

Inferring a Graph's Topology from Games Played on It

Douglas G. Moore and Sara I. Walker

BEYOND Center, Arizona State University, Tempe, AZ 85281
douglas.g.moore@asu.edu
sara.i.walker@asu.edu

Abstract

We consider an iterated model of agents playing a two-player game on a graph. The agents change their strategies as the game progresses based on anticipated payoffs. Using only the time series of the agents' strategies, we determine the pairwise mutual information between all agents in the graph, and use these values as predictors of the graph's topology. From this, we assess the influence of various model parameters on the effectiveness of mutual information at recovering the actual causal structure. It is found that the degree to which the functional connectivity reflects the actual causal structure of the graph strongly depends on which game is being played and how the agents are changing their strategies. Further, there is evidence that the edge density of the graph may also have some impact on the accuracy of the inferred network. This approach allows us to better connect the dynamics of the systems under study with the difference in their functional and actual connectivity, and has broad implications for the interpretation and application of information-based network inference. The methods and analyses described can be generalized and applied to other types of network models.

Introduction

There is a wealth of literature on using information-theoretic tools to infer the network topology of collections of interacting random variables, particularly in neuroscience (Ito et al., 2011; Vicente et al., 2011; Lizier and Rubinov, 2012; Ver Steeg and Galstyan, 2012; Banerji et al., 2013; Sun and Boltt, 2014; Sun et al., 2015). In many of these works, the phrase "functional connectivity" is used to underscore the fact that the actual causal interactions are not always readily inferred from the dynamics of the variables. This is often considered yet another example of the inequivalence of correlation and causation. While it appears that functional connectivity is a useful construct in neuroscience, that may not be the case in other fields. However, it is interesting to explore the conditions under which the actual and functional connectivity may coincide.

We consider a somewhat unorthodox model system to explore whether and when information-theoretic tools recapitulate the actual topology of a causal graph. Agents are situated on the vertices of a graph, and allowed to repeatedly

play a two-player, two-strategy, symmetric game. In each round of the game, the agents may change their strategies based on their anticipated payoff. Treating the strategies of the agents as random variables in time, when can we confidently infer the actual connectivity of the graph? This model system was inspired by conversations with colleagues who explored the dynamics of games when agents' perception of which game they're playing is allowed to change (Antonioni et al., 2018). Such a system is intrinsically interesting to consider, but it has value beyond that; namely in potential applications. Biologists have modeled interacting cells (Hummert et al., 2014), organisms (Smith John Maynard et al., 1979; Smith, 1982; Riechert and Hammerstein, 1983; Nowak and Sigmund, 2004) and even cancer (Bach et al., 2001; Gatenby and Vincent, 2003; Basanta et al., 2008; Dingli et al., 2009) as game-playing agents, and game theory has of course found innumerable applications in economics. In both cases, the strategies of the agents may be readily apparent, while the network itself is obscured. The ability to infer functional or actual connectivity based solely on the strategies could lead to more impactful methods of intervening on these systems.

From a technical perspective, this model offers a number of useful features: the graph's topology can be varied, the rules the agents use to update their strategies can be changed, and the space of graphs is nicely parameterized by two bounded variables. Altogether, this provides a number of ways to tweak the system and observe changes in the efficacy of network inference. Further, because this model is a particular type of Boolean network, it seems reasonable that the results and methods herein could be extrapolated and applied to more generic settings.

In this work, we describe the basic formalism, and present early results from simulations¹. We observe that the efficacy of the network reconstruction is dependent on the game the agents are playing, the rule they are using to update their strategies, and possibly the density of edges in the graph. We conclude with a number of possible future directions.

¹All code used for these simulations is publicly available at <https://github.com/dglmooore/gomen>.

Methods

Consider a finite collection of agents $\{A_1, \dots, A_n\}$. Each agent A_i is situated at a vertex v_i of some *simple*² graph $G = (V, E)$ with vertices $V = \{1, \dots, n\}$ and edges $E \subset V \times V$. The agents then play some, two-player, two-strategy, symmetric game with each of their neighbors. We restrict the games to those characterized by a payoff matrix of the form

$$\mathcal{G} = \begin{array}{c|cc} & 0 & 1 \\ \hline 0 & 1 & S \\ 1 & T & 0 \end{array} \quad (1)$$

with $-1 \leq S \leq 1$ and $0 \leq T \leq 2$. The elements of this matrix \mathcal{G}_{ab} specify the payoff received by the agent should they play strategy a against an opponent playing strategy b . It is assumed throughout that all players are knowingly playing the same game, and that the game does not change. The strategy played by A_i is denoted s_i and can be either 0 or 1, sometimes referred to as *cooperate* or *defect*, respectively. The collection of strategies played by all of the agents is denoted $\mathbf{s} = (s_1, \dots, s_n)$.

The game is played in rounds starting with each agent's strategy selected uniformly at random. In each round, the agents play with their neighbors $N_i = \{j \mid (i, j) \in E\}$ and accumulate payoff

$$P_i(s_1, \dots, s_n) = \sum_{j \in N_i} \mathcal{G}_{s_i s_j} \quad (2)$$

We write this as $P_i(\mathbf{s})$ for compactness. Since the agents know which game they are playing, they may also consider the *counterfactual* payoff they would have received had they played their alternative strategy \hat{s}_i :

$$\bar{P}_i(\mathbf{s}) = P_i(s_1, \dots, \hat{s}_i, \dots, s_n) = \sum_{j \in N_i} \mathcal{G}_{\hat{s}_i s_j}. \quad (3)$$

Based on the difference between counterfactual payoffs $\Delta P_i(\mathbf{s}) = \bar{P}_i(\mathbf{s}) - P_i(\mathbf{s})$, each agent may stochastically change its strategy for the next round with probability

$$q_i(\beta; \mathbf{s}) = \frac{1}{1 + \exp(-\beta \Delta P_i(\mathbf{s}))}, \quad \beta > 0. \quad (4)$$

In words, agent A_i will use strategy \hat{s}_i in the next round with probability $q_i(\beta; \mathbf{s})$. All agents then synchronously update their strategies, and a new round is ready to begin. This process is iterated t times, and a time series of strategies is generated $\{\mathbf{s}_0, \mathbf{s}_1, \dots, \mathbf{s}_t\}$.

The outstanding question is to what degree this time series can be used to infer the topology of the underlying graph. A number of information-theoretic tools have been proposed for assessing connectivity in similar situations. The most popular, particularly among neuroscientists, is transfer entropy or some variation thereof (Ito et al., 2011; Vicente

²Simple here means an undirected graph with no self-loops and at most one edge between any two given vertices.

et al., 2011; Orlandi et al., 2014). However, in this initial work we decided to consider a simpler measure. We start with

$$I(s_i^+, s_j) = \sum_{s_i^+, s_j} p(s_i^+, s_j) \log \frac{p(s_i^+, s_j)}{p(s_i^+)p(s_j)}, \quad (5)$$

the mutual information between agent A_i 's next strategy s_i^+ and agent A_j 's current strategy s_j . Since we know that edges of the graph are undirected, we construct a measure that is symmetric in the agents:

$$\phi_{ij} = \frac{1}{2}(I(s_i^+, s_j) + I(s_j^+, s_i)). \quad (6)$$

This value is computed for every pair of distinct agents, $i \neq j$ for $i, j \in V$, and used as a predictor for the existence of an edge in the graph. If $\phi_{ij} \geq \theta$ for some threshold $0 \leq \theta \leq 1$, we predict that the edge (i, j) exists, and that it otherwise does not. In a binary classification process such as this, one method of characterizing the effectiveness of the classification is in terms of the true-positive rate (TPR) and false-positive rate (FPR)

$$TPR = \frac{TP}{R} \quad \text{and} \quad FPR = \frac{FP}{F}. \quad (7)$$

Here the TPR is the ratio of the number of edges *correctly* predicted to exist (TP) and the actual number of existent edges (R), while the FPR is the ratio of the number of edges *incorrectly* predicted to exist (FP) and the number of possible edges which *do not* exist (F). The value of the threshold can be varied continuously over its range $\theta \in [0, 1]$, yielding a receiver operating characteristic (ROC) curve in the TPR - FPR space (fig. 2). The area under the curve (AUC) acts as a summary statistic for the quality of the classification. A predictor which is no better than uniformly random will yield a perfect diagonal with $AUC = 0.5$. Better predictors will yield curves which deviate from the diagonal with $|AUC - 0.5| \gg 0$. We consider predictors with $AUC \ll 0.5$ to be of high-quality because the ROC curve reflects over the diagonal if we change the classification criterion from $\phi_{i,j} \geq \theta$ to $\phi_{i,j} \leq \theta$. In other words, the predictor was good, but we were using the wrong criterion. It is important to note that this analysis does not tell us what the threshold θ should be, though in principle it can be determined. For now, it is sufficient to know only that such a threshold exists.

Results

We were interested in the relationship between the form of the game \mathcal{G} , the topology of the graph G , and the rule parameter β . To explore this, we considered 50 agents playing on graphs with one of 5 types of topology: a.) Barabási-Albert graphs, b.) Erdős-Rényi graphs, and the c.) cycle, d.) wheel and e.) lattice topologies. The Barabási-Albert model (Barabási and Albert, 1999) generates random graphs using preferential attachment; as new




Figure 1: Samples of the topologies considered. The Barabási-Albert (a-b) and Erdős-Rényi graphs (d-e) are generated non-deterministically according to edge density parameters m and p respectively. The cycle (c), wheel (f) and 5×10 lattice (g) topologies have only one representative each.

nodes are iteratively added to the graph they are randomly connected to m existent nodes with a bias toward those with higher degree; we considered only $m \in \{1, 5\}$ in this work. The Erdős-Rényi model (Erdős and Rényi, 1959) starts with the full set of nodes and adds an edge between each pair of distinct nodes with probability p . We chose $p \in \{0.04, 0.18\}$ to ensure that the edge density of the Barabási-Albert and Erdős-Rényi graphs are comparable. The cycle, wheel and 5×10 non-periodic lattice topologies, on the other hand, are non-random topologies, and are depicted in 1.

Once the graphs were constructed, we considered all games of the form eq. (1) with $S \in \{-1.0, -0.9, \dots, 1.0\}$ and $T \in \{0.0, 0.1, \dots, 2.0\}$. We then generated 100 graphs according to each of the parameters above for each combination of game and topology. For the cycle, wheel and lattice topologies, there is only one graph (up to isomorphism) for a given number of vertices, so we reran the simulations on the same topology 100 times. We simulated 10 rounds of play starting from 100 initial conditions selected uniformly at random. The initial conditions were generated independently for each simulation, i.e. each combination of game, graph and rule parameter. The time series of strategies in each simulation of 10 rounds with 100 replicates was used to generate an ROC curve. The 100 ROC curves for each game




Figure 2: Receiver operating characteristic curves for four canonical games with (T, S) parameterizations of (a) $(1/2, 1/2)$, (b) $(3/2, 1/2)$, (c) $(1/2, -1/2)$ and (d) $(3/2, -1/2)$ as played by 50 agents on Barabási-Albert graphs with $m = 5$. In all cases, the agents employ a threshold function with rule parameter $\beta = 1.0$. Each plot summarizes the results of simulations run on 100 randomly generated graphs. The black curve is the point-wise mean of the 100 resulting ROC curves, while the blue and orange curves are those with the greatest and least AUC, respectively. The dashed gray line represents the curve for a predictor which is no better than chance.

and topology combination (one for each randomly generated graph) were point-wise averaged to yield a mean ROC curve. Finally, the AUC was computed for each mean curve, and a scaled deviation of that value from the worst-case value of 0.5:

$$\overline{AUC} = 2|AUC - 0.5| \quad (8)$$

was plotted as a heatmap on an S - T plot (fig. 3, fig. 4 and fig. 5). Here S and T refer to the parameters which specify a game as in eq. (1). In other words, each point in the plot corresponds to a different game with the color roughly characterizing the effectiveness of ϕ_{ij} as a predictor for the existence of the edge (i, j) in the graph.

Discussion

There are a number of notable features apparent in fig. 3. The plots are roughly symmetric across their diagonals so that opposite quadrants have similar characteristics. The nice thing about these S - T parameterization is that each of the quadrants corresponds to a common two-player, two-strategy, symmetric game. Counter-clockwise from the top-right quadrant, the games are *Hawk-Dove*, *Harmony*, *Stag-Hunt* and *Prisoner's Dilemma*. The interesting point




Figure 3: \overline{AUC} heatmaps for games played on Barabási-Albert random graphs. Each row corresponds to a different connection parameter $m \in \{1, 5\}$, while the columns refer to the rule parameter used by the agents $\beta \in \{1, 10\}$. Each point in the plot corresponds to a different game parameterized as in eq. (1) by S and T . For each game, 100 random graphs were selected and an average ROC curve was generated. The color represents how much the AUC under the mean curve differs from the uniformly random predictor as in eq. (8).

to note is that the games opposite one another have the same number of pure-strategy Nash equilibria: Harmony and Prisoner's Dilemma each have one while Hawk-Dove and Stag-Hunt have two. On its face, this may not appear to be relevant to the problem of inferring the topology of the graph. However, the Nash equilibria play a crucial role given how the agents are updating their strategies: over time, the agents tend to settle on playing an equilibrium strategy. This process happens fairly quickly, typically within the 10 rounds of play simulated in this study. In effect, the simulations of the Harmony and Prisoner's Dilemma reach equilibrium and remain there, though with some amount of noise induced by the agents' strategy update rule. However, the transition to equilibrium takes longer for Hawk-Dove and Stag-Hunt because they have two equilibria. The resulting time series then provides more information about the interactions between agents. One additional factor to consider here is that in the two-equilibrium games, it is possible for domains to form wherein agents in a given domain all play the same equilibrium strategy while agents in neighboring domains play a different strategy. The stability of these domains and the impact they have on the graph inference problem is something to be investigated in future work.

One apparent problem with the above interpretation is it suggests that each of the four quadrants of the heatmap should be filled with the same \overline{AUC} value, e.g. that every harmony game should have the same $\overline{AUC} \sim 0$. Of course, that is clearly not the case (fig. 3(a)). Notice though that the




Figure 4: \overline{AUC} heatmaps for games played on Erdős-Rényi random graphs. Each row corresponds to a different connection probability $p \in \{0.04, 0.18\}$, while the columns refer to the rule parameter used by the agents $\beta \in \{1, 10\}$.

classification improves as the game moves away from the negatively-sloped diagonal within the Harmony and Prisoner's Dilemma quadrants. This effect is due to the agents using a stochastic update rule which is dependent on the difference in payoffs of two strategies. As the game moves further from the negatively-sloped diagonal, the payoff differences make differentiating between the strategies more difficult in games where there is only one equilibrium. As a result, the agents spend more time in transition to equilibrium and consequently provide more information about the topology of the graph. If this analysis is correct, we should expect to see that increasing the rule parameter β allows the agents to differentiate smaller payoff differences, and the quadrants of these heatmaps should tend to become more homogeneous. This is exactly what we see, particularly when edge density of the graphs is low (fig. 3, fig. 4 and fig. 5).

This brings us to the final factor involved in determining the efficacy of the classification: the edge density of the graph. When we compare, for example, the Barabási-Albert graphs with $m = 1$ (fig. 3(a-b)) to the Erdős-Rényi graphs with $p = 0.04$ (fig. 4(a-b)), we see almost identical structure in the \overline{AUC} heatmaps. On average the Erdős-Rényi graphs have the same edge density to their Barabási-Albert counterparts, i.e. they have about the same number of edges. The cycle graph (fig. 5(a-b)) has the same edge density to the Barabási-Albert graph ($m = 1$), and again we see almost identical heatmaps. By increasing the edge parameters for the Barabási-Albert and Erdős-Rényi graphs, m and p respectively, we can see a similar comparison for yet denser graphs (fig. 3(c-d) and fig. 4(c-d)). Again, the Barabási-Albert ($m = 5$) and




Figure 5: \overline{AUC} heatmaps for games played on the cycle, wheel and lattice topologies. Each row corresponds to a different topology while the columns refer to the rule parameter used by the agents $\beta \in \{1, 10\}$. The edge density of the topologies increases down the column, i.e. the cycle has fewer edges than the lattice which has fewer edges than the wheel.

Erdős-Rényi ($p = 0.18$) graphs have similar edge density. What's more, the same qualitative change is seen for a single type of topology when the edge density changes. For example, comparing fig. 3(a) to 3(c) or fig. 4(a) to 4(c), we see fewer harmony and prisoner's dilemma games admit a faithful inference. The same phenomenon is observed in comparing cycles, lattices and wheels (fig. 5), which further supports the claim. The graphs have remarkably distinct topologies even with only 50-nodes. For example, it is difficult to find any similarity between the Barabási-Albert and cycle graphs other than edge density. However, it is entirely plausible that some other topological similarity between the graphs is causing these effects. Increasing the size of the graph should lead to increasingly divergent topological features. That would allow us to test this claim more fully.

Conclusion

From all of these results, however incomplete, we can begin to draw some insight. It seems profoundly important that in order to be confident that the reconstructed topology – that is the inferred causal structure of the system – is the actual topology, an outside observer must be privy to

the inner workings of the agents. What game are they playing, and how are they deciding which strategy to use next? Of course, the simple retort to this is that reconstructions are failing for the one-equilibrium games because the time series are not diverse enough. If the observer could assess this in advance, it might temper or bolster their confidence. However, it cannot be understated that the rule parameter plays some role in the diversity of the time series. If the rule is too stochastic, the time series will be riddled with *biased* noise. One possibility is to use similar information-theoretic tools to infer which game is being played. It's been shown in previous work that dynamical agents operating under different rules show distinct patterns of information transfer (Valentini et al., 2018). If the observer could first apply such a method to determine the game³, rule, and noise profile, they may be able to yield more faithful reconstructions of the topology.

From here there are a number of avenues for further research. The mutual information is a somewhat simplistic measure for applications like this. We chose it in part because of this simplicity. Two alternative measures, for example, might be conditional mutual information (CMI) or transfer entropy (TE), (Schreiber, 2000). In preliminary work, we briefly explored both. The conditional mutual information approach was problematic in that it required more substantial computational overhead than we had available at the time. One method of using CMI is to compute $I(s_i^+, s_j | s_{ij})$ where $s_{ij} = \{s_k | k \neq i, k \neq j\}$, that is to say condition the pair-wise mutual information on the strategies of all other agents. Since our graphs have 50 nodes, that computation was infeasible using direct approaches. Alternative methods of computing CMI, such as the iterative approaches of (Lizier and Rubinov, 2012; Sun et al., 2015), may make it viable for future work. Transfer entropy is a special case of CMI; it would condition on the target agent's previous strategies. Both approaches are worth exploring more rigorously.

In this work, we were interested in methods that required only one only observe a system and not intervene on it, as in many cases only observation is accessible to the researcher. However, it is well understood that observation alone, particularly in the case of large systems wherein only a small portion of the state space can be explored, is inadequate for extracting causal relationships between components of a system. Interventional methods, such as (Ay and Polani, 2008), are much more effective in such cases, provided you have the ability to control aspects of the system. Subsequent work may explore such an approach.

The dynamics of the models we considered here are similar in many respects to the classic Ising model, with β acting much like temperature and the payoff matrix eq. (1) behaving similarly to interaction strengths. It is reasonable to ex-

³It may not be necessary to know precisely which game is being played. Perhaps knowing the number of equilibria is sufficient?

pect that analogues of heat capacity, entropy, and other thermodynamic measures could be related in some way to the effectiveness of the network inference. Comparing the AUC with some of these thermodynamic analogues is a possibility for a future work.

Another advance could be to look at other methods of quantifying the efficacy of the edge classifications. As the reader may have noticed, we never considered complete graphs. The reason for this is that to construct the ROC curves, you have to compute the false-positive rate. This value is undefined if your dataset has no negative cases. In our application, that means there must be pairs of agents that are not connected.

Two of the most egregious assumptions we've made in this work are that every agent is using the same rule to update its strategy, and every pair of agents is playing the same game. There is little reason to expect that real-world systems behave in this manner, and either assumption can be relaxed. For example, we could associate a different game with each edge. We have no expectation for how this would affect the results, but will explore it in the future.

Finally, as with virtually every simulation study, there is room for more fine-grained parameter sweeps. Our sampling of the S - T space of games was rather coarse; we limited ourselves to an unreasonably small number of initial conditions (100 of the 2^{50} possible) for each graph, and we considered only a handful of graph topologies. All of this will be greatly expanded in future work.

Ultimately, we see that the difference in functional and actual connectivity can sometimes be tied directly to the dynamical structure of the system. It may be impossible to know how similar the functional and actual connectivity of these types of networks are without detailed knowledge of the dynamics of the agents in the system. This may not be a surprise to some, but it is important that researchers to acknowledge it explicitly. Without that, it can be all too easy to blindly apply these network inference methods and over interpret the results, a sin of which at least the first author is guilty.

Acknowledgements

D.G.M. and S.I.W. would like to acknowledge the Lifelong Learning Machines program from DARPA/MTO for funding this work. D.G.M also acknowledges Cole Mathis for the conversations about game theory which inspired this work.

References

Antonioni, A., Martinez-Vaquero, L. A., Mathis, C., Peel, L., and Stella, M. (2018). Individual perception dynamics in drunk games. *arXiv:1807.08635*.

Ay, N. and Polani, D. (2008). Information flows in causal networks. *Advances in Complex Systems*, 11(01):17–41.

Bach, L. A., Bentzen, S. M., Alsner, J., and Christiansen, F. B. (2001). An evolutionary-game model of tumour-cell interactions: possible relevance to gene therapy. *European journal of cancer*, 37(16):2116–2120.

Banerji, C. R. S., Severini, S., and Teschendorff, A. E. (2013). Network transfer entropy and metric space for causality inference. *Physical review. E, Statistical, nonlinear, and soft matter physics*, 87(5):052814.

Barabási, A. L. and Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439):509–512.

Basanta, D., Simon, M., Hatzikirou, H., and Deutsch, A. (2008). Evolutionary game theory elucidates the role of glycolysis in glioma progression and invasion. *Cell proliferation*, 41(6):980–987.

Dingli, D., Chalub, F. A. C. C., Santos, F. C., Van Segbroeck, S., and Pacheco, J. M. (2009). Cancer phenotype as the outcome of an evolutionary game between normal and malignant cells. *British journal of cancer*, 101(7):1130–1136.

Erdős, P. and Rényi, A. (1959). On random graphs I. *Publicationes Mathematicae Debrecen*, 6:290–297.

Gatenby, R. A. and Vincent, T. L. (2003). Application of quantitative models from population biology and evolutionary game theory to tumor therapeutic strategies. *Molecular cancer therapeutics*, 2(9):919–927.

Hummert, S., Bohl, K., Basanta, D., Deutsch, A., Werner, S., Theissen, G., Schroeter, A., and Schuster, S. (2014). Evolutionary game theory: cells as players. *Molecular bioSystems*, 10(12):3044–3065.

Ito, S., Hansen, M. E., Heiland, R., Lumsdaine, A., Litke, A. M., and Beggs, J. M. (2011). Extending transfer entropy improves identification of effective connectivity in a spiking cortical network model. *PLoS one*, 6(11):e27431.

Lizier, J. T. and Rubinov, M. (2012). Multivariate construction of effective computational networks from observational data. Preprint, Max Planck Institute for Mathematics in the Sciences.

Nowak, M. A. and Sigmund, K. (2004). Evolutionary dynamics of biological games. *Science*, 303(5659):793–799.

Orlandi, J. G., Stetter, O., Soriano, J., Geisel, T., and Battaglia, D. (2014). Transfer entropy reconstruction and labeling of neuronal connections from simulated calcium imaging. *PLoS one*, 9(6):e98842.

Riechert, S. E. and Hammerstein, P. (1983). Game theory in the ecological context. *Annual Review of Ecology and Systematics*, 14(1):377–409.

Schreiber, T. (2000). Measuring information transfer. *Physical review letters*, 85(2):461–464.

Smith, J. M. (1982). *Evolution and the Theory of Games*. Cambridge University Press.

Smith John Maynard, Maynard Smith J., and Holliday Robin (1979). Game theory and the evolution of behaviour. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 205(1161):475–488.

- Sun, J. and Bollt, E. M. (2014). Causation entropy identifies indirect influences, dominance of neighbors and anticipatory couplings. *Physica D. Nonlinear phenomena*, 267:49–57.
- Sun, J., Taylor, D., and Bollt, E. (2015). Causal network inference by optimal causation entropy. *SIAM journal on applied dynamical systems*, 14(1):73–106.
- Valentini, G., Moore, D. G., Hanson, J. R., Pavlic, T. P., Pratt, S. C., and Walker, S. I. (2018). Transfer of information in collective decisions by artificial agents. *The 2018 Conference on Artificial Life: A Hybrid of the European Conference on Artificial Life (ECAL) and the International Conference on the Synthesis and Simulation of Living Systems (ALIFE)*, pages 641–648.
- Ver Steeg, G. and Galstyan, A. (2012). Information transfer in social media. In *Proceedings of the 21st International Conference on World Wide Web, WWW '12*, pages 509–518, New York, NY, USA. ACM.
- Vicente, R., Wibral, M., Lindner, M., and Pipa, G. (2011). Transfer entropy—a model-free measure of effective connectivity for the neurosciences. *Journal of computational neuroscience*, 30(1):45–67.