PHYSICAL REVIEW E **94**, 040301(R) (2016)

Pairwise network information and nonlinear correlations

Elliot A. Martin,[1] Jaroslav Hlinka,[2,3] and Jörn Davidsen[1,*]

[1]*Complexity Science Group, Department of Physics and Astronomy, University of Calgary, Calgary, Alberta, Canada T2N 1N4*

[2]*Institute of Computer Science, The Czech Academy of Sciences, Pod vodarenskou vezi 2, 18207 Prague, Czech Republic*

[3]*National Institute of Mental Health, Topolová 748, 250 67 Klecany, Czech Republic*

(Received 19 August 2015; revised manuscript received 9 September 2016; published 27 October 2016)

Reconstructing the structural connectivity between interacting units from observed activity is a challenge across many different disciplines. The fundamental first step is to establish whether or to what extent the interactions between the units can be considered pairwise and, thus, can be modeled as an interaction network with simple links corresponding to pairwise interactions. In principle, this can be determined by comparing the maximum entropy given the bivariate probability distributions to the true joint entropy. In many practical cases, this is not an option since the bivariate distributions needed may not be reliably estimated or the optimization is too computationally expensive. Here we present an approach that allows one to use mutual informations as a proxy for the bivariate probability distributions. This has the advantage of being less computationally expensive and easier to estimate. We achieve this by introducing a novel entropy maximization scheme that is based on conditioning on entropies and mutual informations. This renders our approach typically superior to other methods based on linear approximations. The advantages of the proposed method are documented using oscillator networks and a resting-state human brain network as generic relevant examples.

DOI: 10.1103/PhysRevE.94.040301

Introdution. Pairwise measures of dependence such as cross-correlations (as measured by the Pearson correlation coefficient or covariance matrix) and mutual information are widely used to characterize the interactions within complex systems. They are a key ingredient to techniques such as principal component analysis, empirical orthogonal functions, and functional networks (networks inferred from dynamical time series) [1–3]. These techniques are widespread since they provide greatly simplified descriptions of complex systems and allow for the analysis of what might otherwise be intractable problems [4]. In particular, functional networks have been widely applied in fields such as neuroscience [4,5], genetics [6], and cell physiology [7], as well as in climate research [1,8].

In this paper we study how faithfully these measures alone can represent a given system. With the increasing use of functional networks, this topic has received much attention recently, and many technical concerns have been brought to light dealing with the inference of these networks. Previous studies have shown that the estimates of the functional networks can be negatively affected by properties of the time series [9–11], as well as properties of the measure of association, e.g., cross-correlations [12–15]. In this work, however, we address a more fundamental question: How well do pairwise measurements represent a system?

In principle, this can be evaluated using a maximum entropy approach. The corresponding framework was first laid out in Ref. [16] and later applied in Ref. [17], where they assessed the rationale of only looking at the pairwise correlations between neurons. They examined how well the maximum entropy distribution, consistent with all the pairwise correlations described the system. If the system is not well described by this maximum entropy distribution, then we know from the work of Jaynes [18] that other information beyond pairwise relationships would need to be taken into account. Similar analyses have since been applied in neuroscience [19–21], as well as in genetics [22], linguistics [23], and economics [24] and to the Supreme Court of the United States [25].

However, the data to accurately estimate the needed bivariate probability distributions may not be available. To get around this, some researchers have used the first two moments of the variables as constraints instead of the full bivariate distributions [26,27]—effectively using the cross-correlations as their constraints. In the case of binary variables, as in the original work [17], this is equivalent to conditioning on the bivariate distributions. For larger cardinality variables this is only an approximation, though, as the cross-correlation is only sensitive to linear relationships [28]. Systems where larger cardinalities and nonlinear behavior are thought to play a significant role such as in coupled oscillators—which have been used to model systems as diverse as pacemaker cells and crickets [29]—are, however, rather the norm than an exception [28]. In particular, we show here that this plays a significant role in a resting-state human brain network.

In order to retain the attractive properties of the cross-correlation and simultaneously capture a much wider range of relationships we propose using the mutual information. Mutual information can detect arbitrary pairwise relationships between variables and is only nonzero when the variables are pairwise independent, making it the ideal measure [30]. However, while calculating the maximum entropy given the moments of a distribution result in simple equations in the probabilities, using mutual informations as constraints results in transcendental equations which are much harder to solve. We circumvent this problem here using the set-theoretic formulation of information theory [31], which gives us an upper bound on the maximum entropy that is saturated in many cases.

The set-theoretic formulation of information theory allows us to map information-theoretic quantities to the regions of

*davidsen@phas.ucalgary.ca

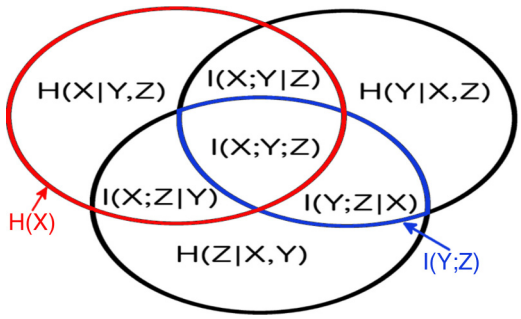

FIG. 1. The information diagram for three variables. It contains seven regions corresponding to the possible combinations of three variables, with their corresponding information-theoretic quantities defined in the text. The univariate entropy $H(X)$ is the sum of all the regions in the red circle, and the mutual information $I(Y;Z)$ is the sum of all the regions in the blue oval.

an information diagram, a variation of a Venn diagram. The information diagram for three variables is shown in Fig. 1 with the associated information-theoretic quantities labeled [32]: entropy, $H(X) = -\sum p(x)\log[p(x)]$; conditional entropy, $H(X|Y,Z) = -\sum p(x,y,z)\log(p(x|y,z))$; mutual information, $I(X;Y) = \sum p(x,y)\log(p(x,y)/(p(x)p(y)))$; conditional mutual information, $I(X;Y|Z) = \sum p(x,y,z)\log(p(x;y|z)/[p(x|z)p(y|z)])$; and multivariate mutual information, $I(X;Y;Z) = I(X;Y) - I(X;Y|Z)$.

We illustrate our method using systems of coupled oscillators, as they commonly occur in nature and are used to model a large variety of systems [29]. In particular, we look at the Kuramoto model [33,34] as a paradigmatic example that is capable of a wide range of dynamics from synchronization to chaos [35] and, hence, provides an excellent test bed for our method.

Method. Given a set of N variables ($\{X\}_N$), we want to know how well the cross-correlation or mutual information between all pairs of variables can encode the state of the system. To do this, we first determine the maximum entropy consistent with the given measure of similarity, $H_m(\{X\}_N)$, which represents a standard variational problem. This means that any model of the system consistent with the $\binom{N}{2}$ values of the similarity measure can have an entropy of at most $H_m(\{X\}_N)$. From the work of Jaynes [18], we also know that any model of the system with a smaller entropy must implicitly or explicitly include information beyond these values. As a result, the true joint entropy, $H(\{X\}_N)$, will always be less than or equal to $H_m(\{X\}_N)$.

If the variables are all independent, then the entropy of the system is the sum of the entropies of the individual variables, $H_I(\{X\}_N) = \sum_i H(X_i)$. The most this uncertainty can be reduced is if the true joint entropy is known and is given by the multi-information (also called total correlation [36]) $I_N(\{X\}_N) = H_I(\{X\}_N) - H(\{X\}_N) \geqslant 0$. We similarly define the measure information $I^m(\{X\}_N) = H_I(\{X\}_N) - H_m(\{X\}_N)$ to be the reduction in uncertainty given a measure. The fraction of information retained by describing the system with a given measure, as opposed to the true joint entropy, is then $0 \leqslant I^m/I_N \leqslant 1$. If the used measure is the bivariate probability distribution, we call I^m the pairwise network information or the second-order connected information as defined in Ref. [16]. This is approximated linearly if the measure used is the cross-correlation and nonlinearly if the measure used is the mutual information.

When using the cross-correlation, estimating H_m is conceptually straightforward, though finding the optimum value can be computationally expensive. Estimates of the first two moments of the variables uniquely determine the cross-correlations and can be used as constraints in a Lagrange multiplier problem solving for H_m. The resulting probability distribution $P_m(\{X\}_N)$ is the Boltzmann distribution $P_m(\{X\}_N) = \exp(\sum_i h_i x_i + \sum_{i\geqslant j} J_{i,j}x_i x_j)$, where h_i and $J_{i,j}$ are the Lagrange multipliers [3].

When using mutual information, estimating H_m with Lagrange multipliers is much harder as the derivatives of the Lagrange function are transcendental functions in $P_m(\{X\}_N)$. Instead, we use the mutual informations and univariate entropies as constraints and draw on the structure of information diagrams. Each univariate entropy and mutual information corresponds to a region in the information diagram that can be written as a sum of a number of *atomic regions (atoms)*. The sum over all atoms is simply $H(\{X\}_N)$. Thus, as seen in Fig. 1, we obtain constraints of the form:

$$\text{const} = I(Y;Z) = I(Y;Z|X) + I(X;Y;Z), \tag{1}$$

$$\text{const} = H(X) = H(X|Y,Z) + I(X;Y|Z) + I(X;Z|Y) + I(X;Y;Z). \tag{2}$$

In general, a system of N variables results in $\binom{N}{1}$ univariate entropy constraints, $\binom{N}{2}$ mutual information constraints, and $A = \sum_{k=1}^{N}\binom{N}{k} = 2^N - 1$ atoms to be determined. In the simplest case of $N = 3$ variables we have six constraints and $A = 7$ regions to specify, see Fig. 1. This means we only have one free parameter, making the maximization process to get $H_m(\{X\}_N)$ particularly easy in this case; in general, there are $\sum_{k=3}^{N}\binom{N}{k}$ free parameters.

Apart from the chosen constraints defined above, there are also general constraints on the values of the subregions ensuring they define a valid information diagram, i.e., that there exists a probability distribution with corresponding information-theoretic quantities. A family of such constraints (so-called Shannon inequalities) can be inferred from the fundamental requirement that, for discrete variables, (conditional) entropies and mutual informations are necessarily non-negative: (A) $H(X_i|\{X\}_N - X_i) \geqslant 0$ and (B) $I(X_i;X_j|\{X\}_K) \geqslant 0$, where $i \neq j$ and $\{X\}_K \subseteq \{X\}_N - \{X_i, X_j\}$ [37]. Each inequality can also be written as a sum of atoms, e.g.,

$$I(X_1;X_2|X_3) = I(X_1;X_2|X_3,X_4) + I(X_1;X_2;X_4|X_3) \geqslant 0. \tag{3}$$

Not so well known, for $N \geqslant 4$, is that there are also inequalities that are not deducible from the Shannon inequalities, the so-called non-Shannon inequalities [31]. In principle, these inequalities may be included in our maximization problem; however, they have not yet been fully described. Therefore, we suggest constructing the diagram with the maximum entropy that satisfies the problem-specific constraints and is

consistent with the Shannon inequalities. As it may violate the non-Shannon inequalities, it may not represent a valid distribution. However, the sum of the atomic regions would still be an upper bound on the entropy H_m and thus provide a *lower* bound on I^m/I_N. Notably, for the particular (and, in our simulations, common) case where all A elements are non-negative—which is always true for $N = 3$—one can prove that the bound is attainable (see Theorem 3.11 [31]).

To summarize, the task of finding the maximum entropy conditioned on the univariate entropies, mutual information, and elemental Shannon inequalities can be solved using linear optimization: Each constraint will take the form of a linear (in-)equality, as in Eqs. (1) and (3), and we maximize the N-variate entropy by maximizing the sum over all A atoms of the information diagram. Thus, we avoid having to perform the maximization over probability distributions.

a. Example of a nonlinear pairwise distribution. We now give an example illustrating how the mutual information can better detect pairwise relationships than the cross-correlation. Consider a set of variables $\{X\}_N$: Each variable is drawn uniformly from the set $\{-1,0,1\}$, and all variables are simultaneously 0 or independently distributed among $\{-1,1\}$. The cross-correlation between any pair of variables is zero and therefore consistent with the hypothesis that all variables are independent. Therefore, the fraction of information captured by the cross-correlation is $I^m/I_N = 0$. However, there is a significant amount of mutual information between the variables.

Since $P(X_i|\{X\}_N - X_i) = P(X_i|X_{j\neq i})\forall i$ and j, all the conditional mutual informations are zero. Therefore, the only nonzero atoms in the information diagram will be the N-variate mutual information $I(X_1;\ldots;X_N) = I(X_1;X_2)$, and the conditional entropies $H(X_i|\{X\}_N - X_i) = 2/3$ bits. This is the maximum entropy diagram consistent with the pairwise mutual information and univariate entropies, so the expected result using the mutual information is $I^m/I_N = 1$. We can see why this is the case by starting with the information diagram for two variables (which is fixed from our conditions) and successively adding new variables. The addition of each new variable adds 2/3 bits to the total entropy—which is the maximal amount consistent with the mutual information.

b. Kuramoto model. The Kuramoto model is a dynamical system of N phase oscillators with all-to-all coupling proportional to K [33,34]. The ith oscillator has an intrinsic frequency of ω_i and a phase of θ_i and its dynamics is given by $\frac{\partial\theta_i}{\partial t} = \omega_i + \frac{K}{N}\sum_{j=1}^{N}\sin(\theta_j - \theta_i) + \eta_i(t)$. Here, we have followed Ref. [38] and added a dynamical noise term to mimic natural fluctuations and environmental effects; $\eta_i(t)$ is drawn from a Gaussian distribution with correlation function $\langle\eta_i(t)\eta_j(t')\rangle = G\delta_{i,j}\delta(t-t')$, where G determines the amplitude of the noise. For values of K above a critical threshold, $K > K_c$, synchronization occurs [29]. In the limit of constant phase differences the dynamics are trivial, and knowledge of one oscillator will specify the phase of all others. Therefore, pairwise information is sufficient to describe the system in this case. Yet, the presence of noise results in random perturbations of the phases and typically prevents constant phase differences [38] such that only $I^m/I_N \lesssim 1$ is expected. In the weak-coupling regime when synchronization is absent, it is nontrivial what I^m/I_N should be.

To estimate I^m/I_N and to establish the importance of the level of discretization or cardinality, we first discretize the phase of each oscillator into n equally likely states [39]. Alternatively, estimators for continuous variables can be used as we discuss in Ref. [40]. To provide clear proofs of principle, we first focus on three-oscillator systems in the following as this is the smallest system size at which the results are nontrivial. Specifically, we consider three different cases: (i) all oscillators have the same intrinsic frequency, (ii) all oscillators have unique intrinsic frequencies and are still synchronized, and (iii) all oscillators have unique intrinsic frequencies and the entire system and all subsystems are unsynchronized ("weak-coupling regime"). For three-oscillator systems, the corresponding parameter regimes in the absence of noise have been carefully documented in Ref. [35].

For each of the three cases examined, we created ensembles of 100 three-oscillator systems, where each element of the ensemble will have randomly sampled frequencies [41]. These ensembles are studied in two different noise regimes, $G = 0.001$ and $G = 0.5$. The same ensemble of frequencies is used in both noise regimes.

In the first case, all oscillators are synchronized with $\omega_1 = \omega_2 = \omega_3$ and $K = 1.65$. Recall that in the synchronized case we expect $E[I^m/I_N] \approx 1$. This is indeed what we see in the low-noise case, $G = 0.001$, Fig. 2(A); though the mutual information preserves slightly more information at larger cardinalities. However, for increased noise, $G = 0.5$, the cross-correlation performs poorly at larger cardinalities, while the mutual information behaves robustly, Fig. 2(D).

For the second case, where the oscillators are synchronized with different intrinsic frequencies, we use $\Delta_1/\Delta_2 = 1.11$, $K/\Delta_2 = 4$, and $K = 2.20$, where $\Delta_1 = \omega_2 - \omega_1$ and $\Delta_2 = \omega_3 - \omega_2$. Now at both noise levels, at cardinalities greater than 2, the cross-correlation fails to capture a significant portion of the available information—as $\hat{E}[I^m/I_N]$ is significantly less than 1 [Figs. 2(B) and 2(E)]. This indicates that even small-amplitude noise can prevent the cross-correlation from accurately encoding information about the system in this case. The mutual information again robustly encodes almost all of the possible information, $\hat{E}[I^m/I_N] \approx 1$, in both noise regimes and across all discretizations analyzed.

In the final case, the weak-coupling regime ($K/\Delta_2 = 0.99$, all other parameters as in the second case), we do not have a strong hypothesis for what $E[I^m/I_N]$ should be. In Figs. 2(C) and 2(F), we can see that the cross-correlation encodes virtually no information about the system for cardinalities greater than 2, $\hat{E}[I^m/I_3] \approx 0$. The mutual information again robustly encodes the vast majority of the multi-information, with $\hat{E}[I^m/I_3] > 0.8$ for all noise levels and discretizations.

Similar overall results hold for larger systems and when only a subset of oscillators is observed. As an example, we consider here a system of 100 nonidentical Kuramoto oscillators in two regimes: (i) All oscillators are synchronized, $K = 4$, and (ii) the oscillators are partially synchronized with more than 20 different synchronized clusters, $K = 1.75$. In both cases we use the same set of intrinsic frequencies (drawn from a normal distribution with mean zero and unit variance) and a noise level $G = 0.001$.

As in the analysis done in Ref. [17], we analyzed the effects of sampling from a larger system by randomly selecting

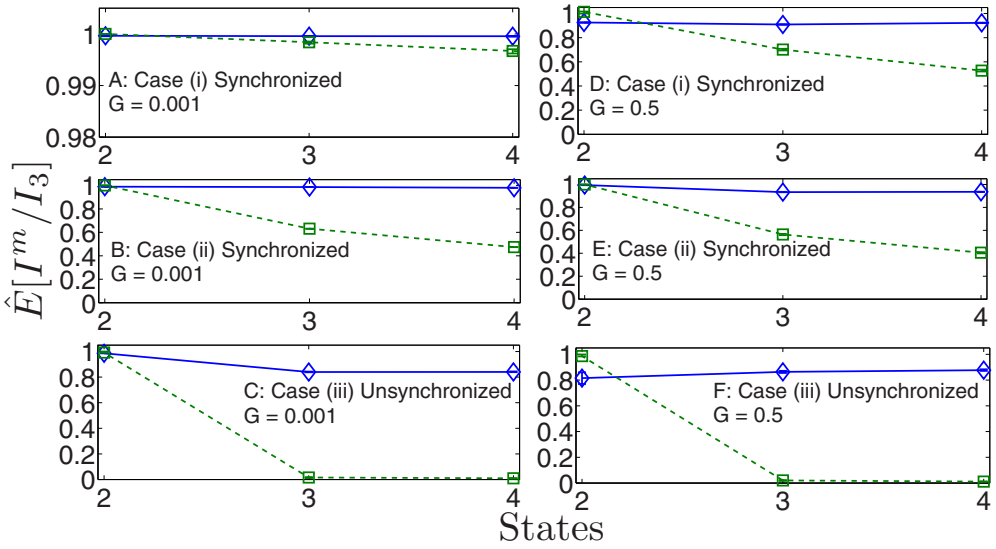

FIG. 2. The fraction of shared information coded by the mutual information (blue diamonds, solid line) and the cross-correlation (green squares, dashed line). Notice the scale only goes from 0.98 to 1 in panel (A) and from 0 to 1 for the rest. The estimated expectations, $\hat{E}[\ldots]$, are averages over the ensemble of 100 realizations where we draw ω_3 from a normal distribution with zero mean and unit variance. Uncertainties corresponding to the 25% and 75% quantiles are smaller than the symbol sizes.

T of the 100 oscillators and calculating I^m/I_T for those oscillators. For each tuple size, T, we repeated this 100 times, using the same sets of tuples in both regimes, and computed $\hat{E}[I^m/I_T]$ as the average of these values. As in our previous examples, our method outperforms the cross-correlation in the synchronized case (see Fig. 3), as well as for weaker coupling (see Fig. 4). Our method results in $\hat{E}[I^m/I_T] \lesssim 1$ in both regimes, and across all discretizations and tuple sizes, while the cross-correlation only does so for binary variables.

c. Resting-state human brain networks. To illustrate the applicability of our methodology in real-world data situations, we apply it to neuroimaging data, in a similar context as in Ref. [19]. In particular, we want to assess to what extent the multivariate activity distribution is determined by purely bivariate dependence patterns. The used data consist of a time series of functional magnetic resonance imaging signals from 96 healthy volunteers measured using a 3T Siemens Magnetom Trio scanner at the Institute for Clinical and Experimental Medicine in Prague, Czech Republic. Average signals from 12 regions of the frontoparietal network were extracted using a brain atlas [42]. After preprocessing and denoising as in Ref. [43,44], the data were temporally concatenated. Each variable was further discretized to two or three states using equiquantal binning. Using our approach, we find $I^m/I_N = 0.88$ for the 2-state and $I^m/I_N = 0.77$ for the 3-state discretizations, suggesting that bivariate dependence patterns capture the dominant proportion of the information. For 2-state discretization, this is smaller than in Ref. [19]. However, for

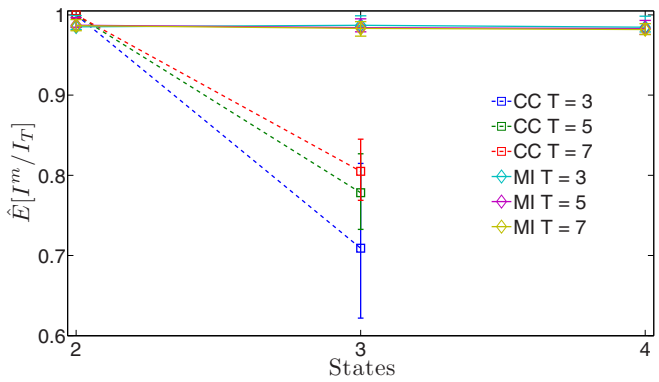

FIG. 3. The fraction of shared information coded by the mutual information (MI, diamonds with solid lines) and the cross-correlation (CC, squares with dashed lines) for a tuple of size T. We simulated 100 nodes with $K = 4, G = 0.001$, and the estimated expectations, $\hat{E}[\ldots]$, are averages over 100 randomly selected tuples of the given size. All oscillators are synchronized, and their intrinsic frequencies are drawn from a normal distribution with zero mean and unit variance. Error bars are 25% and 75% quantiles.

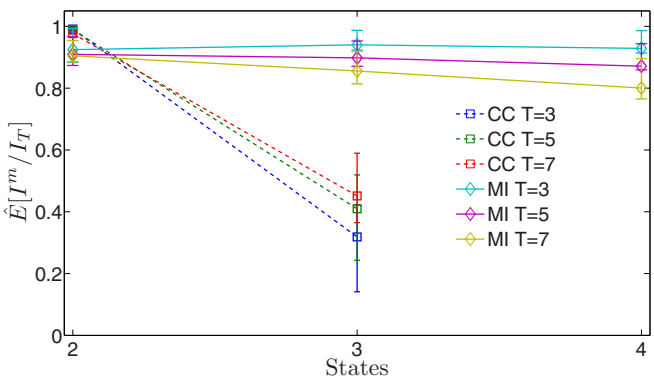

FIG. 4. The fraction of shared information coded by the mutual information (MI, diamonds with solid lines) and the cross-correlation (CC, squares with dashed lines) for a tuple of size T as in Fig. 3 but for weak coupling $K = 1.75$ leading to partial synchronization with more than 20 different clusters.

the 3-state discretization, it provides a much higher estimate of the bivariate dependence role than the method taking into account only correlations, as in the case of the Kuramoto model. This suggests that only when accounting also for nonlinear coupling does the bivariate dependencies provide sufficient data structure approximation to resolve the apparent inconsistency of the results in Ref. [19]. This is also true for other brain networks [40].

Discussion. Our method allows for potential speedups over the maximum entropy calculation when conditioning on the bivariate distributions, as well as when conditioning on the cross-correlations. In both of these cases, solving the associated Lagrange multiplier equations are nonlinear optimization problems. The maximum entropy distribution could also be found using iterative fitting routines like in Ref. [45], but in these cases the problem will still scale like n^N (n is the cardinality of the variables). While there are pathological linear optimization problems that scale exponentially with N, there will always be a slightly perturbed problem such that our method will scale polynomially [46].

Researchers have so far relied on conditioning on the cross-correlations when insufficient data are available to estimate the bivariate distributions. They either coarse grain to binary variables where it is equivalent to conditioning on the distributions [19]—potentially losing important information—or use higher cardinality variables where it is only a linear approximation [26,27]. Our approach based on mutual information can be applied in these cases; the associated entropies can be estimated with as few as $2^{H/2}$ data points [47] (H is measured in bits). While this maximization has previously been prohibitively difficult, our work shows that it is feasible, allowing it to become widely applicable and serve as a starting point before considering multivariate information measures [48–50]. Additionally, if our method returns a small I^m/I_N, then this suggests both that the faithfulness assumption used in causal inference is violated [51–54] and that there is synergy among the variables [55].

Our calculation of H_m for the mutual information is free of distributional assumptions, computing the maximum entropy in the general space of arbitrary cardinality variables. This may result in higher entropy estimates than methods that consider predefined cardinality, e.g., binary variables. Notably, our simulations suggest that estimating H_m in this way provides comparable, or substantially lower, entropy estimates than H_m for the cross-correlation, which explicitly constrains the cardinality. This makes the technique competitive even when a specific cardinality could be reasonably assumed.

Conclusions. In this work we introduced a novel method to determine the importance of pairwise relationships by estimating the maximum entropy conditioned on the mutual information. We showed that by mapping this problem to a linear optimization problem it could also be efficiently computed. Using the generic case of coupled oscillators, we gave a proof-of-principle example where our method was able to widely out-perform conditioning on the cross-correlations. The example of the resting-state brain network showed that this also carries over to real-world applications, highlighting the potential of the method when cardinalities larger than two and nonlinear behavior are important.

Our results indicate that in many relevant cases functional networks based on mutual information can in principle more accurately capture the dynamics of the system than those functional networks based on cross-correlations. These types of analyses should be applied *before* studying functional networks, both to assess the validity of the network paradigm as well as to test the appropriateness of using the given measure of association. Only high values in the fraction of shared information ensure that this is the case. This has not been done in the vast majority of applications in the past. Due to the computational efficiency, our proposed methodology should allow us to revisit this question, especially in areas where functional networks have already been widely applied, such as in climate research [1,8,9,56].

Acknowledgments. This project was financially supported by NSERC (E.M. and J.D.) and by the Czech Science Foundation Project No. 13-23940S and the Czech Health Research Council Project No. NV15-29835A and by the project No. LO1611 with a financial support from the MEYS under the NPU I program (J.H.). All authors thank the MPIPKS for its hospitality and hosting the international seminar program "Causality, Information Transfer and Dynamical Networks", which stimulated some of the involved research. We also thank P. Grassberger for many helpful discussions.

[1] J. F. Donges, I. Petrova, A. Loew, N. Marwan, and J. Kurths, Clim. Dynam. **45**, 2407 (2015).
[2] A. Haimovici, E. Tagliazucchi, P. Balenzuela, and D. R. Chialvo, Phys. Rev. Lett. **110**, 178101 (2013).
[3] M. Timme and J. Casadiego, J. Phys. A **47**, 343001 (2014).
[4] E. Bullmore and O. Sporns, Nat. Rev. Neurosci. **10**, 186 (2009).
[5] V. M. Eguiluz, D. R. Chialvo, G. A. Cecchi, M. Baliki, and A. V. Apkarian, Phys. Rev. Lett. **94**, 018102 (2005).
[6] A. A. Margolin, I. Nemenman, K. Basso, C. Wiggins, G. Stolovitzky, R. D. Favera, and A. Califano, BMC Bioinform. **7**, S7 (2006).
[7] A. Stožer, M. Gosak, J. Dolenšek, M. Perc, M. Marhl, M. S. Rupnik, and D. Korošak, PLoS Comput. Biol. **9**, e1002923 (2013).
[8] J. Runge, V. Petoukhov, J. F. Donges, J. Hlinka, N. Jajcay, M. Vejmelka, D. Hartman, N. Marwan, M. Paluš, and J. Kurths, Nat. Commun. **6**, 8502EP (2015).
[9] E. A. Martin, M. Paczuski, and J. Davidsen, Europhys. Lett. **102**, 48003 (2013).
[10] M. Paluš, D. Hartman, J. Hlinka, and M. Vejmelka, Nonlinear Proc. Geoph. **18**, 751 (2011).
[11] S. Bialonski, M. Wendler, and K. Lehnertz, PLoS ONE **6**, e22826 (2011).
[12] G. Tirabassi, R. Sevilla-Escoboza, J. M. Buldú, and C. Masoller, Sci. Rep. **5**, 10829EP (2015).
[13] J. Hlinka, D. Hartman, and M. Paluš, Chaos **22**, 033107 (2012).
[14] E. A. Martin and J. Davidsen, Nonlin. Proc. Geoph. **21**, 929 (2014).

[15] W. Mader, M. Mader, J. Timmer, M. Thiel, and B. Schelter, Sci. Rep. **5**, 10805 (2015).
[16] E. Schneidman, S. Still, M. J. Berry, and W. Bialek, Phys. Rev. Lett. **91**, 238701 (2003).
[17] E. Schneidman, M. J. Berry, R. Segev, and W. Bialek, Nature **440**, 1007 (2006).
[18] E. T. Jaynes, Phys. Rev. **106**, 620 (1957).
[19] T. Watanabe, S. Hirose, H. Wada, Y. Imai, T. Machida, I. Shirouzu, S. Konishi, Y. Miyashita, and N. Masuda, Nat. Commun. **4**, 1370 (2013).
[20] S. Yu, H. Yang, H. Nakahara, G. S. Santos, D. Nikolić, and D. Plenz, J. Neurosci. **31**, 17514 (2011).
[21] I. E. Ohiorhenuan, F. Mechler, K. P. Purpura, A. M. Schmid, Q. Hu, and J. D. Victor, Nature (London) **466**, 617 (2010).
[22] T. R. Lezon, J. R. Banavar, M. Cieplak, A. Maritan, and N. V. Fedoroff, Proc. Natl. Acad. Sci. USA **103**, 19033 (2006).
[23] G. J. Stephens and W. Bialek, Phys. Rev. E **81**, 066119 (2010).
[24] N. Xi, R. Muneepeerakul, S. Azaele, and Y. Wang, Physica A **413**, 189 (2014).
[25] E. D. Lee, C. P. Broedersz, and W. Bialek, J. Stat. Phys. **160**, 275 (2015).
[26] W. Bialek, A. Cavagna, I. Giardina, T. Mora, E. Silvestri, M. Viale, and A. M. Walczak, Proc. Natl. Acad. Sci. USA **109**, 4786 (2012).
[27] K. Wood, S. Nishida, E. D. Sontag, and P. Cluzel, Proc. Natl. Acad. Sci. USA **109**, 12254 (2012).
[28] H. Kantz and T. Schreiber, *Nonlinear Time Series Analysis* (Cambridge University Press, Cambridge, UK, 1997).
[29] A. Pikovsky, M. Rosenblum, and J. Kurths, *Synchronization: A Universal Concept in Nonlinear Sciences* (Cambridge University Press, Cambridge, 2003), Vol. 12.
[30] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (Wiley, New York, 2006).
[31] R. W. Yeung, *Information Theory and Network Coding* (Springer, Berlin, 2008).
[32] We use the convention $p(x,y,z) = P(X = x, Y = y, Z = z)$.
[33] Y. Kuramoto, in *International Symposium on Mathematical Problems in Theoretical Physics*, edited by H. Araki, Lecture Notes in Physics Vol. 39 (Springer, Berlin, 1975), pp. 420–422.
[34] J. A. Acebrón, L. L. Bonilla, C. J. P. Vicente, F. Ritort, and R. Spigler, Rev. Mod. Phys. **77**, 137 (2005).
[35] Y. Maistrenko, O. Popovych, O. Burylko, and P. A. Tass, Phys. Rev. Lett. **93**, 084102 (2004).
[36] S. Watanabe, IBM J. Res. Dev. **4**, 66 (1960).
[37] This set of equalities is minimal in the sense that no inequality is implied by any combination of the others [31].
[38] H. Sakaguchi, Prog. Theor. Phys. **79**, 39 (1988).
[39] Using n equally sized bins did not noticeably alter our results.
[40] E. A. Martin, J. Hlinka, A. Meinke, F. Děchtěrenko, and J. Davidsen, arXiv:1601.00336 (2016).
[41] For the numerical simulations, we use the Euler-Maruyama method, with a time step $dt = 2^{-6}$, and unless otherwise stated we use an integration time of $T = 2000$ in all of our results. We also discard times up to 50 to remove transient effects, and resample the data only taking every eighth data point.
[42] W. R. Shirer, S. Ryali, E. Rykhlevskaia, V. Menon, and M. D. Greicius, Cereb. Cortex (2011).
[43] J. Hlinka, M. Paluš, M. Vejmelka, D. Mantini, and M. Corbetta, NeuroImage **54**, 2218 (2011).
[44] D. Hartman, J. Hlinka, M. Paluš, D. Mantini, and M. Corbetta, Chaos **21**, 013119 (2011).
[45] J. N. Darroch and D. Ratcliff, Ann. Math. Stat. **43**, 1470 (1972).
[46] R. Vershynin, SIAM J. Comput. **39**, 646 (2009).
[47] I. Nemenman, Entropy **13**, 2013 (2011).
[48] N. Timme, W. Alford, B. Flecker, and J. M. Beggs, J. Comput. Neurosci. **36**, 119 (2014).
[49] B. Kralemann, A. Pikovsky, and M. Rosenblum, New J. Phys. **16**, 085013 (2014).
[50] D. Kugiumtzis, Phys. Rev. E **87**, 062918 (2013).
[51] M. Eichler, Philos. Trans. R. Soc. Lond. A **371**, 20110613 (2013).
[52] J. Sun, C. Cafaro, and E. M. Bollt, Entropy **16**, 3416 (2014).
[53] J. Runge, J. Heitzig, V. Petoukhov, and J. Kurths, Phys. Rev. Lett. **108**, 258701 (2012).
[54] K. Hlaváčková-Schindler, M. Paluš, M. Vejmelka, and J. Bhattacharya, Phys. Rep. **441**, 1 (2007).
[55] V. Griffith and C. Koch, in *Guided Self-Organization: Inception, Emergence, Complexity and Computation*, edited by M. Prokopenko (Springer, Berlin, 2014), Vol. 9, pp. 159–190.
[56] J. Deza, M. Barreiro, and C. Masoller, Eur. Phys. J-Spec. Top. **222**, 511 (2013).