

Causal relationships between frequency bands of extracellular signals in visual cortex revealed by an information theoretic analysis

Michel Besserve · Bernhard Schölkopf ·
Nikos K. Logothetis · Stefano Panzeri

Received: 20 August 2009 / Revised: 27 January 2010 / Accepted: 25 March 2010 / Published online: 16 April 2010
© The Author(s) 2010. This article is published with open access at Springerlink.com

Abstract Characterizing how different cortical rhythms interact and how their interaction changes with sensory stimulation is important to gather insights into how these rhythms are generated and what sensory function they may play. Concepts from information theory, such as Transfer Entropy (TE), offer principled ways to quantify the amount of causation between different frequency bands of the signal recorded from extracellular electrodes; yet these techniques are hard to apply to real data. To address the above issues, in this study we develop a method to compute fast

and reliably the amount of TE from experimental time series of extracellular potentials. The method consisted in adapting efficiently the calculation of TE to analog signals and in providing appropriate sampling bias corrections. We then used this method to quantify the strength and significance of causal interaction between frequency bands of field potentials and spikes recorded from primary visual cortex of anaesthetized macaques, both during spontaneous activity and during binocular presentation of naturalistic color movies. Causal interactions between different frequency bands were prominent when considering the signals at a fine (ms) temporal resolution, and happened with a very short (ms-scale) delay. The interactions were much less prominent and significant at coarser temporal resolutions. At high temporal resolution, we found strong bidirectional causal interactions between gamma-band (40–100 Hz) and slower field potentials when considering signals recorded within a distance of 2 mm. The interactions involving gamma bands signals were stronger during movie presentation than in absence of stimuli, suggesting a strong role of the gamma cycle in processing naturalistic stimuli. Moreover, the phase of gamma oscillations was playing a stronger role than their amplitude in increasing causations with slower field potentials and spikes during stimulation. The dominant direction of causality was mainly found in the direction from MUA or gamma frequency band signals to lower frequency signals, suggesting that hierarchical correlations between lower and higher frequency cortical rhythms are originated by the faster rhythms.

Action Editor: Mark van Rossum

Electronic supplementary material The online version of this article (doi:10.1007/s10827-010-0236-5) contains supplementary material, which is available to authorized users.

This research was supported by the Max Planck Society, by the BMI Project of the Italian Institute of Technology and by the San Paolo Foundation.

M. Besserve (✉) · B. Schölkopf · N. K. Logothetis
Max Planck Institute for Biological Cybernetics,
Spemannstrasse 38, 72076 Tuebingen, Germany
e-mail: michel.besserve@tuebingen.mpg.de

N. K. Logothetis
Division of Imaging Science and Biomedical Engineering,
University of Manchester, Manchester M13 9PT,
United Kingdom

S. Panzeri
Robotics, Brain and Cognitive Sciences Department,
Italian Institute of Technology,
Via Morego 30, 16163 Genova, Italy

Keywords Local field potentials · Transfer Entropy (TE) · Information theory

1 Introduction

A prominent feature of both spontaneous and sensory-evoked cortical activity, revealed by extracellular recordings of Local Field Potentials (LFPs) and spike trains, is the presence of rhythmic activity (Buzsáki 2006; Buzsáki and Draguhn 2004). This rhythmic activity has a complex structure: even within the same recording location and during the same task, fluctuations span a very broad frequency spectrum, ranging from a fraction of a Hz to well over 100 Hz, and these rhythms often interact with each other in a hierarchical fashion (Roopun et al. 2008). The fact that these broadband fluctuations and their interactions, as well as their behavioral correlates, are largely preserved throughout the mammalian evolution has led to suggest that they are supported by universal mechanisms, and that the interplay between different rhythms is crucial to the function of the brain and forms a basis for cortical information processing (Destexhe and Sejnowski 2003; Gray et al. 1989; Llinas and Ribary 1993; Kahana et al. 2001; Bragin et al. 1999; Buzsáki and Draguhn 2004; Buzsáki 2006; Roopun et al. 2008). Understanding which rhythms drives which, and how the causal chain is modulated by the stimulus, is thus important to understand how rhythms are generated and what role they play in sensory function.

The interactions between different rhythms have been mainly studied so far by considering a correlation analysis between features of two rhythms. These studies have revealed a hierarchically organized set of relationships between activity at lower and higher frequencies (Roopun et al. 2008). For example, the phase of slow rhythms (in the theta or delta frequency range) often correlates with the power of the gamma rhythm (Lisman 2005; Canolty et al. 2006; Roopun et al. 2008). However, a problem with a pure correlation analysis is that it cannot tell whether the covariation between rhythms arises from true causal relations between the two rhythms or from other sources. As a result, we do not know whether or not these relationships imply the presence of a leading set of frequencies that drives the others, and if so, which are these leading frequencies.

A more principled and effective approach to establish causal relationships is to use the causality principle formulated by Wiener and Granger (Granger 1969). Using this principle, techniques from information theory, such as TE (Schreiber 2000), can in principle provide measures of the amount of causation between rhythms, which provide meaningful answers even in the presence of strong nonlinearities in the considered signals. If the appropriately applied causal techniques reveal a clearly dominant direction of causality, this

direction can be used to individuate the leading signal. If instead these techniques reveal a similar amount of causation in both directions, their result should be interpreted as individuating the presence of coherency between signals.

Despite their promise, the application of TE to the study of causal interactions from brain recordings have been so far limited, largely because of the technical difficulties in computing any information theoretic quantity from limited samples of neuronal data (Panzeri et al. 2007) and in applying computationally expensive estimation methods on large datasets. The goal of this article is to overcome the estimation problems that have previously limited extensive use of the TE approach to neurophysiology data, and to prove the worth of these techniques by demonstrating the presence and stimulus modulation of causal relationships between rhythms of sensory cortex during naturalistic function. We develop and test computationally efficient mathematical methods for the reliable computation of TE between experimentally recorded oscillatory neural signals, and we then use this approach to investigate which frequency ranges of cortical activity in primary visual cortex (either observed from the same location or from nearby locations) cause each other. To understand how the causal chain of frequency relationships is modulated by the presence of sensory stimuli, we quantify the changes in the amount of causation induced by sensory stimulation compared to spontaneous activity. We consider neural fluctuations and oscillations expressed both at the level of spiking activity and of LFPs, since they express largely independent and complementary aspects of the network activity (Logothetis 2008) and have a largely complementary content in terms of sensory information (Belitski et al. 2008).

This article is organized as follows. We first discuss how to measure causality and introduce the concept of TE; we then consider and address the algorithmic problems arising when computing TE from limited stretches of neurophysiology data; we develop an efficient algorithm for such an estimation; we then apply this algorithm to recordings from primary visual cortex and we compare TE with other methods such as phase coherence analysis; finally we discuss the implications of our findings.

2 Measuring causality

Causality methods compute directional measures of interactions between dynamical systems from their associated time series. This methodology has been established by the pioneering work of Wiener and

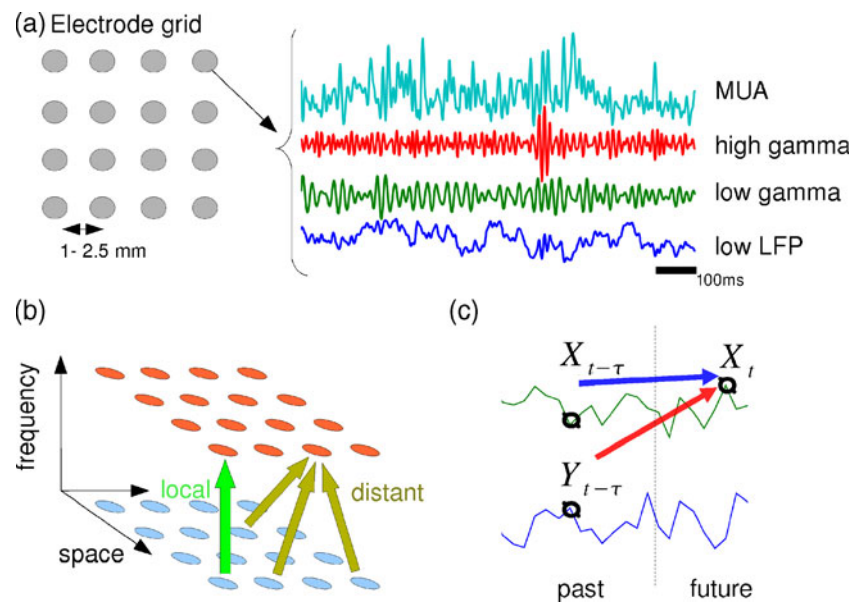


Fig. 1 General scheme of our approach. **(a)** Extracellular potentials were recorded in several sites in V1, electrodes are organized on a grid with interelectrode distance in the 1–2.5 mm range. Extracellular potentials recorded in each site were filtered in several frequency bands. **(b)** Directed interactions were computed between pairs of frequency bands both locally (i.e. from

bands obtained from the same electrode) and from more distant recording sites. **(c)** Principle of causality analysis: Y causes X if the uncertainty on the time course of X is reduced when using the information from Y in the past (a time delay τ before), compared to only using the past of X

Granger (Granger 1969). As illustrated in Fig. 1(c), the definition of causality between two scalar valued time series X and Y observed from systems \mathcal{X} and \mathcal{Y} leans heavily on the idea that the cause occurs before the effect. If there are two time series $\{Y_t\}$ and $\{X_t\}$, and if the knowledge of past values of Y allows a better forecast of the present value of X than the forecast obtained just based on the knowledge of past values of X , then the signal Y is said to be a Granger cause of X . Although the Granger causality principle is general and was formulated originally without any assumption about the linearity or nonlinearity of the systems, practical implementations of measures of Granger causality usually rely heavily on the assumptions of the linearity of the systems and of the interaction between them. This is because the amount of causality is quantified directly from linear multivariate autoregressive models fitted to the two time series (Granger 1969), and statistical testing is often done under the assumption of stationary Gaussian processes. Most previous investigations into causality relied significantly on these assumptions of linearity (Brovelli et al. 2004; Chen et al. 2006; Guo et al. 2008b; Bernasconi et al. 2000; Seth 2005; Seth and Edelman 2007; Roebroek et al. 2005). Several extensions of Granger causality which decompose causations in the frequency domain, such as Directed Transfer Function (Geweke 1982; Kamiński

and Blinowska 1991) and Partial Directed Coherence (Baccalá and Sameshima 2001), though able to lead to interesting results in the quantification of interactions between different brain areas (Bressler et al. 2007; Kayser and Logothetis 2009), also rely on linearity assumptions.

A potential problem with the linear system approach is that neural responses in general, and cortical oscillations in particular, are intrinsically non-linear. For example, the conversion between input rate and oscillation power of network of excitatory and inhibitory neurons is non linear (Brunel and Wang 2003), and so are interactions between rhythms (Chavez et al. 2003; da Silva et al. 1989). Although extensions of Granger causality have been proposed to allow non-linearities in the models of the dynamical systems (Ancona et al. 2004; Ancona and Stramaglia 2006; Marinazzo et al. 2006), the most general way to introduce arbitrary non-linearities in the Granger causality principle is to use information theoretic measures of causality, of which TE (Schreiber 2000) is the most known one. TE has been already applied to intracranial electroencephalography recordings in epileptic patients (Chavez et al. 2003), to study single unit spiking activity in the auditory pathway (Gourévitch and Eggermont 2007) and to functional Magnetic Resonance Imaging data (Hinrichs et al. 2006). In the following, we extend its use to

quantify the causal relationships between rhythms generated in sensory cortex during spontaneous activity and during naturalistic sensory stimulation, and we consider the computational problems arising when computing these quantities from limited datasets of neural data.

3 Calculation of TE

Transfer Entropy (TE) is a measure of causality that stems from information theory and relies on the concepts of entropy and mutual information, which for completeness will be briefly reviewed next.

3.1 Background

Given a discrete random variable X with probability distribution $p(x)$, following Shannon (1948) we define the entropy of X as

$$H(X) = - \sum_{x \in \mathcal{X}} p(x) \log_2(p(x)) \quad (1)$$

where the summation over x stands for the sum over all possible values of X . H is a positive quantity that quantifies the uncertainty (or variability) of the random variable X . The conditional entropy of X given another discrete random variable Y is

$$H(X|Y) = - \sum_{y \in \mathcal{Y}} p(y) \sum_{x \in \mathcal{X}} p(x|y) \log_2(p(x|y)) \quad (2)$$

Then mutual information between X and Y is defined as $I(X; Y) = H(X) - H(X|Y)$. $I(X; Y)$ quantifies the reduction of uncertainty about X gained by the knowledge of Y . If X and Y are independent then $I(X; Y) = 0$, otherwise mutual information is strictly positive.

3.2 Transfer Entropy (TE)

We consider the time series of two simultaneously recorded neurophysiological signals X and Y . The time series of the values of the two signals simultaneously recorded at each sampling time is denoted by (X_t, Y_t) . We assume that this joint time series can be represented by a discrete stationary Markov process of order k . This means that the probability distribution of the signals at time t given the past depends only of the vectors composed of the k previous samples $X_t^{(k)} = (X_{t-1}, \dots, X_{t-k})$ and $Y_t^{(k)} = (Y_{t-1}, \dots, Y_{t-k})$, and not on the

values of the signals at earlier times. Then, following (Schreiber 2000), the TE from Y to X is defined as:

$$T_{Y \rightarrow X} = H(X_t | X_t^{(k)}) - H(X_t | X_t^{(k)}, Y_t^{(k)}) \quad (3)$$

TE is the mutual information between the present value of X and the past values of Y , conditioned on the knowledge of past values of X . As such, TE quantifies the reduction of uncertainty in X_t when the knowledge of the past of Y is added to the past of X itself. A non-zero value for $T_{Y \rightarrow X}$ can be interpreted as “the past values of Y have an effect on the present value of X ”. The conditioning on past values of X makes TE asymmetric with respect to changes between X and Y . This asymmetry of TE is a crucial feature to establish the directionality of information flow between two systems. However, as reported in Schreiber (2000) a direct quantitative comparison of the flow of information in both directions should be avoided when the two systems have fundamentally different characteristics.

In practice, although it is reasonable to model the time series of the neurophysiological signals as Markov processes, the order of the Markov process (i.e. the number k of past delays that influence the current neural response and thus must be considered when computing Eq. (3)) is not known a priori, and must be determined empirically by balancing the following conflicting requirements. On the one hand, it would be desirable to use a large value of k in order to include all possible dependencies between the neural responses. On the other hand, conditioning on many past values of the neurophysiological signals makes it very difficult to sample the probabilities entering Eq. (3), the number of samples needed increasing exponentially with k . Following (Schreiber 2000), the empirical solution we chose is to use only one time delay, but to take it at a variable delay τ , which is the same for both time series X and Y , and whose value is varied parametrically within a range to test the potential effect of causations at different delays. The expression of TE is then

$$T_{Y \rightarrow X} = H(X_t | X_{t-\tau}) - H(X_t | X_{t-\tau}, Y_{t-\tau}) \quad (4)$$

Importantly, choosing the same delay for both time series X and Y requires they vary at comparable time scales, this point will be ensured by the preprocessing described in Section 4.1. We also checked whether the conditioning of TE on a single time delay was sufficient and not inducing false causality values, as follows. We computed TE values when including an additional time delay 2τ , as follows:

$$T_{Y \rightarrow X} = H(X_t | X_{t-\tau}, X_{t-2\tau}) - H(X_t | X_{t-\tau}, X_{t-2\tau}, Y_{t-\tau})$$

We found (results not shown) that the magnitude of TE, the patterns of significance of TE values, and the differences between spontaneous and visual stimulation (see next section) remained similar to those obtained when conditioning on only one previous time point (as in Eq. (4)). This consistency lends credibility to our findings.

Following Gourévitch and Eggermont (2007), we will quantify causal relationships using a Normalized Transfer Entropy value (NTE) defined as the proportion of reduction of entropy compared to the reference entropy $H(X_t|X_{t-\tau})$:

$$NTE_{Y \rightarrow X}(\tau) = \frac{T_{Y \rightarrow X}}{H(X_t|X_{t-\tau})} = 1 - \frac{H(X_t|X_{t-\tau}, Y_{t-\tau})}{H(X_t|X_{t-\tau})} \quad (5)$$

This normalization is very useful because it enables to compare information flows independently of the degree of dependence between X_t and its past (Gourévitch and Eggermont 2007). In that way, it contributes to normalize the measure with respect to the different degree of complexity of the X and Y signals.

3.3 Estimation of TE

We wish to estimate TE between two time series of extracellular potentials, which (unlike spike trains) are analog variables. Calculations of TE between analog variables is possible by using approximations of differential entropies using Kernel density estimation (KDE) or nearest neighbor distance estimation (NND) (Schreiber 2000; Kaiser and Schreiber 2002; Chavez et al. 2003; Victor 2002). However, these techniques require a large amount of neural data to converge unless the underlying probability distributions are sufficiently smooth (Victor 2002; Nelken et al. 2005). Moreover, KDE and NND techniques are computationally expensive, and their use would make it practically unfeasible to analyze such an extensive dataset (containing hours of multichannel recordings from several tens of recordings sites) in a reasonable amount of time on an up-to-date server.

To overcome these difficulties, here we developed a simpler and data robust approach to the estimation of TEs from analog signals. This approach, which is based on a recently developed and successful approach to estimating mutual information between external stimuli and LFPs and EEGs (Belitski et al. 2008; Montemurro et al. 2008; Magri et al. 2009; Kayser et al. 2009), consists in first discretizing the considered analog neural signals into a given number of bins R ; then computing a plug-in estimate of TE (denoted by $\hat{T}_{Y \rightarrow X}$) obtained by

simply plugging the experimentally measured discrete probabilities into the TE equations; and by correcting for the bias of the plugin TE estimate due to limited sampling.

Several strategies are possible for the quantization of the analog signals (see Hlavackova-Schindler et al. 2007 for a review). We used equipopulated binning of the marginal distribution of each signal, because it allows a good sampling of the conditional probabilities of neural signals. Since it equalized the entropies $H(Y)$ and $H(X)$ of the two signals, it is also useful to reduce potential problems arising from the different degrees of complexity of the X and Y signals (Quiroga et al. 2000; Stam and van Dijk 2002). In all the following study we used a discretization into five bins ($R = 5$), because we previously found (Magri et al. 2009) that this discretization is the coarser one which is sufficient to approximate with very high precision the mutual information that the LFPs in the dataset we analyze below (see Section 4.1) carry about the visual features in the movie. Consistently with these previous findings, here we found that increasing the number of bins did not change appreciably the TE and NTE values (results not shown). The estimation and subtraction of the bias due to limited sampling was performed by means of a generalization to specific case of TE of a “shuffling” bias correction procedure originally developed in Montemurro et al. (2007) and Panzeri et al. (2007) for the case of mutual information between stimuli and responses. Details on this bias correction procedure which as we will see greatly increases the convergence of the TE estimation with the sample size are provided in the Appendix. We implemented the required entropy calculations using the Information Breakdown Toolbox (Magri et al. 2009).

We compared the run time complexity of our approach with respect to a KDE with a rectangular window (as proposed in Schreiber 2000). On a personal laptop equipped with an Intel Core 2 duo processor (2.4 GHz), time for computing a single TE value on 50,000 data points takes 50 ms with our approach and 10 s using KDE techniques. The reduction in run time complexity with our approach is thus crucial for the extensive use of TE measures on a large dataset.

4 Computations of TE between different frequency components of the extracellular signals recorded in primary visual cortex

After having defined TE and NTE and having outlined the computation procedure, we now apply it to real

data with the aim of evaluating its convergence properties and consider which interactions between frequency bands it reveals.

4.1 Neurophysiological data

We begin by describing the neurophysiological recordings. These data were described before (Belitski et al. 2008) in the context of the analysis of how different frequencies of neural activity encode naturalistic stimuli. In brief, we recorded with an array of extracellular tungsten electrodes from primary visual cortex of four macaques (monkeys A98, D04, G97, C98) anesthetized with opiates. All procedures were approved by the local authorities (Regierungspraesidium) and were in full compliance with the guidelines of the European Community (EUVD 86/609/EEC) for the care and use of laboratory animals. The electrodes were arranged in a 4×4 square matrix (interelectrode spacing varied from 1 mm to 2.5 mm) and introduced for each experimental session into the cortex through the overlying dura mater by a microdrive array system (Thomas Recording). We refer to the study by Eckhorn and Thomas (1993) for more details. Electrode tips were typically (but not always) positioned in the upper or middle cortical layers. The impedance of the electrode varied from 300 k Ω to 800 k Ω . For each recording site, the extracellular signals were sampled at 20835 Hz and collected in response to either a binocularly presented (3.5–6-min long, depending on recording session) naturalistic color movie or during a 5 min period of spontaneous activity (that is, in absence of visual stimulation). In each session, between 5 and 30 repetitions (trials) of stimulation with the same movie were available, and 5–10 spontaneous-activity trials were also available. Each recording site corresponded to a well-defined V1 visual receptive field within the field of movie projection. From each electrode, we extracted both spiking activity and LFPs as follows.

4.1.1 Extraction of multi-unit-activity (MUA)

Multi-unit-activity (MUA) was extracted by band passing the extracellular signal in the 1,000–3,000 Hz range and extracting the envelope of the resulting oscillations (Gail et al. 2004; Logothetis et al. 2001). The resulting quantity is known to represent a weighted average of the extracellular spikes of all neurons within a sphere of approximately 140–300 μm around the tip of the electrode (Logothetis 2003), and can be thus taken as a good measure of the local spiking activity.

4.1.2 Extraction of LFPs

LFPs were extracted by bandpassing the extracellular signal in the 1–150 Hz frequency range. LFPs obtained in this way reflect the fluctuations in the input and the intracortical processing of the local cortical network, including the overall effect of population synaptic potentials (Mitzdorf 1987; Juergens et al. 1999) and other types of slow activity, such as spike afterpotentials and voltage-dependent membrane oscillations (Harada and Takahashi 1983; Kamondi et al. 1998; Buzsáki 2002; Logothetis 2003).

LFPs were subsequently further decomposed into frequency bands widely used in the literature (Buzsáki 2006). In this study we focused on the following bands: the theta (4–8 Hz) LFP bands, because it is very informative about naturalistic stimuli in both primary visual cortex (Belitski et al. 2008; Montemurro et al. 2008) and auditory cortex (Kayser et al. 2009); the low (40–60 Hz) and high (60–120 Hz) gamma bands, which are also strongly modulated by visual stimuli (Belitski et al. 2008; Berens et al. 2008) and are thought to reflect the rapid cycles of excitation and inhibition in local recurrent networks (Brunel and Wang 2003); and the beta (24–40 Hz) band, which in visual cortex activity has a relatively strong power and has been hypothesized to be mainly driven by stimulus-independent neuromodulatory processes (Belitski et al. 2008; Logothetis 2008).

4.2 Temporal resolutions considered in the analysis

Since signals in different bands may vary on very different time scales, a direct causality analysis of interactions between them may not always be appropriate, especially because we compute TE with the same time delay for both time series under analysis (see Section 3.3), thereby assuming a similar time scale for the dynamics of each series.

The requirement of similarity of time scales of changes in the two signals is partly supported in this dataset by our previous finding (Belitski et al. 2008) that gamma and MUA power variations due to stimulation are as slow as the time variations of the low frequency LFPs. In fact, we found that in this dataset MUA carries most information and has the highest power in the low frequency (1–4 Hz) range (Belitski et al. 2008), possibly reflecting network entrainment to slow regularities in the naturalistic stimulus (Mazzoni et al. 2008). Moreover, the envelope of gamma-range LFPs varies slowly too and covaries with MUA (Belitski et al. 2008).

However, to better control for possible effects due to differences in the time scales of the signals, and to study parametrically the temporal resolution at which frequency bands may cause one another, we decided to low pass all signals to a cut-off frequency F , which was varied parametrically and set the temporal resolution at which causal relationships were considered. The TE analysis was thus performed at 3 different time scales corresponding to low pass cutting frequencies F of 8, 30 and 100 Hz and down-sampled at 80, 300 and 1,000 Hz respectively (see Table 1).

Additionally to the bandpassed LFPs described above, at each considered time scale we computed a broadband signal called “low LFP”, for which the bandwidth corresponds exactly to the cutting frequency of the low pass filter. Because the power spectrum of LFP signals decays with increasing frequency, activity in this band is dominated by lower frequencies. We then computed also the LFP partitioned in the frequency bands defined above. The frequency ranges above the cutting frequency of the low pass filter were rectified. Since the frequency domain of the amplitude of an oscillation can range approximately from 0 to half its bandwidth, rectified oscillations were included in the analysis only if their bandwidth was sufficiently large to reach the cutting frequency of the low pass filter (otherwise the changes in amplitude are too slow for the considered temporal resolution). This explains for example why no rectified low gamma activity is computed at middle temporal resolution (see Table 1).

4.3 Convergence of the estimation of TE with sample size

To ensure our estimation is reliable and not affected by a limited sampling bias, we first studied the convergence properties of NTE for our neurophysiology data.

Table 1 Frequency bands used for causality analysis at three different time resolution

Resolution (F)	Low (8 Hz)	Middle (30 Hz)	High (100 Hz)
Low LFP	[0 8 Hz]	[0 30 Hz]	[0 100 Hz]
Theta	[4 8 Hz]		
Beta	[24 40 Hz] ^R	[14 30 Hz]	
Low gamma	[40 60 Hz] ^R		[40 60 Hz]
High gamma	[60 120 Hz] ^R	[60 120 Hz] ^R	[60 100 Hz]
MUA	[1 3 k Hz] ^R	[1 3 k Hz] ^R	[1 3 k Hz] ^R

Depending on the temporal resolution, each band can be rectified (indicated by an “R” exponent, i.e. absolute value is computed and low pass filtered at the cutoff frequency F indicated above the column)

We estimated NTE in an experiment in one monkey (A98) both during stimulation with a 260 s long movie or during 300 s of stimulus-free (spontaneous) activity. For both stimulation conditions, we used 5 trials, and we computed NTE estimates for an increasing number of samples by using only a fraction of data points.

Results are reported in Fig. 2 for the middle-resolution case (frequency cutoff $F = 30$ Hz). We first tested the performance of the shuffled bias correction technique (developed and described in the Appendix) in removing the upward bias (Panzeri et al. 2007) of information calculations due to limited sampling. A comparison of the convergence with data size of the plug-in estimate (no bias correction) with the bias corrected estimate (Fig. 2) shows that our bias correction clearly helps in reducing the bias in both stimulation conditions: the NTE values are clearly reduced for a number of samples below 10,000, and the bias-corrected NTE reaches a plateau earlier than the non-corrected NTE, especially during spontaneous activity. The number of samples necessary to reach this plateau was approximately 10,000 samples.

A know source of bias for TE and NTE is the correlation between samples close to each other in time (Theiler 1986). To test for this effect, we recomputed bias-corrected NTEs by taking randomly spaced samples rather than continuous data points from the beginning of recordings. We found that the random spacing estimation needed less data points to converge than the procedure taking consecutive data points, although both procedures converged to the same value when using the whole dataset. An even simpler procedure consisting of estimating bias-corrected TE from a down-sampled time series (decimated by a factor 5) to reduce time correlation between samples held a faster convergence (Fig. 2). It should be noted that this fast convergence of the downsampled estimate was obtained only using the bias-corrected estimate in combination with the downsampling, and was not so fast when using either technique in isolation (results not shown). It is also interesting to note that Fig. 2 shows that NTE estimate converge faster with spontaneous than with movie data. In our view, the reason is that the stationarity condition is likely to be more severely violated during movie stimulation (because the stimulus drives larger, non-stochastic changes in the network response).

Results were qualitatively similar when considering high and low temporal resolution, and also when considering activity from other frequency bands (data not shown). In sum, all estimators converged to the same

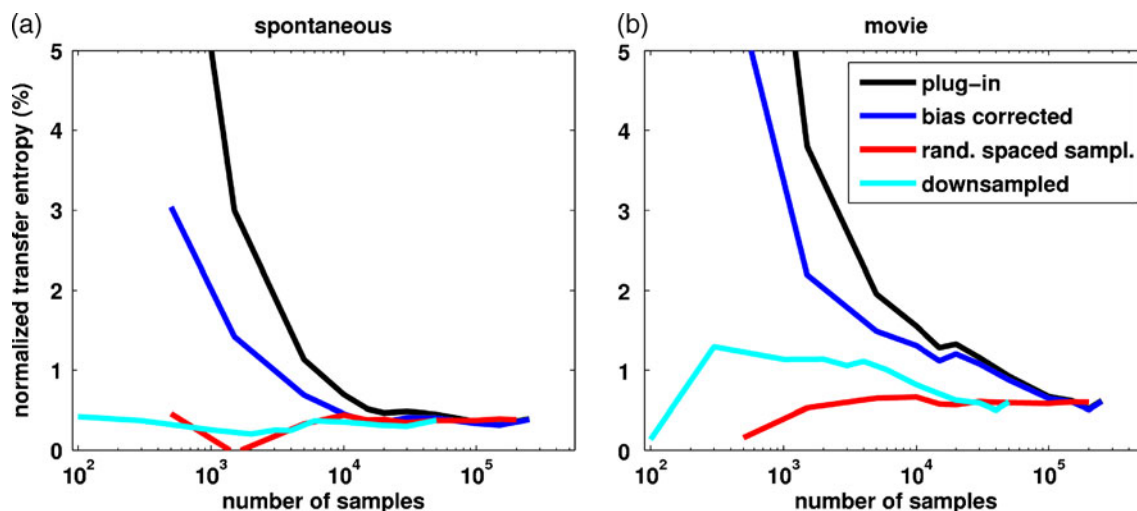


Fig. 2 Convergence of NTE at middle temporal resolution between MUA activities from pairs of sites recorded from V1 of monkey A98. Two stimulus conditions are considered; spontaneous activity (**a**) and movie stimulation (**b**). NTE is estimated from 5 trials of 4 minutes each and the number of samples is modulated by increasing the length of the time interval starting from stimulus onset. Four estimation procedures are considered. *plug-in* estimate (*black curve*) is computed directly from Eq. (5) with

value when either 5 trials of spontaneous activity or 5 trials with presentation of the same movie were used. However, a much faster convergence to the same asymptotic value was obtained by downsampling the time series by decimation of a factor 5 combined with our novel NTE sampling bias correction techniques that compensates for the reduced number of samples. This allowed a substantial reduction of the computational time without deteriorating the sampling properties. Further analysis will thus be done using downsampling combined with bias corrections.

4.4 Bootstrap test of significance of causal relationships above those imposed by a common stimulus drive

After studying the convergence of the method, we used it to investigate whether there are significant causal interactions between the bands of the cortical extracellular potentials. Based on the above convergence results, in the following we used 5 trials of the same experiment to compute one NTE value (for each particular pair of frequencies and electrodes). When more than five trials from the same movie were available, trials for each movie condition were divided into subgroups of 5 consecutive trials, and each subgroup was analyzed separately. We computed TE between different frequency bands as a function of the delay τ used to compute the conditioning with respect to the causing signal

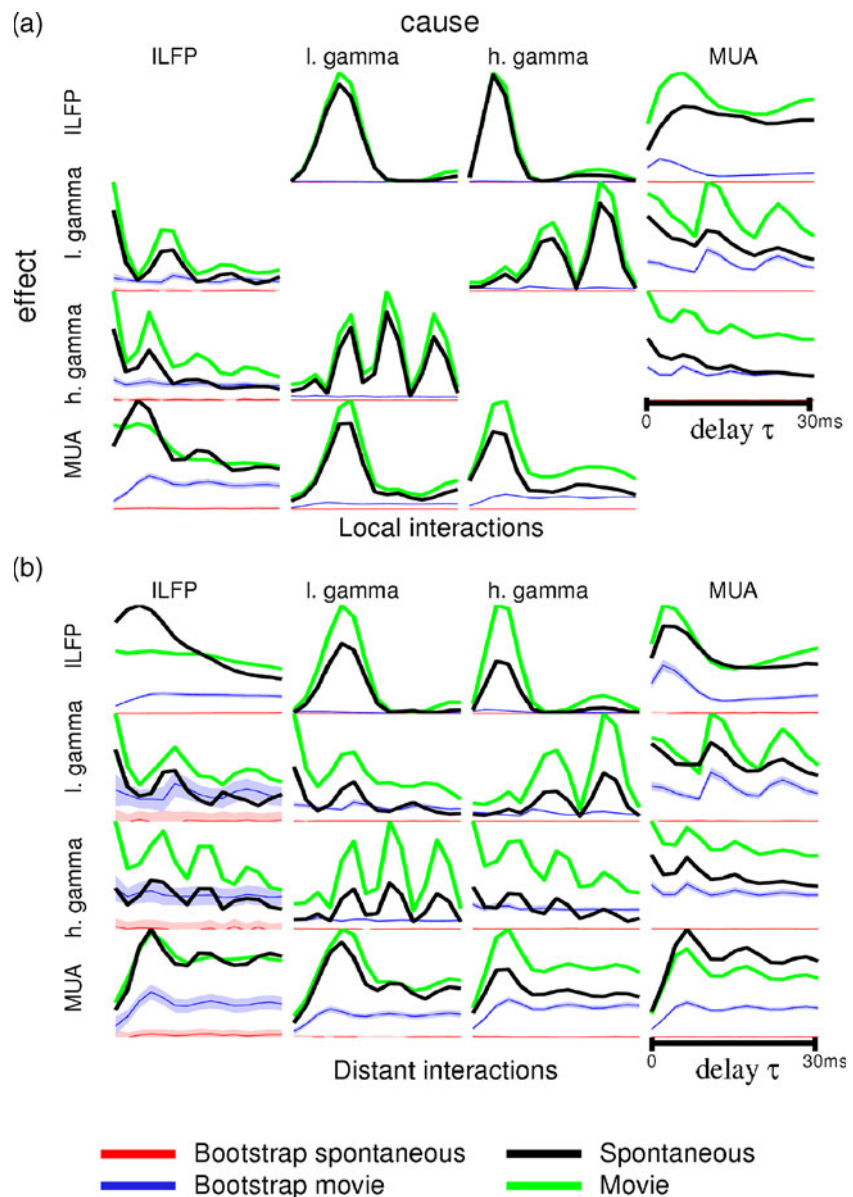
out any bias correction. *bias corrected* estimate is computed by subtracting a bootstrap estimate to cancel the bias (see Eq. (7)). *rand. spaced sampl.* is a corrected estimate, which takes the same amount of samples as the previous estimate, but the samples are spaced randomly on the whole time interval of the trials: this is a way to destroy time correlation between samples. *downsampled* is a bias-corrected estimate using time series down-sampled to 60 Hz (original sampling frequency is 300 Hz)

(Eq. (4)). Unless otherwise stated, results will be reported as average over all animals, recording sites and trial subgroups.

We first investigated whether the measured causal interactions between frequency bands were statistically significant. To assess this significance it is necessary to quantify the distribution of NTE values under a null hypothesis of non-causality \mathcal{H}_0 . We estimated the properties of this distribution from our data using a bootstrap procedure. To compute the distribution under \mathcal{H}_0 of TE from X to Y , we estimated TE from X to Y^* , where the trials of Y^* are drawn randomly without replacement from the trials of Y . Given that the trials were several minutes long, and that correlations between neural signals span a much shorter range, this bootstrapping destroys all causal relationships apart from those arising in the movie condition due to a common stimulation history for both neural signals by the same movie in all different trials. By running 20 times the bootstrap procedure for each subset of experiments, we computed the mean and standard deviation of the bootstrapped NTE values and compared to the original NTE value.

Figure 3 reports, separately for causal interactions computed from the same electrode (panel a) and from a different electrode (panel b), the results of this comparison at high temporal resolution ($F = 100$ Hz). At this resolution, in the case of movie stimulation and for several pairs of frequency bands of the extracellular

Fig. 3 *Bootstrap NTE statistics* for causal interactions between frequency bands at high temporal resolution (low pass filtered to $F = 100$ Hz, cf. Table 1), as a function of the delay parameter (τ). Average value of $NTE_{cause \rightarrow effect}^{cor}$ across the whole dataset (all monkeys, all movies and all recording sites) is plotted as a function of the delay τ for two conditions: movie stimulation (*green curve*) and spontaneous activity (*black curve*). *Blue* and *red* curves indicate the corresponding average of bootstrap estimates for the same subset of electrodes. The *shaded* area indicates the average standard deviation of the statistics around the mean. Due to run-time constraints, the original and the bootstrapped NTE values in this figure were computed from two electrodes per recording session only, and then averaged across sessions. **(a)** *Local interactions*, the average is taken between NTE values of signals from the same electrode. **(b)** *Distant interactions*, the average is taken between NTE values of signals from two different electrodes



signal, we found that the bootstrap values of TE were not distributed around zero, meaning that a part of the causal relationships were due to common movie stimulation. However, for all pairs of LFP bands and for MUA the original causality values remained well above their corresponding bootstrap distribution, implying that causal interactions between frequency bands exist even when discounting for common stimulation history. Moreover, Fig. 3 shows that in most cases the fraction of causal interactions during movie stimulation due to a common stimulus drive was only a small fraction of the total amount of causation between the signals.

Causal interactions at low temporal resolution ($F = 8$ Hz) are reported in Supplementary Fig. 4 (Online Resource 1). At such low resolution ($F = 8$ Hz) causal

interactions from lower-frequency signals (low LFP and theta bands) to higher frequency signals (beta and gamma bands) were not significantly higher than in the bootstrapped condition. On the other hand, some causal interactions both within lower frequencies (low LFP and theta) and within higher frequencies (gamma and MUA) were far from their bootstrap distribution and were thus highly significant. In general, the comparison of Fig. 3 and Supplementary Fig. 4 (Online Resource 1) shows that the ratio between actual values of NTE and bootstrap values for movie stimulation was much lower at low temporal resolution than at high resolution. This implies that the common driving by the stimuli has less of an impact on causality measures at fine temporal resolution. This is consistent with the

fact that these movies had the most power in the low frequency (below 4 Hz) range, which in turn implies that the stimulus drive is mostly at low frequencies (Belitski et al. 2008; Montemurro et al. 2008).

Another result worth commenting is the dependency of NTE value on the time delay τ . In particular, we observe that NTE values in Fig. 3 for interactions in the gamma bands are oscillatory. We investigated whether this shape was specific to our results or a consequence of the oscillatory nature of signals by a simulation study fully reported in online resource 1, Section “Two frequencies, linear system”. The main result was that the observed dependence of NTE on the time delay shape could be obtained for simulated band pass signals, and the pseudo period was a function of the original period of the oscillations. As a consequence, maxima in the curve should not be interpreted as the characteristic time delay of the causal interaction. Nevertheless, simulations also showed the maximal amount of NTE over time was related to the actual causal interaction. Therefore we decided to use the latter parameter as the quantification of causality.

4.5 Modulation of causal interactions by the presence of visual stimuli

The above results indicate that the amount of causal interaction is modulated by the presence of visual stimuli. In this subsection, we examine in detail the strength and statistical significance of these changes of causality due to the presence of the movie stimuli with respect to the stimulus-free (spontaneous) condition.

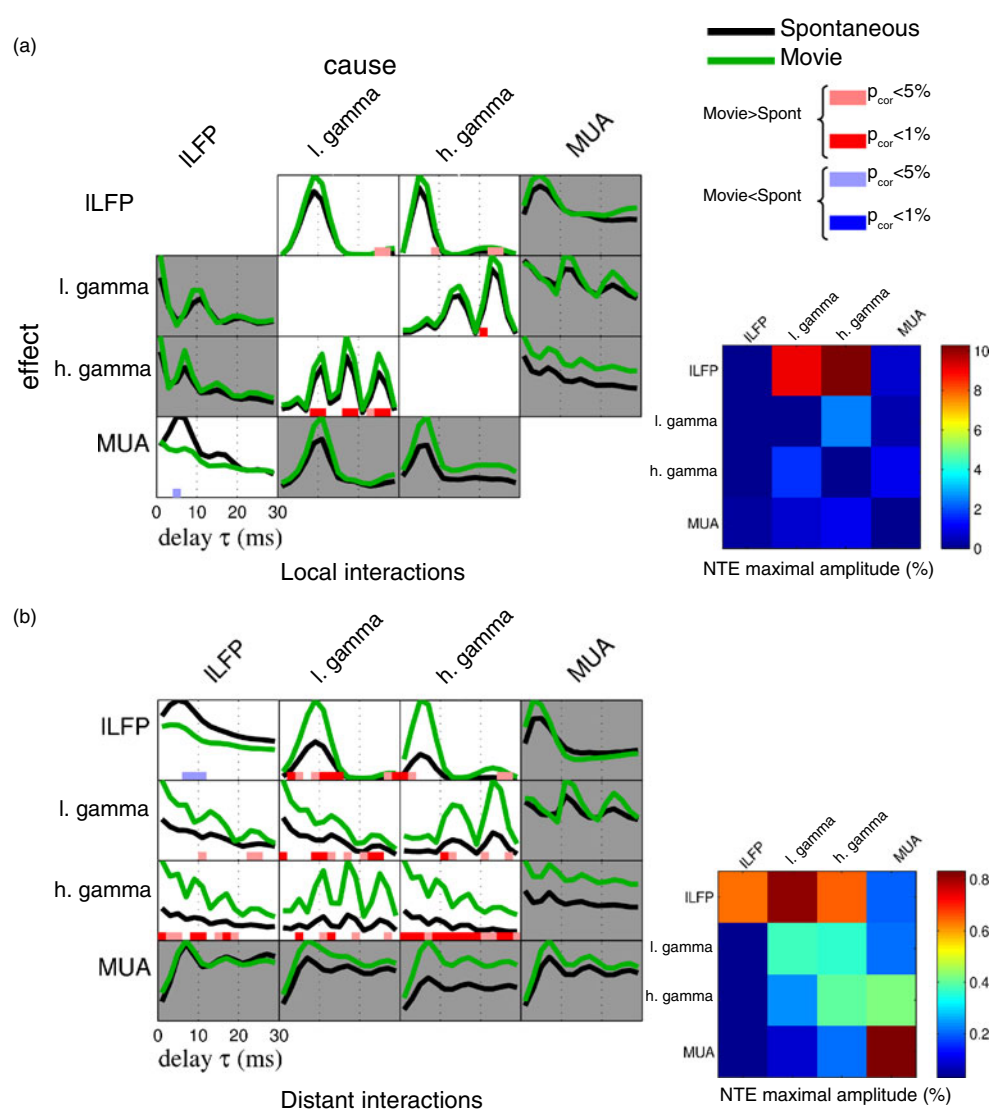
To allow correction for multiple comparison, the statistical significance of the effect of the type of stimuli (movie versus spontaneous activity) was evaluated through all possible delays and couples of frequencies with a permutation test. The chosen approach is similar to the one used by Pantazis et al. (2005) for the case of a T-statistics. In our case, we test the significance level of the F-statistics of the stimulus effect in a two way analysis of variance (ANOVA) using the independent variables “monkey” and “stimulus”. The distribution of this statistics under the null hypothesis was computed as follows: for each monkey, we randomly shuffled the NTE values corresponding to “movie” and “spontaneous” conditions. Then the maximum of the F-statistics for the stimulus effect through all possible time delays and couple of frequencies was computed. Using 300 iterations of this shuffling, the distribution of the maximal F-statistics under the null hypothesis was computed, as well as a threshold corresponding to the desired p-value. Then the F-statistics was evaluated on the original dataset and any delay and couple

of frequency corresponding to a TE value above the threshold was considered significant. This analysis was carried out at each time resolution considering separately: causal interactions within the same electrode, and interactions between different recording sites.

Results for the high temporal resolution case ($F = 100$ Hz) are reported in Fig. 4. For high resolution interactions between signals from different electrodes (denoted as “distant interaction”, Fig. 4(b)), the amount of causation between all pairs of LFP bands (low LFP, low gamma and high gamma band) during movie stimulation was significantly different from that measured in the absence of stimulation. Except for the low LFP to low LFP causality, for which movie stimulation induces a decrease of NTE, all other significant pairs exhibited an increase in causality. There were no significant changes involving MUA spiking activity either as cause or effect. Since the distances between recording sites were in the range 1–7 mm, this suggests that interactions between MUA and LFP bands are more local. We thus computed local causal interactions within the same recording site (i.e. causal interactions between signals from the same electrode; Fig. 4(a)). The NTE values obtained for local interactions were several times larger than those obtained for distant interactions. However, the pattern of NTE changes between movie and spontaneous condition for local causal interactions (Fig. 4(a)) were in most cases consistent with the case of distant interactions: for example we found that interactions between the two LFP gamma bands, and between gamma LFPs and lower-frequency LFP bands also increased during movie presentation. The most notable difference between the local and the distant case was that in the local case we also found causal interactions involving the MUA band: in particular a decrease of low-LFP to MUA causation during movie presentation.

Since the gamma bands are not rectified at this high temporal resolution (see Table 1), the measurements in this band mix envelope (or power) and phase information. We thus did further analysis to disambiguate the causal interactions provided respectively by the phase and envelope of gamma oscillations. Using the Hilbert transform we computed the instantaneous phase and envelope associated to oscillations in the low and high gamma bands. These measures were computed respectively as modulus and angle of the complex time series given by the Hilbert transform of the considered signal. Then causal interactions were recomputed using either the envelope or phase time series for the gamma band. Results are reported in Fig. 5. We found (Fig. 5(a, b)) that gamma amplitude only accounts for increases in causality from high gamma to high gamma and from gamma to low LFP. On the other hand, low and high

Fig. 4 Causal interactions between frequency bands at high temporal resolution ($F = 100\text{ Hz}$, cf. Table 1), as a function of the delay parameter (τ). Average value of $NTE_{cause \rightarrow effect}^{cor}$ across all monkeys and recording sites is plotted as a function of the delay τ for two conditions: movie stimulation (green curve) and spontaneous activity (black curve). Red and blue horizontal lines indicate the delays for which the difference between conditions is significant. The column in the table corresponds to the cause frequency and the row to the effect frequency. Gray background indicates no significant effect for all delays. (a) *Local interactions*, the average is taken between all NTE values of signal from the same electrode. (b) *Distant interactions*, the average is taken between all NTE values of signal from any two different electrodes. For each panel, the colored grid on the right hand side indicates the maximal NTE value of the considered pair of frequencies over the entire τ range shown in the corresponding left hand side plots



gamma phase is involved in causality increases with many frequency bands (Fig. 5(c, d)). Interestingly, in addition to previously observed NTE increases in LFP bands, we observe significant increases of interactions between gamma phase and spiking (MUA) activity, which were not detected in the initial analysis that did not separate phase and amplitude (Fig. 4). Moreover, whereas previous results (Fig. 4) were mainly symmetric (i.e. causality changes were the same from band A to band B and from B to A), when separating out the contribution of gamma phase, some changes were observed in only one direction: namely local interactions from MUA to low gamma, and distant interactions from high gamma to MUA. This shows that the phase/amplitude decomposition gives additional information on the underlying causal structure of our data.

We further investigated whether NTE measures were related to other measures of interactions such as

phase locking value (Lachaux et al. 1999). We computed phase locking value between gamma frequency bands and found that they were positively correlated with NTE measures when looking at interactions in the same frequency band. For interactions between different frequency bands, we computed n:m phase locking value (see Schack et al. 2005 for example) and found no significant correlation with NTE. These results are reported in Online Resource 1 (Section “Linking TE with phase synchrony”).

The study of causal interactions was also carried out for the low ($F = 8\text{ Hz}$) and middle ($F = 30\text{ Hz}$) temporal resolutions. For low resolution (Supplementary Fig. 5 in Online Resource 1), significant decreases in causal interactions in the movie (with respect to the spontaneous) condition were found from low gamma to theta band and from MUA to beta; and significant increases in the movie (with respect to the spontaneous)

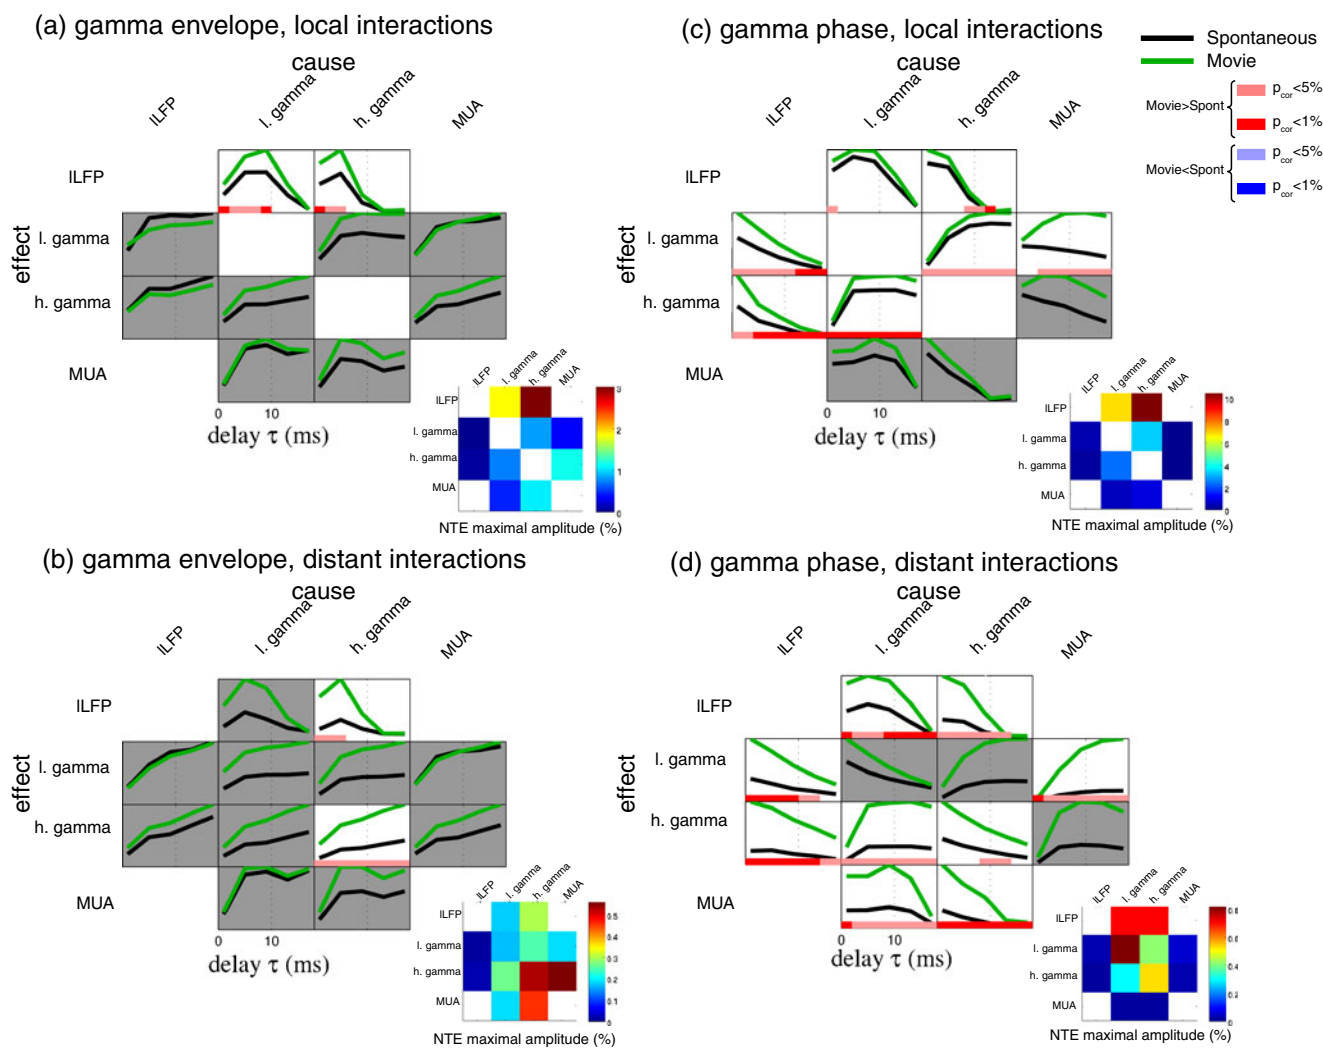


Fig. 5 Causal interactions with phase/envelope of the gamma band at high temporal resolution ($F = 100$ Hz, cf. Table 1), as a function of the delay parameter (τ). Average value of $NTE_{cause \rightarrow effect}^{cor}$ across all monkeys and recording sites is plotted as a function of the delay τ for two conditions: movie stimulation (green curve) and spontaneous activity (black curve). Red and blue horizontal lines indicate the delays for which the difference between conditions is significant. The column in the table corresponds to the cause frequency and the row to the effect frequency.

quency. (a) Average NTE for the *envelope* of gamma oscillations (*local interactions*). (b) Average NTE for the *envelope* of gamma oscillations (*distant interactions*). (c) Average NTE for the *phase* of gamma oscillations (*local interactions*). (d) Average NTE for the *phase* of gamma oscillations (*distant interactions*). For each panel, the colored grid on the bottom right hand side indicates the maximal NTE value of the considered pair of frequencies over the entire τ range shown in the corresponding left hand side plots. Gray background indicates no significant effect for all delays

condition were found within the gamma bands. The magnitude of the estimated NTE was highly dependent on the considered couple of frequencies and was maximal for low frequency (theta) and high frequency (gamma, MUA) bands. In particular, at low temporal resolution, cross-interactions between low and high frequencies have clearly lower NTE magnitudes. Values were also higher when considering local interactions between signals from the same electrode (panel a) compared to distant interactions (panel b). The shape of NTE curves as a function of delay exhibit dissimilarities

depending on the considered couple of frequencies. NTE computed between rectified frequencies (gamma and MUA) tend to be small for small values of τ , to increase rapidly to a maximal value and then slowly decrease. Considering causality measures from low frequency bands to themselves (on the diagonal in the arrays), NTE values are progressively decreasing as a function of the delay τ . At middle temporal resolutions (Supplementary Fig. 5 in Online Resource 1) almost no interactions are significant, thus the following analysis will focus on the two other time resolutions.

4.6 Net causality

The results presented above report that causality between two signals is often found in both directions. These results are difficult to interpret. In particular, it is legitimate to infer a leading causality direction only if causality is stronger in one direction than in the other. If causality is similar in magnitude in both directions, then the results are better interpretable in terms of coherency between the two signals. To investigate whether there was a leading direction of causality among different frequency bands, we therefore computed net NTE values as the difference of NTE values between the two directions $\Delta NTE(1, 2) = NTE(1 \rightarrow 2) - NTE(2 \rightarrow 1)$.

Results of net NTE values in different frequency bands, across all couples of electrodes and sessions of our electrophysiology data are presented in Fig. 6. We first focused on local interactions and computed the net NTE between raw gamma signals and other frequency band signals. (Fig. 6(a)). We observed the following leading causal directions: MUA \rightarrow ILFP, gamma \rightarrow ILFP. Interactions within the gamma band and between MUA and gamma did not exhibit a clear asymmetry. When considering distant interactions (Fig. 6(b)) the driving of ILFP by raw gamma and MUA was preserved. Additionally, the causal direction MUA \rightarrow gamma was found. We then investigated the leading causal directions when considering the gamma phase (rather than the raw gamma-band signal). The results (reported in Fig. 6(c) for the case of local interactions) confirmed the leading causal direction gamma \rightarrow ILFP. Moreover, the leading causal directions h. gamma \rightarrow l. gamma and gamma \rightarrow MUA were found when considering phase. Interestingly, these results were different in the case of distant interactions (Fig. 6(d)) for the case of gamma-phase/MUA interactions: the dominant direction l. gamma \rightarrow MUA was found (similarly to the results using raw gamma) whereas h. gamma/MUA relationships exhibited the same trend but were less clear at small delays. To summarize, ILFP was always driven by gamma and MUA; gamma phase was driving MUA locally and MUA was driving gamma on distant electrodes.

We further checked that the results obtained with the net TE were not simply due to the difference in frequency of the two signals. Using simulations, reported in Online Resource 1 Section “Two frequencies, linear system”, we found that the difference in frequencies does not bias substantially the net TE. However we found out that post-processing by filtering, as it is done in our study to extract frequency bands, can bias the causality measure towards the direction of

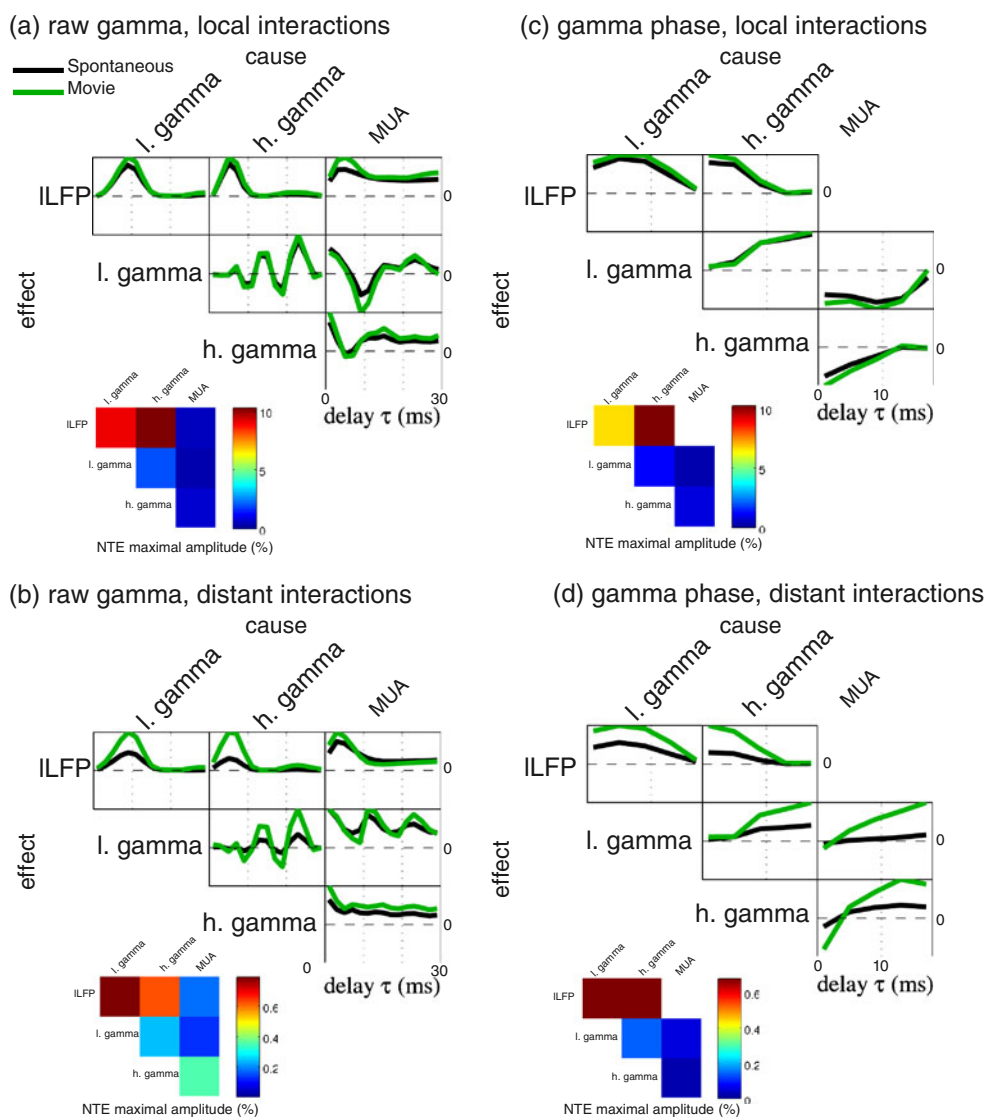
the lowest frequency to the highest. This is consistent with previous reports (e.g. Quiroga et al. 2000) that interactions or causality measures between time series tend to be biased in the direction from low to high frequencies (see e.g. Quiroga et al. 2000, Fig. 2). Since our experimental results report positive net causality in the direction from high to low frequencies, these considerations suggest that the experimental findings do not arise simply because of frequency differences in the signals. This point is further corroborated by the observation (Online Resource 1 Supplementary Fig. 8) that the average power spectrum for MUA, ILFP and gamma amplitudes is also highest at low frequencies (a fact that, as reported in Mazzoni et al. 2008) can be explained by the fact that spiking and gamma activity is influenced by stimulus drive and spontaneous state fluctuations, and both of them vary at slow time scales). This latter observation suggests that the differences in the natural frequencies of signals in the different bands are less pronounced than those suggested by simply looking at the band boundaries.

4.7 Robustness across recording sites and sessions

After presenting the results of the average NTE across the entire dataset, we tested the robustness of causal interactions across all the different recording sites and animals used in the experimental sessions. We focused this robustness analysis mainly on pairs of frequency bands exhibiting significant causal interactions, and for simplicity we report NTE values at the time delay for which the maximal F statistics for the effect of stimulus was reached.

For each experimental session, we computed the proportion of recording sites for which NTE during movie presentation was significantly higher than NTE during spontaneous activity. If this proportion is close to 100% then movie stimulation induces a very robust increase of causality between all couples of recording sites. Conversely, if this proportion is close to 0%, it reflects a robust decrease of causality induced by stimulation. This proportion, plotted as function of the distance between recording sites, is reported for the case of high temporal resolution in Supplementary Fig. 6 (Online Resource 1). At this high temporal resolution, the consistency of the results across sessions and recording sites was very good. In particular for local interactions between low gamma and high gamma, almost 100% of couples of sites exhibited an increase of causality in all sessions. Most of the significant changes were highly consistent. It is noteworthy that the increase of local interaction from high gamma to MUA is also very consistent (although the previously computed statistical

Fig. 6 Average NTE difference between the two directions of causality. The quantity is $\Delta NTE(\text{cause}, \text{effect}) = NTE_{\text{cause} \rightarrow \text{effect}} - NTE_{\text{effect} \rightarrow \text{cause}}$. If the quantity is positive, the causal direction $\text{cause} \rightarrow \text{effect}$ (from column to row) dominates. For each panel, the colored grid on the bottom left hand side indicates the maximal absolute value of NTE difference for the considered pair of frequencies over the entire τ range shown in the corresponding above plots. **(a)** NTE difference for local interactions within the same electrode using raw gamma oscillations. **(b)** NTE difference for distant couples of electrodes using raw gamma oscillations. **(c)** NTE difference for local interactions within the same electrode using phase of gamma oscillations. **(d)** NTE difference for distant couples of electrodes using phase of gamma oscillations



test did not reveal a significant increase, possibly because it is very conservative). Moreover, the decrease of distant interactions within the low LFP band is also very consistent. For distant interactions, we observed only a slight effect of the distance between recording sites, in particular the increases in causality between low gamma and low LFP were less pronounced for large distances (>3 mm). Finally, we observed that distant interactions within the MUA band were highly variable across session and recording sites, which explains why these interactions were not detected in the previous analysis performed on population averages.

Results for the consistency analysis at low temporal resolution ($F = 8$ Hz) are reported in Supplementary Fig. 7 (Online Resource 1). The results are much more variable across sessions than in the high temporal

resolution case, in particular if they are recorded in different monkeys. One of the most striking examples is the distant causal interaction from high gamma to MUA, for which the index is close to 0% for sessions D04nm1 and D04nm2, and close to 100% for sessions C98nm1 and A98nm5. Among the pairs of frequency bands with a significant modulation of causality by the movie, the only pairs of frequency bands showing a good stability of the result across sessions and recording sites are interactions within gamma bands. For local interactions, the robustness is also good for the decrease from low gamma to theta (Supplementary Fig. 7b in Online Resource 1).

In sum, causal interaction exhibit variability across monkeys and recording sites at lower temporal resolution. At high temporal resolution, the pairs of

frequencies showing significant modulation of the NTE by the movie stimulation at the level of population average also exhibit a very good robustness of this result across sessions and animals.

4.8 Effect of the distance between sites

We finally investigated the effect of interelectrode distance on TE. We investigated this by computing NTE at high temporal resolution and by plotting the joint distribution of these NTE values at different interelectrode distances. Smoothed histograms of this distribution are shown on Supplementary Fig. 9 (Online Resource 1) for the most relevant LFP frequency bands. For all considered LFP bands, and both during spontaneous activity and during visual stimulation with movies, NTE values remained high (more than 70% of their maximal value, which was observed when computing them from 1 mm distant electrodes) for interelectrode distances up to 3 mm, and then it decayed rapidly to very low NTE values (<30% of the maximum) for interelectrode distances larger than 5 mm.

5 Discussion

Over the last few years, several theoretical and experimental studies have investigated the origin and potential functional meaning of the rich structure of frequencies present in the time series of extracellular potentials (Buzsáki 2006; Bedard et al. 2006; Pettersen and Einevoll 2008; Mazzone et al. 2008; Belitski et al. 2008; Roopun et al. 2008; Kayser et al. 2009; Nadasdy 2009). Here, we aimed at contributing to the understanding of the relationships between the different frequency components of the extracellular signal by developing a novel, fast and data robust procedure to compute TE between time series of bandpassed neurophysiological signals in various frequency bands. Although it is possible that causality is to some extent a wide-band phenomenon (Nolte et al. 2008), our approach considering the causal relationship between separate frequency bands in the extracellular signal is justified and motivated by the bulk of neurophysiological evidence linking different frequency ranges to different functional states and neural phenomena (Buzsáki and Draguhn 2004; Buzsáki 2006).

We illustrated our technique by computing and comparing causal interactions between different frequency bands of the extracellular signal recorded from primate V1 during spontaneous activity or during binocular visual stimulation with naturalistic movies. Our results individuated causality changes during visual stimula-

tion, which involved specific time scales and frequency bands. The significance of the methods developed here and of the analysis of the neurophysiological data is discussed next.

5.1 Methods for computing TE and their convergence properties

One of the main contributions of our study was to introduce and develop a novel procedure to computing NTE between frequency bands and recording sites of the extracellular signal recorded intracranially with microelectrodes. This technique built on previous progress in computing the sensory information carried by LFP and MUA bands (Belitski et al. 2008; Montemurro et al. 2008; Magri et al. 2009), and was based on discretization of the signal, on corrections for the bias due to limited sampling in information measures (Panzeri et al. 2007), and on downsampling to achieve at the same time a faster speed of computation and a reduction of potential artifacts due to correlation between successive time samples (Theiler 1986). We investigated the convergence properties of the measure with the sample size and found that this method held excellent convergence properties. This fast convergence was crucial in allowing us to estimate reliably NTE values and to compare them across a large dataset and across several frequency bands and time scales within a reasonable computational time. Because of these features, our technique could become valuable to the neurophysiology community for further studies of causation.

Our analysis methods depart from that used in a previous attempt to estimate TE from intracranial recordings, which used an approach based on approximating differential entropies using KDE (Schreiber 2000; Chavez et al. 2003). One reason why we could not use KDE (or NND) techniques in the present study was that, despite their undoubted power and appeal (Grassberger 1988; Kraskov et al. 2004), they were too computationally expensive to be run on such a large dataset as ours. However, an important topic of future methodological research is to compare in detail the relative advantages of KDE/NND and discretization methods with bias corrections in computing information theoretic quantities from analog neural signals, and in trying to integrate the relative strengths of both approaches. In the neurophysiology domain, detailed comparisons between KDE/NND methods and discretization methods based on up-to-date bias correction procedures were so far only performed on spike trains (Nelken et al. 2005) and showed that NND techniques required a large amount of neural data to converge unless the underlying probability distributions were

sufficiently smooth. However, it is quite conceivable that distributions of analog neural signals are much smoother than those of spike trains (Magri et al. 2009), and so KDE methods may give a faster convergence on these datasets. Understanding the relative advantages of both methods, and producing a set of fast publicly available routines to compute them would greatly increase the tools available to experimental laboratories to investigate the chain of causal processes in the nervous system.

5.2 Causality depends on the temporal resolution at which it is considered

Since the time scale of causal interactions was not known, we investigated causal relationships at three different temporal resolutions (low, middle and high) by low pass filtering the signals at 8 Hz, 30 Hz and 100 Hz respectively. Significant changes in TE induced by visual stimulation were mainly observed at low and high resolution. However, the causations observed at high temporal resolution were stronger and more robust across sessions than those observed at low resolution. We therefore focus the rest of this discussion on causal interactions revealed at high temporal resolution.

5.3 Significance and directionality of causal interactions

One of the main results of our analysis was that we established the presence of several highly significant causal relationships between specific frequency bands. We established significance of causal interactions by means of a bootstrap test, which left only causations due to common stimulation history. At high temporal resolutions, the NTE values obtained after bootstrapping were typically much smaller than the NTE values recorded from the non-bootstrapped data. The significance of these results is that they suggest that the causations we observed were due only in small part to the effect of common stimulation history, and is important given the fact that techniques to eliminate confounders in causal inference (Pearl 2000) are well developed for linear measures of causality (Guo et al. 2008a; Chen et al. 2006), but are very difficult to handle in the non-linear case. The study of how to handle them in nonlinear situations is an important topic for future analytical research.

At a high time resolution, there were highly significant robust interactions between virtually all frequency bands considered, both for signals from the same electrode or for different electrodes. The most interesting and robust causality relationships were

those involving the gamma band. Gamma band had a stronger causal effect of all LFP bands at lower frequencies during movie stimulation than during spontaneous activity. These TE changes were significant for very short time delays of a few milliseconds. Theoretical and experimental results relate gamma oscillations to the activity of recurrent local microcircuits of inhibitory and excitatory neurons (Mazzoni et al. 2008; Brunel and Wang 2003; Cardin et al. 2009). The increases in the causal effect that the gamma band exerts on other LFP bands suggests that the local recurrent loops of inhibitory and excitatory circuitry are more prominently activated and play a more central role in controlling the rest of the network activity during the processing of visual stimuli than during spontaneous activity.

Another interesting result afforded by the high temporal resolution analysis was that at this resolution the gamma activity could be quantified without rectification, and this allowed us to disambiguate the causal effects due to gamma phase from those due to gamma amplitude. We found that, although gamma amplitude played a role in causing a number of other LFP bands, gamma phase had a much prominent role in increasing causation with other signals during visual stimulation, in particular with MUA activity originating from different electrodes. This result is in agreement with, and extends to the non-linear causal analysis domain, previous seminal findings obtained using linear correlation analysis to show that gamma phase modulates the communication between neuronal groups (Womelsdorf et al. 2007). Interestingly, previous studies on the same dataset that we analyzed here showed that gamma phase was not reliably locked to the stimulus (Montemurro et al. 2008). The findings that gamma phase is crucially involved in controlling and causing spiking activity in other locations suggests that one reason why it may be functionally convenient not to lock gamma phase to the stimulus time course is that in this way gamma phase is left free to vary across regions receiving the same stimulus dynamics and thus tune the amount of communication between neuronal groups depending on other needs.

The finding of significant interactions at high temporal resolution between frequency bands in both directions raised the question of whether these results imply coherency between the bands or whether it is possible to infer a leading direction of causality between bands. We addressed this problem by considering the net NTE (defined as the difference of NTE values between the two directions). This calculation led to the finding that there is indeed a dominant direction of causality, mainly in the direction from MUA or gamma frequency

band signals to lower frequency signals. Moreover, we found that the dominant direction of causation between gamma and MUA signals depended upon the spatial scale considered. In particular, there was a same-electrode driving of MUA by gamma phase and a cross-electrode driving of gamma by MUA. This, in our view, gives fresh insights on the hierarchical set of correlations commonly found between lower and higher frequency rhythms of cortical activity (Roopun et al. 2008), and suggest that these hierarchical correlations are largely caused by the faster rhythms rather by the slower rhythms.

We further studied systematically how the amount of causation of frequency bands and their changes with the stimulation conditions varied with the inter-electrode distance, in the range of few millimeters. Robust increases in causal interactions during movie stimulation were found in the gamma band between recording sites separated from several millimeters in the primary visual cortex. However, such robust changes across distant electrode were not found when considering MUA spiking activity. Larger spatial correlation between gamma band activities compared to spiking have already been reported (Berens et al. 2008). It has been shown the impedance of extracellular tissue is independent of the frequency (Logothetis et al. 2007). Thus the large distance interactions observed in the gamma band can not be explained by propagation in the tissue, but result from network activities, possibly mediated by lateral connections, which are known to spread on several millimeters (Stettler et al. 2002).

Long range causal interactions between different areas have also been investigated in the literature using Frequency domain Granger causality (Brovelli et al. 2004; Guo et al. 2008b). The results of these previous studies mainly reported significant interactions in low frequency bands (theta and beta bands), which contrasts with our finding of a central role for the gamma band. This difference may be accounted for by the differences in the considered spatial scale of the interactions: we are considering interactions within a same area at a maximal distance of few millimeters whereas previous studies consider interactions between more distant brain areas.

Acknowledgements We are grateful to J. Martinerie, M. Chavez, D. Janzing, Y. Murayama, C. Kayser, N. Ludtke, A. Belitski and C. Magri for many useful discussions and interactions.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

Appendix: Shuffling correction of the limited sampling bias of TE

In this appendix, we explain the so called “shuffling correction procedure” that we developed to reduce the sampling bias affecting transfer entropies.

Sampling bias arises because the probabilities needed to compute the entropies in the TE equations (4) and (5) are not known but have to be measured experimentally from the limited number of data points in which the neurophysiology or imaging data were recorded. The estimated probabilities are subject to statistical error and necessarily fluctuate around their true values. Since the information theoretic probability functional is non-linear, the finite sampling fluctuations lead to a systematic error (bias) in the estimation of the probability functional (Panzeri et al. 2007). This bias is negative when considering entropies and is approximately directly proportional to the cardinality of the probability space to be sampled and inversely proportional to the number of available data points (Panzeri et al. 2007). We have previously proposed a number of algorithms for the elimination of the bias of mutual information between stimuli and neural responses (Montemurro et al. 2007; Panzeri et al. 2007). Here we extend this work to correct for the bias of TE.

In the following, we will write \hat{Q} for the plug-in estimate of the quantity Q (computed from the empirical probability distribution of the binned data). We first rewrite TE as:

$$T_{Y \rightarrow X} = H(X_t | X_{t-\tau}) + H(Y_{t-\tau} | X_{t-\tau}) - H((X_t, Y_{t-\tau}) | X_{t-\tau}) \quad (6)$$

When computing the plug-in estimate $\hat{T}_{Y \rightarrow X}$ from Eq. (6), the term with the worst sampling behavior and the most biased is $H((X_t, Y_{t-\tau}) | X_{t-\tau})$, because (unlike the other two terms in the r.h.s. of Eq. (6)) it needs estimation of a bivariate conditional probability distribution. Thus the bias of $H((X_t, Y_{t-\tau}) | X_{t-\tau})$ (which is negative) dominates the bias of the TE, and as a result TE is biased upward due to limited sampling. Fortunately, the bias of multivariate entropies such as $H((X_t, Y_{t-\tau}) | X_{t-\tau})$ (and thus that of TE) can be greatly reduced at the source by the techniques of Montemurro et al. (2007) and Panzeri et al. (2007). In a nutshell, the idea is to rewrite TE by subtracting and adding to the estimation of TE two terms with exactly the same asymptotic values for large number of trials, but with a bias for finite number of trials, which cancels out the one of the multivariate entropy causing the most sampling problems. In this way, this corrected estimate

converges faster to the true values with the data size. This is done by estimating $\hat{T}_{Y \rightarrow X}^{sh}$ through the following plug-in shuffled estimate:

$$\begin{aligned} \hat{T}_{Y \rightarrow X}^{sh} = & \hat{H}(X_t|X_{t-\tau}) + \hat{H}(Y_{t-\tau}|X_{t-\tau}) \\ & - \hat{H}((X_t, Y_{t-\tau})|X_{t-\tau}) - \hat{H}_{ind}((X_t, Y_{t-\tau})|X_{t-\tau}) \\ & + \hat{H}_{sh}((X_t, Y_{t-\tau})|X_{t-\tau}) \end{aligned} \quad (7)$$

where

$$\hat{H}_{ind}((X_t, Y_{t-\tau})|X_{t-\tau}) = \hat{H}(X_t|X_{t-\tau}) + \hat{H}(Y_{t-\tau}|X_{t-\tau})$$

is the joint conditional entropy under the assumption of conditional independence of $X_t, Y_{t-\tau}$ given $X_{t-\tau}$. The latter assumption makes it possible to rewrite $\hat{H}_{ind}((X_t, Y_{t-\tau})|X_{t-\tau})$ as a sum of entropies of univariate conditional marginal distributions, which in turn means that this term has a very little bias compared to that of the bivariate entropy $H((X_t, Y_{t-\tau})|X_{t-\tau})$. The term $\hat{H}_{sh}((X_t, Y_{t-\tau})|X_{t-\tau})$ is the bivariate entropy computed by shuffling all the samples of $Y_{t-\tau}$ corresponding to a particular value of $X_{t-\tau}$, while samples from X_t values remain unchanged. This is equivalent to sample data from a distribution with the same marginal conditional probabilities $p(X_t|X_{t-\tau})$, $p(Y_{t-\tau}|X_{t-\tau})$, but where the conditional independence assumption mentioned above holds. As shown in Montemurro et al. (2007), for a large number of samples, this shuffled entropy $\hat{H}_{sh}((X_t, Y_{t-\tau})|X_{t-\tau})$ has asymptotically the same value as $\hat{H}_{sh}((X_t, Y_{t-\tau})|X_{t-\tau})$, but its bias is similar in magnitude and scaling to that of $\hat{H}((X_t, Y_{t-\tau})|X_{t-\tau})$ (because they are both bivariate and computed from the same number of trials). Thus adding $\hat{H}_{sh}((X_t, Y_{t-\tau})|X_{t-\tau}) - \hat{H}_{ind}((X_t, Y_{t-\tau})|X_{t-\tau})$ to the TE plug-in computation does not change its asymptotic value for infinite number of trials, but dramatically improves its bias property for finite datasets because it makes it similar to that of univariate conditional entropies. Therefore in the paper we used $\hat{T}_{Y \rightarrow X}^{sh}$ of Eq. (7) as the bias corrected TE estimate.

We note that, in agreement with the results of Panzeri et al. (2007), subtraction of further bias correction techniques (such as those in Panzeri and Treves 1996; Strong et al. 1998) did improve slightly the convergence of the estimates with the dataset (results not shown). However, they also added a very significant increase computational time with a relatively little convergence benefit, and therefore we did not use these additional bias corrections in the extensive analysis of the neurophysiological dataset.

References

- Ancona, N., & Stramaglia, S. (2006). An invariance property of predictors in kernel-induced hypothesis spaces. *Neural Computation*, *18*, 749–759.
- Ancona, N., Marinazzo, D., & Stramaglia, S. (2004). Radial basis function approach to nonlinear granger causality of time series. *Physical Review E*, *70*, 056,221
- Baccalá, L. A., & Sameshima, K. (2001). Partial directed coherence: A new concept in neural structure determination. *Biol Cybernetics*, *84*, 463–474.
- Bedard, C., Kroger, H., & Destexhe, A. (2006). Model of low-pass filtering of local field potentials in brain tissue. *Physical Review E*, *73*, 051,911.
- Belitski, A., Gretton, A., Magri, C., Murayama, Y., Montemurro, M. A., Logothetis, N. K., et al. (2008). Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. *Journal of Neuroscience*, *28*, 5696–5709.
- Berens, P., Keliris, G. A., Ecker, A. S., Logothetis, N. K., & Tolias, A. S. (2008). Comparing the feature selectivity of the gamma-band of the local field potential and the underlying spiking activity in primate visual cortex. *Frontiers in Systems Neuroscience*, *2*, 2.
- Bernasconi, C., von Stein, A., Chiang, C., & König, P. (2000). Bidirectional interactions between visual areas in the awake behaving cat. *Neuroreport*, *11*, 689–692.
- Bragin, A., Engel, J., Wilson, C. L., Fried, I., & Buzsáki, G. (1999). High-frequency oscillations in human brain. *Hippocampus*, *9*, 137–142.
- Bressler, S. L., Richter, C. G., Chen, Y. H., & Ding, M. (2007). Cortical functional network organization from autoregressive modeling of local field potential oscillations. *Statistics in Medicine*, *26*, 3875–3885.
- Brovelli, A., Ding, M., Ledberg, A., Chen, Y., Nakamura, R., & Bressler, S. L. (2004). Beta oscillations in a large-scale sensorimotor cortical network: Directional influences revealed by granger causality. *Proceedings of the National Academy of Science USA*, *101*, 9849–9854.
- Brunel, N., & Wang, X. J. (2003). What determines the frequency of fast network oscillations with irregular neural discharges? I. Synaptic dynamics and excitation-inhibition balance. *Journal of Neurophysiology*, *90*, 415–430.
- Buzsáki, G. (2002). Theta oscillations in the hippocampus. *Neuron*, *33*, 325–340.
- Buzsáki, G. (2006). *Rhythms of the brain*. Oxford University Press.
- Buzsáki, G., & Draguhn, A. (2004) Neuronal oscillations in cortical networks. *Science*, *304*, 1926–1929.
- Canolty, R. T., Edwards, E., Dalal, S. S., Soltani, M., Nagarajan, S. S., Kirsch, H. E., et al. (2006). High gamma power is phase-locked to theta oscillations in human neocortex. *Science*, *313*, 1626–1628.
- Cardin, J. A., Carlén, M., Meletis, K., Knoblich, U., Zhang, F., Deisseroth, K., et al. (2009). Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature*, *459*, 663–667.
- Chavez, M., Martinerie, J., & Le Van Quyen, M. (2003). Statistical assessment of nonlinear causality: Application to epileptic EEG signals. *Journal of Neuroscience Methods*, *124*, 113–128.
- Chen, Y., Bressler, S. L., & Ding, M. (2006). Frequency decomposition of conditional granger causality and application to multivariate neural field potential data. *Journal of Neuroscience Methods*, *150*, 228–237.

- da Silva, F. L., Pijn, J. P., & Boeijinga, P. (1989). Interdependence of EEG signals: Linear vs. nonlinear associations and the significance of time delays and phase shifts. *Brain Topography*, 2, 9–18.
- Destexhe, A., & Sejnowski, T. J. (2003). Interactions between membrane conductances underlying thalamocortical slow-wave oscillations. *Physiological Reviews*, 83, 1401–1453.
- Eckhorn, R., & Thomas, U. (1993). A new method for the insertion of multiple microprobes into neural and muscular tissue, including fiber electrodes, fine wires, needles and micro-sensors. *Journal of Neuroscience Methods*, 49(3), 175–179.
- Gail, A., Brinksmeyer, H., & Eckhorn, R. (2004). Perception-related modulations of local field potential power and coherence in primary visual cortex of awake monkey during binocular rivalry. *Cerebral Cortex*, 14, 300–313.
- Geweke, J. (1982). Measurement of linear dependence and feedback between multiple time series. *Journal of the American Statistical Association*, 77, 304–313.
- Gourévitch, B., & Eggermont, J. J. (2007). Evaluating information transfer between auditory cortical neurons. *Journal of Neurophysiology*, 97, 2533–2543.
- Granger, C. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37, 424–438.
- Grassberger, P. (1988). Finite sample corrections to entropy and dimension estimates. *Physics Letters A*, 128, 369–373.
- Gray, C. M., König, P., Engel, A. K., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338, 334–337.
- Guo, S., Seth, A. K., Kendrick, K. M., Zhou, C., & Feng, J. (2008a). Partial granger causality—eliminating exogenous inputs and latent variables. *Journal of Neuroscience Methods*, 172, 79–93.
- Guo, S., Wu, J., Ding, M., & Feng, J. (2008b). Uncovering interactions in the frequency domain. *PLoS Computational Biology*, 4, e1000087.
- Harada, Y., & Takahashi, T. (1983). The calcium component of the action potential in spinal motoneurons of the rat. *Journal of physiology*, 335, 89–100.
- Hinrichs, H., Heinze, H., & Schoenfeld, M. (2006). Causal visual interactions as revealed by an information theoretic measure and fMRI. *NeuroImage*, 31, 1051–1060.
- Hlavackova-Schindler, K., Palus, M., Vejmelka, M., & Bhattacharya, J. (2007). Causality detection based on information-theoretic approaches in time series analysis. *Physics Reports*, 441, 1–46.
- Juergens, E., Guettler, A., & Eckhorn, R. (1999). Visual stimulation elicits locked and induced gamma oscillations in monkey intracortical- and EEG-potentials, but not in human EEG. *Experimental Brain Research*, 129(2), 247–259.
- Kahana, M. J., Seelig, D., & Madsen, J. R. (2001). Theta returns. *Current Opinion in Neurobiology*, 11, 739–744.
- Kaiser, A., & Schreiber, T. (2002). Information transfer in continuous processes. *Physica D: Nonlinear Phenomena*, 166, 43–62.
- Kamiński, M., & Blinowska, K. J. (1991). A new method of the description of the information flow in the brain structures. *Biological Cybernetics*, 65, 203–210.
- Kamondi, A., Acsády, L., Wang, X. J., & Buzsáki, G. (1998). Theta oscillations in somata and dendrites of hippocampal pyramidal cells *in vivo*: Activity-dependent phase-precession of action potentials. *Hippocampus*, 8, 244–261.
- Kayser, C., & Logothetis, N. K. (2009). Directed interactions between auditory and superior temporal cortices and their role in sensory integration. *Frontiers in Integrative Neuroscience*, 3, 1–11.
- Kayser, C., Montemurro, M. A., Logothetis, N., & Panzeri, S. (2009). Spike-phase coding boosts and stabilizes information carried by spatial and temporal spike patterns. *Neuron*, 61, 597–608.
- Kraskov, A., Stogbauer, H., & Grassberger, P. (2004). Estimating mutual information. *Physical Review E, Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, 69, 066138.
- Lachaux, J., Rodriguez, E., Martinerie, J., & Varela, F. (1999). Measuring phase synchrony in brain signals. *Human Brain Mapping*, 8, 194–208.
- Lisman, J. (2005). The theta/gamma discrete phase code occurring during the hippocampal phase precession may be a more general brain coding scheme. *Hippocampus*, 15, 913–922.
- Llinas, R., & Ribary, U. (1993). Coherent 40-hz oscillation characterizes dream state in humans. *Proceedings of the National Academy of Science USA*, 90, 2078–2081.
- Logothetis, N. K. (2003). The underpinnings of the bold functional magnetic resonance imaging signal. *Journal of Neuroscience*, 23(10), 3963–3971.
- Logothetis, N. K. (2008). What we can do and what we can not do with fMRI. *Nature*, 453, 869–878.
- Logothetis, N. K., Kayser, C., & Oeltermann, A. (2007). *In vivo* measurement of cortical impedance spectrum in monkeys: Implications for signal propagation. *Neuron*, 55, 809–823.
- Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412, 150–157.
- Magri, C., Whittingstall, K., Singh, V., Logothetis, N., & Panzeri, S. (2009). A toolbox for the fast information analysis of multiple-site LFP, EEG and spike train recordings. *BMC Neuroscience*, 10, 81.
- Marinazzo, D., Pellicoro, M., & Stramaglia, S. (2006). Nonlinear parametric model for granger causality of time series. *Physical Review E*, 73, 066,216.
- Mazzoni, A., Panzeri, S., Logothetis, N. K., & Brunel, N. (2008). Encoding of naturalistic stimuli by local field potential spectra in networks of excitatory and inhibitory neurons. *PLoS Computational Biology*, 4, e1000239.
- Mitzdorf, U. (1987). Properties of the evoked potential generators: Current source-density analysis of visually evoked potentials in the cat cortex. *International Journal of Neuroscience*, 33(1–2), 33–59.
- Montemurro, M. A., Rasch, M. J., Murayama, Y., Logothetis, N. K., & Panzeri, S. (2008). Phase-of-firing coding of natural visual stimuli in primary visual cortex. *Current Biology*, 18, 375–380.
- Montemurro, M. A., Senatore, R., & Panzeri, S. (2007). Tight data-robust bounds to mutual information combining shuffling and model selection techniques. *Neural Computation*, 19, 2913–2957.
- Nadasdy, Z. (2009). Information encoding and reconstruction from the phase of action potentials. *Frontiers in Systems Neurosci*, 3, 6.
- Nelken, I., Chechik, G., Mscis-Flogel, T. D., King, A. J., & Schnupp, J. W. H. (2005). Encoding stimulus information by spike numbers and mean response time in primary auditory cortex. *Journal of Computational Neuroscience*, 19, 199–221.
- Nolte, G., Ziehe, A., Nikulin, V., Schlögl, A., Krämer, N., Brismar, T., et al. (2008). Robustly estimating the flow direction of information in complex physical systems. *Physical Review Letters*, 100, 234,101.

- Pantazis, D., Nichols, T. E., Baillet, S., & Leahy, R. M. (2005). A comparison of random field theory and permutation methods for the statistical analysis of MEG data. *NeuroImage*, *25*, 383–394.
- Panzeri, S., Senatore, R., Montemurro, M. A., & Petersen, R. S. (2007). Correcting for the sampling bias problem in spike train information measures. *Journal of Neurophysiology*, *98*(3), 1064–1072.
- Panzeri, S., & Treves, A. (1996). Analytical estimates of limited biases in different information measures. *Network*, *7*, 87–107.
- Pearl, J. (2000). *Causality—models, reasoning, and inference*. Cambridge, UK: Cambridge University Press.
- Pettersen, K. H., & Einevoll, G. T. (2008). Amplitude variability and extracellular low-pass filtering of neuronal spikes. *Biophysical Journal*, *94*, 784–802.
- Quiroga, R. Q., Arnhold, J., & Grassberger, P. (2000). Learning driver-response relationships from synchronization patterns. *Physical Review E, Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, *61*(5 Pt A), 5142–5148.
- Roebroeck, A., Formisano, E., & Goebel, R. (2005). Mapping directed influence over the brain using granger causality and fMRI. *Neuroimage*, *25*, 230–242.
- Roopun, A. K., Kramer, M. A., Carracedo, L. M., Kaiser, M., Davies, C. H., Traub, R. D., et al. (2008). Temporal interactions between cortical rhythms. *Frontiers in Neuroscience*, *2*, 145–154.
- Schack, B., Klimesch, W., & Sauseng, P. (2005). Phase synchronization between theta and upper alpha oscillations in a working memory task. *International Journal of Psychophysiology*, *57*, 105–114.
- Schreiber, T. (2000). Measuring information transfer. *Physics Review Letters*, *85*, 461–464.
- Seth, A. K. (2005). Causal connectivity analysis of evolved neural networks during behavior. *Network*, *16*, 35–55.
- Seth, A. K., & Edelman, G. (2007). Distinguishing causal interactions in neural populations. *Network*, *19*, 910–933.
- Shannon, C. E. (1948). A mathematical theory of communication. *AT&T Bell Labs Technical Journal*, *27*, 379–423.
- Stam, C. J., & van Dijk, B. W. (2002). Synchronization likelihood: An unbiased measure of generalized synchronization in multivariate data sets. *Physica D*, *163*(3), 236–251.
- Stettler, D. D., Das, A., Bennett, J., Gilbert, C. D. (2002). Lateral connectivity and contextual interactions in macaque primary visual cortex. *Neuron*, *36*, 739–750.
- Strong, S. P., Koberle, R., de Ruyter van Steveninck, R. R., & Bialek, W. (1998). Entropy and information in neural spike trains. *Physical Review Letters*, *80*, 197–200.
- Theiler, J. (1986). Spurious dimension from correlation algorithms applied to limited time-series data. *Physical Review A*, *34*, 2427–2432.
- Victor, J. D. (2002). Binless strategies for estimation of information from neuronal data. *Physical Review E*, *66*, 51,903–51,918.
- Womelsdorf, T., Schoffelen, J. M., Oostenveld, R., Singer, W., Desimone, R., Engel, A. K., et al. (2007) Modulation of neuronal interactions through neuronal synchronization. *Science*, *316*, 1609–1612.