Reading out neural populations

Shared variability, global fluctuations and information processing

Iñigo Romero Arandia

TESI DOCTORAL UPF / 2017

Thesis Supervisor

Prof. Rubén Moreno Bote,
Department of Information and Communication Technologies

A la amatxi,

"There are things known and there are things unknown, and in between are the doors of perception."
Aldous Huxley

"Whatever we call reality, it is revealed to us only through the active construction in which we participate."
Ilya Prigogine

"Science may set limits to knowledge, but should not set limits to imagination."
Bertrand Russell

"La nasa sirve para coger peces;
cogido el pez, olvídate de la nasa.
La trampa sirve para cazar conejos;
cazado el conejo, olvídate de la trampa.
La palabra sirve para expresar la idea:
comprendida la idea, olvídate de la palabra."
Zhuang zi

"Si contemplas los distintos tipos de sufrimiento,
encontrarás que, en última instancia,
todos ellos están enraizados en el estado de la mente."
Dalai Lama

"Al Norte y al Sur, al Este y al Oeste, el hombre serrucha, con
delirante entusiasmo, la rama donde está sentado."
Eduardo Galeano

Esker onak / Agraïments / Agradecimientos

Cuando llega el momento de terminar una etapa tan larga, es difícil recordar a toda la gente que ha contribuido de una manera o de otra. Asumiendo que probablemente me dejaré a alguien, y pidiendo disculpas por adelantado, haré el intento.

El primero que merece aparecer aquí es Rubén, por su amabilidad en el trato y su paciencia durante estos años y por todo lo que he aprendido a su lado. Y por supuesto Ramón, compañero infatigable de problemas, dudas existenciales, y de alegrías también. Siempre dispuesto a escuchar y a ofrecer una palabra amable. Thanks as well to the other members of the lab, for all the lunches and conversations together: Alex, Gabriela, Phillip, Thijs, Irene, Belen, David.

I want to acknowledge our experimental collaborators, Adam Kohn and Seiji Tanabe, for their important contribution to this work. I also want to thank Jan, for his invaluable help with statistical methods used in the projects presented here, and its patience to explain them. I really appreciate the hospitality you show during my stay in Geneva. Alex Pouget also deserves a position here, for allowing me to enjoy amazing discussions spending three months in his lab.

Me acuerdo de Luca, y también de Andrés, por aparecer en el momento justo en el lugar adecuado, por tantas jornadas de escalada aderezadas con filosofía y análisis político mientras nos perdíamos por los Alpes como niños pequeños. Je tiens aussi à remercier à Cyril pour ces jolies coins qu'il m'a fait decouvrir aux Alpes.

I want to thank all the developers that contributed to create and maintain Python and all the packages and programs I used during this PhD. Open source software is an exemplary movement that shows that self-organization works, and that it worths the effort of working for a

common goal, beyond individualism and economic profit.

Nire esker ona izan ditudan irakasle guztiei. Ama Guadalupekoa ikastetxeko Jose Karlosi, fisikaren mundua baino zerbait gehiago irekitzeagatik. Enekori ere, historiaren garrantzia eta interesa kontagiatzeagatik. EHUko Josu Igartuak aipamen berezia merezi du, 2. legea ezagutarazteagatik, eta beti entzun eta laguntzeko prest egoteagatik. Aitor Bergarari doktoretza bilatzerakoan eskainitako laguntzagatik. A Manuel Tello. I al Conrad Pérez, per tot el que vaig aprofundir a la física estadística, per obrir-me la porta de les neurociències, i ajudar-me a travessar-la. Olatzek ere merezi du hemen tartetxoa, biokimikatik ere nire jakinmina sustatzeagatik eta motibatzeagatik.

Por supuesto, no me puedo olvidar de toda la gente del Hospital Sant Joan de Deu, del grupo de psicologia experimental y del apal, por tantas comidas (y algunas cenas) compartidas. Gracias a Claudia, Joel, Irune, Jordi, Laura, Jorgina, Helena, Maria, Ferrán, Nagore, Carles, Gemma, Sole, Guillem, Elia, Raquel, ... espero que los puntos suspensivos compensen mi mala memoria.

Gracias también a la gente del CBC, y en especial a la pecera, por acogerme con los brazos abiertos y por su alegre compañía. A Txema y a Ana por estar tan pendientes de mí en la recta final.

Mención especial merece la gente del IDIBAPS y su entorno, por tantos seminarios compartidos, por ayudarme a desarrollar ese pensamiento crítico tan importante, tanto dentro de la ciencia como fuera: Albert, Jaime, Juan Pablo, Joao, Dani, Maira, Ainhoa, Klaus, Ernest, Bernat, Genis, Maria, ...

Y por supuesto a Marina Vegué, por su coherencia y su coraje, y por sus ganas de aprender y de compartir conocimiento.

Gràcies també a l'Ester i la Rosa per acollir-me quan vaig tornar a Barcelona a fer el doctorat. I a la Cèlia per tants moments espe-

cials en aquella època. A Saül y Gian, por tantos días de montaña compartidos, tantas vías de escalada entre risas. Vaya tela, ...

A la Clara Ariza, per compartir tant, per tot el que he après al teu costat, i per ser una font inesgotable d'energia revolucionària.

I gràcies al barri de Sants i la seva gent per acollir-me durant tants anys i fer-me sentir com a casa, a les cooperatives i projectes autogestionàris. I a tota la gent amb la qual he compartit les ganes i l'esforç de millorar col·lectivament: les assemblees del 15M de la plaça de Sants, la Universitat Indignada, CanBatlló, CanVies, la Guerxa, la xarxa d'estudis, Cantonal i la gent de Vallbona, el Seminari de Economia Crítica Taifa, el grup ConCiencias ... gràcies per compartir tant i mostrar que, entre totes, un altre món és possible.

A tota la gent que ha passat pel pis de la Bordeta, per tants moments, tants àpats, postes de sol, tants sentiments. Gràcies al Dani, la Neus, la Mireia, l'Eki, el Javi, la Guayedra, la Irene, ...

Evidenment, especial menció mereixen els "molones" de Sants, per formar la meva "cuadrilla" a Barcelona, per acompanyar-me sempre i no deixar que em senti sol: Irene, Mireia, Oscar, Isra, Dani, Javi. Sense oblidar a la Neus i el Pepe.

Muchas gracias a Li Ping, Adrià, Aroa, ... por tantas facilidades, por esas ganas de enseñar y compartir. Por mostrarme otra forma de interpretar la naturaleza y el ser humano desde la medicina china.

Al loco de Pablo, por cuidarme tanto y mostrame otra manera de ver la vida, siempre con alegría y buen humor.

Mil gràcies a Clara Sendra, per aparèixer de cop i quedar-se, per compartir tant en tant poc temps. Eta eskerrak Maialeni ere, etxetik hurbil sentiarazteagatik.

Milesker Enekori, edozein arazoren aurrean beti entzuteko eta laguntzeko prest egoteagatik. Ikerri eta koadrila guztiari ere, azaltzen naizen guztietan harrera beroa emateagatik. Bidasoako aktibistak ezin

ahaztu, bereziki Mikel eta Leire, zuon indar iraultzailea elkarbanatzeagatik, koherentzia mailak eredu izateraino eramateagatik eskerrak zuei ere. Gracias a Tali, por tantos días de escalada, y por compartir su inagotable energía cada vez que vuelvo a casa.

Esker mila Txomin eta Agurtzaneri, elkarbanatutako denbora eta energia guztiarengatik, beti alai, beti konten.

Nire eskerrik beroenak Etxalarrerantz doaz, Miren eta Paxkurentzat. Natura beste modu batean ezagutzen laguntzeagatik, ziur aski zientziak eskainitakoa baino harreman orekatuagoa eta osasuntsuagoa bultzatzeagatik. Beti elkarbanatzeko prest egoteagatik.

Y para acabar, gracias a la family, por estar siempre ahí, después de cada viaje, dispuestos a ayudar con lo que haga falta, a Begoña, Javi, Asun, Juan. Al aita, y a los bros, por soportar tantas comidas de tarro, por tantas salidas a la montaña, por tener la mente abierta y sangre en las venas. A la amona, y a los aitonas, que tanto me enseñaron cada uno a su manera, y no tuve tiempo de valorarlo y agradecerlo en su día. A la amatxi, que siempre hizo más de lo que debía, y nunca tuvo la recompensa merecida; espero que aunque tarde, le llegue mi agradecimiento infinito de alguna manera. Y evidentemente a Marina, la amatxo, por todo su esfuerzo y dedicación durante tanto tiempo, por todo su amor, que nunca podrá ser suficientemente valorado.

Bakarmena amalurrari,
bere maitasun iturri amaiezinari,
tesi hau egiteko xahututako energiarengatik,
merezi bezala zaindu ez izanagatik.
Eskerrik asko itsasoari, mendiari, zeruari,
euren handitasun zirraragarriari,
txiki sentiarazteagatik,
amalurran barne bizi naizela oroitarazteagatik.

Laburpena

Populazio neuronalek erakusten duten aktibitatearen jatorria eta funtzioa ulertzea, eta aktibitate hau nola lotzen den kanpoko estimulu sentsorialekin, erabakiekin edo ekintza motorrekin, erronka garrantzitsua da neurozientzietan. Lan honetan, tximinoen ikusmen garun-azal primarioan erregistratutako neuronen aktibitatea aztertu dugu, tximinoei orientazio ezberdinetako sareta sinusoidalak aurkezten zitzaizkien bitartean. Populazio osoaren aktibitatearen bidez neurtutako sarearen fluktuazio orokorrek neuronen selektibitatea modulatzen dutela ikusi dugu, modu biderkor eta batukorrean. Gainera, fluktuazio hauek neurona talde txikien informazioa ere aldatzen dute, neurona horien selektibitatean duten eraginaren arabera. Populazio osoak daraman informazioa, baina, ez da aldatzen aktibitate orokorarrekin. Bigarren zatian datu murritzetan korrelazio diferentzialak neurtzeko modu bat garatu dugu, eta datu experimentaletara aplikatuz informazioa mugatzen duten korrelazio hauen lehen behin-behineko estimazioa lortu dugu. Erdietsitako emaitzek informazioaren kodifikazioaren ulermenean aurreratzeko balio dute, eta aldi berean galdera gehiago sustatzen dituzte populazio neuralen aktibitateak informazioa prozesatu eta garraiatzen duen moduaren gainean.

Hitz gakoak: neurozientzia konputazionala, kodifikazio neurala, populazio neuronala, aldakortasun neuronala, fluktuazio orokorrak, informazio prozesamendua, orientazioa, prozesamendu sentsoriala, dekodifikazioa, korrelazioak, pertzepzio visuala.

Resum

Entendre l'origen i la funció de l'activitat de poblacions neuronals, i com aquesta activitat es relaciona amb els estímuls sensorials, les decisions o les accions motores és un gran repte per les neurociències. En aquest treball hem analitzat l'activitat de desenes de neurones enregistrades a l'escorça visual primària de micos mentre se'ls presentaven escletxes sinusoïdals en diferents orientacions. Hem trobat que les fluctuacions globals de la xarxa mesurades mitjançant l'activitat de la població modulen la selectivitat de les neurones de forma multiplicativa i additiva. A més, l'activitat de la població també afecta la informació present en grups petits de neurones, depenent de la modulació que ha provocat a la selectivitat d'aquestes. La informació de la població sencera, però, no canvia amb aquestes fluctuacions. A la segona part hem desenvolupat un mètode per mesurar 'correlacions diferencials' amb dades limitades. En aplicar-ho a les dades experimentals hem aconseguit la primera estimació preliminar de la grandària d'aquestes correlacions que limiten la informació. Els nostres resultats contribuexen a l'avenç en la comprensió de la codificació d'informació en poblacions neuronals, i alhora generen noves preguntes sobre com aquestes processen i transmeten informació.

Paraules clau: neurociència computacional, codificació neural, població neuronal, variabilitat neuronal, fluctuacions globals, processament d'informació, orientació, processament sensorial, decodificació, correlacions, percepció visual.

Resumen

Entender el origen y la función de la actividad de poblaciones neuronales, y cómo está actividad se relaciona con los estímulos sensoriales, las decisiones o las acciones motoras es un gran desafío en neurociencia. En este trabajo hemos analizado la actividad de decenas de neuronas registradas en la corteza visual primaria de monos mientras rejillas sinusoidales en diferentes orientaciones eran presentadas. Hemos encontrado que las fluctuaciones globales de la red medidas mediante la actividad de la población modulan la selectividad de las neuronas de manera multiplicativa y aditiva. Además, la actividad de la población también afecta a la información presente en grupos pequeños de neuronas, dependiendo de la modulación que ha provocado en la selectividad de estas neuronas. La información en la población completa, sin embargo, no varía con estas fluctuationes. En la segunda parte hemos desarrollado un método para medir 'correlaciones diferenciales' con datos limitados. Al aplicarlo a los datos experimentales hemos obtenido la primera estimación preliminar del tamaño de estas correlaciones que limitan la información. Nuestros resultados contribuyen al avance del entendimiento sobre la codificación de la información en poblaciones neuronales, y al mismo tiempo generan más preguntas sobre cómo éstas procesan y transmiten información.

Palabras clave: neurociencia computacional, codificación neural, población neuronal, variabilidad neuronal, fluctuaciones globales, procesamiento de información, orientación, procesamiento sensorial, decodificación, correlaciones, percepción visual.

Abstract

Understanding the sources and the role of the spiking activity of neural populations, and how this activity is related to sensory stimuli, decisions or motor actions is a crucial challenge in neuroscience. In this work, we analyzed the spiking activity of tens of neurons recorded in the primary visual cortex of macaque monkeys while drifting sinusoidal gratings were presented in different orientations. We found that global fluctuations of the network measured by the population activity affect the tuning of individual neurons both multiplicatively and additively. Population activity also has an impact in the information of small ensembles, which depends on the kind of modulation that the tuning of those neurons undergoes. Interestingly, the total information of the network is not altered by these fluctuations. In the second part, we developed a method to measure 'differential correlations' from limited amount of data, and obtained the first, although preliminary, estimate in experimental data. Our results have important implications for information coding, and they open new questions about the way information is processed and transmitted by the spiking activity of neural populations.

Keywords: computational neuroscience, neural coding, neural populations, neural variability, global fluctuations, information processing, orientation, sensory processing, decoding, correlations, visual perception.

Contents

List of Figures

Chapter 1

GENERAL INTRODUCTION

"Science is nothing but perception."
Plato

Brain function generates a huge curiosity, but we are far from understanding how it works. People participating in the famous Ironman races are able to swim, run and cycle for endless hours, even beyond their physical limits, thanks to a vast physical and mental preparation. Other people is able to reproduce the conditions in the universe many thousands of years ago building a 27km long circular tunnel and accelerating particles till velocities close to that of the light. Some others are able to write poems that express profound emotions, and those words can intensely affect the feelings and the behavior of many readers. It's not very strange a person that loves someone's life more than its own life, and acts accordingly. And almost everyone has memories of events that happened years ago. The capacities of the human being are impressive, but the fact that we are conscious of the myriad of complex things we are able to do, think, or feel is the striking prop-

erty that makes studying ourselves an amazing task. We know a lot about physiology and the role that different elements and substances play, but we are still very far from understanding how our bodies (and mind?) work. From previous studies we know that many of the most exciting abilities of the human being depend on the brain. Therefore, understanding how the brain works is a crucial challenge to learn more about ourselves and the nature we are part of.

One of the most salient features of humans is their ability to adapt. Humans are able to live in freezing areas close to the poles and in extremely warm deserts; in densely populated cities and in the wild jungle. In order to develop good adaptability is necessary to obtain information about the environment, and we receive this external information through our senses. See, hear, taste, touch and smell are the five different ways we have to capture information about the outside, and thus they are responsible of our perception of the world. Understanding how perception and the senses work would be an enormous advance in neuroscience, as perception is the first step to be able to communicate with the environment and respond to its variations. To reach this goal it is necessary to understand how physical stimuli like light or pressure change are transformed into electrical activity by sensory neurons, and how that information is processed in the cerebral cortex. This problem can be addressed by studying the neural code(s), the language(s) that the brain uses not only to map sensory information into neural activity, but also for motor action, decision-making or other cognitive functions.

In order to understand the neural code a basic property of cortical activity that must be taken into account is shared variability. The activity of the neurons is noisy, and part of this noise is common among neurons. For example, if the activity of a sensory neuron is recorded in two trials while the same stimulus is presented and all controllable

conditions are maintained fixed, the neural response recorded will be different. However, comparing this trial-to-trial variability among different simultaneously recorded neurons there are some similarities in the responses of the different neurons in each trial. This striking feature makes reproducibility impossible to achieve, and therefore other strategies must be developed to understand the mapping between neural responses and external stimuli. Although shared variability has been extensively studied, a clear explanation of its sources and its role is missing yet, and understanding its function (or functions) seems crucial to decipher the neural code and increase our knowledge of the brain function.

1.1 Orientation selectivity

In the 1950s, advances in experimental psychology and electrophysiology brought the two fields closer, and together with novel theoretical ideas about neuronal processing (Hebb, 1949), motivated research on the neural mechanisms underlying perception. The convergence of these two fields culminated with the discovery of orientation selectivity in the primary visual cortex of cats by Hubel and Wiesel (1959). Since then, orientation selectivity became one of the most studied properties of sensory processing and a model of cortical computation.

The physical interaction of photons in the retina is transducted into electrical activity by ganglion cells in the retina. These sensory neurons project onto the LGN nucleus in the thalamus, where neurons still do not show orientation selectivity. Then, the thalamic LGN neurons connect with the primary visual cortex, also called V1. Orientation selectivity arises in the primary visual cortex of mammals, at the first stage of cortical processing, as a consequence of the spatial organization of receptive fields of excitatory LGN cells (Chapman

et al., 1991; Hubel and Wiesel, 1962; Jin et al., 2011; Reid and Alonso, 1995), the influence of intracortical inhibitory connections (Ringach et al., 1997; Sillito, 1975; Xing et al., 2011), and the nonlinearity introduced by the spiking threshold (Carandini and Ferster, 2000). How these three elements combine to give rise to orientation selectivity and all its properties is still unknown though. Specifically, the factors that generate intracortical activity are likely to be related to global fluctuations and may play an important role for sensory processing. Supporting this hypothesis, it is known that, for example, orientation selectivity is modulated by attention (McAdams and Maunsell, 1999), but the mechanisms underpinning this effect are not understood.

1.1.1 Tuning curves

The tuning curves are one of the main tools to study sensory coding, and it is widely used to analyze orientation selectivity. It is computed by averaging, for each different stimulus, the response of a neuron across many trials with the same stimulus presentation. The tuning curve of a neuron gives a sense of how its response is modulated by the sensory variable, the orientation in this case, removing the "noise" or variability. This can be understood taking into account that each point of the tuning curve comes from a distribution, the different trials recorded while the same stimulus was shown (see figure 1.3). V1 neurons present bell-shaped tuning curves for orientation, but with a variety of width, baseline and maximum amplitudes. Tuning diversity is an extremely important property for information coding (Ecker et al., 2011; Goris et al., 2015). Bell-shaped tuning curves as those obtained in orientation-selective neurons have been found for other variables as the pitch of a tone, or a velocity of a moving bar, suggesting that the results obtained in this area are likely to generalize to other sensory domains. The similarities among different cortical areas (Douglas and Martin, 2004; Harris and Shepherd, 2015) suggest

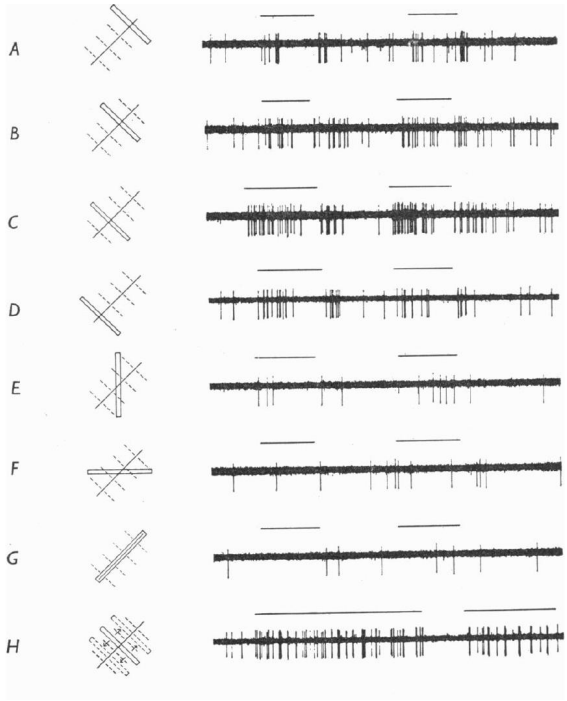

Figure 1.1: Neural responses (A-G right) of an orientation-selective cell in the cat primary visual cortex to oriented bars in different positions and orientations (A-G left). Horizontal lines in right panels indicate the presention of the stimulus on the left. Note the different responses for different orientations (D-G). By comparing the activity between the two presentations of the same stimulus (two horizontal lines on top of the activity) in each row trial-to-trial variability can be seen. Extracted from (Hubel and Wiesel, 1962)

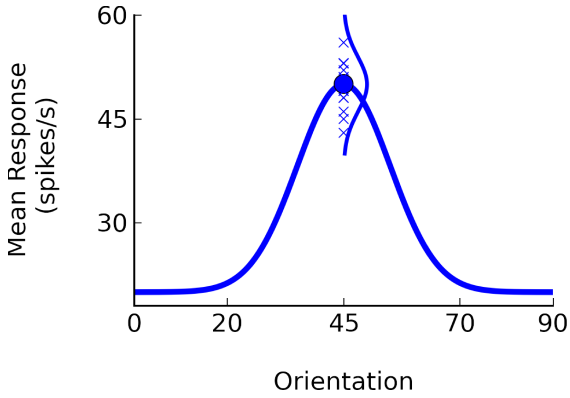

Figure 1.2: Tuning curve. Each point is the mean activity across trials with the same presented stimulus, while the cross indicates the activity in individual trials.

a common mechanism to process information, and this fact helped orientation selectivity to become an excellent model to study cortical processing.

Usually, the peak of the tuning curve is employed to characterize a neuron, and it is intuitive to think that information is encoded in this peak. However, some theoretical works found that information was better encoded in the regions of the tuning curve with larger slope, as a small change in the stimulus would produce a very different response and thus help discriminating between two similar stimuli (Abbott and Dayan, 1999; Seung and Sompolinsky, 1993). Thus, the width of the tuning curve received large attention as it could be directly related to information coding (Pouget et al., 1999; Seriès et al., 2004; Zhang and Sejnowski, 1999). Both possibilities have been found to be correct, with more information on the peak or on the slope depending on the experimental conditions and the amount of noise (Butts and Goldman, 2006).

Nevertheless, it is clear that an individual V1 neuron with neither a narrow nor a broad tuning curve cannot convey complete information about the orientation of an external stimulus. Therefore, the activity of a neuronal population, with neurons encoding different orientations is necessary to extract information about any stimulus. In order to describe the interaction among neurons with different tuning properties it is extremely important to take into account variability and correlations.

1.2 What is variability? Properties of neural responses

Facing the problem of neural code basically consists in trying to understand the encoding and decoding of information in neural responses, that is, the mapping between an external -sensory, motor or decision-variable and the neural activity. In the case of sensory stimuli, the problem reduces to decipher the sort of dictionary that transforms the physical quantities that affect our sensors (e.g., photons or pressure difference) in spikes or action potentials of neural populations. If we aim at understanding the neural code, it is important to describe first the main features of neuronal responses, and without a doubt variability is the most striking one.

1.2.1 Irregularity

The spiking activity of cortical neurons is very irregular, both during stimulus presentation and in the absence of stimulus or task. This variability can be measured using the Inter-Spike Intervals (ISI), $\Delta t_i \equiv \tau = t_i - t_{i-1}$, the temporal interval between consecutive spikes, and the Coefficient of Variation (CV). The CV is defined as the standard deviation of ISI divided by the mean: $C_v = \sigma_\tau / < \tau >$ (Dayan

et al., 2001). The distribution of ISI has been shown to follow and exponential (or Gamma) function (Softky and Koch, 1993).

Spikes fired by the same neuron are usually correlated in a non-trivial way, and the autocorrelation function gives a measure of the timescale of this correlation (Bair et al., 2001; Deger et al., 2012). Moreover, some neurons tend to generate spikes exactly at the same time. This spiking synchrony is identified as a peak in the cross-correlogram, which is a function used to quantify the correlation among spikes from two different neurons. Synchrony has been proposed as possible code used by the brain (Singer, 1999), but whether simultaneous spikes from different neurons convey extra information is not clear yet.

The firing rate approximation

One of the first approaches to deal with this variability was to assume that it is just noise. To get rid of spiking irregularity the firing rate was proposed as the meaningful quantity to study the neural code REFS. This quantity can be defined in different ways, but the most intuitive one is the discretization of time, counting the number of spikes n (spike count) fired by a neuron in a certain time window Δt: $r = n/\Delta t$. The instantaneous firing rate can be obtained as the inverse of the ISI: $r(t) = 1/\tau$. This quantity is not affected by jittering, and thus it is robust to small changes in the exact timing of the spikes. The rate-based approach is one of the most common approximations employed in the field and there is some evidence supporting this view (London et al., 2010). However, it neglects the information contained in the exact timing of the spikes, which has been shown to be important for coding in many areas (Panzeri et al., 2001; VanRullen et al., 2005). The work by Brette (2015) offers a nice dicussion of both approaches. A recent paper, however, tries to reconcile these supposed oposed views finding that both time and rate code are important and complementary (Zuo

et al., 2015).

In this work we are going to take advantage of this approximation to compute information and correlations in the neural populations.

1.2.2 Trial-to-trial variability

On the other hand, neural responses show trial-to-trial variability. If exactly the same stimulus is presented in different trials while the activity of the same neuron is recorded over and over, the spiking activity of this neuron will be different in each trial. As a consequence, there is a lack of reproducibility. Trial-to-trial variability usually is quantified computing the Fano factor, the variance of the spike count divided by the mean spike count, $\sigma_n^2 / < n >$ (Dayan et al., 2001).

This variability makes very difficult to describe the neural response to a stimulus using a deterministic function of that stimulus. The relationship between stimulus and response could be stochastic, chaotic, or/and dependent on an uncontrollable variable like attentional state or arousal. A common approach is the stochastic one, although there is no clear evidence supporting it. Spikes have been described as random events of an underlying stochastic process, also called point process. The probability of a neuron to fire a spike could depend in all the previous spikes, but assuming that just the preceding spike can affect the firing probability the spiking process can be described by a renewal process. If all the spikes are considered to be statistically independent from each other the spike train can be modeled with a Poisson process, a very common assumption in the field. The Fano factor of a Poisson process is one, and for any renewal process the Fano factor computed over long time intervals approaches C_v^2.

1.2.3 Shared variability

When recording more than one neuron simultaneously, it turns out that trial-to-trial variability is correlated among neurons (Kohn and Smith, 2005; Zohary et al., 1994). This shared variability is usually measured computing the Pearson's correlation coefficient between the spike counts of two neurons in many trials with the same stimulus presentation or under the same behavioral condition:

$$\rho_{ij} = \frac{< (r_i - < r_i >)(r_j - < r_j >) >}{\sigma_i \sigma_j} \tag{1.1}$$

It is called spike count or noise correlations, or simply correlations in this text from now on. The exact magnitude of the correlations is still a matter of debate, as recordings in the same area by different groups yielded different results (Ecker et al., 2010; Kohn and Smith, 2005). Typically, noise correlations are small and positive, with values ranging between $0.01 - 0.2$.

Another important property found experimentally is that neurons that are nearby with similar functional properties (i.e., similar tuning or signal correlations) usually have stronger correlations. Limited-range correlation structure have been found in many cortical areas, such as V1 (Ecker et al., 2010; Smith and Kohn, 2008), V4 (Cohen and Maunsell, 2009), or MT (Bair et al., 2001; Zohary et al., 1994) (see figure 1.3). This property of correlations establishes a strong constrain for the connectivity, reducing drastically the amount of plausible neural networks.

Correlations are usually lower in motor areas than in the sensory domain and it has been shown that correlations of pairs recorded from different hemispheres are very low (Cohen and Maunsell, 2009). This later finding suggests that correlations are more likely to be generated by common fluctuations in the responses of certain subsets of neu-

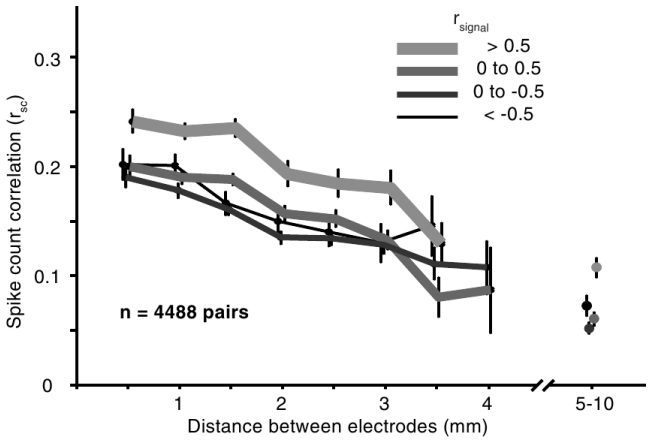

Figure 1.3: Dependence of noise correlations (r_{sc}) on distance for pairs of neurons with different tuning similarities (r_{signal}) recorded in monkey V1. Extracted from Smith and Kohn (2008).

rons, rather than reflect fluctuations that affect all neurons (Cohen and Kohn, 2011). We have focused on the second-order correlations, although higher order correlations could also be important for coding (Montani et al., 2009) (but see (Schneidman et al., 2006)).

From this section we can conclude that neural activity is conveying information through irregular spike trains that change from trial-to-trial in an structured way even if the external conditions are maintained fixed. A priori, this does not seem the best approach to process information in a reliable and efficient way, but this is likely to reflect simply how far we are from understanding the mechanisms underlying sensory processing.

1.3 How is variability arising? Sources of variability

An obvious question that arises after the description of neural variability is: where is this variability coming from? If variability is important for coding, it seems reasonable to look for the sources that control the magnitude and structure of shared variability, and thus may modulate information content in a neural population.

Probably the first and easier hypothesis is that variability is due to intrinsic neuronal noise (see the review by Faisal et al. (2008)). In this case, only individual cell properties would be the responsible of the neuronal variability. When it was found that variability is shared, it was still interpreted as noise, but with a private source from individual neurons and a common source that would induce correlations. This is a usual approach when the neural responses are described through a stochastic process. The Fano factor ($Fn = \frac{mean}{variance}$) gives a measure of neuronal variability across trials, but it does not take into account temporal structure or statistics above second order. When the Fano factor is zero, there is no variability, and in a random Poisson process the Fano factor is one. Different values of Fano factors have been found in different neurons, and even in the same area, from values close to zero (Deweese and Zador, 2004; Gur et al., 1997; Kara et al., 2000) to Fano factors larger than one (Shadlen and Newsome, 1998; Tolhurst et al., 1983). This differences may also depend on the stimulus conditions (Churchland et al., 2010; van Steveninck et al., 1997; Warzecha and Egelhaaf, 1999). But even if the statistics of spiking activity might resemble those of random processes, for example the match between mean and variance of spike counts measured through the Fano factor, that does not imply that responses come from a random process.

Many cellular processes are involved in the generation and transmission of action potentials, and most of them are noisy: protein production and degradation, opening and closing of ion channels, changes in the membrane resistance, creation and fusion of synaptic vesicles, difussion and binding of molecules to receptors ... Averaging over many stochastic elements is often assumed to remove variability. However, biochemical or electrochemical fluctuations could induce very different responses due to neuronal nonlinearities. In addition, cellular nonlinearities could also affect neural correlations (De La Rocha et al., 2007; Shea-Brown et al., 2008; Tchumatchenko et al., 2010). Nevertheless, experiments in slices show that single neuron responses are way less variable when neurons are isolated than when they are part of a network (Bryant and Segundo, 1976; Mainen and Sejnowski, 1995). Therefore, even if intrinsic noise is present in different cellular mechanisms, its impact on neuronal variability seems to be generally low. The incoherence between the amount noise in the different stages of information processing and the reliability observed in the brain could be explained taking into account that the structures and mechanisms in the brain were developed under the constraints imposed by the noise present at each level.

Another possible source of variability could be synaptic noise (Calvin and Stevens, 1967, 1968). The release of neurotransmitter vesicles is another noisy process, and it also happens in the absence of presynaptic input (Fatt and Katz, 1950). The probability of a vesicle to fuse with the membrane releasing neurotransmitters has also been proposed as a candidate signal for information processing (Abbott and Regehr, 2004).

One of the challenges when building a neural network is to reproduce similar variability as that of cortical responses. Specifically, it is difficult to generate Poisson-like variability over a broad range of rates. Moreno-Bote (2014) showed that adding probabilistic synaptic

transmission to neural networks the expected Poisson-like behavior is recovered for different magnitudes of the firing rates. This work suggests that probabilistic synapses could be an important mechanism to explain part of the neural variability present in brain activity.

The data-processing inequality makes extremely important the first stage of sensory processing, and thus the noise in the stimulus and in the transduction is very relevant for the later transmission of information. Every external stimuli is noisy, either because of thermodynamical fluctuations as in chemical sensing (smell and taste), or due to quantum mechanical effects, as the arrival of photons to photoreceptors following a Poisson process. Consequently, the transduction of energy from the sensory signal into electrical activity and its amplification will be affected by noise in the external stimulus. Due to this constrain, many organisms concentrate a large proportion of their energy resources in the first stage of processing.

1.3.1 Shared variability

From the analysis of the different sources of neuronal variability we can conclude that they are not enough to justify the amount of variability present in cortical data. Therefore, the correlated nature of variability must be taken into account if we aim at understanding the sources and the functional role of variability. The fact that neural activity is correlated opens a door to study connectivity and information transmission by analyzing shared variability. Many different possibilities have been proposed to explain the shared variability found in cortical responses, not all of them mutually exclusive. We will review some of them with a special focus on global fluctuations, as they are able to explain a large part of the variability (Ecker et al., 2014; Goris et al., 2014) with important implications for coding (Arandia-Romero et al., 2016; Pachitariu et al., 2015).

The outstanding work by Shadlen and Newsome (1998) proposed common or shared input as a potential candidate to explain pair-wise correlations. However, common inputs do not always generate strong correlations. In balanced networks composed by randomly connected excitatory and inhibitory neurons (van Vreeswijk and Sompolinsky, 1996), fluctuations in the excitatory neurons are tracked by fluctuations in the inhibitory ones (Renart et al., 2010). Therefore, correlated excitatory inputs are largely cancelled by correlated inhibitory responses (Graupner and Reyes, 2013; Salinas and Sejnowski, 2000). This mechanism could explain the low correlations found in cortex, which could be understood as a recurrent network where neurons that are close to each other receive large amounts of common inputs. In fact, the balanced networks have been proposed to describe the decorrelated or asynchronous cortical state (Ecker et al., 2010; Renart et al., 2010; Tan et al., 2014) (but see also (Bujan et al., 2015)).

Shared variability could also be a consequence of correlations in the inputs (Rosenbaum and Josić, 2011). Irregularities in the firing similar to those found in vivo have been generated in vitro with correlated inputs. For example, an image has correlations, as some areas share the same colour or the same tone, and they could generate correlations in the activity of retina ganglion cells that project to LGN and V1 afterwards. Independent Poisson inputs summed, however, did not produce supra-Poisson variability, even if excitation is balanced with inhibition (Chance et al., 2002; Stevens and Zador, 1998).

Despite the usefulness of randomly connected balanced networks, the connectivity in the cortex is clustered, and this property is likely to contribute to correlations as well. Connectivity is different among different neurons (Song et al., 2005; Yoshimura and Callaway, 2005), and those cells with similar functional properties usually show stronger connections (Bosking et al., 1997; Cossell et al., 2015). Clustering will induce large correlations inside each cluster and lower among neu-

rons from different cluster, which could partly explain the presence of limited-range correlations in sensory cortex.

Global fluctuations

The intrinsic dynamics exhibited by cortical networks can induce coordinated large-scale fluctuations that could explain part of the correlated variability found in cortical responses (Ecker et al., 2014; Goris et al., 2014; Okun et al., 2015; Pachitariu et al., 2015; Sakata and Harris, 2009; Schölvinck et al., 2015). For example, Schölvinck et al. (2015) found that the variability of one neuron can be explained in part by the summed activity of the rest of the neurons, the so-called population activity. However, Okun et al. (2015) found large diversity in the coupling strength between individual neurons and the whole population. This coupling is independent of tuning properties and stimulus conditions, suggesting that it is a property of the individual neuron. Moreover, highly coupled neurons are modulated by non-sensory behavioral variables more strongly. This work proposes that the coupling of individual neurons with the population could be a signature of synaptic connectivity.

Trial-to-trial variability is defined during the presentation of a stimulus, but in the absence of stimulation cortical neurons still show spiking activity. Therefore, there is variability in this spontaneous or ongoing activity, and it can be very useful to study intrinsic global fluctuations. The spontaneous activity has been useful to explain part of the variability (Arieli et al., 1996) and in some cases it resembled the stimulus-driven activity (Han et al., 2008; Hoffman and McNaughton, 2002; Ji and Wilson, 2007; Kenet et al., 2003; Tsodyks et al., 1999) (but see (Schölvinck et al., 2015)). This activity has structure (Luczak et al., 2007) and it is modulated by previous stimulus-evoked activity (Contreras et al., 2013). Furthermore, spontaneous activity may be

playing a role in memory recall and consolidation (Battaglia et al., 2011; Carr et al., 2011; Hoffman et al., 2007; O'Neill et al., 2010). These studies show that the structure of spontaneous activity and the way it shapes stimulus-evoked activity is tightly linked to the brain state. Therefore, the sensory system can not be explained just with a simple feedforward model, and neural responses during stimulus presentation depend both on the stimulus and on the brain state (Buonomano and Maass, 2009; Curto et al., 2009; Harris and Thiele, 2011; Renart and Machens, 2014).

There are many works suggesting that the classical view of two discrete states is likely to reflect two extremes of a continuum, with different states depending on the level of fluctuations in network activity (Luczak et al., 2007, 2009; Okun et al., 2010; Petersen et al., 2003; Poulet and Petersen, 2008). Furthermore, different behavioral and experimental conditions, such as attention or anesthesia, can lead to different cortical states. Comparing anesthetized to awake data, Ecker et al. (2014) found that accounting for cortical state in the anesthetized condition they recovered the pairwise correlations observed during awake recordings. Supporting this view, Schölvinck et al. (2015) showed that pairwise correlations were reduced and almost independent of the state once the global fluctuations were taken into account. These results suggest that different states could explain at least in part the different measures of the magnitude of correlations found in the literature (Cohen and Kohn, 2011).

In addition to correlations, cortical state also affects sensory processing (Curto et al., 2009), with, for example, more precise, reliable and informative responses during the desynchronized state of auditory cortex compared to the synchronized one (Pachitariu et al., 2015). An important variable that may change the cortical state, and at the same time affect sensory processing and correlations is attention (Harris and Thiele, 2011). Attentional mechanisms have been proposed to be similar to those used by cortical desynchronization but at a local level.

Responses to attended stimuli are larger and show lower variability and reduced correlations (Cohen and Maunsell, 2009), the same properties that characterize the desynchronized state. Our work supports this idea, as we found that tuning properties of V1 neurons are modulated by the fluctuations of population activity in a similar fashion as by attention, involving multiplication (McAdams and Maunsell, 1999; Treue and Maunsell, 1996), but also addition of the tuning curves. Adding a multiplicative factor to a Poisson model allowed Goris et al. (2014) to explain large part of the variability in different stages of the visual cortex that was independent of the stimulus. A more complete model of neural responses that takes into account feature and stimulus attention, and population activity as well has been proposed recently by Ecker et al. (2016). At the same time, Lin et al. (2015) showed that including additive offsets to a common multiplicative gain of the neurons also helps to better explain the variability in the neural responses of cortical sensory neurons. However, they applied the same multiplicative gain to the whole population, while we found that individual multiplicative factors in each neuron improved the accuracy of the model.

As explained before, the activity of the cortex is likely to have common inputs, and these can generate correlations. However, the cortex can also show very low correlations in behaving primates (Ecker et al., 2010) and in rats under urethane anesthesia (Renart et al., 2010), implying that the cortex should have a mechanism to decorrelate activity despite common inputs. One possibility is that if the excitatory and inhibitory inputs of a neuron are correlated, they can cancel each other avoiding positive correlations. Renart et al. (2010) showed theoretically that if inhibition is fast and strong enough, the fluctuations due to shared connections are tracked by inhibitory interneurons, producing small correlations.

Therefore, global fluctuations and state modulation are extremely important to understand neural processing and computation, and it

also gives important insights on the circuit and cellular mechanisms that could underlie the complex patterns recorded in neural populations.

1.4 Why is variability so important? The functional role of correlations

Why is variability important for neural coding? The nervous system is likely to face the famous problem of "reliable computation with unreliable components" (Von Neumann, 1956). Neurons seem unreliable systems, as their response to the same stimulus is different in each trial. However, they can be reliable in certain cases (Bullock, 1970), and individual neurons can contain as much information as the whole subject (Britten et al., 1992). The communication among different neurons seems unreliable as well, as it depends on noisy synapses. However, when combining responses of many neurons the nervous system is able to respond very accurately.

Shared variability is a general property of cortical responses and uncovering the possible role that correlations could play for information processing is an important challenge to understand the reliable population coding capabilities of neural responses. Solving this problem may have far-reaching consequences. For example, if noise correlations do not affect information, there would be no need for populations of simultaneously recorded neurons, and cells from different sessions could be pooled together. On the other hand, if correlations are important to encode information, the brain could have developed mechanisms to modify the size or the structure of correlations, depending for example on the level of attention or the arousal state.

One possibility is that activities from different neurons work in a synergistic way to increase the amount of information conveyed by the

sum of individual neurons, implying that correlations increase information. Another option is that correlations reflect redundancy in the neural code, and thus shared variability would reduce the information content of the sum of individual neurons. Or it could also be that correlations do not alter information at all.

The pioneering work by Britten et al. (1992) brought correlations to the center of the information coding problem. According to this work, the sensitivity of MT neurons is similar to the perceptual sensitivity of psychophysical observers. One possibility to explain this finding is that signals are carried by a small subset of neurons if responses in the population of neurons are independent. The alternative is that information is carried by a larger pool of neurons with partially correlated activity. The later was supported by Zohary et al. (1994), who recorded pairs of neurons that turned out to be correlated. If the variability of each neuron is not independent, all the noise cannot be averaged out pooling responses from different neurons. Therefore, this work suggests that information would saturate for populations of ~ 100 neurons due to correlations.

However, the important theoretical work by Abbott and Dayan (1999) showed that correlated neurons could also work in a synergistic manner and contain more information than an ensemble of independent neurons, depending on the correlation structure.

An intuitive way to analyze the effect of correlations is to compare the amount of information that can be extracted from a neural population with and without pairwise correlations. It may seem a huge simplification to analyze only second order correlations, but it can be a first step to dig into the problem and it received some support (Schneidman et al., 2006). There are two main approaches to study the effect of correlations, one focused on the encoding and another one from the decoding perspective (see Averbeck et al. (2006)

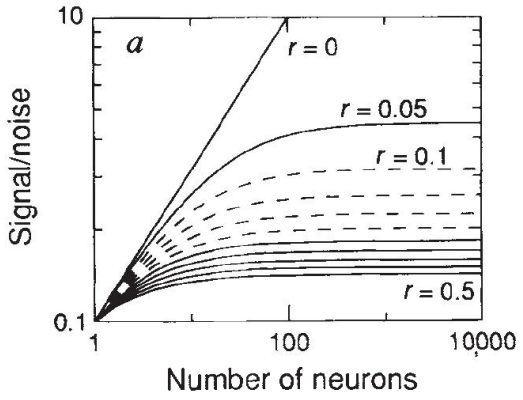

Figure 1.4: Dependence of information with population size on the level of correlations. Extracted from Zohary et al. (1994).

for a review). The encoding point of view can be analyzed comparing the information against the information of the neural response shuffling trials randomly, $\Delta I_{shuffle}$. The shuffling removes the correlations among neurons maintaining the other statistical properties of the responses unchanged, that is, the single neuron responses are the same in both cases. Therefore, this quantity $\Delta I_{shuffle}$ tell us whether adding correlations in the encoding affects information. The effect of the shuffling depends on the shape and magnitude of both signal and noise correlations (see figure 1.5).

From the point of view of the readout and computations, the goal is to obtain a decoder that is able to extract all the information present in the correlated neural responses: an optimal decoder. The measure employed in this case is the difference between the information in a population code and the information that would be extracted with a decoder trained on the shuffled data but applied to the real correlated data. This is called ΔI_{diag} when working with Fisher information

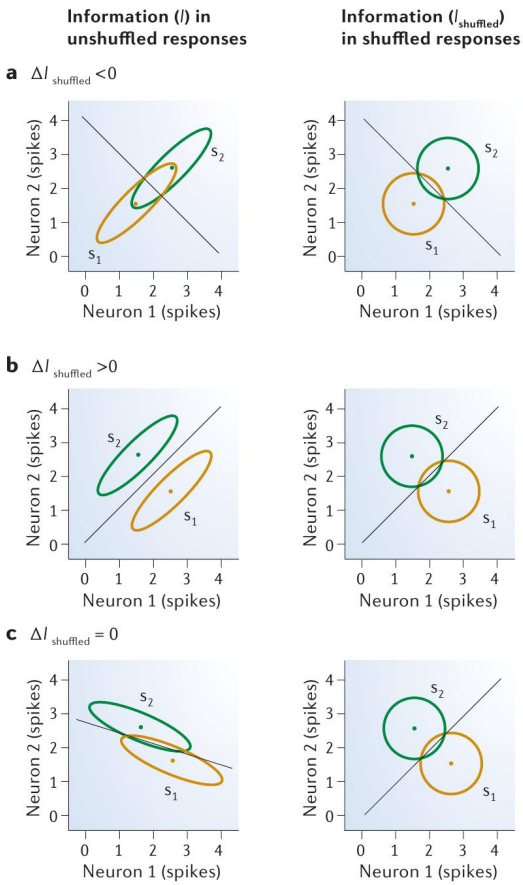

Figure 1.5: Effects of correlations on information encoding. Response distributions (ellipses show 95% confidence intervals) of two neurons that respond to different stimuli in three different cases, with correlations reducing (top), enhancing (middle) or leaving encoded information unaffected. The left column shows the correlated case and the right column the uncorrelated shuffled data. $\Delta I_{shuffled} = I - I_{shuffled}$. Extracted from Averbeck et al. (2006).

(Averbeck and Lee, 2003), and ΔI (Latham and Nirenberg, 2005) or $I_{cor-dep}$ when using mutual information (Pola et al., 2003). The role of shared variability can also be studied comparing simple decoding techniques that ignore correlations with complex algorithms that take into account high-order patterns (Graf et al., 2011; Pitkow et al., 2015). Most works found 10% loss when ignoring correlations in rat barrel cortex (Petersen and Sakmann, 2001), mouse retina (Nirenberg et al., 2001), cat V1 (Golledge et al., 2003), monkey supplementary motor area (Averbeck and Lee, 2003, 2006) and salamander retina (Oizumi et al., 2010). (Ince et al., 2010), on the contrary, found that although ΔI was small for pairs or triplets of neurons in the barrel cortex, it reached 40% when increasing the size of the population to 8 neurons. Supporting the idea that for larger populations correlations become important (Shamir and Sompolinsky, 2004), Ganmor et al. (2011) found in a population of 100 retinal ganglion cells that decoding speed was multiplied by 3 when correlations were taken into account. Graf et al. (2011) also found in V1 populations of tens of neurons that correlations increased the decoding accuracy compared to a decoder blind to correlations.

From these studies we can conclude that the results obtained for small ensembles do not always generalize to larger populations, and therefore the interpretation must be done carefully. In general we can say that correlations seem important for coding, but the mechanisms underlying this effect are still unknown.

1.5 Information, correlations and tuning properties

Information processing is affected by many variables, some of which are properties of individual neurons while others arise when analyzing populations of a certain size. To make the problem more difficult there are some variables that are very difficult to account for, such as the arousal state. As a consequence, it is difficult to disentangle the contribution of the different factors that modulate the amount of information in a neural population. Despite the effort of experimental studies on information coding to control as much variables as possible, it seems impossible to account for all possible hidden variables, such as the commented arousal state, for example. In addition, the experimental conditions already set some limitations, like the population size, or the amount of recorded trials. Moreover, the interpretation of the findings used to rely on theoretical studies that are also prone to errors. Theoretical works usually apply simplifications and/or approximations to derive their results, which in some cases turned out to affect drastically the problem under study. The distinct -sometimes contradictory- conclusions of different works are just reflecting the difficulty of the challenge, but as time goes ahead the scientific community has been able to increase the knowledge on the topic solving some of the inconsistencies.

(Zohary et al., 1994) suggested that correlations are detrimental for information, assuming that activity from different neurons is pooled together. If neural activity is correlated there is part of the variability that cannot be averaged out, and therefore information must saturate. Averaging the activity of each neuron, however, does not seem the most intelligent strategy to follow, as different neurons show different sensitivities and may play different roles for coding (Butts and

Goldman, 2006; Földiák, 1993; Jazayeri and Movshon, 2006; Salinas and Abbott, 1994; Sanger, 1996; Seung and Sompolinsky, 1993). Employing linear decoders to combine information from different neurons allows to extract more information from neural populations than averaging, and at the same time it is a biologically plausible strategy. The performance of linear decoders can be quantified by the linear Fisher information through the Cramer-Rao bound (Dayan et al., 2001; Seung and Sompolinsky, 1993).

The fact that correlations could either increase, decrease or leave information unaffected was already shown by Abbott and Dayan (1999). This work went against previous papers that suggested that correlations would generally decrease information (Shadlen et al., 1996; Zohary et al., 1994). The solution to this debate lies in the assumptions made by each work.

Theoretically, for neuronal populations with homogeneous tuning curves and positive limited-range correlations as those found in cortex, increasing correlations reduces information content (Abbott and Dayan, 1999; Sompolinsky et al., 2001). Therefore, it is not that correlations are always detrimental as Zohary et al. (1994) proposed; instead it seems that correlations that are stronger among neurons with similar properties reduce information content. These theoretical results match with experimental works showing lower correlations when information should be higher, as during attention (Cohen and Maunsell, 2009; Mitchell et al., 2009), adaptation (Gutnisky and Dragoi, 2008) or after learning (Gu et al., 2011).

The assumption of tuning curve homogeneity, however, is not likely to be accomplished by cortical data, and correlations in large populations with heterogeneous tuning curves actually increases information (Ecker et al., 2011; Shamir and Sompolinsky, 2006). The effect of correlations in experimental data under attention, adaptation or learning seems to contradict this finding. Nevertheless, it must be noted that other variables that also vary with, for example attention, could also

affect information, as the reduction in Fano factor or the increase in firing rates (Cohen and Maunsell, 2009; Mitchell et al., 2009). Larger rates increase signal, and lower Fano factors and correlations may reflect a reduction in common noise, which reduces noise entropy and therefore increases information (Ecker et al., 2011).

The role of the tuning curve width for coding has received large attention (Pouget et al., 1999; Seriès et al., 2004; Zhang and Sejnowski, 1999) but actually it is the heterogeneity in the tuning curves of the population, together with the correlation structure, what is more important than the individual tuning width (Ecker et al., 2011; Goris et al., 2015). Tuning diversity has been found to allow efficient encoding in the visual cortex: highly selective neurons encode better single orientations, while less selective neurons are responsible of encoding mixtures more accurately (Goris et al., 2015).

Another common assumption in many works has been that correlations do not depend on stimulus, but these approximation has been challenged by several works showing that the stimulus dependence of correlation could be important (Franke et al., 2016; Josić et al., 2009; Moreno-Bote et al., 2014; Pola et al., 2003; Ponce-Alvarez et al., 2013; Zylberberg et al., 2016a) and that this dependence could even increase information.

A recent work shed some light on this topic explaining that it is not the magnitude of the pair-wise correlations what matters, and suggesting that variability in a certain direction of the space defined by the activity of the different neurons in the population is the main responsible of limiting information (Moreno-Bote et al., 2014).

1.5.1 Information limiting correlations

While most of the works have focused on the effect of the magnitude of correlations on information, Moreno-Bote et al. (2014) draw the attention towards the structure of shared variability, showing than only mean fluctuations along the direction of the derivative of the tuning curve ($f'(\theta)$) in the space of neural responses limit information.

If correlations reduce information as was previously thought (Zohary et al., 1994), it would seem useful to decorrelate neural responses. This question has received large attention (Cohen and Maunsell, 2009; Mitchell et al., 2009; Renart et al., 2010) but these works assumed certain correlation structures, and the fact that information was coming from another variable spike train was not taken into account. This last condition implies that the amount of input information is finite, and according to the data-processing inequality, the output information can never be larger than the information introduced in the system. As a consequence, decorrelating activity will not always increase information, and for the same reason increasing the amount of neurons will not add extra information forever. Renart et al. (2010), for example, assumed that the information in the input came from independent neurons, and thus the fact that correlations disappear for large networks does not violate the data processing inequality as the information present in the input was already infinite. As there are less neurons in sensory organs than in primary sensory areas, the information must saturate with the number of neurons, and, as a consequence, Moreno-Bote et al. (2014) predicted the presence of differential correlations.

Moreno-Bote et al. (2014) showed that the covariance matrix can be written as:

$$\Sigma = \Sigma_0 + \epsilon f'^T(\theta) f'(\theta) \tag{1.2}$$

where the first term contains the covariance matrix with information growing with N, and the second term reflects the so-called differen-

tial correlations, those which have a huge effect on the information contained by the population. These differential correlations can be understood as a shift in the population hill (see figure 1.6). If the activity of one neuron fluctuates towards one side, due to the limited-range correlations, the neurons with similar tuning properties will be more likely to fluctuate in the same direction. Unfortunately, the decoder is blind to this shift originated by the correlated fluctuation of neurons with similar properties, as it is impossible to disentangle between a shift due to differential correlations or a change in the external stimulus. Therefore, differential, information-limiting, correlations are just a consequence of the correlation structure present in neural data.

The amount of neurons in sensory organs is smaller than in later cortical stages. This representational expansion, together with the data-processing inequality, is a possible explanation for the correlations found among cortical neurons (Kohn et al., 2016). As coding accuracy is lower in, for example retina than V1, information cannot grow linearly with the amount of neurons, and it must thus saturate (Kanitscheider et al., 2015b). As a consequence, differential correlations must be present in cortical populations.

Another source of differential correlations is suboptimal computation. Suboptimal computation, by definition, will always reduce the amount of information, and many computations in the brain have been found to be suboptimal (Beck et al., 2012). This property can be understood analyzing a variable that is encoded in a population composed of independent neurons. If the readout neuron only receives activity from a few neurons in the population, increasing the amount of downstream neurons will not increase information linearly and forever, but instead will saturate at some point.

Differential correlations can be very small compared to the magni-

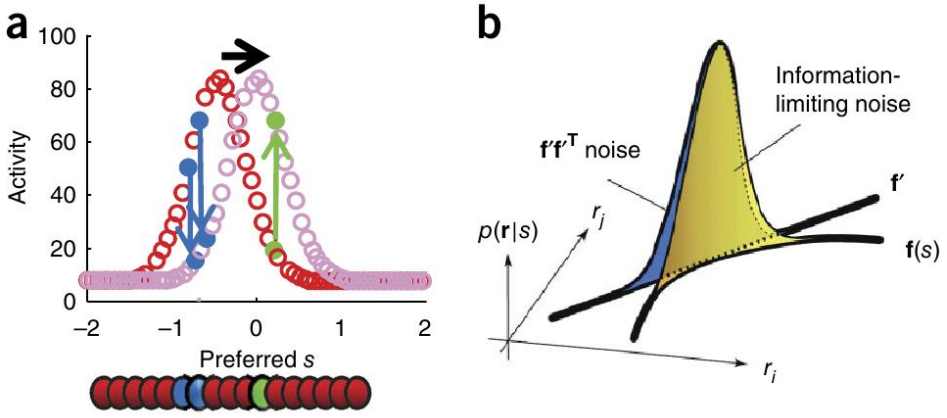

Figure 1.6: Differential correlations induced by shifting hill. (a) Population activity for neurons with translation invariant tuning curves ranked according to their preferred stimulus (s). The red and pink curves correspond to the population response to the same input on two different trials. They assumed that the variability is such that the hill simply translated sideways from trial to trial. When this was the case, correlations were mostly proportional to the product of derivatives of the tuning curves. The two blue neurons were positively correlated because the derivatives of their tuning curves were negative for both neurons. In contrast, the green neuron was negatively correlated with either of the blue because its derivative had the opposite sign as theirs. (b) Pure information-limiting noise looks like a sideways shift of the hill, corresponding to movement along the curve that represents the mean population activity ($f(s)$). If the variability is small compared to the curvature of the manifold, the distribution of the activity (yellow) can be approximated by the blue distribution that lies along the tangent to the curve ($f'(s)'$), with a resultant covariance matrix proportional to $f'(s)f'^T$. Extracted from Moreno-Bote et al. (2014).

tude of the pairwise correlations, but their effect on information can be very large. This does not mean that non-differential correlations do not affect information, but their influence is much smaller. The small size of differential correlations has two important consequences for information coding: on one hand, measuring the magnitude of these correlations in experimental data is quite difficult because they suppose just a small contribution to the total noise correlations; on the other hand, extracting meaningful conclusions about information from measures of noise correlations does not seem very reliable, as the amount of differential correlations could be very different for similar values of noise correlations. Thus, being able to obtain a measure of differential correlations would be a great advance in the study of population codes and the role that shared variability plays on information coding.

Summaryzing, we aimed at gaining some knowledge about perception by studying neural activity in the primary visual cortex. In this work we tried to analyze by parts two main sources contributing to shared variability in V1. In the first part, we studied the effect that global fluctuations had on the tuning of individual neurons. When looking for the functional roles that may come along with this tuning changes, surprisingly we found that the information encoded in the network was not affected by the global fluctuations. Information in smaller ensembles, however, is influenced by the within-state fluctuations in different ways, depending on the multiplicative or additive modulation that the tuning of their neurons undergo with these fluctuations. In this manner we offer interesting insights about the way information is encoded in the primary visual cortex of higher mammals, and the role that these global fluctuations -common to different brain areas or more local and probably coming from higher cortical areas- may play on visual information processing.

We combined this study about the top-down modulations affecting V1 with a second one that focused on the sensory information. We analyzed the impact that noise associated with the external stimuli due to anatomical and physical limitations had on the information that reaches the first stage of cortical processing. In this feedforward approach, we wanted to asses the magnitude of differential correlations with the aim of getting some insight about the way the noise and amount of neurons in retina and LGN affect the information content in V1. Finally, we combined both feedforward and top-down influences to build a model of neural activity that better resembles the scaling of information encoded in neuronal ensembles with its size in experimental recordings.

Chapter 2

MULTIPLICATIVE AND ADDITIVE MODULATION OF NEURONAL TUNING WITH POPULATION ACTIVITY AFFECTS ENCODED INFORMATION

This work has been performed in collaboration with Seiji Tanabe and Adam Kohn, who collected the data, Jan Drugowitsch, and Rubén Moreno-Bote, and it has been already published: Arandia-Romero, I., Tanabe, S., Drugowitsch, J., Kohn, A., Moreno-Bote, R. (2016). Multiplicative and Additive Modulation of Neuronal Tuning with Population Activity Affects Encoded Information. Neuron 89, 1-12.

2.1 Introduction

Neuronal activity fluctuates at both the single neuron and the population levels. These activity fluctuations can limit the reliability of neuronal codes because a given response can arise from several distinct sensory stimuli (Shadlen and Newsome, 1998; Tolhurst et al., 1983). Fluctuations in stimulus-evoked responses have been generally viewed as harmful noise that needs to be averaged out to extract the desired signal (Cohen and Maunsell, 2009; Mitchell et al., 2009; Shadlen and Newsome, 1998). Recent work has shown, however, that population activity fluctuations modulate single-cell stimulus-evoked responses in additive and multiplicative manners (Ecker et al., 2014; Goris et al., 2014; Lin et al., 2015; Schölvinck et al., 2015), suggesting that they are highly structured and hence might have a computational role. However, the role, if any, of fluctuations of total activity in neuronal populations on sensory neuronal tuning and encoding has not been demonstrated.

We studied the influence of population activity fluctuations on the responses of single neurons and small neuronal ensembles in primary visual cortex (V1) of both anesthetized and awake monkeys. We found that the tuning for stimulus orientation of orientation-selective neurons changes multiplicatively or additively with the total, stimulus-evoked activity of the neuronal population that embeds these individual neurons, while leaving their tuning width and orientation preference mostly unaffected. While distributed on a continuum, neurons with strong multiplicative effects tended to have weak additive effects and vice versa, suggesting some specificity of the modulation across neurons. Consistent with a multi-gain model of neuronal responses, we found that neurons and small neuronal ensembles with strong multiplicative effects became more informative with stronger population activity, whereas those with strong additive effects became less infor-

mative. Population activity before stimulus onset was also predictive of both tuning modulation and changes in encoded information, but to a lesser degree than stimulus-evoked population activity. Importantly, we found that population activity does not substantially alter total sensory information in the recorded population. Rather, it routes how this information is represented, in an antagonist way, into multiplicatively and additively modulated neurons and neuronal ensembles. These results suggest that intrinsic fluctuations in the activity of neuronal populations may act as a 'traffic light' that modulates the tuning of individual neurons and can differentially redistribute sensory information in the neuronal population.

2.2 Methods

2.2.1 Animal preparation

The experimental data analyzed for this work was obtained by Adam Kohn and Seiji Tanabe at the Albert Einstein College of Medicine at Yeshiva University (NY, USA). They recorded data from five adult male macaque monkeys (Macaca fascicularis), four anesthetized and one performing a fixation task. The techniques used in anesthetized animals have been previously described (Smith and Kohn, 2008).

Briefly, anesthesia was induced with ketamine (10 mg/kg) and maintained during preparatory surgery with isoflurane (1.5–2.5% in 95% O_2). Sufentanil citrate (6–24 μg/kg/h, adjusted as needed for each animal) was used to maintain anesthesia during recordings. Eye movements were suppressed with vecuronium bromide (0.15 mg/kg/h). Drugs were administered in normosol with dextrose (2.5%) to maintain physiological ion balance. We monitored physiological signs (ECG, blood pressure, SpO_2, end-tidal CO_2, EEG, temperature, urinary output and osmolarity) to ensure adequate anesthesia and animal well-

being. Temperature was maintained at 36–37 °C.

For experiments involving the awake monkey, the animal was implanted with a head post and then trained to fixate in a 1 deg window. Eye position was monitored with a high-speed infrared camera (Eyelink, 1000 Hz). 500 ms after the establishment of fixation, a drifting grating appeared over the aggregate receptive field of the recorded units. If the animal broke fixation, the trial was aborted and the data discarded. The animal was rewarded with a drop of water for successfully completed trials, typically 500-800 per session. All procedures were approved by the Institutional Animal Care and Use Committee of the Albert Einstein College of Medicine at Yeshiva University and were in compliance with the guidelines set forth in the United States Public Health Service Guide for the Care and Use of Laboratory Animals.

2.2.2 Visual stimuli

For anesthetized animals, we presented full contrast drifting sinusoidal grating for 1280 ms, with an interstimulus period of 1500 ms. Gratings of 8 different orientation were each shown 300-400 times. For the awake animal, we used 12 different orientations and each was presented for 350 ms, with an interstimulus period of 50 ms. After every 4 stimuli, the monkey was rewarded. We recorded 50 trials for each stimulus orientation.

Visual stimuli were presented on a CRT monitor, at a resolution of 1024 x 768 pixels and a video frame rate of 100 Hz. The display had a mean luminance of 40 cd/m^2, and was placed 110 cm away from the animal where it subtended 20° of visual angle. We generated stimuli with custom software based on OpenGL (EXPO). The spatial (1.3-2 cpd) and temporal frequency (6-6.25 Hz) of the gratings were

chosen to correspond to the usual preference of parafoveal V1 neurons. Stimuli were shown in a circular aperture surrounded by a gray field of average luminance. Receptive fields in both awake and anesthetized animals were 2-4 degrees from the fovea. The size of the gratings (2-4 degrees in diameter) was chosen to cover the receptive fields of all the neurons.

2.2.3 Recording methods and data preprocessing

We recorded in the superficial layers of primary visual cortex (V1), using a Utah array (96 microelectrodes, 1 mm length, 400 μm spacing; 48 electrodes in the awake animal). Events crossing a user-defined threshold were digitized (30 kHz), saved, and sorted offline. We quantified spike waveform quality using a simple signal-to-noise ratio metric (SNR; Kelly et al. (2007)). We defined multiunits (multi-unit activity; MUA) to be units with SNR>2, which corresponds to clusters with a small number of single units. Well-isolated units were defined to have SNR>3.5, a conservative threshold. A small number of MUA sites with SNR<2 were also used only to compute the population activity.

We only analyzed blocks of trials in which responses did not change markedly over time. While the data from anesthetized animals were quite stable, those from the awake preparation showed strong evidence of adaptation. We therefore removed from this dataset the first 200 trials, leaving 400 trials during which responses were stable. To avoid biases due to a different number of trials per orientation in the decoding performance analysis, we (randomly) selected 30 trials for each orientation. Due to the low number of trials compared to the 'anesthetized' datasets, we merged trials with adjacent orientations in pairs to obtain 6 orientations with 60 trials per orientation to offer more comparable results. For all datasets, we measured responses beginning 60 ms after stimulus onset to account for V1 response latencies, as in Graf et al. (2011).

2.2.4 Autocorrelation of spontaneous population activity

The autocorrelation of spontaneous population activity fluctuations is well-approximated by an exponential with a time constant of around 300 ms. We computed the autocorrelation (black lines in figure 2.7) of population activity during the last second of spontaneous activity before stimulus onset, excluding the 500 ms period after stimulus offset to avoid transients. The autocorrelation was computed for trials with the same previous stimulating grating, and then averaged across stimulus orientations, following the method used by (Kohn and Smith, 2005):

$$ACG_j(\tau) = \frac{N_j}{N_j - 1} \frac{1}{(T - |\tau|)\overline{R}} \{ \frac{1}{N_j} \sum_{k}^{N_j} \sum_{t}^{T-\tau} R^k(t) R^k(t+\tau) - \frac{1}{N_j - 1} \sum_{k}^{N_j - 1} \sum_{t}^{T-\tau} R^k(t) R^{k+1}(t+\tau) \} \tag{2.1}$$

where $R^k(t)$ is the population activity of trial k at time t (sum of spikes of orientation selective units in that millisecond), $\overline{R}$ is the mean population activity, j is the orientation, and N_j is the number of trials with orientation j. The autocorrelations were well-fit by exponential functions (red lines in figure 2.7) with time constant τ of: D1 337 ms, D2 300 ms, D3 251 ms, D4 546 ms.

2.2.5 Dependence of tuning curves on population activity

Tuning conditioned to population activity was computed for every orientation-selective unit using the following model-free approach. Orientation-selective neurons were defined as those with tuning well

fitted by a von Mises function ($r^2 \geq 0.75$) (Graf et al., 2011); remaining neurons were termed non-selective. For each trial, we computed the population activity as the average number of spikes per second across all other neurons in that time window. Population activity was based on all neurons except the one for which tuning modulation was computed. Trials corresponding to a given stimulus orientation were sorted as a function of mean population activity, and then split at their median into subgroups of 'low' and 'high' activity. We computed the 'high' and 'low' population activity tuning, denoted $f^{high}(\theta)$ and $f^{low}(\theta)$, as the firing rate of the chosen neuron as a function of stimulus orientation in trial subgroups with 'high' and 'low' population activity, respectively. For comparison, we also computed population activity after z-scoring the responses of each neuron across trials. This approach yielded similar results, because trial ranking was nearly identical with the two methods, with only a few close-to-median trials changing category.

2.2.6 Multiplicative and additive modulation of tuning

To estimate the multiplicative and additive gains, we performed a type II weighted linear regression between the 'low' and 'high' tuning, using the model $f^{high}(\theta) = gf^{low}(\theta) + s$, where g is the multiplicative factor, and s is an additive offset. The multiplicative factor is unitless by definition. To obtain a comparable unit-less additive factor, we normalized s by the mean activity across orientations. Neurons with fit values outside the range $[0.3, 3]$ for multiplicative factors and $[-1, 1]$ for additive factors were excluded from analysis. Results do not qualitatively depend on the exclusion of these few outliers (5% of cases).

For each dataset, we estimated the gains' significance by a permutation test that sampled the null hypothesis. We randomly assigned trials to build 'high' and 'low' tuning, instead of ranking trials by population activity. Then, we obtained the multiplicative and additive factors for each neuron by linear regression, as described before, and computed the median factors across all neurons. We repeated this procedure 1000 times. We defined the probability that these medians were larger or smaller than the real median across neurons by the fraction of samples below or above the real population median. The reported two-tailed p-values were twice that fraction.

To estimate the correlation between multiplicative and additive factors (Figure 2.10D) from the linear fits described above we proceeded as follows. Because the linear model contains two regressors, corresponding to the multiplicative and additive factors, a linear fit based on the same data points can create artificial correlations between the values of these two regressors (Donahue and Lee, 2015). To avoid this artifact, we estimated the factors from randomly divided trials into two halves with the same number of trials for each orientation. A Pearson correlation coefficient sample was obtained by computing this coefficient between the multiplicative factors obtained from fits in the first trial subgroup and the additive factors obtained from fits in the second subgroup. Using two non-overlapping groups of trials ensured that the factors were not trivially correlated (Donahue and Lee, 2015). This procedure was repeated 1000 times (thus obtaining 1000 samples) to build a distribution of correlations. The reported correlation was the median across samples (the displayed distribution of multiplicative and additive factors in figure 2.10D is originated from a randomly chosen subdivision of the trials). The two-tailed p-value for the reported correlation was computed as twice the fraction of samples that are below or above zero, whichever is the smaller quantity. To improve the signal to noise ratio, the analysis used responses measured

during the full trial (excluding the first 60 ms).

2.2.7 Broadening and displacement of tuning curves

We determined the change in width and preferred orientation with population activity using Von Mises function fits to each neuron for trials with 'low' and 'high' population activity:

$$f(\theta) = a + b \exp\left[k(\cos\left(2(\theta - \theta_{pref}\right) - 1\right] \tag{2.2}$$

We fit the function by minimizing the weighted squared error with bounded parameters to ensure physiologically plausible parameters (minimize function from lmfit python package with the following constraints: $\theta \in [0, \pi]$, $k > 0.0001$, $max(f_\theta) < 1.3 \cdot max(\overline{r_\theta})$, $min(f_\theta) > 0.7 \cdot min(\overline{r_\theta})$, where $(\overline{r_\theta})$ is the mean response across trials for each orientation). To evaluate broadening, we used the squared width of the von Mises distribution, defined as $\sigma^2 = 1 - (I_1(k))/(I_0(k))$, where $I_n(k)$ is the modified Bessel function of the first kind of order n evaluated at k. The broadening factor was computed as the ratio between the widths of the 'high' and 'low' tuning $(\sigma^{high}/\sigma^{low})$.

The displacement of the tuning was defined as the absolute difference between the preferred orientation of the 'high' and 'low' tuning $(\Delta\theta_{pref} = |\theta^{high}_{pref} - \theta^{low}_{pref}|)$. We tested for significant broadening and displacement with a permutation test by sampling, as described above, using two-tailed and one-tailed p-values respectively. In an additional analysis, we used the von Mises fits to compute the multiplicative and additive modulation of tuning, which yielded similar results to those reported in the main text.

2.2.8 Validation of the method to estimate multiplication, addition and broadening of tuning curves

In order to be sure that our method reliably detects the multiplicative and additive modulation of tuning curves found in experimental data, we generated neural activity using neurons with the same tuning curves as in the data, but modifying the contribution of the multiplicative, additive and broadening effects using different models. First we show that positive pair-wise correlations do not necessarily imply the presence of multiplication and/or addition. We obtained positive pairwise correlations with three different models: a purely multiplicative, a purely additive and another one with only broadening of the tuning curves (See figure 2.1).

Once that we rejected the possibility to trivially obtain multiplicative and additive effects from positive pair-wise correlations we wanted to validate the procedure employed to estimate the values of the different possible modulations that tuning curves can undergo: the multiplicative factor, the additive factor, and the broadening. We generated neural activity in populations with neurons having the same tuning curves as in each dataset, maintaining similar average correlation coefficients as well. Under these constraints, we manipulated the influence of the different factors (multiplication, addition and broadening) contributing to global fluctuations. For instance, we first simulated data using a model (the homogeneous multiplicative model; first column) with purely multiplicative factors and no addition or broadening. Our estimation method, when applied to this simulated data, correctly discovers significant multiplicative factors (top panel) and correctly does not find either additive factors (middle) or broadening (bottom). We repeated this procedure with several combinations of factors to

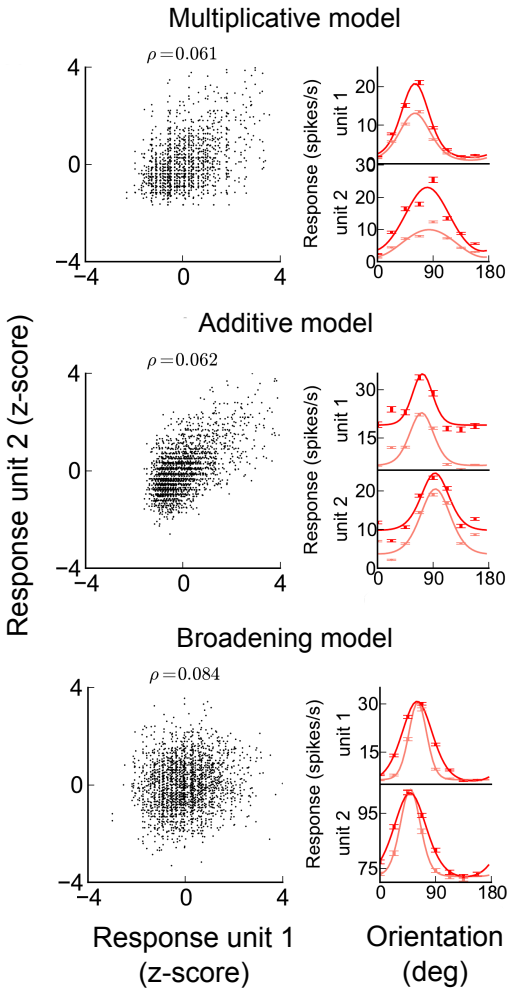

Figure 2.1: Positive pair-wise correlations can be generated either with a purely multiplicative (top panel), a purely additive (center) or a broadening (bottom) model. These models generate positive correlations in the responses of pairs of neurons. The tuning curves of two simulated pairs for each model are modulated by population activity (small panels), like in the data. Across orientations, there is a correlation between their z-scored responses.

validate our estimation method in a broad set of conditions, always showing the general validity of our method (see details of the models next). Mathematically, these models are specified by the mean firing rate of neuron i given the stimulus and the modulatory global factor. For each model, the mean firing rate is:

i Homogeneous multiplicative model: $f_i(\theta, g) = (1 + g)h_i(\theta)$ (based on Goris et al. (2014)).

ii Heterogeneous multiplicative model: $f_i(\theta, g) = (1 + \alpha_i g)h_i(\theta)$.

iii Heterogeneous additive model: $f_i(\theta, g) = h_i(\theta) + \beta_i$

iv Multi-gain model: $f_i(\theta, g) = (1 + \alpha_i g)h_i(\theta) + \beta_i g$ (multiplicative and additive as Lin et al. (2015), but with different factors for each neuron).

v Broadening model: $f_i(\theta, g) = h_i(\theta, g)$, where the global modulatory factor enters inside the tuning curve and rescales its width. Because the fluctuations of the width happen simultaneously across all neurons, this effect introduces purely positive pair-wise correlations.

For each model, spike counts in a time window of 100 ms (mimicking the time window 160-260 ms used in the data analysis of figure 2.10) were drawn from an independent Poisson distribution across neurons with mean as specified above. The global modulatory factor g is shared among neurons, which introduced pair-wise correlations. The mean of the modulatory factor g was zero. The values of g were drawn from a gamma distribution with unit mean (the variance is specified below), followed by subtracting one, to keep g bounded from below by -1 and its mean at zero. The parameters α_1 and β_i, which control the heterogeneity in multiplicative and additive factors, were drawn from Gaussian distributions with mean and standard deviation respectively

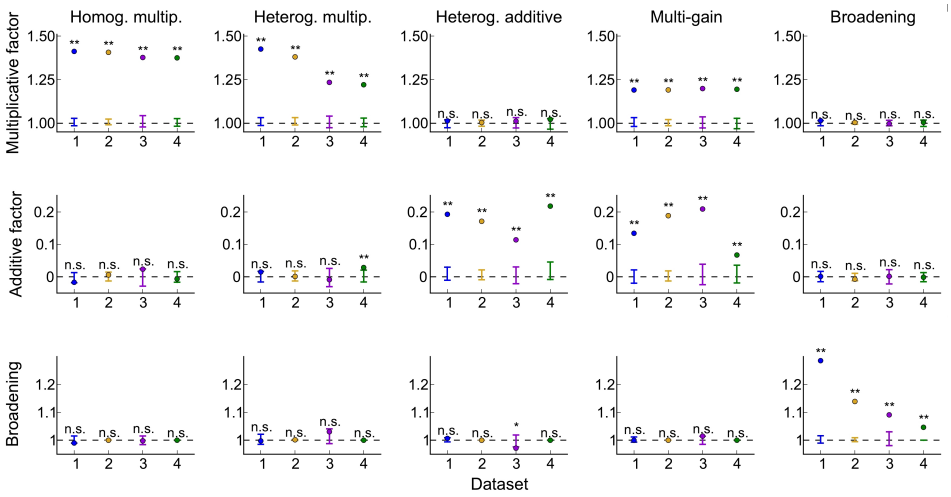

Figure 2.2: Validation of the estimation method of multiplicative and additive factors and broadening. The method reliably detects the presence of significant multiplicative, additive or broadening effects, and rejects their presence if they are absent in simulated data. Each column provides simulations of neuronal populations with tuning curves as in datasets D1-D4 (with similar average correlation coefficients) for different models. Each row shows the outcome of the estimation of multiplicative (top), additive (middle) and broadening (bottom) factors for each case. From left to right, we consider the following models: (i) homogeneous multiplicative model, with same gain for all neurons, (ii) heterogeneous multiplicative model, characterized by having a potentially different multiplicative factor for each neuron, (iii) heterogeneous additive model, with different additive factors for each neuron (iv) multi-gain model where multiplicative and additive factors are independently and heterogeneously assigned to the neurons in the populations, but with the same global modulation for both factors and (v) broadening model, characterized by having fluctuations of the width of the tuning curve and no other modulatory factor. *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$, n.s.: $p > 0.05$.

for each model (ii) $< \alpha_i >= 0.25$, $\sigma_\alpha = 0.25$, (iii) $< \beta_i >= 0.35$, $\sigma_\alpha = 0.2$, and (iv) $< \alpha_i >= 0.13$, $\sigma_\alpha = 0.2$, $< \beta_i >= 0.8$, $\sigma_\beta = 0.8$. The parameters α_i and β_i where drawn once per neuron, and did not subsequently vary across trials for a given model. The modulatory global factor, in contrast, was re-drawn for each trial, and so varied across trials. For all the models, the average tuning curves constituted those computed from real data. The variances (σ_g^2) of the global modulatory factor g took the values (i) 0.05, (ii) 2.5, (iii) 2.5, (iv) 0.33, (v) 0.4, chosen to provide values of positive average correlations comparable to those in our datasets. The median correlation coefficient obtained in this 100 ms bin across datasets for each model was: (i) 0.07, (ii) 0.06, (iii) 0.08, (iv) 0.07, (v) 0.003, while in real data it was 0.07.

Finally, we wanted to make sure that the method also detected reliably the correlation between the multiplicative and additive factors found in the experimental data. Thus, we computed the multiplication, addition, and the correlation between both factors for two models: one with correlations between multiplication and addition (C model), $f(\theta) = (1+\alpha_i g)h_i(\theta)+\beta_i g$, where $\beta_i = b_i(c_i - \alpha_i)$, and parameters $b_i = 24$, $c_i = 0.5$, $< \alpha >= 0.26$, $\sigma_\alpha = 0.8$, $\sigma_g^2 = 0.05$; and another model with uncorrelated factors (I model), with different modulations for multiplication and addition $f(\theta) = (1 + \alpha_i g_1)h_i(\theta) + \beta_i g_2$, where $< \alpha >= 0.15$, $\sigma_\alpha = 0.5$, $< \beta >= 2.5$, $\sigma_\beta = 1$, $\sigma_{g_1}^2 = \sigma_{g_2}^2 = 0.5$. The pair-wise correlations are 0.16 and 0.18 for the C and I models, respectively, close to the values for the real data (0.21). These models were simulated using a longer time window to allow direct comparison with the data analysis performed in figure 2.10D. The correlations between multiplicative and additive factors was negative and significant for the model with negatively correlated factors ($\rho_C = -0.34$, non-parametric bootstrap $p = 0.002$) and not significant for the one with independent factors ($\rho_I = 0.048$, non-parametric bootstrap $p = 0.4$).

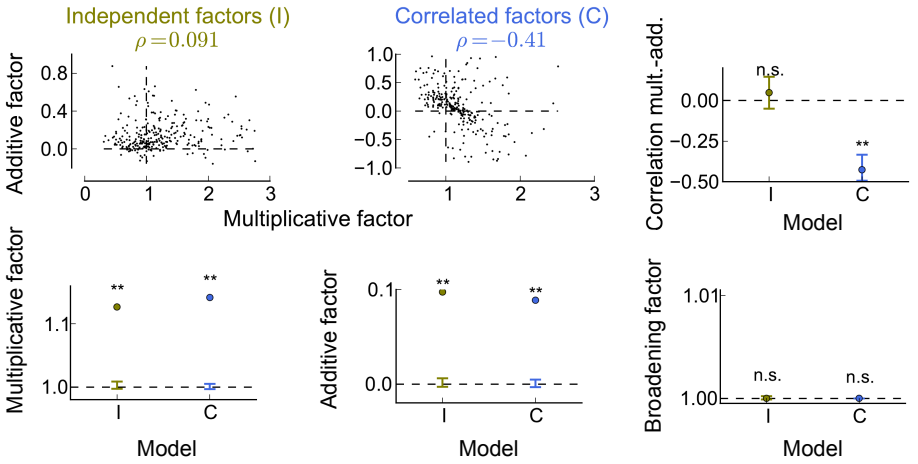

Figure 2.3: The method reliably detects the presence of negative correlations between multiplicative and additive factors, and rejects its presence if such a correlation does not exist in the simulated data. We compared the correlation between the multiplicative and additive factors in two simulated models having the same tuning curves as in the real data, both with similar multiplicative, additive and broadening factors (bottom row), but with different level of correlation between the multiplication and the addition. Our estimation method founds non-significant correlation for the independent model (I, top-left, top-right), but correctly assigns a significant correlation to the model with correlated factors (C, top-middle, top-right). *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$, n.s.: $p > 0.05$.

2.2.9 Model comparison

We computed the log-likelihood of different models to study which one was able to explain better the data we are analyzing. The models we selected are the following ones:

i Independent model: $f_i(\theta) = h_i(\theta)$.

ii Homogeneous multiplicative model: $f_i(\theta, g) = (1 + g)h_i(\theta)$ (based on Goris et al. (2014)).

iii Affine model: $f_i(\theta, g_1, g_2) = g_1 h_i(\theta) + \beta_i g_2$ (following Lin et al. (2015))

iv Multi-gain model: $f_i(\theta, g) = (1 + \alpha_i g)h_i(\theta) + \beta_i g$

We applied cross-validation to avoid overfitting. First we trained each model on 80% of the data by optimizing its parameters to maximize likelihood of this part of the data. Then, we computed the log-likelihoods for the remaining 20% of the data, averaging over 10 independent divisions of the data. We performed this analysis using the same time window of 100 ms (160-260 ms after stimulus onset) as in most of the other figures throughout the paper. Using the explained cross-validated procedure to compute the log-likelihood, it is not necessary any extra step to take into account the different number of parameters for each model, as cross-validation already controls for model complexity.

For the affine and multi-gain models the optimization procedure was as follows: first we maximized the log-likelihood of the model on the training set by coordinate ascent, maximizing the likelihood with respect to each of the parameters (the β_i-s, the g_1-s, and the g_2-s for the affine; and the α_i-s, the β_i-s and g-s for the multi-gain model) alternatively while keeping the other parameters fixed. For the multi-gain model first we maximized the log-likelihood with the α_i-s and

β_i-s fixed, and obtained the g-s. Then, we maintained the g-s and the β_i-s constant while maximizing the log-likelihood to get the new α_i-s. Finally, we maximized the likelihood again but allowing only the β_i-s to change. This procedure was repeated many times and then we selected the amount of iterations that maximized the likelihood of the test trials. A similar procedure was follow for the affine model. Note that while in figure 2.2 the mean values of the modulatory factors g were zero, in this analysis the mean values are fitted. We performed each maximization step using the Large-scale Bound-constrained Optimization (L-BFGS-B) method. Once we obtained these parameters, we computed for each neuron the modulatory factors (g_1-s, and the g_2-s for the affine, g-s for the multi-gain model) in the test trials, removing from the population in each case the neuron whose modulatory factor was being computed, and maximizing the likelihood of each trial with the Sequential Least SQuares Programming (SLSQP) method. Then, we computed the log-likelihood of the test set using the parameters from the train set, and the global modulatory factors from the test set.

2.2.10 Extracting visual information from population recordings

We defined decoding performance (Figures 2.16- 2.20) as the fraction of trials where stimulus orientation was correctly predicted by a trained decoder. Results shown are for multinomial logistic regression (MLR) (Bishop, 2006), which outperformed a linear SVM decoder (an obvious alternative). In brief, in MLR the probability of orientation θ_j is modeled as: $Pr(\theta_j|\vec{r}) = \exp(\vec{w_j} \cdot \vec{r})/(\Sigma_k \exp(\vec{w_k} \cdot \vec{r}))$, where $\vec{r} = (r_1, r_2, ...)$, is a vector of the individual firing rates of the simultaneously recorded neurons, defined as the spike count in a given time window per trial divided by the time window: $r_i = counts_i/time$. The

sets of vector parameters $\vec{w_k}$] are learned using maximum likelihood estimation using the function (smlr) from the PyMVPA Python package (Hanke et al., 2009). For each trial, the predicted stimulus orientation is given by $\theta_j^{pred} = argmax_j[Pr(\theta_j|\vec{r})]$. The decoded (predicted) orientation was considered correct if it matched the true orientation.

We used 10-fold cross-validation (CV) to avoid over fitting the data; reported performance is the average performance across the 10 sets of left-out data.

2.2.11 Decoding performance as a function of population activity

We computed the decoding performance in the time window from 160-260 ms after stimulus onset using simultaneously recorded small ensembles of N= 1, 2, 3, 5, 10, or 15 neurons. Decoding performance was independently analyzed for each orientation and each ensemble (Fig. 2.20) using MLR with the population vector formed by the firing rates of the neurons of the ensemble. Trials for each orientation were divided into 'low' and 'high' population activity trials, and their averages across trials were computed (Fig. 2.16). As for the tuning analysis, population activity was computed using all neurons but excluding the N neurons of the ensemble. Performance change per orientation was defined as the difference in decoding performance between high and low population activity trials for the MLR decoder trained in the two conditions and across orientations. We report the Pearson correlation coefficient and its two-tailed p-value, and we also plot a linear regression to highlight the relationship between performance change and mean factors in the ensemble.

2.3 Results

We recorded neuronal populations in the superficial layers of V1 in four anesthetized (datasets 1-4, D1-D4) and one awake (D5) macaque monkeys. We measured responses to gratings drifting in 8 (12 for D5) equally spaced directions. Gratings were presented for 1,280 ms (350 ms) each, interleaved with a 1,500 ms (50 ms) blank screen and repeated 300 or 400 (50) times in random order. We analyzed the activity of 567 single neurons and multiunits, which we refer to together as 'units'. We analyzed 122, 106, 73, 161, and 18 simultaneously-recorded units in datasets D1 to D5, respectively. We also analyzed separately a subset of 83 well-isolated single neurons (27, 14, 7, 31 and 4 from D1-D5, of which 12, 12, 4, 15, 2 were orientation selective; see Methods, section 2.2). The firing rate of many V1 neurons is tuned to the orientation of a drifting grating (illustrated in figure 2.4A). Since neurons are embedded in a local network and are correlated (median of pair-wise spike count correlations: $\rho = 0.21$ across all anesthetized datasets), the summed total activity of that local population (called population activity) might modulate this tuning (Figure 2.4B). This modulation could involve multiplicative or additive effects, or both, as well as broadening and displacement (Figure 2.4C). Similarly, positive correlations among neurons can arise in multiple ways, such as additive modulation, multiplicative modulation, broadening of tuning, or a combination of these effects or others (Figure 2.1). Therefore, the existence of correlations does not specify how tuning is modulated. We thus developed an analysis that could distinguish how tuning is modulated with population activity fluctuations.

In our data, population activity showed substantial fluctuations across trials for a fixed stimulus condition (several representative trials shown in Fig. 2.5A, left), consistent with previous reports (Arieli et al., 1996; Ecker et al., 2014; Schölvinck et al., 2015). The distribution of spike counts during the stimulation period across trials for one

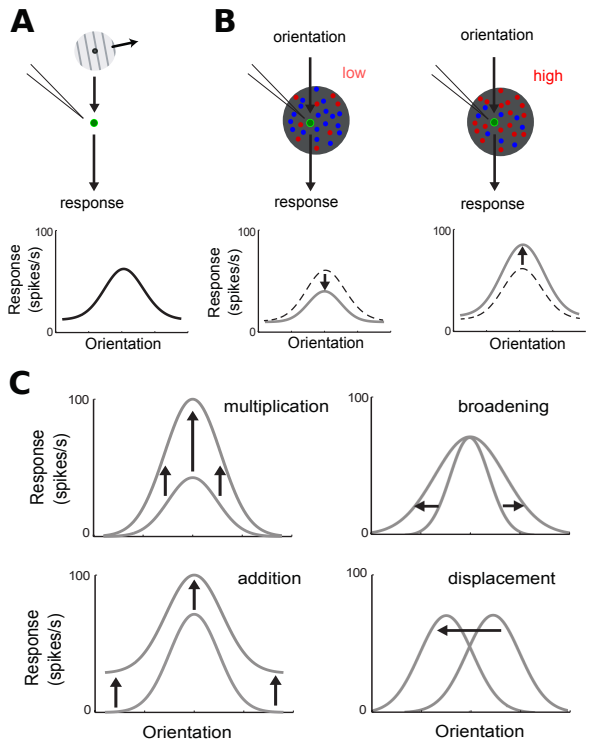

Figure 2.4: Hypothetical modulation of sensory tuning with population activity. (A) The 'classical' tuning of a neuron, computed without conditioning on population activity. (B) The firing rate of the neuron can depend on population activity. When population activity is low, tuning could have a lower gain (gray line, left); when population activity is high, tuning might have a higher gain (right). If the activity of the neuron is independent of population activity, the tuning for the two cases would be identical to the 'classical' tuning (dashed lines). (C) Tuning can be modulated in several ways with fluctuations in population activity, including multiplicative and additive effects, or both, and broadening and displacement.

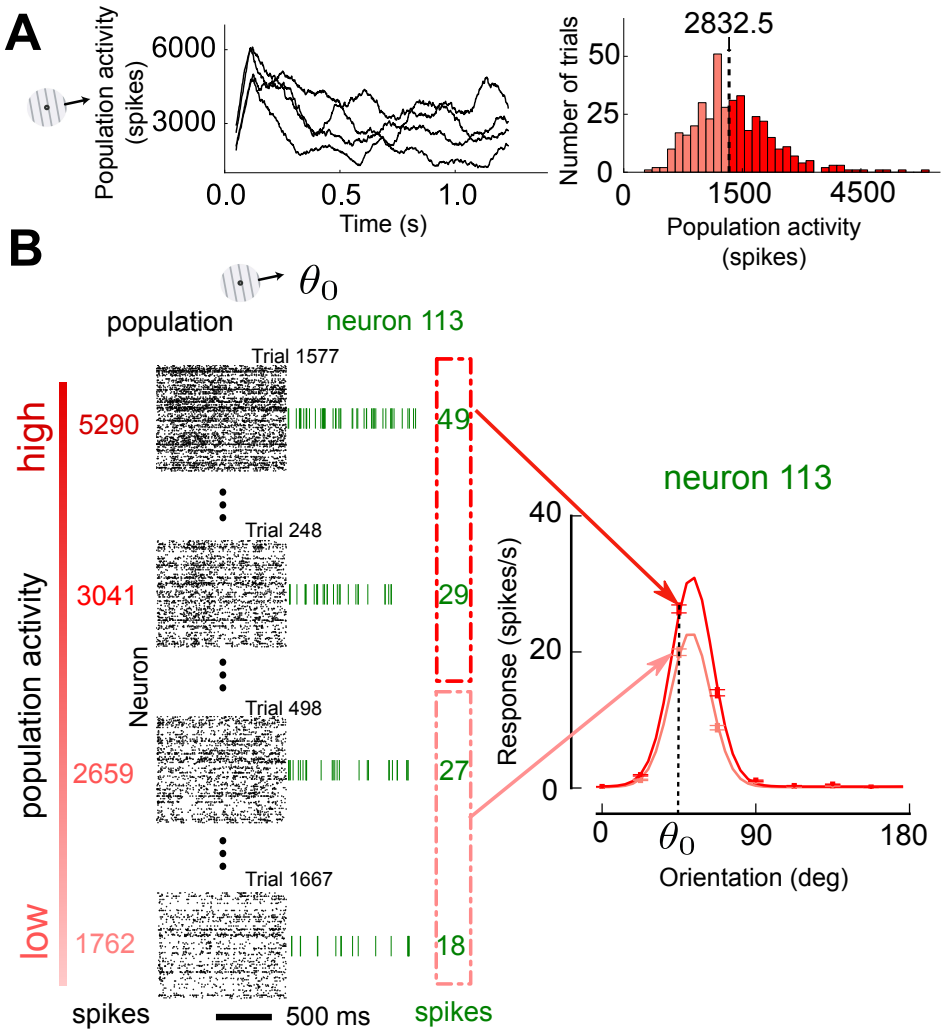
A
Population activity (spikes)
6000
3000
0.0
0.5
1.0
Time (s)
Number of trials
2832.5
50
25
0
0
1500
4500
Population activity (spikes)
B
θ_0
population
neuron 113
Trial 1577
Trial 248
Trial 498
Trial 1667
high
population activity
low
5290
3041
2659
1762
Neuron
49
29
27
18
spikes
500 ms
spikes
neuron 113
Response (spikes/s)
40
20
0
0
θ_0
90
180
Orientation (deg)

stimulus was roughly unimodal and broad (Figure 2.5A, right; similar unimodal distributions were obtained in all datasets, figure 2.6). We characterized the timescale of the fluctuations using the spontaneous activity periods. Fluctuations in population activity were correlated with a timescale of a few hundreds of milliseconds (Figure 2.7), consistent with previous reports on single neuron activity in V1 (Ecker et al., 2014; Kohn and Smith, 2005). Population activity was negatively correlated with LFP signals (see figure 2.8), as previously reported (Okun et al., 2015). We tested how neuronal tuning varies with fluctuations in population activity using a model-free approach, by comparing responses of a single neuron when the activity of the rest of the recorded neurons was either high (defined as the half of trials in which the summed population activity was greatest) or low (remaining trials). We used all recorded units to define periods of high and low population activity, excluding the neuron whose tuning was being characterized to avoid artifacts. Therefore, any observed modulation of tuning must arise from network effects, and would not be observed

Figure 2.5 ***(preceding page)*:** Sensory tuning depends on population activity in an example single neuron. (A) Stimulus-evoked population activity fluctuates across trials for the same stimulus (left panel). Four trials are shown. Distribution of population activity (sum of spikes across all neurons in the recorded population) across trials for a given stimulus (right). (B) The tuning of a single neuron (neuron 113 in D4) is strongly modulated with population activity (population activity is defined here as the sum of spikes across all neurons excluding the activity from the neuron for which tuning is being characterized). Population activity was ranked from high (top left) to low (bottom left) for each stimulus orientation θ_0. The activity of the selected neuron (green spike trains) was averaged across either the top (red box) or bottom (light red box) 50th percentile of trials, and the averages were plotted as a function of stimulus orientation (rightmost panel). The tuning was modulated with population activity (red vs. pink lines), with stronger responses during periods of high population activity. Points and error bars are mean responses and s.e.m., respectively; lines are von Mises fits.

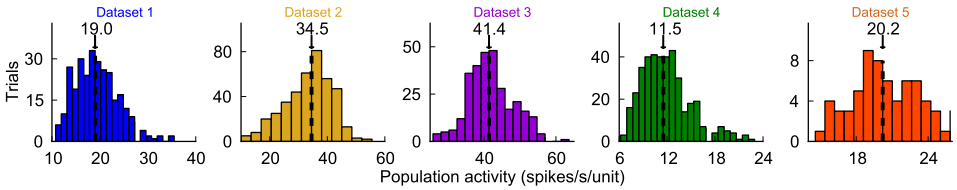

Figure 2.6: Distributions of population activity for all datasets, computed as in figure 2.5A. The distributions tend to be unimodal and wide.

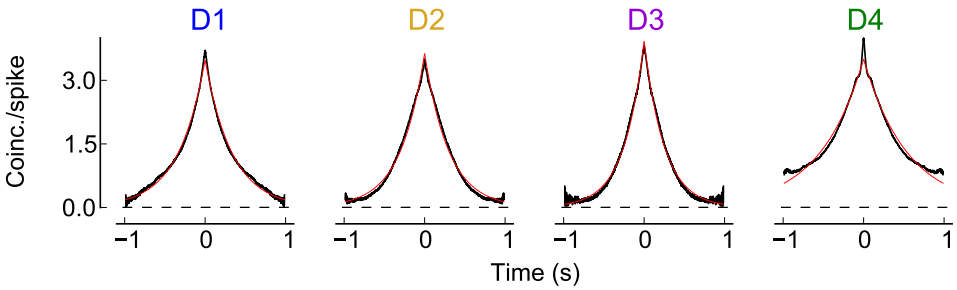

Figure 2.7: The autocorrelation of spontaneous population activity fluctuations is well-approximated by an exponential (red lines) with a time constant of around 300 ms. We computed the autocorrelation (black lines) of population activity during the last second of spontaneous activity before stimulus onset, excluding the 500 ms period after stimulus offset to avoid transients, following the method used by Kohn and Smith (2005). See methods, section 2.2.

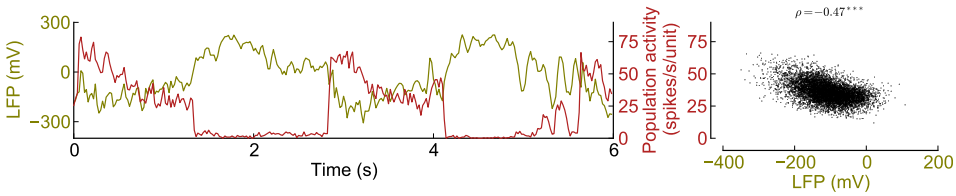

Figure 2.8: Population activity is negatively correlated with the LFP. The left panel shows the mean LFP and population activity computed in time intervals of 20 ms across a few consecutive trials, with shaded area indicating the stimulus presentation periods. In the right panel the correlation between the LFP and the population activity computed in 6 time bins of 200 ms each during evoked activity across all orientations is shown. The observed correlation was negative and significant ($\rho = 0.47$,two-tailed t-test $p < 10^{-20}$).

for uncorrelated neural populations. Both tuning and population activity were measured during the entire duration of the evoked activity period (shorter periods are considered below).

The tuning of an example single neuron depended clearly on population activity (Figure 2.5B): responses were stronger when population activity was high (dark red box, figure 2.5B) compared to when it was low (light red). To characterize how tuning was altered, we first determined whether there was substantial broadening (where tuning width was defined as the distance between peak to half-peak) or displacement of tuning with population activity. To quantify these effects, we fit the tuning of each neuron with a von Mises function (see Methods, section 2.2). Across single neurons, we found a small (2% relative change) widening of tuning when population activity was high compared to low, but this effect was not significant (Figure 2.9A; Mann-Whitney $U = 855$, $p = 0.2$). Tuning preference was also only weakly modulated with population activity (Figure 2.9B; median absolute displacement = 1.0 degrees; permutation test $p < 0.002$), a small shift when compared to the typical tuning width. Therefore, we conclude that changes in tuning width and preference are small, and that the

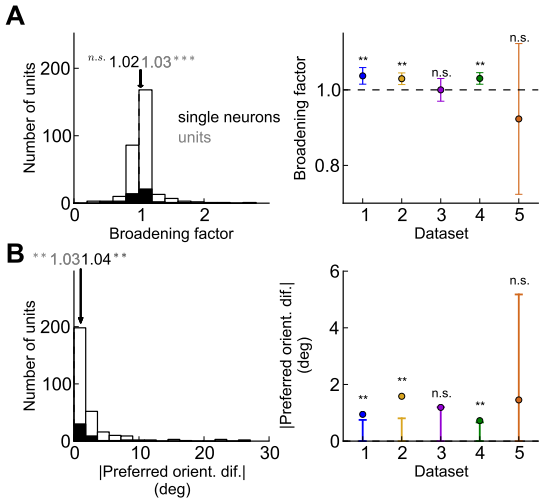

Figure 2.9: Tuning width and preferred orientation are only weakly modulated by fluctuations in population activity. (A) Histogram (left) of the broadening factor (ratio of standard deviations of high and low population activity tuning curves, BF) for all single neurons (black) and all units (white), and median values for each dataset (right), as obtained from von Mises fits to the tuning measured over the entire stimulus period (see Methods, section 2.2). We found a small (2% for single neurons and 3% relative change for all units) widening of tuning, as a function of population activity (left panel; Mann-Whitney $U = 855$, $p = 0.2$, single neurons; Mann-Whitney $U = 310^4$, $p < 10^{-11}$, all units). This effect was present in three datasets out of five (right, all units; median $BF = 1.04$, permutation test $p < 0.002$, D1; median $BF = 1.04$, $p < 0.002$, D2; median $BF = 1.0$, $p = 0.5$, D3; median $BF = 1.03$, $p < 0.002$, D4; median $BF = 0.92$, $p = 0.5$, D5). Error bars correspond to 95% confidence intervals. (B) Same as (A), but for the absolute difference in preferred orientation (DPO) between the high and low activity tuning. Tuning preference was weakly but significantly modulated by population activity (left; median shift of 1 degree for single neurons and units, permutation test $p < 0.001$ in both cases). This effect was significant in 3 of the 5 datasets (right; median $DPO = 0.91$, permutation test $p < 0.001$, D1; median $DPO = 1.6$, $p < 0.001$, D2; median $DPO = 1.2$, $p = 0.05$, D3; median $DPO = 0.75$, $p < 0.001$, D4; median $DPO = 1.4$, $p = 0.3$, D5). Error bars show the 95% confidence intervals of the null hypothesis. *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$, n.s.: $p > 0.05$.

influence of population activity can only involve multiplicative and additive modulation of tuning.

2.3.1 Multiplicative and additive modulation of tuning with population activity

We sought to determine the extent to which tuning was multiplicatively and additively modulated with population activity. In the following analysis (Figure 2.10A-C), both tuning and evoked population activity were measured from 160 to 260 ms after stimulus onset. This brief time period was chosen such that we could, on one hand, analyze the data from awake and anesthetized animals in the same way, and, on the other hand, study the temporal dynamics of the modulatory effects of population activity. The results for other time periods are discussed further below. Tuning varied strongly with population activity (Figure 2.10A, four examples shown). For each single neuron, we characterized its multiplicative and additive modulation with population activity by performing linear regression on the average response to each orientation, when population activity was high compared to when it was low (Figure 2.10B). The slope of the linear fit indicates how tuning scales multiplicatively with population activity (termed hereafter the multiplicative factor (MF)). The intercept of the fit, on the other hand, describes the additive shift to tuning with population activity. To obtain a relative measure of the additive shift, like the multiplicative factor, we defined the additive factor (AF) as the ratio between this intercept and the mean firing rate of the neuron across orientations. Thus, neurons with purely multiplicative factors will feature a fit with slope larger than one that passes through the origin, whereas neurons with purely additive factors will have fits with slope one and a positive intercept. In a separate analysis, we confirmed that estimates of the multiplicative and additive factors from von Mises fits

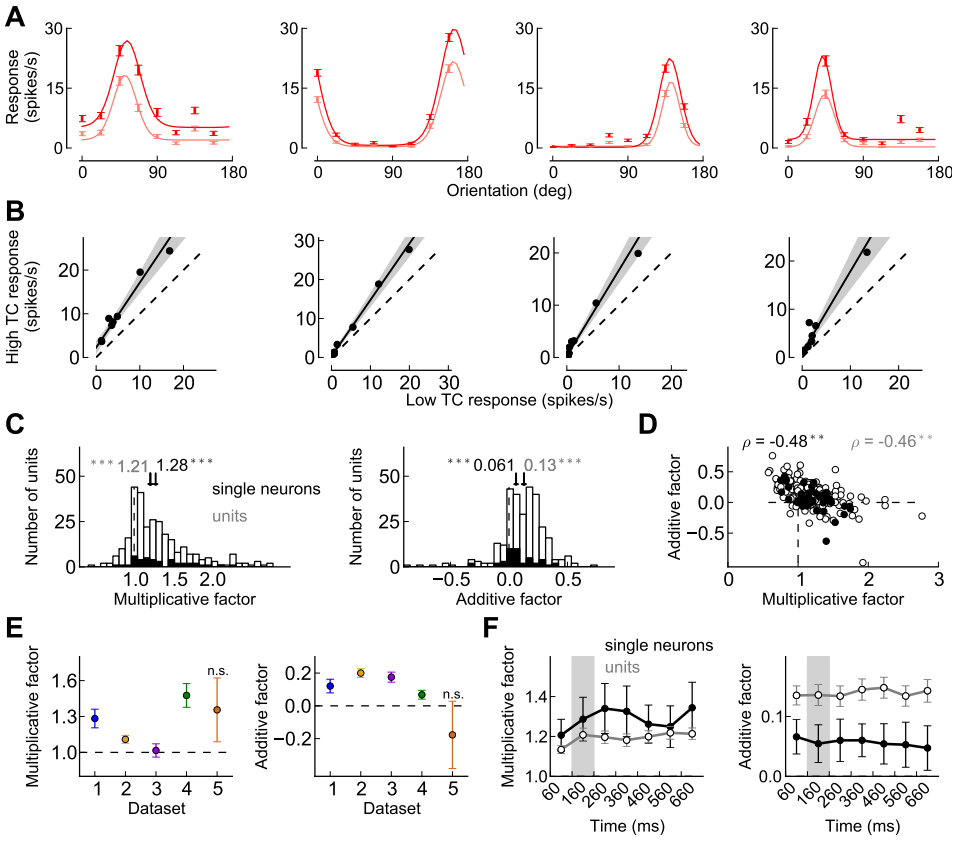
A
Response (spikes/s)
Orientation (deg)
B
High TC response (spikes/s)
Low TC response (spikes/s)
C
Number of units
*** 1.21 1.28 ***
single neurons
units
Multiplicative factor
*** 0.061 0.13 ***
Additive factor
D
ρ = -0.48 **
ρ = -0.46 **
Additive factor
Multiplicative factor
E
Multiplicative factor
n.s.
Dataset
Additive factor
n.s.
Dataset
F
single neurons
units
Multiplicative factor
Time (ms)
Additive factor
Time (ms)

to the tuning gave similar results (not shown).

For the example single neuron of Fig 2B, tuning was modulated multiplicatively ($MF = 1.4$, permutation test $p < 0.002$) with little additive modulation ($AF = 0.009$, $p = 0.01$). The four example neurons of figure 2.10A displayed different levels of multiplicative and additive modulation. For instance, the neuron in the 2nd panel was modulated in a mostly multiplicative manner ($MF = 1.5$, permutation test $p < 0.002$; $AF = 0.046$, $p = 0.04$), and the remaining neurons displayed a combination of multiplicative and additive effects. When we calculated tuning in more finely binned sets of trials, we observed that the modulation varied smoothly with the population

Figure 2.10 ***(preceding page)*:** Sensory tuning undergoes multiplicative and additive modulation as a function of population activity. (A) Modulation of sensory tuning with population activity in four single neurons, computed as in Fig. 2.5B. Error bars indicate s.e.m. (B) Mean response of single neurons when population activity is high (ordinate, corresponding to red lines in (A)) vs. low (abscissa, light red lines in (A)). Each dot is the mean response to a different stimulus orientation. Shaded areas around the lines correspond to 95% confidence intervals. (C) Histograms of multiplicative (left panel) and additive factors (right) for all orientation-selective single neurons (black; N = 45) and orientation-selective units (white; N = 293). Median multiplicative factor across single neurons is 1.28, and across units is 1.21, shown in bold and non-bold formats respectively. Median additive factor across single neurons is 0.061, and across units is 0.13. (D) Additive and multiplicative factors for both single neurons (black circles) and units (open circles; all) are negatively correlated. (E) For individual datasets, multiplicative factors (left panel) are typically significantly larger than one. For individual datasets, additive factors (right) are significantly larger than zero for all except for one dataset. (F) Median multiplicative and additive factors as a function of time (100 ms time windows), relative to stimulus onset (time zero) across single neurons (black line) and units (gray). Shaded areas indicate time window (160-260 ms) used to compute population activity and tuning curves in panels (A-E). Error bars correspond to 95% confidence intervals ($2.91 \cdot m.a.d/\sqrt{N}$, where $m.a.d$ is median absolute deviation). *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$, n.s.: $p > 0.05$.

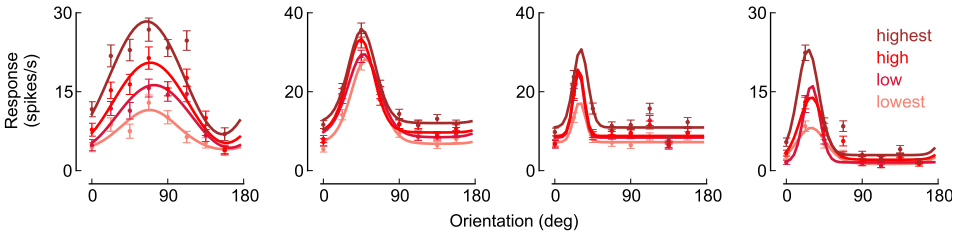

Figure 2.11: Tuning modulation as a function of population activity divided into four bins (25^{th}, 50^{th}, 75^{th} and 100^{th} percentiles; quartiles), ranging from low (lighter lines) to high population activity (darker lines). Population activity was computed in the 160-260 ms period. Error bars correspond to s.e.m and lines are von Mises fits. Same neurons as in Fig. 2.10A.

activity level (Figure 2.11). Statistically-significant multiplicative and additive factors were found in a substantial fraction (26/45 for multiplication, 15/45 for addition) of orientation-selective single neurons. The median multiplicative factor across all single neurons was 1.28, significantly larger than one (Figure 2.10C left, black; Mann-Whitney $U = 215$, $p = 10^{-10}$). This corresponds to a change of 28% in the firing rate, which occurs with a 35% increase in population activity between low-activity and high-activity trials. The median additive factor was also significantly larger than zero (Figure 2.10C right, black; median $AF = 0.06$, Mann-Whitney $U = 387$, $p = 10^{-8}$), indicating a 6% increase relative to the neuron's mean firing rate at low population activity. Importantly, we found that there was a negative correlation between the multiplicative and additive factors across single neurons (Fig. 2.10D, black dots; $\rho = -0.48$, non-parametric bootstrap $p < 0.002$, see Methods).

The results described thus far were based on well-isolated single neurons, but remained qualitatively unchanged if we included activity from multiunits. Across all units, the median multiplicative and additive factors were significantly larger than one and zero, respectively

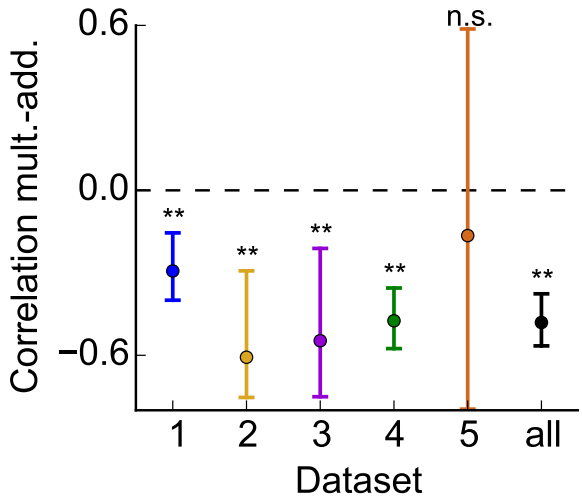

Figure 2.12: Correlation between multiplicative and additive factors dataset by dataset. The correlation is negative for all datasets, and significant in 4 out of 5 datasets. Error bars show 95% confidence intervals computed by a non-parametric bootstrap test. *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$, n.s.: $p > 0.05$.

(Figure 2.12, white; median $MF = 1.21$, Mann-Whitney $U = 210^4$, $p = 10^{-48}$; median $AF = 0.13$, Mann-Whitney $U = 210^5$, $p = 10^{-44}$). The negative correlation between multiplicative and additive factors was also apparent across this large set of units (Figure 2.10D, white dots: $\rho = -0.46$, $p < 0.002$). Although multiplicative and additive factors formed a continuum rather than distinct groupings, the negative correlation between multiplicative and additive factors across both single neurons and all units indicates a partial division of multiplicative and additive modulation with population activity. To test whether the finding of both multiplicative and additive modulation with virtually no broadening was not due to artifacts in our estimation method, we applied the same method to simulated population activity with tuning identical to the one observed in the data (Figure 2.1). We

created neuronal populations with purely multiplicative modulations, purely additive modulations, or purely broadening effects and tested whether our method discovered the true modulation while rejecting other types of modulation. The method reliably estimated the correct type of modulation in each simulated dataset (Figure 2.2). We furthermore confirmed that the negative correlation between multiplicative and additive factors found in our data was not an artifact of our estimation method, as our method could reliably detect or reject the presence of correlations between these factors in simulated data (Figure 2.3).

We also evaluated the significance of modulatory factors in each dataset separately, in part to test whether there are substantial differences between anesthetized and awake preparations. We found strong and significant multiplicative and additive factors in most individual datasets (Figure 2.10E). The median multiplicative factors were significant in four of five datasets, including the awake dataset (Figure 2.10E right; median $MF = 1.28$, permutation test $p < 0.002$, D1; median $MF = 1.10$, $p < 0.002$, D2; median $MF = 1.02$, $p = 0.4$, D3; median $MF = 1.48$, $p < 0.002$, D4; median $MF = 1.36$, $p = 0.04$, D5). Significant positive additive factors were found in all datasets (Figure 2.10E right; median $AF = 0.12$, permutation test $p < 0.002$, D1; median $AF = 0.2$, permutation test $p < 0.002$, D2; median $AF = 0.18$, $p < 0.002$, D3; median $AF = 0.07$, $p < 0.002$, D4), except the awake dataset, for which there was a non-significant negative trend, presumably due to the lower number of neurons and trials recorded when compared to the anesthetized datasets (Figure 2.10E right; median $AF = -0.18$, $p = 0.4$, D5). We also confirmed that the negative correlation between multiplicative and additive factors was present in all anesthetized datasets separately (Figure S3D), indicating that this correlation did not emerge from aggregating data with different mean values.

Substantial multiplicative and additive effects with no broaden-

ing, and a negative correlation between multiplicative and additive factors, were also observed when, instead of using direct measures of population activity, we used the projection of the population activity vector onto the first PCA component on a trial-by-trial basis (Figure 2.18A,B). Thus, our findings are not sensitive to the specific definition of population activity used but generalize to other sensible alternative measures of population activity strength.

Finally, we tested whether the tuning modulation was also present during other response epochs than the window 160-260 ms after stimulus onset, considered above. We repeated our analyses measuring both neuron tuning and population activity in 100 ms windows spanning the range from 60 to 1260ms. We found that the modulation of neuronal tuning with population activity was robust in these other epochs as well (Figure 2.10F).

2.3.2 Modulation of tuning with pre-stimulus population activity

We have thus far considered how tuning changes with fluctuations in evoked population activity. These population fluctuations vary slowly under spontaneous conditions, over a timescale of hundreds of milliseconds (Figure 2.7), and are well-documented (Arieli et al., 1996; Ecker et al., 2014; Fiser et al., 2004; Kenet et al., 2003; Kohn and Smith, 2005; Smith and Kohn, 2008; Tsodyks et al., 1999). Spontaneous activity fluctuations have been shown to influence subsequent evoked responses (Arieli et al., 1996; Tsodyks et al., 1999). Thus, we sought to determine how tuning during stimulus presentation varies with the strength of population activity before stimulus onset. We measured population activity in the 100 ms preceding stimulus onset, and tuning from 60 to 160 ms after stimulus onset. We excluded the data of the awake preparation, as the short inter-stimulus interval (50 ms) made

a reliable estimation of pre-stimulus activity impossible. Orientation tuning depended on the strength of pre-stimulus population activity. We found significant positive multiplicative and additive factors (Figure 2.13A; for single neurons, median $MF = 1.03$, Mann-Whitney $U = 645$, $p = 0.01$; median $AF = 0.093$, $U = 430$, $p = 510^{-6}$; for all units, median $MF = 1.06$, $U = 310^{4}$, $p = 210^{-16}$; median $AF = 0.084$, $U = 210^{5}$, $p = 310^{-42}$). The modulation with pre-stimulus population activity was significantly smaller than that based on fluctuations in stimulus-evoked population activity (MFs: Wilcoxon sign-rank test, $p < 0.01$; AFs: $p < 0.01$). Furthermore, the modulation with pre-stimulus activity was most evident when tuning was measured shortly after stimulus onset (60-160 ms; figure 2.13B). The factors typically declined over time, as one would expect from the spike correlation times of a few hundreds of milliseconds found in our data and usually reported for V1 (Arieli et al., 1996; Ecker et al., 2014; Fiser et al., 2004; Kenet et al., 2003; Kohn and Smith, 2005; Smith and Kohn, 2008; Tsodyks et al., 1999).

2.3.3 Model comparison

In the last years, different models have been proposed to explain the shared variability found in cortical data. (Goris et al., 2014) were able to explain the supra-Poisson variability in visual areas adding a multiplicative factor to the tuning curves of the different neurons. Another work Lin et al. (2015) suggested that both multiplication and addition help explaining the nature of shared variability. Their analysis favors a model with different additive factors for each neuron, but a common multiplicative gain factor. Nevertheless, we found that the tuning curves recorded in the primary visual cortex of macaques show both multiplication and addition, each neuron having different contributions from each factor. Therefore, we wanted to compare the different models to know which one is able to explain better the cortical vari-

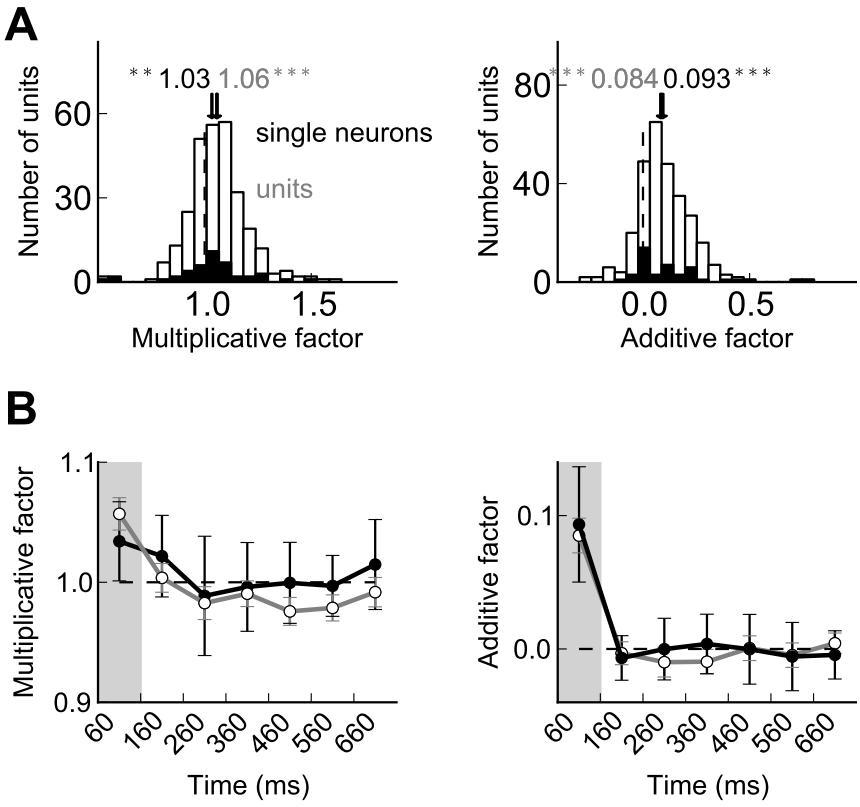

Figure 2.13: Sensory tuning depends on pre-stimulus population activity. (A) Histograms of multiplicative (left) and additive (right) factors across single neurons (black) and units (white). Median multiplicative factor across single neurons was 1.03 and 1.06 across units. Median additive factor across single neurons was 0.093 and 0.084 across units. (B) Median multiplicative and additive factors as a function of time (time windows of 100 ms) for single neurons (black line) and all units (gray). Shaded areas indicate time window (160-260 ms) used to compute statistics in panel (A), while population activity was computed during the pre-stimulus period (100 ms before stimulus onset). Error bars correspond to 95% confidence intervals ($2.91 \cdot m.a.d/\sqrt{N}$, where $m.a.d$ is median absolute deviation). *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$, n.s.: $p > 0.05$.

ability recorded in macaque V1. For each model, the gain-conditioned, Poisson firing rates were as follows:

i Independent model: $f_i(\theta) = h_i(\theta)$.

ii Homogeneous multiplicative model: $f_i(\theta, g) = (1 + g)h_i(\theta)$ (based on Goris et al, 2014).

iii Affine model: $f_i(\theta, g_1, g_2) = g_1 h_i(\theta) + \beta_i g_2$ (following Lin et al, 2015)

iv Multi-gain model: $f_i(\theta, g) = (1 + \alpha_i g)h_i(\theta) + \beta_i g$

The models were trained on 80% of the data by optimizing model parameters to maximize likelihood of this training set. The reported log-likelihoods are for predicting the remaining 20% of the data, averaged over 10 independent splits of the data. This analysis was performed using the same time window of 100 ms (160-260 ms after stimulus onset) as in most of the other figures throughout the paper. As we use cross-validated log-likelihood, it is not necessary to take into account the different number of parameters for each model, as cross-validation already controls for model complexity.

2.3.4 A multi-gain model predicts how tuning modulation affects information encoding

Our analysis revealed that tuning undergoes both multiplicative and additive modulation as a function of population activity, and that across neurons there is a negative correlation between these two types of modulation. To what extent does this tuning modulation influence encoded sensory information? Does information depend on whether the modulation was multiplicative or additive? To address these questions, we considered an idealized model with both multiplicative and

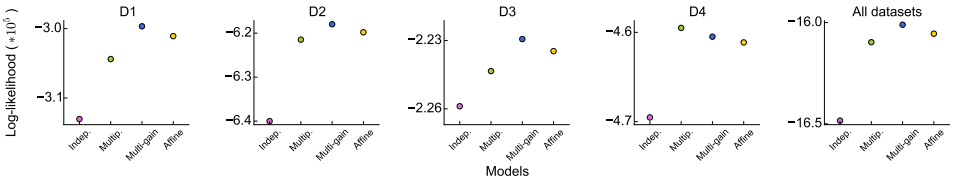

Figure 2.14: A model comparison between independent, homogeneous multiplicative, multi-gain and affine models favors multi-gain model, which allows each neuron to have different multiplicative and additive factors. The cross-validated log-likelihood of the data given the model is plotted individually for each simulated dataset (D1-D4, first 4 columns) and combined across datasets (last column). A higher log-likelihood means that the model better predicted the data of a hold-out set.

additive tuning modulations. We assumed that the mean response of a neuron in the population depends on a global modulatory factor g as:

$$f_i(\theta, g) = g_{m,i}(g)h_i(\theta) + g_{a,i}(g) = (1 + \alpha_i g)h_i(\theta) + \beta_i g \tag{2.3}$$

The first term in the sums corresponds to the multiplicative modulation of tuning, and the second term corresponds to its additive modulation. The normalized tuning function $h_i(\theta)$ describes the tuning of the neuron with respect to the sensory variable θ, which is modulated by the neuron-specific multiplicative and additive factors, $g_{m,i}$ and $g_{a,i}$, respectively. These factors relate to a global modulatory factor, g, linearly by $g_{m,i}(g) = 1 + \alpha_i g$ and $g_{a,i}(g) = \beta_i g$. For instance, a neuron with a purely multiplicative factor corresponds to $\alpha_i > 0$ and $\beta_i = 0$. The global modulatory factor g, assumed to be shared by all neurons in the population, generates correlations between neurons. Firing of each neuron, conditioned on the global modulatory factor, is assumed to be Poisson with the rate dictated by equation 2.3. This multi-gain model, with arbitrary mixtures of multiplicative and additive factors across neurons, is a generalization of recently introduced models with

purely multiplicative modulation of neuronal variance and pair-wise covariance (Goris et al., 2014), or with purely additive modulation to describe state-transitions in neuronal populations (Ecker et al., 2014). Our model generalizes also the recent affine model (Lin et al., 2015), which allows arbitrary additive factors for each neuron but features a multiplicative factor identical for all neurons. In our model, in contrast, each neuron can have a different multiplicative factor α_i (see equation 2.3), as our data suggest (Figure 2.10). Using this more complex model was justified by its ability to better predict neural activity of a hold-out set than alternative models (Figure S5). Most of the models described above have not been used to make predictions about information encoding, and the predictions that have been made were not tested experimentally. The multi-gain model provides specific predictions about how sensory information in neural data should depend on the multiplicative and additive modulation of tuning, which we tested. From our multi-gain model described in equation 2.3 it is straightforward to compute its Fisher information, which is a measure of discriminability between two nearby stimulus orientations (Ma et al., 2006; Seung and Sompolinsky, 1993). Because we are interested in how neurons' information about orientation depends on population activity, we conditioned information on the global modulatory factor g, resulting in:

$$I_i(\theta, g) = \frac{g_{m,i}^2(g) h_i'^2}{g_{m,i}(g) h_i(\theta) + g_{a,i}(g)} \tag{2.4}$$

where the prime denotes a derivative with respect to the stimulus (i.e. $h'(\theta)$ is proportional to the tuning slope). This equation captures the information provided by each neuron if there is no change in the relationship between response magnitude and variability. Consistent with this assumption, we found little difference in Fano factors between trials with low or high population activity (Figure 2.15). We also found similar correlations for the two sets of trials (Figure 2.15).

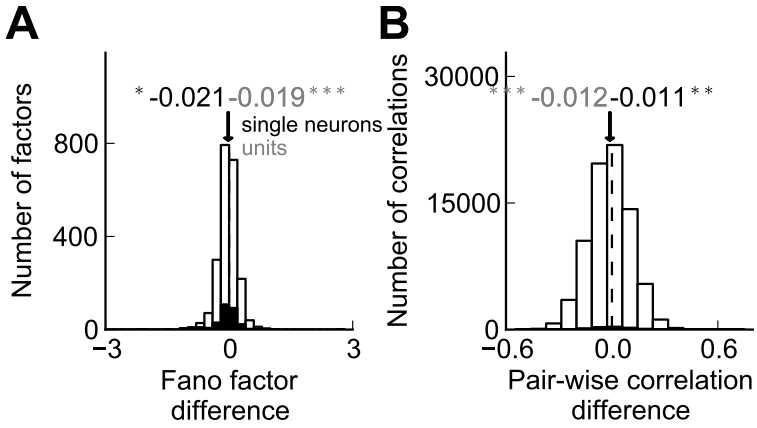

Figure 2.15: Fano factors of the spike counts and pairwise correlations depend weakly on population activity. (A) Median Fano factor across was slightly but significantly smaller when population activity was high compared to when it was low (single units, in black: median $F_n = -0.021$, Mann-Whitney $U = 310^3$, $p = 0.02$, multi-units, in white: median $F_n = -0.020$, Mann-Whitney $U = 210^6$, $p < 10^{-4}$). (B) Same as before for the difference in spike count correlations (single units: median $cor = -0.011$, Mann-Whitney $U = 610^5$, $p = 0.001$; multi-units: median $cor = -0.012$, Mann-Whitney $U = 310^9$, $p < 10^{-4}$).*: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$, n.s.: $p > 0.05$.

Eq. 2.4 predicts that a neuron's information about stimulus orientation increases with multiplicative gain (as their effect is dominated by the numerator), but decreases with additive modulation (as they only appear in the denominator). Intuitively, a multiplicative modulation increases the tuning slope, and thus information grows; in contrast, an additive modulation increases the response variance without altering slope, and thus information decreases. For instance, the neuron in the second panel of 2.10A had a pure multiplicative gain, and therefore its responses to different orientations became more distinct with increasing population activity, potentially increasing the sensory information encoded. In contrast, the neuron in the first panel had also

a large additive modulation, which could result in a drop in the information it encodes (since the response variance will be higher for the stronger responses). Thus, how information is affected by fluctuations in population activity will depend in part on the relative prevalence of multiplicative and additive modulation in the ensemble.

2.3.5 The information encoded by neurons depends on the strength of population activity

We tested the predictions of our model with our data. As an illustration, we first selected an orientation-selective neuron that had a strong multiplicative factor ($MF = 1.8$, permutation test $p < 0.002$; Figure 2.16). The prediction of the multi-gain model is that the information encoded by this neuron about stimulus orientation should increase with population activity. As a proxy for information we used the decoding performance (fraction of correctly predicted stimulus orientation) of a multivariate logistic regression decoder (see Methods, section 2.2.10), cross-validated on hold-out trials that were not used to train the decoder. Better decoding performance corresponds to an increase in sensory information (Moreno-Bote et al., 2014). Although our decoder was trained on all orientations simultaneously, we split the performance into each orientation and obtained a separate decoding performance per orientation, as the non-uniformity of tuning curves caused some orientations to be better encoded than others. In addition, for each stimulus orientation we split the data into trials with either high or low population activity to characterize how population activity modulated information. When performing this analysis, we measured population activity as the summed activity of all recorded units, excluding the unit (or ensemble of units, see below) for which information was computed, just as when we characterized tuning modulation. For the selected unit, decoding performance increased sub-

stantially with population activity, by 44 and 9 percentage points for the two illustrated orientations (Figure 2.16, top panel). This example shows that the sensory information encoded by neurons can vary substantially with the overall population activity.

2.3.6 Sensory information in neuronal ensembles

The multi-gain model predicts that neurons with stronger multiplicative effects should provide better performance with higher population activity than neurons with weaker multiplicative effects. Consistent with this prediction, we found a significant positive correlation across all units between the magnitude of the multiplicative factor and the performance change when moving from low to high population activity (Figure 2.17A, left panel; $\rho = 0.6$, t-test, $p < 10^{-28}$; analysis based on responses measured 160-260 ms after stimulus onset). As also predicted by the multi-gain model, units with stronger additive effects had a larger negative performance change (Figure 2.17A, right panel; $\rho = -0.32$, t-test, $p = 310^{-7}$). In summary, units with strong multiplicative effects provide more information as population activity increases, whereas units with additive modulation provide less information.

The multi-gain model also predicts that the balance of multiplicative and additive factors should determine how the information encoded by small neuronal ensembles, not just by units, should vary with the population activity level. This is true if responses are conditionally independent given the global modulatory factor (equation 2.3) such that information in the ensemble becomes the sum of the "units" contributions (equation 2.4). We tested this prediction by grouping orientation-selective units, including both single neurons and multiunits, in ensembles of size N (N = 1, 2, 3, 5, 10 and 15) as follows. Within each dataset, we ordered orientation-selective units by their multiplicative (or additive) factors, and then split them into

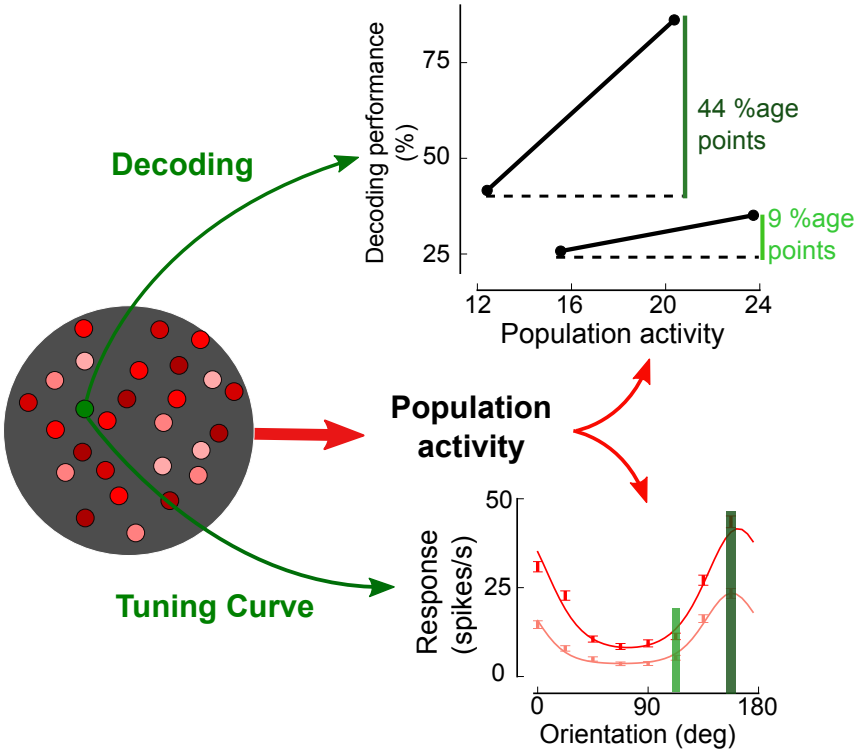

Figure 2.16: Information increases with population activity in a neuron with strong multiplicative modulation. This unit (from D1) had a large multiplicative factor ($MF = 1.8$, permutation test $p < 0.002$; $AF = 0.16$, $p < 0.002$). Decoding performance per orientation (cross-validated probability of correctly predicting the orientation on a trial by trial basis) for the selected unit increases with evoked population activity for two sample orientations (top right panel; performance changes are indicated for the two orientations).

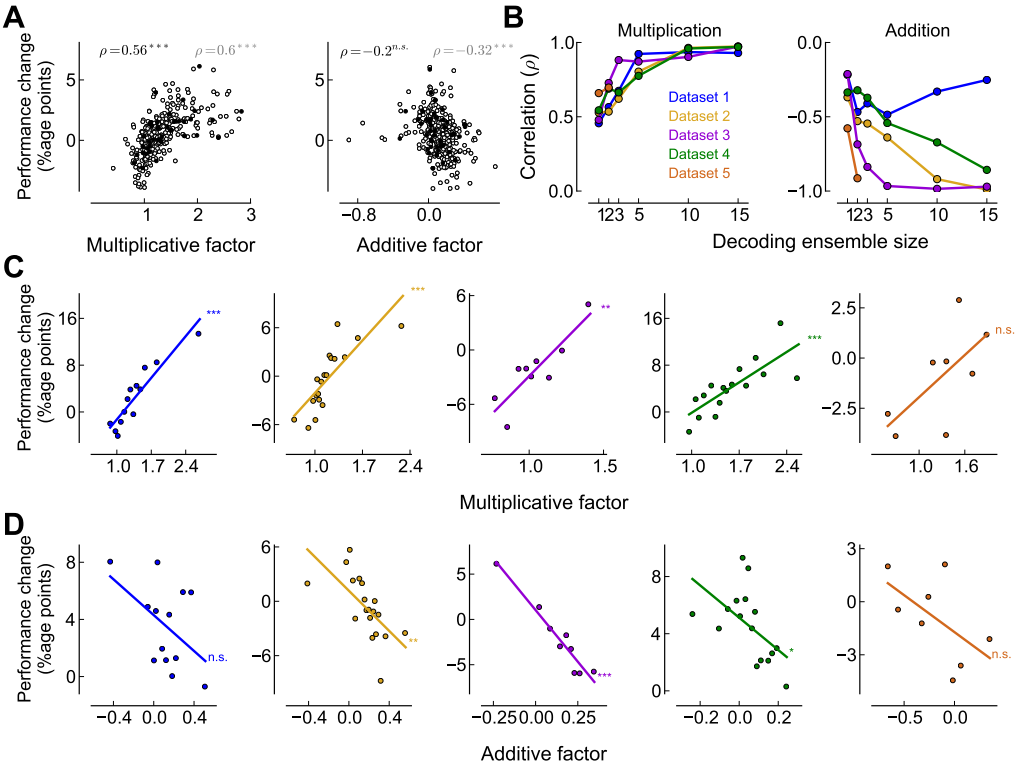
A
B
C
D
Performance change (%age points)
Multiplicative factor
Additive factor
Multiplication
Addition
Correlation (ρ)
Dataset 1
Dataset 2
Dataset 3
Dataset 4
Dataset 5
Decoding ensemble size
n.s.

non-overlapping ensembles of N units that preserved this ordering (Methods). For all datasets, the correlation between performance change and average multiplicative factor of the ensemble was positive and increased rapidly for larger sizes N of the ensemble (Figure 2.17B, left panel). Examples of these correlations are shown in figure 2.17C for ensembles of size N=5 (except D5, where individual units are shown). The datasets from anesthetized animals featured a strong positive correlation between performance change and the average multiplicative factor (Pearson's $\rho = 0.92$, $p = 710^{-6}$, D1; $\rho = 0.80$, $p = 410^{-5}$, D2; $\rho = 0.87$, $p = 0.002$, D3; $\rho = 0.78$, $p = 410^{-4}$, D4). The dataset from an awake animal (Fig. 2.17C, last panel) showed the same trend but did not reach significance (Pearson's $\rho = 0.66$, $p = 0.08$), most likely due to the small number of orientation-selective units available (8 neurons). On average across datasets, the perfor-

Figure 2.17 ***(preceding page)*:** Performance change increases for units and ensembles of units with strong multiplicative modulation, and decreases for those with strong additive modulation when population activity is higher. (A) Performance change as a function of multiplicative (left) and additive (right) factors for all orientation-selective single neurons (black circles) and units (open circles; all). (B) Correlation between performance change and multiplicative (left) or additive factors (right), as a function of the number of units N in the ensemble for each dataset (same color code as in figure 2.10E). Decoding is based on the entire ensemble, and the multiplicative and additive factors refer to the average factors in the ensemble. Population activity was measured after excluding the ensemble used to decode stimulus orientation. The correlation between performance change and average multiplicative factor in the ensemble increases with ensemble size and then asymptotes. In contrast, the correlation between performance change and average additive factor drops with ensemble size, as predicted by the multi-gain model. (C) Performance change increases with the strength of multiplicative modulation for each dataset individually (ensembles of N=5, except N=1 for D5). (D) Performance change decreases with the average additive factor of the ensemble (same sizes as panel C), and can even become negative. *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$, n.s.: $p > 0.05$.

mance change was roughly 10 percentage points for the ensembles of N=5 neurons with strongest multiplicative modulation, from a baseline performance of 44% correct (where chance performance is 12.5% in anesthetized data). In contrast, the correlation between performance change and the average additive factor in the ensemble showed the opposite pattern: performance change was typically more negative for larger ensembles (Figure 2.17B, right). Examples of these negative correlations are shown in figure 2.17D, following the same conventions as in figure 2.17C. In 3 out of 4 anesthetized datasets, the correlation between performance change and additive factors was significantly negative (Pearson's $\rho = -0.48$, $p = 0.09$, D1; $\rho = -0.63$, $p = 0.003$, D2; $\rho = -0.97$, $p = 210^{-5}$, D3; $\rho = -0.54$, $p = 0.03$, D4), while in the awake dataset the correlation was negative but not significant ($\rho = -0.58$, $p = 0.1$, D5). On average across datasets, the performance change was roughly -2% percentage points for the ensembles of N=5 neurons with strongest additive modulation. When, instead of using population activity, we repeated the analysis described above with the projection of the population activity vector onto the first PCA component, we again found that information was differentially modulated in ensembles with strong multiplicative and additive effects (Figure 2.18C-F). We found similar but weaker results when information in the evoked response was conditioned on the strength of population activity measure just before stimulus onset (Figure 2.19), consistent with the modulation of tuning with pre-stimulus activity described in figure 2.13.

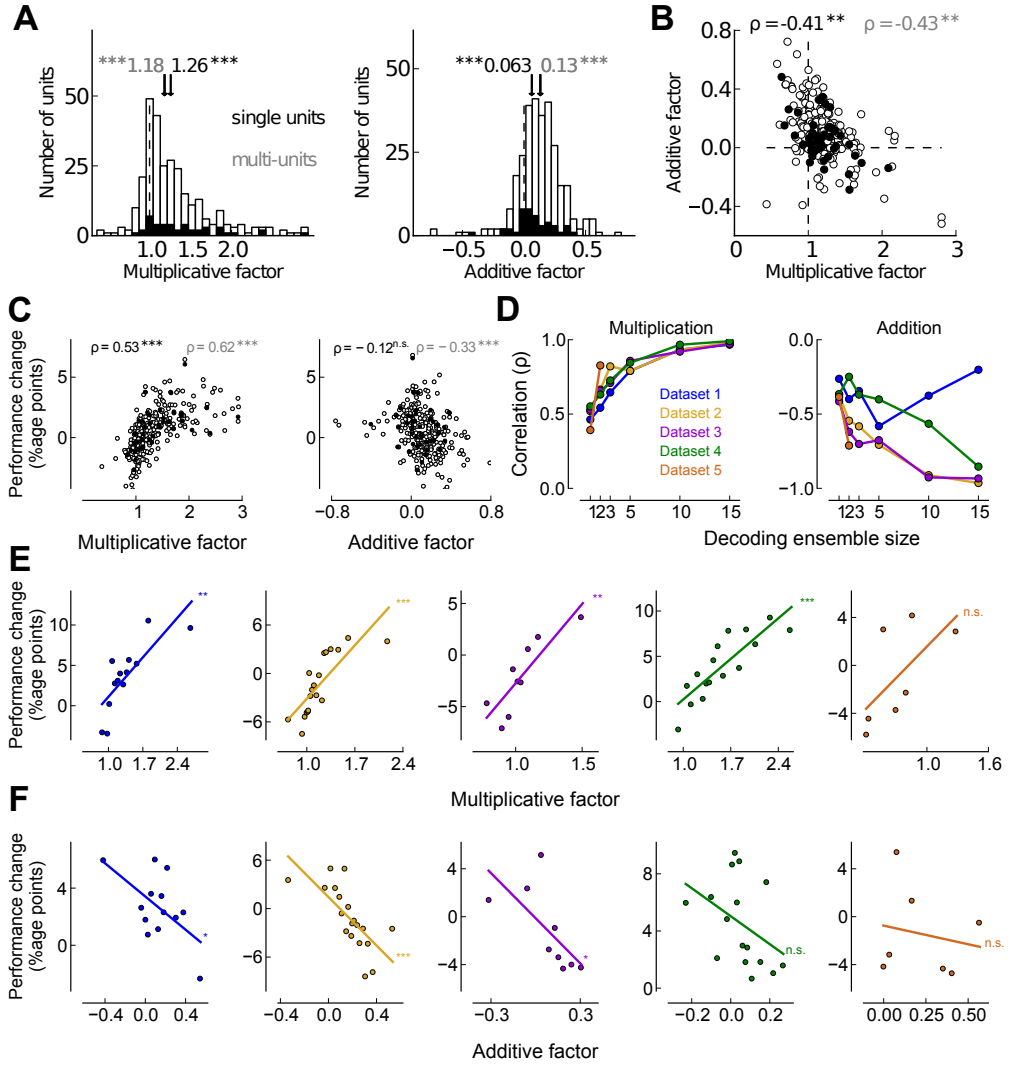
A
1.18 1.26
single units
multi-units
Number of units
Multiplicative factor
0.063 0.13
Additive factor
B
ρ = -0.41**
ρ = -0.43**
Additive factor
Multiplicative factor
C
ρ = 0.53***
ρ = 0.62***
ρ = −0.12n.s.
ρ = −0.33***
Performance change (%age points)
Multiplicative factor
Additive factor
D
Multiplication
Addition
Correlation (ρ)
Dataset 1
Dataset 2
Dataset 3
Dataset 4
Dataset 5
Decoding ensemble size
E
Performance change (%age points)
Multiplicative factor
n.s.
F
Performance change (%age points)
Additive factor
n.s.

2.3.7 Population activity does not substantially change total information, but redirects information into additively and multiplicatively modulated neuronal ensembles

Thus far we have shown that information increases for multiplicatively-modulated ensembles and decreases for additively-modulated ensembles, when population activity is stronger. But what is the net dependence of information on the strength of population activity? To address this question, we randomly selected units to form neuronal en-

Figure 2.18 ***(preceding page)*****:** The projection of the evoked population activity vector onto the first principal component (PC) obtained by PCA modulates tuning curves and information much like population activity does. (A) The projection of population activity vector (vector whose elements are the spike counts of each neuron in the time period 160ms to 260ms) onto the first PC (vector capturing the largest amount of variance of the data, obtained by PCA) modulates V1 neurons both multiplicatively and additively. We applied PCA to a data matrix where each row corresponds to the population activity vector for each trial, and where the mean of each column has been subtracted. Factors were computed as in figure 2.5C while classifying trials as 'high' or 'low' based on the value of the projection of the population activity vector onto the first principal component. Color codes and time bins used are as in figure 2.5C. (B) The multiplicative and additive factors obtained using the first PC projection are negatively and significantly correlated. (C) Performance change from 'high' to 'low' activity (computed from the first PC projection) as a function of multiplicative (left) and additive (right) factors (from first PC projection) for all orientation-selective single neurons (black circles) and units (open circles; all). (D) Correlation between performance change and multiplicative (left) or additive factors (right), as a function of the number of units N in the ensemble for each dataset. (E) Performance change increases with the strength of multiplicative modulation for each dataset individually (ensembles of N=5, except N=1 for D5). (F) Performance change decreases with the average additive factor of the ensemble. *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$, n.s.: $p > 0.05$.

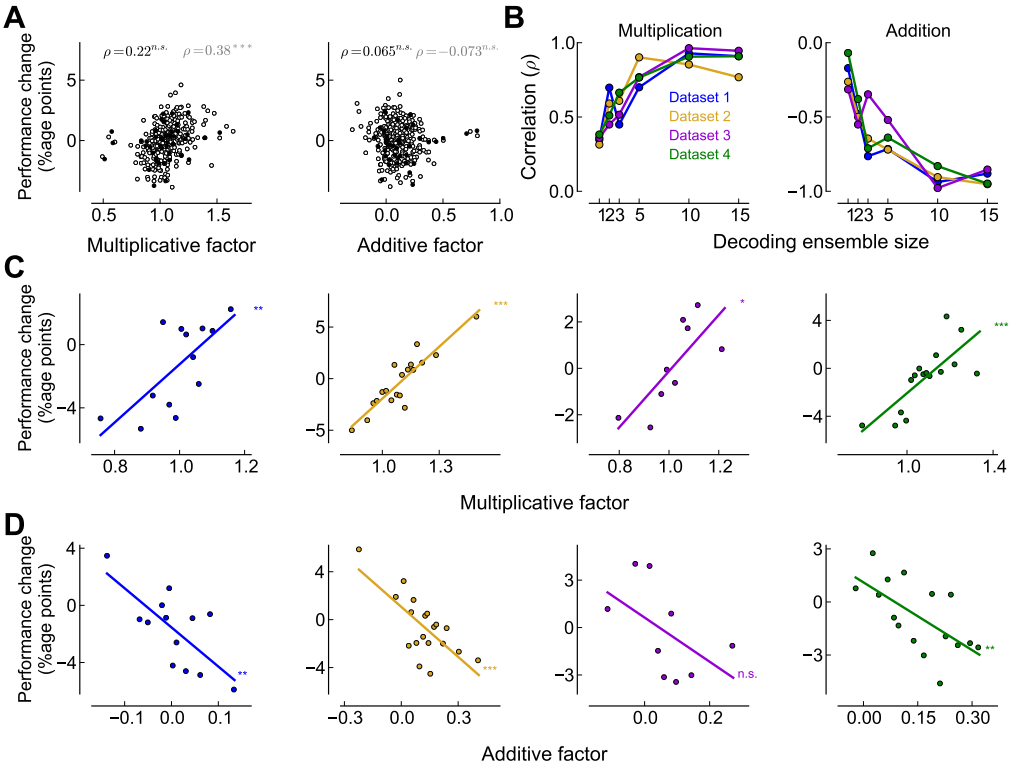

Figure 2.19: Performance increases for ensembles with strong multiplicative modulation and decreases for ensembles with strong additive modulation, when comparing trials with low versus high pre-stimulus population activity, much like for the same analysis using population activity measured during stimulus presentation. Panels are as in Figure 2.17, except that awake data is not shown because of its short inter-stimulus intervals. Pre-stimulus population activity was measured in the 100 ms preceding stimulus onset, and performance changes were computed from 60 to 160 ms after stimulus onset. *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$, n.s.: $p > 0.05$.

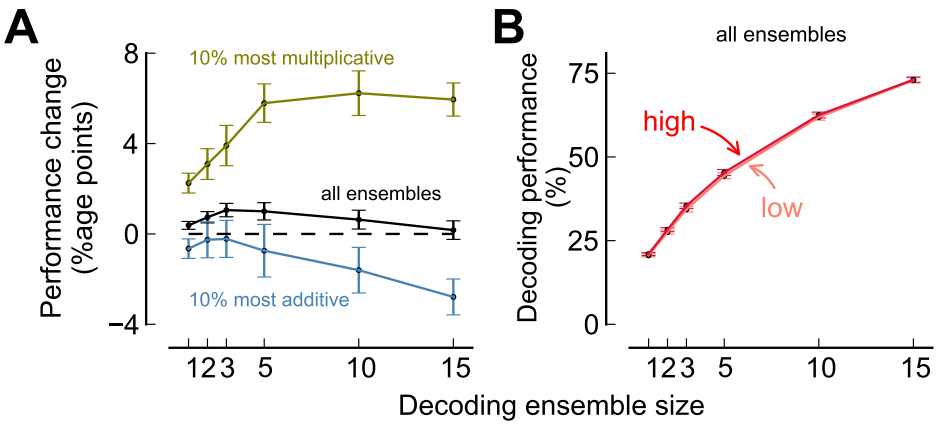

Figure 2.20: Population activity does not substantially change total information, but rather it differentially redirects information into multiplicatively and additively modulated neuronal ensembles. (A) Change in decoding performance averaged across randomly chosen ensembles of varying size (solid line), and for the top 10% ensembles with strongest average multiplicative (green) and additive (blue) factors. (B) Lack of modulation of performance with population activity as a function of ensemble size for randomly chosen ensembles.

sembles of varying sizes (N = 1, 2, 3, 5, 10 and 15), instead of choosing subsets of neurons based on their modulation as we did previously. We computed the performance change between low and high population activity, averaged across many ensembles. We found that there was little change in performance on average (Figure 2.20A, black line). However, when we selected from these randomly-generated ensembles the 10% of cases with the strongest overall multiplicative modulation, we found that performance change was large and saturated as a function of ensemble size (green line), consistent with our previous analysis. Similarly, when we selected the 10% of ensembles with the strongest additive modulation, we found that performance change was consistently negative (blue line). These results suggest that in randomly sampled neuronal ensembles, the effect of population activity on information is negligible. In fact, when we computed decoding performance in these populations at low and high population activity we did not find a visible modulation (Figure 2.20B; the two lines overlay). Therefore, population activity does not substantially modulate the information present in these populations, but rather it modulates which ensembles have more information about the stimulus at different times: when population activity is high, the information encoded by multiplicatively-modulated ensembles is enhanced; when population activity is low, the information provided by additively-modulated ensembles is more important.

2.4 Discussion

We found that intrinsic fluctuations of stimulus-evoked and ongoing population activity are associated with multiplicative and additive modulation of the tuning of orientation-selective neurons in monkey V1. Neurons that showed strong multiplicative modulation tended to display weak additive modulation, and vice versa. These forms

of modulation affected the sensory information encoded by neurons and small neuronal ensembles. As predicted by a multi-gain model, we found that sensory information increased with population activity for neuronal ensembles with strong multiplicative gains. However, sensory information decreased with greater population activity for ensembles with strong additive modulation. Importantly, we found that these effects largely offset each other, so that intrinsic fluctuations of population activity do not strongly affect total sensory information. Rather, the strength of population activity seems to act as a 'traffic light' that differentially redirects information into different subsets of neurons. Previous work (Arieli et al., 1996; Tsodyks et al., 1999) found that pre-stimulus ongoing population activity has an additive effect on evoked responses, and others have reported either additive (Ecker et al., 2014) or multiplicative (Goris et al., 2014) modulations of evoked responses with population activity fluctuations. Recent work has shown that both multiplicative and additive modulations are present in mice and cat neuronal populations (Lin et al., 2015) and reported that a model with a single multiplicative factor across all neurons and different additive factors is favored. Instead, by observing how population activity affects tuning curves, we found that separate multiplicative and additive factors per neuron are required to describe our monkey data, a result further supported by a model comparison analysis (Figure 2.14). Importantly, we found that multiplicative and additive effects are not randomly intermixed across neurons. Rather, neurons with strong modulation of one type tend to show weak modulation of the other. In addition, we found that the strength of spontaneous activity preceding stimulus onset induces not only an additive modulation of stimulus-evoked responses as described in previous work (Arieli et al., 1996), but also a multiplicative effect. However, the influence of pre-stimulus population activity on tuning was weaker than that of stimulus-evoked population activity fluctuations, presumably because activity fluctuations have a timescale of a

few hundreds of milliseconds. Several recent studies have addressed how network state affects sensory responses and encoding, generally defining states based on the degree to which activity is synchronized across neurons or based on LFP measurements (Luczak et al., 2013; Mochol et al., 2015; Pachitariu et al., 2015; Schölvinck et al., 2015). In these works it is often reported that correlations are coupled with population activity measurements (Mochol et al., 2015; Pachitariu et al., 2015; Schölvinck et al., 2015). Much less attention has been paid, however, to the question of how within-state, across-trials fluctuations in the strength of population activity affect neuronal tuning and encoded information, although these fluctuations have been well-documented (Arieli et al., 1996; Kenet et al., 2003; Tsodyks et al., 1999). In our data, fluctuations of population activity strength do not correspond to changes in the degree of network synchronization, because neither variability nor correlations change substantially when going from low to high population activity (Figure 2.15). This may be because fluctuations of population activity in our data correspond to within-state fluctuations, rather than to across-state fluctuations. In fact, our analysis shows that the distributions of population activity are unimodal, suggesting a single state (Figure 2.5A). Overall, by performing an analysis in which population activity was the central quantity to condition on, we were able to reveal that across-trials fluctuations in the strength of population activity affect sensory tuning and the information encoded in distinct subsets of neurons.

Interestingly, our results show that the multiplicative effects on orientation tuning are as large in the awake animal as in the anesthetized preparation. This similarity in modulation occurred despite differences in the magnitude of pairwise correlations between our datasets from awake and anesthetized animals (median pair-wise spike count correlations in 160-260 ms window for anesthetized data: $\rho = 0.073$, D1; $\rho = 0.091$, D2; $\rho = 0.043$, D3; $\rho = 0.061$, D4; and for awake:

$\rho = 0.013$, D5). However, the smaller pairwise correlations observed in the awake preparation nevertheless involved substantial shared fluctuation in the full population, which were clearly evident when we conditioned on the population activity of ~ 20 units. Therefore, although the magnitude of pairwise correlations might vary across experimental preparations (e.g. brain state, cortical areas, layers, etc.) (Cohen and Kohn, 2011; Ecker et al., 2010, 2014; Kohn and Smith, 2005), their net effect on the population can be similar. Indeed, recent work has emphasized that the magnitude of pairwise correlations is not informative about their functional impact: even tiny correlations of a particular form called differential correlations can have massive effects on population information, whereas large correlations with a different structure can have little effect (Moreno-Bote et al., 2014).

A modulation of sensory tuning similar to the one that we report has been observed with optogenetic stimulation of specific V1 neuronal subpopulations. Optogenetic stimulation of layer 6 in mouse primary visual cortex induces divisive (i.e. multiplicative) gain modulation of orientation-selective neurons in the upper layers (Olsen et al., 2012). Similarly, optogenetic stimulation of inhibitory neurons in rat primary visual cortex has been shown to cause divisive or subtractive changes in the tuning of target neurons, depending on the inhibitory subpopulation that is stimulated (Wilson et al., 2012a). More recently, antidromic spikes generated by optogenetic stimulation of distal V1 locations have been shown to additively and divisively modulate layer 2/3 neuronal responses in the mouse (Sato et al., 2014). These effects are similar to those we report, although future work will need to determine whether they provide a mechanistic explanation for the effects we observe under stimulus-driven conditions. An alternative explanation is that additive and multiplicative modulation can arise from balanced excitatory and inhibitory inputs (Chance et al., 2002). Specifically, multiplicative modulation arises from excitatory and in-

hibitory currents that are tightly balanced, whereas additive modulation might involve a slight imbalance in these currents. In this context, our results suggest that the balance of excitation and inhibition varies across neurons.

One might be tempted to equate the modulation of single neuron activity that we observe to that induced by the allocation of attention. Indeed, attention has been shown to modulate tuning in multiplicative and additive manners, similar to the modulation of tuning that what we have observed with population activity (Baruni et al., 2015; McAdams and Maunsell, 1999; Thiele et al., 2009; Treue and Trujillo, 1999). However, attention has been also shown to reduce response variability and pairwise correlations (Cohen and Maunsell, 2009; Mitchell et al., 2009) (but see (Ruff and Cohen, 2014)), whereas we found little change in these measures with fluctuations in population activity (Figure 2.15). Multiplicative modulation of tuning is also evident with manipulations of stimulus contrast (Carandini et al., 1994; Finn et al., 2007; Priebe and Ferster, 2012). It is possible that the multiplicative modulation of tuning we report here shares similar mechanisms to those that occur with manipulations of stimulus contrast. In this regard, it is worth noting that the similarity of the multiplicative modulation we report to variations caused by altering contrast suggests that fluctuations in population activity limit information about stimulus contrast in V1, perhaps explaining limitations on perceptual contrast discriminability. This is because fluctuations that are identical to those generated by stimulus variations are the ones that limit information about the stimulus (Moreno-Bote et al., 2014). The neuron-specific modulation of tuning with population activity fluctuations that we have characterized might govern important aspects of sensory processing, as these fluctuations affect the amount of sensory information that can be read out from small neuronal ensembles. For instance, the modulation might contribute to the trafficking

of information in primary visual cortex, because an increase in overall activity tends to boost information in multiplicatively-modulated neurons while impoverishing it in additively-modulated neurons. Although we have shown that population activity does not substantially change total information in the recorded population, population activity through its neuron-specific multiplicative and additive modulations may act as a global context or 'traffic light' which influences which neuronal ensembles convey more information about the stimulus. In a speculative vein, efficient synaptic plasticity in small neuronal assemblies requires that their responses carry information about relevant internal and external variables (Fusi et al., 2007; Urbanczik and Senn, 2009), so modulating their information about those variables can also gate plasticity. Therefore, population activity might also control important aspects of learning.

Chapter 3

LOOKING FOR DIFFERENTIAL CORRELATIONS

This work is part of an ongoing project developed in collaboration with Ramon Nogueira, Jan Drugowitsch, Adam Kohn and Rubén Moreno-Bote. The experimental datasets analyzed for this project were obtained by Adam Kohn at the Albert Einstein College of New York (USA).

3.1 Introduction

Information is a central concept in nature. Obtaining information from the external world is fundamental to modify the behavior accordingly and better adapt to the environment. However, as central and as present in nature as information we find another important concept: noise. Noise is ubiquitous in nature, and it just changes its shape depending on the scale or the system we are looking at: thermal fluctuations, quantum effects, stochastic processes, uncontrollable or hidden variables, ... The noise has a bad reputation, as it is commonly view as the opposite to information, but it is not always harmful. For example, phenomena like stochastic resonance benefit from the presence of noise and it has been described in many physical and biological systems (Gammaitoni et al., 1998; McDonnell and Abbott, 2009) and it has also been applied to neuroscience and sensory processing (Moss et al., 2004). Both information and noise are extremely important for the nervous system, which has the huge responsability of obtaining information about the environment, planning movements and making decisions. In order to accomplish these goals information about sensory variables, motor actions or decision variables should be encoded in the neural activity. Together with information, noise is also widely present in the nervous system. Neural activity shows large variability, and finding its sources and its role is one of the key challenges for neuroscience, as it would be a major step to increase our knowledge about information coding, learning and memory.

A great effort was put on developing methods to measure information from neural activity and separate the signal from the noise (Salinas and Abbott, 1994; Seung and Sompolinsky, 1993). These new tools allowed to study neural populations with different statistical properties or under different stimulation or experimental conditions, aiming at understanding how information is encoded and decoded in the brain. Most of these works focused on the information content, and

mainly analyzed the way tuning properties (Butts and Goldman, 2006; Pouget et al., 1999; Seriès et al., 2004; Zhang and Sejnowski, 1999), variability and correlations (Nirenberg and Latham, 2003; Nirenberg et al., 2001; Panzeri et al., 1999; Schneidman et al., 2006) could modulate this information. During many years correlations were thought to be a consequence of shared inputs and thus detrimental for coding (Shadlen and Newsome, 1998; Zohary et al., 1994). If correlations reduce information, it seems meaningful to think that the cortex may have mechanisms to decorrelate neural activity Ecker et al. (2010); Renart et al. (2010), depending on the arousal or attentional state for example. However, theoretical work found that information can either increase or decrease information, or leave it unaffeted Averbeck et al. (2006). It is important to note, though, that sensory information always depends on the sensory noise, independently of whether it is measured in the same retina or in a higher cortical area. Taking into account that the amount of neurons in the visual cortex is much larger than in the retina, visual information cannot grow with the size of the neural population in the primary visual cortex and must therefore saturate. Moreno-Bote et al. (2014) suggested that correlations in a specific direction, the direction of the derivative of the tuning curve in the space of neuronal population responses, are those responsible of limiting information -the so-called differential correlations. Pairwise orrelations have been shown to depend on the stimulus (Josić et al., 2009; Ponce-Alvarez et al., 2013), but recent works went beyond this observation and highlight the impact of stimulus dependent correlations for coding in the retina (Franke et al., 2016; Zylberberg et al., 2016a), supporting the idea that correlations in the direction of the derivative of the tuning are detrimental. Differential correlations are stimulus specific by definition, or in other words, the noise that is harmful is the one that looks similar to the variable encoded in the neural population.

Despite their impact on information is huge, the size of differential correlations is very tiny compared to the noise correlations measured in cortical data. As a consequence, it is very difficult to track the evolution of differential correlations, or even confirm their presence in experimental data unless large neuronal populations are recorded. Drawing any conclusion about the possible role they may play from common measures of noise correlations seems unfeasible at this point. Very large populations simultaneously recorded and many trials seemed necessary to obtain a reliable measure of differential correlations. Developing a method to estimate the magnitude of differential correlations with widely available datasets would be very helpful to get some insight into the role that variability in general and differential correlations specially play on information coding.

Although many works focused on the information content, much less attention received the transmission of information (Renart and van Rossum, 2012; Zylberberg et al., 2016b). One of the few works that studied information transmission focused on the connectivity that would allow the propagation of information across layers (Renart and van Rossum, 2012), but not on the noise or the correlation structure. Taking the visual system as an example, it is well known that photoreceptors impact retina, neural activity from the retina projects onto the LGN nucleus in the thalamus, and from here information reaches the primary visual cortex (V1). To better understand how information is encoded in V1 or later stages along the visual pathway it is crucial to understand the information flow, and specially how the noise present in the first layer propagates to the following areas in the hierarchy. Goris et al. (2014) found that the variability increases along the visual pathway, but it is not clear whether this increase is due to the propagation of the sensory noise. Other factors such as global fluctuations or internal noise could also have an impact on variability in higher cortical structures. In fact, global fluctuations have been pre-

dicted to reduce information, but not limit as differential correlations (Kanitscheider et al., 2015b). What is clear from the data-processing inequality is that the information measured in the cortex cannot be larger than the information encoded in the very first layer, the sensory one. Very recently the differential correlations that limit information (Moreno-Bote et al., 2014) have been suggested to be crucial for reliable information transmission, allowing an optimal propagation of information with minimum corruption by extra noise (Zylberberg et al., 2016b).

In addition to limiting information in neural populations and allowwing robust information transmission, differential correlations can also account for realistic psychophysical thresholds and choice probabilities (Kanitscheider et al., 2015b). Therefore, confirming the presence of differential correlations and estimating their size would be an important step to increase our knowledge about population coding and sensory processing. In this work, we present a method to estimate the amount of differential correlations in a neural population. Our method is useful even if the differential correlations are masked with other kind of correlations, such as limited-range exponential correlations, correlations from global fluctuations, or uniform correlations. As in other theoretical methods, the only limitation of this technique is the size of the population and the amount of trials required to obtain good enough results. In this case, however, using realistic generative models, this method offers a good estimation of the introduced differential correlations for populations of just ~ 80 neurons if the amount of trials is large. For less trials the estimation shows a bias, but it can be partially corrected using the bias corrected Fisher information developed by (Kanitscheider et al., 2015a) and substracting to the estimation of the magnitude of differential correlations the same value obtained for the shuffled uncorrelated data. Our next step should be an exploration of the relationship between the amount of trials and the error and the bias in the estimation. These constrains, thus, make

all datasets recorded with the Utah array potentially useful for this method, given that enough trials have been recorded. Finally we applied this procedure to two datasets of monkey V1 data in populations of $\sim$ 80 neurons and despite the existence of the bias expected from the models for lower amount of trials, our estimation is equivalent to a discrimination threshold of around 1.5 degrees. This result, although preliminary, is plausible and encourage us to refine the method and estimate better the accuracy of the values obtained.

3.2 Methods

In this section we will describe, first, the experimental datasets used for this project, and how tuning curves are computed from these experimental recordings. Then, we will explain the generative models built to validate our method, explaining how the tuning curves are computed and the way different kind of correlations are included. Afterwards, the methods selected to compute information are described, together with the different ways utilized to estimate the magnitude of the differential correlations. Finally, we show the measures of uncorrelated data that will be used as a baseline for comparison.

3.2.1 Experimental recordings

The experimental recordings used for this project have been described in the previous chapter (section 2.2). Very briefly, the activity of anesthetized macaque monkeys' V1 neurons were recorded with a Utah array that allowed simultaneous recordings of populations of tens of neurons. Drifting sinusoidal gratings were presented in 8 different orientation randomly selected during 1.28 seconds, and there was an inter-stimulus interval with a black screen of 1.5 seconds. From all the datasets described that used in the previous chapter, we only applied

our method to the two datasets that exhibited the larger amount of trials per condition (400) and the larger populations ($\sim$ 80 neurons).

3.2.2 Tuning curves

We employed von Mises equation to characterize tuning curves (Ecker et al., 2011; Graf et al., 2011):

$$f(\theta) = a + b \exp\left[k(\cos\left(2(\theta - \theta_{pref}\right) - 1\right] \tag{3.1}$$

In the case of experimental recordings, we obtained the parameters fitting the von Mises function to the neuronal activity by minimizing the weighted squared error with bounded parameters to ensure physiologically plausible tuning curves (*minimize* function from *lmfit* python package with the following constraints: $\theta \in [0, \pi]$, $k > 0.0001$, $max(f_\theta) < 1.3 \cdot max(\overline{r_\theta})$, $min(f_\theta) > 0.7 \cdot min(\overline{r_\theta})$, where $(\overline{r_\theta})$ is the mean response across trials for each orientation). We only employed neurons with a good fit as in Graf et al. (2011) ($r^2 > 0.75$). The time to compute the tuning curves was one second after the first 60 ms from stimulus onset, that were not taken into account due to the latency of V1 neurons.

The tuning curves used in the models were randomly selected from the experimental ones. We selected randomly a group of three parameteres that characterize the tuning curve (a, b and k), and then assigned different orientations drawn from a uniform distribution between 0 and 180 degrees.

3.2.3 Generative models

3.2.3.1 Neuronal activity model with differential correlations

We sought to build a model of neuronal activity that exhibited similar properties as the recorded V1 data. This model would allow us to validate the method developed to estimate the magnitude of differential correlations and at the same time it should be useful to get some insight into the role of shared variability for information processing. Differential correlations can be introduced by adding sensory noise. We can write the activity of neuron i in one trial as the value of the tuning curve ($f(\theta)$) at a value slightly different to the real stimulus ($\theta+\delta\theta$). The deviation from the real stimulus (θ) can be interpreted as sensory noise or irreducible uncertainty given the system under study. This noise generates differential correlations and therefore limits information (Moreno-Bote et al., 2014).

$$\lambda_i(\theta) = \vec{f}_i(\theta + \delta\theta) + M\vec{z} \tag{3.2}$$

The activity of the unit i is composed by the sum of the tuning curve plus a second term that can include additional variability and also pair-wise correlations, depending on the matrix M. The tuning curves were extracted from experimental data, but using heterogeneous artificial tuning curves does not qualitatively change the results. The matrix M will introduce some variability and it will control the kind of correlations present in the model. $\vec{z}$ is a vector drawn from the standard normal distribution with the size of the amount of neurons in the population. If M is diagonal, then the second term just adds some noise to the model, while non-diagonal elements add pair-wise correlations to the model. M can be understood as the cholesky decomposition of the covariance matrix: $\Sigma = MM^T$.

We generated different correlation matrices that showed a variety of properties (see below). Once we had the correlation matrix, we obtained the covariance matrix for each stimulus multiplying each value ρ_{ij} by the square root of the geometric mean activity of each pair for that stimulus: $\Sigma_{ij} = \rho_{ij}\sqrt{f_i(\theta)f_j(\theta)}$. If this were the only source of variability we would generate a model with Poisson-like firing, obtaining a fano factor around one. However, we included other steps that add noise into the model. If the variability generated in the model was larger than the variability in the experimental data we used just a fraction of the mean activity as normalizing constant to get the covariance matrix.

$$\Sigma_{ij}(\theta) = \rho_{ij}\sqrt{af_i(\theta)}\sqrt{af_j(\theta)} \tag{3.3}$$

$\rho_{ij}(\theta)$ is the pair-wise correlation between neurons i and j for orientation θ, $f_i(\theta)$ is the tuning curve of neuron i, and a controls the amount of variability included in the model through this term. If a is one, then Poisson-like variability is introduced. Note that if we reduce the variability with $a < 1$, then the correlations generated by this term are also reduced.

In order to compare the effect on information of differential and non-differential correlations we built models with different correlation structures.

- Exponential correlations: we generated the correlation matrix with limited range exponential correlations as in previous works (Ecker et al., 2011; Sompolinsky et al., 2001).

  $$\rho_{ij} = (1 - \delta_{ij})\rho_0 \exp\left[(\theta_i - \theta_j)/\Delta\theta\right] + \delta_{ij} \tag{3.4}$$

  where δ_{ij} is the Kroencker delta (1 if $i = j$ and 0 otherwise). θ_i is the preferred orientation of neuron i, $\Delta\theta$ controls the correlation

length and the parameter ρ_0 modulates the size of the correlations. In the results presented here we used values of $\Delta\theta = 1$ and $\rho_0 = 0.2$ to recover similar correlations and the experimental V1 data.

- Uniform correlations: we used a model with uniform correlations, having all pairs the same amount of noise correlations. The covariance matrices obtained from this structure, however, do depend on the stimulus, as the uniform correlations are normalized by the mean activity of each neuron for each stimulus.

$$\rho_{ij}(\theta) = c_0 \rightarrow \Sigma_{ij}(\theta) = c_0\sqrt{f_i(\theta)f_j(\theta)}, i \neq j \tag{3.5}$$

When we used the tuning curves extracted from the real data we took the variance of the data instead of the mean activity to obtain the covariance matrix from the correlation matrix. As before, we controlled the level of variability introduced with a factor that allows to introduce variability lower than that of a Poisson process, as this is not the only source of variability in this model.

$$\Sigma_{ij}(\theta) = \rho_{ij}\sqrt{a\sigma_i^2(\theta)}\sqrt{a\sigma_j^2(\theta)} \tag{3.6}$$

In order to model the deviation from the real stimulus, we computed a Taylor expansion of the tuning curve around the real value (θ) : $f(\theta + \delta\theta) \sim f(\theta) + \delta\theta f'(\theta) + \frac{1}{2}\delta\theta^2 f'^2(\theta) + ...$ Neglecting terms of second order and above the model of neuronal activity becomes

$$\lambda_i(\theta) = \vec{f}_i(\theta) + \delta\theta\vec{f}_i'(\theta) + M\vec{z} \tag{3.7}$$

At this point we applied a rectifying nonlinearity to avoid negative values. Finally, we introduced a Poisson step to $\lambda_i(\theta)$.

3.2.3.2 Neuronal activity model with differential correlations and global modulations

Global fluctuations have been shown to explain a large part of the cortical variability. In order to model the neuronal activity we decided to include global gains to the model previously described. We tried different versions of the model with global gains, with contributions from heterogeneous or homogeneous multiplicative and/or additive modulations.

$$\lambda_i^{gain}(\theta) = (1 + \alpha_i g)\lambda_i(\theta) + \beta_i g \tag{3.8}$$

In this case, $\lambda_i(\theta)$ is the activity from the model with differential correlations explained in equation 3.7. g is the global gain which is different for each trial, and is obtained from a gamma distribution with mean equals one ($\mu_g = 1$) and variance given by σ_g^2. α_i and β_i allow for heterogeneous multiplication and addition, respectively. Homogeneous multiplication as in Goris et al. (2014) and Lin et al. (2015) can be obtained applying the same $\alpha_i = 1$ to all neurons. The values of α_i and β_i are drawn from a gamma distribution with mean μ_α and μ_β, and variance σ_α^2 and σ_β^2.

The covariance matrix obtained in this model can be computed analitically, as a function of the covariance matrix in the absence of global modulations, the variance of global gain (σ_g^2) and the parameters controlling the levels of multiplicative and additive contributions (μ_α and μ_β):

$$\Sigma_{ij}^g(\theta) = (1 + \mu_\alpha^2\sigma_g^2)\Sigma_{ij} + (2\mu_\beta + \mu_\alpha)\mu_\alpha\sigma_g^2 f(\vec{\theta}) f(\vec{\theta})^T + \mu_\beta^2\sigma_g^2 + diag(\vec{f}) \tag{3.9}$$

3.2.4 Estimating information in neural populations

The linear Fisher information is one of the most used quantities to measure the amount of information present in a neural population (Moreno-Bote et al., 2014; Seriès et al., 2004; Seung and Sompolinsky, 1993), and it is related to the inverse of the variance in the error of the optimal decoder through the Cramer-Rao Bound (Dayan et al., 2001).

$$I(\theta) = f'\vec{(}\theta)\Sigma^{-1}f'\vec{(}\theta) \tag{3.10}$$

However, this quantity is prone to underestimate the true information, and it suffers from the sampling bias problem (Panzeri et al., 2007). In order to overcome this limitation, Kanitscheider et al. (2015a) developed a bias corrected version of the Fisher information for fine-discrimination tasks:

$$I_{BC}(\theta) = I(\theta)\frac{2M - N - 3}{2M - 2} - \frac{2N}{Md\theta^2} \tag{3.11}$$

The bias correction term, thus, depends on the difference among the two stimuli ($d\theta$), the amount of trials (M) and the size of the population (N). Although our data and our models imply
coarse-discrimination tasks, with $\Delta\theta = 22.5$ degrees, we also tested whether this correction helped to measure information more accurately in this case.

3.2.5 Estimation of the magnitude of differential correlations (ϵ)

From the work by Moreno-Bote et al. (2014) we know that the information content in a population (I) depends on the magnitude of differential correlations (ϵ) and the information that would be present

in that population in the absence of correlations (I_0):

$$I = \frac{I_0}{1 + \epsilon I_0} \tag{3.12}$$

The information that a neural population without differential correlation contains should grow with the population size, due to the absence of information-limiting noise. For large N this information can be assumed to be proportional to the size of that population: $I_0 \propto N$, and therefore we can write $I_0 = \alpha N$. Introducing this assumption into equation 3.12 we obtain:

$$I = \frac{\alpha N}{1 + \epsilon \alpha N} \tag{3.13}$$

We performed a nonlinear fit to this equation to compute the α and the ϵ using the kmpfit function of the Kapteyn package (Terlouw and Vogelaar, 2015).

If we invert 3.13 and taking the limit when $N \to \infty$:

$$\frac{1}{I} = \epsilon + \frac{1}{\alpha N} \Rightarrow \ell_{N\to\infty} \frac{1}{I} \to \epsilon \tag{3.14}$$

Obviously, we do not have access to infinitely large populations, but with the only assumption that the information of the population without differential correlations grows linearly with N we can compute the linear fit of equation 3.14 ($1/I$ vs $1/N$) and interpret the intercept as the value of ϵ, and the slope as the proportionality constant between I_0 and N (α). We estimated ϵ using again the kmpfit function of the Kapteyn package (Terlouw and Vogelaar, 2015) in python, applying different weights to each point according to the error in the estimation of information.

A common concern when estimating information is the possibility of having a bias in the method, and we are not free from this problem

when estimating ϵ. To avoid a possible bias with this method we used the same procedure to the shuffled version of the data, and computed the ϵ as the difference among the value computed using the correlated data and the shuffled one: $\epsilon = \epsilon^{raw} - \epsilon^{shuffle}$. To obtain the shuffled, uncorrelated, neuronal activity, for each orientation and for each neuron, we randomly reassigned the neuronal responses to different trials. This reassignment of trials was different for each neuron, and therefore the activity of the neurons in each new trial was not recorded simultaneously. In this way, we destroyed any correlation structure present in the activity, but maintaining other statistical properties present in the data.

3.2.5.1 Estimating ϵ from the covariance matrix

According to the work by Moreno-Bote et al. (2014), the covariance matrix of a population response can be decomposed in two terms, with and without differential correlations:

$$\Sigma = \Sigma_0 + \epsilon f'(\theta) f'^T(\theta) \tag{3.15}$$

Σ_0 would be the covariance matrix in the absence of differential correlations, $f'(\theta)$ is the derivative of the tuning curve, and ϵ measures the size of the differential correlations. As the differential correlations are, by definition, the only kind of correlations that limit information in large networks, one straightforward approach to measure them would be to substract the largest possible $\epsilon f'(\theta) f'^T(\theta)$ term from the empirical covariance matrix that still maintains a semi-definite positive matrix. Nevertheless, this simple measure does not reflect the sensory noise present in the system. Differential correlations can arise due to sensory noise, suboptimal computations, or because of the finite size of the population. The influence of the size of the population on the magnitude of differential correlations can be explained from

equation 3.12, that can be rewritten as:

$$\frac{1}{I} = \epsilon + \frac{1}{I_0} = \epsilon + \frac{1}{\alpha N} = \epsilon + \epsilon_0 \tag{3.16}$$

where $\epsilon_0 = 1/I_0 \propto 1/N$ is the fraction of epsilon that is due to a population of size N. Taking the limit when N goes to ∞ we actually reach asymptotically $\epsilon \sim 1/I$. One option to compute the information present in a population is to remove the largest possible $\epsilon f'(\theta) f'^T(\theta)$ factor from the covariance matrix that still maintains the matrix properly defined as semi-definite positive. The inverse of this estimated ϵ can be interpreted as the information content in the population, but this does not reflect only differential correlations, as explained above (it would if $N \to \infty$). We found that this technique gives a very similar measure to the linear Fisher information explained above, and it can be employed to extrapolate to large N and obtain an estimate of the differential correlations in the same way (not shown).

3.2.6 Baseline for comparison: uncorrelated data

We know that measuring information in neural populations is prone to bias (refsss panzeri quiroga). Therefore we wanted to compare the estimation of the magnitude of differential correlations to some baseline that we could trust. We employed two different measures that helped us to interpret the results and asses the reliability of our estimation methods. One option to obtain uncorrelated data is to shuffle the trials randomly and in a different way for each neuron. This procedure destroys the correlations among neurons but keeps the other statistical properties of the neural population and it is a standard way to compare the effect of noise correlations on information. When the size of the population is small, however, this method might have problems to erase all the correlations. Shuffling the trials randomly, the goal is to achieve new shuffled trials where the activity of the each neuron

has been recorded in different trials, and therefore there are no correlations. By chance, some neurons might be assigned the same real trial, and therefore show correlations in the shuffled trial. To avoid these coincidences and its subsequent correlations we performed the shuffling randomly, but only accepted those configurations that did not show correlations.

We also compared our measures of information and the quantities derived from them with the information computed in a true independent population, summing up the information of individual neurons: $I_{indep} = \sum_i I_i$. Theoretically both the information in the shuffled data ($I_{shuffle}$) and the information in the independent population (I_{indep}) should be similar.

3.3 Results

The goal of this project is to develop a method to quantify the information-limiting correlations described theoretically by Moreno-Bote et al. (2014) that could be applied to experimental data and be useful to measure the accuracy of the perceptual system and gain knowledge about the transmission of information and the computations performed by cortical neurons. The covariance matrix of a neural population can be decomposed in two terms, the information-limiting one and the remaining one that does not limit information:

$$\Sigma = \Sigma_0 + \epsilon f'(\theta) f'^T(\theta) \tag{3.17}$$

The part proportional to $f' f'^T$ limits information, and the magnitude of these differential correlations is controlled by the parameter ϵ. The other part (Σ_0) does not limit information. From this decomposition of the covariance matrix, Moreno-Bote et al. (2014) derived an expression

for the information of the population as a function of the ϵ and the Fisher information in the absence of differential correlations (I_0):

$$I = \frac{I_0}{1 + \epsilon I_0} \tag{3.18}$$

The information without differential correlations, by definition, cannot limit information, and therefore, it should grow with the size of the population. At this point, we assumed that the neurons convey a similar amount of information, and thus take I_0 to be proportional to the amount of neurons in the population: $I_0 = \alpha N$. Then we could obtain the value of epsilon computing the Fisher information for different population sizes and computing a nonlinear fit of:

$$I = \frac{\alpha N}{1 + \epsilon \alpha N} \tag{3.19}$$

Inverting both sides of 3.13 we can see that there is linear relationship between the inverse of information and the inverse of the population size. Therefore, a linear fit of these quantities give us the value of the ϵ that we were looking for in another way, easier in the computational sense. In order to improve the estimation we also tried with a linear fit of $1/I$ vs $1/N$, whose intercept equals ϵ:

$$\frac{1}{I} = \epsilon + \frac{1}{\alpha N} \tag{3.20}$$

As expected by the limiting nature of differential correlations, which are responsible of the saturation of information for large population sizes, in the limit of $N \to \infty$ the inverse of information goes assymptotically to ϵ (see Methods).

3.3.1 Validation of the method with generative models

In order to validate these methods we built realistic generative models that included differential correlations but also other non-differential correlations such as limited-range exponential correlations, uniform correlations or global fluctuations (see Methods in section 3.2). In figure 3.1 we used neurons with heterogeneous tuning curves (bottom-left panel) to build a population with exponential correlations that decay with the distance between the preferred directions of the neurons, as widely reported in cortical neurons (bottom-middle panel). In addition to these limited-range correlation we also included differential correlations with $\epsilon = 10^{-3}$, which correspond to a discrimination threshold of 2 degrees. The variability of the neurons in this model is supra-Poisson (panel bottom-right), with similar mean value as those measured in the experimental V1 data and as previously reported (Goris et al., 2014). In the top panel of figure 3.1 we show how information grows with the size of the population (N) in this model that included both differential and exponential limited-range correlations. This figure shows a clear difference in the scaling of information between the correlated population that includes differential correlations (dark green) and the uncorrelated populations (light green). The non-linear fit in each case is represented by the dashed (correlated and shuffled population) and dash-dotted (independent population) lines. In the extrapolation to very large N we can see that the saturation point is a bit below the theoretical one. The middle row shows the inverse of information vs the inverse of the population size, and the linear fit performed following equation 3.14. In this we can interpret that the intercept, the point were the lines cut the y-axis, is reflecting the ϵ, that is, the inverse of the information for infinitely large networks. In this plot, the estimation of this method is a bit larger than the differential correlations introduced in the model. This overesti-

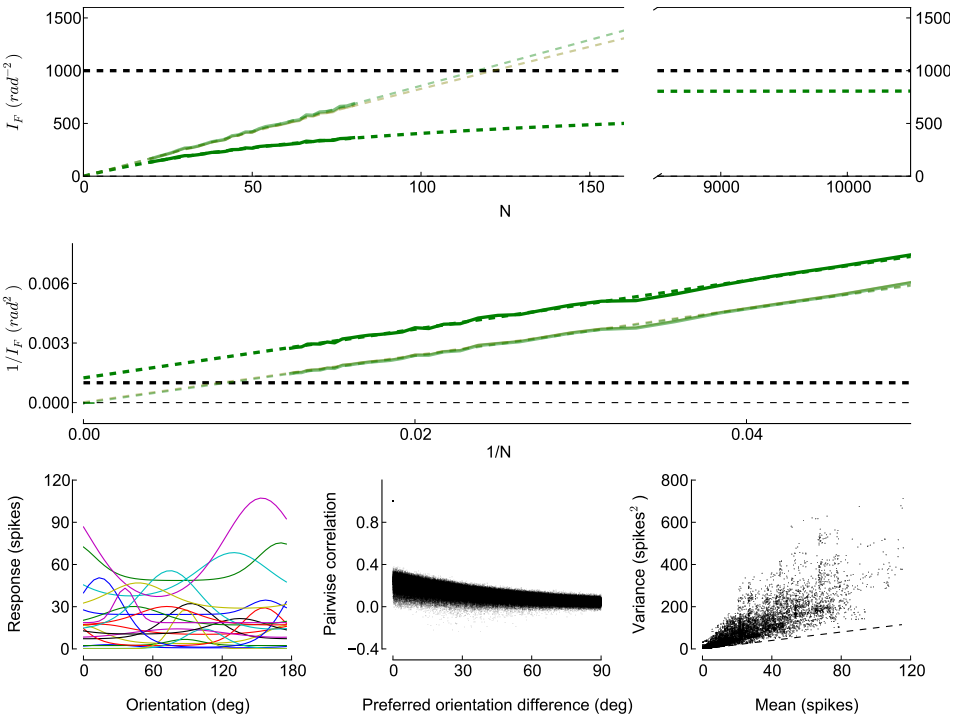
1500
1000
500
0
I_F (rad^{-2})
0
50
100
150
9000
10000
N
0.006
0.003
0.000
$1/I_F$ (rad^2)
0.00
0.02
0.04
1/N
120
90
60
30
0
Response (spikes)
0
60
120
180
Orientation (deg)
0.8
0.4
0.0
−0.4
Pairwise correlation
0
30
60
90
Preferred orientation difference (deg)
800
600
400
200
0
Variance (spikes^2)
0
40
80
120
Mean (spikes)

mation is likely to be due to the exponential correlations, that might also have a component in the $f'f'^T$. If instead of exponential correlations, we mask the same amount of differential correlations with uniform correlations the estimation of ϵ diminishes, going below the true amount of differential correlations introduced in the model (see figure 3.2). Therefore, in this case the method underestimates the amount of information-limiting correlations. Taking into account that the estimation of our method was not completely accurate in the uniform case, we decided to perform the very same analysis using the method developed by Kanitscheider et al. (2015a) to correct for the bias in the estimation of the Fisher information (I_{BC}). As the bias should depend on the number of trials, in figure 3.3 we compared the performance of our method in two identical models with exponential and differential correlations that only differed in the amount of trials per condition: large amount of trials (3000, as before) in the left column, and the amount of trials per condition we had in the experimental recordings (400) in the right column. The overlap of the lines for the Fisher information (I, green) and the bias corrected Fisher

Figure 3.1 ***(preceding page)*:** Estimation of differential correlations in surrogate data with exponential and differential correlations. The top panel shows how the Fisher information I_F scales with N for the model (solid dark line), for the trial-shuffled version of the model (solid light green line) and for the independent population (solid olive line). The dashed colored lines show the extrapolation with the nonlinear fit of equation 3.13 for its case. The middle panel follows the same color code as the top panel, but the fits are linear according to equation 3.14. The bars in the y-axis show the error-bars (s.e.m) from the shuffled estimation of the intercept (ϵ), which are very small. The black thick dashed lines indicate the introduced differential correlations, ϵ in the middle row and the saturation point ($1/\epsilon$) in the top row. The bottom panels show some tuning curves used in the model (bottom-left), the pair-wise correlations as a function of the difference in the preferred direction of each neuron (bottom-middle), and the variance against the mean (bottom-right).

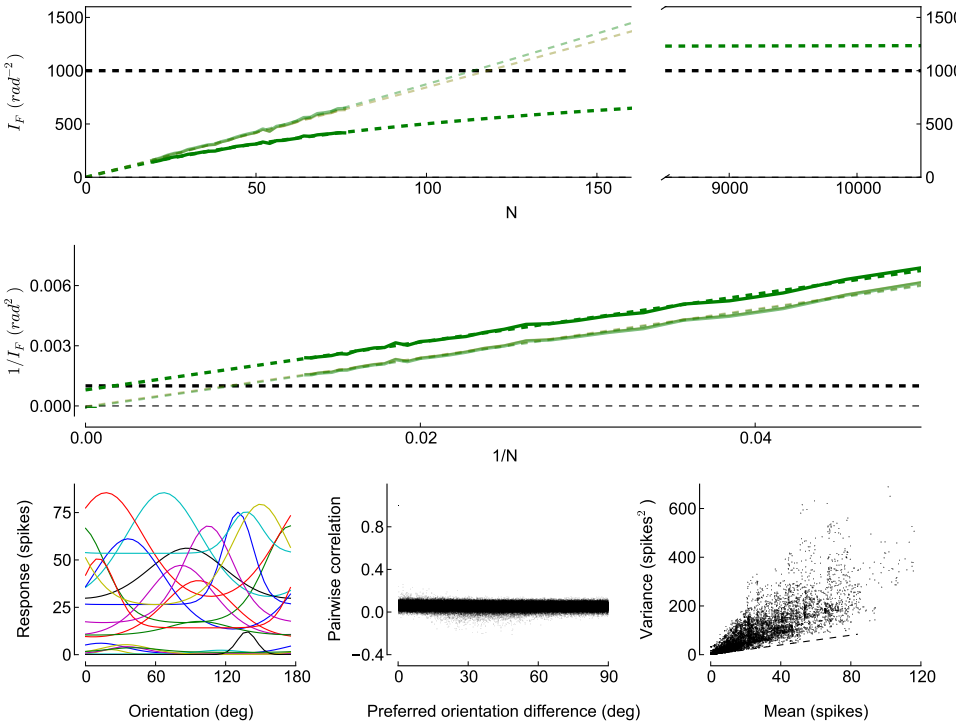

Figure 3.2: Estimation of differential correlations in surrogate data with uniform and differential correlations. Exactly the same figure as 3.1 but with uniform correlations instead of the exponential ones.

information (I_{BC}, red) in the left columns means that the correction of the bias does not improve the estimation of the information in this case. For low amount of trials (right column), however, the green and red line do not overlap anymore, meaning that the bias correction term is playing a role. Indeed, the estimation of ϵ is better when the bias correction is applied. It is important to note that the estimation without the bias correction (green) changes from the left to the right column, from a small overestimation to underestimate the true value of differential correlations, that was the same in both models. This means that the estimation of ϵ is biased for low amount of trials. Moreover, the shuffled version also shows a bias, with a small althoug negative intercept, that goes to zero as expected in the condition with many trials. The bias correction improves the estimation of the ϵ in the shuffled case but it is still negative. The bias appearing in the case of the shuffled population led us to compute the estimate of the ϵ as the difference between the raw measure and the value obtained from the shuffle. As both populations share the same statistical properties but the correlations, with covariance matrices estimated from the same amount of data, we hypothesized that the bias should affect in a similar way to both estimations, and therefore substracting the shuffled version could help while we lack an explanation of this bias. In fact, we confirmed that in most of the conditions that we explored removing the estimation from the shuffled population improved the estimation of the ϵ. In addition to the comparison between the raw measure, and difference with the shuffled case, we also sought to compare the accuracy of the estimation using the linear and the nonlinear fit on one hand, and using the Fisher information with and without the bias. A preliminary analysis suggests that the nonlinear fit is usually better than the linear one. Moreover, the bias correction improves the estimation, specially for low amount of trials, and substracting the difference from the estimation in the shuffled population also helps. From this preliminary study we can conclude that the best method

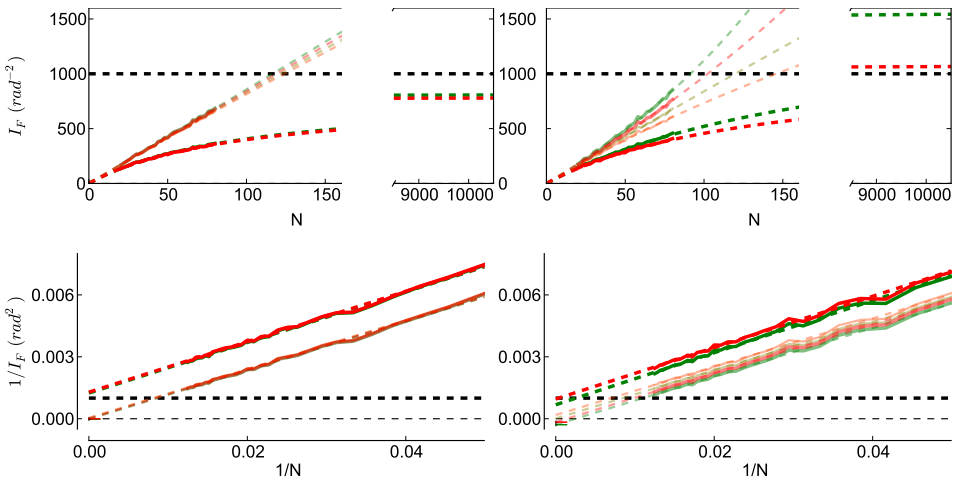

Figure 3.3: Comparison between Fisher information (I_F, green) and bias corrected Fisher information (I_{BC}, red) in two identical models with exponential and differential correlations but with different amount of trials. For large amount of trials per condition (3000, left column) the green and red lines overlap and thus the bias correction does not affect the estimation of information. For small amount of trials per condition (400, right column), however, the green line shows a clear bias in the estimation of ϵ (low row) compared to the estimation for large amount of trials. The estimation of ϵ is clearly improved by using the bias corrected version of Fisher information (Kanitscheider et al., 2015a). The lines corresponding to the shuffled data (light green and light red) and to the independent populations (olive and orange) differ as well for the case with small amount of trials, meaning that information is overestimated both in the correlated and uncorrelated data in the low trial condition.

to estimate differential correlations consists in performing a nonlinear fit following equation 3.13, that the bias correction is useful when the amount of trials is not very large, and that substracting the estimation obtained from the shuffle population usually improves the estimation.

3.3.2 Estimating the magnitude of differential correlations in real data

Once we have assesed the validity of the method for different correlation structures and we have an idea of how it behaves in realistic generative models, we sought to compute for the first time the magnitude of the mysterious differential correlations in real data. With that goal, we applied the same procedure described before to a population of V1 neurons recorded in anesthetized monkeys while drifting sinusoidal gratings were preesented in 8 randomly selected orientations. In this recordings each stimulus was shown 400 times, and therefore we are more likely to obtain better measures of ϵ by using the bias corrected version of information (Kanitscheider et al., 2015a). The results can be shown in figure 3.4.

The raw estimation of ϵ using the bias corrected information is $\epsilon_{BC}^{raw} = 3.9 \pm 0.2\ 10^{-4}$ rad^2 with both the linear (equation 3.14) and the nonlinear fit (equation 3.13, which is equivalent to a discrimination threshold of 1.1 degrees. Interestingly, if we compute the the ϵ as the difference between the raw estimate and the estimation of the shuffled version of the data, the result is the same with or without applying the bias correction, $\epsilon_{BC} = \epsilon_{BC}^{raw} - \epsilon_{BC}^{shuffle} = 5.5 \pm 0.2\ 10^{-4}$ rad^2, in the case of the nonlinear fit. This estimation corresponds to a discrimination threshold of 1.3 degrees. The same procedure applied to the other dataset that allow this method to be applied yielded a value of: $\epsilon_{BC} = 7.9 \pm 0.2\ 10^{-4}$ rad^2, which approximately are 1.6 degrees. Further analysis using bootstrap techniques are required

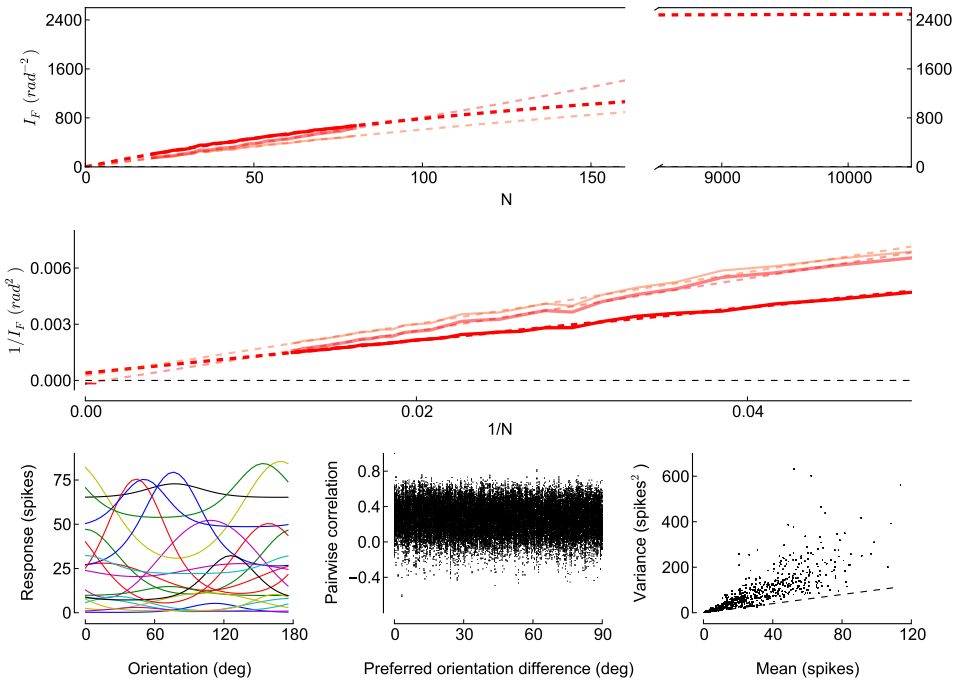

Figure 3.4: Estimation of differential correlations in real data using the bias corrected Fisher information. The figure follows the same color code as previous plots. Due to the low amount of trials there is a bias in the estimation, as can be seen from the negative prediction of ϵ for the shuffled data. Therefore, we corrected our estimation with the estimation of the shuffled uncorrelated population.

to compute more reliable confidence intervals of these measures, but in a first approximation we can conclude that the neural activity in the primary visual cortex of macaque cortex shows a discrimation threshold around 1.5 degrees.

3.3.3 Synergy OR redundancy?

One of the most striking differences when comparing the scaling of information with the network size for the real data and for the models is that the real correlated population contains more information than the uncorrelated ones (the trial shuffled version and the independent). We know that population with differential correlations must be redundant for large N, but the scaling of information with N for small ensembles favours clearly a synergistic code. Another important difference between the real data and our models is that the pair-wise correlations, although having the same mean, show larger hetereogeneity in the experimental data. To further investigate these differences, we built more realistic models including global fluctuations in addition to the differential and the exponential correlations. In our previous work (chapter 2, (Arandia-Romero et al., 2016)) we have seen that V1 tuning curves undergo multiplicative and additive modulations depending on the activity of the whole network. Moreover, we found that these global fluctuations may play important roles for information processing as well. These results motivated us to include global fluctuations in our models to further asses the validity of our method to measure differential correlations, and get more insights about the difference in the scaling of information with the network size between our models and the experimental data. Therefore, before the Poisson step, we added these heterogeneous multiplicative and additive modulations controlled by a global gain g. As in previous models, we maintained the similar mean pair-wise correlations and Fano factors as in the real data. Regarding the estimation of the differential correlations, global

fluctuations do not limit information (Kanitscheider et al., 2015b) and therefore they should not affect the magnitude of the differential correlations.

Looking at the top panel of figure 3.5, we can see that the correlated population (red) is always below the uncorrelated ones (lighter red). This plot corresponds to a generative model with small differential correlations ($\epsilon = 0.0002$) and with exponential correlations. The middle row shows the same model but with global modulations, both multiplicative and additive. In this case we can see that for small ensembles the correlated population contains more information than the uncorrelated ones. Interestingly, this is the behavior that can be observed in real data (bottom panel). Therefore, it seems that the synergistic effect that occurs for small neural ensembles can be recovered by generative models that include a global modulation.

3.4 Discussion

We developed a method to estimate the magnitude of the differential correlations, those that are information-limiting. Using generative models with different kinds of correlations and under different conditions we found that the nonlinear fit of equation 3.13, substracting the estimation from the shuffled population, and applying the bias correction suggested by Kanitscheider et al. (2015a) if the amount of trials was not very high, allows to estimate the ϵ in neural populations of 80 neurons better than other alternatives, although the accuracy of this measure should be computed yet. We applied this procedure to two datasets were the activity of macaque monkeys' V1 neurons were recorded while drifting sinusoidal gratings were presented in different orientations. Both datasets yield a value of ϵ corresponding to a discrimation threshold of around 1.5 degrees, although the accuracy of these estimations should be computed with a bootstrap procedure to

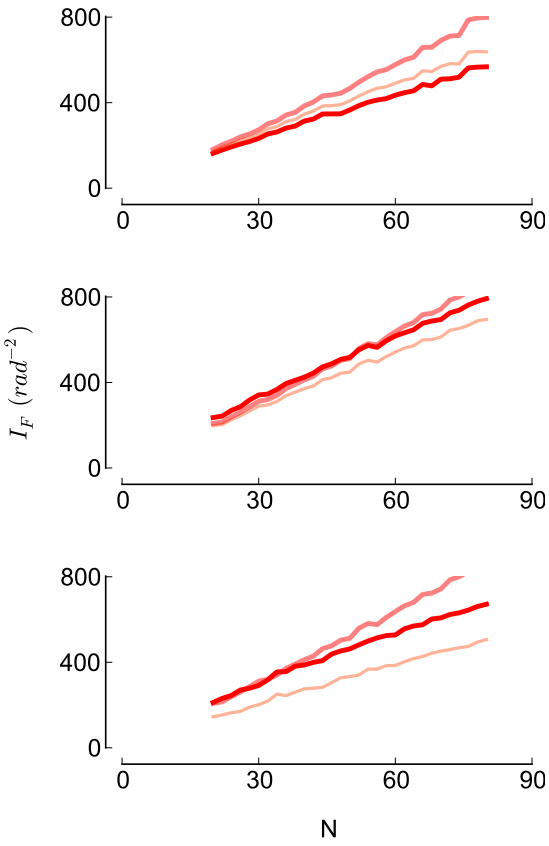

Figure 3.5: Bias corrected Fisher information for the true correlated population (red), for the trial-shuffled (light red) and for the independent (orange) populations, for a generative model with exponential and differential correlations (top row), the same model with global modulations (middle row) and for the real data (bottom row). Introducing global modulations (middle row), ensembles smaller than 50 neurons show larger information if correlations are taken into account, as observed in real data for small ensembles (bottom row)

confirm the results. Other theoretical methods to compute information, or gaining knowledge on the origin and dependence of the bias in the estimation are other lines of research that could also be helpful to improve the accuracy of this method, specially for the very low amount of trials usually available in behavioral experiments.

One of the main sources of differential correlations is representational expansion. The first sensory layer always has less neuron that following cortical regions, and therefore even if information could grow with the amount of neurons in the retina, it cannot scale in the same way in V1, for example. Therefore, the sensory noise that affects retina limits the amount of information that can be retrieved in the subsequent layers, and the differential correlations will depend on the architecture as well. The ratio of the population sizes in two following layers is going to determine the amount of differential correlations in the second layer, although other sources might contribute with extra differential correlations.

Estimating the size of differential correlations, thus, is very important to study the encoding of information in different areas and specially to understand how this information is transmitted. The variability has been shown to increase along the visual pathway (Goris et al., 2014), but still there is no evidence for the possible impact that this increased variability could have on information. A recent theoretical paper suggested that noise structure more robust to extra noise in subsequent layers is the one with differential correlations (Zylberberg et al., 2016b). Surprisingly, the noise structure that is more harmful for information content in one layer is the one that maximizes the transmission of information along the hierarchy. The measure of differential correlations developed here might help to confirm this interesting hypothesis. Nowadays, the experimental development allows to record large amount of neurons from different brain areas simultaneously. As our method works quite well with just ~ 80 neurons,

recording for example with one Utah array in V1 and another one in V2 would enough to start investigating these new questions.

Another possible source of differential correlations is suboptimal computation, and therefore our method could also be applied to studies about optimality (Pitkow et al., 2015). Once the differential correlations coming from sensory noise are known for different animals and for the different sensory domains, changes in differential correlations are likely to be very informative about other modulations that affect information and computations.

The extra assumption made to be able to compute the magnitude of differential correlations has been that information in the absence of these information-limiting correlations should grow with the size of the population. This assumption is likely to be true for very large networks, but for smaller it does not necessary hold. In fact, in real data we found that the information grows faster with N than in the independent or trial-shuffled population. By definition, networks with differential correlations should be redundant, as information must saturate. This saturation point, however, might be reached for very large amount of neurons. We found that for neuronal ensembles of less than 80 neurons, correlations seem to add information, promoting a synergistic code for small ensembles. Despite it may be surprising, there is more evidence pointing in this direction (Ince et al., 2013; Olshausen and Field, 1997; Panzeri et al., 2015). Some works challenge the simplification of considering that most of the neurons have a similar role regarding information, and propose that small subsets of neurons could convey almost all the information present in the network. The way information is distributed among the different neurons in the network is still an active area of research. Interestingly, we have been able to reproduce this different scaling of information with N found in experimental data in more realistic generative models that included mul-

tiplicative and additive global modulations (Arandia-Romero et al., 2016). This interesting feature challenges the main view about population codes, that should be redundant due to differentia correlations, and opens new questions about the role that correlations play for sensory processing, that might be different depending on the scale where we are looking at.

The results presented here must be taken with caution, as they are still very preliminary and a more exhaustive research is necessary to refine the method, and asses its performance for different population sizes and for different amount of trials. Moreover, a bootstrap procedure is likely to give more realistic error-bars than those obtained from the standard error of the fit. Finally, studying the bias in the estimation of the ϵ would be important to better understand the method, its applications and its limitations. At least using the shuffled version of the data as a baseline to compensate for this bias, we can obtained a first measure of the size of differential correlations in the primary visual cortex. Our estimation is around 1.5 degrees, which is a bit below the human discrimination threshold in perceptual tasks with this stimulus (Dosher and Lu, 1999). If this result is confirmed with a more detailed analysis and a better estimation of the errorbars, that would imply that the sensory noise is not the only limitation that affects performance, and that further processing is limiting the accuracy in these tasks.

This kind of information-limiting correlations can be interpreted as a consequence of having a network with much larger coding capacities than the amount of sensory information that is supposed to encode (Kanitscheider et al., 2015b). In fact, this should not be surprising. The cortical capabilites go far beyond the encoding of sensory information, and the limitations that we find when studying sensory processing are likely to be largely compensated by many other com-

putations that the brain is able to perform. Therefore, characterizing the differential correlations would be a great step to gain knowledge on the way information is encoded and transmitted, and the understanding of sensory processing might help to get insights about some of the other myriad of fascinating functions that the brain exhibits.

Chapter 4

GENERAL DISCUSSION

There is no whole system without an interconnection of its parts and there is no whole system without an environment.
Francisco Varela

In this thesis we tried to gain some knowledge on the role that the primary visual cortex plays in sensory perception. We faced this challenge analyzing the activity of simultaneously recorded neurons from V1 macaque monkeys obtained while drifting sinusoidal gratings were presented in different directions. On one hand we studied the limitation imposed by the previous structures taking part in the visual pathway and how they limit the information that can be retrieved in V1. We combined this work with another study that focused on the effect that global fluctuations (measured by the summed activity of the whole recorded population) had on the tuning and the information of small ensembles. If we aim at understanding cortical activity it seems necessary to consider the different sources contributing to the still misterious shared variability.

In this sense, this double-sided approach combining the feedforward sensory information with the global fluctuations that are likely

to reflect top-down influeces, helped to characterize and better understand the activity and the role of the primary visual cortex. Nevertheless, the recordings analyzed here only took into account visual information about orientation, while perception in the real world involves an integration of information coming from different sources, in a changing environment and with the possibility to move and interact. Therefore, if the goal is to understand perception, the point of view adopted for this work should be complemented with research on perception from other perspectives. Experiments in the absence of external stimulation, such as while dreaming, imagining or hallucinating may allow to obtain useful information about the intrinsic dynamics and top-down influences. In addition to exteroceptive perception, propioception and interoception should be investigated as well, as they affect the way of perceiving both the external world and the internal state. Facing the challenge of studying perception requires a broad perspective and the combination of different fields including philosophy, psychology and neuroscience, as it implies dealing with ancient and difficult problems such as the objective-subjective representation, the division between the self and the environment and consciousness.

4.1 Potentialities of the feedforward approach

Our world is restricted by our perceptual capabilities; or in other words, we are limited to what we can perceive. We have five different sensors that allow us to obtain information from light, sounds, smells, touch and tastes. However, we know that the amount of signals present out there is much larger. Even the sensors we have are able to convey information in a certain range only (of wavelengths or frequencies in the case of vision and hearing, for example) and with

certain accuracy, missing most of the signals generated in nature.

In order to overcome this limitation, we have developed tools to enhance our capabilities and even detect signals that our sensors overlook. One simple example could be the graduate glasses or the telescope, which expand the range and accuracy of visual perception. These tools somehow reduce the noise, or modify the signal in such a way that they shift the range of the stimuli making it detectable by our visual system. But in any case, the restriction imposed by the first sensory layer, the retina, remains there and determines the posterior stages of visual information processing. There have been many works devoted to understand how visual information is represented by the retina (Field and Chichilnisky, 2007) and other visual areas (Graf et al., 2011), but the way information propagates along the visual pathway received much less attention (Renart and van Rossum, 2012; Zylberberg et al., 2016b). In chapter 3 we have studied how sensory noise affects information in the first cortical stage of visual processing, the primary visual cortex (V1). The amount of visual information that can reach any brain area will be limited by the information encoded in the very first layer. Afterwards information cannot be created, and, at most, it will not be reduced.

Measuring information in neural populations implies dealing with shared variability, and thus it has been linked to pair-wise correlations during many years (Shadlen and Newsome, 1998; Zohary et al., 1994). However the work by Moreno-Bote et al. (2014) directed the focus towards differential correlations, those fluctuations that look like the signal and therefore limit information. These correlations can arise due to representational expansion or suboptimal computation, but they are very small compared to pair-wise correlations coming from other sources. As a consequence, they are masked when measuring pair-wise correlations, and getting information about them is difficult.

We presented a method that is able to predict the amount of differ-

ential correlations in realistic generative models with only 80 neurons provided the amount of trials is large. We applied this method in simultaneously recorded V1 neurons, in two datasets with populations around 80 neurons, and with more realistic amount of trials. In this way, we obtained the first known measure of the magnitude of differential correlations, equivalent to a discrimination threshold of around 1.5 degrees in both datasets. These results, however, are still preliminary and a boostrap procedure should be performed to compute more accurately the error-bars and the validity of the estimation.

What we can learn from this measure of differential correlations is very different than what we get from raw measures of information. Information can only be limited by differential correlations (Moreno-Bote et al., 2014), but it can be reduced by other kind of correlations, such as those arising from global fluctuations (Kanitscheider et al., 2015b). Therefore, this measure allows to know the very first limitation, the one related to the noise in the first layer and the amount of neurons in the first and subsequent layers. In a recent work Zylberberg et al. (2016b) suggest that the differential correlations, those which reduce drastically the amount of available information, are the most robust ones against additional noise from other layers. They claim that in order to avoid loosing extra information with additional noise sources in the following layers, the structure of the noise in the first layer should be proportional to $f'f'^T$, i.e. differential correlations. Our method could be useful to check this attractive theoretical prediction and get more insight about information propagation and the robustness of information encoded in neural populations.

Suboptimal computation is another well-known source of differential correlations, and the study of these correlations may be useful to gain knowledge about the optimality of certain computations (Moreno-Bote et al., 2014; Pitkow et al., 2015). If the following layer

is not reading out the whole population information will be lost, even if the amount of neurons in both layers is the same. Therefore, a measure of differential correlations, combined with anatomical information about the architecture of the cortical regions involved in the codification of a particular variable might give some light on the role of shared variability and the way information is processed. It would be extremely interesting to record two consecutive areas of the visual pathway and compute the differential correlations present in each one. If they are similar, we could say that the transmission of information is nearly optimal, while a difference would suggest an additional source of information limiting correlations or a different noise structure.

4.2 The influence of top-down projections on early sensory areas

When studying differential correlations the focus is directed towards the feedforward component of sensory processing, but this is clearly not enough to explain sensory perception. The phenomenon of perceptual bistability clearly shows that perception is not determined by the external input. Looking at the famous Necker cube or the vase-face image our percept can fluctuate among the two possibilities, depending on the person and the moment, but it won't never merge into an unique percept (Huguet et al., 2014; Moreno-Bote et al., 2007). Another interesting example where the external input elicits different percepts can be found in the blinks. Blinks are not noticed, while the lack of a frame in a movie with the same duration of a blink is perceived (Gawne and Martin, 2002; Golan et al., 2016), meaning that the visual system processes information in such a way that erases the lack of information due to blinks. In fact, perception is influenced by expectation, arousal state or attention, but many works show that

top-down projections are necessary to explain the cortical activity of primary sensory areas as well (Goris et al., 2014; Lin et al., 2015). This prediction has been confirmed by feedback connections found anatomically (Thomson and Bannister, 2003).

In chapter 2 we show that global fluctuations play an important role in V1, modulating the tuning of individual neurons and modifying the way information is distributed in the network. These global fluctuations reflect the activity of the whole network, in different trials but in the same session, and they should be interpreted as within-state fluctuations, in contrast to changes in the cortical states studied elsewhere (Ecker et al., 2014; Harris and Thiele, 2011). The fluctuations of the population activity described here could be due to a global cortical modulation, but as our data is restricted to one area we cannot discard that it represents a local modulation.

4.2.1 Arousal state, cortical/brain state and attention

Brain states and attention are tightly linked, but their exact relationship and the role each one plays is still unknown (see the works by Harris and Thiele (2011) and Gilbert and Sigman (2007) for reviews from different perspectives). It has been suggested that state-dependent processing and attention may share similar mechanisms, enhancing the representation of certain features (Harris and Thiele, 2011). It should be noticed an important difference when talking about cortical states. On one hand, we can separate the anesthetized, from the awake, or the slow wave sleep as global states. On the other hand, with-in state slow fluctuations are sometimes referred to as different states, and they might be more related to attention or arousal. Despite our datasets are likely to be recorded during a single cortical

state in each session (anesthetized or awake), we are tempted to speculate that the population activity employed in our analysis could be related to attention as well. This relationship is in part motivated by the gain modulation that we found in V1 neurons with population activity, very similar to the effects of attention on the tuning of orientation-selective neurons (McAdams and Maunsell, 1999; Rabinowitz et al., 2015; Reynolds et al., 2000; Treue and Trujillo, 1999; Williford and Maunsell, 2006). The main hypothesis to explain attentional effects on tuning is that the increase in the gain of the selected neurons would imply an increase in the signal-to-noise ratio, improving the accuracy on perceptual tasks (Ardid et al., 2007; Compte and Wang, 2006; McAdams and Maunsell, 1999; Rabinowitz et al., 2015).

In contrast with this hypothesis, we did not find an increase in the total information of the network with population activity (see figure 2.20). Interestingly, we found an increase in information only for ensembles of multiplicatively modulated neurons (see figure 2.17). In addition to the multiplicative ones, however, we also found many neurons that underwent an additive modulation with population activity. In this latter case, the information content was decreased for larger population activity, and finally it compensated the increased information in the multiplicative neurons. Therefore, our work could challenge the common view that attention improved performance just increasing the neural gain, and further research seems necessary to understand the mapping from sensory activity to perceptual decisions and the role that these slow fluctuation might be playing.

In this line, it would be helpful to apply the methods to disentangle between multiplicative and additive modulations developed in this work to data recorded during attention-mediated tasks. In this way, we could confirm if the modulation induced by attention is purely multiplicative as previously suggested (McAdams and Maunsell, 1999) or it involves additive effects as well (divisive and substractive as in

(Wilson et al., 2012b)) and is similar to the one we describe by population activity.

A famous controversy when dealing with cortical states is whether there are two clearly distinct states or they just change along a continuum. Cortical states have been closely related to levels of synchronization, and since the discovery of the up and down states in anesthetized rodents (Steriade et al., 1993; Stern et al., 1997) a lot of effort has been devoted to understand these slow oscillations (Compte et al., 2003; Cossart et al., 2003; Destexhe et al., 2007; Haider et al., 2006; Petersen et al., 2003; Sanchez-Vives and McCormick, 2000; Vyazovskiy and Harris, 2013) and to find a similar bimodality in other species (Lampl et al., 1999), or other conditions, such as awakefulness (Poulet and Petersen, 2008; Steriade and Timofeev, 2003). Very recently Engel et al. (2016) claim that they found signatures of bimodality in awake monkeys using data recorded during a spatial location task. In contrast with this literature, our work favors the hypothesis of a continuum of states, as the modulation of the tuning by population activity does not show any all-or-none effect. Actually, we find that the tuning of V1 neurons undergoes a gradual modulation with population activity (see figure 2.11). Thus, our results favor a modulation that changes the tuning of V1 neurons along a continuum rather than showing different discrete states.

The population activity used here and the attentional state could seem a bit too far to be compared. However, there is growing evidence that support this hypothesis (Reimer et al., 2014; Vinck et al., 2015). The arousal state is very related to attention, and it has been shown to enhance visual processing (Vinck et al., 2015). At the same time the arousal state has been found to be closely related to the pupil size (Reimer et al., 2014; Vinck et al., 2015), which is widely applied in psychophysical experiments with humans (Hess and Polt, 1960; Kah-

neman and Beatty, 1966). Therefore, tracking the level of dilation of the pupil while performing perceptual tasks is likely to help understanding how attention and cortical-states affect sensory processing, and it is experimentally cheap and relatively easy. In the same line, it would also be interesting to investigate further the relationship between the population activity used here, or the LFP that is tightly locked to it (see figure 2.8 and (Okun et al., 2015)), and the pupil size. In this regard, Vinck et al. (2015) already found some links between the arousal state, the LFP and the neural responses. Interestingly, Reimer et al. (2014) show that dilation of the pupil correlates with the activation of VIP+ and inhibition of SOM+ interneurons, which at the same time are involved in increasing the gain of visual responses during running (Fu et al., 2014). In fact, there are many studies suggesting that locomotion and behavioral state could also generate a gain modulation of V1 responses (Bennett et al., 2013; Niell and Stryker, 2010; Polack et al., 2013; Reimer et al., 2014). Therefore, the behavioral state seems an important feature that should be taken into account to better understand the responses of the primary visual cortex, the gain modulation, and the roles that the population activity might be playing.

The relationship between synchronization and cortical states also motivated research on noise correlations. If asynchronous states are related to awake and more active states, it would make sense to try to decorrelate neural activity to further increase the performance. This fact lead to many studies investigating the role that noise correlations could play in attention. Correlations have been found to be reduced by attention (Cohen and Maunsell, 2009; Herrero et al., 2013; Mitchell et al., 2009) and other conditions such as learning (Gu et al., 2011; Jeanne et al., 2013), cognitive challenge (Ruff and Cohen, 2014), adaptation (Gutnisky and Dragoi, 2008), task engagement (Downer et al., 2015), or wakefulness (Poulet and Petersen, 2008). These results sug-

gest that attention could also modify the underlying connectivity of the network. Contrary to what could be expected from these works, we did not find a clear reduction in noise correlations with population activity. However, the increase in the mean response mediated by attention in V1 has been found to be due to cholinergic input (Disney et al., 2007; Herrero et al., 2008, 2013), while the decrease in noise correlations seems to be dependent on NMDA receptors (Herrero et al., 2013). Therefore, the analysis performed here is likely to reflect just one part of the story concerning attention, and the modulation of pair-wise correlations could arise in another way. In any case, the theoretical prediction that lower pair-wise correlation should enhance information is too simplistic, as information can be affected by the fano factors or the mean responses as well. Moreover, as explained in chapter 3, the magnitude of pair-wise correlations does not need to correlate with information. It is the structure of the noise correlations, the direction of the noise, which is more important when measuring information. It has been shown that correlations coming from global fluctuations just reduce information, but they do not limit it as differential correlations do (Kanitscheider et al., 2015b). If the global fluctuations are related with attention, it would make sense to reduce global fluctuations during attended conditions to enhance the information content.

Theoretical models of attention also suggest the importance of the top-down influences to better describe the rich phenomena that has been observed in this topic (Deco and Rolls, 2004, 2005).

In addition to attention, there are other top-down influences that are relevant when studying perception. Expectation (Summerfield and De Lange, 2014) or the context where the stimulus is presented also affect the results of perceptual tasks, and the attentional effects of V1 have been shown to be increased by more complex and larger context (Ito and Gilbert, 1999; Li and Gilbert, 2002; Motter, 1993).

4.2.2 Multiplicative and additive modulation of tuning

The multiplicative and additive modulations by population activity described here are quite similar to the effects found by (Wilson et al., 2012b). They found that parvalbumin-expressin (PV) neurons mainly divide responses while somatostatin-expresion (SOM) just subtract responses maintaining their selectivity unaltered. We do not have information about the class of each recorded neuron, but it seems that different cell types respond in different ways to stimulation. The functional role that we show for the different ensembles, with different modulation of information for multiplicative or additive groups, makes extremely interesting to test the amount of information present in PV or SOM neurons in different stimulation conditions. What we can conclude from these works is that the tuning curve is not a property of the neuron that remains constant, but it is a dynamical feature that is modulated by the inputs it receives. In the hippocampus it has been shown that the individual tuning of the neurons changes from day to day, while the representation of the environment is maintained by the whole population (Ziv et al., 2013). The tuning curves of V1 neurons, however, do not seem to change in a similar way across days (Montijn et al., 2016), and although modulated, they are likely to be more estable. The ultimate role of this tuning modulation, however, is still unknown. We propose that it could reflect different ways to distribute the information controlled by the population activity, but more research is necessary to test this hypothesis.

Changes in tuning have been linked to plasticity, specially those shifts related to adaptation (Dragoi et al., 2000; Kohn and Movshon, 2004). An interesting work by del Mar Quiroga et al. (2016) was able to reproduce similar shifts in the tuning using a recurrent network in the absence of plasticity rules. Therefore, sensory adaptation could

emerge just from the neural dynamics of the recurrent network, and plasticity does not need to be present to reproduce the effects of adaptation. In the same vein, it would be interesting to build a recurrent network that reproduces the multiplicative and additive modulation described here. It might help to understand how these global fluctuations affect the tuning of individual neurons and under which conditions multiplicative and additive effects can be generated. There have been many studies trying to understand state-dependent mechanisms and computations from a theoretical perspective (see Buonomano and Maass (2009); Doiron et al. (2016) for reviews) and this line of research is likely to be important in the future as well.

The different effects of population activity in different cells, with more or less multiplication and/or addition, and how this modulation affects information is a novel result, and as such opens new questions. The flexibility of the tuning may allow to convey information through different pathways, in such a way that it makes either easier or more difficult the readout by the following layer. This feature of V1 neurons may allow a constant readout by the next layer, but there is no evidence so far to discard a flexible readout yet. At the same time, our results leave open the door to a distributed code, with different neurons playing different roles for sensory processing (Ince et al., 2013; Panzeri et al., 2015).

4.3 Information processing: combining the bottom-up and top-down approaches

The study of information processing is very old but it is still very controversial. At this point, it can be helpful to compare what we learned from the two distinct points of view employed here to study this issue, the feedforward and the feedback approach. From previous work it

is known that differential correlations limit information (Moreno-Bote et al., 2014), suggesting that the neural population with differential correlations should be redundant. The amount of neurons in V1 is much larger than in retina which implies that if information is proportional to the number of neurons in retina it cannot grow linearly with the population size in V1 (unless V1 neurons are much noisier in such a way that each one carries just a small piece of information, which is not the case). In other words, information is limited by the first layer of the visual system, and afterwards it cannot be created. Therefore, the data-processing inequality supports the presence of differential correlations in V1, and thus it should be a redundant population. The models of neural activity developed here with differential correlations support this idea (see figure 3.4), showing larger information for the independent or shuffled populations compared to the true correlated one.

In contrast with this result, when analyzing real experimental recordings, we found that small neuronal ensembles carry more information than the uncorrelated ones. This means that correlations somehow increase the amount of information in the ensemble, challenging the redundancy of population codes with differential correlations. Surprisingly, when adding global fluctuations to the model with differential correlations, we are able to recover a similar behavior of information with the population size, with larger information in the correlated ensembles. Global modulations do not limit information as differential correlations, but they reduce the amount of information that can be retrieved from the population (Kanitscheider et al., 2015b). From this work, we could expect a limitation of information when including differential correlations in the model, and a further decrease without changing the saturation point when adding global fluctuations. This predictions, however, are not fulfilled by our models.

Obviously, for larger populations information in the correlated population must be lower than for the uncorrelated one, as it must saturate due to the differential correlations. But the fact that for small ensembles correlations increase information supports the idea that perception is driven by small ensembles of neurons (Houweling and Brecht, 2007). A recent review article (Panzeri et al., 2015) also favors this hypothesis based on previous works showing that small subset of neurons are able to carry almost all the information present in the whole observed population (Ince et al., 2013; Olshausen and Field, 1997). They claim as well that the sparseness observed in cortical data (Barth and Poulet, 2012) is another fact to take into account this view.

The results described in 2 can also be interpreted in such a way that support this hypothesis. We found that information in small ensembles is modulated in different ways by population activity, depending on the effect that the tuning of the neurons in the ensemble underwent. Interestingly, and despite the small ensembles are clearly modulated by information (see 2.17), the total amount of information in the population remains constant and independent of the population activity (see 2.20). From these results we could speculate that there is always an ensemble or a few ensembles carrying a large part of information in the network, and that the most informative ensemble or ensembles in the network change according to the level of population activity.

4.4 The role of the primary visual cortex

The interpretion of V1 as the first cortical stage of visual processing that receives the visual input from the LGN was based on the feedforward model, but as we have shown the top-down influences are very important as well. In fact, there is growing evidence in favor of V1 playing different roles in the absence of external stimula-

tion. The spontaneous activity has been shown to be very structured (Arieli et al., 1996) and it has been related to the internal model of the environment (Berkes et al., 2011). Nonstimulated V1 may also convey information about surrounding context (Smith and Muckli, 2010). It seems that this region may also contribute to some hallucinations (Bressloff et al., 2002). Moreover, anatomical studies show that the primary visual cortex receives more connections from "higher" areas than from the LGN (Budd, 1998). Therefore, with the evidence we have so far it is difficult to establish a difference in importance between the bottom-up and the top-down influences in V1, and it has even suggested that V1 may be important for higher cognitive functions (Muckli, 2010). Supporting this ideas, it has been shown that V1 activity is modulated by changes in the surroundings of the stimulus that control its predictability (Alink et al., 2010). Thus, in addition to the gain control discussed earlier, top-down influences may also reach V1 by the mechanism of predictive coding (Rao and Ballard, 1999). Then, instead of the starting point of the feedforward pipeline, V1 has been suggested to be a highly interacting region that integrates information coming both from bottom-up sensory stimulation and from top-down connections, more like a buffer that performs calculations (Lee and Mumford, 2003; Mumford, 1991) or as an importat area in the predictive coding framework (Rao and Ballard, 1999; Rauss et al., 2011; Spratling, 2010).

Top-down influences could reach V1 through the strong connections it has with V2 and V4, or mediated by the weaker projections from inferotemporal cortex (Salin and Bullier, 1995), which can receive feedback signals from higher areas and transmit them towards earlier areas in the hierarchy (see the work by Deco and Rolls (2004) for a model of visual attention taking into account these areas). As explained before, the primary visual cortex is modulated by attention and it is also necessary for visual awareness (Tong, 2003), as shown

by blindsight studies in patients with lesions in V1 (Leopold, 2012; Weiskrantz, 1996).

Experiments studying perceptual learning also found differences in the individual properties of V1 neurons, highlighting the functional changes that individual neurons can undergo: when using orientation-discrimination tasks a sharpening of the tuning curves has been observed (Schoups et al., 2001), while contextual influences are affected by shape-discrimination tasks (Crist et al., 2001; Li et al., 2004). These works support the idea that the changes we found in V1 neurons are very likely to play a role for sensory processing. The main learning mechanism is the Hebbian rule, which basically modifies the strength of the synapses depending on how often they are used. In this way, highly correlated neurons tend to be more correlated, while weakly correlated neurons tend to weaken their connections. Gilbert and Sigman (2007) proposed an alternative to this view based on the gating controlled by feedback projections. Certain ensembles of top-down connections could gate subsets of horizontal connections in a selective manner. Then, different horizontal inputs can modulate the response of individual neurons facilitating the performance of the required task. In this hypothesis, the learning would involve linking properly the subsets of inputs that are actually useful for each task with the feedback signal conveying information about the task that needs to be performed. This hypothesis is compatible with the modulation we found in V1, but of course a similar analysis should be performed in task-mediated cortical data to see if the effects we found actually support this view.

The results presented in this work leave open the possibility that the population activity together with the readout are controlling the information that goes through V1. In this line, it would be interesting to study whether external information in cortical areas connected to

V1, like V2 or V4, are affected by the network activity in V1. Population activity seems to control the distribution of information, and this could happen in such a way that it is more or less accessible for the next layer. Whether population activity in V1 affects information content in any of the layers V1 projects to, or whether facilitates the readout in the following layers is a topic for another research though.

Fortunately, the development of experimental techniques permits nowadays to record different areas simultaneously. Recording two following areas would allow to ask new questions about cortical activity and computation. How does information represented in one area reach the next one? How the second region reads out information from the first one? Does the second area add extra noise? Which is the structure of this noise? Is it harmful? Does it change the information content from one layer to the next? These are just some examples of questions that can now be tackled experimentally and may help to better understand the results described here and the role of V1 in visual perception.

4.5 Perception: a broad overview

Facing the problem of perception is a great challenge and it requires a broad strategy. The approach taken here can help unraveling the neural correlates of sensory perception, but I think it should be complemented with research from different perspectives if we really aim at understanding perception deeply. So, in the following lines I will review some implicit assumptions made in our works that are quite common in the field but that are important to understand the results obtained and their limitations. The reductionist approach is ruling science nowadays, but as important as the results of the part where we focus is to know which other parts we are neglecting, which simplifications have been applied, and how they affect the whole system of

study. The assumptions taken for this work are not the only possible ones, neither the only necessary ones, and thus I find important to be aware of them and offer alternatives to complement our work.

4.5.1 Information flow and cortical hierarchy

It is very common to assume that the information flow starts in the sensory organ, which captures the external stimuli and sends an electric signal to the following layer. The sequential processing paradigm states that information goes from lower to higher areas, depending on the distance to the stimulation, from the pheriphery to the central regions. In this scheme we can talk about bottom-up sensory information that goes up along the hierarchy and top-down information that is supposed to modulate this information. However, as discussed earlier, the somehow secondary role given to the top-down inputs is not able to explain some phenomena.

In this regard, it seems fruitful to analyze the spontaneous activity of the visual cortex and other regions, and specially the activity in the absence of external stimulation but under other conditions: while dreaming (Braun et al., 1998; Horikawa et al., 2013; Nir and Tononi, 2010), while imagining different scenes or situations (Kosslyn et al., 1993; Stokes et al., 2009) or during hallucinations (Howard et al., 1998; Jardri and Denève, 2013; Jardri et al., 2016). These processes allow to study intrinsic mechanisms to generate percepts, and they may be useful to disentangle the feedforward component due to the external input from the effect that top-down projections from higher cortical areas may have in earlier sensory areas. Of course, in each case the cortical state must be taken into account, as the spontaneous activity will be different while doing a mental imagery task or while sleeping. Moreover, understanding the mechanisms responsible of visual hallucinations could be very useful to reduce them when they have a negative impact, as in some psicosis. The investigation of hallucinations and

visual illusions might help to gain knowledge on other diseases like schizophrenia (Notredame et al., 2015). In this sense, it is always much more effective to treat the root generating the symptoms, than treating the symptoms alone, in the same way it is more effective to use antibiotics for an infection by bacteria than using painkillers for the headache originated by that infection.

As discussed earlier, the difference in importance of the information coming from bottom-up projections compared to the one carried by feedback mechanisms may not be the best picture to explain neither the activity in V1, nor the perceptual processing. Even if the use of the term hierarchy, along with higher and lower cortical areas, is quite extended, there are different theories that do not assume a starting and an ending point for the information flow (Deco and Rolls, 2004, 2005; Gilbert and Sigman, 2007; Lee and Mumford, 2003; MacKay, 1956; Sporns et al., 1991). Instead, it is proposed that perception is a consequence of a reverberation or resonance between bottom-up and top-down information. This resonance could be activated by an external stimulus or by changes in the internal state, allowing different brain states depending on the sensory evidence and the task requirements in each moment.

From the Helmholtz machine (Dayan et al., 1995) to the framework of predictive coding (Rao and Ballard, 1999) it has been suggested that the brain could build generative models on the hidden causes of sensory signals. The prediction of these models travels through feedback connections and it is compared to the bottom-up sensory information. A mismatch between the expected and the obtained signal is called the prediction error and is employed to update the internal model. The bayesian version of the probabilistic inference performed in this framework has been applied to the visual cortex (Lee and Mumford, 2003), and it has also been proposed as a general framework to study

brain functioning (Knill and Pouget, 2004) that could be implement by probabilistic population codes (Ma et al., 2006). The free-energy principle (Friston, 2010) also offers an interesting view under a similar idea of minimizing the mismatch between the expected and the received stimulation.

Therefore, the approach taken here applying an external stimulus and analyzing the activity of the primary visual cortex as the first step of visual processing in the cortex, although useful, can tell us just part of the story and it is obviously not enough to understand visual perception.

4.5.2 Beyond objectivity: alternatives to the external observer when studying perception

In order to study the information about orientation present in V1 we assumed the information-processing paradigm based on the work by Shannon on communication (Shannon, 1949). In this theory, there is an emitter sending a message to a receiver, with the message being altered following a code. We did not use strictly the measure of information proposed by Shannon, but we assumed that there was a quantity (the orientation) that was "encoded" in neural activity and which we tried to decode. Then, we faced what is called "the inverse problem", we tried to reconstruct the orientation from the neural activity, given that the neural code is unknown. One of the problems of this approach is that we actually don't know if the information we measure is actually used by the animal. Combinining studies on perception with behavior seems a good idea to overcome this limitation (Panzeri et al., 2017).

In our research, then, we assumed that the orientation is an objective property that exists outside the animal and we aimed at extracting

this information from her neural activity. This perspective of an external stimulus internally represented by spiking activity was suggested by the pioneering work of Marr (1982). In this work, Marr defined the problem of perception as knowing "what is where in the world", and suggested that the brain should be able to somehow represent this information. In this influential work, Marr (1982) proposed three different levels of analysis of the perceptual systems. The first level is the computational one, and it deals with the question: What does the system do? The second is the algorithmic or representational level, which tries to understand how does the system perform the tasks investigated in the computational level. Finally, the physical level studies how are implemented the algorithms analyzed in the previous level. According to this separation, our research involves mainly the algorithmic or representational level, assuming that what the perceptual system does is to represent the orientation of the stimulus. Usually the perceptual system would be the whole animal, but in this case the monkey has been reduced to the spiking activity of some neurons in V1. This is a huge simplification, and although it is very useful to get insights about the roles that V1 may play, it should be taken into account when trying to combine the knowledge obtained in this way with other research. Dividing the perceptual system of the animal in different domains has been a very fruitful strategy to gain knowledge on each sensory domain, but perception involves the combination of information from different domains. Despite some attempts to explain the integration of inputs from different modalities in a single experience, the so-called binding problem (Singer, 1999; Treisman, 1996) remains so far unsolved.

When trying to extract the orientation of the stimulus from the neural activity we are assuming that the orientation is something objective that exists outside the animal, meaning that an external observer is necessary. This has been termed "computational objectivism"

by Thompson et al. (1992), in a reference to the first level of analysis of perceptual systems proposed by Marr (1982) explained earlier. Instead, they favor the "enactive" approach, very related to embodied theories, which aims at unraveling the sensory-motor patterns that are responsible of the visual guidance of the animal (Varela, ThomsonVarela92Vision).

In this regard, Brette (2013) offered a new definition for the computation level of perceptual systems from the point of view of the "perceiver embedded in its environment", instead of relying on objective descriptions that depend on an external observer. In this essay, Brette extends ideas from other explanations of perception (Gibson, 1979; O'Regan and Noë, 2001) in what he called subjective physics. He compares the perceptual system that tries to know its environment with the scientist that aims at unraveling physical laws from nature by making experiments and measurements. In contrast with the traditional approach in computational neuroscience where the system tries to achieve a specific transformation from a given input, the perceptual system explores the environment by making voluntary actions, which at the same time shape the environment, with the goal of finding "sensory laws" or "invariant structures". Instead of the notion of information used in the framework of communication, Brette (2013) proposes the subjective structure of the world as the goal of the perceptual system, a notion that is closer to the notion of information as scientific knowledge. The system tries to discover these subjective structures in the form of laws, which are not just collections of observations, but they also offer predictive power about future observations. As in science, predictability is what makes knowledge useful. These laws contain knowledge about the world if they offer the possibility to be falsified in the future, following the notion of falsifiability proposed by Popper (1959). Falsifiability is an important concept in philosophy of science that can be used to separate scientific knowledge (that could potentially be falsified by an experiment) from metaphys-

ical knowledge (there is no experiment to falsify it). Then, subjective physics is developed by "ecological reduction", that is, removing any prior knowledge about the world on favour of the perceptual system and applying falsifiability to the "experiments" that the perceiver performs in the environment to learn the relationships between actions and sensory inputs. Statistically, this approach is related to the active learning scheme (Cohn et al., 1996).

As a consequence of removing metaphysical knowledge from the perceptual system, applying this framework to computational neuroscience would favour autonomous models that perform actions (similar to Brooks (1991) in the field of robotics, or the work by Thompson and Stapleton (2009) in philosophy of mind) rather than the usual neural models that represent particular properties of the world.

However, coming back to our data, if we take the point of view of the perceptual system (V1 in our case), presenting the very same stimulus s many times will be perceived as an stimulus $s + \delta s$. This mismatch between the "real" and the perceived stimulus in each trial is characterized by δs, and it is related to differential correlations in neural activity, noise that looks like the signal and that the perceptual system cannot distinguish. Thus, our approach allow us to study the limitation of the perceptual system that would be impossible, or at least very difficult, in the absence of an external observer.

4.5.3 Integrating with sense: the interacting self, the environment and consciousness

Most of the research on perception focuses on exteroceptive perception, trying to understand the effect of the external stimuli. However, propioception and interoception are also part of our world, and they can actually affect the way we perceive external stimuli. Both exte-

roception and propioception are very related and they are important for the interaction with the environment. In fact, with appropiate stimulation the body representation can be modified, as in the rubber hand illusion (Botvinick and Cohen, 1998; Tsakiris and Haggard, 2005), in out-of-body experiences (Blanke et al., 2004) or in other settings (Sanchez-Vives et al., 2010; Slater et al., 2010). Therefore, the representation of the body is quite plastic, and the clear separation between the self and the environment gets difuminated under certain stimulation conditions. Some experiments investigating the experience of body ownership show that predictive multisensory integration is necessary to generate the body representation (Aspell et al., 2013; Suzuki et al., 2013). These experiments show that perception is also closely related to the interaction with the environment, and to the notion of presence and consciousness. These are difficult problems even to define, but there have been interesting approaches using virtual reality (Sanchez-Vives and Slater, 2005). Despite all the difficulties and unknows on these topics, it seems that the integration of sensory information from different sources, both external and internal, plays a crucial role.

In this regard, it seems that synaesthesia, where the stimulation of a sensory domain generates experiences in other sensory modality, offers a good opportunity to study multisensory integration. The famous technique developed by Ramachandran and Rogers-Ramachandran (1996) to treat patients with phantom limb pain is one of the big succeses in the practical application of synaesthesia, but it offers an extremely valuable window to study perception (Blakemore et al., 2005; Ramachandran and Hubbard, 2001a,b). Synaesthesia has been used inside the predictive processing framework (Seth, 2014) to test sensorymotor theories of perception (Gibson, 1979; O'Regan and Noë, 2001) which are the basis of the subjective physics framework (Brette, 2013) explained before.

Summaryzing, studying perception involves dealing with a vast amount of different phenomena, which are difficult to understand together. As a consequence, there are different ways to reduce or simplify the problem, and each one offers interesting insights about the way perception works. In this research we investigated the role that the primary visual cortex may play in coding information about the orientation of an external stimulus. We found that the global fluctuations measured by the population activity affect the tuning and the information of neuronal ensembles, supporting the idea that top-down influences are important to understand V1. Moreover, we developed a method to quantify the sensory noise which might be applied to the study of information transmission between different areas. Our doble-sided approach combining the bottom-up and the top-down influences also opens new questions about the way populations encode information and the role that different noise structures may play to encode and transmit information in reliable manner. Future research aiming at cracking the neural code for sensory perception is likely to make progress by studying the information that is actually employed by the animal, combining the statistical analysis of sensory information under different behavioral conditions (Panzeri et al., 2017).

4.6 Concluding remarks

Despite everyone is aware of -at least some of- its perceptual capabilites, how they actually work is still unknown. The problem of perception has been faced by different cultures throught the history and all around the world, with different goals and in very different ways: from the ancient rituals using ayahuasca around the amazonas jungle to enhance perception, to the modern psychotropic drugs nowadays used to have fun. In the western civilization it has appeared in the fields of philosophy, psychology, neuroscience, machine learn-

ing, robotics, ... and even literature. Therefore, different perspectives are likely to provide complementary insights and a multidisciplinary approach seems crucial to advance in its understanding.

This problem is so deep that it reaches questions of philosophy of science like: What is knowledge? From the perspective of who? How can this knowledge be acquired? Research on propioception and body ownership opened new questions about the limit between the self and the environment, something that seemed clear from Descartes. Moreover, the study of perception also challenges the common scientific paradigm based on objective properties of the world, making difficult to disentagle between the subjective and the objective information. Information may depend on the perspective as well, that is, whether the external observer or the subject is collecting the data. This objective-subjective paradox is specially clear in the case of interoception, the less attended side of perception (at least by western science) that recently started to be studied in a scientific way.

Finally, perception is also very related to consciousness and although there is a growing interest in this topic there is still no agreement in the definition of the problem (Chalmers, 1995; Crick, 1994; Koch et al., 2016; Thompson and Varela, 2001; Tononi et al., 2016; Varela, 1996) or even in the existence of the problem itself (Dennett, 1993). The combination of the fields of philosophy, psychology, and neuroscience -at least- seems crucial to address these deep and ancient problems.

Nevertheless, these topics might be highlighting some problems of the current scientific paradigm, very much inherited from the success of physics in the XXth century, and very likely influenced by the wave of physicists and mathematicians that moved to neuroscience in the last decades. The reductionist approach, despite has proven to be useful to gain knowledge in neuroscience, imposes important limitations. The way science and society is organized at this moment seems to fa-

vor more normal, "puzzle-solving", science than a scientific revolution, borrowing the terms from Kuhn (1962/2012). The growing difficulty not only to explain but to define and tackle many problems of neuroscience could be a signature of the limitation of the current point of view, and looking in other directions or from other perspectives may help at this point . A change in paradigm does not necesarily imply denying all the knowledge obtained so far. Indeed, the current paradigm is accepted because it is able to explain many things and its techniques still are useful to discover many others. It could be even that another paradigm is not able to explain all the concepts that the current one allows to understand, but it can offer a different perspective and new tools to face these problems in another way (Kuhn, 1962/2012).

A succesful application of borrowing knowledge and methods from another paradigm can be found in the Nobel prize award in Physiology and Medicine in 2015. Tu Youyou took advantage of knowledge from Traditional Chinese Medicine that applied to her scientific (in the Western sense) research, using pharmaceutical tools to discover a treatment against malaria (Tu, 2011). Without going that far, one example of a change in perspective could be the study of first-person data, as Lutz et al. (2002) and Varela and Shear (1999) made when trying to gain knowledge on consciousness, instead of relying in an external observer. In any case, there is no doubt that knowledge from different fields, attitude to share and discuss ideas, an open mind and a good dose of creativity are helpful ingredients to continue exploring this fascinating world.

Epilogue

"Not to be absolutely certain is, I think, one of the essential things in rationality."
Bertrand Russell

When I started doing research I still had the positivist conception of science as a method to extract objective knowledge from the study of nature, with the romantic goal of generating knowledge that would contribute to solve the problems of society and foster its development. Some years later, with more life and research experiences, and after some readings on the history and philosophy of science, the position from where I interpret science has changed considerably. Because of this, I thought it would be coherent to end the thesis with some thoughts and concerns about science and its interaction with nature and society that I developed in this period.

The common view in science is to identify scientific knowledge as the result of applying logical reasoning to empirical and experimental data, which are obtained independently of the experimenter. In this way, the human being and nature are separated as the subject that studies and the object to be studied. The final goal of science defined in this way would be to reach a single, unified, objective theory to explain all natural phenomena by applying the scientific method. Nevertheless, this point of view was challenged fifty years ago by some

historians of science and philosophers like Kuhn (1962/2012) or Feyerabend (1974). They stressed the idea that scientific observations are not objective, and that they are always affected in an inevitable manner by theoretical commitments. Kuhn argues that being scientist is something that requires learning, and that the training inevitably shapes the way to do science. In this sense, Kuhn employs the word *paradigm* as the set of scientific practices and conceptual schemes accepted by a group of scientists during a given period of history. Each group that agrees on some activities and concepts forms a scientific community. In this regard, it is important to note that, according to Kuhn, these practices and conceptual schemes cannot be assessed in terms of truth or false. In order to explain better his notion of paradigm, Kuhn uses the term "disciplinary matrix", which he defines as: "the entire constellation of beliefs, values, techniques and so on shared by the members of a given [scientific] community." Therefore, committing to a paradigm implies a particular way of doing science, and different paradigms would allow different ways of obtaining knowledge.

According to this view, before the invention of a paradigm, there is a period of pre-paradigm research that cannot be considered science, and which is characterized by the application of different techniques to different phenomena, making difficult to compare results or to reach agreements on fundamental questions. A specially great achievement or a theory that is clearly successful would convince some people to accept it, and then the paradigm is established. This paradigm defines a set of practices and a theory, and some problems that can be faced with them. Other problems are considered metaphysical, matter of another discipline or just a waste of time for the researches that accept these conditions. For Kuhn, the focus on just one part of the phenomena is fundamental to allow scientists to investigate a limited range of problems in great detail, and it is crucial for scientific

progress, although Feyerabend (1974) does not agree on this. The onset of the paradigm, thus, promotes specialization and contact among researchers from the same community. Scientists start to work extending the current theory and bringing more phenomena into agreement with it. This period is called "normal science", or "puzzle-solving" by Kuhn. These investigations, however, cannot lead to the discovery of new sets of phenomena or to the development of new theories. The explanation of nature that normal science can provide is constrained by the "preformed and relatively inflexible box that the paradigm supplies".

At some point, the paradigm starts to be applied far beyond the initial applications that lead to its formation. As a consequence, anomalies start to appear, novel phenomena that are in contradiction with the theory or that cannot be included adapting the current theory. Then, different researchers propose different versions. This proliferation of branches and the lack of agreement may result in a crisis, similar to the pre-paradigm period. Crisis might be generated by a scientist with a different background or coming from another field, or with a different training that allows her to use different tools or to face problems from a different angle. One option to end with the crisis is what Kuhn termed a "scientific revolution", with the generation of a new paradigm that is able to explain those phenomena that could not be brought into agreement with the previous paradigm.

Nevertheless, the new paradigm does not necessarily need to explain all the phenomena that the previous one did, and different paradigms could potentially survive together. Each one may provide different sets of practices that can be applied to explain different phenomena. As a result of a scientific revolution the meaning of some familiar terms can be altered. An example can be found in the change of meaning of the word mass, from being a property in classical mechanics to be a relation after the revolutionary work by Einstein. It is important to note that this way of understanding science goes against the

accumulative explanation, implying that knowledge, as the paradigms, is temporary. There are periods of "normal" cummulative progress, interleaved with periods of crisis and scientific revolutions when new paradigms arise.

This way of understanding science might be surprising, as our experience is usually confined to the paradigm where we have been trained. The explanation of Kuhn is that the content of scientific books is the final result of a scientific revolution, a whole body of works performed by normal science to allow teaching. Its development and the history of this science, by contrast, does not appear in those books, and it is therefore difficult to get this view of paradigm changes and scientific revolutions. Summarizing, Kuhn explains science and its historical development by different paradigms, which are an agreement on a set of practices and a theory that allow researchers committed to one paradigm to explore some phenomena in depth. Different paradigms may explain different phenomena, and therefore, the objective nature does not exist anymore, as different conceptual schemes allow accessing the world in different ways, or, in other words, different paradigms offer access to different worlds.

Although these arguments are not accepted by everyone, this revolutionary work motivated many studies analyzing not only the way intellectual context affects scientific research, but also the influence that the social context could have in science. An interesting view on this topic has been offered by Haraway (1991), who claims, following the idea of Kuhn, that nature is constructed, not discovered. But Haraway goes beyond that, and by analyzing the investigation in primates in the previous century, she uncovers the effect that gender has in culture and in science. The society where scientists are raised, educated and trained is going to have an inevitable impact on the questions that are going to be asked, the tools that are going to be used, the

way the results will be interpreted and how these results are going to be explained and taught. Despite the utopian version of science relies on an objective observer that collects data and looks agnostically for the most plausible theory to explain the experimental observation, for Haraway the personal experiences are very important. They shape the way we interact with nature and how we interpret the results.

An interesting example could be the different evolutionary theories proposed by Darwin and Kropotkin, very much influenced by their political ideas, the society where they lived and the region they studied. Many other examples can be found in studies about gender, clearly biased by the patriarchal society (Bian et al., 2017; Keller and Longino, 1996).

Haraway goes against objectivity but without falling in relativism, highlighting the importance of the partial perspective. She argues in favor of embodied situated knowledge, criticizing both the essentialism and the relativism: "Relativism and totalization are both 'god-tricks' promising vision from everywhere and nowhere equally and fully, common myths in rhetorics surrounding Science. But it is precisely in the politics and epistemology of partial perspectives that the possibility of sustained, rational, objective inquiry rests".

One way to show the relationship between knowledge production and society is to analyze some social, economic and political changes in the last decades, and compare them to the evolution of scientific research. With this goal, I will follow the work by Pestre (2003, 2008), mainly focused in the United States and the Western societies, to show how deeply related is science to society, and how they affect each other.

The main goal of science is to generate knowledge, and knowledge implies power. As a consequence, the history of science has been tightly linked to politics and economics. An easy example is the success of the Manhattan project in building the atomic bomb,

that lead to a new global political order. The Space Race that came after the Second World War is another example of science used as a political tool, with a huge propagandistic exhibition of the scientific achievements as part of the Cold War against soviet communism. The economic growth in this period was accompanied by a huge scientific development, mainly promoted by the army, the universities that were funded and coordinated by the Governments and by industry. The latter showed a clear tendency to privatize knowledge using patents, but the national interests and a strong Government compensated for this with appropriate laws. Private companies, like Bell, Kodak, or IBM, invested in basic research, but the Governments protected science and the most fundamental discoveries were made publicly available in favor of the national interests. Therefore, there was an equilibrium between public and private research, which fostered the advancement of science keeping fundamental discoveries as a public good. Science was admired and the society acknowledged its progresses, natural sciences enjoyed an impressive authority in this period.

This situation, however, has radically changed. Science is seen by many people as another institution controlled by the establishment, whose role is far from being solving the problems of the society. The reaction to this feeling in the United States is scary, with many people denying the global warming or the climate change, and with policies that act in the same direction. This, of course, is not desirable, but it can be understood in the present social and economic context, with more social inequalities than 50 years ago, an economic crisis without recovery, and the predominance of individualism.

The economic development after the Second World War had a positive economic impact on the society, with the Governments redistributing, at least partially, the wealth generated with the contribution of science. With the rise of neoliberalism, however, the Governments have lost an important part of their power to control the economy

in favor of the markets and other supranational non-democratic organizations, like the World Trade Organization (WTO), the Internation Monetary Found (IMF) or the World Bank (Stiglitz and Chemla, 2002).

At the same time, the ways to produce knowledge have changed as well. The university is not the central institution for scientific research. Instead, other research centers, mainly with private funding, generate an important part of the scientific knowledge. As a consequence, the freedom of the scientist to select the topics to study is very limited and the problems that are economically profitable, either at a shorter or a longer timescale, are favored against others that might be necessary for the society.

A clear example of this can be found in the biomedical sciences, which only promote research on those diseases that are going to give profits. Moreover, drugs are patented even if that prevents people that could save their lives with them from getting access. That might be legal but the moral concerns in these cases are difficult to obviate.

In this regard, it is important to note the changes in the intellectual property laws, which are responsible of controlling the conditions of the patents according to their usefulness. Their relaxation allowed an increased commercial use of scientific knowledge (Boyle, 2002; Hesse, 2002; Kevles, 2002), and boosted restricting access to more basic research, that is patentable nowadays. These are just consequences of the transition that global economy is supposed to be undergoing towards a "knowledge economy", where knowledge can be treated as a business-product.

Therefore, the preference given to knowledge that generates benefits over scientific research on topics more important to society is likely to contribute to the new relationship between science and society. Furthermore, the disaster of Fukushima, oil spills due to accidents in oil tankers, or the atmospheric pollution do not help to recover confidence

on science and technology.

Another domain where the reductionist economy-based approach of science failed for society appears in agriculture. A seed can be genetically modified in different ways, and finally the one that is supposed to be the best is selected. The definition of "best", however, is likely to be different for the farmer and for the company in charge of the research process. Their goal is to obtain economic profits, and they often might prefer a seed that requires a particular fertilizer, even if it contaminates the water in the area or have a negative impact on the local community. Many times, the fauna, the environment and the local community are not variables taken into account to select the appropriate seed. The goal of the company is to increase productivity, not to improve the quality of life of the farmers.

The reduction of public research also reduced the development of many sciences potentially useful. Coming back to agriculture, we can find an example of a marginalized science in pedology, the study of the soils, which is not appealing for private companies.

These examples show that there is a problem in the way science and society are organized, and that it affects nature. Science is based on logical reasoning, and the reason can be interpreted as what distinguishes the human being from nature. The supremacy of reason in Western societies strengthened this distinction, and thus favored a relationship to nature based on domination. When science started to develop in this way, dominating nature could seem something positive, and it was socially accepted. At this moment of the history, with the global warming, the climate change, ozone depletion, atmospheric pollution, soil erosion, ... I think there is enough evidence that shows the degradation of nature in the last decades, which is not likely to be independent of the human being. Therefore, in the same way we decided to interact with nature in a particular way, I support the idea that

the relationship should be rebuild in a more environmentally friendly manner (Hubbard et al., 1979), which, hopefully, may also contribute to rebuild social relationships in the absence of dominations.

The development of machine learning and artificial intelligence, the peak oil and the lack of resources, the climate change, soil erosion, ... novel social, ethical and environmental problems are arising right now. The current social and economic organization must face important challenges in the near future. In fact, there are researchers proposing that a paradigm change is necessary in economy as well, and that the economy growth that controlled politics from 1945 should be modified by a new paradigm that permits the development of societies without requiring an increase in the consumption of resources (Schneider et al., 2010). The proposal of economic degrowth is different from a depression where degrowth is generated by a crisis. In this case, it should be a voluntary transition to another society requiring less resources and thus environmentally friendly and sustainable.

Science is a necessary tool to find out alternatives to obtain resources, and to set the limitations to the amount of resources that we can use if we aim at building a sustainable society. If we take a bit of perspective to analyze which are the main topics of scientific research, and we compare them with the current problems that the society is facing, we are going to notice a clear mismatch. Therefore, I think that the topics that are crucial for the development of the society should be a priority for science, and they must be socially selected and socially funded. This is specially important when the scientific and technical solutions might be risky or with irreversible effects, as in the case of nuclear plants or in genetic agriculture. In front of new possibilities, with a potential dangerous effect, caution should play an important role.

Science has a huge responsibility to guide the decisions of the society. However, it is prone to accumulate power based on authority

arguments, and it is thus important that science is controlled and organized by the whole society. In my opinion, the role of science should be to provide as many different perspectives, and as deep as possible, on the different topics to inform properly the community, and encourage discussion and debates that will finally help to make appropriate decisions. The society should not delegate decisions on "the experts", concentrating the power just in a few hands. Indeed, they are likely to be very implicated on the issue, and they are manipulable as well.

Science is not innocent, and the knowledge it generates has a huge impact on nature and society. As a consequence, it is crucial for the development of a sustainable relationship with nature but also for the construction of a healthy participative democracy.

Science being affected by society is not bad *per se*, it is just an inevitable consequence of both science and society being human constructs. Although this vision can ruin the ideal of science, I prefer to think that realizing about this fact helps to make better science and to build more adequate structures for scientific research and for political organization. Science is an extremely powerful tool that we must exploit, but we should employ it carefully.

All human beings should keep in mind that their reason is limited, and therefore so is the knowledge they can generate. There are many things we do not know, and many others that we will never know. Modesty should be a value to cultivate.

From my perspective, the goal of science should be to learn about nature, or to construct nature, in such a way that it improves life, reducing the domination in different domains present nowadays, from gender to race, from nature to culture. I really think that science plays a crucial role in shaping the social, political and economic devel-

opments, and thus it has a huge responsibility in the current situation to contribute to build a better future. In any case, a necessary first step is to start talking about these topics.

I would like to conclude with a quote by Bertrand Russell, a philosopher, mathematician, and political activist, and one of the most brilliants and coherent intellectuals of the XXth century: "The good life is one inspired by love and guided by knowledge". Following this idea, I think that science should seek knowledge inspired by love and guided and realized by the community.

Bibliography

Abbott L. and Dayan P. . The effect of correlated variability on the accuracy of a population code. *Neural Computation*, 11(1):91–101, 1999.

Abbott L. and Regehr W. G. . Synaptic computation. *Nature*, 431 (7010):796–803, 2004.

Alink A. , Schwiedrzik C. M. , Kohler A. , Singer W. , and Muckli L. . Stimulus predictability reduces responses in primary visual cortex. *Journal of Neuroscience*, 30(8):2960–2966, 2010.

Arandia-Romero I. , Tanabe S. , Drugowitsch J. , Kohn A. , and Moreno-Bote R. . Multiplicative and additive modulation of neuronal tuning with population activity affects encoded information. *Neuron*, 89(6):1305–1316, 2016.

Ardid S. , Wang X.-J. , and Compte A. . An integrated microcircuit model of attentional processing in the neocortex. *Journal of Neuroscience*, 27(32):8486–8495, 2007.

Arieli A. , Sterkin A. , Grinvald A. , and Aertsen A. . Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science*, 273(5283):1868, 1996.

Aspell J. E. , Heydrich L. , Marillier G. , Lavanchy T. , Herbelin B. , and Blanke O. . Turning body and self inside out visualized heartbeats alter bodily self-consciousness and tactile perception. *Psychological Science*, page 0956797613498395, 2013.

Averbeck B. , Latham P. , and Pouget A. . Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7 (5):358–366, 2006.

Averbeck B. B. and Lee D. . Neural noise and movement-related codes in the macaque supplementary motor area. *Journal of Neuroscience*, 23(20):7630–7641, 2003.

Averbeck B. B. and Lee D. . Effects of noise correlations on information encoding and decoding. *Journal of Neurophysiology*, 95(6):3633–3644, 2006.

Bair W. , Zohary E. , and Newsome W. T. . Correlated firing in macaque visual area mt: time scales and relationship to behavior. *The journal of Neuroscience*, 21(5):1676–1697, 2001.

Barth A. L. and Poulet J. F. . Experimental evidence for sparse firing in the neocortex. *Trends in neurosciences*, 35(6):345–355, 2012.

Baruni J. K. , Lau B. , and Salzman C. D. . Reward expectation differentially modulates attentional behavior and activity in visual area v4. *Nature neuroscience*, 2015.

Battaglia F. P. , Benchenane K. , Sirota A. , Pennartz C. M. , and Wiener S. I. . The hippocampus: hub of brain network communication for memory. *Trends in cognitive sciences*, 15(7):310–318, 2011.

Beck J. M. , Ma W. J. , Pitkow X. , Latham P. E. , and Pouget A. . Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron*, 74(1):30–39, 2012.

Bennett C. , Arroyo S. , and Hestrin S. . Subthreshold mechanisms underlying state-dependent modulation of visual responses. *Neuron*, 80(2):350–357, 2013.

Berkes P. , Orbán G. , Lengyel M. , and Fiser J. . Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*, 331(6013):83, 2011.

Bian L. , Leslie S.-J. , and Cimpian A. . Gender stereotypes about intellectual ability emerge early and influence children's interests. *Science*, 355(6323):389–391, 2017.

Bishop C. . *Pattern recognition and machine learning*, volume 4. springer New York, 2006.

Blakemore S.-J. , Bristow D. , Bird G. , Frith C. , and Ward J. . Somatosensory activations during the observation of touch and a case of vision–touch synaesthesia. *Brain*, 128(7):1571–1583, 2005.

Blanke O. , Landis T. , Spinelli L. , and Seeck M. . Out-of-body experience and autoscopy of neurological origin. *Brain*, 127(2):243–258, 2004.

Bosking W. H. , Zhang Y. , Schofield B. , and Fitzpatrick D. . Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *Journal of neuroscience*, 17(6):2112–2127, 1997.

Botvinick M. and Cohen J. . Rubber hands' feel'touch that eyes see. *Nature*, 391(6669):756, 1998.

Boyle J. . Fencing off ideas: enclosure & the disappearance of the public domain. *Daedalus*, 131(2):13–25, 2002.

Braun A. R. , Balkin T. J. , Wesensten N. J. , Gwadry F. , Carson R. E. , Varga M. , Baldwin P. , Belenky G. , and Herscovitch P. . Dissociated pattern of activity in visual cortices and their projections during human rapid eye movement sleep. *Science*, 279(5347): 91–95, 1998.

Bressloff P. C. , Cowan J. D. , Golubitsky M. , Thomas P. J. , and Wiener M. C. . What geometric visual hallucinations tell us about the visual cortex. *Neural Computation*, 14(3):473–491, 2002.

Brette R. . Subjective physics. *arXiv preprint arXiv:1311.3129*, 2013.

Brette R. . Philosophy of the spike: rate-based vs. spike-based theories of the brain. *Frontiers in systems neuroscience*, 9:151, 2015.

Britten K. H. , Shadlen M. N. , Newsome W. T. , and Movshon J. A. . The analysis of visual motion: a comparison of neuronal and psychophysical performance. *The Journal of Neuroscience*, 12(12): 4745–4765, 1992.

Brooks R. A. . Intelligence without representation. *Artificial intelligence*, 47(1-3):139–159, 1991.

Bryant H. L. and Segundo J. P. . Spike initiation by transmembrane current: a white-noise analysis. *The Journal of physiology*, 260(2): 279, 1976.

Budd J. M. . Extrastriate feedback to primary visual cortex in primates: a quantitative analysis of connectivity. *Proceedings of the Royal Society of London B: Biological Sciences*, 265(1400):1037–1044, 1998.

Bujan A. F. , Aertsen A. , and Kumar A. . Role of input correlations in shaping the variability and noise correlations of evoked activity in the neocortex. *Journal of Neuroscience*, 35(22):8611–8625, 2015.

Bullock T. H. . The reliability of neurons. *The Journal of general physiology*, 55(5):565, 1970.

Buonomano D. V. and Maass W. . State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, 10(2):113–125, 2009.

Butts D. A. and Goldman M. S. . Tuning curves, neuronal variability, and sensory coding. *PLoS Biol*, 4(4):e92, 2006.

Calvin W. H. and Stevens C. F. . Synaptic noise as a source of variability in the interval between action potentials. *Science*, 155:842–844, 1967.

Calvin W. H. and Stevens C. F. . Synaptic noise and other sources of randomness in motoneuron interspike intervals. *J Neurophysiol*, 31 (4):574–587, 1968.

Carandini M. and Ferster D. . Membrane potential and firing rate in cat primary visual cortex. *Journal of Neuroscience*, 20(1):470–484, 2000.

Carandini M. , Heeger D. J. , et al. Summation and division by neurons in primate visual cortex. *Science-AAAS-Weekly Paper Edition-including Guide to Scientific Information*, 264(5163):1333–1335, 1994.

Carr M. F. , Jadhav S. P. , and Frank L. M. . Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature neuroscience*, 14(2):147–153, 2011.

Chalmers D. J. . Facing up to the problem of consciousness. *Journal of consciousness studies*, 2(3):200–219, 1995.

Chance F. S. , Abbott L. , and Reyes A. D. . Gain modulation from background synaptic input. *Neuron*, 35(4):773–782, 2002.

Chapman B. , Zahs K. R. , and Stryker M. P. . Relation of cortical cell orientation selectivity to alignment of receptive fields of the geniculocortical afferents that arborize within a single orientation column in ferret visual cortex. *J Neurosci*, 11(5):1347–1358, 1991.

Churchland M. M. , Byron M. Y. , Cunningham J. P. , Sugrue L. P. , Cohen M. R. , Corrado G. S. , Newsome W. T. , Clark A. M. , Hosseini P. , Scott B. B. , et al. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature neuroscience*, 13 (3):369–378, 2010.

Cohen M. and Kohn A. . Measuring and interpreting neuronal correlations. *Nature neuroscience*, 14(7):811–819, 2011.

Cohen M. R. and Maunsell J. H. . Attention improves performance primarily by reducing interneuronal correlations. *Nature neuroscience*, 12(12):1594–1600, 2009.

Cohn D. A. , Ghahramani Z. , and Jordan M. I. . Active learning with statistical models. *Journal of artificial intelligence research*, 4(1): 129–145, 1996.

Compte A. and Wang X.-J. . Tuning curve shift by attention modulation in cortical neurons: a computational study of its mechanisms. *Cerebral Cortex*, 16(6):761–778, 2006.

Compte A. , Sanchez-Vives M. V. , McCormick D. A. , and Wang X.-J. . Cellular and network mechanisms of slow oscillatory activity

(¡ 1 hz) and wave propagations in a cortical network model. *Journal of neurophysiology*, 89(5):2707–2725, 2003.

Contreras E. J. B. , Schjetnan A. G. P. , Muhammad A. , Bartho P. , McNaughton B. L. , Kolb B. , Gruber A. J. , and Luczak A. . Formation and reverberation of sequential neural activity patterns evoked by sensory stimulation are enhanced during cortical desynchronization. *Neuron*, 79(3):555–566, 2013.

Cossart R. , Aronov D. , and Yuste R. . Attractor dynamics of network up states in the neocortex. *Nature*, 423(6937):283–288, 2003.

Cossell L. , Iacaruso M. F. , Muir D. R. , Houlton R. , Sader E. N. , Ko H. , Hofer S. B. , and Mrsic-Flogel T. D. . Functional organization of excitatory synaptic strength in primary visual cortex. *Nature*, 518(7539):399–403, 2015.

Crick F. . The Astonishing Hypothesis. *The Journal of Consciousness Studies*, 1(1):10–16, 1994.

Crist R. E. , Li W. , and Gilbert C. D. . Learning to see: experience and attention in primary visual cortex. *Nature neuroscience*, 4(5): 519–525, 2001.

Curto C. , Sakata S. , Marguet S. , Itskov V. , and Harris K. D. . A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex. *The Journal of Neuroscience*, 29(34):10600–10612, 2009.

Dayan P. , Hinton G. E. , Neal R. M. , and Zemel R. S. . The helmholtz machine. *Neural computation*, 7(5):889–904, 1995.

Dayan P. , Abbott L. , and Abbott L. . Theoretical neuroscience: Computational and mathematical modeling of neural systems. 2001.

De La Rocha J. , Doiron B. , Shea-Brown E. , Josić K. , and Reyes A. . Correlation between neural spike trains increases with firing rate. *Nature*, 448(7155):802–806, 2007.

Deco G. and Rolls E. T. . A neurodynamical cortical model of visual attention and invariant object recognition. *Vision research*, 44(6): 621–642, 2004.

Deco G. and Rolls E. T. . Attention, short-term memory, and action selection: a unifying theory. *Progress in neurobiology*, 76(4):236–256, 2005.

Deger M. , Helias M. , Boucsein C. , and Rotter S. . Statistical properties of superimposed stationary spike trains. *Journal of Computational Neuroscience*, 32(3):443–463, 2012.

Mar Quiroga M. del , Morris A. P. , and Krekelberg B. . Adaptation without plasticity. *Cell reports*, 17(1):58–68, 2016.

Dennett D. C. . *Consciousness explained*. Penguin UK, 1993.

Destexhe A. , Hughes S. W. , Rudolph M. , and Crunelli V. . Are corticothalamic 'up'states fragments of wakefulness? *Trends in neurosciences*, 30(7):334–342, 2007.

Deweese M. R. and Zador A. M. . Shared and private variability in the auditory cortex. *Journal of neurophysiology*, 92(3):1840–1855, 2004.

Disney A. A. , Aoki C. , and Hawken M. J. . Gain modulation by nicotine in macaque v1. *Neuron*, 56(4):701–713, 2007.

Doiron B. , Litwin-Kumar A. , Rosenbaum R. , Ocker G. K. , and Josić K. . The mechanics of state-dependent neural correlations. *Nature neuroscience*, 19(3):383–393, 2016.

Donahue C. H. and Lee D. . Dynamic routing of task-relevant signals for decision making in dorsolateral prefrontal cortex. *Nature neuroscience*, 18(2):295–301, 2015.

Dosher B. A. and Lu Z.-L. . Mechanisms of perceptual learning. *Vision research*, 39(19):3197–3221, 1999.

Douglas R. J. and Martin K. A. . Neuronal circuits of the neocortex. *Annu. Rev. Neurosci.*, 27:419–451, 2004.

Downer J. D. , Niwa M. , and Sutter M. L. . Task engagement selectively modulates neural correlations in primary auditory cortex. *Journal of Neuroscience*, 35(19):7565–7574, 2015.

Dragoi V. , Sharma J. , and Sur M. . Adaptation-induced plasticity of orientation tuning in adult visual cortex. *Neuron*, 28(1):287–298, 2000.

Ecker A. , Berens P. , Keliris G. , Bethge M. , Logothetis N. , and Tolias A. . Decorrelated neuronal firing in cortical microcircuits. *Science*, 327(5965):584–587, 2010.

Ecker A. S. , Berens P. , Tolias A. S. , and Bethge M. . The effect of noise correlations in populations of diversely tuned neurons. *The Journal of Neuroscience*, 31(40):14272–14283, 2011.

Ecker A. S. , Berens P. , Cotton R. J. , Subramaniyan M. , Denfield G. H. , Cadwell C. R. , Smirnakis S. M. , Bethge M. , and Tolias A. S. . State dependence of noise correlations in macaque primary visual cortex. *Neuron*, 82(1):235–248, 2014.

Ecker A. S. , Denfield G. H. , Bethge M. , and Tolias A. S. . On the structure of neuronal population activity under fluctuations in attentional state. *Journal of Neuroscience*, 36(5):1775–1789, 2016.

Engel T. A. , Steinmetz N. A. , Gieselmann M. A. , Thiele A. , Moore T. , and Boahen K. . Selective modulation of cortical state during spatial attention. *Science*, 354(6316):1140–1144, 2016.

Faisal A. A. , Selen L. P. , and Wolpert D. M. . Noise in the nervous system. *Nature reviews neuroscience*, 9(4):292–303, 2008.

Fatt P. and Katz B. . Some observations on biological noise. *Nature*, 166(4223):597–598, 1950.

Feyerabend P. . Against method: Outline of an anarchistic theory of knowledge. 1974.

Field G. D. and Chichilnisky E. . Information processing in the primate retina: circuitry and coding. *Annu. Rev. Neurosci.*, 30:1–30, 2007.

Finn I. M. , Priebe N. J. , and Ferster D. . The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. *Neuron*, 54(1):137–152, 2007.

Fiser J. , Chiu C. , and Weliky M. . Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature*, 431(7008):573–578, 2004.

Földiák P. . The 'ideal homunculus': statistical inference from neural population responses. In *Computation and neural systems*, pages 55–60. Springer, 1993.

Franke F. , Fiscella M. , Sevelev M. , Roska B. , Hierlemann A. , and Silveira R. A. da . Structures of neural correlation and how they favor coding. *Neuron*, 89(2):409–422, 2016.

Friston K. . The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.

Fu Y. , Tucciarone J. M. , Espinosa J. S. , Sheng N. , Darcy D. P. , Nicoll R. A. , Huang Z. J. , and Stryker M. P. . A cortical circuit for gain control by behavioral state. *Cell*, 156(6):1139–1152, 2014.

Fusi S. , Asaad W. F. , Miller E. K. , and Wang X.-J. . A neural circuit model of flexible sensorimotor mapping: learning and forgetting on multiple timescales. *Neuron*, 54(2):319–333, 2007.

Gammaitoni L. , Hänggi P. , Jung P. , and Marchesoni F. . Stochastic resonance. *Reviews of modern physics*, 70(1):223, 1998.

Ganmor E. , Segev R. , and Schneidman E. . The architecture of functional interaction networks in the retina. *Journal of Neuroscience*, 31(8):3044–3054, 2011.

Gawne T. J. and Martin J. M. . Responses of primate visual cortical neurons to stimuli presented by flash, saccade, blink, and external darkening. *Journal of neurophysiology*, 88(5):2178–2186, 2002.

Gibson J. J. . *The ecological approach to visual perception*. Psychology Press, 1979.

Gilbert C. D. and Sigman M. . Brain states: top-down influences in sensory processing. *Neuron*, 54(5):677–696, 2007.

Golan T. , Davidesco I. , Meshulam M. , Groppe D. M. , Mégevand P. , Yeagle E. M. , Goldfinger M. S. , Harel M. , Melloni L. , Schroeder C. E. , et al. Human intracranial recordings link suppressed transients rather than'filling-in'to perceptual continuity across blinks. *eLife*, 5:e17243, 2016.

Golledge H. D. , Panzeri S. , Zheng F. , Pola G. , Scannell J. W. , Giannikopoulos D. V. , Mason R. J. , Tovée M. J. , and Young M. P. . Correlations, feature-binding and population coding in primary visual cortex. *Neuroreport*, 14(7):1045–1050, 2003.

Goris R. L. , Movshon J. A. , and Simoncelli E. P. . Partitioning neuronal variability. *Nature neuroscience*, 17(6):858–865, 2014.

Goris R. L. , Simoncelli E. P. , and Movshon J. A. . Origin and function of tuning diversity in macaque visual cortex. *Neuron*, 88 (4):819–831, 2015.

Graf A. B. , Kohn A. , Jazayeri M. , and Movshon J. A. . Decoding the activity of neuronal populations in macaque primary visual cortex. *Nature neuroscience*, 14(2):239–245, 2011.

Graupner M. and Reyes A. D. . Synaptic input correlations leading to membrane potential decorrelation of spontaneous activity in cortex. *Journal of Neuroscience*, 33(38):15075–15085, 2013.

Gu Y. , Liu S. , Fetsch C. R. , Yang Y. , Fok S. , Sunkara A. , DeAngelis G. C. , and Angelaki D. E. . Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron*, 71(4): 750–761, 2011.

Gur M. , Beylin A. , and Snodderly D. M. . Response variability of neurons in primary visual cortex (v1) of alert monkeys. *Journal of Neuroscience*, 17(8):2914–2920, 1997.

Gutnisky D. A. and Dragoi V. . Adaptive coding of visual information in neural populations. *Nature*, 452(7184):220–224, 2008.

Haider B. , Duque A. , Hasenstaub A. R. , and McCormick D. A. . Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition. *Journal of Neuroscience*, 26 (17):4535–4545, 2006.

Han F. , Caporale N. , and Dan Y. . Reverberation of recent visual experience in spontaneous cortical waves. *Neuron*, 60(2):321–327, 2008.

Hanke M. , Halchenko Y. O. , Sederberg P. B. , Olivetti E. , Fründ I. , Rieger J. W. , Herrmann C. S. , Haxby J. V. , Hanson S. J. , and Pollmann S. . PyMVPA: a unifying approach to the analysis of neuroscientific data. *Frontiers in Neuroinformatics*, 3(00003), 2009. ISSN 1662-5196.

Haraway D. J. . Simians, cyborgs and women the reinvention of nature. 1991.

Harris K. D. and Shepherd G. M. . The neocortical circuit: themes and variations. *Nature neuroscience*, 18(2):170–181, 2015.

Harris K. D. and Thiele A. . Cortical state and attention. *Nature reviews neuroscience*, 12(9):509–523, 2011.

Hebb D. O. . *The organization of behavior: A neuropsychological approach*. John Wiley & Sons, 1949.

Herrero J. L. , Roberts M. , Delicato L. , Gieselmann M. A. , Dayan P. , and Thiele A. . Acetylcholine contributes through muscarinic receptors to attentional modulation in v1. *Nature*, 454(7208):1110–1114, 2008.

Herrero J. L. , Gieselmann M. A. , Sanayei M. , and Thiele A. . Attention-induced variance and noise correlation reduction in macaque v1 is mediated by nmda receptors. *Neuron*, 78(4):729–739, 2013.

Hess E. H. and Polt J. M. . Pupil size as related to interest value of visual stimuli. *Science*, 132(3423):349–350, 1960.

Hesse C. . The rise of intellectual property, 700 bc-ad 2000: An idea in the balance. *Daedalus*, 131(2):26–45, 2002.

Hoffman K. and McNaughton B. . Coordinated reactivation of distributed memory traces in primate neocortex. *Science*, 297(5589): 2070–2073, 2002.

Hoffman K. L. , Battaglia F. P. , Harris K. , MacLean J. N. , Marshall L. , and Mehta M. R. . The upshot of up states in the neocortex: from slow oscillations to memory formation. *Journal of Neuroscience*, 27(44):11838–11841, 2007.

Horikawa T. , Tamaki M. , Miyawaki Y. , and Kamitani Y. . Neural decoding of visual imagery during sleep. *Science*, 340(6132):639–642, 2013.

Houweling A. R. and Brecht M. . Behavioural report of single neuron stimulation in somatosensory cortex. *Nature*, 450(7172), 2007.

Howard R. , Brammer M. , David A. , Woodruff P. , Williams S. , et al. The anatomy of conscious vision: an fmri study of visual hallucinations. *Nature neuroscience*, 1(8):738–742, 1998.

Hubbard R. , Henifin M. S. , and Fried B. . Women look at biology looking at women a collection of feminist critiques. 1979.

Hubel D. and Wiesel T. . Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology*, 160(1):106–154, 1962.

Hubel D. H. and Wiesel T. N. . Receptive fields of single neurones in the cat's striate cortex. *The Journal of physiology*, 148(3):574–591, 1959.

Huguet G. , Rinzel J. , and Hupé J.-M. . Noise and adaptation in multistable perception: Noise drives when to switch, adaptation determines percept choice. *Journal of vision*, 14(3):19–19, 2014.

Ince R. A. , Senatore R. , Arabzadeh E. , Montani F. , Diamond M. E. , and Panzeri S. . Information-theoretic methods for studying population codes. *Neural Networks*, 23(6):713–727, 2010.

Ince R. A. , Panzeri S. , and Kayser C. . Neural codes formed by small and temporally precise populations in auditory cortex. *Journal of Neuroscience*, 33(46):18277–18287, 2013.

Ito M. and Gilbert C. D. . Attention modulates contextual influences in the primary visual cortex of alert monkeys. *Neuron*, 22(3):593–604, 1999.

Jardri R. and Denève S. . Computational models of hallucinations. In *The neuroscience of hallucinations*, pages 289–313. Springer, 2013.

Jardri R. , Hugdahl K. , Hughes M. , Brunelin J. , Waters F. , Alderson-Day B. , Smailes D. , Sterzer P. , Corlett P. R. , Leptourgos P. , et al. Are hallucinations due to an imbalance between excitatory and inhibitory influences on the brain? *Schizophrenia bulletin*, 42(5):1124–1134, 2016.

Jazayeri M. and Movshon J. A. . Optimal representation of sensory information by neural populations. *Nature neuroscience*, 9(5):690–696, 2006.

Jeanne J. M. , Sharpee T. O. , and Gentner T. Q. . Associative learning enhances population coding by inverting interneuronal correlation patterns. *Neuron*, 78(2):352–363, 2013.

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

Jin J. , Wang Y. , Swadlow H. A. , and Alonso J. M. . Population receptive fields of on and off thalamic inputs to an orientation column in visual cortex. *Nature neuroscience*, 14(2):232–238, 2011.

Josić K. , Shea-Brown E. , Doiron B. , and Rocha J. de la . Stimulus-dependent correlations and population codes. *Neural computation*, 21(10):2774–2804, 2009.

Kahneman D. and Beatty J. . Pupil diameter and load on memory. *Science*, 154(3756):1583–1585, 1966.

Kanitscheider I. , Coen-Cagli R. , Kohn A. , and Pouget A. . Measuring fisher information accurately in correlated neural populations. *PLoS Comput Biol*, 11(6):e1004218, 2015a.

Kanitscheider I. , Coen-Cagli R. , and Pouget A. . Origin of information-limiting noise correlations. *Proceedings of the National Academy of Sciences*, 112(50):E6973–E6982, 2015b.

Kara P. , Reinagel P. , and Reid R. C. . Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron*, 27(3):635–646, 2000.

Keller E. F. and Longino H. E. . *Feminism and science*. Oxford University Press, USA, 1996.

Kelly R. C. , Smith M. A. , Samonds J. M. , Kohn A. , Bonds A. , Movshon J. A. , and Lee T. S. . Comparison of recordings from microelectrode arrays and single electrodes in the visual cortex. *The Journal of neuroscience*, 27(2):261–264, 2007.

Kenet T. , Bibitchkov D. , Tsodyks M. , Grinvald A. , and Arieli A. . Spontaneously emerging cortical representations of visual attributes. *Nature*, 425(6961):954–956, 2003.

Kevles D. . A history of patenting life in the united states with comparative attention to europe and canada: a report to the european group on ethics in science and new technologies. Technical report, Office for Official Publications of the European Communities, 2002.

Knill D. C. and Pouget A. . The bayesian brain: the role of uncertainty in neural coding and computation. *Trends in neurosciences*, 27(12): 712–719, 2004.

Koch C. , Massimini M. , Boly M. , and Tononi G. . Neural correlates of consciousness: progress and problems. *Nature Reviews Neuroscience*, 17(5):307–321, 2016.

Kohn A. and Movshon J. A. . Adaptation changes the direction tuning of macaque mt neurons. *Nature neuroscience*, 7(7):764–772, 2004.

Kohn A. and Smith M. . Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *The Journal of neuroscience*, 25(14):3661–3673, 2005.

Kohn A. , Coen-Cagli R. , Kanitscheider I. , and Pouget A. . Correlations and neuronal population information. *Annual review of neuroscience*, 39:237–256, 2016.

Kosslyn S. M. , Alpert N. M. , Thompson W. L. , Maljkovic V. , Weise S. B. , Chabris C. F. , Hamilton S. E. , Rauch S. L. , and Buonanno F. S. . Visual mental imagery activates topographically organized visual cortex: Pet investigations. *Journal of Cognitive Neuroscience*, 5(3):263–287, 1993.

Kuhn T. S. . *The Structure of Scientific Revolutions*. The University of Chicago Press, 4rd edition, 1962/2012. ISBN 9780226458113.

Lampl I. , Reichova I. , and Ferster D. . Synchronous membrane potential fluctuations in neurons of the cat visual cortex. *Neuron*, 22(2):361–374, 1999.

Latham P. E. and Nirenberg S. . Synergy, redundancy, and independence in population codes, revisited. *Journal of Neuroscience*, 25 (21):5195–5206, 2005.

Lee T. S. and Mumford D. . Hierarchical bayesian inference in the visual cortex. *JOSA A*, 20(7):1434–1448, 2003.

Leopold D. A. . Primary visual cortex: awareness and blindsight. *Annual review of neuroscience*, 35:91–109, 2012.

Li W. and Gilbert C. D. . Global contour saliency and local colinear interactions. *Journal of neurophysiology*, 88(5):2846–2856, 2002.

Li W. , Piëch V. , and Gilbert C. D. . Perceptual learning and top-down influences in primary visual cortex. *Nature neuroscience*, 7 (6):651–657, 2004.

Lin I.-C. , Okun M. , Carandini M. , and Harris K. D. . The nature of shared cortical variability. *Neuron*, 87(3):644–656, 2015.

London M. , Roth A. , Beeren L. , Häusser M. , and Latham P. E. . Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature*, 466(7302):123–127, 2010.

Luczak A. , Barthó P. , Marguet S. L. , Buzsáki G. , and Harris K. D. . Sequential structure of neocortical spontaneous activity in vivo. *Proceedings of the National Academy of Sciences*, 104(1):347–352, 2007.

Luczak A. , Barthó P. , and Harris K. D. . Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*, 62(3):413–425, 2009.

Luczak A. , Bartho P. , and Harris K. D. . Gating of sensory input by spontaneous cortical activity. *The Journal of Neuroscience*, 33(4): 1684–1695, 2013.

Lutz A. , Lachaux J.-P. , Martinerie J. , and Varela F. J. . Guiding the study of brain dynamics by using first-person data: synchrony patterns correlate with ongoing conscious states during a simple visual task. *Proceedings of the national academy of sciences*, 99(3): 1586–1591, 2002.

Ma W. J. , Beck J. M. , Latham P. E. , and Pouget A. . Bayesian inference with probabilistic population codes. *Nature neuroscience*, 9(11):1432–1438, 2006.

MacKay D. M. . Towards an information-flow model of human behaviour. *British Journal of Psychology*, 47(1):30–43, 1956.

Mainen Z. F. and Sejnowski T. J. . Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503, 1995.

Marr D. . Vision: A computational investigation into the human representation and processing of visual information, henry holt and co. *Inc., New York, NY*, 2:4–2, 1982.

McAdams C. and Maunsell J. . Effects of attention on orientation-tuning functions of single neurons in macaque cortical area v4. *The Journal of Neuroscience*, 19(1):431–441, 1999.

McDonnell M. D. and Abbott D. . What is stochastic resonance? definitions, misconceptions, debates, and its relevance to biology. *PLoS Comput Biol*, 5(5):e1000348, 2009.

Mitchell J. F. , Sundberg K. A. , and Reynolds J. H. . Spatial attention decorrelates intrinsic activity fluctuations in macaque area v4. *Neuron*, 63(6):879–888, 2009.

Mochol G. , Hermoso-Mendizabal A. , Sakata S. , Harris K. D. , and Rocha J. de la . Stochastic transitions into silence cause noise correlations in cortical circuits. *Proceedings of the National Academy of Sciences*, 112(11):3529–3534, 2015.

Montani F. , Ince R. A. , Senatore R. , Arabzadeh E. , Diamond M. E. , and Panzeri S. . The impact of high-order interactions on the rate of synchronous discharge and information transmission in somatosensory cortex. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 367(1901):3297–3310, 2009.

Montijn J. S. , Meijer G. T. , Lansink C. S. , and Pennartz C. M. . Population-level neural codes are robust to single-neuron variability from a multidimensional coding perspective. *Cell reports*, 16(9): 2486–2498, 2016.

Moreno-Bote R. . Poisson-like spiking in circuits with probabilistic synapses. *PLoS Comput Biol*, 10(7):e1003522, 2014.

Moreno-Bote R. , Rinzel J. , and Rubin N. . Noise-induced alternations in an attractor network model of perceptual bistability. *Journal of neurophysiology*, 98(3):1125–1139, 2007.

Moreno-Bote R. , Beck J. , Kanitscheider I. , Pitkow X. , Latham P. , and Pouget A. . Information-limiting correlations. *Nature neuroscience*, 17(10):1410–1417, 2014.

Moss F. , Ward L. M. , and Sannita W. G. . Stochastic resonance and sensory information processing: a tutorial and review of application. *Clinical neurophysiology*, 115(2):267–281, 2004.

Motter B. C. . Focal attention produces spatially selective processing in visual cortical areas v1, v2, and v4 in the presence of competing stimuli. *Journal of neurophysiology*, 70(3):909–919, 1993.

Muckli L. . What are we missing here? brain imaging evidence for higher cognitive functions in primary visual cortex v1. *International Journal of Imaging Systems and Technology*, 20(2):131–139, 2010.

Mumford D. . On the computational architecture of the neocortex. *Biological cybernetics*, 65(2):135–145, 1991.

Niell C. M. and Stryker M. P. . Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron*, 65(4):472–479, 2010.

Nir Y. and Tononi G. . Dreaming and the brain: from phenomenology to neurophysiology. *Trends in cognitive sciences*, 14(2):88–100, 2010.

Nirenberg S. and Latham P. E. . Decoding neuronal spike trains: how important are correlations? *Proceedings of the National Academy of Sciences*, 100(12):7348–7353, 2003.

Nirenberg S. , Carcieri S. M. , Jacobs A. L. , and Latham P. E. . Retinal ganglion cells act largely as independent encoders. *Nature*, 411(6838):698–701, 2001.

Notredame C.-E. , Pins D. , Deneve S. , and Jardri R. . What visual illusions teach us about schizophrenia. *Neurovision: Neural bases of binocular vision and coordination and their implications in visual training programs*, page 152, 2015.

Oizumi M. , Ishii T. , Ishibashi K. , Hosoya T. , and Okada M. . Mismatched decoding in the brain. *Journal of Neuroscience*, 30 (13):4815–4826, 2010.

Okun M. , Naim A. , and Lampl I. . The subthreshold relation between cortical local field potential and neuronal firing unveiled by intracellular recordings in awake rats. *Journal of neuroscience*, 30 (12):4440–4448, 2010.

Okun M. , Steinmetz N. A. , Cossell L. , Iacaruso M. F. , Ko H. , Barthó P. , Moore T. , Hofer S. B. , Mrsic-Flogel T. D. , Carandini M. , et al. Diverse coupling of neurons to populations in sensory cortex. *Nature*, 521(7553):511–515, 2015.

Olsen S. R. , Bortone D. S. , Adesnik H. , and Scanziani M. . Gain control by layer six in cortical circuits of vision. *Nature*, 483(7387): 47–52, 2012.

Olshausen B. A. and Field D. J. . Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research*, 37(23): 3311–3325, 1997.

O'Regan J. K. and Noë A. . A sensorimotor account of vision and visual consciousness. *Behavioral and brain sciences*, 24(05):939–973, 2001.

O'Neill J. , Pleydell-Bouverie B. , Dupret D. , and Csicsvari J. . Play it again: reactivation of waking experience and memory. *Trends in neurosciences*, 33(5):220–229, 2010.

Pachitariu M. , Lyamzin D. R. , Sahani M. , and Lesica N. A. . State-dependent population coding in primary auditory cortex. *The Journal of Neuroscience*, 35(5):2058–2073, 2015.

Panzeri S. , Schultz S. R. , Treves A. , and Rolls E. T. . Correlations and the encoding of information in the nervous system. *Proceedings of the Royal Society of London B: Biological Sciences*, 266(1423): 1001–1012, 1999.

Panzeri S. , Petersen R. S. , Schultz S. R. , Lebedev M. , and Diamond M. E. . The role of spike timing in the coding of stimulus location in rat somatosensory cortex. *Neuron*, 29(3):769–777, 2001.

Panzeri S. , Senatore R. , Montemurro M. A. , and Petersen R. S. . Correcting for the sampling bias problem in spike train information measures. *Journal of neurophysiology*, 98(3):1064–1072, 2007.

Panzeri S. , Macke J. H. , Gross J. , and Kayser C. . Neural population coding: combining insights from microscopic and mass signals. *Trends in cognitive sciences*, 19(3):162–172, 2015.

Panzeri S. , Harvey C. D. , Piasini E. , Latham P. E. , and Fellin T. . Cracking the neural code for sensory perception by combining statistics, intervention, and behavior. *Neuron*, 93(3):491–507, 2017.

Pestre D. . *Science, argent et politique: un essai d'interprétation*. Editions Quae, 2003.

Pestre D. . *Ciència, diners i política*. 2008.

Petersen C. C. and Sakmann B. . Functionally independent columns of rat somatosensory barrel cortex revealed with voltage-sensitive dye imaging. *Journal of Neuroscience*, 21(21):8435–8446, 2001.

Petersen C. C. , Hahn T. T. , Mehta M. , Grinvald A. , and Sakmann B. . Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex. *Proceedings of the National Academy of Sciences*, 100(23):13638–13643, 2003.

Pitkow X. , Liu S. , Angelaki D. E. , DeAngelis G. C. , and Pouget A. . How can single sensory neurons predict behavior? *Neuron*, 87(2): 411–423, 2015.

Pola G. , Thiele A. , Hoffmann K. , and Panzeri S. . An exact method to quantify the information transmitted by different mechanisms of correlational coding. *Network-Computation in Neural Systems*, 14 (1):35–60, 2003.

Polack P.-O. , Friedman J. , and Golshani P. . Cellular mechanisms of brain state-dependent gain modulation in visual cortex. *Nature neuroscience*, 16(9):1331–1339, 2013.

Ponce-Alvarez A. , Thiele A. , Albright T. D. , Stoner G. R. , and Deco G. . Stimulus-dependent variability and noise correlations in cortical mt neurons. *Proceedings of the National Academy of Sciences*, 110 (32):13162–13167, 2013.

Popper K. R. . *The Logic of Scientific Discovery*. Routledge classics. Routledge, 2 edition, Mar. 1959. ISBN 0415278449.

Pouget A. , Deneve S. , Ducom J. , and Latham P. . Narrow versus wide tuning curves: What's best for a population code? *Neural Computation*, 11(1):85–90, 1999.

Poulet J. F. and Petersen C. C. . Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature*, 454(7206):881–885, 2008.

Priebe N. J. and Ferster D. . Mechanisms of neuronal computation in mammalian visual cortex. *Neuron*, 75(2):194–208, 2012.

Rabinowitz N. C. , Goris R. L. , Cohen M. , and Simoncelli E. P. . Attention stabilizes the shared gain of v4 populations. *Elife*, 4: e08998, 2015.

Ramachandran V. S. and Hubbard E. M. . Psychophysical investigations into the neural basis of synaesthesia. *Proceedings of the Royal Society of London B: Biological Sciences*, 268(1470):979–983, 2001a.

Ramachandran V. S. and Hubbard E. M. . Synaesthesia–a window into perception, thought and language. *Journal of consciousness studies*, 8(12):3–34, 2001b.

Ramachandran V. S. and Rogers-Ramachandran D. . Synaesthesia in phantom limbs induced with mirrors. *Proceedings of the Royal Society of London B: Biological Sciences*, 263(1369):377–386, 1996.

Rao R. P. and Ballard D. H. . Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1):79–87, 1999.

Rauss K. , Schwartz S. , and Pourtois G. . Top-down effects on early visual processing in humans: A predictive coding framework. *Neuroscience & Biobehavioral Reviews*, 35(5):1237–1253, 2011.

Reid R. C. and Alonso J.-M. . Specificity of monosynaptic connections from thalamus to visual cortex. *Nature*, 378(6554):281, 1995.

Reimer J. , Froudarakis E. , Cadwell C. R. , Yatsenko D. , Denfield G. H. , and Tolias A. S. . Pupil fluctuations track fast switching of cortical states during quiet wakefulness. *Neuron*, 84(2):355–362, 2014.

Renart A. and Machens C. K. . Variability in neural activity and behavior. *Current opinion in neurobiology*, 25:211–220, 2014.

Renart A. and Rossum M. C. van . Transmission of population-coded information. *Neural computation*, 24(2):391–407, 2012.

Renart A. , De La Rocha J. , Bartho P. , Hollender L. , Parga N. , Reyes A. , and Harris K. D. . The asynchronous state in cortical circuits. *science*, 327(5965):587–590, 2010.

Reynolds J. H. , Pasternak T. , and Desimone R. . Attention increases sensitivity of v4 neurons. *Neuron*, 26(3):703–714, 2000.

Ringach D. L. , Hawken M. J. , and Shapley R. . Dynamics of orientation tuning in macaque primary visual cortex. *Nature*, 387(6630): 281, 1997.

Rosenbaum R. and Josić K. . Mechanisms that modulate the transfer of spiking correlations. *Neural computation*, 23(5):1261–1305, 2011.

Ruff D. A. and Cohen M. R. . Attention can either increase or decrease spike count correlations in visual cortex. *Nature neuroscience*, 17 (11):1591–1597, 2014.

Sakata S. and Harris K. D. . Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex. *Neuron*, 64 (3):404–418, 2009.

Salin P.-A. and Bullier J. . Corticocortical connections in the visual system: structure and function. *Physiological reviews*, 75(1):107–155, 1995.

Salinas E. and Abbott L. . Vector reconstruction from firing rates. *Journal of computational neuroscience*, 1(1-2):89–107, 1994.

Salinas E. and Sejnowski T. J. . Impact of correlated synaptic input on output firing rate and variability in simple neuronal models. *Journal of neuroscience*, 20(16):6193–6209, 2000.

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

Sanchez-Vives M. V. and Slater M. . From presence to consciousness through virtual reality. *Nature Reviews Neuroscience*, 6(4):332–339, 2005.

Sanchez-Vives M. V. , Spanlang B. , Frisoli A. , Bergamasco M. , and Slater M. . Virtual hand illusion induced by visuomotor correlations. *PloS one*, 5(4):e10381, 2010.

Sanger T. D. . Probability density estimation for the interpretation of neural population codes. *Journal of Neurophysiology*, 76(4):2790–2793, 1996.

Sato T. K. , Häusser M. , and Carandini M. . Distal connectivity causes summation and division across mouse visual cortex. *Nature neuroscience*, 17(1):30–32, 2014.

Schneider F. , Kallis G. , and Martinez-Alier J. . Crisis or opportunity? economic degrowth for social equity and ecological sustainability. introduction to this special issue. *Journal of cleaner production*, 18 (6):511–518, 2010.

Schneidman E. , Berry M. , Segev R. , and Bialek W. . Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, 2006.

Schölvinck M. L. , Saleem A. B. , Benucci A. , Harris K. D. , and Carandini M. . Cortical state determines global variability and correlations in visual cortex. *The Journal of Neuroscience*, 35(1):170–178, 2015.

Schoups A. , Vogels R. , Qian N. , and Orban G. . Practising orientation identification improves orientation coding in v1 neurons. *Nature*, 412(6846):549–553, 2001.

Seriès P. , Latham P. , and Pouget A. . Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations. *Nature neuroscience*, 7(10):1129–1135, 2004.

Seth A. K. . A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cognitive neuroscience*, 5(2):97–118, 2014.

Seung H. S. and Sompolinsky H. . Simple models for reading neuronal population codes. *Proceedings of the National Academy of Sciences*, 90(22):10749–10753, 1993.

Shadlen M. N. and Newsome W. T. . The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *The Journal of Neuroscience*, 18(10):3870–3896, 1998. URL `http://www.jneurosci.org/content/18/10/3870.abstract`.

Shadlen M. N. , Britten K. H. , Newsome W. T. , and Movshon J. A. . A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *Journal of Neuroscience*, 16 (4):1486–1510, 1996.

Shamir M. and Sompolinsky H. . Nonlinear population codes. *Neural computation*, 16(6):1105–1136, 2004.

Shamir M. and Sompolinsky H. . Implications of neuronal diversity on population coding. *Neural computation*, 18(8):1951–1986, 2006.

Shannon C. E. . Communication theory of secrecy systems. *Bell Labs Technical Journal*, 28(4):656–715, 1949.

Shea-Brown E. , Josić K. , La Rocha J. de , and Doiron B. . Correlation and synchrony transfer in integrate-and-fire neurons: basic

properties and consequences for coding. *Physical review letters*, 100 (10):108102, 2008.

Sillito A. . The contribution of inhibitory mechanisms to the receptive field properties of neurones in the striate cortex of the cat. *The Journal of physiology*, 250(2):305, 1975.

Singer W. . Neuronal synchrony: a versatile code for the definition of relations? *Neuron*, 24(1):49–65, 1999.

Slater M. , Spanlang B. , Sanchez-Vives M. V. , and Blanke O. . First person experience of body transfer in virtual reality. *PloS one*, 5(5): e10564, 2010.

Smith F. W. and Muckli L. . Nonstimulated early visual areas carry information about surrounding context. *Proceedings of the National Academy of Sciences*, 107(46):20099–20103, 2010.

Smith M. and Kohn A. . Spatial and temporal scales of neuronal correlation in primary visual cortex. *The Journal of Neuroscience*, 28(48):12591–12603, 2008.

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

Sompolinsky H. , Yoon H. , Kang K. , and Shamir M. . Population coding in neuronal systems with correlated noise. *Physical Review E*, 64(5):051904, 2001.

Song S. , Sjöström P. J. , Reigl M. , Nelson S. , and Chklovskii D. B. . Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol*, 3(3):e68, 2005.

Sporns O. , Tononi G. , and Edelman G. M. . Modeling perceptual grouping and figure-ground segregation by means of active reentrant connections. *Proceedings of the National Academy of Sciences*, 88 (1):129–133, 1991.

Spratling M. W. . Predictive coding as a model of response properties in cortical area v1. *Journal of Neuroscience*, 30(9):3531–3543, 2010.

Steriade M. and Timofeev I. . Neuronal plasticity in thalamocortical networks during sleep and waking oscillations. *Neuron*, 37(4):563–576, 2003.

Steriade M. , McCormick D. A. , and Sejnowski T. J. . Thalamocortical oscillations in the sleeping and aroused brain. *SCIENCE-NEW YORK THEN WASHINGTON-*, 262:679–679, 1993.

Stern E. A. , Kincaid A. E. , and Wilson C. J. . Spontaneous subthreshold membrane potential fluctuations and action potential variability of rat corticostriatal and striatal neurons in vivo. *Journal of Neurophysiology*, 77(4):1697–1715, 1997.

Stevens C. F. and Zador A. M. . Input synchrony and the irregular firing of cortical neurons. *Nature neuroscience*, 1(3):210–217, 1998.

Stiglitz J. E. and Chemla P. . *La grande désillusion*, volume 108. Fayard Paris, 2002.

Stokes M. , Thompson R. , Cusack R. , and Duncan J. . Top-down activation of shape-specific population codes in visual cortex during mental imagery. *Journal of Neuroscience*, 29(5):1565–1572, 2009.

Summerfield C. and De Lange F. P. . Expectation in perceptual decision making: neural and computational mechanisms. *Nature Reviews Neuroscience*, 15(11):745–756, 2014.

Suzuki K. , Garfinkel S. N. , Critchley H. D. , and Seth A. K. . Multisensory integration across exteroceptive and interoceptive domains modulates self-experience in the rubber-hand illusion. *Neuropsychologia*, 51(13):2909–2917, 2013.

Tan A. Y. , Chen Y. , Scholl B. , Seidemann E. , and Priebe N. J. . Sensory stimulation shifts visual cortex from synchronous to asynchronous states. *Nature*, 509(7499):226–229, 2014.

Tchumatchenko T. , Malyshev A. , Geisel T. , Volgushev M. , and Wolf F. . Correlations and synchrony in threshold neuron models. *Physical Review Letters*, 104(5):058102, 2010.

Terlouw J. P. and Vogelaar M. G. R. . *Kapteyn Package, version 2.3*. Kapteyn Astronomical Institute, Groningen, Mar. 2015. Available from `http://www.astro.rug.nl/software/kapteyn/`.

Thiele A. , Pooresmaeili A. , Delicato L. S. , Herrero J. L. , and Roelfsema P. R. . Additive effects of attention and stimulus contrast in primary visual cortex. *Cerebral Cortex*, 19(12):2970–2981, 2009.

Thompson E. and Stapleton M. . Making sense of sense-making: Reflections on enactive and extended mind theories. *Topoi*, 28(1): 23–30, 2009.

Thompson E. and Varela F. J. . Radical embodiment: neural dynamics and consciousness. *Trends in cognitive sciences*, 5(10):418–425, 2001.

Thompson E. , Palacios A. , and Varela F. J. . Ways of coloring: Comparative color vision as a case study for cognitive science. *Behavioral and Brain Sciences*, 15(01):1–26, 1992.

Thomson A. M. and Bannister A. P. . Interlaminar connections in the neocortex. *Cerebral cortex*, 13(1):5–14, 2003.

Tolhurst D. J. , Movshon J. A. , and Dean A. F. . The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision research*, 23(8):775–785, 1983.

Tong F. . Primary visual cortex and visual awareness. *Nature Reviews Neuroscience*, 4(3):219–229, 2003.

Tononi G. , Boly M. , Massimini M. , and Koch C. . Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience*, 17(7):450–461, 2016.

Treisman A. . The binding problem. *Current opinion in neurobiology*, 6(2):171–178, 1996.

Treue S. and Maunsell J. H. . Attentional modulation of visual motion processing in cortical areas mt and mst. *Nature*, 382(6591):539, 1996.

Treue S. and Trujillo J. C. M. . Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, 399(6736): 575–579, 1999.

Tsakiris M. and Haggard P. . The rubber hand illusion revisited: visuotactile integration and self-attribution. *Journal of Experimental Psychology: Human Perception and Performance*, 31(1):80, 2005.

Tsodyks M. , Kenet T. , Grinvald A. , and Arieli A. . Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science*, 286(5446):1943–1946, 1999.

Tu Y. . The discovery of artemisinin (qinghaosu) and gifts from chinese medicine. *Nature medicine*, 17(10):1217–1220, 2011.

Urbanczik R. and Senn W. . Reinforcement learning in populations of spiking neurons. *Nat Neurosci*, 12(3):250–252, 2009.

Steveninck R. R. d. R. van , Lewen G. D. , Strong S. P. , Koberle R. , and Bialek W. . Reproducibility and variability in neural spike trains. *Science*, 275(5307):1805–1808, 1997.

Vreeswijk C. van and Sompolinsky H. . Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293): 1724, 1996.

VanRullen R. , Guyonneau R. , and Thorpe S. J. . Spike times make sense. *Trends in neurosciences*, 28(1):1–4, 2005.

Varela F. J. . Neurophenomenology: A methodological remedy for the hard problem. *Journal of consciousness studies*, 3(4):330–349, 1996.

Varela F. J. and Shear J. . First-person methodologies: What, why, how. *Journal of Consciousness studies*, 6(2-3):1–14, 1999.

Vinck M. , Batista-Brito R. , Knoblich U. , and Cardin J. A. . Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding. *Neuron*, 86(3):740–754, 2015.

Von Neumann J. . Probabilistic logics and the synthesis of reliable organisms from unreliable components. *Automata studies*, 34:43–98, 1956.

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

Warzecha A.-K. and Egelhaaf M. . Variability in spike trains during constant and dynamic stimulation. *Science*, 283(5409):1927–1930, 1999.

Weiskrantz L. . Blindsight revisited. *Current opinion in neurobiology*, 6(2):215–220, 1996.

Williford T. and Maunsell J. H. . Effects of spatial attention on contrast response functions in macaque area v4. *Journal of neurophysiology*, 96(1):40–54, 2006.

Wilson N. R. , Runyan C. A. , Wang F. L. , and Sur M. . Division and subtraction by distinct cortical inhibitory networks in vivo. *Nature*, 488(7411):343–348, 2012a.

Wilson N. R. , Runyan C. A. , Wang F. L. , and Sur M. . Division and subtraction by distinct cortical inhibitory networks in vivo. *Nature*, 488(7411):343–348, 2012b.

Xing D. , Ringach D. L. , Hawken M. J. , and Shapley R. M. . Untuned suppression makes a major contribution to the enhancement of orientation selectivity in macaque v1. *Journal of Neuroscience*, 31(44):15972–15982, 2011.

Yoshimura Y. and Callaway E. M. . Fine-scale specificity of cortical networks depends on inhibitory cell type and connectivity. *Nature neuroscience*, 8(11):1552–1559, 2005.

Zhang K. and Sejnowski T. . Neuronal tuning: To sharpen or broaden? *Neural Computation*, 11(1):75–84, 1999.

Ziv Y. , Burns L. D. , Cocker E. D. , Hamel E. O. , Ghosh K. K. , Kitch L. J. , El Gamal A. , and Schnitzer M. J. . Long-term dynamics of ca1 hippocampal place codes. *Nature neuroscience*, 16(3):264–266, 2013.

Zohary E. , Shadlen M. N. , and Newsome W. T. . Correlated neuronal discharge rate and its implications for psychophysical performance. 1994.

Zuo Y. , Safaai H. , Notaro G. , Mazzoni A. , Panzeri S. , and Diamond M. E. . Complementary contributions of spike timing and spike rate

to perceptual decisions in rat s1 and s2 cortex. *Current Biology*, 25 (3):357–363, 2015.

Zylberberg J. , Cafaro J. , Turner M. H. , Shea-Brown E. , and Rieke F. . Direction-selective circuits shape noise to ensure a precise population code. *Neuron*, 89(2):369–383, 2016a.

Zylberberg J. , Pouget A. , Latham P. E. , and Shea-Brown E. . Robust information propagation through noisy neural circuits. *arXiv preprint arXiv:1608.05706*, 2016b.