What do neural travelling waves tell us about information flow?

Andrea Alamia[a], Antoine Grimaldi[a], Frederic Chavane[b], Martin Vinck[c]

[a] *Cerco, CNRS Université de Toulouse, Toulouse 31059, France*
[b] *Institut de Neurosciences de la Timone, CNRS Aix-Marseille Université, Marseille 13385, France*
[c] *Ernst Strüngmann Institute (ESI) for Neuroscience, in Cooperation with the Max Planck Society, 60528 Frankfurt am Main, Germany*

Abstract

In many behavioral conditions, neural activity propagates within and across brain regions as traveling waves, revealing the importance of analyzing spatiotemporal dynamics in electrophysiological data. Most methods quantify such propagation by measuring spatial phase gradients, i.e., monotonic and ordered phase changes through space. Here, we demonstrate that the phase ordering in travelling waves is insufficient to determine the effective flow of information unambiguously. We demonstrate that, in some specific cases, the phase gradient indicates information propagation in the opposite direction than indicated by methods for causal inference. Using autoregressive modeling, we further show that such a discrepancy between the effective waves and the apparent waves measured via phase-based methods can, for example, be predicted by the sign of the projection from the lower to the higher nodes in the hierarchy. Together with an input signal in the lowest node, inhibitory bottom-up connections produce apparent waves propagating in the opposite, top-down direction. As a methodological solution, we show that Granger causality analysis can recover the information flow and its underlying causal structure, which can be used to disambiguate the "effective" flow.

1. Significant statement

Neural activity in the cerebral cortex shows a highly complex spatiotemporal organization that varies systematically across behavioral, cognitive, and physiological states, as recently described by the concept of neural travelling waves Yet, the exciting surge in research on traveling waves comes with some significant methodological limitations and interpretational issues concerning their patterns and the methods used to quantify them. Most analytical approaches are based on spatial phase gradients, which estimate the phase change through cortical space. Here, we demonstrate that the phase ordering in travelling waves is insufficient to determine the effective flow of information unambiguously and, in some cases, can produce results that are opposite to the actual propagation, as measured via causal methods (such as Granger analysis).

2. Introduction

Recent experimental work demonstrates that neural activity propagates between and within brain regions as traveling waves at macroscopic and mesoscopic scales (Muller et al. (2018); Ermentrout and Kleinfeld (2001); Bhattacharya et al. (2022a); Rubino et al. (2006); Hatsopoulos et al. (1998); Troy and Shusterman (2007); Benucci et al. (2007); Richardson et al. (2005); Campbell et al. (2024); Pesaran et al. (2018)). It is therefore important to consider the spatial and temporal component in neural recordings to characterize the putative role of such activity thoroughly (Zhang et al. (2018); Lubenov and Siapas (2009); Patel et al. (2012); Davis et al. (2020); Zeng et al. (2024); Galas et al. (2025); Luo et al. (2021); Ye et al. (2023); Bhattacharya et al. (2022b); Muller et al. (2014); Chemla et al. (2019a)). Several methods have been proposed in the literature to identify and quantify the waves' propagation direction (Das et al. (2023); Alexander et al. (2013); Alamia and VanRullen (2019); Schwenk and Alamia (2024)). Most of these methods rely on phase estimation to quantify the spatial gradients in a given frequency band and have provided relevant insights on the mechanisms associated with traveling waves (Davis et al. (2020); Zanos et al. (2015); Muller et al. (2014); Das et al. (2023)).

However, when considering traveling waves, it is important to acknowledge the distinction between the actual activity effectively propagating through the medium, the *effective wave*, and the propagation detected at the sensor level, the *apparent wave*. Inferring the effective wave from the apparent one is an under-determined problem with a non-unique solution (Budzinski et al. (2023); Kulkarni et al. (2020); Ermentrout (1998); Ermentrout and Kleinfeld (2001). In particular, one might wonder to what extent the phase gradients measured at the sensor level are indicative of the actual flow of information, defined here as the effective wave of activity that causally changes activity from one region to another.

Here, we use vector autoregressive (VAR) modeling to reveal cases of a systematic discrepancy between effective and apparent wave propagation direction. Previous work successfully applied VAR approaches to model neural responses and physiological signals (Seth et al., 2015; Vinck et al., 2015; Pourahmadi and Noorbaloochi, 2016; Gor-

rostieta et al., 2013; Ghosh et al., 2019; Marinazzo et al., 2011; He et al., 2019; Dowdall and Vinck, 2023), but this approach has never been applied for traveling waves analysis. In this study, we test different connections between nodes in VAR models, with or without temporal delays, and having a different number of nodes. All models were fed with an input to the lowest node, which, in principle, should generate waves propagating from the lower to the higher node, producing gradual phase shifts. We compared these phase shifts with the actual flow of information revealed with Granger Causality analysis, a technique for causal inference widely used in the neurosciences (Brovelli et al., 2004a; Bastos et al., 2015; Seth et al., 2015; Vinck et al., 2015; Ding et al., 2006).

The following sections of the paper are organized as follows: we first describe the mathematical formalism used to define the VAR models and the four connectivity conditions we investigated (i.e., bottom-up and top-down, either excitatory or inhibitory). We then define cases with different temporal delays and one model inspired by a biologically plausible implementation. In all cases, we consistently compare the phase-based analysis results with the one obtained with Granger causality analysis to highlight discrepancies between the two methods.

3. Methods

3.1. General VAR models

Consider a k-variate VAR model of order T, described by the system of discrete-time equations:

$$X_t = \begin{bmatrix} x_{1,t} \\ \vdots \\ x_{k,t} \end{bmatrix} = \sum_{\tau=1}^{T} A_\tau \cdot X_{t-\tau} + \epsilon_t \tag{1}$$

Here, X_t is a vector describing the evolution of a set of k variables. A_τ represents a $(k \times k)$-coefficient matrix that describes the influence of the different past states, at lag τ, on the current state. The innovation term ϵ_t is a vector of the same dimension as X_t where each component is assumed to be white noise with zero mean and covariance matrix Σ.

Next, consider a simplification of this system where the nodes are sequentially organized, such that there are only direct interactions from a node i with the nodes $i-1$ and $i+1$. In this case, $\forall \tau \in [0, T], A_\tau$ is a tridiagonal matrix with a specific set of coefficients defined as:

$$a_{i,j}^\tau = \begin{cases} \alpha_i^\tau, & \text{if } i = j \\ \mathrm{s}_{FF}^\tau \beta_{i,j}^\tau, & \text{if } j = i-1 \\ \mathrm{s}_{FB}^\tau \beta_{i,j}^\tau, & \text{if } j = i+1 \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

where $\forall \tau \in [0, T]$, α_i^τ, $\beta_{i,j}^\tau \in (\mathbb{R}^+)^2$ are random variables drawn from $\mathcal{N}_{\text{trunc+}}(0, 2)$, a Gaussian distribution truncated for all values below zero, and s_{FF}^τ, $\mathrm{s}_{FB}^\tau \in \{-1, 1\}^2$ determine the sign of the feedforward and feedback connections.

If we take a 2-variate VAR model, this yields in total 4 possible cases illustrated in Figure 1A, conditional on the value of s_{FB}^τ. In addition, all parameters connecting subsequent nodes have the same absolute value, such that:

$$\beta_{i,j}^\tau = \beta_{j,i}^\tau$$

In the following, we will first consider the simplest VAR(1) case. Then, we will introduce models with higher orders, i.e., from VAR(2) to VAR(4). The last section will describe the biologically plausible cortical model.

3.2. The VAR(1) case

Starting from the general case described above, we first consider the condition without given temporal delays, i.e. using only the previous state to describe the current one :

$$X_t = A_1 \cdot X_{t-1} + \epsilon_t \tag{3}$$

In this case, we test models with k=3 or k=7 nodes. For each of the four cases described above (figure 1A), we run 1000 simulations of 500 time steps each. In each simulation, we determine the stability of the model by assessing that all eigenvalue modules associated with the companion matrix were $|\lambda_i| < 1$. We discard all runs in which the system was unstable.

The first node is fed an input signal, a series of independent, pseudo-random values drawn from a Gaussian distribution $\mathcal{N}(0, 1)$.

3.3. From VAR(2) to VAR(4)

We also consider the case with fixed temporal delays, i.e., a VAR model with order $p > 1$. In this case, the VAR model is described by equation 1.

Specifically, we investigated four different conditions. Starting from the VAR(1) case, we considered three cases in which $\tau \in \{2, 3, 4\}$. In addition, we also considered the VAR model described by the following equation:

$$X_t = \sum_{\tau=1}^{4} A_\tau \cdot X_{t-\tau} + \epsilon_t \tag{4}$$

All the other parameters (i.e., number of simulations, time-steps, nodes, and stability assessment) were the same as in the case without temporal delays.

3.4. Simulations for a relevant cortical model

We also apply our analysis to a biologically plausible model. Based on the generic VAR model given in equation 1, we can introduce a specific set of parameters to match a given connectivity matrix or other existing models. Here, we set the parameters to match the dynamical system described in Alamia and VanRullen (2019).

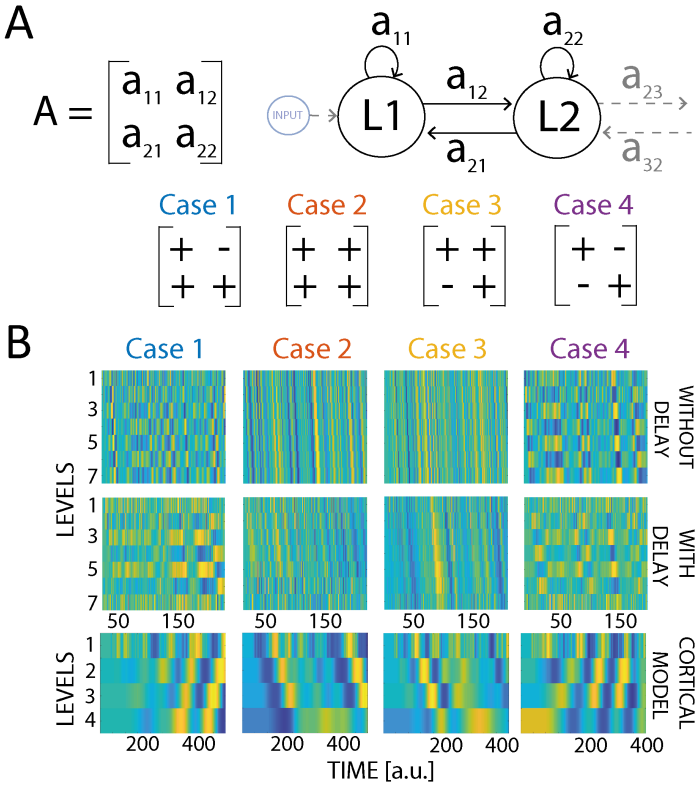

Figure 1: Figure 01 A) Schematic representation of the VAR models. We considered four cases in which the forward and backward connection have either positive or negative signs. B) 2D representations of the activity in the four cases in all implementations. The x-axis and the y-axis represent the time and the levels (i.e., nodes), respectively. The color (in arbitrary units) reflects the activity over time in each node.

We started from the equations of a two layers system as described in (Alamia and VanRullen, 2019):

$$x_{1,t} - x_{1,t-1} = \frac{1}{\tau_{fw}} e_1(t-\Delta t) + \frac{1}{\tau_{bw}}(x_{2,t-\Delta t} - x_{1,t-1}) \quad (5)$$

where τ_{fw} and τ_{bw} are time constants, $e_1(t)$ is the prediction error in the first layer:

$$e_1(t) = x_{0,t} - x_{1,t-\Delta t}$$

and $x_{0,t}$ is the bottom-up input to the network at time t. This, eventually, simplifies to:

$$\begin{aligned} x_{1,t} = (1 - \frac{1}{\tau_{bw}}) \cdot x_{1,t-1} + \frac{1}{\tau_{fw}} \cdot x_{0,t-\Delta t} \\ + \frac{1}{\tau_{bw}} x_{2,t-\Delta t} - \frac{1}{\tau_{fw}} \cdot x_{1,t-2\Delta t} \end{aligned} \quad (6)$$

For this dynamical system, we have an equivalent VAR($2\Delta t$) model where all coefficient matrices have zero values except for $A^1, A^{\Delta t}$, and $A^{2\Delta t}$. Thus, we have:

$$X_t = A_1 \cdot X_{t-1} + A_{\Delta t} \cdot X_{t-\Delta t} + A_{2\Delta t} \cdot X_{t-2\Delta t} + \epsilon_t \quad (7)$$

To match the dynamical system given by equation 5, we need A_1 to be a diagonal matrix with the diagonal coefficients equal to $(1 - \frac{1}{\tau_{bw}})$, $A_{2\Delta t}$ to be another diagonal matrix with the diagonal coefficients equal to $-\frac{1}{\tau_{fw}}$ and $A_{\Delta t}$ to be a tridiagonal matrix with 0 on the diagonal, $\frac{1}{\tau_{fw}}$ on the lower diagonal and $\frac{1}{\tau_{bw}}$ on the upper diagonal. As in Alamia and VanRullen (2019), we set $\tau_{fw} = 30$ ms and $\tau_{bw} = 200$ ms, these values granted a stable behavior and can be interpreted from a biological perspective as time constants. From this VAR($2\Delta t$) model, one can compute the frequency and oscillatory behavior, which depend on its parameters τ_{fw}, τ_{bw} and Δt, as demonstrated in Alamia and VanRullen (2019). Note that equation 7 describes the 2-layer model given by equation 6 but does not hold any assumption on the dimensionality of X. It can be easily generalized to k nodes by increasing the size of vector X and of the coefficient matrices A with the same structure as the one described in this case.

Considering this implementation, we tested models having k=4 nodes. For each of the four cases described above (Fig.1A), we run 300 simulations of 500-time steps with independent Gaussian noise $\epsilon(t)$.

3.5. Quantifying waves' direction of propagation

In all simulations in which the model was stable (i.e., all eigenvalues $|\lambda_i| < 1$), we used two different approaches to quantify the activity's direction of propagation. The first method was based on the estimates of the signal's instantaneous phase; the second one applied Granger Causality to assess the flow's direction.

3.5.1. Phase-based analysis

At the end of each simulation, we computed the Hilbert transform $H(x_k)$ of each signal x_k for each k node and calculated the complex vector obtained by computing the phase difference between subsequent nodes at any time point t:

$$z_{k,t} = exp(i * [\psi(H(x_{k,t})) - \psi(H(x_{k+1,t}))]) \quad (8)$$

where $\psi(H(x_{k,t}))$ is the phase of the Hilbert-transformed signal x_t at the node k. We then sum all complex vectors over time and compute the average phase difference between nodes:

$$\Psi_k = \sum_{t=1}^{T} \psi(z_{k,t}) \quad (9)$$

The sign of Ψ_k provides the phase shift between subsequent nodes: if positive, it indicates a wave propagating forward (from node k to node $k+1$); vice versa, a negative value reveals activity propagating from the higher to the lower nodes.

3.5.2. Granger analysis

We used the multivariate Granger causality analysis (MVGC) toolbox in MATLAB (MathWorks) to compute the Granger causality between subsequent nodes in all cases (Barnett and Seth (2014)). For each condition and stable models, we first fit the model's parameters, considering each simulation as a trial. We estimated the best model's order by optimizing the Aikaike Information Criterion (we obtained similar results using the Bayesian Information Criterion). Besides estimating the model's parameters, such fit provides the covariance matrix of the residuals, which is next used to calculate the autocovariance sequence according to the estimated VAR model. Lastly, the pair-wise time-domain Granger causalities are computed for each pair of x_i from the covariance sequence. Significant effects were estimated against the null hypothesis of zero causality based on a theoretical asymptotic null distribution (Barnett and Seth (2014)).

4. Results

4.1. The VAR(1) case

We first assessed the stability of the 1000 randomly generated models for each of the four conditions. As shown in details in table 1, around 35% of simulations were stable for $k = 3$ and 25% for $k = 7$ for the four conditions. We considered only stable models for the subsequent analyses.

4.1.1. Phase-based analysis

As shown in figure 3A, we found consistent results for k=3 and k=7 nodes. Specifically, irrespective of the number of nodes k, we found a negative phase difference between all subsequent nodes in the first and fourth case (i.e., when then $a_{i,i+1}$ coefficient is negative), indicating a propagation backward from higher to lower nodes. Conversely, we found the opposite direction flow in the second and third cases (having a positive sign in the $a_{i,i+1}$ coefficient), in which we observed a positive phase difference between subsequent nodes. This result was supported by a Von Mises test (V-test) for non-uniformity of circular data considering the mean direction in any pair of nodes (in all conditions and pair of nodes, we obtained for k=3 all V-value > 250 and $p < 0.0001$; and for k=7 all V-value > 180 and $p < 0.0001$).

k=3	Case1	Case2	Case3	Case4
$\tau = 1$	40.9	27.5	38.7	31.4
$\tau = \{1, 2\}$	22.1	16.3	22.1	14.0
$\tau = \{1, 3\}$	24.9	15.5	25.4	13.6
$\tau = \{1, 4\}$	24.8	14.8	26.2	15.3
$\tau = \{1, .., 4\}$	12.8	7.5	11.7	8.0
k=7	Case1	Case2	Case3	Case4
$\tau = 1$	30.4	19.7	31.4	21.4
$\tau = \{1, 2\}$	18.3	10.8	15.5	11.7
$\tau = \{1, 3\}$	16.5	12.3	15.0	12.7
$\tau = \{1, 4\}$	19.0	13.3	20.2	11.4
$\tau = \{1, .., 4\}$	8.8	5.3	8.9	5.9

Table 1: Models' stability. The table reports the percentage of stable simulations for each condition (columns) with different temporal delays (rows). $\tau = 1$ denotes the case without temporal delay.

All in all, the results from the phase-based analysis suggest that the direction of propagation, as measured by the phase difference between subsequent nodes, depends on the sign of the $a_{i,i+1}$ coefficient. A positive influence from the i to the $i+1$ node determines a flow propagating from lower to higher nodes, whereas a negative modulation produces a direction flowing in the opposite direction, from higher to lower regions.

4.1.2. Granger analysis

We then applied a Granger analysis to all nodes' activity to assess the propagation direction and reveal the causal information flow. We fitted the model's order for k=3 and k=7 nodes, optimizing Akaike's criterion, and we obtained a fit of the third and fifth order, respectively, for all conditions. The results reported below don't change qualitatively by imposing the model's order to equal 1 before fitting the model's parameters. Figure 3A shows the time-domain Granger Causality (GC) value between all nodes. Asterisks reveal a significant effect when compared against a theoretical asymptotic null distribution (Barnett and Seth (2014)). For k=3, we found GC values of 2.566 ± 0.38 from node 1 to node 2 and 1.019 ± 0.06 from node 2 to node 3. All other GC values were smaller than 0.08. Similarly, when k=7, we found the larger GC value from node 1 to node 2 (2.880 ± 0.31), whereas all other GC values from node i to node $i+1$ (with $2 \leq i \leq 7$) were between 0.925 ± 0.14 and 0.988 ± 0.15. All the GC values

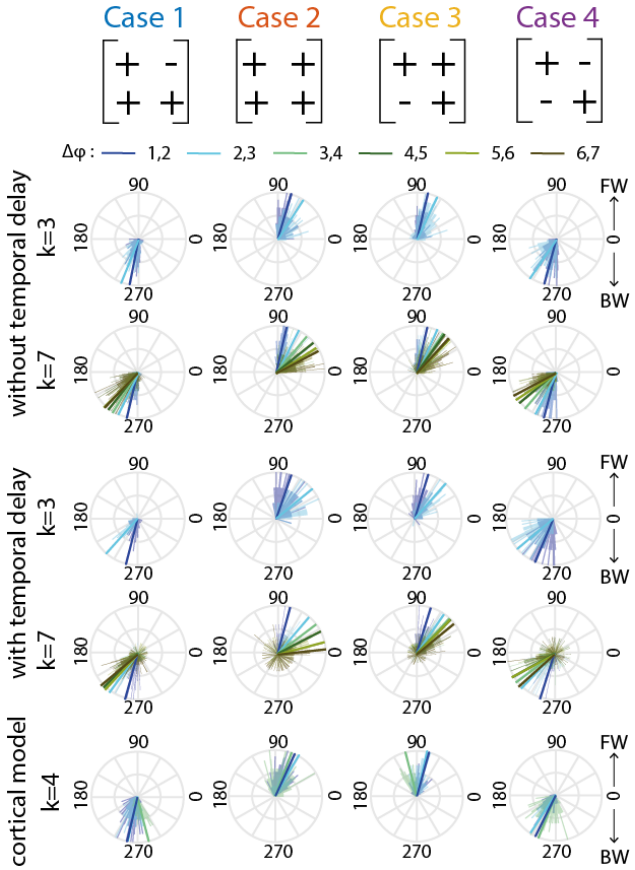

Figure 2: Results of the phase analysis. Each panel shows the distribution of phase differences between consequent nodes in the four cases (represented in each column). In cases 2 and 3, differences between 0 and 180 degrees indicate waves propagating from lower to higher regions (i.e., forward). In cases 1 and 4, values between 180 and 360 degrees indicate propagation from higher to lower nodes. This result holds for all simulations with and without temporal delays and in the biologically plausible cortical model. The case without temporal delay corresponds to the case when $\tau = 1$, whereas the case displayed with temporal delay is with $\tau = 1, 2$. Very similar results are obtained when considering the other temporal delays (i.e., $\tau = \{1, 2, 3, 4\}$

.

connecting nodes in the opposite, top-down direction, that is from node i to node $i+1$ (with $2 \leq i \leq 7$), were smaller than 0.02 ± 0.02.

In sum, the results based on the Granger analysis reveal a stronger propagation in the opposite direction than the phase-based analysis when the $a_{i,i+1}$ coefficient is negative (cases 1 and 4 in figure 1A). This discrepancy suggests that the Granger and the phase-based analysis are sensitive to different effects: the first one reveals the causal flow of information, propagating the input information from lower to higher regions, whereas the second one captures the absolute sign of the connectivity of the feedforward system, independent of the transfer of input information.

4.2. From Var(2) to Var(4)

We then tested whether the effect could be generalized to cases with temporal delays. First, we assessed the stability of the 1000 simulations we performed with randomized coefficients. As reported in table 1, around 20% and 15% of the simulations were stable for $k = 3$ and $k = 7$, respectively. When considering several delays at once (last row, $\tau = [1, 4]$), the stability dropped to less than 10% of the simulations, leaving, however, more than 50 runs in the lowest cases (cases 2 and 4). As previously, we considered only stable models for the subsequent analyses.

4.2.1. Phase-based analysis

As in the case without temporal delays, we found consistent results for k=3 and k=7 nodes (figure 2). In particular, we found a backward propagation in the first and fourth cases, as confirmed by a negative phase difference between all subsequent nodes, but the opposite direction in the second and third cases, corroborating our previous results. These results were very consistent across the different delays we considered, and, for the sake of space, we reported in figure 2 only those for $\tau = 2$. As in the previous analysis, a Von Mises test (V-test) for non-uniformity of circular data considering the mean direction in any pair of nodes confirmed the results: in all conditions and pair of nodes, we obtained for k=3 all V-value > 120 and $p < 0.0001$ when considering $\tau = [1, 4]$, and V-value > 65 and $p < 0.0001$ when having several delays. Similarly, for

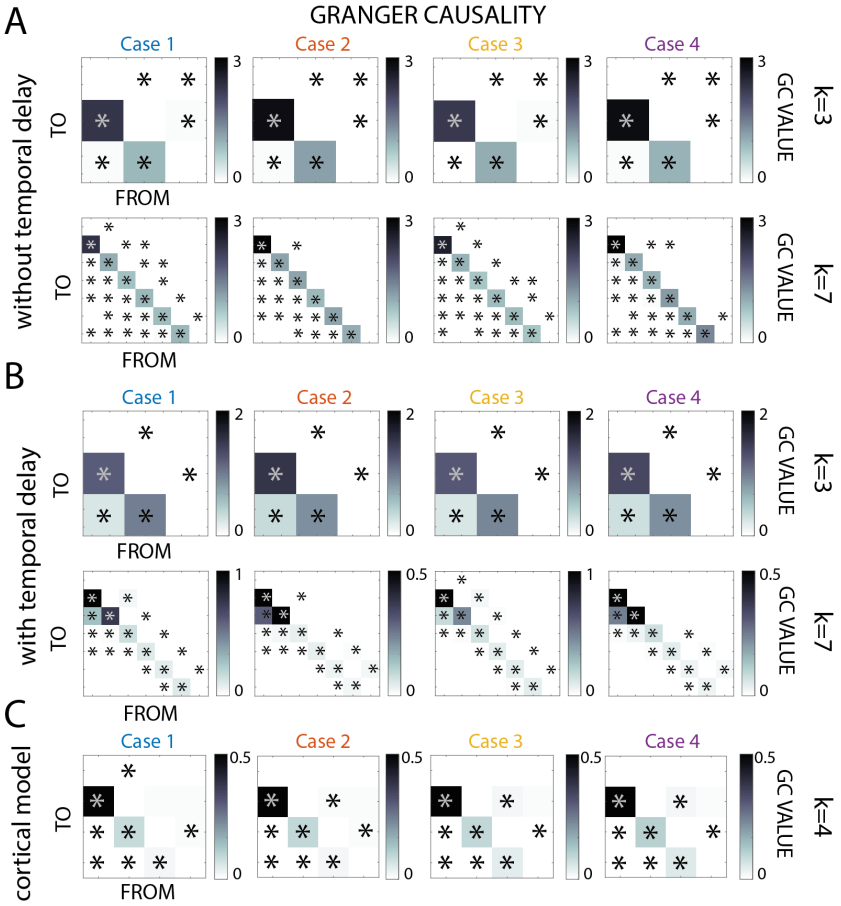

Figure 3: Results of the Granger analysis. Pairwise Granger causalities between all the nodes in the four cases, without and with temporal delays (panels A and B, respectively), and in the biologically plausible cortical model (panel C). In all cases, different from the phase analysis, the results reveal a stronger directionality in the bottom-up direction, from lower to higher nodes. Asterisks denote significant effects against the null hypothesis of non-causality.

k=7 and $\tau = [1, 4]$, we obtained all V-value > 50 (with larger values ≈ 110 for lower nodes, but invariably significant) and $p < 0.0001$, and V-value > 20 when considering several delays.

As in the case without temporal delays, the results of our simulations corroborated the finding that a positive sign from the node i to node $i + 1$ produces a phase difference consistent with a forward flow, whereas a negative one leads to a phase difference corresponding to a flow from higher to lower regions.

4.2.2. Granger analysis

The Granger analysis was applied to all nodes' activity to estimate the causal propagation direction. We fitted the model's order for k=3 and k=7 nodes for each temporal delay, optimizing Akaike's criterion. For a single temporal delay, i.e., $\tau \in \{2, 3, 4\}$, we obtained orders between 3 and 5, and we obtained an order equal to 8 when considering all the temporal delays, i.e., $\tau = \{1, .., 4\}$.

Considering the k=3 case, we found very similar time-domain Granger Causality (GC) values when having a single temporal delay. For instance, when $\tau = 2$, we found 1.377 ± 0.10 from node 1 to node 2, and 0.88 ± 0.20 from node 2 to node 3, and 0.284 ± 0.04 from node 1 to node 3. All other GC values were smaller than 0.05. When considering the case with all temporal delays, we obtained relatively smaller values: 0.804 ± 0.04 from node 1 to node 2, and 0.56 ± 0.05 from node 2 to node 3, and 0.10 ± 0.02 from node 1 to node 3.

Regarding k=7 and $\tau = 2$ (very similar values were observed in the other cases having one single temporal delay), we found the larger GC values in lower nodes: progressively from node 1 to node 7 we reported 1.433 ± 0.31, 0.768 ± 0.11, 0.16 ± 0.06, and then all other values were smaller than 0.07 ± 0.03. Similarly, when $\tau = \{1, .., 4\}$, we found value decreasing from lower to higher nodes: 0.876 ± 0.08, 0.346 ± 0.03, 0.097 ± 0.2, and then all other values were smaller than 0.05 ± 0.02.

Altogether, the Granger analysis confirms that the causal flow of information propagates the input information from lower to higher regions irrespective of the sign of the $a_{i,i+1}$ coefficient. Next, we investigate whether this effect, which proved consistent with and without temporal delays, holds true with a more biologically plausible model, grounded in the predictive coding framework (Alamia and

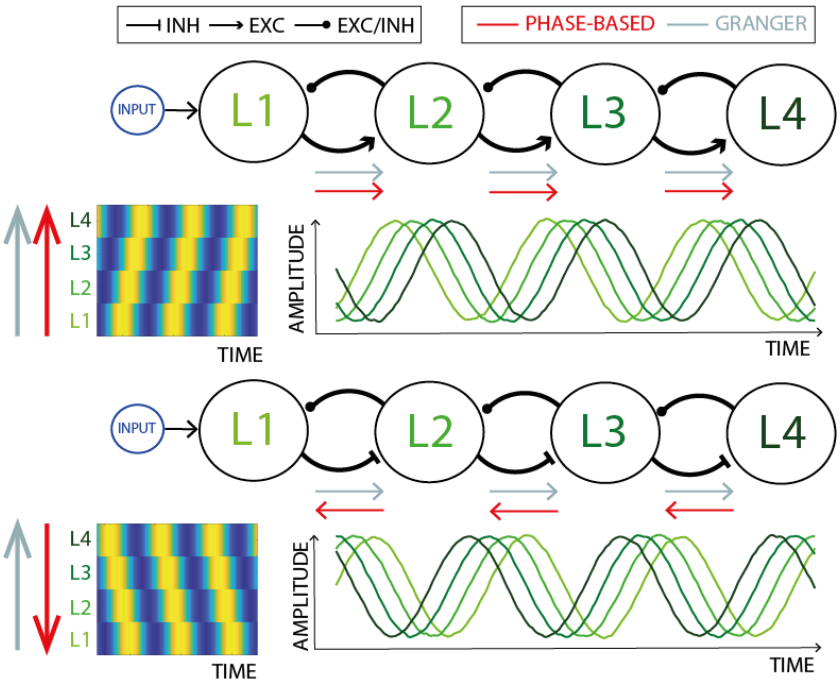

Figure 4: Summary of the results. The upper and lower plots outline the different cases from Figure 1. When the 'forward' connections are inhibitory, the phase-based and the Granger-based analyses provide different results.

VanRullen (2019)) and including temporal delays between subsequent nodes.

4.3. Simulations for a relevant cortical model

We simulated 300 instances of the model for 500 time steps for each of the four conditions. The model had $k = 4$ nodes, and its parameters were chosen to be biologically interpretable and with a stable behavior (see Methods).

4.3.1. Phase-based analysis

The results based on the phase analysis support the same conclusions as in the previous simualtions, as shown in figure 3B. In particular, we observed opposite propagation directions depending on the sign of the $a_{i,i+1}$ coefficient. In the first and fourth cases, when the $a_{i,i+1}$ coefficient is negative, we observe a consistent, negative phase difference across subsequent nodes, revealing a propagation from higher to lower nodes. On the other hand, a positive $a_{i,i+1}$ coefficient (i.e., second and third case) generates the opposite, forward direction flow from lower to higher areas, as confirmed by a positive phase difference between subsequent nodes. These results were confirmed by a V-test for non-uniformity of circular data performed over the 300 simulations and considering the mean direction in any pair of nodes (in all conditions and pair of nodes V-value$>$ 280 and $p < 0.0001$).

4.3.2. Granger analysis

We then estimated the Granger causalities between all subsequent nodes. We assessed the best model's order for each condition, having an order of 19 for the first condition and 18 for all the other ones. Figure 3C illustrates the time-domain Granger Causality (GC) values between all nodes. Asterisks reveal a significant effect when compared against a theoretical asymptotic null distribution (Barnett and Seth (2014)). Across conditions, we found decreasing GC values across the hierarchy: 1.407 ± 0.36 from node 1 to node 2, 0.340 ± 0.03 from node 2 to node 3, and 0.05 ± 0.01 from node 3 to node 4. All other GC values (i.e., in the opposite direction and between non-subsequent nodes) were smaller than 0.02 ± 0.01.

5. Discussion

5.1. Summary of the results

Neural activity propagates as traveling waves within and between brain regions (Muller et al. (2018); Ermentrout and Kleinfeld (2001); Hatsopoulos et al. (1998); Troy and Shusterman (2007); Zhang et al. (2018)), and different methods have been proposed to reliably measure the waves' propagation direction (Das et al. (2023); Alexander et al. (2013); Alamia and VanRullen (2019); Muller et al. (2014)). In this work, we implemented simulations based on vector auto-regressive (VAR) modeling, which has been used in several studies for neural simulations and provides a good approximation of neural dynamics (Seth et al. (2015); Vinck et al. (2015); Pourahmadi and Noorbaloochi (2016); Gorrostieta et al. (2013); Ghosh et al. (2019); Marinazzo et al. (2011); He et al. (2019)). Specifically, we compared the outcomes of two methods to estimate the direction of propagation with different connectivity matrices: one method based on the phase estimate

and one based on Granger causal connectivity. Our simulations show that these two methods provide different results in specific conditions, as summarized in figure 4. In particular, we found that when the 'forward' connection from node i to node $i+1$ is excitatory, both methods report the same direction from lower to higher areas, irrespective of the sign of the top-down connection. However, when the forward connection is negative (i.e., inhibitory), we observed a divergence between the two methods' results, with the phase-based one reporting a propagation in the opposite, top-down direction, i.e., from higher to lower areas. This discrepancy reveals that, despite the effective flow of information propagating in the bottom-up direction, as confirmed by the Granger analysis, the gradient of phases indicates a propagation in the opposite, top-down direction, possibly providing a misleading interpretation of the results.

5.2. Granger Causality and phase-based analyses

Granger Causality analysis has been widely used in the neurosciences for many applications (Marinazzo et al. (2011); Brovelli et al. (2015, 2004b); Marinazzo et al. (2008); Shojaie and Fox (2022); Cekic et al. (2018); Stokes and Purdon (2017); Seth et al. (2015); Ding et al. (2006)); however, it has not been applied with a focus on detecting traveling waves before. When applied to a linear system, it provides a well-defined solution with uniquely separated contributions (Seth et al. (2015); Marinazzo et al. (2008)), and it reliably captures the causal directionality following the propagation of the input provided to the system. On the other hand, phase-based methods have been widely used to quantify traveling waves, with various differences between specific implementations (Das et al. (2023); Alexander et al. (2013); Alamia and VanRullen (2019); Muller et al. (2014)), spanning from macroscale to mesoscale recordings. Considering the distinction between effective waves, i.e., the actual activity propagating and significantly affecting the activity of the recipient structures through the medium, and the apparent wave, i.e., the propagation detected at the sensor level, we can consider that the Granger analysis reveals the former, whereas phase-based methods the latter. Interestingly, we noticed that this distinction shares similarities with the difference between effective and functional connectivity (Friston (2011); Stephan et al. (2009)), with the former resting on a causal model generating the observed data and the latter revealing statistical correlations between brain regions' activities. Our simulations reveal a difference that can arise between the two approaches: differently than the phase-based approach, Granger Causality proved invariant to sign reversal in the connectivity matrix. Importantly, this discrepancy between methods can prove informative of the system's connectivity, as it emerges specifically with an inhibitory connection in the same direction as the effective wave.

5.3. Biological plausibility of our results

Our results reveal that, under certain conditions, a significant difference can arise in the identified direction of propagation between the two methods. Even though the results proved to be consistent across different implementations with and without temporal delays and in a model grounded in the Predictive Coding framework (Alamia and VanRullen (2019)), our results are based on relatively simple linear models. Given that the cortical dynamics are supposedly generated by non-linear systems (Vinck et al., 2023; Freeman and Skarda, 1985), one may wonder whether such a discrepancy can be generally found in more biologically plausible systems, and whether Granger-causality methods still provides a solution in those cases. We argue this may be the case for two reasons: First, linear dynamics may, in many cases, dominate inter-areal connectivity measures due to the strong influence of afferent connection on synaptic potentials in downstream areas (Pesaran et al., 2018; Dowdall et al., 2023; Dowdall and Vinck, 2023; Vinck et al., 2023; Schneider et al., 2021) Second, local oscillatory network dynamics may also be well approximated by linear dynamics, e.g. a damped harmonic oscillatory driven by noise (Spyropoulos et al., 2022). In addition, phase direction can be an emergent property of complex and dense set of multi-direction propagation Budzinski et al. (2023).

We argue that the other assumption, namely that the connections between nodes show heterogeneous signs (i.e. both inhibitory and excitatory), is realistic: Relatively well-structured inhibitory connections can be found in the brain, for example, in the cerebellum (Rieubland et al., 2014). Regarding cortical networks, while feedforward and feedback projections are excitatory (Callaway (2004); Shao and Burkhalter (1996)), we note that feedforward connections target both inhibitory and excitatory neurons, generating both IPSPs and EPSPs in principal cells in a downstream area. It is now well documented that paradoxical emergent signs of activation can arise from balanced E/I networks Tsodyks et al. (1997), as well demonstrated by inhibitory-stabilized networks Ozeki et al. (2009), with long-range excitation leading to decrease of activation Chemla et al. (2019b) In particular, high-frequency inputs have been suggested to predominantly target fast-spiking inhibitory interneurons through feedforward connections (Schneider et al. (2023); Spyropoulos et al. (2024)). In light of these results, our conclusions can generalize to relevant experimental conditions and neural recordings that may lead to wrong assumptions about the effective flow of information.

Actually, experimental evidence suggested that feedback activity may have a predominantly inhibitory influence on the targeted areas ((Olsen et al., 2012; Isaacson and Scanziani, 2011; Vezoli et al., 2020)). Taking this into account, it is possible to consider our implementation from a symmetrically opposite point of view, in which the input consists of endogenous activity propagating backward

from higher regions, e.g., in line with Hierarchical Predictive Coding theories Rao and Ballard (1999); Rauss et al. (2011). In this case, the discrepancy between Granger-based and phase-based methods would occur when feedback connections have a net inhibitory effect (Briggs, 2020; Murphy et al., 1999; Wilmes and Clopath, 2019), thus generating misleading results. In other words, feedback-effective inhibition may result in an apparent traveling wave propagating in the opposite, feedforward direction.

Further experimental work and analyses are necessary to investigate whether previously observed travelling waves indicate effective information flow or not. For example, given that previous work reported discrepancies between phase-based and Granger-causal methods in the beta-frequency range (Brovelli et al., 2004a) and beta is associated with inhibitory top-down influences (Bastos et al., 2020), it is plausible that the direction of beta travelling waves differs from the effective information flow.

5.4. Limitations and future work

The conclusions from our simulations may be relevant for most papers investigating traveling waves with phase-based methods. However, it is important to consider our work's methodological limitations. On the one hand, as mentioned above, we investigated linear VAR models, which capture neural dynamics well (Seth et al. (2015); Vinck et al. (2015); Pourahmadi and Noorbaloochi (2016); Gorrostieta et al. (2013); Marinazzo et al. (2011)), but may fail in capturing the whole spectra of neural activity. Similarly, it is possible that models with a broader range of parameters, such as temporal delays or long-range connections, may reveal cases in which the discrepancy between Granger and phase-based methods occurs due to excitatory rather than inhibitory connections. Future work will be needed to explore more realistic modeling, which would allow us to better capture the richer cortical dynamics while also fully exploring the range of parameter spaces, including the strength of the connection and different temporal delays. This more comprehensive approach would allow us to take one step closer to proper inference and provide a complete outlook on the methods to quantify waves' propagation. In principle, our conclusion could also generalize to systems outside of the neurosciences: any system with negative feedback would supposedly provide similar misleading results. In any case, the purpose of this study is to show that it is technically possible to have a dissociation between the direction of effective and apparent traveling waves.

5.5. Conclusion

All in all, this study highlights the importance of distinguishing between *effective* traveling waves, reflecting the actual flow of information, and *apparent* traveling waves, as measured indirectly via phase-based analysis. Our results reveal that apparent waves can propagate in the opposite direction compared to the effective one, and we call for caution when interpreting results based on phase methods alone. We advocate for using additional methods to better infer the underlying causal structure generating the measured phase gradient. One practical recommendation is to complement the phase-based analysis by fitting Granger-causality models to the data: integrating both approaches provides a more comprehensive view of the flow of information.

Acknowledgements

AA, FC and MV conceived of the study. AA and AG analyzed data. AA and MV wrote the paper, with comments from the other authors. MV was supported by an ERC starting grant (850861) SPATEMP; DFG VI Grants (908/5-1 and 908/7-1); an NWO VIDI Grant; and the Dutch Brain Interface Initiative (DBI2). AA was supported by ERC starting grant (OSCI-PRED, no. 101075930). Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the European Research Council (ERC). Neither the European Union nor the granting authority can be held responsible for them.

References

Alamia, A., VanRullen, R., 2019. Alpha oscillations and traveling waves: Signatures of predictive coding? PLoS Biology 17, e3000487.

Alexander, D.M., Jurica, P., Trengove, C., Nikolaev, A.R., Gepshtein, S., Zvyagintsev, M., Mathiak, K., Schulze-Bonhage, A., Ruescher, J., Ball, T., et al., 2013. Traveling waves and trial averaging: the nature of single-trial and averaged brain responses in large-scale cortical signals. Neuroimage 73, 95–112.

Barnett, L., Seth, A.K., 2014. The mvgc multivariate granger causality toolbox: a new approach to granger-causal inference. Journal of neuroscience methods 223, 50–68.

Bastos, A.M., Lundqvist, M., Waite, A.S., Kopell, N., Miller, E.K., 2020. Layer and rhythm specificity for predictive routing. Proceedings of the National Academy of Sciences 117, 31459–31469.

Bastos, A.M., Vezoli, J., Bosman, C.A., Schoffelen, J.M., Oostenveld, R., Dowdall, J.R., De Weerd, P., Kennedy, H., Fries, P., 2015. Visual areas exert feedforward and feedback influences through distinct frequency channels. Neuron 85, 390–401.

Benucci, A., Frazor, R.A., Carandini, M., 2007. Standing waves and traveling waves distinguish two circuits in visual cortex. Neuron 55, 103–117.

Bhattacharya, S., Brincat, S.L., Lundqvist, M., Miller, E.K., 2022a. Traveling waves in the prefrontal cortex during working memory. PLoS computational biology 18, e1009827.

Bhattacharya, S., Donoghue, J.A., Mahnke, M., Brincat, S.L., Brown, E.N., Miller, E.K., 2022b. Propofol anesthesia alters cortical traveling waves. Journal of Cognitive Neuroscience 34, 1274–1286.

Briggs, F., 2020. Role of feedback connections in central visual processing. Annual review of vision science 6, 313–334.

Brovelli, A., Chicharro, D., Badier, J.M., Wang, H., Jirsa, V., 2015. Characterization of cortical networks and corticocortical functional connectivity mediating arbitrary visuomotor mapping. Journal of Neuroscience 35, 12643–12658.

Brovelli, A., Ding, M., Ledberg, A., Chen, Y., Nakamura, R., Bressler, S.L., 2004a. Beta oscillations in a large-scale sensorimotor cortical network: directional influences revealed by granger causality. Proceedings of the National Academy of Sciences 101, 9849–9854.

Brovelli, A., Ding, M., Ledberg, A., Chen, Y., Nakamura, R., Bressler, S.L., 2004b. Beta oscillations in a large-scale sensorimotor cortical network: directional influences revealed by granger causality. Proceedings of the National Academy of Sciences 101, 9849–9854.

Budzinski, R.C., Nguyen, T.T., Benigno, G.B., oàn, J., Mináč, J., Sejnowski, T.J., Muller, L.E., 2023. Analytical prediction of specific spatiotemporal patterns in nonlinear oscillator networks with distance-dependent time delays. Physical Review Research 5, 013159.

Callaway, E.M., 2004. Feedforward, feedback and inhibitory connections in primate visual cortex. Neural Networks 17, 625–632.

Campbell, J.M., Davis, T.S., Anderson, D.N., Arain, A., Davis, Z., Inman, C.S., Smith, E.H., Rolston, J.D., 2024. Macroscale traveling waves evoked by single-pulse stimulation of the human brain. bioRxiv , 2023–03.

Cekic, S., Grandjean, D., Renaud, O., 2018. Time, frequency, and time-varying granger-causality measures in neuroscience. Statistics in medicine 37, 1910–1931.

Chemla, S., Reynaud, A., Di Volo, M., Zerlaut, Y., Perrinet, L., Destexhe, A., Chavane, F., 2019a. Suppressive traveling waves shape representations of illusory motion in primary visual cortex of awake primate. Journal of Neuroscience 39, 4282–4298.

Chemla, S., Roux, S., Reynaud, A., Chavane, F., VanRullen, R., 2019b. Revealing α oscillatory activity using voltage-sensitive dye imaging in monkey v1. bioRxiv , 810325.

Das, A., Zabeh, E., Jacobs, J., 2023. How to detect and analyze traveling waves in human intracranial eeg oscillations?, in: Intracranial EEG: A Guide for Cognitive Neuroscientists. Springer, pp. 487–505.

Davis, Z.W., Muller, L., Martinez-Trujillo, J., Sejnowski, T., Reynolds, J.H., 2020. Spontaneous travelling cortical waves gate perception in behaving primates. Nature 587, 432–436.

Ding, M., Chen, Y., Bressler, S.L., 2006. Granger causality: basic theory and application to neuroscience. Handbook of time series analysis: recent theoretical developments and applications , 437–460.

Dowdall, J.R., Schneider, M., Vinck, M., 2023. Attentional modulation of inter-areal coherence explained by frequency shifts. NeuroImage 277, 120256.

Dowdall, J.R., Vinck, M., 2023. Coherence fails to reliably capture inter-areal interactions in bidirectional neural systems with transmission delays. NeuroImage 271, 119998.

Ermentrout, B., 1998. Neural networks as spatio-temporal pattern-forming systems. Reports on progress in physics 61, 353.

Ermentrout, G.B., Kleinfeld, D., 2001. Traveling electrical waves in cortex: insights from phase dynamics and speculation on a computational role. Neuron 29, 33–44.

Freeman, W.J., Skarda, C.A., 1985. Spatial eeg patterns, non-linear dynamics and perception: the neo-sherringtonian view. Brain Research Reviews 10, 147–175.

Friston, K.J., 2011. Functional and effective connectivity: a review. Brain connectivity 1, 13–36.

Galas, L., Donovan, I., Dugué, L., 2025. Attention rhythmically shapes sensory tuning. Journal of Neuroscience .

Ghosh, S., Khare, K., Michailidis, G., 2019. High-dimensional posterior consistency in bayesian vector autoregressive models. Journal of the American Statistical Association .

Gorrostieta, C., Fiecas, M., Ombao, H., Burke, E., Cramer, S., 2013. Hierarchical vector auto-regressive models and their applications to multi-subject effective connectivity. Frontiers in computational neuroscience 7, 159.

Hatsopoulos, N.G., Ojakangas, C.L., Paninski, L., Donoghue, J.P., 1998. Information about movement direction obtained from synchronous activity of motor cortical neurons. Proceedings of the National Academy of Sciences 95, 15706–15711.

He, B., Astolfi, L., Valdés-Sosa, P.A., Marinazzo, D., Palva, S.O., Bénar, C.G., Michel, C.M., Koenig, T., 2019. Electrophysiological brain connectivity: theory and implementation. IEEE transactions on biomedical engineering 66, 2115–2137.

Isaacson, J.S., Scanziani, M., 2011. How inhibition shapes cortical activity. Neuron 72, 231–243.

Kulkarni, A., Ranft, J., Hakim, V., 2020. Synchronization, stochasticity, and phase waves in neuronal networks with spatially-structured connectivity. Frontiers in computational neuroscience 14, 569644.

Lubenov, E.V., Siapas, A.G., 2009. Hippocampal theta oscillations are travelling waves. Nature 459, 534–539.

Luo, C., VanRullen, R., Alamia, A., 2021. Conscious perception and perceptual echoes: a binocular rivalry study. Neuroscience of Consciousness 2021, niab007.

Marinazzo, D., Liao, W., Chen, H., Stramaglia, S., 2011. Nonlinear connectivity by granger causality. Neuroimage 58, 330–338.

Marinazzo, D., Pellicoro, M., Stramaglia, S., 2008. Kernel-granger causality and the analysis of dynamical networks. Physical Review E—Statistical, Nonlinear, and Soft Matter Physics 77, 056215.

Muller, L., Chavane, F., Reynolds, J., Sejnowski, T.J., 2018. Cortical travelling waves: mechanisms and computational principles. Nature Reviews Neuroscience 19, 255–268.

Muller, L., Reynaud, A., Chavane, F., Destexhe, A., 2014. The stimulus-evoked population response in visual cortex of awake monkey is a propagating wave. Nature communications 5, 3675.

Murphy, P.C., Duckett, S.G., Sillito, A.M., 1999. Feedback connections to the lateral geniculate nucleus and cortical response properties. Science 286, 1552–1554.

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

Ozeki, H., Finn, I.M., Schaffer, E.S., Miller, K.D., Ferster, D., 2009. Inhibitory stabilization of the cortical network underlies visual surround suppression. Neuron 62, 578–592.

Patel, J., Fujisawa, S., Berényi, A., Royer, S., Buzsáki, G., 2012. Traveling theta waves along the entire septotemporal axis of the hippocampus. Neuron 75, 410–417.

Pesaran, B., Vinck, M., Einevoll, G.T., Sirota, A., Fries, P., Siegel, M., Truccolo, W., Schroeder, C.E., Srinivasan, R., 2018. Investigating large-scale brain dynamics using field potential recordings: analysis and interpretation. Nature neuroscience 21, 903–919.

Pourahmadi, M., Noorbaloochi, S., 2016. Multivariate time series analysis of neuroscience data: some challenges and opportunities. Current Opinion in Neurobiology 37, 12–15.

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

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

Richardson, K.A., Schiff, S.J., Gluckman, B.J., 2005. Control of traveling waves in the mammalian cortex. Physical review letters 94, 028103.

Rieubland, S., Roth, A., Häusser, M., 2014. Structured connectivity in cerebellar inhibitory networks. Neuron 81, 913–929.

Rubino, D., Robbins, K.A., Hatsopoulos, N.G., 2006. Propagating waves mediate information transfer in the motor cortex. Nature neuroscience 9, 1549–1557.

Schneider, M., Broggini, A.C., Dann, B., Tzanou, A., Uran, C., Sheshadri, S., Scherberger, H., Vinck, M., 2021. A mechanism for inter-areal coherence through communication based on connectivity and oscillatory power. Neuron 109, 4050–4067.

Schneider, M., Tzanou, A., Uran, C., Vinck, M., 2023. Cell-type-specific propagation of visual flicker. Cell Reports 42.

Schwenk, J.C., Alamia, A., 2024. A hierarchical multiscale model of forward and backward alpha-band traveling waves in the visual system. bioRxiv , 2024–11.

Seth, A.K., Barrett, A.B., Barnett, L., 2015. Granger causality analysis in neuroscience and neuroimaging. Journal of Neuroscience 35, 3293–3297.

Shao, Z., Burkhalter, A., 1996. Different balance of excitation and inhibition in forward and feedback circuits of rat visual cortex. Journal of Neuroscience 16, 7353–7365.

Shojaie, A., Fox, E.B., 2022. Granger causality: A review and recent advances. Annual Review of Statistics and Its Application 9, 289–319.

Spyropoulos, G., Saponati, M., Dowdall, J.R., Schölvinck, M.L., Bosman, C.A., Lima, B., Peter, A., Onorato, I., Klon-Lipok, J., Roese, R., Neuenschwander, S., Fries, P., Vinck, M., 2022. Spontaneous variability in gamma dynamics described by a damped harmonic oscillator driven by noise. Nature communications 13, 1–18.

Spyropoulos, G., Schneider, M., van Kempen, J., Gieselmann, M.A., Thiele, A., Vinck, M., 2024. Distinct feedforward and feedback pathways for cell-type specific attention effects. Neuron .

Stephan, K., Friston, K., Squire, L., 2009. Functional connectivity .

Stokes, P.A., Purdon, P.L., 2017. A study of problems encountered in granger causality analysis from a neuroscience perspective. Proceedings of the national academy of sciences 114, E7063–E7072.

Troy, W.C., Shusterman, V., 2007. Patterns and features of families of traveling waves in large-scale neuronal networks. SIAM Journal on Applied Dynamical Systems 6, 263–292.

Tsodyks, M.V., Skaggs, W.E., Sejnowski, T.J., McNaughton, B.L., 1997. Paradoxical effects of external modulation of inhibitory interneurons. Journal of neuroscience 17, 4382–4388.

Vezoli, J., Magrou, L., Wang, X.J., Knoblauch, K., Vinck, M., Kennedy, H., 2020. Cortical hierarchy and the dual counterstream architecture. bioRxiv , 2020–04.

Vinck, M., Huurdeman, L., Bosman, C.A., Fries, P., Battaglia, F.P., Pennartz, C.M., Tiesinga, P.H., 2015. How to detect the granger-causal flow direction in the presence of additive noise? Neuroimage 108, 301–318.

Vinck, M., Uran, C., Spyropoulos, G., Onorato, I., Broggini, A.C., Schneider, M., Canales-Johnson, A., 2023. Principles of large-scale neural interactions. Neuron 111, 987–1002.

Wilmes, K.A., Clopath, C., 2019. Inhibitory microcircuits for top-down plasticity of sensory representations. Nature communications 10, 5055.

Ye, Z., Bull, M.S., Li, A., Birman, D., Daigle, T.L., Tasic, B., Zeng, H., Steinmetz, N.A., 2023. Brain-wide topographic coordination of traveling spiral waves. bioRxiv , 2023–12.

Zanos, T.P., Mineault, P.J., Nasiotis, K.T., Guitton, D., Pack, C.C., 2015. A sensorimotor role for traveling waves in primate visual cortex. Neuron 85, 615–627.

Zeng, Y., Sauseng, P., Alamia, A., 2024. Alpha traveling waves during working memory: Disentangling bottom-up gating and top-down gain control. Journal of Neuroscience 44.

Zhang, H., Watrous, A.J., Patel, A., Jacobs, J., 2018. Theta and alpha oscillations are traveling waves in the human neocortex. Neuron 98, 1269–1281.