during a decision-making task where intervals between auditory stimuli were categorized. Neuronal responses during and after evoked activity were compared in engaged versus idle states. Their firing rate, mutual information and variability were also quantified. Out of eighty six neurons recorded in the auditory cortex of the awake freely moving rat, auditory responses were evoked in forty nine neurons, a proportion similar to that in [58]. We refer to task- engagement since we consider that the animal needs to be engaged to do the task correctly. However, during trials of engagement we cannot rule out the participation of other mechanisms like expectation [52]. We found that neuronal firing rate during engagement was more often up than down-regulated during auditory responses. We cannot rule out a bias of extracellular recordings towards more active neurons, influencing our observed

impact of task-engagement on the firing rate. Ongoing activity recorded in the intervals in between auditory stimuli during the same task is in some cases also significantly modulated by engagement, being usually increased [10]. A prominent decrease in neuronal variability during both sensory-evoked and non- evoked activity was detected during engaged versus idle listening. Finally, information content in auditory-evoked spike trains was higher in engaged than in idle states, in particular in those evoked by the task-relevant stimulus.

In general, Fano factor reduction can be associated with incremented encoding capabilities only under strong assumptions. Indeed, a neuronal network can have a very low Fano factor (almost identical spike trains in multiple trials), but zero coding precision (identical spike trains for multiple stimuli). We showed that, in our case, the reduction of the Fano factor is indeed directly associated with an increment of the encoding/processing of the discrimination capability evidenced in the behavioural response. Furthermore, mutual information is more powerful because it is defined by the measurement of different sources of variability, namely an entropy term that characterizes the neuronal variability in general, and another "conditional" entropy that measures the specific variability observed for a given condition (or behavioural response). Let us note, that the increase of MI observed is not only due to a decrease of the conditional entropy term but to the combination of both, total response entropy, and conditional noise entropy terms. Indeed, the total response entropy term increases too, so that the increased information acts synergistically with changes in the neural representation. In order to complement this view we also studied the direct reduction of the variability *per se*. We thus studied the Fano factor reduction for a specific condition. Ff is particularly useful because, contrary to the mutual information, it can be computed in small sliding windows during the whole trial. Indeed, by doing this, we were able to show for the first time that a reduction of variability is observed in the absence of external stimulation (between the stimuli) but also in a relevant time region. In this study, we have demonstrated that the reduction of variability (Fano factor) observed during stimulus presentation (including the interval between both stimuli), in particular the larger reduction due to engagement, is in fact associated with increased encoding capabilities for discrimination. We show this by complementing the Fano factor variability measurement with a direct information-theoretical measurement of encoding capabilities via mutual information between neuronal activity and behavioural responses.

3.5.1 Firing variability of single units in A1

Sensory processing during the processing of task relevant information has been linked to enhanced responses [72–74] and also to decreased ones [51]. Moreover, evoked responses in primary auditory cortex can be modulated as a result of temporal expectation [52]. Increased inhibition has also been suggested to play an important role in cortical responses to relevant stimuli [89, 90]. Task-engagement also induces tonotopic changes [91–94] and tuning shifts of the same neurons towards the target stimulus [95]. However, there are no studies describing how engagement affects response variability of single units in the auditory cortex of the behaving animal.

Earlier studies have suggested that a decline in response variability is a widespread phenomenon in the cortex that spans different areas, animal species, and that always occurs to the onset of stimuli presented, irrespective of the brain state of the animal [77, 96–98]. Furthermore, the neuronal variability over trials declines in particular in situations where the encoded information serves to guide behaviour [77, 96, 99, 100]. This has been experimentally demonstrated in recorded neurons of the visual area V4 [101, 102] in the context of an attentional paradigm. In these studies it has been shown that the mean-normalized variance (Fano factor) of the spiking activity is reduced by attention, consequently increasing the sensitivity of neurons towards relevant aspects of stimuli. Neuronal response variability may also depend on the type of neuron, i.e., narrow or broad spiking [80], while some authors suggest that the attentional effects on variability may reflect an intrinsic property of neural circuits [103]. Moreover, stimulus-induced trial-to-trial variability may be explained by the same dynamics of ongoing spontaneous activity [104]. Our results suggest that the external stimulation or the behavioural requirements of an interval-discrimination task stabilize the dynamics in a controlled way such that neuronal variability is reduced. Thus, the signal-to-noise ratio is increased, yielding the basis for an improved encoding of the stimulus information. The possible dependence of Ff on firing rate deserves to be considered. It has been suggested that a reduction in neuronal response variability could be correlated with an increase in firing rate [77–80]. Some studies have shown that decreased variability is not due to an increase in firing rate [77]. In our current study, we find that in bins where the firing rate was equal between the idle and engaged states, variability was still reduced in the later (see Fig. 3.6). This is consistent with the result reported by [80] in visual cortex (V4), where lower variability during engagement was observed when bins with equal firing rate were compared.

3.5.2 Information content in single units of A1

Quantification of information content in spike patterns has provided important insights in the understanding of key features of sensory processing [24, 81–83, 105–108]. Previous studies [108] have suggested that information content in multiunit activity of auditory cortex is higher in the interspike interval, than during the firing rate or event-locked spikes [109] when repetitive stimulation is presented to anesthetized animals. Similarly, [83] quantified the information present in temporal spike patterns and the phase of population firing, suggesting that these combine information for encoding natural sounds in the auditory cortex. Therefore, the combination of different neuronal codes could enrich auditory stimuli representation, and increase robustness against noise. Looking into the mutual information between the stimulus and neuronal response, [110] suggested that spike precision enhances the encoding of information about extended complex sounds. Also, information can be carried by spike timing in case of sparse acoustic events, while firing rate-based representations encode rapidly occurring acoustic events [81]. The study of auditory activity during a task where a monkey compared the relative frequency of two auditory stimuli showed that stimulus-locked responses, and in particular firing rate, only correlated with performance during stimulus presentation [12]. This was not the case though during delay periods, as it would be the case if it was related to working memory or decision-making. The authors suggested that the auditory cortex may serve to encode information of sensory stimuli, mostly by means of firing rate, with no cognitive function related to decision-making or memory. Our analysis shows that MI is particularly enhanced to the relevant stimulus during task-engagement, although we do not definitely demonstrate its association to performance. Altogether, we think that this is an evidence of the role of auditory cortex in temporal discrimination during a decision-making task, even when its origin may be a top-down influence.

3.6 Methods

Cluster cutting (isolating single units from the multiunit recording data) was performed using an Off-Line Spike Sorter (OFSS, Plexon). Waveforms were sorted as in [75]. Single units exhibited a recognizable refractory period (>1 ms) in their ISI histograms.

Analysis of peristimulus histogram (PSTH) were performed using 10 ms bins to estimate responses to auditory events accurately. Frequency response histograms

were obtained by averaging the spiking activity within each bin during the whole recording. The onset of each stimulus presentation was aligned to 0. Raster plots illustrate the timings of individual spikes in individual trials during the whole recording. Only correct trials were selected, comprising 180-200 responses per each side.

We refer generally to "spontaneous activity" along the manuscript to that firing rate of the neuron occurring whenever there was no auditory stimulation. Therefore we include under the term "spontaneous activity" neuronal firing rate that may as well correspond to prolonged responses to the stimuli or to modulation due to cognitive and behavioural states (expectation, engagement, attention, etc).

3.6.1 Fano Factor

The Fano factor (Ff= $\frac{\sigma^2}{\mu}$), was computed as the ratio between spike count variance ($sigma^2$) across trials and mean spike count (μ). Ff was calculated in 10 ms bins along the trial duration (from - 200 ms to $S2$ + 600-750 ms) for correct trials (180-200 at each response side). For each cell we compared the minimum Ff value during the interstimulus interval (ISI) (-200 ms to 0) and the minimum Ff value in the interval from 0 to $S2$ + 600-750 ms. We defined a decreased Ff when a neuron had three consecutive bins with Ff lower than the minimum value obtained during the ISI (whether short or long ISI). Sparseness of neuronal activity has been shown to affect information theoretic measurements [24]. To avoid this, we selected neurons with the highest firing rate while the rat performed correct trials. In all cases a minimum of 180 trials per side were considered. Given that the animal made few errors plus the fact that the firing was sparse, there were not enough spikes fired during wrong trials to allow for independent analysis of correct vs incorrect trials. The response to all correct trials was represented in a PSTH in 10-ms bins. The mean firing frequency and its standard deviation during the 200 ms preceding $S1$ was calculated. Neurons were considered to have significant evoked-spiking responses if after $S1$ onset the evoked firing during the 50 ms stimulation was 5 standard deviations over the spontaneous frequency [57, 111]. Twenty-one neurons crossed the threshold during $S1$ presentation (50 ms).

3.6.2 Information Analysis

To estimate the information content carried by the neuron's firing rate, we performed Mutual information (MI) analysis between "spike rate" and "category of

ISI (ISI)". The MI was calculated following the equation 1.3. In that specific case, s is spike rate and r is category of the ISI (either short or long). As we explain in the introduction, the calculations were based on frequency estimates of the probabilities ($p(r,s), p(s)$, and $p(r)$) we used spike counts in 50-ms time windows during each stimulus presentation.

We estimated the value of the MI by using the direct method for the MI estimate and the [112] method for the bias correction since the computation of the information content is subject to statistical errors given that the MI is based on the estimation of probabilities. This method corrects for the bias by means of decomposing the mutual information in different factors and then removing the ones affected by the bias or noise.

Significance of the MI

In order to establish the significance of the estimated MI we applied the surrogates method [29]. We tested if the estimated MI significantly rejected the null hypothesis, i.e., non information content ($I(r;s) = 0$). We tested this null hypothesis by generating 1000 surrogates of the spike activity during stimulus duration, which by construction should not contain information. Thus, each surrogate is generated by shuffling the assignment between stimuli and response. We computed for each surrogate the MI value and we compared all these bootstrapped information values with the real value of the original data. We calculated the statistical significance of the estimated MI by computing the area in the null hypothesis distribution (MI of the surrogate) below the MI value corresponding to the original data. We considered that if the area of the null hypothesis was larger than 85% of the total area, the estimated MI value of the original data was significant. This *criterion* means that the Null hypothesis (no significant MI) can be rejected with a 85% probability, i.e., the estimated MI is significant at 0.15 level ($p < 0.15$) Let us note that even for a smaller p-value of 0.05 still 7 neurons passed the test. Nevertheless, we took a p-value of 0.15 in order to increase the amount of neurons that passed the test and have a more reliable statistics. All the calculated MI values for both original data and surrogates were bias corrected using this method.

3.6.3 Experimental Details

Ethics approval and Surgical Procedure

The project was approved by the animal Ethics Committee of the University of Barcelona. Rats were cared for and treated in accordance with the Spanish regulatory laws (BOE 256; 25-10-1990) which comply with the EU guidelines on protection of vertebrates used for experimentation (EUVD 86/609/EEC).

Recordings were obtained from two Lister Hooded rats (250-350 grs) that were chronically implanted with tetrodes in their primary auditory cortex. Animals were trained for 21 days. After a week of water and food *ad libitum*, a microdrive holding the tetrodes was implanted. To perform the surgery, anaesthesia was induced using intraperitoneal injections of ketamine (60 mg/kg) and medetomidine (0.5 mg/kg). The animals were then mounted in a stereotaxic frame and their skulls exposed. A 3 mm diameter craniotomy was made, with its center at -5.3 mm anterior-posterior, and 6.6 -7 mm medium-lateral from bregma [113]. These coordinates were used in order to position the microdrive dorsally, which made it more stable than entering laterally over the auditory cortex. Body temperature was monitored through a rectal thermometer and maintained (36-38^o) using an electric blanket. Heart rate and blood oxygen levels were monitored. Reflexes were regularly checked during surgery to assure deep anaesthesia. Other drugs were given during surgery and recovery period to prevent infection, inflammation and as analgesia: antibiotics (enrofloxacin; 10mg/kg; s.c.) and topical application of neomycin and bacitracin in powder (Cicatrin ®), analgesic (buprenorphine; 0.05 mg/kg; s.c.), anti-inflammatory (methylprednisolone; 10 mg/kg; i.p.), and atropine (0.05 mg/kg, s.c.) to prevent secretions during surgery. Once the animals went through all experimental sessions, humane killing was performed by means of an overdose of pentobarbital (0.8 ml).

Tetrodes and Microdrives

Each tetrode was made from four twisted strands of HM-L-coated 90% platinum-10% iridium wire of 17 diameters (California Fine Wire, Grover Beach, CA). Gold plating decreased their impedance to *ca.* 300-500 $\kappa\Omega$. Four tetrodes were held by a cannula attached to a microdrive supplied by Axona Ltd, St Albans, UK. This microdrive allowed for dorsal to ventral tetrode movement to search for new units. Microdrives were attached to the skull with dental cement and 7 stainless steel screws. The auditory cortex was reached by vertical descent, and the tetrodes

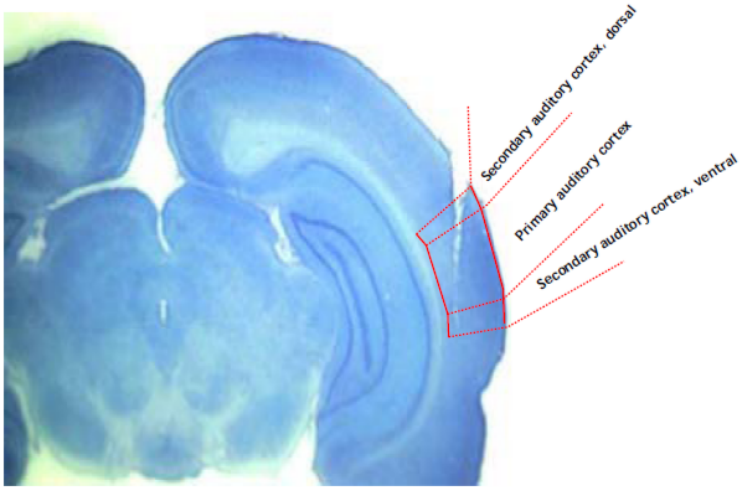

FIGURE 3.11: Picture of a coronal rat's brain slice after tinction. Tetrodes socket can be seen in the auditory cortex at the right side of the picture.

were lowered $300\mu m$ during the surgery. Vertical descent performed after surgery was of $50\mu m$ per day until an auditory response was observed. All the recordings included in this study corresponded to A1 [114] . This estimation is based on the depth of the included recordings and on the histological reconstruction of the electrode's tracks (see Fig. 3.11). The auditory latencies were typically 10-20 ms, which are also characteristic of A1 [115–117].

Electrophysiological recordings from awake freely moving rats

During the training period, animals lived in large cages of 28 x 42 x 30cm (Charles River) in a rich environment, under a 12 hr light/dark cycle, and with food *ad libitum* and water restriction. Before training and after a week of postoperative recovery period, the animals were accustomed to the recording chamber. The electrode wires were AC-coupled to unity-gain buffer amplifiers. Lightweight hearing aid wires (2-3 m) connected these to a preamplifier (gain of 1000), and to the filters and amplifiers of the recording system (Axona, St. Albans, UK). Signals were amplified (x15000-40000), high pass filtered (360 Hz), and acquired using software from Axona Ltd (St Albans, UK). Each channel was continuously monitored at a sampling rate of 48kHz. Action potentials were stored as 50 points per channel (1 msec, 200μsec pre-threshold; 800μsec post-threshold) whenever the signal from any of the prespecified recording channels exceeded a threshold set by the experimenter for subsequent off-line spike sorting analysis. Data were excluded if any drift was detected. Before each experimental session, tetrodes were screened for

FIGURE 3.12: The earphone used for the experiment.

neuronal activity. Once spikes could be well isolated from background noise, the experimental protocol started.

Experimental set up

The recordings were performed inside a box built in black acrylic of 22 x 25.5 x 35 cm. This box was placed inside two wooden boxes placed one inside the other. Between each box, two isolating foam rubbers (4 and 2 cm thick) were placed to soundproof for low and high frequencies. A wooden cover and soundproof foams closed the entire recording chamber, with only a hole to allow the entry of a recording wire (2 mm thick) connected to the preamplifier. Water valves were placed outside the recording chamber. The animals poked their noses into three different sockets (2 cm wide and separated by 3 cm each, and with no cover in the top part to avoid being hit by the microdrive). Recordings were obtained in darkness, and the experiment was filmed with an infrared camera placed above the recording chamber.

Presentation of Sound Stimuli

The protocols of stimulation were controlled through MATLAB®, a National Instrument card (BNC-2110), and a breakout box (FS 300 kHz). Sound triggers had μs precision. Sound stimuli were delivered through earphones (ER.6i Isolator, Etymotic Research Inc.) which were screwed in each recording session to the earphone (see Fig. 3.12) holders, chronically attached to the animal skull with dental cement. The earphones were adjusted inside the ear with silicone tips with a separating distance of 1 mm from the ear canal. Similarly sound calibration

was performed inside the acoustic isolation box with a microphone (MM1, Beyerdynamic) placed 1 mm away from the earphone and using a preamplifier (USB Dual Pre, Applied Research and Technology). The sound stimuli during idle and engaged recording stages had a duration of 50 ms, with an intensity of 80 dBs SPL pure tones of 5322 Hz, and a 6 ms rise/fall cosine ramps. It was identical for both the first and second stimulus. Interstimuli intervals were 150 or 300 ms and both had the same amount of trials (180-200). Similarly, the total number of correct trials in the engagement stage was the same as in the idle one (180-200). The intertrial interval also had a similar duration in the engagement and idle stages (2-3 s).

Chapter 4

Neural correlates of decision confidence in PMv

The work presented in this chapter is submited in *Journal of Neuroscience*: The authors are: Marina Martinez-Garcia, Andrea Insabato, Mario Pannunzi, Jose L. Pardo-Vázquez, Carlos Acuña and Gustavo Deco.

4.1 Introduction

Decision-confidence, the sensation of the correctness of a choice, is a first step towards introspection and it is crucial for learning in perceptual decision-making. Although confidence has been a topic of investigation in psychophysics since its beginning [118–122], and is generally being considered as one of the most meaningful ways to measure behavior (together with accuracy and reaction time), very little, in fact, is known about its neurophysiological and computational correlates.

Using two-alternative forced-choice (2AFC), psychophysical experiments have established a direct relation between confidence and the discriminability of the stimulus in correct trials [118, 123]. Inversely, confidence as a function of evidence for the decision decreases in error choices [1, 119, 122]. This modulation produces an x-shaped pattern, which is considered a hallmark of confidence (for a review see [9]).

4.1.1 Neurophysiological studies

Recently, two neurophysiological studies [124, 125] have shed some light on the neural basis of confidence. The first was an olfactory discrimination task, [124]

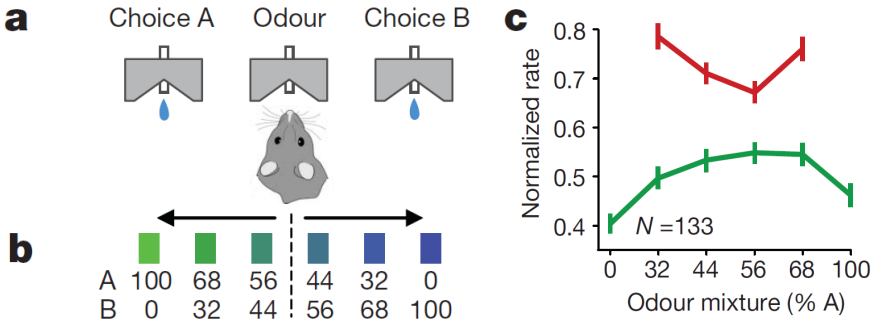

FIGURE 4.1: **A**, Schematic of the olfactory discrimination task . The rat enters the central socket and after a pseudorandom delay of 0.2–0.5 s a mixture of odours is delivered. It indicates left or right reward delivery, respectively. **B** The possible mixture of the two odours, A and B. design. **C**, Mean normalized firing rate of OFC neurons, for the negative outcome selective neurons. In red, error trials. In green correct trials. Both, the correct and the incorrect trials form the x-shaped pattern. Figure adapted from [124]-

(see Fig. 4.1). The authors trained rats to perform a two-alternative forced-choice (2AFC). In this task, subjects initiated a trial by entering a central socket. Then the stimulus was delivered, a mixture of two odours (A and B see Fig. 4.1B). The subject decided which of the two odours was the dominant one by going to the socket responding to A or B (see Fig. 4.1A). The difficultly of the task can be manipulated by varying the ratio of the two odours, as is indicated in Fig. 4.1B. Correct choices were rewarded after a delay period of 0.3-2 s.

The authors showed that the mean firing-rate of single neurons in a rat's Orbitofrontal Cortex (OFC) systematically increases as a function of the evidence in correct trials, and decreases in error trials, producing the above mentioned x-shaped pattern (see Fig. 4.1C). They have therefore suggested that OFC-neurons encode the uncertainty (or conversely the confidence) of a decision process.

The second was a visual discrimination task, [125], a random-dot motion task. In this task the subjects have to decide the net direction (to the right/left) of the dots, where some of them are moving coherently while the others just do it randomly. After a delay period, monkeys have to to indicate decision choice by a saccadic movement to one of the two targets (see Fig. 4.2). In a random half of the trials the authors introduced a third possibility, the "sure target". With this option the subjects got a small but certain reward. The "sure target" appeared during the

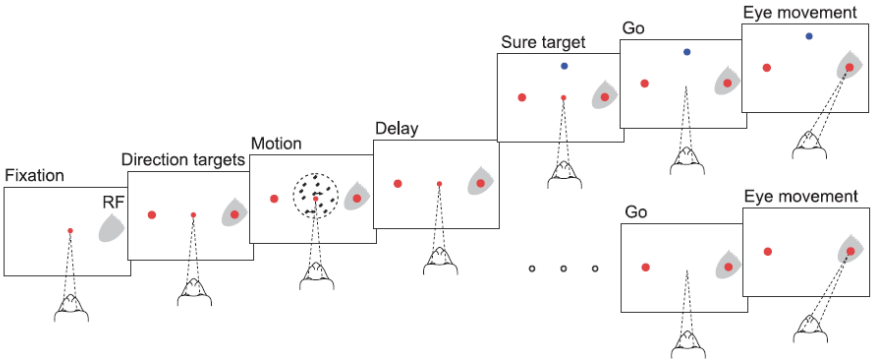

FIGURE 4.2: The sequence of events in the task. After fixation point (small red circle), two direction targets (large red circles) appeared on the screen. The motion stimulus appeared after a short delay. Bottom row: the delay persisted until the fixation point was turned off, which served as a "go" signal, then the monkey makes a saccade on the one target direction. A correct response were rewarded. Top row: a third target was presented after extinction of the motion. Choosing this sure target (Ts, blue circle) led to a smaller reward (80% of correct reward). On these trials, the monkey could choose Ts or a direction choice. The two trial types were randomly interleaved. Figure adapted from [125].

delay period, at least 500 ms after the random-dot motion was extinguished. The task's difficulty was controlled by varying the percentage of coherently moving dots.

They recorded neurons from Lateral Intraparietal sulcus (LIP) because it is well-known that in motion-direction discrimination-tasks single neurons in the primate's LIP-cortex fire with persistent levels of activity which predicts the saccade towards the neuron's response field, and indicates the subject's choice [6, 126, 127].

With this task Kiani et al. [125] have shown that the average firing-rate of LIP decision neurons in uncertain trials reaches only an intermediate level of activity, as compared to confident trials. Based on that they suggest that the LIP-sulcus neurons are involved in confidence encoding.

4.1.2 Neuronal code of confidence

However, all these results still leave open the question as to how the neural signal of confidence is encoded in single trials. In particular, it remains unclear whether neurons encode confidence in a continuous manner, or in a discrete manner. The results of [124] seem to suggest that OFC encodes confidence in a continuous way, which has also been suggested by theoretical studies based on drift-diffusion models (DDM) [1, 128, 129]. These theoretical studies seem to indicate that

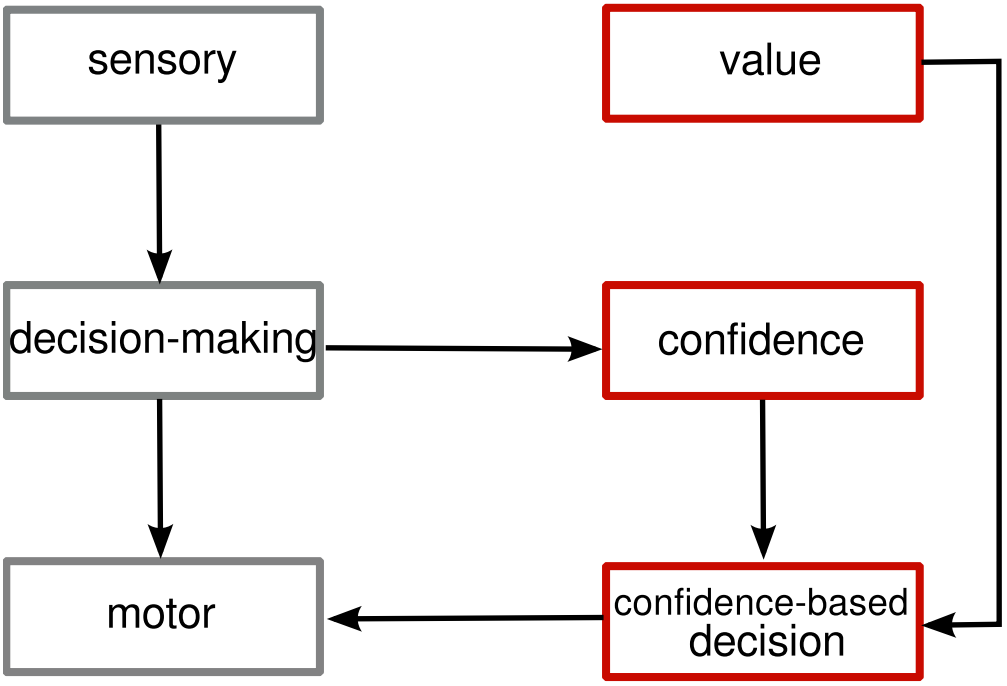

FIGURE 4.3: Pipeline of decision and confidence processing. The left part (in gray) represents the simplified sensory-motor path of perceptual decision-making. On the right (in red): modules involved in confidence estimation and confidence-related decisions. The "confidence" module receives input from the "decision-making" stage, thereby implementing a sort of monitoring of the decision process. This module represents the reliability of the decision process. The "confidence-based decision" module makes a judgment based on the confidence representation and value signals about the given options and transmits this second decision to the "motor" stage.

confidence is encoded by the position of accumulators at the moment of choice (race models), or by the time it takes the system to reach a decision (diffusion models). A binary coding of decision confidence, on the other hand, is suggested by the activity of LIP; Here neurons present a very different firing-rate in confident trials, as compared to uncertain trials [125]. Furthermore, it is interesting to note that [130] have shown that human subjects' confidence ratings present a highly bimodal distribution. A recent biologically realistic attractor model [131] combines continuous and discrete encoding into a two-stage model. The authors suggested that continuous confidence signals encoded in the decision-making neurons are translated into a binary response of confidence neurons, which discriminates high from low confidence. A similar mechanism has also been proposed within the DDM-framework of [132].

In order to embed these results in a general picture, we outline (in Fig. 4.3) the

essential pipeline of a simple decision task involving confidence computations. The left part of the graph represents the usual perceptual decision. In this context, sensory neurons encode the relevant information about the stimulus and inform decision neurons. Hence, once the decision has been computed, the motor plan can be elaborated by the "motor" module. When the decision confidence is going to have a role in the behavioral output one need to consider also the right part of the graph. A new module ("confidence") can compute the confidence in the decision by monitoring the activity of the decision area. Then the confidence information can be compared with informations about the value of different options and a new, confidence-based decision can be taken (e.g. the post-decision wagering experiments well summarized by [9]. In this outline it is reasonable that the "confidence" module would encode in a continuous manner the decision confidence. However, if the decision confidence is ever going to have an influence on the behavior, at some point in the sensory-motor path this information need to be discretized, in order to select one course of action (this is represented in Fig. 4.3 by the module for making confidence-based decisions).

In this chapter we analysed the data presented in the Introduction (1.1), in this task we controlled the difficulty by varying the difference between the bars orientation. Our aim in this chapter is to shed light on the encoding mechanisms of decision confidence in the primate brain. PMv neurons seem well suited to evaluate decision confidence, since previous studies have shown a central role of premotor cortex in the conversion of a decision into an action [8, 133]. More precisely, it has been found that PMv neurons encode higher cognitive processes, such as decision-making [7, 134], and performance-monitoring [7, 19], they also found that, besides the response, these neurons also encode the difficulty of the decision. These neurons can be the basis for an encoding of decision confidence since confidence is modulated by the difficulty of the decision.

We chose this task to analyse confidence, 1) Because the recorded area is well-suited to look for confidence encoding neurons. 2) Because the moment in which the decision is made is well marked off, from S2 presentation until the saccade. We have analyzed neuronal activity both across trials using Linear Models, and in single trials, using a Hidden Markov model (HMM).

We found in the primate brain a pool of neurons whose firing-rate continuously encoded confidence, in a very similar way to the results in rats OFC [124]. These neurons are not merely encoding the difficulty since their activity presents a mirror modulation in correct versus error trials, which is considered a hallmark of confidence. Interestingly, we also found neurons that showed two distinct levels

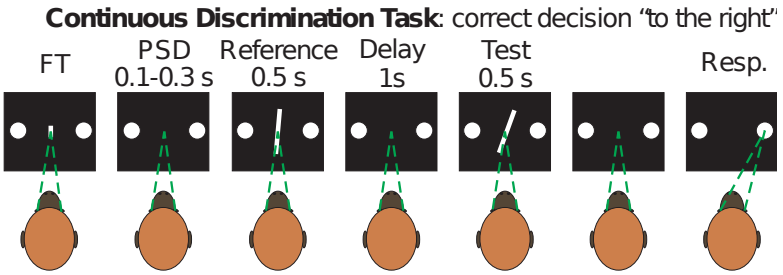

FIGURE 4.4: Experimental paradigm. The trial starts when the monkey fixes the gazeto the central target (FT). A brief pre-stimulus delay follows (PSD). The reference bar is presented for 500 ms with one of three possible orientations (Reference). During the subsequent delay the subject has to maintain fixation (delay). The test bar is shown tilted to the left or to the right in relation to the reference bar. The orientation of the test bar relative to the reference (TRO) manipulates the difficulty of the trial. When the test bar disappear the subject decides whether the test was tilted right or left in relation to the reference bar by making a saccade towards the right or left choice targets respectively. For more details see [7].

of activity over trials. Our results thus suggest that both continuous and discrete coding schemes for confidence are active in the brain.

4.2 Results

4.2.1 PMv neurons encode decision confidence

We studied the decision-process in the primate brain during a simple binary decision task. Two male monkeys (Macaca mulatta) performed a two-interval two-alternative discrimination task. They had to compare the orientation of a reference bar, presented during the first interval, with that of a test bar, presented during the second interval. They then had to decide whether the test bar was tilted right or left as compared to the reference bar (see Fig. 4.4 and section 4.4 for details). The level of difficulty of the task was controlled by varying the difference between the orientation of the first and the second bar, i.e the test bar's relative orientation (TRO). The TRO was varied from one up to four degrees and in both directions.

Single cells from PMv were recorded while monkeys performed the task. For a more detailed description of the task, behavioural results and neural recordings see section 4.4 and [7].

Our principal objective was to find neural signatures of decision confidence computations in this area of the primate brain. It is plausible that confidence-related

computations take place in the same area as where the decision is encoded, given the dependence of decision confidence processes on decision-making processes. In addition, [7] found decision-making neurons in PMv that encode the difficulty of the decision, a computation that is fundamental to confidence processing. And [125] found a correlate of decision confidence in the same neurons that encode the choice in monkeys' LIP. We therefore analyzed the activity of PMv activity recorded during the decision task above described.

Our analysis was restricted to a subset of the recorded neurons (336 neurons, see [7]), comprising the cells that were relevant to the decision task. Unless specified otherwise, in the following we will only describe correct trials, since error trials were enough only for few neurons. We identified a population of neurons (49 cells) whose firing-rate was high for both right and left decisions (as can be seen in the raster plot of a single neuron in Fig. 4.5A). Even if not predictive of the choice, the firing-rate of these neurons encoded the difficulty of the task, independently of the subject's choice, as revealed by the linear regression model (LRM) we used (see section 4.4). Fig. 4.5B shows the evolution in time of the coefficient of the LRM for a single neuron. The shaded area marks the time-window where the coefficient (d_1) of the LRM was significant ($p < 0.05$), giving an estimate of the encoding time-window of the neuron. For the neuron in Fig. 4.5 the encoding window spanned approximately 300 ms. As shown in Fig. 4.5C, during this period the firing-rate of the neuron, as a function of TRO, increases for both positive and negative values of TRO, producing a v-shaped pattern. Moreover, by pooling data from all the neurons we were able to analyze firing-rates during error trials. When the behavioral response was incorrect the firing-rate of this population showed an inverse pattern compared to correct trials. Overall, the normalized firing-rate separated in correct and error trials formed an x-shaped pattern (see Fig. 4.7), which had already been described by [124] in rats' OFCs as a correlate of decision confidence. Therefore these neurons are not merely encoding the difficulty but the confidence of the decision. As has already been remarked, consistent neural recordings during error trials were rare, but in order to confirm the pattern of the pooled responses at the level of single neuron we analyzed fifteen of the 49 neurons for which we had enough error trials recordings (see section 4.4 for details). To check whether the x-pattern was present in the single-cell activity we looked for neurons having the LRM coefficient significant, but with a different sign in error trials compared to correct ones. According to this analysis the coefficient was statistically significant for five of the fifteen neurons.

The x-pattern associated with confidence is poorly understood in terms of the computations from which it could be said to arise. In the following, we will try to

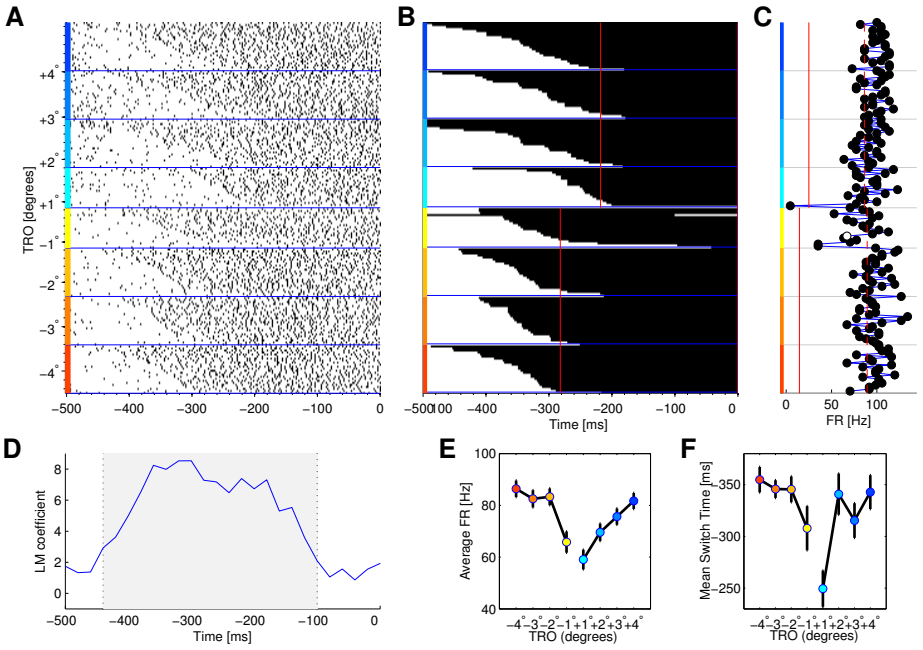

FIGURE 4.5: Single neuron from PMv cortex encoding confidence in a continuous way. To each separate test relative orientation (TRO) is assigned a different color: colors from dark blue (red) to light blue (yellow) correspond to right (or left) responses with increasing difficulty, i.e., relative orientation from 4° (−4°) to 1° (−1°). Trials are aligned to the saccade; around -510 ms the second bar was shown. **A**: Raster plot. The trials are sorted by the TRO and according to the timing of the state switch (as indicated by the HMM). **B**: Time course of state switchings, according to HMM and for the same neuron. Every row represents a trial. State one is represented by the color white, while black represents state two. Trials start in state one and later change to state two (indicating that they have increased their firing-rate). The vertical red line indicates when the 90% of the trials have changed state. **C**: Time averaged firing-rate of the singles trials taken in the window from the red line of panel B to the saccade; red lines represent the mean firing-rate over trials for "down" state (solid) and "up" state (dashed). **D**: The time evolution of the LRM coefficient. Shaded area corresponds to the period in which the coefficient was significant ($p < 0.05$ t statics). **E**: Average firing-rate as a function of TRO. The time average is taken during the time window marked by the shaded area D. Error bars represent SEM. **F**: Mean switch time (according to HMM analysis) as a function of TRO.

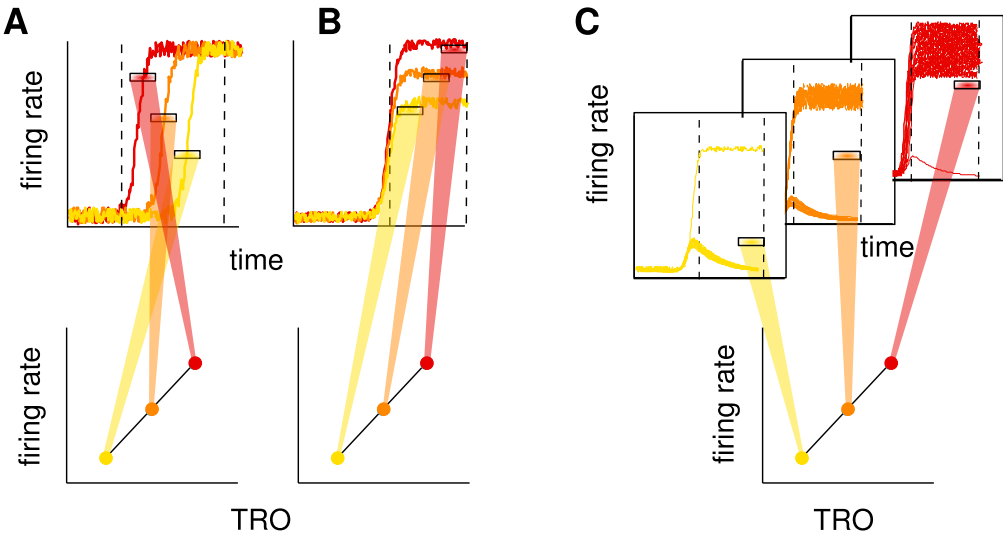

FIGURE 4.6: Pictorial representation of possible mechanisms underlying the confidence x-shaped pattern. The figure only illustrates a linear relation between firing-rate and test relative orientation (TRO), since it is the basis of the x-shaped pattern. **A**: switch timing code. In the upper panel the time evolution of the firing-rate is shown for three trials (one for each TRO; the color code is the same of the bottom panel). Each trial presents a switch from a state of low activity to a state of high activity. The three horizontal marks show the time averaged firing-rates taken in the window enclosed in the vertically dashed lines. In the bottom panel the average firing-rates are shown as a function of TRO. The different switch times of the trials produce different firing-rates. **B**: rate code; each trial reaches a different level of firing-rate in the high activity state (upper panel) and this is reflected in the mean firing-rate (bottom panel). **C**: binary code; only some trials switch to the high activity state while others remain in the "down" state. The number of trials that switch states depends on TRO. When many trials are in an "up" state (red trials) the mean-over-trials of the time averaged firing-rate is higher in respect to the case of many of the trials in the "down" state (yellow trials).

shed light on this matter. However, we only analyze correct trials because of the small number of errors in most of the recordings. The increased firing-rate as a function of the absolute value of TRO, i.e., the v-shaped pattern associated with correct trials, can arise from at least three distinct mechanisms (for a pictorial representation see Fig. 4.6). 1) Rate coding (panel A): neurons increase the firing-rate respect to the baseline in proportion to the confidence in the decision. 2) Switch time coding (panel B): neurons increase the firing-rate, switching from a low to a high activity state, with a different timing according to the confidence, and with the average rate reflecting this timing. 3) Binary coding (panel C): neurons have a binary response, i.e., they increase the firing-rate only in high confidence trials (whilst when confidence is low they remain in a "down" state). In this last scenario the proportion of high confidence trials depends on its level of difficulty, and mixing trials of high and low activity produces the v-shaped

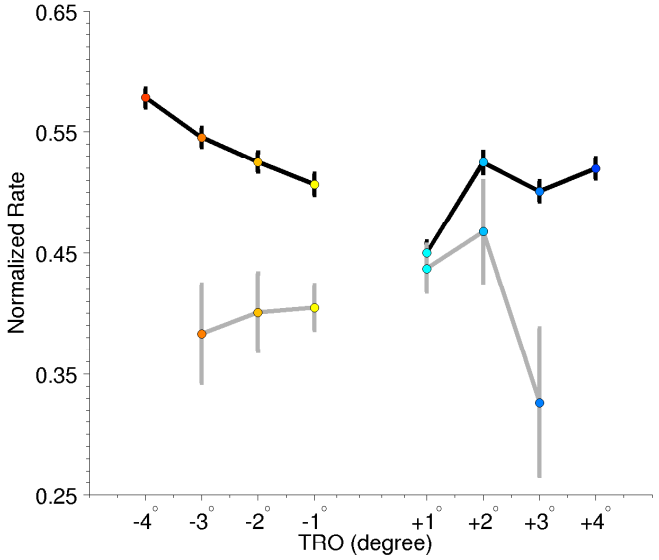

FIGURE 4.7: Average population activity. Normalized firing-rate as a function of TRO of the confidence neurons that present a positive relation between firing-rate and difficulty. Black lines: correct trials; gray lines: error trials. Every neuron firing-rate was normalized to its own maximum firing-rate. Bars represent SEM.

pattern of average firing-rates [131]. In order to identify neurons implementing each of these mechanisms we used different statistical techniques. Although we present them here as separated mechanisms, we do not rule out the possibility that they could all appear at the same time.

We started by verifying whether the switch timing had any relevant effect in our data. To do so we used a Hidden Markov Model (HMM) which is able to detect when a system switches from one state of activity to another (see section 4.4 for details). In Fig. 4.5B we show a summary of the two-state HMM analysis for one confidence neuron (each row represents a trial). The color of the row changes from white to black when the neuron goes from a low to a high-activity state. This neuron exhibits a lot of variability in the switch timing, changing state from just a few milliseconds up to 300 ms after stimulus onset. The timing of the change was correlated with the difficulty of the trial (Kendall's correlation coefficient $\tau = 0.18, p < 0.05$). Fig. 4.5F represents the mean switch time as a function of TRO. Once we had determined when a neuron changes its state we were then able to assess the relevance of the rate coding mechanism. The firing-rate after the state switch is represented in Fig. 4.5C for the same confidence neuron. Each dot represents the time-averaged firing-rate of one trial, color-coded according to the state assigned by HMM (for comparison the dashed and solid red lines represent the firing-rate of the high and low states founded by HMM respectively).

To estimate whether the increase in firing-rates was proportional to the difficulty (i.e., the rate coding mechanism of Fig. 4.6B), we first calculated the average firing-rate from when 90% of the trials switched states (red vertical line in Fig. 4.5B), until the coefficient of LRM had a significant value ($p < 0.05$, shaded region in Fig. 4.5D). Then we effectuated a correlation analysis between the level of difficulty and the average firing-rate. We obtained a significant correlation coefficient for the neuron in Fig. 4.5 ($\tau = 0.24$, Kendall's correlation, $p < 0.05$), which suggests that it could be the firing-rate of the neuron in the "up" state that encodes the trial's level of difficulty.

To summarize, we found that twelve neurons presented a significant impact on the timing in the formation of the pattern, while nine neurons increased the firing-rate proportionally to the difficulty of the trial, thereby implementing the rate coding mechanism. There were also five neurons that presented both switch timing code and rate code (see Fig. 4.9 for a graphical representation of all classes of neurons). We note that we could apply this method only to 28 out of 49 confidence neurons, as we considered the HMM analysis was only reliable under some constraints (see section 4.4).

4.2.2 Discrete confidence encoding

The binary mechanism postulated above corresponds to a discrete confidence encoding like the one hypothesized in [131]. Although a continuous representation of confidence is probable at some stage of the sensory-motor integration, a discretization stage is needed to account for the behavioral effect of the confidence computation [132]. Indeed, both the usual confidence ratings and confidence-related decisions [124] require a selection between different alternatives. Therefore this mechanism is not merely bounded to the difficulty of the task but to the confidence in the decision.

In order to identify neurons with a binary response we hypothesized that the distribution over trials of the mean firing-rate as calculated during the test-bar presentation, has to consist of two different distributions. The resulting distribution is not necessarily bimodal but it should differ substantially from the expected Poisson distribution [135–137]. For each trial, therefore, we took the average firing-rate over a 200 ms time-window, ending at the time of subject's response. Then we fitted these mean firing-rates to the average of two gamma distributions, parametrically varying the shape and the mean of the distributions:

$$\mathrm{B} = (F_\Gamma(x; k_1, \mu_1) + F_\Gamma(x; k_2, \mu_2))/2,$$

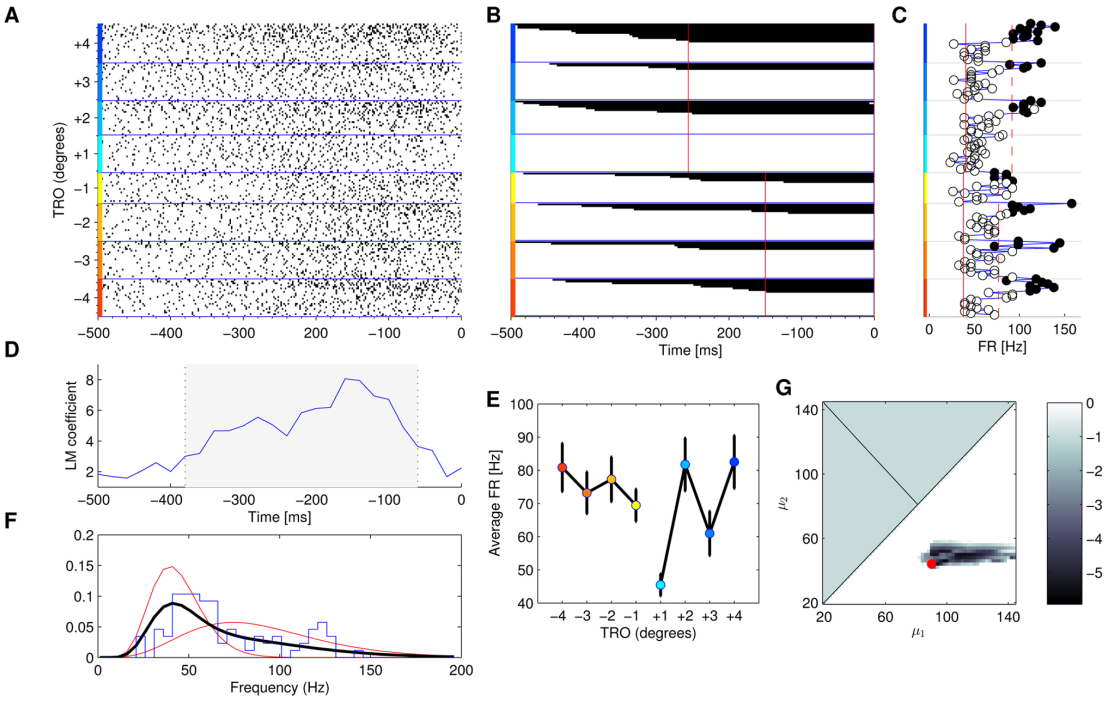

FIGURE 4.8: Single neuron from PMv, implementing a binary confidence encoding. Color and line convention is the same as Fig. 4.5. **A:** Raster plot. The trials are sorted as in 4.5. **B**: Time course of state switching, according to HMM. Same conventions as in Fig. 4.5B. In contrast to the neuron depicted in Fig. 4.5, not every trial of this neuron changed from state one to state two. This clearly indicates that the neuron has two distinct behaviors: in some trials it increases the firing-rate, while in others it remains at a lower level of activity. **C**: Average firing-rates of single trials. **D**: LRM coefficient value. The shaded area corresponds to the period in which the coefficient was significant ($p < 0.05$ t statics). **E**: Average firing-rate versus TRO. The time average was calculated during the period marked by the shaded area in D. **F**: Histogram of the firing-rate during the 200 ms previous to the saccade (in blue). In black the distribution that best fits the data. This distribution is the average of two gamma distributions shown in red. The fitting has been computed separately for positive and negative TRO; the picture corresponds to negative TRO values. **G**: χ^2 goodness-of-fit. The fitting was done in the four dimensions of the model parameters but here we show the results in the plane of μ_1, μ_2. Each point in the plane correspond to a pair of μ_1, μ_2 values. The color of the point represents the difference between the highest acceptable p-value (0.05) of the χ^2 and the p-value obtained with that set of parameter values. See section 4.4 for details. The white area indicates regions of the parameter values that give non-significant results. The red point represents the best fit.

where $k_{1,2}$ is the shape parameter and $\mu_{1,2}$ the mean. We used a gamma distribution because of its broad generality. We used a chi-squared test (χ^2) to evaluate the goodness of fit for each set of parameters (see section 4.4 for details). Fig. 4.8F shows the histogram of the firing-rate of one single neuron (blue line) and the best fit model (black line), composed of the two gamma distributions (red lines).

In our analysis we took all the possible combinations of the parameters of the two distributions into consideration. Therefore, it is possible that, if the actual empirical firing-rate distribution is unimodal (e.g. Poisson, normal, etc.), it could be well fitted by the mean of two gamma distributions with similar parameters. In order to eliminate this possibility we tested whether a model with four parameters (two for each distribution) was more adequate than a model with only two parameters (only one gamma distribution). To this aim we used the Bayesian Information Criterion (BIC) that, while comparing the likelihood function of the two models, corrects the result by penalizing for the number of free parameters. Therefore, even if the likelihood of the single distribution model were equal to that of the double distribution model, the BIC would always prefer the simpler model (or, conversely, a double distribution model would be preferable only if it was able to explain much more than the single distribution model).

In conclusion, we consider a neuron to have given a binary response if the chi-squared test gives a significant result and the model with a double distribution is better than the one with only a single distribution according to the BIC. Fig. 4.8G represents the goodness of fit (for significant values only) in the space of the means of the two gamma distributions (μ_1, μ_2). In this space the color of each point represents the difference between the highest acceptable value of the probability of the chi-square statistics ($p < 0.05$) and the actual value of this probability obtained for the combination of μ_1 and μ_2 of that point. It is interesting to note that the points where we get significant values do not lay on the diagonal (where the two means are identical). We ran this analysis on the entire confidence related population (49 neurons) and we found that eighteen neurons displayed a binary encoding of decision confidence in the case of at least one behavioral response (e.g., "left") (see Fig. 4.9 for a graphical representation of all classes of neurons). In these neurons the v-shaped pattern of the firing-rate is the result of the fact that the proportion of trials with high firing-rate correlates with the difficulty of the trial. We also reasoned that the neurons showing a binary behavior should also lead to a characteristic pattern showing up in the HMM analysis: they should present a state switch only on a subset of trials. And indeed this pattern can be seen in Fig. 4.8B. In Fig. 4.8C the mean firing-rate of the neurons is shown for each trial, separated in high firing-rate trials (filled dots) and low firing-rate trials

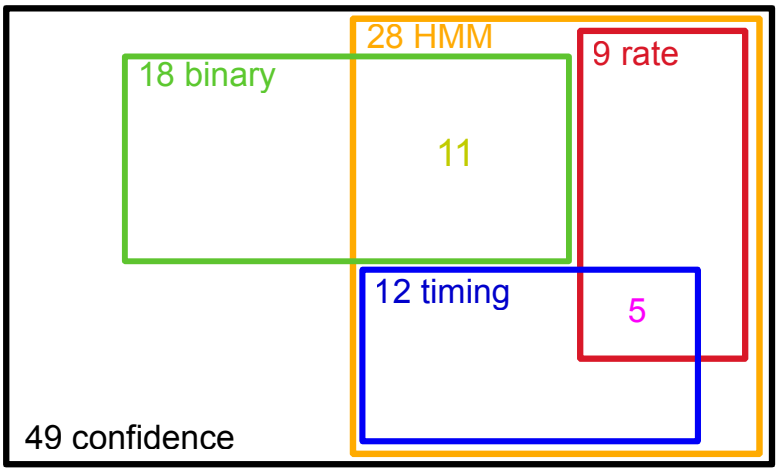

FIGURE 4.9: Graphical representation of the different classes of neurons. The label and number in each rectangle indicate the class (rate code, switch timing code and binary code) and the number of neurons in that set. Five neurons are in the intersection between the rate and switch time populations.

(open dots). The separation of firing-rates can be clearly seen when compared to the continuous confidence encoding neurons (see Fig. 4.5C).

4.2.3 PMv encode decision and task difficulty

Our aim with this section is to complete the vision of the data-base analised. In this section we sumarize the dinamics of the neurons which activity is related with the behavioural response of the monkey and the task difficultly at the same time.

In order to find these neurons we used a variation of the linear regression model (LRM) used before (see section 4.4). As already reported, [8, 19], we found that the firing rate of decision neurons in VPM cortex encodes the difficulty of the task. We identified 192 neurons (57% of all recorded neurons, n=336) whose activity correlates with the decision of the monkey during the presentation of the second bar which also encode the task difficulty. Out of these, 80 neurons had a high firing rate when the subject took the decision "right", while they had a low firing rate in "left" decision. We call this the "right selective population". Another population (112 neurons) was selective for "left" responses.

In Fig. 4.10 we show the activity of a single neuron encoding both decision and trial difficulty. This neuron was selective for "right" decisions, as can be observed in the raster plot and the PSTH (panels A and B). Fig. 4.10C shows the coefficient of the LRM as a function of time. The shaded area marks the time window where the coefficient was significant. In this same time window we took the average firing

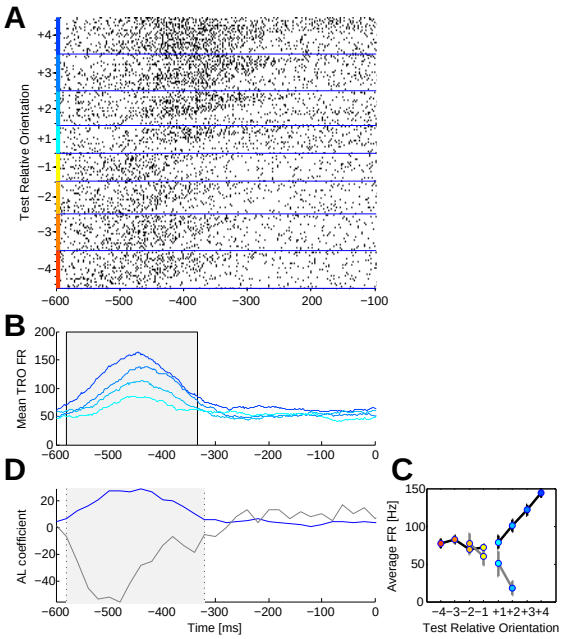

FIGURE 4.10: Single neuron from PMv, whose activity correlates with the task difficulty. **A:** Raster plot during the visual discrimination task. This neuron fires when monkey chose the right target, its activity increases depending on the difficulty of the task. Each row of points is a correct trial, and each point is a spike. The trials are sorted by the TRO. **B:** Time course of mean firing rates sorted by TRO. **C:** Resulting coefficients values of LRM, (c_1) for correct/incorrect trials for neuron in A, as functions of time. In blue c1 correponding to the correct trials, in grey c_1 for the incorrect. Shaded area corresponds to the period in which the c_1 for the correct trials was significative ($p < 0.05\,, t-tstatics$). **D:** Average Firing rate during the shaded area period in B/D, versus TRO, in blue for the correct trials in grey for the incorrect trials (error bars indicate SEM).

rates shown in Fig. 4.10D. This procedure gives us the advantage of choosing the precise interval where the neuron is responsive for the feature of interest. Indeed the firing rate during the significant time window plotted as a function of TRO (see Fig. 4.10D) shows a clear dependence on stimulus difficulty. The corresponding population activity is shown in Fig. 4.11 *Upper Rigth*, confirming the activation pattern at the population level. Error trials are very few to be analyzed on a single neuron bases, therefore we grouped the activity of all neurons in order to get more statistical power. When the responses of all neurons are normalized and pooled together they show indeed a clear tendency to decrease when stimulus gets easier (see Fig. 4.11).

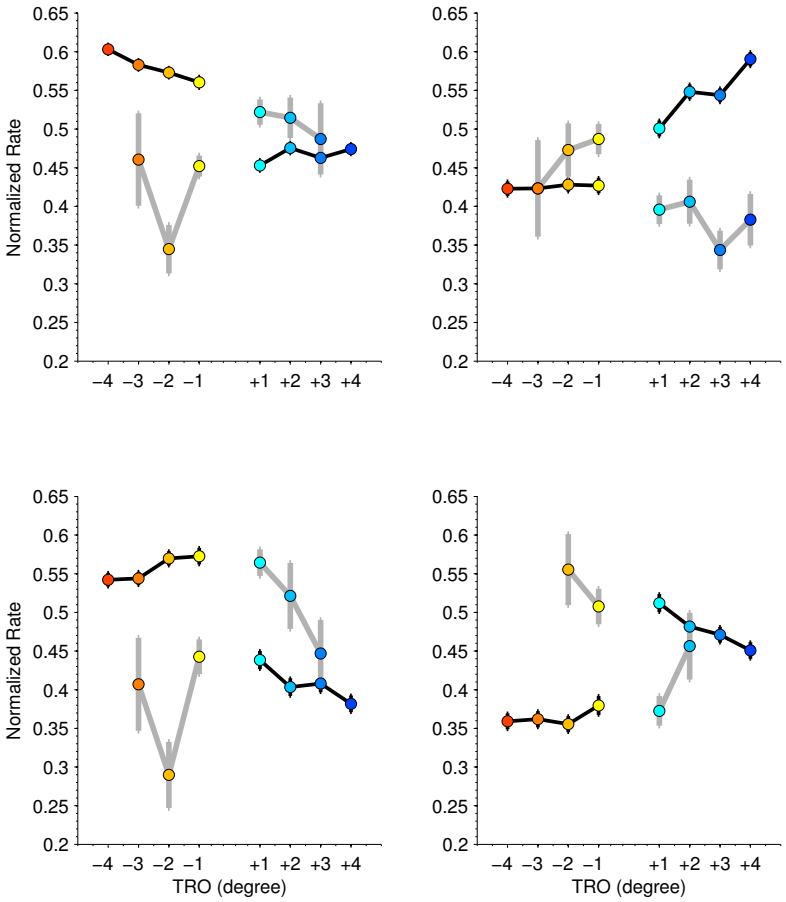

FIGURE 4.11: Average population activity. Normalized firing-rate as a function of TRO of the decision neurons, which also encode the task difficultly. We represented the 4 possible population: positive/negative relation between firing-rate and task difficulty for both decisions left/tigth. Black lines: correct trials; gray lines: error trials. Every neuron firing-rate was normalized to its own maximum firing-rate. Bars represent SEM.

4.3 Discussion

In this study we address the question of how neurons in the primate brain encode decision confidence in single trials. While evidence found in rat OFCs seems to suggest that confidence is encoded by the continuously varying firing-rate of neurons, recordings made in monkey LIPs show that confident trials and uncertain trials do not display the same pattern. Continuous encoding schemes and discrete encoding schemes involve different computations and probably serve different functions; it is therefore important to know which of the two (or both) is implemented in the brain.

We have demonstrated that, during correct trials, neurons in primate PMv increase their firing-rate as a function of stimulus discriminability (in our experiments: the relative orientation of a bar), whereas in error trials the firing-rate decreased. This

peculiar pattern has already been described by [124] as a correlate of confidence in the rat OFCs. Here, for the first time, a similar result in monkey PMv cortex is presented.

It is worth noting that the pattern emerges when the firing-rates of neurons over several trials, and with the same discriminability, are averaged together. Nonetheless, different computations performed by neurons in single trials can produce the same pattern of average firing-rates. We suggested three hypothesis: 1) The switch time coding: when the activity of the neuron changes, the confidence of the decision, is encoded in the timing of the change; 2) The rate coding: the confidence is encoded in the firing-rate, after the change has taken place; or 3) The binary coding: the neuron only changes activity in high (or low) confidence trials and the proportion of high confidence trials changes according to the discriminability of the stimulus. The first two alternatives correspond to a continuous encoding of confidence, whereas the last one is a form of discrete encoding. We found that, in fact, all three mechanisms are at work in monkey PMv neurons. For certain neurons the timing and firing-rate mechanisms work together, i.e., a neuron that changed state earlier on less difficult trials will also have a higher firing-rate after the change. Other neurons present a binary response (increasing activity only in some trials), which suggests a possible role in confidence judgments.

An important question is: why should neurons use different schemes to encode confidence? Our hypothesis is that confidence neurons carry-out more than one function in the sensory-motor path. It is logical that a "confidence" module would encode decision confidence on a continuous scale, since confidence is a graded sensation. However, if decision confidence is to have behavioral relevance, the information about confidence needs to be discretized (see Fig. 4.3). Our hypothesis is that, while certain neurons encode confidence in a continuous manner, other neurons read-out this scale and transform it into a discrete quantity in order to produce consistent behavior.

This idea has been partially implemented in the biologically realistic attractor neural network of [131] and in the DDM framework of [132]. Indeed, [131] have shown that the sum of the firing-rate of the decision neurons is a good representation of the confidence on a continuous scale. This representation is given as input to another decision-making network that then makes a choice based on the confidence estimation. Therefore, our results of binary confidence neurons confirm the predictions of the model about a confidence-related decision-making module. However, we must note that the model does not explicitly consider the continuous confidence representation as found in PMv neurons. In the model, confidence is

represented implicitly by the sum of decision-neurons' firing-rates, as no neuron seems to encode it directly.

Despite the results as presented in this article, it could be objected that, like in [124], the task has no confidence measurement, and that reasonable doubts could thus be cast on the interpretation of the results. Although we think that a set-up where both choice and confidence are recorded would better serve the scope, we note that, rather more intriguingly, we found confidence encoding neurons in a context where confidence estimation was not relevant. Indeed, it stands to reason that a mechanism that estimates confidence and that makes decisions based on this would stay in place even when not in use. We therefore expect that a more suited task will lead to the recording of even stronger signals of confidence. Moreover the population activity, as shown in Fig. 4.7, presents the hallmark of confidence, excluding that these neurons were merely encoding the difficulty of the trial. In addition, although we demonstrate the three encoding schemes only for correct, the binary neurons present two different states of activity that can not be explained as a simple encoding of the difficulty. Conversely these binary neurons can be interpreted as confidence neurons of the type predicted by the model of [131].

Most of our results depend on a linear model of the firing-rate. But does this relation have to be linear (and not, for example, logarithmic or sigmoidal)? Firstly we note that linear functions have been extensively used to model the relation between the firing-rates of neurons and certain task features (e.g. [7, 14, 138]). Yet it is possible for the relation not to be linear. Indeed, we consider the linear function as a first probable approximation.

The three mechanisms underlying the confidence x-shaped pattern that we have suggested, raise the question of whether PMv neurons change their firing-rate gradually, or whether they jump from a low to a high activity state. This question, which has often been raised concerning the decision neurons of the lateral intraparietal sulcus (LIP), has been bothering the scientific community for some time now [6, 96, 139, 140]. Recently, [141] have provided reliable evidence for the hypothesis that LIP neurons display a gradual ramp. Although our analysis was aimed at differentiating single trial mechanisms, we did not address this issue. We do note that all three proposed mechanisms are compatible with both a gradual and an abrupt transition of states.

In this article we have abstained from directly defining decision confidence. We agree with [9] that a formal computational foundation of the phenomenon is far more valuable than a semantic definition. The semantics surrounding concepts involving subjective experience (such as consciousness and confidence) is often

inadequate to define the underlying phenomena [142]. A computational account thus aims at characterizing the system in terms of functions and computational units [143]. As it has been presented here, the concept of decision confidence probably needs revising. We may thus separate the neural representation of the reliability of a decision, call it 'decision confidence, from the classification of this sensation, be it a verbal rating or a post-decision wagering, which we can call 'confidence judgment'. Decision-confidence would be a continuous quantity that depends on the decision process and on sensory input. Confidence judgment, then, would be of a discrete nature that depends on decision confidence as well as on the constraints of the required judgment.

4.4 Methods

All analyses were performed using custom-made programs in MATLAB®. Unless noted otherwise, all statistical analyses were applied to the firing-rates of single neurons during the 500 ms preceding the saccade. In fact, the second stimulus was presented during this period, and therefore the decision-making process was expected to take place during this time window.

Our first aim was to found any existing neurons whose activity relates to:

1. Difficulty of the task: Neurons whose mean firing-rate for the correct trials increases linearly with the difficulty. When plotted against the TRO the activity of these neurons shows the v-pattern.

2. Confidence: Confidence measures present a characteristic x-shaped pattern when plotted against the signed difficulty of the task [1, 9, 124]. Therefore, when a difficulty neuron had enough error trials to be analyzed and when its firing-rate in error trials showed a mirror modulation of the difficulty respect to correct trials we considered it a confidence neuron (i.e., neurons that are a subset of the difficulty neurons)

In order to accomplish this we used a linear regression analysis (LRM) [144]. We applied a linear measure instead of a non linear measure of correlations (like MI), because at the end we are looking for a v-paterns and x-patterns, which in both cases are a linear relationship between the rate and the difficulty.

Our second aim is to try to shed light into the mechanism by which these neurons produce the above modulation of the difficulty neurons. We individuated three

possible neural mechanisms responsible for modulations; a simplified representation of these mechanisms is presented in Fig. 4.6. In order to understand which difficulty neuron belongs to each of the three categories, we applied two methods:

1. In order to find neurons that switch states with a timing dependent on the difficulty (i.e., signaled the decision with a change in activity), we used the Hidden Markov Model (HMM) analysis [145]. Indeed, the HMM was able to cluster the spiking activity of individual neurons into periods of 'stationary' activity (the states) within a single trial. Hence the switch time between states could be estimated.

2. In order to find neurons whose activity after the change correlates whith the difficulty we calculated the correlation between the mean activity and the difficulty of the task. The mean activity was calculated in the time window starting at the time bin where the 90% of the trials had passed from one state to the other and ending at the last significant time bin marked by the LRM.

3. In order to find the neurons whose activity could be explained as a compound of high and low firing-rate states we fitted (χ^2 goodness-of-fit test) the firing-rate distribution to the average of two gamma distribution functions.

4.4.1 Linear regression analysis

The firing-rate (FR) of the last 500 ms before the saccade was computed by averaging the spike count in a sliding window of 100 ms slided with a step of 20 ms. In this way we got for each trial and each neuron a time series $r(t)$ of the firing-rate, where t is time discretized in 25 time bins.

To individuate the neurons presenting a modulation of task difficulty (v-shaped modulation) the following LRM analysis was used, $r(t) = d_1 \, |S1 - S2| + d_2(t)$, where d_1, d_2 are the parameters to be fitted. Activity was considered linearly dependent on the difficulty, $|S1 - S2|$, if the coefficient d_1 was different from 0 ($p < 0.05 \, , t - statistics$) and no sign switch occurred during a time interval (T) with a length of at least four consecutive bins (140 ms). Details as to how T was chosen are given below.

The population activity in Fig. 4.7 shows that these neurons don't encode merely the difficulty but the confidence in the decision. Indeed their activity has the opposite modulation in error trials which can not be explained by a difficulty

encoding mechanism. However to confirm that this pattern is verified by single neurons we applied the above LRM analysis to both correct and error trials, but we kept the two analyses separate. Therefore, the pairs of parameters d_1, d_2 are different for error (d_1^e,d_2^e) and correct (d_1^c,d_2^c) trials. We looked for neurons whose firing-rate presented a mirror modulation of difficulty in error trials compared to correct trials, hence we considered a neuron as encoding confidence if the sign of d_1^c was the opposite of that of d_1^e. Due to the low number of error trials, it was only possible to analyze fifteen of the 49 difficulty neurons; only five satisfied all of the constraints described above. In order to produce Fig. 4.7 the firing-rate of each neuron was normalized to its maximum value and then the activity of all neurons was averaged together.

Error trials

In order to identify confidence neurons we independently applied the LRM to correct and error trials. Unfortunately, the number of error trials was not enough to analyze all the neurons. As experiments were done using animals that were awake it was very difficult to record single neurons over a long period. We recorded approximately 10 trials per monkey and stimulus conditions (i.e., orientation of $S1$, $S2$). Hence only few error trials were recorded under easier conditions ($|S1 - S2| = 3, 4$). In the end only 124 neurons had at least one error for difficult categories ($|S1 - S2| = 1, 2$). Therefore, we were only able to run the LRM methods (described below) on this subset.

Minimal time window (T)

In order to find the minimum length of T we proceeded as follows. Given that linear regression has a p-value cutoff at 0.05 in each bin there is a probability of 0.05 to get a false positive. We wanted to know what the probability P_n is of getting n consecutive false positives. We then selected n such that $P_n < 0.05$.

In order to calculate P_n we proceeded as follows. A statistical test with $p < \alpha = 0.05$ applied to a time series produced a time series of significant and non-significant bins. The vector X representing the time series is generated by:

$$X = \{x_1, \ldots, x_{25} \,|\, x_i \in [0, 1] \,, P(x_i = 1) = \alpha, P(x_i = 0) = 1 - \alpha\}$$

where x_i takes value 1 when the ith bin is significant by chance and 0 otherwise. X has length 25 since our time series ($r(t)$) has 25 bins. We generated 10^6 vectors with this procedure and then evaluated the probability P_n of having n consecutive ones. P_n is thus the probability that an i exists such that $x_i + \ldots + x_{i+(n-1)} = n$. Since we ran the test on a large number of neurons we corrected P_n for the family wise error-rate: $P_n^N = (1 - (1 - P_n)^N)$, where N is the number of neurons (336). Then we found n such that $P_n < 0.05$. We found that the minimum number of consecutive bins needed to get a significant result was $n = 4$. The applied method gave the same results of the more common Bonferroni correction.

Decision Neurons

To identify the decision neurons we applied a LRM using the equation $r(t) = b_1(t)S1 + b_2(t)S2 + b_3(t)$. A neuron was considered as decision during a time interval T the coefficients b_1, b_2 were different from 0 ($p < 0.05$, t-statistics), of opposite sign ($sign(b_1) = -sign(b_2)$) and they never switched their sign during the whole T. Note that if a neuron satisfies these constrains the information about the response can be extracted by comparing the coefficients signs, i.e, the sign coefficients determine the selectivity of the neuron to the behavioral response: It is classified as "left selective" if $b_1 < 0, b_2 > 0$ and "right selective" if $b_1 > 0, b_2 < 0$. We impose the sign constant because we require neurons whose selectivity remain constant in time. Once we have individuated the decision neurons, we want to test if they also encode the difficulty. For this purpose we run a LRM analysis, with equation $r(t) = c_1(t)(S1 - S2) + c_2(t)$. Where c_1, c_2 are the parameters to be fitted. The main difference with respect to the previous analysis is that here we used for each neuron only those trials that are in correspondence to its selectivity, i.e. for 'left selectivity' neurons we used only trials to-the-left. Neuron response was considered linearly dependent on difficulty if the coefficient c_1 was different from 0 ($p < 0.05\,, t-statistics$), and during T it never switched its sign. The mean firing rate of the trials whichbelong to the neuron's selectivity, increased/decreased the activity level with the difficulty during T. If c_1 is positive the mean firing rate vs difficulty produces just a branch of the v-shaped (see Fig. 4.10C).

4.4.2 V-mechanism methods: Hidden Markov Model

To analyze the single-trial activity of the recorded neurons we used the HMM that clusters the spiking activity of individual neurons into periods of stationary activity

within a single trial. The HMM technique has been successfully applied to characterize the single-trial activity of cortical neuronal ensembles during movement with holding and preparation [146, 147], taste processing [148], and perceptual decision-making [140]. Here, we briefly review some aspects of the HMM analysis; more details about the algorithms can be found in previous works [140, 146, 148].

Within the HMM, the activity of a recorded neuron at time t is assumed to be in one of a (predetermined) number (Q) of hidden firing-rate states. In each state q, the discharge of a neuron is assumed to be a Poisson process of intensity λ_q, which defines the instantaneous firing probability E_q, i.e., the probability of firing a spike within one time bin, equal to 2ms throughout this study. States are said to be hidden because they are not directly measured; instead, we observe the stochastic realizations of the state-dependent Poisson process (observation sequences). The state variable changes from state i to state j with fixed probabilities that defined a transition matrix A, given by $A_{ij} = P(q_{t+1} = j|q_t = i)$, where q_t is the state at time t and $i, j \in \{1, \ldots, Q\}$. The entire process is a Markov chain: the transition probabilities A_{ij} are independent of time, i.e., they depend only on the identities of states i and j, which means that the state sequence at time t only depends on the state at time $t - 1$. In summary, for a single neuron the HMM is fully characterized by the spike-emission probabilities (E) and the transition matrix (A). These model parameters are estimated from the data, using a likelihood expectation-maximization algorithm [140, 146, 148].

Briefly explained, the procedure starts with random values for E and A and re-estimates the parameters to maximize the probability of observing the data given the model. After optimization of the model parameters, the Viterbi algorithm is used to find the most likely sequence of hidden states given, for each single trial, the model and the observation sequence [140]. In the present study we used the HMM to detect the transitions between a state of low and a state of high activity. For this reason, the number of states was set to $Q = 2$. For each neuron, the data was divided into two subsets, composed of trials corresponding to each behavioral response (left or right). For each subset, a HMM was estimated using the activity of 80% of the trials (randomly selected) during the period within the last 500 ms before the saccade. After optimization the most likely state sequence was stored for all trials.

Unfortunately, a HMM analysis was not reliable for all the neurons. We only considered the HMM reliable if 1) The duration of both states was at least 25 ms. (i.e., we do not take into account states with very brief duration) B) The number of state-switches per trial was three or less; or (i.e., we do not take into account bursting neurons). C) At least five of both the left and right oriented

trials had a state-switch (i.e., we want neurons with 2 different states). We found 45 confidence neurons (out of 49) whose HMM was interpretable. For this subset, we wanted to distinguish between the three v-shaped mechanisms, to do so we analyzed the state-switch time. For each trial the HMM gave the time in which it changed from a low to a high state (or vice versa).

4.4.3 V-mechanism methods: Bimodality vs Unimodality

Our aim was to investigate whether the firing-rate distribution of neurons during correct trials was better described using a bimodal than a unimodal function distribution. The procedure we applied was the following:

1. Trials were divided into two sets, depending on their behavioral responses ("left" or "right"). We calculated the average firing-rate for each trial in a 200 ms time window that ended at the time of subject's behavioral response. We called the empirical distribution functions of the average firing-rate $F(\nu_L)$ and $F(\nu_R)$ respectively. The distributions were fitted with a function B, which was the average of two gamma distributions: $B = (F_\Gamma(x;\kappa_1,\mu_1) + F_\Gamma(x;\kappa_2,\mu_2))/2$. Gamma distribution was chosen because is one of the most general function distributions with positive support. The gamma distribution is given by:

$$F_\Gamma(x;\kappa,\mu) = \frac{1}{(\mu/\kappa)^\kappa}\frac{1}{\Gamma(\kappa)}x^{\kappa-1}e^{(-\frac{x\kappa}{\mu})}$$

   for $x,\kappa,\mu > 0$; Γ is the gamma function; κ is the shape parameter and μ is the mean.

2. We looked for best fit (in terms of χ^2 goodness-of-fit, $p < 0.05$) using the following parameter space: $\mu_d \in [min(F(\nu_d)), max(F(\nu_d))], \kappa_d \in [0.1, 10]$, for $d \in (1, 2)$.

   In order to apply the goodness-of-fit test, we discretized the firing-rate distribution in bins of 5 Hz. The χ^2 goodness-of-fit probability test is valid under the assumption that the number of events in each bin is greater than five. Whenever this condition was not satisfied we enlarged the bin on the right until the event count was at least five.

3. Finally, in order to verify whether a bimodal model explains the data better than a model with just one mean, we fitted the data to a single Gamma

distribution and then we compared the two models using the Bayesian Information Criterion (BIC) [149, 150], which is given by:

$$BIC = -2 \cdot \ln L + p \ln(T)$$

where L is the maximized value of the likelihood function for the estimated model; p the number of free parameters of the model (2 or 4); and T the length of the observation data (the number of bins). This means that the BIC method penalized the model likelihood by a measure of its complexity (i.e., the number of free parameters). The single mode model has two free parameters, while the bimodal model has four free parameters. Therefore, in order to have a better score in the BIC, the higher complexity due to the second mode should really be well balanced by a better ability to explain the data. Hence we considered that a neuron has a binary response if the BIC favored the model with two modes.

4.4.4 Experimental Details

The Visual Discrimination task

Experiments were made using two male monkeys (Macaca mulatta). Animals (BM5, 8 kg; and BM6, 6 kg) were handled according to the standards of the European Union (86/609/EU), Spain (RD 1201/2005), and the Society for Neuroscience Policies and Use of Animals and Humans in Neuroscience Research. The experimental procedures were approved by the Bioethics Commission of the University of Santiago de Compostela (Spain).

The monkeys' heads were immobilized during the task and looked binocularly at a monitor screen placed 114 cm away from their eyes (1 cm subtended 0.5 to the eye). The room was isolated and soundproofed. Two circles (1° in diameter) were horizontally displayed 6° at the right and 6° at the left of the fixation point (a vertical line; 0.5° length, 0.02° wide) displayed in the screen center. The monkeys used right and left circles to signal with an eye movement the orientation of visual stimuli to the right and to the left, respectively. Orientation Discriminations Task: the monkeys were trained to discriminate up to their psychophysical thresholds in the visual discrimination task sketched in Fig. 4.4A (training lasted for approx. 11 months). The stimuli were presented in the center of the monitor screen and eye movements larger than 2.5° aborted the task. The orientation discrimination task was a two-interval, two-alternative forced-choice task. A masking white noise

signaled the beginning of the trial and then the fixation target (FT) appeared in the center of the screen (see Fig. 4.4A). The monkey was required to fixate the FT. If fixation was maintained for 100 ms, the FT disappeared, and, after a variable pre-stimulus delay (100 - 300 ms), two stimuli ($S1$ and $S2$), each of 500 ms duration, were presented in sequence, with a fixed inter-stimulus interval (1 s). At the end of the second stimulus, the subject made a saccadic eye movement, in a 1200 ms time window, to one of the two circles, indicating whether the orientation of the second stimulus was clockwise or counterclockwise to the first. Trials lasted approx. 3.5 s separated by a variable intertrial interval ($1.5 - 3$ s). Fifty milliseconds after the correct response, a drop of liquid was delivered as a reward. A modulation of the masking noise signaled the errors; the modulation started 50 ms after the incorrect response and lasted for 75 ms.

Monkeys' weights were measured daily to control hydration, and once a week the animals had access to water ad libitum. The level of training was assessed by the psychometric functions. Once trained, the monkeys performed around 1000 trials per day. The lines were stationary, subtending 8° length and 0.15° wide. Three different S1 orientations were used for each monkey during the recordings: 87°, 90°, and 93° (BM5) and 84°, 90° and 96° (BM6); all angles referred to the horizontal axis. Different $S2$, eight per $S1$, were presented, four clockwise and four counterclockwise to $S1$ in steps of 1° (BM5) and 2° (BM6). More details can be found in [7].

Recordings

Neuronal population: extracellular single-unit activity was recorded with tungsten micro-electrodes (epoxylite insulation, 1.5-3.5M, catalog # UEWMGCLMDNNF; FHC) in the posterior bank of the ventral arm of the sulcus arcuatus and adjacent surface in the ventral premotor cortex in the four hemispheres of the two monkeys (see [7], for a detailed description of the recording sites). In this work, we studied the responses of a subset (336) of the recorded neurons. This subset was selected with a ROC analysis of firing-rate respect to the choice (see [7] for details).

Chapter 5

Causal correlation paths in decision-making

The work presented in this chapter was carried out by Adrià Tauste Campo and the author of this thesis in collaboration with: Verónica Nácher, Ranulfo Romo and Gustavo Deco.

5.1 Introduction

A fundamental problem in neuroscience is to understand how neural activity encodes, integrates and communicates information across different brain areas. An ideal paradigm to study this problem is the *perceptual discrimination task*.

As we described in the introduction, Chap. 1, this tasks involves a number of processes which require communication across the cortical areas; from the areas which perceive the stimuli, to the areas which make the decision and to the areas that report the decision. The purpose of this chapter is to study the causal paths that arise between sensory and motor areas within the realization of the vibrotactile decision-making task and also during the *control task*.

To study this problem we resort again to the *vibrotactile discrimination* task designed by Romo [8, 18] (see sections 1.1, 2 and Fig. 5.1). Briefly, the subjects must compare two different frequencies (f_1, f_2) and then report their decision ($f_1 < f_2$ or $f_1 > f_2$) by a saccade. During the task the brain must store in memory the traces of the first frequency, then perceive the second one and compare both to make a decision. The processes of perception, memory and comparison are not

overlapped in time, and we can thus study them independently. In this chapter we also analyse the *control task*. In this task, the monkey receivep both mechanical vibrations at the same time periods of the discrimation task, but was requested to remain still upon reward (the reward time was not fixed, it arrived at different time instances after f_2).

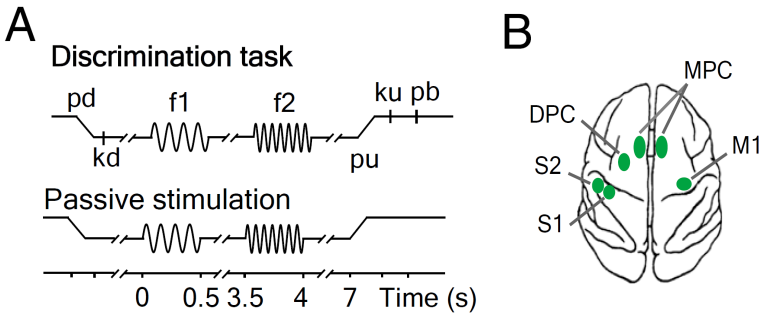

FIGURE 5.1: **A:** Sequence of events of the vibrotactile discrimination and for the passive stimulation tasks. **B:** Top view of macaque brain highlighting recorded cortical areas. (f1, first stimulus; f2, second stimulus; kd, key down; ku, key up; pb, push button; pd, probe down; pu, probe up; Primary somatosensory cortex (S1), secondary somatosensory cortex (S2), medial premotor cortex (MPC), dorsal premotor cortex (DPC), and primary motor cortex (M1). Figure adapted from [151].

Previous studies on this dataset have mainly focused on the correlation of single-neuron responses with different aspects of the task such as stimuli, memory or decision ([8, 18], see sections 2 and 1.2.1). The results show that, while stimuli are mostly encoded in sensory areas, the comparison between stored and ongoing sensory information is reported to take place in a distributed fashion, and in particular, can be identified as the unique site of decision-making.

We measure the causal correlation between sets of simultaneous spike trains in sensory and motor areas. Our approach can be summarized into two steps. First, we model spike-train responses as binary sequences. By using this model, we characterize both the amount of neuronal activity and the inter-spike interval (ISI) of spike trains. Second, we make use of a non-parametric method to estimate the causal correlation between binarized neural responses based on an information-theoretic measure called directed information. In the next section, we describe the estimation method.

The main features of our method are:

- The data available was simultaneously recorded.

- We model the neural spike trains as binary series without making any assumption on the neural.
- Our correlation measure is non-linear and directional.
- The algorithm used to estimate the correlation measure follows a Bayesian non-parametric approach.

5.2 Directed information as a measure of causal correlation

Since the later 90's, information theory concepts and techniques have been extensively used in neuroscience [152, 153]. In most cases, the measure of mutual information has been employed either to measure non-linear correlations between pairs of neural responses [154] or between a neural response and a stimulus ([153, 155], and sections 2, 3). As it has been earlier introduced in this thesis, the mutual information between any pair of variables tells us indistinctively the added knowledge that one has about one variable by knowing the other. It is therefore a symmetric function of the pair of variables [21]. However, this property is not suitable to characterize the feed-forward and feedback correlations that arise in neural systems. Instead, a more natural approach is to use a non-symmetric correlation measure such as the *directed information* (DI) [156].

The directed information definition is built upon from the ideas of Granger [157] on causal interactions, and it is philosophically grounded on the following principle: the extent to which X^T statistically causes Y^T is measured by how much help the causal side information of process X^T provides when predicting the future of Y^T, given knowledge of Y^T's past [157].

5.2.1 Definition

Directed information was placed in an information-theoretic perspective by Massey [156] to give a meaningful notion of directionality to the information flow through a channel. He defined the directed information of a sequence $X^T = (X_1, \ldots, X_T)$

to $Y^T = (Y_1, \ldots, Y_T)$ as

$$I(X^T \to Y^T) \triangleq \sum_{t=1}^{T} I(Y_t; X^t | Y^{t-1}) \tag{5.1}$$

$$= \mathbb{E}\left[\sum_{t=1}^{T} \log \frac{P_{Y_t|Y^{t-1},X^t}(Y_t|Y^{t-1}, X^t)}{P_{Y_t|Y^{t-1}}(Y_t|Y^{t-1})}\right]. \tag{5.2}$$

Similarly to the mutual information (see section 1.2.2), directed information can be expressed as a difference of two entropies:

$$I(X^T \to Y^T) = H(Y^T) - H(Y^T \| X^T), \tag{5.3}$$

where $H(\cdot)$ is the usual vector entropy function, defined by

$$H(Y^T) \triangleq \sum_{t=1}^{T} H(Y_t|Y^{t-1}), \tag{5.4}$$

and $H(\cdot\|\cdot)$ is the causal conditional entropy [158], defined by

$$H(Y^T \| X^T) \triangleq \sum_{t=1}^{T} H(Y_t|X^t, Y^{t-1}), \tag{5.5}$$

where X^T and Y^T are assumed to be finite-memory Markov chains that are stationary and ergodic.

Directed Information and Mutual Information The mutual information measure between two time series can be defined as:

$$I(X^T; Y^T) \triangleq \sum_{t=1}^{T} I(Y_t; X^T | Y^{t-1}). \tag{5.6}$$

The main difference between MI and DI is that the MI is a symmetric measure while the DI is not. If we compare Eq.(5.1) and Eq.(5.6),

$$I(X^T \to Y^T) = \sum_{t=1}^{T} I(Y_t; X^t | Y^{t-1})$$
$$I(X^T; Y^T) = \sum_{t=1}^{T} I(Y_t; X^T | Y^{t-1})$$

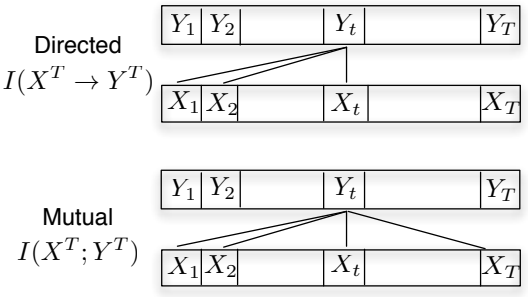

FIGURE 5.2: Relationship directed information vs mutual information.

we realise that the difference lies on the following fact: the MI takes into account the past and the future of X^T, while the DI only take into account the past of the X^T. By inspection of (5.6), it can be checked that $0 \leq I(X^T \to Y^T) \leq I(X^T; Y^T)$ since $I(Y_t; X^t | Y^{t-1}) \leq I(Y_t; X^T | Y^{t-1})$, for every $t = 1, \ldots, T$.

Moreover, it can be proved that [159]

$$I(X^T; Y^T) = I(X^T \to Y^T) + I(Y^{T-1} \to X^T), \tag{5.7}$$

i.e., the mutual information can be expressed as a sum of directed information from X^T to Y^T and a *reversed* directed information from the delayed sequence Y^{T-1} to X^T.

5.2.2 Related causal correlation quantities

Directed Information and Granger Causality Directed information generalizes upon previous approaches to quantify causal dependencies between pairs of time series. The original computation of causal dependency due to Granger assumes that the two time series are stationary Gaussian and are linearly correlated [157]. Then, Granger's causality computation results from applying a linear multivariate autoregressive model fitted to the two time series.

DI and Transfer Entropy A more closely related measure to the directed information is the transfer entropy developed by Schreiber [160]. Unlike the directed information, the transfer entropy is sustained on two assumptions. First, it assumes that the processes X^T and Y^T satisfy the Markov property

$$P_{Y_t|Y^{t-1},X^{t-1}}\left(y_t\middle|y^{t-1},x^{t-1}\right)=P_{Y_t|Y_{t-J}^{t-1},X_{t-K}^{t-1}}\left(y_t\middle|y_{t-J}^{t-1},x_{t-K}^{t-1}\right), \tag{5.8}$$

for every $t=\max(J,K)+1,\ldots,T$, which means that we require that the value of Y_t only depends on the past J values of its process and the past K values of X^T at time t. Second, it assumes that

$$I(X_t;Y_t|X^{t-1},Y^{t-1})=0, \tag{5.9}$$

for every $t=1,\ldots,T$, i. e., there is no instantaneous dependency conditioned on the past of each process. Then, the transfer entropy from X^T and Y^T is defined as [160]

$$\begin{aligned}\mathrm{TE}(X^T\to Y^T)&\triangleq I\left(Y_t;X_{t-K}^{t-1}|Y_{t-J}^{t-1}\right) &(5.10)\\ &=H\left(Y_t|Y_{t-K}^{t-1}\right)-H\left(Y_t|X_{t-J}^{t-1},Y_{t-K}^{t-1}\right), &(5.11)\end{aligned}$$

for every $t=\max(J,K)+1,\ldots,T$. By expressing the directed information under the same assumptions as the transfer entropy it can be checked that

$$I(X^T\to Y^T)\leq \mathrm{TE}(X^T\to Y^T), \tag{5.12}$$

with the equality holding when $I(Y^t;Y^{t-J-1}|Y_{t-J}^{t-1})=0$ for every $t=\max(J,K)+1,\ldots,T$. Finally, we note that a number of works [161, 162] use a normalized version of the transfer entropy, defined by

$$\mathrm{NTE}(X^T\to Y^T)\triangleq\frac{H\left(Y_t|Y_{t-K}^{t-1}\right)-H\left(Y_t|X_{t-J}^{t-1},Y_{t-K}^{t-1}\right)}{H\left(Y_t|Y_{t-K}^{t-1}\right)}. \tag{5.13}$$

The main argument for using (5.13) is that it normalizes (5.10) with respect to the different degrees of complexity of each process [161, 162].

5.2.3 Causal measures in neuroscience

Recently, directed information (DI) and transfer entropy (TE) have been applied to analyze interactions between simultaneous neural recordings. We next provide a brief review of a few examples.

- Gourévitch et al [161] estimated the TE between simultaneous spike trains in the cat's auditory cortex, under acoustic stimulation. The TE's asymmetry also allows feedback evaluations. In this paper the authors show the properties of the measure by simulation. When they applied it to the real data their results suggest that: 1) the common history between pairs of neurons would account for roughly 20% of information transfer values if conditioning was not performed; 2) there exist pairs of neurons that are transferring information but are poorly synchronized, or in the opposite direction; 3) the activity may be integrated over a larger interval than that strictly associated with the mean delay between neuronal firings.

- Besserve et al [162] estimate the normalized transfer entropy (NTE) between local field potentials from primary visual cortex of anaesthetized macaques, during spontaneous activity and during binocular presentation of naturalistic color movies. They show that the visual stimulus modulates the causality between gamma bands, suggesting a strong role of the gamma cycle in processing naturalistic stimuli. They also found that the dominant direction of causality was mainly found in the direction from gamma band signals to lower frequency signals, suggesting that hierarchical correlations between lower and higher frequency cortical rhythms are originated by the faster rhythms.

- So et al [163] applied a variation of the DI method to rodent and monkeys recordings. On the rodent data, their goal was to measure the influence of intracortical micro-stimulation applied to the primary motor cortex on spiking activity in the dorsolateral striatum (DLS). They show that DI can correctly infer the degree of influence of intracortical stimulation in M1 on neural spiking activity in DLS. They also show that their method can approximate the conduction delay between M1 and DLS.

  On the monkeys data, they analyzed neural recordings while the subjects were performing a center-out reaching task under manual and neuro-prosthetic modes of operation. They examined the degree of in/out connections, and they found that neurons with high out-degree influenced the activity of many

	f_1 vs f_2	f_1 vs resp	f_2 vs resp
Pearson correlation	-0.000 (0.997)	0.679 (0.000)	0.633 (0.000)
MI	0.738 (0.000)	0.399 (0.000)	0.343 (0.000)
$MI - < MI_{sur} >$	0.388	0.374	0.318

TABLE 5.1: Relationhip between variables, f_1, f_2, and the behavioural response. The values in brackets are the p-values of the corresponding measures.

others, suggesting that these neurons are important drivers of the observed network activity.

- Quinn et al [164] estimated the DI from simultaneous spike trains to estimate neural connectivities. The authors analysed single neurons from M1 from monkeys, while the subjects were performing a series of trials involving contralateral arm movement tasks. They could distinguish three patterns in the connections: 1) neurons which had several incoming and outgoing connections, 2) neurons which had more incoming, 3) neurons which had very few, if any, incoming or outgoing connections.

5.3 Results

5.3.1 Correlation with task parameters

The aim of this section is to reproduce the results obtained by Hernández et al. (Fig. 5 [8], see section 1.2.1 for a summary) using a binary non-parametric model of spike-trains and a non-linear measure of correlation. We calculated the mutual information between the spike trains, the parameters f_1, f_2, and the behavioural response.

In order to directly apply the mutual information, the task variables must be a priori independent. Although f_1 and f_2 are linearly independent, they are dependent in a non-linear sense (See Table 5.1). We therefore use the conditional mutual information to remove any indirect association between spike-train responses and frequency variables (see section 5.5.2.1 for more details).

Despite the differences in the model and the correlation measure, the outcomes of both analyses are similar, as shown in Fig. 5.3 (MI, A; LRM, B). This suggests that the information about the task parameters is mainly encoded in the spike count, and also that the amount of activity is fundamentally linearly related to the task variables.

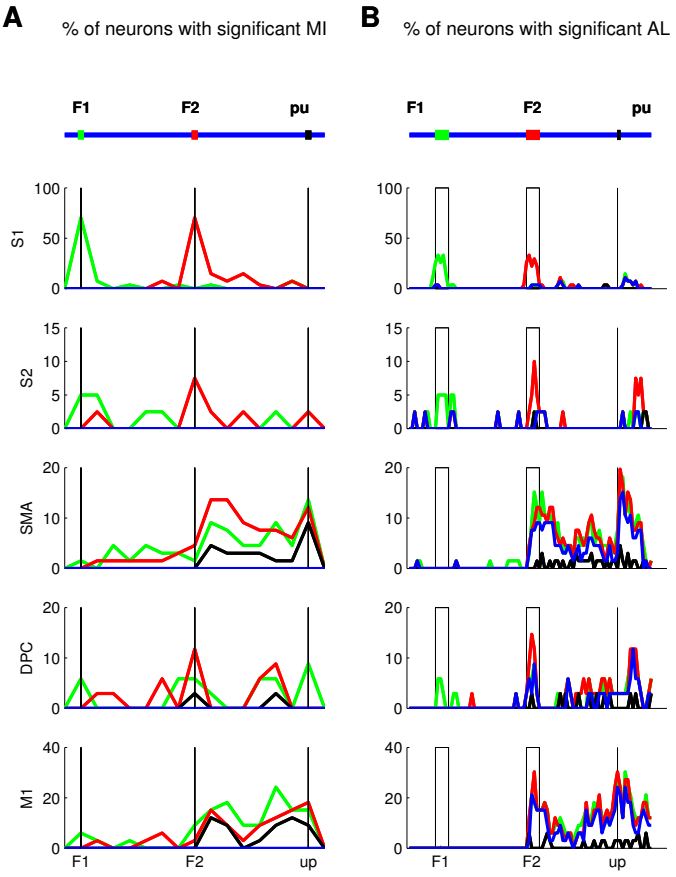

FIGURE 5.3: Information carried by neurons across areas. **A:** Percentage of neurons with significant MI values as a function of time. Green points correspond to $I(X; f_1|f_2)$, red points correspond to $I(X; f_2|f_1)$, and black traces correspond to the $I(X$;response) **B:** Percentage of neurons with linear significant coefficients (see section 1.2.1) as a function of time. Green and red dots correspond to a_1 and a_2 coefficients (see section 1.2.1), respectively. Black dots indicate percentage of neurons with a_1 and a_2 coefficients of opposite sign but of different magnitudes. Blue dots show the percentage of neurons with coefficients a_1 and a_2 of opposite sign but similar magnitude. Primary somatosensory cortex (S1), secondary somatosensory cortex (S2), medial premotor cortex (MPC), dorsal premotor cortex (DPC), and primary motor cortex (M1).

5.3.2 Causal correlations during the discrimination task

To investigate how interactions arise during the discrimination task, we only considered correct trials of similar psychophysical performance (see section 5.6). The central measure of our analysis is the directed information, which is a non-linear measure of causal correlation between any pair of random variables.

To estimate this measure in practice we applied a non-parametric method to pairs of simultaneous binary time series. This method sequentially learns the joint distribution of the sequences when the pair is a realization of a finite-memory stationary and ergodic process [165]. The estimation of the joint probability is based on the context-tree weighting (CTW), a well-known algorithm in the field of data compression [166]. To apply this method to neuronal recordings, we first converted the spike-trains into binary time series using bins of 2 ms [167]. Second, we divided the task into 16 adjacent time intervals of 500 ms from the period prior to the f_1 stimulation to the lift of the sensory cue (*Probe up*), and concatenated the binary series at each time interval for trials recorded under the same conditions (session, stimulation, and pair). The concatenation was done chronologically across trials simultaneously recorded neurons (see section 5.6). Finally, we obtained an estimate of the causal correlation between any pair of simultaneously recorded neurons at a given interval by running the directed information estimator over the corresponding concatenated binary time series (see section 5.5). To avoid the effect of instantaneous coupling artifacts, we only considered inter-neuronal correlations at time lags larger than 10 ms.

For each ordered area pair and time interval, we tested the significance of each estimate against a null hypothesis of complete causal independence. Local tests were used to threshold each interneuronal directional correlation and obtain a binarized functional connectivity matrix. Then, we clustered the underlying connectivity graph into its connected components. Finally, we tested the significance of each cluster using the aggregated sum of causal correlations over its directed paths as the test statistic (see section 5.6). For each one of the 25 inter-area comparison, we computed both the percentage of paths lying in a significant cluster (Green curves in Fig. 5.4) and the average value of the test statistic over significant clusters (Green curves in Fig. 5.4) as a function of the 16 time intervals. In general, clusters of causal correlations were present in all studied areas from f_1 stimulation to the PU period and were negligible in the interval preceding the first stimulation. In inter-area comparisons involving S1 the size of significant clusters also dropped during stimulation intervals while for S2 the drop was only evident during f_2 stimulation. Another drop was found in the number of incoming paths

to DPC during the interval preceding the PU. These patterns were consistent across all sessions and stimulation pairs. An inspection of Fig. 5.4 also reveals that comparisons with opposite direction were of the same magnitude, which may indicate that feedforward and feedback interactions are equally present along the task (Green curves in Fig. 5.4).

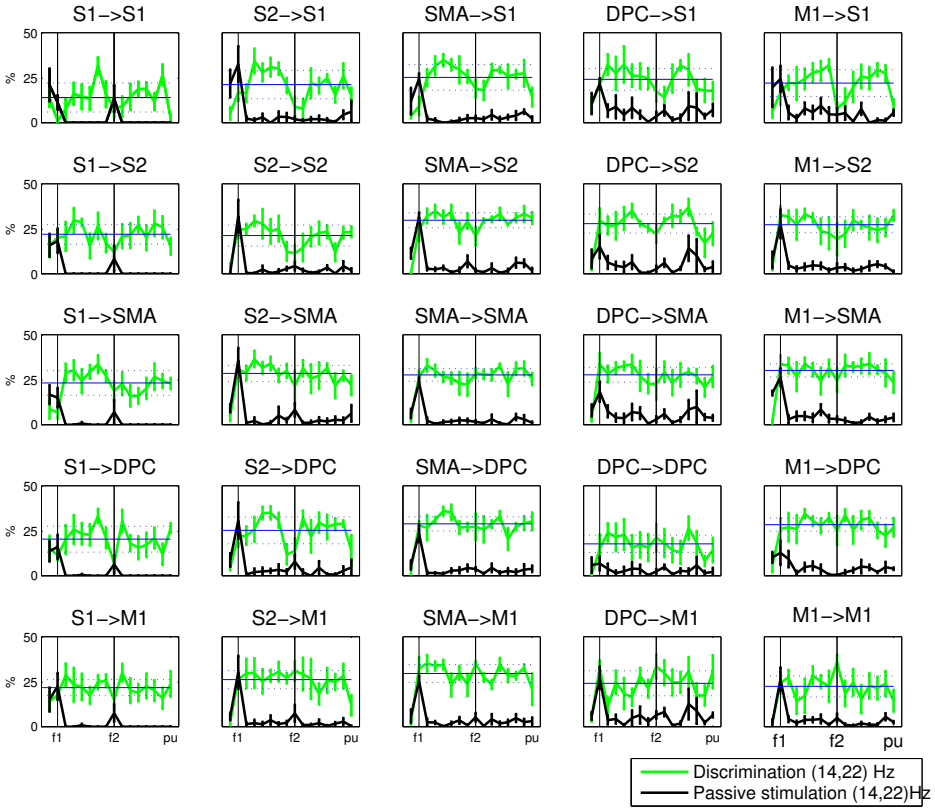

FIGURE 5.4: In green, percentage of significant cluster paths for every directional inter-area comparison during the time course of the discrimination task with $f_1 = 14$ Hz and $f_2 = 22$ Hz. In black, the same percentage for the passive stimulation task with $f_1 = 14$ Hz and $f_2 = 22$ Hz. Percentages are averaged over 5 sessions and error bars denote the standard deviation across sessions.

5.3.3 Causal correlations are inherent to decision-making

To test whether these correlations were directly associated with the discrimination task, we estimated the causal correlation for the same simultaneous neuron pairs in a control task, in which the monkey received both mechanical vibrations identically to the original task but was requested to remain still upon a reward that arrived at variable time (passive stimulation). Compared with the discrimination task, the percentage of paths in significant clusters abruptly decreased after the first

stimulation and remained almost negligible for the rest of the experiment (See black curves in Fig. 5.5). The only exceptions were found in the arriving paths to S1 where the percentages peaked again during the second stimulation and in the arriving paths to DPC in which cluster sizes fluctuated during the second delay. In summary, active and passive task percentages were almost identical during the first two intervals and became clearly differentiated after f_1 stimulation for all inter-area comparisons except for S1 as a destination area, in which the percentages were again similar during f_2 stimulation. These findings were consistent across the five sessions and stimulation pairs except the increase during f_2, which was more evident for the stimulation pair (14:22) Hz than for the pair (30:22).

Since the directed information is made up of two entropy terms (Eq. 5.3), we corss-validated our findings with the estimation of the entropy for each neuron. Specifically, we investigated whether the remarkable differences between both tasks after the first stimulation were due to a difference in the conditional entropy on the stimulation pair.

To test this hypothesis, we estimated the entropy of single neurons in each area under study. In both tasks, significant values of entropy were found in each studied areas, but were not significantly different between each other (Fig. 5.5) in periods after f_1 stimulation.

5.4 Discussion

In this chapter we used a novel methodology to compute causal correlations between spike-trains responses and applied it to simultaneous recordings to investigate the interactions between nearby and distant cells during a decision-making task.

The first question of this chapter was: What is the neuronal variability that can be explained by the task parameters?

To address this question we computed the mutual information between the spike trains and the task parameters (see Fig. 5.3). This result suggests that the information about the task parameters is mainly encoded in the spike count, and also that the amount of activity is fundamentally linearly related to the task variables.

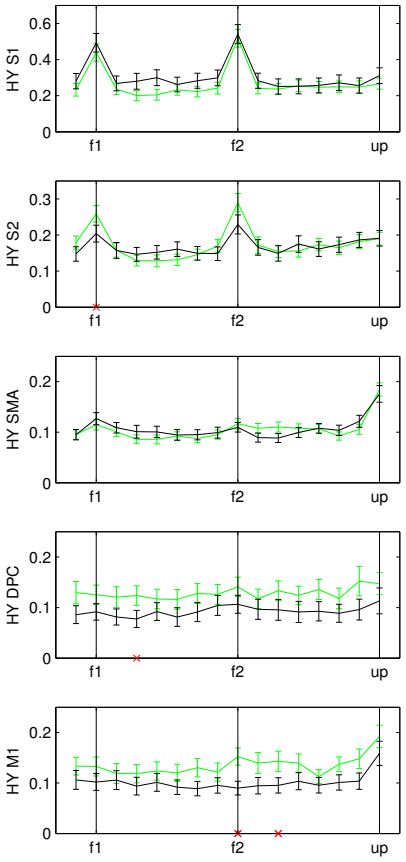

FIGURE 5.5: Mean entropy vs time. Green lines correspond to the discrimination task. Black lines correspond to the passive task. Errors bars SEM. Red crosses on the X-axis mean that the passive and the discrimination task are statistically different.

The second question arising in this chapter was: What is the neuronal information that can be explained by the interactions among neurons? And in particular, is thisadress variability task-dependent?

To solved this question in the second part of the chapter we cnditioned on the external frequencies. Our aim was to capture the variability that can not be directly explained by the task parameters. This approach *per se* was novel because to the classical way to study the information carried by the neurons (under a 2AFC task) was varyring the task parameters.

When we compared that the percentage of significant path during the passive task and discrimation task (Fig. 5.4), surprisingly, the estimated correlations were highly distributed across distant and nearby areas as well as task periods and

were equally present in feedforward and feedback interactions between sensory and motor areas. The observed correlations were task-specific for two reasons. First, they were not present during the period previous to the first stimulation when task-related information is not yet available (Fig. 5.4). Second, these interactions vanish in the passive stimulation task once the first stimulus has been perceived. Interestingly, none of these findings could be devised using an entropic measure over single-cell recordings (Fig. 5.5).

Our results indicate that a great part of the internal neural variability in the discrimination task can be explained by the activity of other neurons, in distant and nearby cortical regions. But a natural question arisest: What is difference between both tasks? The difference may come from some processes present in the discrimination which they are not in the passive task, such as attention or memory.

More generally, our findings strongly suggest that analyzing neuronal interactions on simultaneous recordings opens novel and promising lines of study that may unravel task-related process which have not been yet described.

5.5 Methods: Information Theory quantities

A great number of methods to estimate the information-theoretical quantities (H, MI, DI, TE, etc), involve the estimation of the underlying joint probability distribution of a pair of sequences. Gao et al [168] summarise the properties of a number of entropy estimators. The most commonly used method to estimate the distributions in neuroscience, the *plug-in estimator*, is a frequentist approach used by authors like Panzeri [23, 162]. The *plug-in estimator* approach estimates the underlying joint probability distribution $\hat{P}(X^T, Y^T)$ based on the frequency of string occurrences. All the works presented in the section 5.2.3 used different versions of the *plug-in estimator*: In [161, 162], TE (5.10) has been estimated by using a maximum likelihood estimator of the joint probability, whereas in [163, 164] DI is estimated by assuming an underlying point processes.

The main drawback of the *plug-in estimator* is the undersampling problem. As this method assumes equal distribution of all the possible combinations, it needs a high number of trials to ensure convergence. For example, if we want to estimate $\hat{P}_X$, where X is a time series of alphabet two and length six, we have $2^6 = 64$ possible combinations, and we would need 64 strings for a good estimation of $\hat{P}_X$. In our case, we have a time series of length $T = 250$, so we would need 2^{250} samples. It is worth noting that some bias reduction strategies have been proposed

to increase the convergence of the *plug-in estimator* [84, 162], and we have used them when appropriate (see Chapter 3).

Instead, we used a Bayesian non-parametric approach to estimate the directed information. The estimator is based on methods which have recently been developed [165]. Here we used the non-parametric and sequential estimator based on the context-tree weighting method (CTW) [165], which originated in the field of data compression. This method improves the bias problem. In practice, to estimate these quantities, it is typically assumed that processes (i.e., spike trains in our case) are stationary and ergodic. However, these assumptions are challenging to test.

In the netx sections we dicuss the implementation of the CTW algorithm to estimate the DI.

5.5.1 CTW algorithm

Tree source model

The propierties of the method and the implementation details can be found in [165, 166]. Here we summarise the general idea. Although the first version of the algorithm only applied for a binary alphabets, we here present the generalisation to a M-ary alphabet. Throughout the chapter, we used and modified conveniently when appropriate the MATLAB® code available here.

We consider that sequences of a M-ary alphabet (in our case M=2) are generated by a *tree source* of bounded memory D. The generation of a symbol depends on the past D symbols. More formally stated, the probability of the generated sequence is defined by the model $(\mathcal{S}, \Theta_{\mathcal{S}})$, where $\mathcal{S}$ is the *suffix set* consisting of M-ary strings of length no longer than D, and

$$\Theta_{\mathcal{S}} = (\boldsymbol{\theta}_s; s \in \mathcal{S}) \tag{5.14}$$

is the parameter space where $\boldsymbol{\theta}_s \triangleq (\theta_{0,s}, \theta_{1,s}, \ldots, \theta_{M-2,s})$. The suffix set is required to be *proper* -suffixes in the set are not suffixes of other elements of $\mathcal{S}$- and *complete* -every sequence has a suffix in $\mathcal{S}$-. Then, we can define a mapping $\beta_S(\cdot)$ by which every sequence is mapped onto a unique suffix $s \in \mathcal{S}$. To each suffix, there corresponds a parameter vector $\boldsymbol{\theta}_s$ that determines the next symbol probability in the sequence as

$$\Pr\left\{X_t = i | x_{t-D}^{t-1}, \mathcal{S}, \Theta_S\right\} = \theta_{i,\beta_S(x_{t-D}^{t-1})} \tag{5.15}$$

for $i = 0, \ldots, M-2$, and

$$\Pr\left\{X_t = M-1 | x_{t-D}^{t-1}, \mathcal{S}, \Theta_S\right\} = 1 - \prod_{i=0}^{M-2} \theta_{i,\beta_S(x_{t-D}^{t-1})}. \tag{5.16}$$

The name of tree source comes from the fact that the suffix set can be characterized by a tree. The goal is to find an algorithm that estimates the probability of any sequence generated by a tree source without knowing the underlying model $(\mathcal{S}, \Theta_S)$, i.e, without knowing neither the suffix set $\mathcal{S}$ nor the parameter space Θ.

Example

Let $M = 2$, $D = 2$ and consider the suffix set $\mathcal{S} = \{00, 10, 1\}$. Then, the probability of the sequence $x_1^7 = 0110100$ given the past symbols 10 can be evaluated as $\Pr\{x_1^7 | \mathcal{S}, \theta_{00}, \theta_{10}, \theta_1\}$:

$$\begin{aligned}\Pr(0110100|10) =& P(0|10) \cdot P(1|00) \cdot P(1|01) \cdot P(0|11) \cdot P(1|10) \cdot P(0|01) \cdot P(0|10) \\ =& (1-\theta_{10}) \cdot \theta_{00} \cdot \theta_1 \cdot (1-\theta_1) \cdot \theta_{10} \cdot (1-\theta_1) \cdot (1-\theta_{10})\end{aligned}$$

where we used the mapping $\beta_{\mathcal{S}}(10) = 10$, $\beta_{\mathcal{S}}(00) = 00$, $\beta_{\mathcal{S}}(01) = 1$ (the sufix 01 is not in the set of suffixes $\mathcal{S}$, and we thus map it to the suffix one $\beta_{\mathcal{S}}(11) = 1$.

Bayesian approach

The context-tree weighting is a method of approximating the true probability of a T-length sequence x_1^T (under the true model $(\mathcal{S}^\star, \boldsymbol{\theta}^\star)$) with the mixture probability

$$\hat{P}(x_1^T) = \sum_{(\mathcal{S},\Theta_{\mathcal{S}})} w(\mathcal{S}, \Theta_{\mathcal{S}}) P_{\mathcal{S},\Theta_{\mathcal{S}}}(x_1^T), \tag{5.17}$$

where $w(\cdot)$ is a weighting function over the models and $P_{\mathcal{S},\Theta_S}(x_1^T)$ is the probability of generating the sequence x_1^T under the tree model $(\mathcal{S}, \Theta_S)$.

To approximate (5.17), we first make use of the concept of context tree. The context tree is a set of nodes where each node is an M-ary string s with length $l(s)$, and where $l(s)$ is upper-bounded by a given memory D. Each node s splits into M (child) nodes $0s, 1s, \ldots, (M-1)s$. To each node there corresponds a

vector of counts $\boldsymbol{a}_s = (a_{0,s}, a_{1,s}, \ldots, a_{M-1,s})$. For a parent node s and its children $0s, 1s, \ldots, (M-1)s$, the counts must satisfy $a_{i,s} = \sum_{j=0}^{M-1} a_{i,js}$ for every symbol $i = 0, \ldots, M-1$. Then, for every node with string s we estimate the probability that a sequence is generated with the counts $\boldsymbol{a}_s$.

The idea of the algorithm is to update the counts in each node by keeping track of how many times a symbol $i = 0, \ldots, M-1$, is preceded by the corresponding string.

In general, the probability that a memoryless source with parameter vector $\boldsymbol{\theta}$ generates a given sequence follows a multinomial distribution. By weighting this probability over all θ_i using a prior Dirichlet distribution we obtain the well-known Krichevsky-Trofimov (KT) estimate. A useful property of this estimator is that it can be computed sequentially as $P_e^s(0, 0, \ldots, 0) = 1$ and

$$P_e^s(a_{0,s}, a_{1,s}, \ldots, a_{i-1,s}, a_{i,s}+1, a_{i-1,s}, \ldots, a_{M-1,s}) = \frac{a_{i,s} + \frac{1}{2}}{a_{0,s} + a_{1,s} + \ldots + a_{M-1,s} + \frac{M}{2}}. \tag{5.18}$$

Finally, we assign a probability to each node, which is the weighted combination of the estimated probability and the weighted probability of its children:

$$P_w^s = \begin{cases} P_w^s = \alpha P_e^s(\boldsymbol{a}_s) + (1-\alpha) \prod_{i=1}^{M} P_w^{is}, & 0 \leq l(s) < D \\ P_e^s(\boldsymbol{a}_s), & l(s) = D, \end{cases} \tag{5.19}$$

where α is typically chosen to be $\frac{1}{2}$.

Original Algorithm for $M-$ary alphabet

For every $t = 1, \ldots, T$, we use the context x_{t-D}^{t-1} and the value of x_t. Then, we track nodes from the leaf to the root node along the path determined by x_{t-D}^{t-1}.

- **Leafs**: Identify the leaf s that corresponds to x_{t-D}^{t-1} in the context tree. Then

  1. *Counts update*
     Based on the value of x_t, update $\boldsymbol{a}_s$.
  2. *Estimated probability*
     Compute $P_e^s(\boldsymbol{a}_s)$ using the Krichevsky-Trofimov estimator, which is defined recursively as $P_e(0, 0 \ldots 0) = 1$ and for $a_{i,s} \geq 0$, $i = (1 \ldots M-1)$,

$$P_e(a_{0,s}, a_{1,s}, \ldots, a_{i-1,s}, a_{i,s}+1, a_{i-1,s}, \ldots, a_{M-1,s}) = \frac{a_{i,s} + \frac{1}{2}}{a_{0,s} + a_{1,s} + \ldots + a_{M-1,s} + \frac{M}{2}}.$$

3. *Weighted probability*
   For the leaf nodes, $P_w^s = P_e^s(\boldsymbol{a}_s)$.

- **Internal nodes**: Using the path determined by the context x_{t-D}^{t-1}, REPEAT

  1. *Parent search*
     Identify the parent s of the previously tracked node.
  2. *Counts update*
     Based on the value of x_t, update $\boldsymbol{a}_s$.
  3. *Estimated probability*
     Compute $P_e^s(\boldsymbol{a}_s)$ using $\boldsymbol{a}_s$ and the Krichevsky-Trofimov estimator.
  4. *Weighted probability*
     Compute P_w^s as

$$P_w^s = \alpha P_e^s(a_{0,s}, a_{1,s}, \ldots, a_{M-1,s}) + (1-\alpha)\prod_{i=1}^{M} P_w^{is},$$

     where α is typically chosen to be $\frac{1}{2}$.
     UNTIL the root node is tracked.

- **Probability assignment**: Let λ denote the root node of the context tree. Then, $\hat{P}(x_1^t) \equiv P_w^\lambda(x^n)$ is the universal probability assignment in the CTW algorithm. As a result, we also obtain the conditional probability $\hat{P}(x_1^t|x_1^{t-1})$ as:

$$\hat{P}(x_1^t|x_1^{t-1}) = \frac{P_w^\lambda(x_1^t)}{P_w^\lambda(x_1^{t-1})}.$$

Example

Consider the sequence $x_t = 1011011$ with past 101, where $M = 2$ and $D = 3$. Suppose that we are at time instance $t = 7$ and thus, the context is x_{t-D}^{t-1}. Then, we follow the algorithm steps as shown in Figs. 5.6 and 5.7.

5.5.2 Estimation of the Directed Information

We resort to an estimator of the DI based on the CTW algorithm [165]. This estimator therefore applies to finite-memory Markov chains X^T and Y^T defined

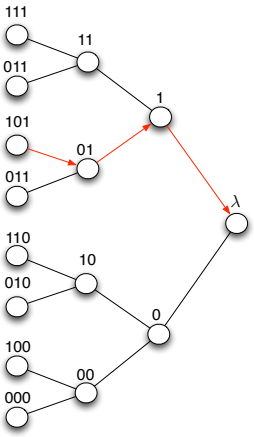

FIGURE 5.6: Step 1: Context tree, in which the path determined by $t = 1$, $x_{-3} = 101$ is marked. Step 2: Counts update of the leaf index for $t = 1, x_1 = 1$ with past $x_{-3} = 101$: update the leaf index corresponding to node 101 from (0,0) to (0,1); the leaf index corresponding to 01 from (0,0) to (0,1); the leaf index corresponding to 1 from (0,0) to (1,0); the leaf index of λ from (0,0) to (0,1). We update until $t = 7$.

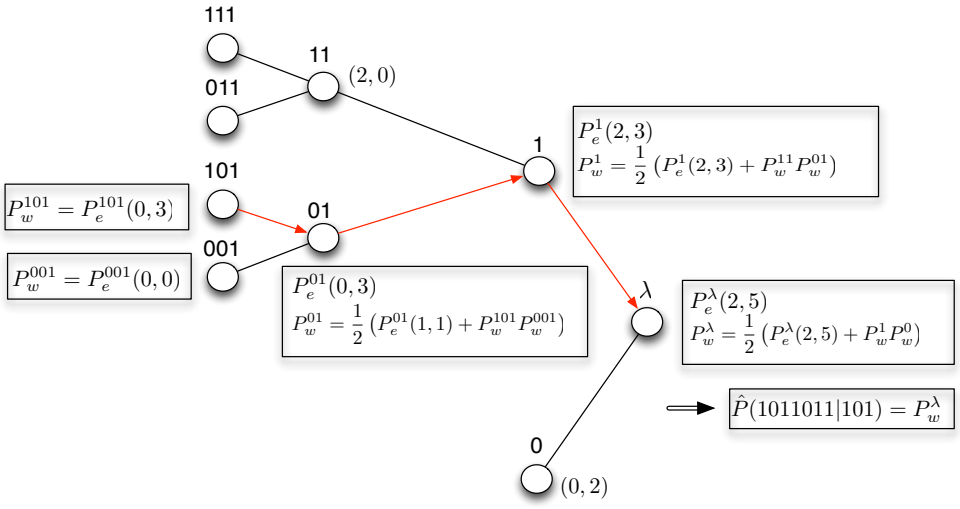

FIGURE 5.7: Steps 3 and 4: Estimation of $P_e^s(\boldsymbol{a}_s)$ and P_w^s

over a finite alphabet that are stationary and ergodic. The estimator is computed using the following alternative definition of the DI:

$$\begin{aligned} I(X^T \to Y^T) &\triangleq \frac{1}{T}\sum_{t=1}^{T} D\left(P(Y_t|X^tY^{t-1})\|P(Y_t|Y^{t-1})\right) \\ &= \frac{1}{T}\sum_{t=1}^{T}\sum_{y_t} P(y_t|X^tY^{t-1}) \log \frac{P(y_t|X^tY^{t-1})}{P(y_t|Y^{t-1})} \end{aligned} \tag{5.20}$$

In summary, the steps to compute the DI are the following:

1. Estimate $P(Y)$ by using the CTW algorithm with memory D:

$$P(Y_t|Y^{t-1}). \tag{5.21}$$

2. Estimate $P(X^T, Y^T)$ by using the CTW algorithm with memory D:

$$P(X_t, Y_t|X^{t-1}Y^{t-1}). \tag{5.22}$$

3. Obtain the marginal probability from (5.22).

$$P(X_t|X^{t-1}Y^{t-1}) = \sum_{y_t} P(X_t, y_t|X^{t-1}Y^{t-1}) \tag{5.23}$$

4. Apply Bayes theorem using (5.22) and (5.24) to compute

$$P(Y_t|X^tY^{t-1}) = \frac{P(X_t, Y_t|X^{t-1}Y^{t-1})}{P(X_t|X^{t-1}Y^{t-1})} \tag{5.24}$$

5. Apply (5.20) using (5.21) and (5.24) to obtain the final result.

5.5.2.1 Mutual Information application

To compute the mutual information between the neuronal spike trains and task parameters we compute the conditional mutual information, we proceed as follows:

$$I(X^T; f_1|f_2) = H(X^T|f_2) - H(X^T|f_1, f_2)$$

We compared results obtained with the MI estimation and the multi regression lineal model (LRM) explained in section 1.2.1 to the same set of neurons. There

are two main differences in the application of the two methods: for the LRM case we used a spike count code within a sliding window of 100 ms duration moving in steps of 20 ms for performing the analysis, while for the MI analysis we used non-overlapping binary time series of 500 ms.

5.6 Methods: Application to spike-train data

5.6.1 Data pre-processing

Preliminary selection

We selected five out of twenty recorded populations. In the following we summarise the neurons per area and population:

Area/population	S1	S2	SMA	DPC	M1	
Pop 1	8/5	8/8	14/13	4/4	8/8	
Pop 2	6/6	7/7	12/12	9/9	9/9	
Pop 3	6/ 5	12/12	13/13	9/9	6/6	
Pop 4	5 /5	4/4	12/11	8/8	5/5	
Pop 5	2 /1	9/9	15/15	4/3	5/5	
TOTAL	27 /22	40/40	66/64	34/33	33/33	200/192

TABLE 5.2: Number of neurons per area and population, for *discrimination/passive* taks.

For each population, we analysed the following frequency pairs:

$$(f_1 : f_2) = \{(14 : 22), (22 : 14), (30 : 22), (22 : 30)\}.$$

We chose the pairs according to two criteria. The first criterion was to maintain the distance between the frequency pairs constant ($|f_1 - f_2| = 8$) to neglect differences among the pairs due to the task difficulty. The second criteria was to keep f_1 or f_2 fixed in at least two opposite responses, namely, for fixed f_1 ((22:14), (22:30)), and for fixed f_2 ((14:22), (30:22)). We only used correct trials in the discrimination task.

Binarization of spike-train trials

To quantify the information carried by spikes, we discretized each trial into non-overlapping bins of 2 ms [167], assigning the value 1 to each bin with at least one spike and the value 0 otherwise (see Fig. 5.8).

We test the goodness of our bin choice by counting the number of times that more than one spike occur in one bin. Although the number of ISI shorter than 2 ms is high (18 spikes each 8.5 seconds), the number of spikes losses is much lower (14% of 18). The results are summarised in Tab. 5.6.1.

Area	S1	S2	SMA	DPC	M1
$\#ISI < 2ms$ per trial	18.6	6.3	1.0	0.4	1
$\frac{\#negleted\,spikes}{\#ISI<2ms}$	14.5 %	10.8%	10.6%	5.5%	8.3 %

TABLE 5.3: Mean number of ISIs less than 2 ms per area in a trial (8.5 s). Ratio of neglected spikes to number of ISIs less than 2 ms per area.

The approximate duration of each trial was 8500 ms, hence we obtain time series of 4250 bins.

By doing this binarization we preserve in the model the most evident pieces of information,i.e., the number and distance between spikes.

FIGURE 5.8: An example of a binary dicretization of a continuous signal.

Trial partition

We subdivided each binarized trial, of 4250 bins in 16 consecutive and non-overlapping intervals of 250 bins (see Fig. 5.9). Each interval is of 500 ms, which coincides with the stimulation period. In particular, the first interval corresponds to the 500 ms previous to the first stimulation, the second interval corresponds to

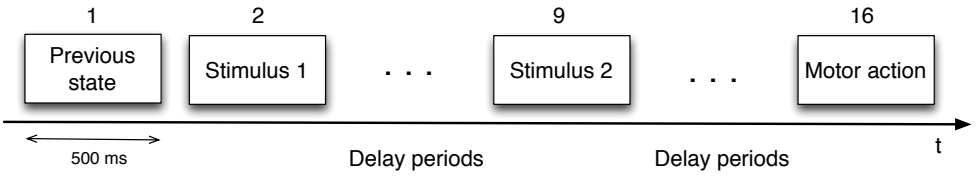

FIGURE 5.9: A simplified representation of the division of one trial into 16 intervals of 500ms, where the second interval corresponds to the first stimulation, the ninth interval corresponds to the second stimulation and the sixteenth interval corresponds to the probe up.

the presentation of the f_1, the ninth interval corresponds to the the presentation of the f_2, and the sixteeenth interval corresponds to the KU.

At the end of this process we obtain one binary time series X^T for each interval, neuron, and pair of stimulation frequencies. The length of this time series is 250 xm bins., where m equals the number of trials under a given stimulation pa At the end of this process we obtain, for each neuron, pair of frequencies, and task period, one binary time series X^T:

$$X^T_{(f_1:f_2),interval} = \{X_1, \ldots, X_T\}$$

Trial concatenation

To obtain a good estimation of the joint distribution we need the more bins the better. For each neuron, we concatenate all trials recorded under the same stimulation pair $(f_1 : f_2)$ and at each interval i, $(i = 1, \cdots, 16)$. This results in sixteen time series per neuron and stimulation pair, sixteen time series, one for each period under study (see Fig. 5.10). The concatenation preserves chronology across neurons, as the trials at interval i have been simultaneously recorded.

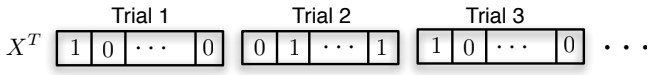

FIGURE 5.10: A scheme of the trial concatenation procedure for a given neuron, interval and frequency pair.

Delays

A central question to our approach is the time scale at which interactions occur. Results on inter-area delays during decision making are scarce in the literature.

Instead, the concept of latency, the average time that an area takes to respond to a certain stimulus, has been used to approximate the computation of delays during the whole discrimination task [133]. Using these results, we consider the inter-area and intra-area delays within the interval of $[10, 140]$ ms.

Memory

The CTW algorithm has a free parameter that we must choose, the maximum depth used, D, which can be interpreted as the memory of the Markov chain. The computational cost of the algorithm grows exponentially with D, and hence, D becomes a critical parameter since we have to deal with a large number of comparisons. To estimate the memory of the neurons, we calculate the entropy of all neurons in one population for values of D ranging from 0 to 9 in all considered task intervals. As showed in Fig. 5.11, the memory of the analyzed spike trains may vary across intervals and areas. By inspecting the general pattern of the average entropy in each area under study (see Fig. 5.11), and trading it off with the computational cost of the estimation, we decided to choose a memory of 2 bins (4 ms).

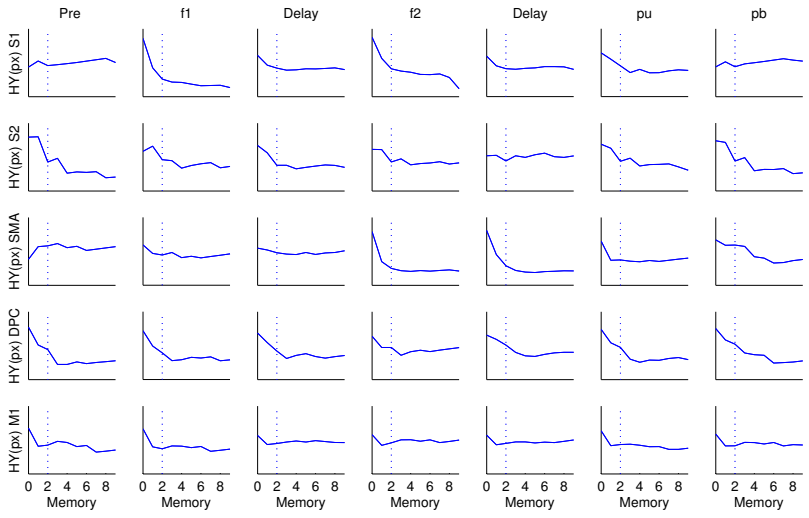

FIGURE 5.11: Averaged entropy of neurons estimated with the CTW algorithm as a function of the memory parameter.

5.6.2 Convergence of the CTW algorithm

In order to observe the convergence of the CTW algorithm on real data, we show the estimated probability distribution $\hat{P}_X$ of a neuron in the primary somatosensory cortex (S1) (see Fig. 1.4). In this example we studied the neuron during the

first vibrotactile stimulation (see Fig. 1.1). The response of S1 neurons under the vibrotactile stimulation is shown in Fig. 5.12 (*Top*). It shows weak rate modulation, which reflects the periodicity of the stimulus (f_1). In Fig. 5.12 (*Middle*) we can see the value of $\hat{P}_X$ for each 2 ms bin. We clearly observe that the algorithm is able to differentiate four distributions conditioned on the pasts 00, 01,10,11, when the number of bins is approximately larger than 300. This is represented in Fig. 5.12 (*middle* plot) by 4 paths that originate from a common root.

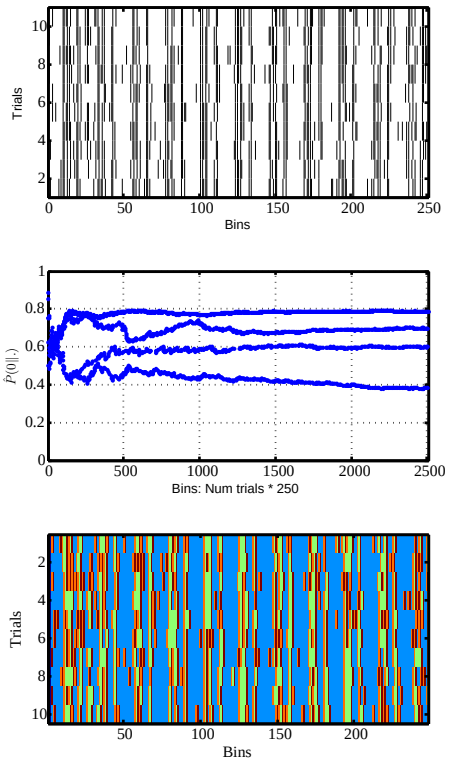

FIGURE 5.12: Application of CTW to real neurons from S1 during a vibrotactile stimulation. *Top*) Binarization of the neuron in bins of 2 ms. Each row corresponds to a trial with $f_1 = 22$. Each black point is a spike. *Middle*) Conditional distribution obtained with the CTW algorithm bin by bin ($D = 2$). All the above trials were concatenated in one trial in order to compute $P(0|-,-)$. The number of bins are represented on the X-axis while the estimated value $\hat{P}(0||-,-)$ is represented on the Y-axis. We can see 4 different paths which correspond to the 4 possible pasts of length-2 (00,01,10,11) that precede any symbol. *Bottom*) All the bins shown in the top figure have been plotted depending on the value of the two preceding bins: 00 corresponds to blue, 01 to orange, 10 to red, and 11 to green.

5.6.3 Statistical procedures

5.6.3.1 Testing the significance: Surrogates procedure

Our goal is to test causal correlations between neurons based on the entropy, causal conditional entropy and DI measures. In practice, to assess statistical significance of these measures, we test whether the measure computed for two simultaneous time series is significantly different from the measure computed when the simultaneity is violated. To do so, we used a Monte-Carlo randomization test. CITE! To construct H_0 we used the shuffling procedure. We estimated the measure (entropy, causal conditional entropy, directed information) from X^T to Y^{*T}. Where Y^{*T} was constructed as the randomly concatenation of the T correct trials. For example if $T = 3$, one possible randomization will be $Y^3 = \{Y_2, Y_3, Y_1\}$ (see Fig. 5.13 for a schematic representation). Concatenations violated simultaneity across all trials. This boostrapping procedure destroys all simultaneous dependencies but preserves the statistics of individual concatenated trials, which may be still affected by common input factors. We obtain an approximation of H_0 by generating 30 randomization for each comparison.

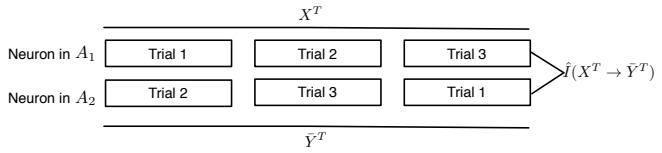

FIGURE 5.13: An example of the shuffling procedure between two time series X^T, Y^{*T}

5.6.3.2 Clustering procedure to correct for multiple comparisons

We applied a multiple test procedure based on clustering locally significant results [169]. The calculation of this test involves:

1. For each $(A_1, A_2, interval$, and $f_1 : f_2)$ and delays from 10 to 140 ms in steps of 10 ms, we run a different Monte-Carlo randomization test [170], over the conditional entropy $H(Y|X)$ at 97%. We obtain a set of triplets (neuron 1, neuron 2, delay) for each interval and stimulation pair.

2. For each neuron pair and interval, we average the directed information over all significant delays and do another Monte-Carlo randomization test over this statistic at level of 97%. We now obtain a set of locally neuron pairs for

each interval and stimulation pair. Based on the previus test, construct the bipartited graph from A_1 neurons to A_2 neurons.

3. Cluster the connected components of the bipartite graph.

4. For each connected component, calculate cluster-level statistics by sum the DI over all the edges within the same component belong to the component, and test at a significance level of 97%.

5. Finally, correct the values with the Šidák multiple test correction.

Chapter 6

Conclusions

We divide the last chapter of the thesis in three parts. In the first part we summarise the main results presented in the previous chapters. In the second part, we compare the methods used throughout the thesis discussing their advantages and drawbacks.

6.1 Summary of the results

In this thesis we have investigated the neuronal processes that occur during a decision-making task based on a perceptual classification judgment. For this purpose we have analysed three different experimental paradigms (sensory, visual, and auditory) in two different animals species (monkey and rat), with the goal of studying the information carried by the neurons.

In particular, we focused on how the information is preserved in the underlying neuronal substrate in time. Furthermore we considered how the decision, the stimuli, and the confidence are encoded in memory and, when the experimental paradigm allowed it, how attention modulates these features. Finally, we went one step further, and we performed an analysis of the interactions between areas during a decision-making task.

In the first part of the thesis, we worked on the mechanisms that enable decisions to be postponed for a period after the evidence has been provided. To tackle this issue we used the mutual information between the rate and the behavioural response.

We showed that information about the forthcoming action becomes available from the activity of neurons in the medial premotor cortex in a sequential decision-making task after the second stimulus is applied. It provided the information for a decision about whether the first or second stimulus is higher in vibrotactile frequency. The information then decays in a 3s delay period in which the neuronal activity declines before the behavioural response can be made. The information then increases again when the behavioral response is required.

We modeled this neuronal activity using an attractor decision-making network in which information reflecting the decision is maintained at a low level during the delay period, and is then selectively restored by a nonspecific input when the response is required. One mechanism for the short-term memory is synaptic facilitation, which can implements a procedure for postponed decisions that can be corrected even when there is little neuronal firing during the delay period before the postponed decision. Another mechanism is graded firing rates by different neurons in the delay period, with restoration by the nonspecific input of the low-rate activity from the higher-rate neurons still firing in the delay period. These mechanisms can account for the decision-making and for the memory of the decision before a response can be made, which are evident in the activity of neurons in the medial premotor cortex.

In chapter 3, we focused on variability and information content in A1 during an interval-discrimination task (of temporal information). First, we studied the neuronal firing rates patterns in this area, and we compared them across brain states. We found neurons that showed prominent responses during the intervals between stimuli, with firing rates that either increased or decreased toward the second stimulus. These patterns of spontaneous activity were often modulated in the attentive versus passive trials. Modulation of the spontaneous firing rate during the task was observed not only between auditory stimuli, but also in the interval preceding the stimulus. These slow modulatory components could be locally generated or the result of a top-down influence originated in higher associative association areas. Such a neuronal discharge may be related to the computation of the interval time and contribute to the perception of the auditory stimulus.

In the second part of chapter 3, we analysed the variability and information content. We found that spike firing variability measured with the Fano factor was

markedly reduced, not only during stimulation, but also in between stimuli in engaged trials. We next explored if this decrease in variability was associated with an increased information encoding. Our theoretical analysis revealed an increased information content in auditory responses during engagement when compared with idle states, in particular in the responses to task-relevant stimuli. Altogether, we demonstrated that task-engagement significantly modulates coding properties of auditory cortical neurons during an interval-discrimination task.

In the third part of the thesis, we studied the neural correlates of decision confidence. In this chapter we studied neurons from Ventral Premotor cortex while the subjctes were performing a decision-making task. In this area we showed the existence of neurons that encode the confidence in the decision. Moreover we found that both continuous and discrete encoding of decision confidence are present in the primate brain. In particular we showed that different neurons encode confidence through three different mechanisms: 1) switch time coding, 2) rate coding, and 3) binary coding. This rich representation of decision confidence indicates the basis of different functional aspects of confidence in the making of a decision.

In the fourth and last part of this thesis, we analysed the interaction between different cortical areas to characterise the causal correlation paths that may arise in the discrimination task. In the analysis, we measured the causal dependence between sets of simultaneously recorded neurons during a discrimination task and a control task of passive stimulation. To quantify these correlations, we made use of a non-symmetrical information-theoretic measure called directed information (DI), which quantifies causal correlations between a pair of sequences without making any linearity assumption. Our results support the role of S1 as a prominently sensory area. We More generally, we provide evidences that the discrimination task conclude that the discrimination task and more specifically, the stages of working memory, comparison and decision involve the existence of feed-forward and feedback paths across sensory and motor areas.

Chap.	Species	Sensory modality	Cortical Area/s	Study	Neuronal Code	Method & feature
2	Monkey	Somatosensory	MPC	Neuron vs response	Spike count	MI: Escalar, Non-linear
3	Rat	Auditory	A1	Neuron vs response	Spike count	MI: Escalar, Non-linear
				Neuron variability	Spike count	$Ff : \frac{\sigma^2}{\mu}$
4	Monkey	Visual	VPC	Neuron vs $\lvert S1 - S2 \rvert$	Spike count	LRM: Captures linear relationships
				States of firing rates vs confidence	Spike train	HMM: Captures transitions in the activity of each trial
5	Monkey	Somatosensory	S1 S2 SMA DCP M1	Neuron vs $(S1, S2,$ and response)	Spike train	MI: Vectorial, Non-linear, Symmetric
				Neuron 1 vs neuron 2	Spike train	DI: Vectorial, Non-linear, Asymmetric

TABLE 6.1: Summary of the experimental data and the methods used along the thesis.

6.2 Summary of the methods

In this section we want to compare the different methodologies used along the thesis. The main objective of this thesis was to establish correlations between the activity of single neurons and other variables, i.e., explain the variability of a neuron based on its relationship to some parameter or neuron. In Table 6.1 we summarise the principal features of the analysis performed in each chapter: neural code, method used, species, sensory modality, cortical area analysed, objective of the study neuronal activity or parameter of interest, neuronal code, method used, and its principal features.

In chapters 2 and 3, we used the mutual information measure to quantify the correlations between the neuronal activity modeled by the spike count and the behavioural response of the animals. Fig. 6.1 shows a scheme of the process.

In chapter 3 we also analysed the interval variability across trials. For that purpose we measured the Fano factor $\left(Ff = \frac{\sigma^2}{\mu}\right)$. In our case, the reduction of Ff was associated with an increment to the encoding/processing of the discrimination capability evidenced in the behavioural response.

In chapter 4, we used a linear regression model to quantify the correlation between the neuronal activity modeled by the spike count and the difficultly of the task, modeled by the absolute difference between both stimuli ($|S1 - S2|$). Fig. 6.2 shows a scheme of the process. We used a linear model instead of a non-linear measure because the hallmark of confidende is the increasing monotonic relationship between difficulty and rate. Confidence, as a function of evidence for the decision, decreases in erroneous choices and increases in correct choices [1, 119, 122], producing an x-shaped pattern. In our paradigm we had four different levels of difficulty and the linear relationship approach is a good aproximation of the increasing monotonic relationships.

In chapter 4, we also wanted to approximate the moment in time when the activity of the single neurons switched dynamics from one state of activity to another, and to relate that moment with the level of confidence of a trial. Note that for that purpose we needed a measure that gave a different output for each trial, instead of a general measure of the whole neuron. To do that we applied a Hidden Markov Model to all the trials of a given neuron. To apply the method we modeled the neurons' activity as a discrete time series of zeros and ones. In this method, instead

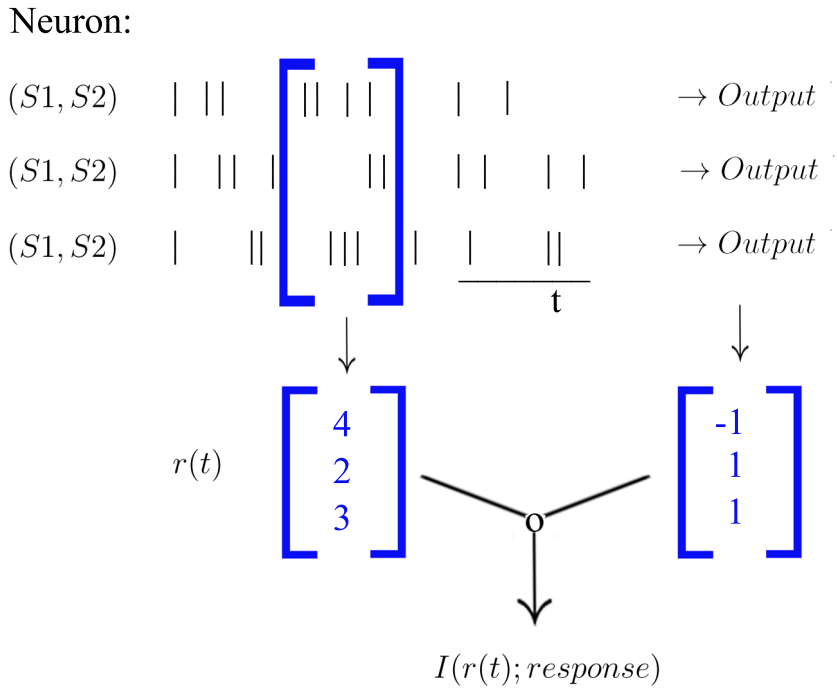

FIGURE 6.1: Scheme of the MI application to the data, where the output was binarised into the two types of responses $(-1/1)$.

of sliding the time window along the analysed period, we applied it to the whole period, so that it could capture the changes along time for each trial. Fig. 6.3 shows a scheme of that process.

In chapter 5, our aim was to study the relationship between two different neurons and, in particular, to measure th e causal dependence between them. We used the and information-theoretic measure called directed information (DI, see Fig. 6.5. The main difference between the DI and the MI is that the DI is a non-symmetric measure. Similarly to the MI, we first estimated the joint probability distribution $P(X^T, Y^T)$ to compute the DI.

Throughout the thesis we used serval times the MI as a measure of non-linear correlation. There are two main difference between the approach taken in chapter 5 as compared to previous chapters (Chap. 2, 3). The first difference is that the variables in chapter 5 are time series while in previous chapters are scalars (see both schemes in Figs. 6.4, 6.1). The second difference lies in the estimation method. In chapters 2 and 3 we used a frequentist approach to estimate information-theoretic quantities and used several methods to remove the intrinsic estimation bias. In

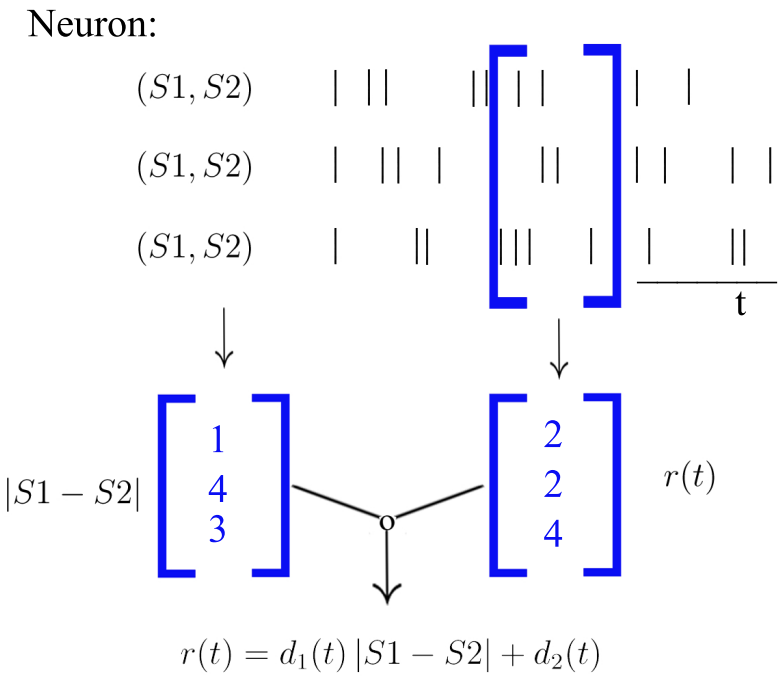

FIGURE 6.2: Scheme of the linear regression model application to the data. In it we related the task difficulty ($|S1-S2| \in [1 \ldots 4]$) with the spike count ($r(t)$). We moved the time window (t) in several steps during all the period of study.

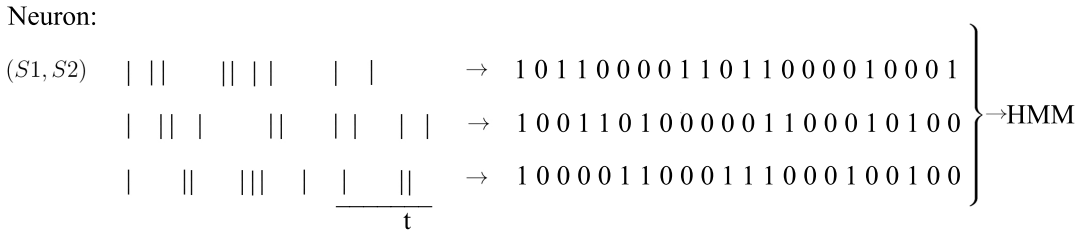

FIGURE 6.3: Scheme of the HMM application to the data.

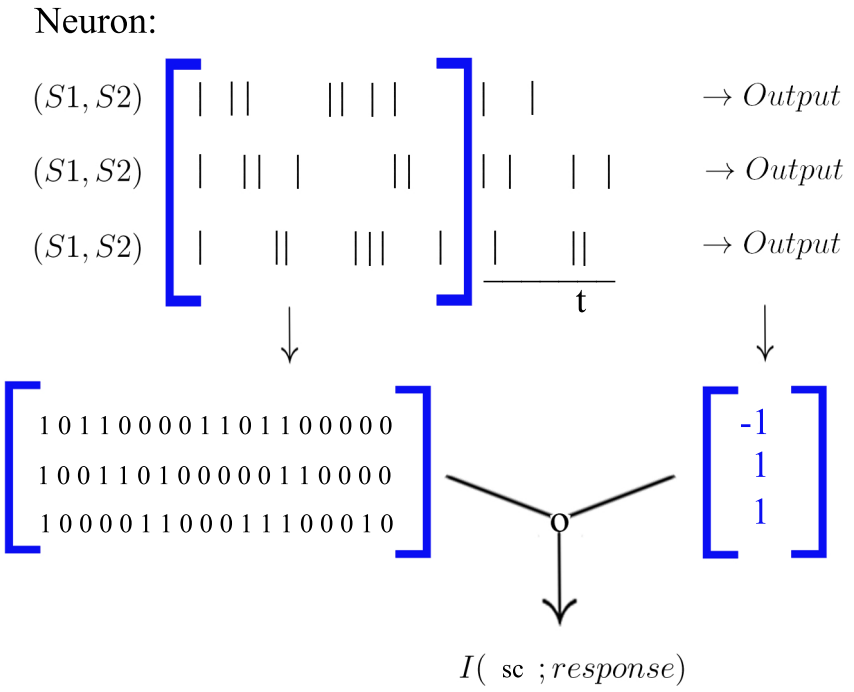

FIGURE 6.4: Scheme of the MI application between a neuron (through its spike train) and the behavioural response: $I(sc\,;n)$.

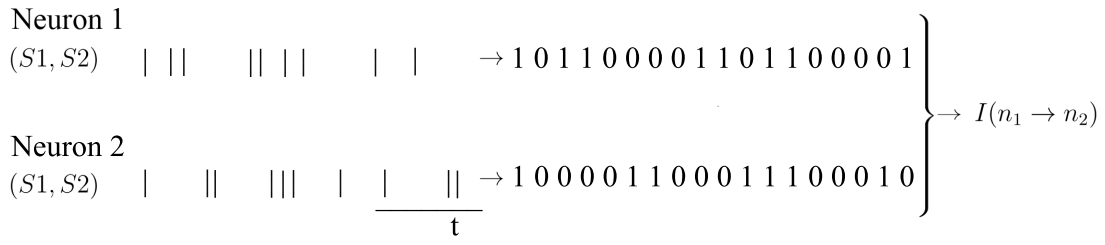

FIGURE 6.5: Scheme of the DI application to two neurons: $I(n_1 \to n_2)$.

contrast, in chapter 5 we followed a Bayesian approach based on the context-tree algorithm.

6.2.1 Feature comparisons

Linear vs Non-linear: The major strength of the information-theoretic measures is that they can capture non-linear relationships between variables, and they have

flexible definitions that can be applied to several neuronal codes (spike train, spike count).

However, the main drawback of these measures is that their computation involves the underlying probability distributions, $\hat{P}_{X,Y}$, of the data, which in most of the cases is difficult to compute. To solve this problem there are two main approaches: the frequentist and the Bayesian. The frequentist approach needs a large amount of data to converge, but some bias reduction strategies have been proposed to increase the convergence [84, 162].

The Bayesian approach minimizes the samplig bias problem (e.g., the CTW algorithm [165, 168]) but its main problem is that it is computationally intense both in space as in time.

The major strengths of the linear regression model (LRM) are: 1) It usually gives a clear interpretation of the results [8], 2) Their application to data is normally simple and is computationally cheaper than the computation of underlying probability measures. However, the LRM can only capture linear relationships.

Spike count vs spike train: The spike train model preserves the most evident pieces of information, the number and distance between spikes (inter-spikes interval), whereas the spike count only accounts for the amount of activity. Another important difference is that the spike train can capture temporal aspects of the time series, which can be potentially correlated with task parameters.

Despite these differences between both approaches, the results from Chapter 5 support a widespread understanding that information about the task parameters is mainly encoded in the spike count, and that the amount of activity is fundamentally linearly related to the task variables. However, to the best of our knowledge, attempts to describe neuronal interactions with spike counts and linear methods (personal communication, Romo's lab) have not been able to characterize the activity of a cognitive task using a control task. We therefore believe that information transfer across cortical areas strongly relies on temporal patterns of spike trains. In conclusion, the choice of method and measure may depend on the research problem at hand.

Bibliography

[1] D. Vickers. *Decision Processes in Visual Perception*. Academic Press, New York, 1979.

[2] AL. Hodking and AF. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol*, 117(4):500–44, 1952.

[3] TN. Wiesel and DH. Hubel. Single-cell responses in striate cortex of kittens deprived of vision in one eye. *J Neurophysiol*, 26:1003–1017, 1963.

[4] DH. Hubel and TN. Wiesel. Receptive fields of cells in striate cortex of very young, visually inexperienced kittens. *J Neurophysiol*, 26:994 –1002, 1963.

[5] G. Fechner. *Elements of psychophysics*. Rinehart and Winston, New York, 1966.

[6] JD. Roitman and MN. Shadlen. Responce of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of Neuroscience*, 22:9475–9489, 2002.

[7] JL. Pardo-Vázquez, V. Leboran, and C. Acuña. Neural correlates of decisions and their outcomes in the ventral premotor cortex. *Journal of Neuroscience*, 28:12396–12408, 2008.

[8] A. Hernández, V. Nácher, R. Luna, A. Zainos, L. Lemus, M. Alvarez, Y. Vázquez, L. Camarillo, and R. Romo. Decoding a perceptual decision process across cortex. *Neuron*, 66(2):300–314, 2010.

[9] A. Kepecs and ZF. Mainen. A computational framework for the study of confidence in humans and animals. *Phil. Trans. R. Soc. B*, 367:1322–1337, 2012.

[10] JM. Abolafia, M. Martinez-Garcia, G. Deco, and MV. Sanchez-Vives. Slow modulation of ongoing activity in the auditory cortex during an interval discriminationtask. *Front Integr Neurosci*, 5:60, 2011.

[11] VB. Mountcastle, MA. Steinmetz, and R. Romo. Frequency discrimination in the sense of flutter: psychophysical measurements correlated with postcentral events in behaving monkeysdiscrimination in the sense of flutter: New psychophysical measurements in monkeys. *J Neurosci*, 10:3032–3044, 1990.

[12] L. Lemus, A. Hernández, R. Luna, A. Zainos, V. Nácher, and R. Romo. Neural correlates of a postponed decision report. *Proc Natl Acad Sci USA*, 104:17174–17179, 2007.

[13] R. Luna, A. Hernández, C D. Brody, and R. Romo. Neural codes for perceptual discrimination in primary somatosensory cortex. *Nature neuroscience*, 8(9):1210–1219, 2005.

[14] R. Romo, A. Hernández, A. Zainos, L. Lemus, and C.D. Brody. Neuronal correlates of decision-making in secondary somatosensory cortex. *Nat Neurosci*, 5:1217–1225, 2002.

[15] A. Hernández, A. Zainos, and R. Romo. Temporal evolution of a decision-making process in medial premotor cortex. *Neuron*, 33(6):959–72, 2002.

[16] R. Romo, A. Hernández, and A. Zainos. Neuronal correlates of a perceptual decision in ventral premotor cortex. *Neuron*, 41:165–73, 2004.

[17] J. K. Jun, P. Miller, A. Hernández, A. Zainos, L. Lemus, C.D. Brody, and R. Romo. Heterogenous population coding of a short-term memory and decision task. *J Neurosci*, 30(3):916–929, 2010.

[18] R. Romo and E. Salinas. Flutter discrimination: neural codes, perception, memory and decision making. *Nature reviews neuroscience*, 4(3):203–218, 2003.

[19] JL. Pardo-Vázquez, V. Leboran, and C. Acuña. A role for the ventral premotor cortex beyond performance monitoring. *Proc Natl Acad Sci, USA*, 106:18815–18819, 2009.

[20] C. Shannon. A mathematical theory of communication. *Bell Syst. Tech. J.*, 27:379–423 and 623–656, July and Oct. 1948.

[21] TM. Cover and JA. Thomas. *Elements of Information Theory*. John Wiley and Sons, New York, 1991.

[22] ET. Rolls. *Memory, Attention, and Decision-Making: A Unifying Computational Neuroscience Approach*. Oxford Univ Press, Oxford, 2008.

[23] S. Panzeri, N. Brunel, NK. Logothetis, and C. Kayser. Sensory neural codes using multiplexed temporal scales. *Trends Neurosci*, 33:111–120, 2010.

[24] I. Nelken and G. Chechik. Information theory in auditory research. *Hear Res*, 229:94–105, 2007.

[25] G. Mongillo, O. Barak, and M. Tsodyks. Synaptic theory of working memory. *Science*, 319:1543–1546, 2008.

[26] G. Deco, ET. Rolls, and R. Romo. Synaptic dynamics and decision making. *Proc Natl Acad Sci USA*, 107:7545–7549., 2010.

[27] O. Barak, M. Tsodyks, and R. Romo. Neuronal population coding of parametric working memory. *J Neurosci*, 30:9424–9430, 2010.

[28] T. Webb, ET. Rolls, G. Deco, and J. Feng. Noise in attractor networks in the brain produced by graded firing rate representations. *PLoS ONE*, 6: e23630, 2011.

[29] T. Schreiber and A. Schmitz. Surrogate time series. *Physica D*, 142:346–382, 2000.

[30] S. Holm. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6:65–70, 1979.

[31] ET. Rolls, A. Treves, MJ. Tovée, and S. Panzeri. Information in the neuronal representation of individual stimuli in the primate temporal visual cortex. *J Comput Neurosci*, 4:309–333, 1997.

[32] XJ. Wang. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36:955–968, 2002.

[33] ET. Rolls and G. Deco. *The Noisy Brain: Stochastic Dynamics as a Principle of Brain Function*. Oxford Univ Press, Oxford, 2010.

[34] RS. Zucker and WG. Regehr. Short-term synaptic plasticity. *Annu Rev Physiol*, 64:355–405, 2002.

[35] L. Franco, NC. Rolls, ET. anb Aggelopoulos, and JM. Jerez. Neuronal selectivity, population sparseness, and ergodicity in the inferior temporal visual cortex. *Biol Cybern*, 96:547–560, 2007.

[36] Y. Wang. Heterogeneity in the pyramidal network of the medial prefrontal cortex. *Nat Neurosci*, 9:534–542, 2006.

[37] JM. Fuster and GE. Alexander. Neuron activity related to short-term memory. *Science*, 173:652–654, 1971.

[38] PS. Goldman-Rakic. Cellular basis of working memory. *Neuron*, 14:477–485, 1995.

[39] MJ. Tovée, ET. Rolls, A. Treves, and RP. Bellis. Information encoding and the responses of single neurons in the primate temporal visual cortex. *J Neurophysiol*, 70:640–654, 1993.

[40] ET. Rolls, A. Treves, RG. Robertson, P. Georges-François, and S. Panzeri. Information about spatial view in an ensemble of primate hippocampal cells. *J Neurophysiol*, 79:1797–1813, 1998.

[41] ET. Rolls, NC. Aggelopoulos, L. Franco, and A. Treves. Memory, attention, and decision-making: A unifying computational neuroscience approach. *Biol Cybern*, 90:19–32, 2004.

[42] MCA. Booth and ET. Rolls. View-invariant representations of familiar objects by neurons in the inferior temporal visual cortex. *Cereb Cortex*, 8: 510–523, 1998.

[43] N. Brunel and XJ. Wang. Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *Journal of Computational Neuroscience*, 11:63–85, 2001.

[44] N. Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, 8:183–208, 2000.

[45] JJ. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci USA*, 79:2554–2558, 1982.

[46] J. Hertz, A. Krogh, and RG. Palmer. *Introduction to the Theory of Neural Computation*. Addison Wesley, Wokingham, U.K, 1991.

[47] C. Kayser, CI. Petkov, and NK Logothetis. Visual modulation of neurons in auditory cortex. *Cereb. Cortex*, 18:1560–1574, 2008.

[48] L. Lemus, A. Hernández, R. Luna, A. Zainos, and R. Romo. Do sensory cortices process more than one sensory modality during perceptual judgments? *Neuron*, 67:335–348, 2010.

[49] DH. Hubel, CO. Henson, A. Rupert, and R. Galambos. Attention units in the auditory cortex. *Science*, 129:1279–1280, 1959.

[50] Y. Gottlieb, E. Vaadia, and M. Abeles. Single unit activity in the auditory cortex of a monkey performing a short term memory task. *Brain Res*, 74: 139–148, 1989.

[51] GH. Otazu, LH. Tai, Y. Yang, and AM. Zador. Engaging in an auditory task suppresses responses in auditory cortex. *Nat. Neurosci*, 12:646–654, 2009.

[52] S. Jaramillo and A. M. Zador. The auditory cortex mediates the perceptual effects of acoustic temporal expectation. *Nat. Neurosci*, 14:246–251, 2011.

[53] M.G. Shuler and M.F. Bear. Reward timing in the primary visual cortex. *Science*, 311:1606–1609, 2006.

[54] E. Selezneva, H. Scheich, , and M. Brosch. Dual time scales for categorical decision making in auditory cortex. *Curr. Biol.*, 16:2428–2433, 2006.

[55] M. Brosch, E. Selezneva, and H. Scheich. Formation of associations in auditory cortex by slow changes of tonic firing. *Hear. Res*, 271:66–73, 2011.

[56] LM. Romanski and PS. Goldman-Rakic. An auditory domain in primate prefrontal cortex. *Nat. Neurosci*, 5:15–16, 2002.

[57] GH. Recanzone. Response profiles of auditory cortical neurons to tones and noise in behaving macaque monkeys. *Hear Res*, 150:104–118, 2000.

[58] T. Hromadka, M. R. Deweese, and A. M. Zador. Sparse representation of sounds in the unanesthetized auditory cortex. *PLoS Biol*, 6:e16, 2008.

[59] S. Hocherman, D. A. Benson, M. H. Jr. Goldstein, H. E. Heffner, Hienz, and R. D. Evoked unit activity in auditory cortex of monkeys performing a selective attention task. *Brain Res.*, 117:51–68, 1976.

[60] BE. Pfingst, TA. O'Connor, and JM. Miller. Response plasticity of neurons in auditory cortex of the rhesus monkey. *Exp. Brain Res*, 29:393–404, 1977.

[61] D. A. Benson and R. D. Hienz. Single-unit activity in the auditory cortex of monkeys selectively attend-ing left vs. right ear stimuli. *Brain Res*, 159: 307–320, 1978.

[62] JB. Fritz, SA. Shamma, M. Elhilali, and D Klein. Rapid task-related plasticity of spectrotemporal receptive fields in primary auditory cortex. *Nat. Neurosci*, 6:1216–1223, 2003.

[63] JW. Schnupp, TM. Hall, RF. Kokelaar, and B. Ahmed. Plasticity of temporal pattern codes for vocalization stimuli in primary auditory cortex. *J. Neurosci.*, 26:4785–4795, 2006.

[64] P. Yin, M. Mishkin, M. Sutter, and J. B. Fritz. Early stages of melody processing: stimulus sequence and task-dependent neuronal activity in monkey auditory cortical fields a1 and r. *J. Neurophysiol*, 100:3009–3029, 2008.

[65] Y. Liu, L. Qin, X. Zhang, C. Dong, and Y. Sato. Neural correlates of auditory temporal-interval discrimination in cats. *Behav. Brain Res.*, 215: 28–38, 2010.

[66] C. Durif, C. Jouffrais, and E. M. Rouiller. Single-unit responses in the auditory cortex of monkeys performing a conditional acousticomotor task. *Exp. Brain Res*, 153:614–627, 2003.

[67] BH. Scott, BJ. Malone, and MN. Semple. Effect of behavioral context on representation of a spatial cue in core auditory cortex of awake macaques. *J. Neurosci*, 27:6489–6499, 2007.

[68] Y. Sakurai. Cells in the rat auditory system have sensory-delay correlates during the performance of an auditory working memory task. *Behav. Neurosci.*, 104:856–868, 1990.

[69] RR. Metzger, NT. Greene, KK. Porter, and JM. Groh. Effects of reward and behavioral context on neural activity in the primate inferior colliculus. *J. Neurosci.*, 26:7468–7476, 2006.

[70] Y. Komura, R. Tamura, T. Uwano, H. Nishijo, and T. Ono. Auditory thalamus integrates visual inputs into behavioral gains. *Nat. Neurosci*, 8:1203–1209, 2005.

[71] M. Brosch, E. Selezneva, and H. Scheich. Representation of reward feedback in primate auditory cortex. *Front. Syst. Neurosci*, 10, 2011.

[72] S. Atiani, M. Elhilali, SV. David, JB. Fritz, and SA. Shamma. Task difficulty and performance induce diverse adaptive patterns in gain and shape of primary auditory cortical receptive fields. *Neuron*, 61:467–480, 2009.

[73] DT. Blake, MA. Heiser, M. Caywood, and MM. Merzenich. Experience-dependent adult cortical plasticity requires cognitive association between sensation and reward. *Neuron*, 52:371–381, 2006.

[74] JB. Fritz, M. Elhilali, and SA. Shamma. Active listening: taskdependent plasticity of spectrotemporal receptive fields in primary auditory cortex. *Hear. Res*, 206:159–176, 2005.

[75] JM. Abolafia, R. Vergara, MM. Arnold, R. Reig, and MV. Sanchez-Vives. Cortical auditory adaptation in the awake rat and the role of potassium currents. *Cereb. Cortex*, 21:977–990, 2011.

[76] N. Ulanovsky, L. Las, D. Farkas, and I. Nelken. Multiple time scales of adaptation in auditory cortex neurons. *J Neurosci*, 24:10440–10453, 2004.

[77] MM. Churchland, BM. Yu, JP. Cunningham, LP. Sugrue, MR. Cohen, GS. Corrado, WT. Newsome, AM. Clark, P. Hosseini, BB. Scott, DC. Bradley, MA. Smith, A. Kohn, JA. Movshon, KM. Armstrong, T. Moore, SW. Chang, LH. Snyder, SG. Lisberger, NJ. Priebe, IM. Finn, D. Ferster, SI. Ryu, G. Santhanam, M. Sahani, and KV. Shenoy. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci*, 13:369–378, 2010.

[78] P. Kara, P. Reinagel, and RC. Reid. Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron*, 27:635–646, 2000.

[79] CJ. McAdams and JH. Maunsell. Effects of attention on the reliability of individual neurons in monkey visual cortex. *Neuron*, 23:765–773, 1999.

[80] JF. Mitchell, KA. Sundberg, and JH. Reynolds. Differential attention-dependent response modulation across cell classes in macaque visual area v4. *Neuron*, 55:131–141, 2007.

[81] T. Lu and X. Wang. Information content of auditory cortical responses to time-varying acoustic stimuli. *J. Neurophysiol.*, 91:301–313, 2004.

[82] I. Nelken, G. Chechik, TD. MrsicFlogel, AJ. King, and JW. Schnupp. Encoding stimulus information by spike numbers and mean response time in primary auditory cortex. *J Comput Neurosci*, 19:199–221, 2005.

[83] C. Kayser, MA. Montemurro, NK. Logothetis, and S. Panzeri. Spike-phase coding boosts and stabilizes information carried by spatial and temporal spike patterns. *Neuron*, 61:597–608, 2009.

[84] S. Panzeri, R. Senatore, MA. Montemurro, and RS. Petersen. Correcting for the sampling bias problem in spike train information measures. *J Neurophysiol*, 98:1064–1072, 2007.

[85] R. Brasselet, S. Panzeri, NK. Logothetis, and C. Kayser. Neurons with stereotyped and rapid responses provide a reference frame for relative temporal coding in primate auditory cortex. *J Neurosci*, 32:2998–3008, 2012.

[86] SN. Baker and RN. Lemon. Precise spatiotemporal repeating patterns in monkey primary and supplementary motor areas occur at chance levels. *J Neurophysiol*, 84:1770–1780, 2000.

[87] A. Ponce-Alvarez, BE. Kilavik, and A. Riehle. Comparison of local measures of spike time irregularity and relating variability to firing rate in motor cortical neurons. *J Comput Neurosci*, 29:351–365, 2010.

[88] MP. Nawrot, C. Boucsein, V. Rodriguez-Molina, A. Riehle, A. Aertsen, and S. Rotter. Measurement of variability dynamics in cortical spike trains. *J Neurosci Methods*, 169:374–390, 2008.

[89] EE. Galindo-Leon, FG. Lin, and RC. Liu. Inhibitory plasticity in a lateral band improves cortical detection of natural vocalizations. *Neuron*, 62:705–716, 2009.

[90] I. Nelken. Inhibitory plasticity in auditory cortex. *Neuron*, 62:605–607, 2009.

[91] KM. Bieszczad and NM. Weinberger. Representational gain in cortical area underlies increase of memory strength. *Proc Natl Acad Sci USA*, 107:3793–3798, 2010.

[92] DB. Polley, EE. Steinberg, and MM. Merzenich. Perceptual learning directs auditory cortical map reorganization through top-down influences. *J Neurosci*, 26:4970–4982, 2006.

[93] RG. Rutkowski and NM. Weinberger. Encoding of learned importance of sound by magnitude of representational area in primary auditory cortex. *Proc Natl Acad Sci U S A*, 102:13664–13669, 2005.

[94] CE. Schreiner and JA. Winer. Auditory cortex mapmaking: principles, projections,and plasticity. *Neuron*, 56:356–365, 2007.

[95] M. Brown, DR. Irvine, and VN. Park. Perceptual learning on an auditory frequency discrimination task by cats: association with changes in primary auditory cortex. *Cereb Cortex*, 14:952–965, 2004.

[96] AK. Churchland, R. Kiani, R. Chaudhuri, XJ. Wang, A. Pouget, and MN. Shadlen. Variance as a signature of neural computations during decision making. *Neuron*, 69:818–831, 2011.

[97] MN. Shadlen and WT. Newsome. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci*, 18:3870–3896, 1998.

[98] D. Sussillo and LF. Abbott. Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63:544–557, 2009.

[99] MM. Churchland, BM. Yu, SI. Ryu, G. Santhanam, and KV. Shenoy. Neural variability in premotor cortex provides a signature of motor preparation. *J Neurosci*, 26:3697–3712, 2006.

[100] C. Hussar and T. Pasternak. Trial-to-trial variability of the prefrontal neurons reveals the nature of their engagement in a motion discrimination task. *Proc Natl Acad Sci U S A*, 107:21842–21847, 2010.

[101] MR. Cohen and JH. Maunsell. Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci*, 12:1594–1600, 2009.

[102] JF. Mitchell, KA. Sundberg, and JH. Reynolds. Spatial attention decorrelates intrinsic activity fluctuations in macaque area v4. *Neuron*, 63:879–888, 2009.

[103] G. Deco and E. Hugues. Neural network mechanisms underlying stimulus driven variability reduction. *PLoS Comput Biol*, 8:e1002395, 2012.

[104] C. Curto, S. Sakata, S. Marguet, V. Itskov, and KD. Harris. A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex. *J Neurosci*, 29:10600–10612, 2009.

[105] G. Chechik, MJ. Anderson, O. Bar-Yosef, ED. Young, N. Tishby, and I. Nelken. Reduction of information redundancy in the ascending auditory pathway. *Neuron*, 51:359–36, 2006.

[106] MR. DeWeese, M. Wehr, and AM. Zador. Binary spiking in auditory cortex. *J Neurosci*, 23:7940–7949, 2003.

[107] DD. Gehr, H. Komiya, and JJ. Eggermont. euronal responses in cat primary auditory cortex to natural and altered species-specific calls. *N. Hear Res*, 150:27–42, 2000.

[108] K. Imaizumi, N. J. Priebe, T. O. Sharpee, S. W. Cheung, and C. E. Schreiner. Encoding of temporal information by timing, rate, and place in cat auditory cortex. *PLoS ONE*, 5:e11531, 2010.

[109] S. Furukawa and JC. Middlebrooks. Cortical representation of auditory space: information-bearing features of spike patterns. *J Neurophysiol*, 87: 1749–1762, 2002.

[110] C. Kayser, NK. Logothetis, and S. Panzeri. Millisecond encoding precision of auditory cortex neurons. *Proc. Natl. Acad. Sci. U.S.A*, 107:16976–16981, 2010.

[111] M. Sakai, S. Chimoto, L. Qin, and Y. Sato. Differential representation of spectral and temporal information by primary auditory cortex neurons in awake cats:relevance to auditory scene analysis. *D. Brain Res*, 1265:80–92, 2009.

[112] S. Panzeri, G. Biella, ET. Rolls, WE. Skaggs, and A. Treves. Speed, noise, information and the graded nature of neuronal responses. *Network*, 7:365–370, 1996.

[113] G. Paxinos and C. Watson. The rat brain in stereotaxic coordinates. *San Diego CA: Academic*, 1998.

[114] NN. Doron, JE. Ledoux, and MN. Semple. Redefining the tonotopic core of rat auditory cortex: physiological evidence for a posterior field. *J. Comp. Neurol*, 453:345–360, 2002.

[115] MS. Malmierca. The structure and physiology of the rat auditory system: an overview. *Int Rev Neurobiol*, 56:147–211, 2003.

[116] I. Nelken, A. Fishbach, L. Las, N. Ulanovsky, and D. Farkas. Primary auditory cortex of cats: feature detection or something else? *Biol Cybern*, 89:397–406, 2003.

[117] H. Ojima and K. Murakami. Intracellular characterization of suppressive responses in supragranular pyramidal neurons of cat primary auditory cortex in vivo. *Cereb Cortex*, 12:1079–1091, 2002.

[118] CS. Pierce and J. Jastrow. On small differences in sensetion. *Proceedings of the National Academy of Sciences of the USA*, 3:73–83, 1884.

[119] DM. Johnson. Confidence and speed in the two-category judgment. *Archs Psychol*, 34:1–53, 1939.

[120] L. Festinger. Studies in decision: I. decision-time, relative frequency of judgment and subjective confidence. *Journal of Experimental Psychology*, 32:291–306, 1943.

[121] R. Pierrel and C. S. Murray. Some relation between comparative judgment confidence and decision-time in weight lifting. *American Journal of Psychology*, 76:28–38, 1963.

[122] D. Vickers and J. Packer. Effects of alternating set for speed or accuracy on response time, accuracy and confidence in a unidimensional discrimination task. *Acta Psychologica*, 50:179–197, 1982.

[123] HE. Garret. A study of the relation of accuracy to speed. *Archs Psychology*, 56:1–105, 1922.

[124] A. Kepecs, N. Uchida, H. A. Zariwala, and Z. F. Mainen. Neural correlates, computation and behavioural impact of decision confidence. *Nature*, 455: 227–231, 2008.

[125] R. Kiani and M. N. Shadlen. Representation of confidence associated with a decision by neurons in the parietal cortex. *Science*, 324:759–764, 2009.

[126] J. Gold and M. Shadlen. Representation of a perceptual decision in developing oculomotor commands. *Nature*, 404:390–394, 200.

[127] MN. Shadlen and WT. Newsome. Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *J Neurophysiolgy*, 86: 1916–1936, 2001.

[128] R. Moreno-Bote. Decision confidence and uncertainty in difusion models with partially correlated neuronal integrators. *Neural Computation*, 7:1786–1811, 2010.

[129] J. Drugowitsch, R. Moreno-Bote, AK. Churchland, MN. Shadlen, , and A. Pouget. The cost of accumulating evidence in perceptual decision making. *The Journal of Neuroscience*, 32:3612–3628, 2012.

[130] M. Graziano and M. Sigman. The spatial and temporal construction of confidence in the visual scene. *PLoS ONE*, 4:e4909, 2009.

[131] A. Insabato, M. Pannunzi, ET. Rolls, and G. Deco. Confidence-related decision-making. *Journal of Neurophysiology*, 104(1):539–47, 2010.

[132] TJ. Pleskac and JR Busemeyer. Two-stage dynamic signal detection: a theory of choice, decision time and confidence. *Psychological Review*, 1: 864–901, 2010.

[133] V. de Lafuente and R. Romo. Neural correlate of subjective sensory experience gradually builds up across cortical areas. *Proc Natl Acad Sci, USA*, 103:14266–14271, 2006.

[134] R. Romo, A. Hernández, and A. Zainos. Neuronal correlates of a perceptual decision in ventral premotor cortex. *Neuron*, 41:165–73, 2004.

[135] W.R. Softky and C. Koch. The highly irregular firing of cortical cells is inconsistent with temporal integration of random epsps. *The Journal of Neuroscience*, 13(1):334–350, 1993.

[136] W. Bair, C. Koch, W. Newsome, and K. Britten. Power spectrum analysis of bursting cells in area mt in the behaving monkey. *The Journal of neuroscience*, 14(5):2870–2892, 1994.

[137] MN. Shadlen and WT. Newsome. Noise, neural codes and cortical organization. *Current opinion in neurobiology*, 4(4):569–579, 1994.

[138] R. Romo, A. Hernández, A. Zainos, C.D. Brody, and L. Lemus. Sensing without touching: psychophysical performance based on cortical microstimulation. *Neuron*, 1:273–278, 2000.

[139] GD. Horwitz and WT. Newsome. Target selection for saccadic eye movements: Prelude activity in the superior colliculus during a direction-discrimination task. *Journal of Neurophysiology*, 86:2543–2558, 2001.

[140] A. Ponce-Alvarez, V. Nácher, R. Luna, A. Riehle, and R. Romo. Dynamics of cortical neuronal ensembles transit from decision making to storage for later report. *The Journal of Neuroscience*, 32(35):11956–11969, 2012.

[141] A. Bollimunta, D. Totten, and J. Ditterich. Neural dynamics of choice: Singletrial analysis of decision-related activity in parietal cortex. *J Neurosci*, 32:12684–12701, 2012.

[142] T. Metzinger. Neural correlates of consciousness: Empirical and conceptual questions. *MIT press*, 2000.

[143] D. Marr. *Vision*. Freeman, New York, 1982.

[144] NR. Draper and H. Smith. *Applied Regression Analysis*. John Wiley and Sons, New York, 1966.

[145] L. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of IEEE*, 77:716723, 1989.

[146] E. Seidemann, I. Meilijson, M. Abeles, H. Bergman, and E. Vaadia. Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *Journal Neuroscience*, 16:752–768, 1996.

[147] C. Kemere, G. Santhanam, Byron MY., A. Afshar, SI. Ryu, TH. Meng, and KV. Shenoy. Detecting neural state transitions using hidden markov models for motor cortical prostheses. *Journal Neurophysiol*, 100:2441–2452, 2008.

[148] LM. Jones, A. Fontanini, BF. Sadacca, P. Miller, and DB. Katz. Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proc Natl Acad Sci U S A*, 104:18772–18777, 2007.

[149] G. Schwarz. Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464, 1978.

[150] H. Akaike. new look at the statistical model identiffication. *IEEE Transactions on Automatic Control*, 19(6):716723, 1974.

[151] V. Nácher, A. Ledberg, G. Deco, and R. Romo. Coherent delta-band oscillations between cortical areas correlate with decision making. *Proc. Natl. Acad. Sci. U.S.A*, 110(37):15085–15090, 2013.

[152] A. Borst and F.E. Theunissen. Information theory and neural coding. *Nature neuroscience*, 2:947–958, 1999.

[153] R.Q. Quiroga and S. Panzeri. Extracting information from neuronal populations: information theory and decoding approaches. *Nature reviews neuroscience*, 10(3):173–185, 2009.

[154] S. Nirenberg, SM. Carcieri, AL. Jacobs, and P.E. Latham. Retinal ganglion cells act largely as independent encoders. *Nature*, 411(6838):698–701, 2001.

[155] S.P Strong, R. Koberle, R.R. de Ruyter van Steveninck, and W. Bialek. Entropy and information in neural spike trains. *Physical review letters*, 80 (1):197, 1998.

[156] J. Massey. Causality, feedback and directed information. In *Proc. Int. Symp. Inf. Theory Applic*, 1990.

[157] C WJ. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, 37(3):424–438, 1969.

[158] Gerhard Kramer. *Directed information for channels with feedback*. PhD thesis, University of Manitoba, Canada, 1998.

[159] J.L. Massey and P.C. Massey. Conservation of mutual and directed information. In *Proc. Int. Symp. Inf. Theory*, 2005.

[160] Thomas Schreiber. Measuring information transfer. *Phys. Rev. Lett.*, 85: 461–464, 2000.

[161] B Gourévitch and JJ. Eggermont. Evaluating information transfer between auditory cortical neurons. *Journal of Neurophysiology*, 97(3):2533–2543, 2007.

[162] M. Besserve, B. Schölkopf, NK. Logothetis, and S. Panzeri. Causal relationships between frequency bands of extracellular signals in visual cortex revealed by an information theoretic analysis. *Journal of computational neuroscience*, 29(3):547–566, 2010.

[163] K. So, A. C. Koralek, K. Ganguly, M. C. Gastpar, and J. M. Carmena. Assessing functional connectivity of neural ensembles using directed information. *Journal of neural engineering*, 9(2):026004, 2012.

[164] C.J. Quinn, T.P. Coleman, N. Kiyavash, and N.G. Hatsopoulos. Estimating the directed information to infer causal relationships in ensemble neural spike train recordings. *Journal of computational neuroscience*, 30(1):17–44, 2011.

[165] J. Jiao, HH. Permuter, L. Zhao, YH. Kim, and T. Weissman. Universal estimation of directed information. *IEEE Trans. Inf. Theory*, 59(10):6220–6242, 2013.

[166] FMJ. Willems, YM. Shtarkov, and TJ. Tjalkens. The context-tree weighting method: Basic properties. *IEEE Trans. Inf. Theory*, 41(3):653–664, May 1995.

[167] H. Tuckwell. *Introduction to Theoretical Neurobiology*. Cambridge University Press, Cambridge, UK, 1988.

[168] Y. Gao, I. Kontoyiannis, and E. Bienenstock. Estimating the entropy of binary time series: Methodology, some theory and a simulation study. *Entropy*, 10(2):71–99, 2008.

[169] E. Maris and R. Oostenveld. Nonparametric statistical testing of eeg- and meg-data. *Journal of Neuroscience Methods*, 164(3):177–190, 2007.

[170] B. Phipson and G. K Smyth. Permutation p-values should never be zero: calculating exact p-values when permutations are randomly drawn. *Statistical Applications in Genetics and Molecular Biology*, 9:346–382, 2010.

keep swimming just keep swimming