

Bounded Degree Approximations of Stochastic Networks

Christopher J. Quinn

*School of Industrial Engineering
Purdue University
West Lafayette, Indiana 47907, USA*

CJQUINN@PURDUE.EDU

Ali Pinar

*Data Science & Cyber Analytics Department
Sandia National Laboratories¹
Livermore, CA 94551, USA*

APINAR@SANDIA.GOV

Negar Kiyavash

*Department of Industrial and Enterprise Systems Engineering
University of Illinois
Urbana, Illinois 61801, USA*

KIYAVASH@ILLINOIS.EDU

Editor:

Abstract

We propose algorithms to approximate directed information graphs. Directed information graphs are probabilistic graphical models that depict causal dependencies between stochastic processes in a network. The proposed algorithms identify optimal and near-optimal approximations in terms of Kullback-Leibler divergence. The user-chosen sparsity trades off the quality of the approximation against visual conciseness and computational tractability. One class of approximations contains graphs with specified in-degrees. Another class additionally requires that the graph is connected. For both classes, we propose algorithms to identify the optimal approximations and also near-optimal approximations, using a novel relaxation of submodularity. We also propose algorithms to identify the r -best approximations among these classes, enabling robust decision making.

Keywords: probabilistic graphical models, network inference, causality, submodularity, approximation algorithms

1. Introduction

Many fields of the sciences and engineering require analysis, modeling, and decision making using networks, typically represented by graphs. Social networks, financial networks, and biological networks are a few categories that are relevant not only academically but also in daily life. A major challenge for studying networks is identifying a concise topology, such as who strongly influences whom in a social network. Real world networks are often large—there are tens of thousands of human genes and trillions of connections in the human brain. Such scales make human visual processing of and decision making with the whole network prohibitive. This paper investigates algorithms to identify provably good approximations of the network topology, which capture important system dynamics while significantly reducing the number of edges to enable tractable analysis.

(1) Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Across different domains, edges are used to model various kinds of ties such as physical connections or dynamic relationships. For instance, in depicting a computer network, an edge might correspond to a physical wire or a packet exchange between a sender and a receiver. For the human brain, edges could correspond to information flow between different cells or brain regions (Takahashi et al., 2015; Kim et al., 2014). For online social networks, edges might represent user-defined relationships or pairs of users that frequently message each other (Ver Steeg and Galstyan, 2012, 2013). Edges that represent dynamics often must be inferred from activity, in some cases statistically.

There is a large literature on graphs whose edges depict statistical relationships. Markov and Bayesian networks are well-known probabilistic graphical models whose edges represent correlation. Many methods have been proposed to infer and approximate the networks from i.i.d. data, often relying on heuristics. For an overview, see Chapters 18 and 20 in Koller and Friedman (2009). For applications involving agents interacting with each other over time, whether in finance, biology, social networks, or other domains, there is interest in identifying and representing causal influences between the agents, not just correlation.

One approach uses known families of models, such as with structural equation modeling, to distinguish cause and effect (Zhang et al., 2015, 2014; Chen et al., 2012). A recent work uses belief propagation to infer directionality (Chang et al., 2014). Alternatively, under appropriate conditions such as with expert labeling and no feedback, Bayesian networks can depict causal relationships using Pearl’s interventional calculus (Pearl, 2009). We consider the general setting when such conditions or modeling assumptions might not hold.

Recently, directed information graphs were introduced to address this issue (Quinn et al., 2011; Amblard and Michel, 2011). Edges in directed information graphs depict statistical causation between non-i.i.d. time-series. In this work, “statistical causation” is in the sense of Granger causality (Granger, 1969), where a process \mathbf{X} statistically causes \mathbf{Y} if in sequentially predicting Y_t , knowledge of the past X^{t-1} helps in prediction even when Y^{t-1} and the past of all the other processes are already known. These graphs use directed information, an information theoretic quantity, which is well-defined for any class of stochastic processes. Directed information has been applied to a range of settings, such as neuroscience (Quinn et al., 2011; Kim et al., 2011; So et al., 2012; Kim et al., 2014), gene regulatory networks (Rao et al., 2007, 2008), and online social networks (Ver Steeg and Galstyan, 2012, 2013; Quinn et al., 2012).

For networks with thousands or millions of edges, directed information graphs become too complicated for direct human analysis. A major approach to simplifying the graphs is to only keep a few edges which together best approximate the dynamics of the system. For example, a directed tree is among the simplest graphs. See Figure 1. Each node has only one parent. Trees have the fewest number of edges possible while being connected. There is a root node and a path from the root to every other node. The graph is concise, facilitating human analysis and decision making. A recent work proposed an efficient algorithm to identify the best directed tree approximation, where goodness of approximation is measured by Kullback-Leibler (KL) divergence from the full joint distribution to the distribution induced by the directed tree (Quinn et al., 2013a). In addition to being computationally efficient, the algorithm in Quinn et al. (2013a) only uses joint statistics for pairs of processes and does not require the full joint distribution to find the best approximation.

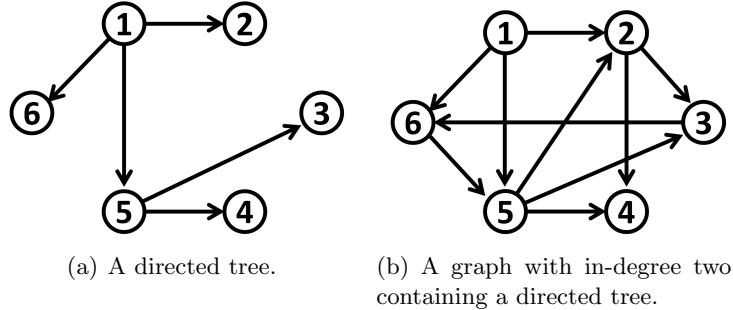


Figure 1: Diagrams for two network approximations.

Though directed tree approximations are easy to comprehend and efficient to construct, they cannot depict feedback. Feedback is essential in many networks, such as in the brain and gene regulatory networks. Thus, for some applications, it is necessary to consider higher order approximations. For instance, a graph with in-degrees two and three and containing a directed spanning tree as a subgraph would trade-off some simplicity and computational efficiency in order to capture more complex relationships in the network.

1.1 Our Contributions

We propose an algorithm to identify the optimal connected bounded in-degree approximations. The algorithm requires only low-dimensional statistics, similar to the algorithm for directed tree approximations. The user decides how complex to make the approximations, changing the in-degrees to trade off visual and computational simplicity against the accuracy of the approximation.

Identifying optimal approximations becomes prohibitive for large in-degrees. For situations where a near-optimal approximation would suffice, we propose algorithms using a greedy search. We identify sufficient conditions, namely a relaxed form of submodularity, that ensure near-optimality.

Additionally, having multiple, good approximations can aid in understanding network dynamics. Instead of just having the best approximation, having the five or ten best approximations in order can yield insight into which edges are most important—those that persist in the top approximations—and those that are less significant. Being able to identify the top- r approximations also enables the user to identify the best approximation of more restricted classes of topologies. For example, suppose that the best directed tree approximation for a network had a height of six. If the user desires the best directed tree approximation with height less than four, he/she can look among the top- r approximations until finding a tree with height less than four and it would necessarily be the best such approximation. We develop algorithms to identify the top- r approximations with similar complexity as finding the optimal approximation.

Lastly, we use simulations to validate the quality of the approximations found.

1.2 Related Work

There is a large body of work on approximating Bayesian and Markov networks. One well known result is an algorithm to identify optimal tree approximations (Chow and Liu, 1968). The algorithm finds a maximum weight spanning tree using mutual information for weights and only requires distributions of pairs of variables.

In general, identifying more complex approximations cannot be done in a computationally efficient manner. Bayesian networks are NP-hard to approximate for topologies with specified in-degree larger than one (Chickering, 1996) and even polytrees with in-degree two (Dasgupta, 1999). Some works have focused on identifying optimal approximations of subclasses of polytrees. One work finds the best bounded in-degree approximation that preserves the statistical dependencies in the best tree approximation (Carvalho and Oliveira, 2007). Another work finds an optimal polytree that can be converted to a tree with a bounded number of edge or node deletions (Gaspers et al., 2012).

Other approaches to approximating graphical models include using l_1 -regularized regression to identify sparse Ising models for Markov networks with binary variables (Ravikumar et al., 2010). Another approach proposes a linear programming relaxation coupled with branch and bound to find an optimal approximation (Jaakkola et al., 2010). Annealed importance sampling is used in Niinimäki and Koivisto (2013); see references therein for Markov chain Monte Carlo based techniques. The performance of a forward-backward greedy search for Markov networks in a high-dimensional setting is studied in Jalali et al. (2011). In Pernkopf and Bilmes (2010), an algorithm is proposed to first identify an variable ordering and then greedily select parents.

There has been much less work developing approximations for directed information graphs. In Quinn et al. (2013a), an algorithm is proposed to identify the best directed spanning tree approximation for directed information graphs. In Quinn et al. (2012), several algorithms are introduced for inferring the exact topology. One of the algorithms can be also used to compute the best approximation where the only topological constraints are user-specified in-degrees. That is discussed here as Algorithm 1 in Section 4. Several works investigated sparse approximations using lasso and related penalties when processes are jointly autoregressive with Gaussian noise (Charbonnier et al., 2010; Haufe et al., 2010; Bolstad et al., 2011; Jung et al., 2014; Basu et al., 2015).

In our preliminary work Quinn et al. (2013b), we developed an algorithm to identify the optimal bounded in-degree approximation containing a directed spanning tree subgraph. This appears here as Algorithm 2. Also, a sufficient condition for a greedy search to return near-optimal approximations was identified in Quinn et al. (2013b), presented here as Definition 10.

There has been research in the graphical models literature for finding the top- r solutions for problems such as the MAP realizations for Bayesian or Markov networks (Nilsson, 1998; Yanover and Weiss, 2004; Fromer and Globerson, 2009; Flerova et al., 2012; Batra et al., 2012). The present work focuses on finding the top- r solutions for structure learning.

1.3 Paper Organization

The paper is organized as follows. Definitions and notations are introduced in Section 2. Section 3 reviews directed information graphs. Section 4 presents algorithms to identify the

optimal bounded in-degree approximations. Section 5 identifies a sufficient condition for the greedy search to construct near optimal approximations. Section 6 describes an algorithm to find the top- r approximations. Algorithmic complexity is discussed in Section 7. The algorithms are empirically evaluated in Section 8. Section 9 concludes the paper. Proofs are in the appendix.

2. Notation and Information-Theoretic Definitions

We now define notation. We use “ \coloneqq ” for denoting.

- For a sequence a_1, a_2, \dots , denote (a_i, \dots, a_j) as a_i^j and $a^k := a_1^k$. Let $[m] := \{1, \dots, m\}$ and the power set $2^{[m]}$ on $[m]$ to be the set of all subsets of $[m]$.
- We consider m finite-alphabet, discrete-time random processes over a horizon n . Let X denote the alphabet and $\mathcal{P}(\mathsf{X})$ the space of probability measures on X . Denote the i th random variable at time t by $X_{i,t}$, the i th random process as $\mathbf{X}_i = (X_{i,1}, \dots, X_{i,n})^\top$, the whole collection of all m random processes as $\underline{\mathbf{X}} = (\mathbf{X}_1, \dots, \mathbf{X}_m)^\top$, and a subset of K processes indexed by $A \subseteq [m]$ as $\underline{\mathbf{X}}_A = (\mathbf{X}_{A(1)}, \dots, \mathbf{X}_{A(K)})^\top$.

Remark 1 *We consider the finite-alphabet setting to simplify the presentation. The results extend to more general cases.*

- Conditional and *causally conditioned* distributions (Kramer, 1998) of \mathbf{X}_i given \mathbf{X}_j are

$$P_{\mathbf{X}_i|\mathbf{X}_j}(\mathbf{x}_i|\mathbf{x}_j) := \prod_{t=1}^n P_{X_{i,t}|X_i^{t-1}, X_j^n}(x_{i,t}|x_i^{t-1}, x_j^n) \quad (1)$$

$$P_{\mathbf{X}_i\|\mathbf{X}_j}(\mathbf{x}_i\|\mathbf{x}_j) := \prod_{t=1}^n P_{X_{i,t}|X_i^{t-1}, X_j^{t-1}}(x_{i,t}|x_i^{t-1}, x_j^{t-1}). \quad (2)$$

Note the similarity between (1) and (2), though in (2) the present and future, $x_{j,t}^n$, is not conditioned on. In Kramer (1998), the present $x_{j,t}$ was conditioned on. The reason we remove it will be made clear in Remark 2.

- Consider the set of processes $\underline{\mathbf{X}}_A$ for some $A \subseteq [m] \setminus \{i\}$. Next consider two sets of causally conditioned distributions $\{P_{\mathbf{X}_i\|\underline{\mathbf{X}}_A=\mathbf{x}_A} \in \mathcal{P}(\mathsf{X}) : \underline{\mathbf{x}}_A \in \mathsf{X}^{|A|n}\}$ and $\{Q_{\mathbf{X}_i\|\underline{\mathbf{X}}_A=\mathbf{x}_A} \in \mathcal{P}(\mathsf{X}) : \underline{\mathbf{x}}_A \in \mathsf{X}^{|A|n}\}$ along with a marginal distribution $P_{\underline{\mathbf{X}}_A} \in \mathcal{P}(\mathsf{X}^{|A|n})$. Then the conditional Kullback-Leibler (KL) divergence between causally conditioned distributions is given by

$$\begin{aligned} & D(P_{\mathbf{X}_i\|\underline{\mathbf{X}}_A} \| Q_{\mathbf{X}_i\|\underline{\mathbf{X}}_A} | P_{\underline{\mathbf{X}}_A}) \\ &:= \sum_{t=1}^n \sum_{\underline{x}_A^{t-1}} D\left(P_{X_{i,t}|\underline{\mathbf{X}}_A^{t-1}=\underline{x}_A^{t-1}} \| Q_{X_{i,t}|\underline{\mathbf{X}}_A^{t-1}=\underline{x}_A^{t-1}}\right) P_{\underline{\mathbf{X}}_A^{t-1}}(\underline{x}_A^{t-1}). \end{aligned} \quad (3)$$

- Let $i, j \in [m]$ and $A \subseteq [m] \setminus \{i, j\}$. The mutual information, *directed information* (Marko, 1973), and causally conditioned directed information (Kramer, 1998) are

$$I(\mathbf{X}_j; \mathbf{X}_i) := D(P_{\mathbf{X}_i, \mathbf{X}_j} \| P_{\mathbf{X}_i} P_{\mathbf{X}_j}) = D(P_{\mathbf{X}_i | \mathbf{X}_j} \| P_{\mathbf{X}_i} | P_{\mathbf{X}_j}) \quad (4)$$

$$\begin{aligned} &= \sum_{t=1}^n I(X_j^n; X_{i,t} | X_i^{t-1}) \\ I(\mathbf{X}_j \rightarrow \mathbf{X}_i) &:= D(P_{\mathbf{X}_i \| \mathbf{X}_j} \| P_{\mathbf{X}_i} | P_{\mathbf{X}_j}) \\ &= \sum_{t=1}^n I(X_j^{t-1}; X_{i,t} | X_i^{t-1}) \\ I(\mathbf{X}_j \rightarrow \mathbf{X}_i \| \underline{\mathbf{X}}_A) &:= D(P_{\mathbf{X}_i \| \underline{\mathbf{X}}_{A \cup \{j\}}} \| P_{\mathbf{X}_i \| \underline{\mathbf{X}}_A} | P_{\mathbf{X}_{A \cup \{j\}}}) \\ &= \sum_{t=1}^n I(X_j^{t-1}; X_{i,t} | X_i^{t-1}, \underline{\mathbf{X}}_A^{t-1}). \end{aligned} \quad (5)$$

While mutual information quantifies statistical correlation (in the colloquial sense of statistical interdependence), directed information quantifies statistical *causation* in the sense of Granger causality (Quinn et al., 2012; Amblard and Michel, 2012). Note that $I(\mathbf{X}_j; \mathbf{X}_i) = I(\mathbf{X}_i; \mathbf{X}_j)$, but $I(\mathbf{X}_j \rightarrow \mathbf{X}_i) \neq I(\mathbf{X}_i \rightarrow \mathbf{X}_j)$ in general.

Remark 2 In (2) and (5), there is no conditioning on the present $X_{j,t}$. This follows Marko's definition (Marko, 1973) and is consistent with Granger causality (Granger, 1969). Massey (1990) and Kramer (1998) later included conditioning on $X_{j,t}$ for the specific setting of communication channels.

3. Directed Information Graphs

In this section, we briefly review directed information graphs (Quinn et al., 2011; Amblard and Michel, 2011).

Definition 3 A directed information graph is a probabilistic graphical model where each node represents a process \mathbf{X}_i and an edge $\mathbf{X}_j \rightarrow \mathbf{X}_i$ is drawn if

$$I(\mathbf{X}_j \rightarrow \mathbf{X}_i \| \underline{\mathbf{X}}_{[m] \setminus \{i,j\}}) > 0.$$

It follows immediately that directed information graphs are unique for a given distribution $P_{\underline{\mathbf{X}}}$. Under certain conditions, the directed information graph corresponds to a particular factorization of the joint distribution. By the chain rule, the joint distribution $P_{\underline{\mathbf{X}}}$ factorizes over time as $P_{\underline{\mathbf{X}}}(\underline{\mathbf{x}}) = \prod_{t=1}^n P_{\underline{\mathbf{x}}_t | \underline{\mathbf{x}}^{t-1}}(\underline{\mathbf{x}}_t | \underline{\mathbf{x}}^{t-1})$. If given the full past $\underline{\mathbf{x}}^{t-1}$, the processes $\{\mathbf{X}_1, \dots, \mathbf{X}_m\}$ at time t are mutually independent, $P_{\underline{\mathbf{X}}}$ can be further factorized as

$$P_{\underline{\mathbf{X}}}(\underline{\mathbf{x}}) = \prod_{t=1}^n \prod_{i=1}^m P_{X_{i,t} | \underline{\mathbf{x}}^{t-1}}(x_{i,t} | \underline{\mathbf{x}}^{t-1}), \quad (6)$$

and $P_{\underline{\mathbf{X}}}$ is said to be *strictly causal*. Equation 6 can be written using causal conditioning notation (2) as $P_{\underline{\mathbf{X}}}(\underline{\mathbf{x}}) = \prod_{i=1}^m P_{\mathbf{X}_i \| \underline{\mathbf{X}}_{[m] \setminus \{i\}}}(\mathbf{x}_i \| \underline{\mathbf{x}}_{[m] \setminus \{i\}})$. A distribution $P_{\underline{\mathbf{X}}}$ is said to be *positive* if $P_{\underline{\mathbf{X}}}(\underline{\mathbf{x}}) > 0$ for all $\underline{\mathbf{x}} \in \mathcal{X}^m$.

Theorem 4 (Quinn et al., 2012) For a joint distribution $P_{\underline{\mathbf{X}}}$, if $P_{\underline{\mathbf{X}}}$ is positive and strictly causal, then the parent sets $\{A(i)\}_{i=1}^m$ in the directed information graph are the unique, minimal cardinality parent sets such that $D(P_{\underline{\mathbf{X}}} \parallel \prod_{i=1}^m P_{\mathbf{X}_i \parallel \underline{\mathbf{X}}_{A(i)}}) = 0$.

A graphical separation criteria, similar to d-separation for Bayesian networks, applies to directed information graphs (Eichler, 2012).

4. Optimal Bounded In-Degree Approximations

When the exact topology is not necessary or is prohibitive to learn, approximations can be useful. Approximations with simple topologies facilitate visual comprehension and in some cases can be efficient to identify. We investigate algorithms to identify optimal approximations for two settings. Goodness of the approximations is measured by the KL divergence between the full joint distribution and the distribution induced by the approximation. The researcher specifies the in-degrees, controlling the complexity. Also, the optimal approximations will be identified using low dimensional statistics, not the whole joint distribution.

We consider approximations of the form

$$\hat{P}_{\underline{\mathbf{X}}}(\underline{\mathbf{x}}) := \prod_{i=1}^m P_{\mathbf{X}_i \parallel \underline{\mathbf{X}}_{A(i)}}(\mathbf{x}_i \parallel \underline{\mathbf{x}}_{A(i)}), \quad (7)$$

where the $A(i) \subseteq [m] \setminus \{i\}$ are candidate parent sets. Let \mathcal{G} denote the set of such approximations. The goal is to find the $\hat{P}_{\underline{\mathbf{X}}} \in \mathcal{G}$ that minimizes the KL divergence $D(P_{\underline{\mathbf{X}}} \parallel \hat{P}_{\underline{\mathbf{X}}})$. The following theorem characterizes an important decomposition property for evaluating the quality of an approximation $\hat{P}_{\underline{\mathbf{X}}}$. The approximation that minimizes the KL divergence is the one that maximizes a sum of directed informations from parent sets to children.

Theorem 5 (Quinn et al., 2013a) For any distribution $P_{\underline{\mathbf{X}}}$,

$$\arg \min_{\hat{P}_{\underline{\mathbf{X}}} \in \mathcal{G}} D(P_{\underline{\mathbf{X}}} \parallel \hat{P}_{\underline{\mathbf{X}}}) = \arg \max_{\hat{P}_{\underline{\mathbf{X}}} \in \mathcal{G}} \sum_{i=1}^m I(\mathbf{X}_{A(i)} \rightarrow \mathbf{X}_i). \quad (8)$$

This decomposition property will be important for the following results.

4.1 An Unconstrained Formulation

Consider finding an optimal approximation where the only constraint is that the in-degrees are $K \geq 1$. Specifically, consider approximations of the form (7) where $|A(i)| = K$. We assume uniform K for simplicity. The results hold if K is a function of i . Let \mathcal{G}_K denote the set of all such approximations. The formula (8) simplifies.

Corollary 6 (Quinn et al., 2012) For any distribution $P_{\underline{\mathbf{X}}}$, the parent sets $\{A^*(i)\}_{i=1}^m$ corresponding to an optimal approximation $\hat{P}^* \in \arg \min_{\hat{P}_{\underline{\mathbf{X}}} \in \mathcal{G}_K} D(P_{\underline{\mathbf{X}}} \parallel \hat{P}_{\underline{\mathbf{X}}})$ satisfy

$$A^*(i) \in \arg \max_{A(i):|A(i)|=K} I(\mathbf{X}_{A(i)} \rightarrow \mathbf{X}_i).$$

Algorithm 1. OPTIMALGENERAL (Quinn et al., 2012)**Input:** $\mathcal{DI}_{\text{BndInd}}, K, m$

1. **For** $i \in [m]$
2. $A(i) \leftarrow \emptyset$
3. $\mathcal{B} \leftarrow \{B : B \subseteq [m] \setminus \{i\}, |B| = K\}$
4. $A(i) \leftarrow \arg \max_{B \in \mathcal{B}} I(\underline{\mathbf{X}}_B \rightarrow \mathbf{X}_i)$
5. **Return** $\{A(i)\}_{i=1}^m$

Thus, finding the optimal structure is equivalent to finding the best individual parent sets for each node. The process is described in Algorithm 1. A modified Algorithm 1 for exact structure learning was presented in Quinn et al. (2012). Algorithm 1 takes as input the following set of directed information values,

$$\mathcal{DI}_{\text{BndInd}} = \left\{ I(\underline{\mathbf{X}}_{B(i)} \rightarrow \mathbf{X}_i) : i \in [m], B(i) \subseteq [m] \setminus \{i\}, |B(i)| = K \right\}.$$

Theorem 7 (Quinn et al., 2012) *Algorithm 1 returns an optimal approximation $\hat{P}_{\underline{\mathbf{X}}} \in \mathcal{G}_K$.*

We next consider a more specific class of graph structures.

4.2 Finding a Connected Graph

Algorithm 1 might return an unconnected graph. For situations where information or influence propagates in the network, it can be better to work with connected structures. Directed trees, the simplest connected structure, were investigated in Quinn et al. (2013a). While visually simple and computationally easy to identify, they cannot depict complex dynamics such as feedback. We next consider a balance between the properties of unconstrained bounded in-degree approximations and directed trees. The new approximations contain a directed spanning tree as a subgraph and have user-specified in-degrees. See Figure 1(b). Note that the root node has no parents. Remark 9 will explain how to obtain graphs where the root also has parents.

Let $\tilde{\mathcal{G}}_K$ be the set of all graphs containing a spanning tree and all nodes except the root have in-degree $K \geq 1$. Let $\tilde{A}(i, j)$ be the best set of K parents for \mathbf{X}_i that contains the edge $\mathbf{X}_j \rightarrow \mathbf{X}_i$,

$$\tilde{A}(i, j) = \arg \max_{A(i) : A(i) \subseteq [m] \setminus \{i\}, j \in A(i)} I(\underline{\mathbf{X}}_{A(i)} \rightarrow \mathbf{X}_i). \quad (9)$$

Then assign weight $I(\mathbf{X}_{\tilde{A}(i, j)} \rightarrow \mathbf{X}_i)$ to edge $\mathbf{X}_j \rightarrow \mathbf{X}_i$ in the complete graph and run a maximum weight directed spanning tree (MWDST) algorithm. Each edge $\mathbf{X}_j \rightarrow \mathbf{X}_i$ in the spanning tree induces the corresponding parent set $\tilde{A}(i, j)$ for \mathbf{X}_i . This process is described in Algorithm 2.

Theorem 8 *Algorithm 2 returns an optimal approximation $\hat{P}_{\underline{\mathbf{X}}} \in \tilde{\mathcal{G}}_K$.*

Algorithm 2. OPTIMALCONNECTED**Input:** $\mathcal{DI}_{\text{BndInd}}, K, m$

1. **For** $i \in [m]$
2. $A(i) \leftarrow \emptyset$
3. **For** $j \in [m] \setminus \{i\}$
4. $\mathcal{B} \leftarrow \{B : B \subseteq [m] \setminus \{i\}, |B| = K, j \in B\}$
5. $\tilde{A}(i, j) \leftarrow \arg \max_{B \in \mathcal{B}} I(\underline{\mathbf{X}}_B \rightarrow \mathbf{X}_i)$
6. $\{a(i)\}_{i=1}^m \leftarrow \text{MWDST}(\{I(\underline{\mathbf{X}}_{\tilde{A}(i, j)} \rightarrow \mathbf{X}_i)\}_{1 \leq i \neq j \leq m})$
7. **For** $i \in [m]$
8. $A(i) \leftarrow \tilde{A}(i, a(i))$
9. **Return** $\{A(i)\}_{i=1}^m$

The proof is in Appendix A.

Remark 9 *The approximations $\hat{P}_{\underline{\mathbf{X}}} \in \tilde{\mathcal{G}}_K$ have root nodes with no inward edges. Algorithm 2 can be modified to find the best approximation where all nodes have in-degree K and there is a directed spanning tree as a subgraph. Namely, create a dummy node \mathbf{X}_0 , set edge weights $I(\mathbf{X}_j \rightarrow \mathbf{X}_0) \leftarrow -\infty$ and $I(\mathbf{X}_0 \rightarrow \mathbf{X}_j) \leftarrow -1$ for all $j \in [m]$. Note that all the other edge weights are directed informations, which are KL divergences and hence non-negative. Then Algorithm 2 will set \mathbf{X}_0 as the root with a single outward edge.*

Algorithms 1 and 2 find optimal approximations in terms of KL divergence $D(P_{\underline{\mathbf{X}}} \parallel \hat{P}_{\underline{\mathbf{X}}})$. They only need distributions over $K + 1$ processes, not the full joint distribution. However, they compute $m \binom{m-1}{K}$ directed informations involving K processes. If K is large, this could be computationally difficult. For some applications, instead of reducing K , it is better to efficiently identify near-optimal approximations.

5. Near-Optimal Bounded In-Degree Approximations

We next find sufficient conditions to identify near-optimal approximations in time polynomial in K .

5.1 Greedy Submodularity

Consider the following greedy procedure to select a parent set for \mathbf{X}_i . Initially, set \mathbf{X}_i 's parent set as the best individual parent $\mathbf{Z} = \arg \max_j I(\mathbf{X}_j \rightarrow \mathbf{X}_i)$. Then look for the second best parent $\arg \max_j I(\mathbf{X}_j \rightarrow \mathbf{X}_i \parallel \mathbf{Z})$. Repeat this $K - 2$ times, adding one parent at each iteration.

In general, greedy methods are not provably good. We next describe sufficient conditions to guarantee near-optimality.

Definition 10 *A joint distribution $P_{\underline{\mathbf{X}}}$ is called greedily-submodular if there exists an $\alpha > 0$, such that for any process \mathbf{Y} and any subset $\underline{\mathbf{X}}_{\mathbf{W}}$ of other processes,*

$$I(\mathbf{X}_j \rightarrow \mathbf{Y} \parallel \mathbf{X}_1, \dots, \mathbf{X}_{j-2}, \mathbf{X}_{j-1}) \leq \alpha I(\mathbf{X}_{j-1} \rightarrow \mathbf{Y} \parallel \mathbf{X}_1, \dots, \mathbf{X}_{j-2}), \quad (10)$$

for all $1 \leq j < |\underline{W}|$ where the processes in $\underline{\mathbf{X}}_{\underline{W}}$ are indexed according to the order in which they are selected by the greedy algorithm.

This is a weaker condition than submodularity, a discrete analog of concavity (Nemhauser et al., 1978). If $P_{\underline{\mathbf{X}}}$ has submodular directed information values, then for all pairs of processes $\{\mathbf{X}_j, \mathbf{Y}\}$ and sets of other processes $\underline{\mathbf{X}}_S \subseteq \underline{\mathbf{X}}_{S'} \subseteq \underline{\mathbf{X}} \setminus \{\mathbf{X}_j, \mathbf{Y}\}$,

$$I(\mathbf{X}_j \rightarrow \mathbf{Y} \mid \underline{\mathbf{X}}_{S'}) \leq I(\mathbf{X}_j \rightarrow \mathbf{Y} \mid \underline{\mathbf{X}}_S). \quad (11)$$

Submodularity implies conditioning does not increase directed information.

Corollary 11 *If $P_{\underline{\mathbf{X}}}$ is submodular, it is also greedily-submodular with $\alpha \leq 1$.*

Proof Let $S = \{1, \dots, j-2\}$ and $S' = S \cup \{j-1\}$, and let the processes be labeled in the order they are selected in a greedy search for parents for \mathbf{Y} . Then if $P_{\underline{\mathbf{X}}}$ is submodular,

$$I(\mathbf{X}_j \rightarrow \mathbf{Y} \mid \mathbf{X}_1, \dots, \mathbf{X}_{j-2}, \mathbf{X}_{j-1}) \leq I(\mathbf{X}_j \rightarrow \mathbf{Y} \mid \mathbf{X}_1, \dots, \mathbf{X}_{j-2}), \quad (12)$$

$$\leq I(\mathbf{X}_{j-1} \rightarrow \mathbf{Y} \mid \mathbf{X}_1, \dots, \mathbf{X}_{j-2}), \quad (13)$$

where (12) would hold by (11) and (13) would hold because \mathbf{X}_j is picked after \mathbf{X}_{j-1} in a greedy search. Thus, (11) holds with $\alpha \leq 1$. \blacksquare

Entropy is submodular (Fujishige, 1978). However, in general mutual information and directed information are not, as shown in the following example.

Example 1 Consider a network of three zero-mean, jointly Gaussian processes $\{\mathbf{X}, \mathbf{Y}, \mathbf{Z}\}$. Let \mathbf{N} denote noise. Let $\{\mathbf{N}, \mathbf{X}, \mathbf{Z}\}$ be mutually independent i.i.d. Gaussian processes. Let $Y_{t+1} = X_t + Z_t + N_t$. Then using stationarity (Cover and Thomas, 2006, pg. 256),

$$\begin{aligned} I(\mathbf{X} \rightarrow \mathbf{Y}) = I(X_1; Y_2) &= \frac{1}{2} \log \left(1 + \frac{\text{var}(X_1)}{\text{var}(Z_1) + \text{var}(N_1)} \right) \\ &< \frac{1}{2} \log \left(1 + \frac{\text{var}(X_1)}{\text{var}(N_1)} \right) \\ &= I(X_1; Y_2 | Z_1) = I(\mathbf{X} \rightarrow \mathbf{Y} \mid \mathbf{Z}). \end{aligned}$$

Since conditioning can increase directed information, it is not submodular.

Remark 12 The authors are not aware of this property being discussed in the literature previously. Two other conditions that are weaker than submodularity are discussed in Cevher and Krause (2011) and Das and Kempe (2011). The former uses submodularity up to an additive error. The latter uses submodularity up to multiplicative error. Both measure the increase in conditioning of the terms in (11), unlike (10) which only bounds sequential increases while greedily selecting a parent set.

Assumption 1 We assume that $P_{\underline{\mathbf{X}}}$ is greedily-submodular.

When Assumption 1 holds, the greedy search yields a near-optimal approximation. Let A denote the set of indices for an optimal set of K parents and B the indices for the greedily selected set of $L \leq K$ parents.

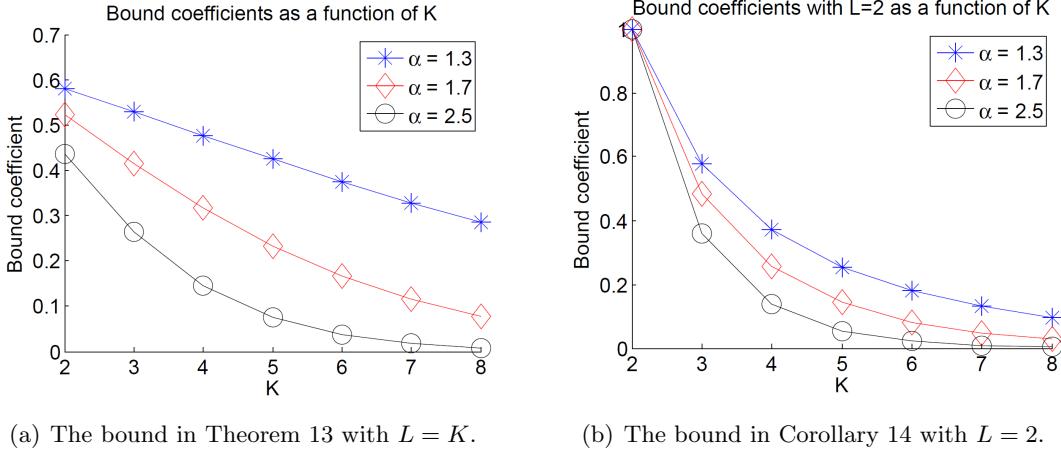


Figure 2: Plots for the bounds in Theorem 13 and Corollary 14 respectively.

Theorem 13 *Under Assumption 1,*

$$I(\underline{\mathbf{X}}_B \rightarrow \mathbf{Y}) \geq \left(1 - \exp\left(\frac{-L}{\sum_{i=0}^{K-1} \alpha^i}\right)\right) I(\underline{\mathbf{X}}_A \rightarrow \mathbf{Y}).$$

The proof is in Appendix B.

Recall from Theorem 5 that the larger the sum of directed information values from parent sets to children is, the better the approximation is. Theorem 13 implies that greedy approximations are near-optimal approximations.

Figure 2(a) shows the bound coefficient in Theorem 13 for $\alpha \in \{1.3, 1.7, 2.5\}$. In Example 1, if the variances were equal, $\alpha = 1.71$ would suffice.

We can also bound how close an optimal parent set A_L with in-degree L is to an optimal parent set A_K with in-degree $K > L$.

Corollary 14 *Under Assumption 1, with $\alpha \neq 1$,*

$$I(\underline{\mathbf{X}}_{A_L} \rightarrow \mathbf{Y}) \geq \left(\frac{\alpha^L - 1}{\alpha^K - 1}\right) I(\underline{\mathbf{X}}_{A_K} \rightarrow \mathbf{Y}).$$

The proof is in Appendix C. The bound coefficient is plotted in Figure 2(b).

5.2 Near-Optimal Solutions for the Unconstrained Problem

We next consider Algorithm 3 which uses a greedy search to find a near-optimal solution to (8) where the only constraints are in-degree bounds, similar to Algorithm 1. Since the greedy search is adaptive, directed information values will be computed as needed. Let $\{B(i)\}_{i=1}^m$ and $\{A(i)\}_{i=1}^m$ denote the parent sets returned by Algorithms 3 and 1 respectively.

Theorem 15 *Under Assumption 1,*

$$\sum_{i=1}^m I(\underline{\mathbf{X}}_{B(i)} \rightarrow \mathbf{X}_i) \geq \left(1 - \exp\left(\frac{-L}{\sum_{i=0}^{K-1} \alpha^i}\right)\right) \sum_{i=1}^m I(\underline{\mathbf{X}}_{A(i)} \rightarrow \mathbf{X}_i).$$

Algorithm 3. NEAR-OPTIMALGENERAL**Input:** L, m

1. **For** $i \in \{1, \dots, m\}$
2. $B(i) \leftarrow \emptyset$
3. **While** $|B(i)| < L$
4. **For** $l \in \{1, \dots, m\} \setminus \{B(i) \cup \{i\}\}$
5. **Compute** $I(\underline{\mathbf{X}}_l \rightarrow \mathbf{X}_i \parallel \underline{\mathbf{X}}_{B(i)})$
6. $B(i) \leftarrow B(i) \cup \arg \max_l I(\underline{\mathbf{X}}_l \rightarrow \mathbf{X}_i \parallel \underline{\mathbf{X}}_{B(i)})$
7. **Return** $\{B(i)\}_{i=1}^m$

Proof The proof follows from Theorem 13 holding for each $i \in \{1, \dots, m\}$. \blacksquare

Remark 16 The edge weight $I(\underline{\mathbf{X}}_{B(i)} \rightarrow \mathbf{X}_i)$ can be computed using a chain rule (Kramer, 1998). Let $\{j_1, j_2, \dots, j_L\}$ denote $B(i)$. Then

$$I(\underline{\mathbf{X}}_{B(i)} \rightarrow \mathbf{X}_i) = \sum_{l=1}^L I(\mathbf{X}_{j_l} \rightarrow \mathbf{X}_i \parallel \mathbf{X}_{j_1}, \dots, \mathbf{X}_{j_{l-1}}).$$

5.3 Near-Optimal Solutions for Finding Connected Graphs

We now propose Algorithm 4 which uses a greedy search to find a near-optimal connected solution to (8). Similar to Algorithm 3, it precomputes parent sets for each possible directed edge. Then a MWDST algorithm is called. In Algorithm 4, $\tilde{B}(i, j)$ is the set of parents for \mathbf{X}_i with \mathbf{X}_j as one of the parents, selected in a greedy fashion. The value $I(\underline{\mathbf{X}}_{\tilde{B}(i, j)} \rightarrow \mathbf{X}_i)$ is the weight of edge $\mathbf{X}_j \rightarrow \mathbf{X}_i$ given to the MWDST algorithm. Let $\{A(i)\}_{i=1}^m$ denote the parent sets returned by Algorithm 2.

Theorem 17 Under Assumption 1, for Algorithm 4,

$$\sum_{i=1}^m I(\underline{\mathbf{X}}_{B(i)} \rightarrow \mathbf{X}_i) \geq \left(1 - \exp\left(\frac{-L}{\sum_{i=0}^{K-1} \alpha^i}\right)\right) \sum_{i=1}^m I(\underline{\mathbf{X}}_{A(i)} \rightarrow \mathbf{X}_i).$$

The proof is in Appendix D.

6. Best r Bounded In-Degree Approximations

For a given $P_{\underline{\mathbf{X}}}$, Algorithms 1–4 each return a single solution. For many applications, knowing the r -best approximations, where r might be five, twenty, or a hundred, can be advantageous. For instance, there is no guarantee of uniqueness for optimal approximations. Additionally, when data is limited or noisy, the actual optimal solution might appear as second or third best due to estimation errors. Also, for more complex constraints, such as directed trees with depth at most five, it might be easier to find the r -best solutions to the more general problem of directed tree approximations, and then pick the highest ranking

Algorithm 4. NEAR-OPTIMALCONNECTED**Input:** L, m

1. **For** $i \in \{1, \dots, m\}$
2. $B(i) \leftarrow \emptyset$
3. **For** $j \in \{1, \dots, m\} \setminus \{i\}$
4. $\tilde{B}(i, j) \leftarrow \{j\}$
5. **While** $|\tilde{B}(i, j)| < L$
6. **For** $l \in \{1, \dots, m\} \setminus \{\tilde{B}(i, j) \cup \{i\}\}$
7. **Compute** $I(\mathbf{X}_l \rightarrow \mathbf{X}_i \parallel \underline{\mathbf{X}}_{\tilde{B}(i, j)})$
8. $\tilde{B}(i, j) \leftarrow \tilde{B}(i, j) \cup \arg \max_l I(\mathbf{X}_l \rightarrow \mathbf{X}_i \parallel \underline{\mathbf{X}}_{\tilde{B}(i, j)})$
9. $\{b(i)\}_{i=1}^m \leftarrow \text{MWDST } (\{I(\underline{\mathbf{X}}_{\tilde{B}(i, j)} \rightarrow \mathbf{X}_i)\}_{1 \leq i \neq j \leq m})$
10. **For** $i \in \{1, \dots, m\}$
11. $B(i) \leftarrow \tilde{B}(i, b(i))$
12. **Return** $\{B(i)\}_{i=1}^m$

solution that satisfies the extra constraint. Lastly, edges that persist in all of the r -best approximations are likely to be important.

We next discuss methods to identify the r -best bounded in-degree approximations. For simplicity, we will focus on altering Algorithm 1 and then discuss differences for modifying the other algorithms. A strategy for identifying the r -best solutions for assignment problems is discussed in Lawler (1972) based on branching candidate solutions. The method would be impractical to apply here. We develop an alternative algorithm, which, like that in Lawler (1972), will be based on a branching of candidate solutions.

6.1 Optimal Solutions for the Unconstrained Problem

Recall from Theorem 5 that the sum of directed information values from parent sets to children corresponds to how good an approximation is. When the only constraint is a user specified in-degree K , the parent sets can be chosen independently, as in Algorithm 1. This property simplifies searching for the r -best approximations. For instance, the second best approximation will only differ from the first by one parent set.

Algorithm 5 identifies the r -best bounded in-degree approximations in order. It maintains a list of candidate approximations. It instantiates that list by calling Algorithm 1. Each time an approximation is selected from the list, Algorithm 6 generates new candidate approximations from that “seed.” Algorithm 6 finds m solutions by keeping all but one parent set, replacing it with the next best one.

To simplify the presentation, we will make the following assumption to avoid ties.

Assumption 2 *For a given joint distribution $P_{\underline{\mathbf{X}}}$, for any process \mathbf{X}_i , no two disjoint parent sets $A(i), B(i) \subseteq [m] \setminus \{i\}$, with $|A(i)| = |B(i)|$ have identical directed information values, $I(\underline{\mathbf{X}}_{A(i)} \rightarrow \mathbf{X}_i) \neq I(\underline{\mathbf{X}}_{B(i)} \rightarrow \mathbf{X}_i)$.*

For cases where Assumption 2 does not hold, line 4 in Algorithm 6 can be modified to check not only values, but elements of parent sets as well.

Algorithm 5. TOPRGENERAL**Input:** $\mathcal{DI}_{\text{BndInd}}, K, m, r$

1. $Top \leftarrow \emptyset$
2. $l \leftarrow 0$
3. $\mathcal{S} \leftarrow \text{OPTIMALGENERAL}(\mathcal{DI}_{\text{BndInd}}, K, m)$
4. **While** $l < r$
5. $l \leftarrow l + 1$
6. $Top(l) \leftarrow \arg \max_{\hat{P}_{\mathbf{X}} \in \mathcal{S}} \sum_{i=1}^m I(\underline{\mathbf{X}}_{A(i)} \rightarrow \mathbf{X}_i)$
7. $\mathcal{S} \leftarrow \mathcal{S} \cup \text{GETNEWSOLN}(\mathcal{DI}_{\text{BndInd}}, K, m, Top(l))$
8. $\mathcal{S} \leftarrow \mathcal{S} \setminus Top$
9. **Return** Top

Algorithm 6. GETNEWSOLNS**Input:** $\mathcal{DI}_{\text{BndInd}}, K, m, \{A'(i)\}_{i=1}^m$

1. $\mathcal{S} \leftarrow \emptyset$
2. **For** $i \in \{1, \dots, m\}$
3. $\{A(i)\}_{i=1}^m \leftarrow \{A'(i)\}_{i=1}^m$
4. $\mathcal{B} \leftarrow \{B : B \subseteq [m] \setminus \{i\}, |B| = K, I(\underline{\mathbf{X}}_B \rightarrow \mathbf{X}_i) < I(\underline{\mathbf{X}}_{A'(i)} \rightarrow \mathbf{X}_i)\}$
5. $A(i) \leftarrow \arg \max_{B \in \mathcal{B}} I(\underline{\mathbf{X}}_B \rightarrow \mathbf{X}_i)$
6. $\mathcal{S} \leftarrow \mathcal{S} \cup \{A(i)\}_{i=1}^m$
7. **Return** \mathcal{S}

Theorem 18 *Under Assumption 2, Algorithm 5 returns the r -best bounded in-degree approximations.*

The proof is in Appendix E. We also provide a discussion for how to index approximations for efficient implementation of lines 6–8 of Algorithm 5 in Appendix F.

6.2 Optimal Solutions to Find Connected Graphs

Identifying the r -best *connected* approximations is more complicated than the general case in Section 6.1. Algorithm 5 would not need to be modified, but Algorithm 6 would. Recall from Algorithm 2 line 5 that each edge $\mathbf{X}_j \rightarrow \mathbf{X}_i$ in the complete graph is assigned an edge weight $I(\underline{\mathbf{X}}_{\tilde{A}(i,j)} \rightarrow \mathbf{X}_i)$. If the edge $\mathbf{X}_j \rightarrow \mathbf{X}_i$ is selected by the MWDST algorithm in line 6, then $\tilde{A}(i,j)$ is the parent set assigned to \mathbf{X}_i . To find the r -best connected approximations, as in Algorithm 6, “seed” approximations should be modified to generate new candidate approximations. However, Algorithm 6 will not work properly. Some candidate approximations might be identical to the seed.

To see this, let $\{j_1, j_2, \dots, j_K\}$ denote $\tilde{A}(i, j_1)$. Suppose $\tilde{A}(i, j_1) = \tilde{A}(i, j_2)$ and $\mathbf{X}_{j_1} \rightarrow \mathbf{X}_i$ was an edge in the MWDST that induced the seed approximation. Thus, $\tilde{A}(i, j_1)$ is the parent set selected for \mathbf{X}_i . In generating new approximations, even if $\mathbf{X}_{j_1} \rightarrow \mathbf{X}_i$ is given

a smaller weight, $I(\underline{\mathbf{X}}_{\tilde{A}'(i,j_1)} \rightarrow \mathbf{X}_i)$, edge $\mathbf{X}_{j_2} \rightarrow \mathbf{X}_i$ might be selected by the MWDST algorithm instead, yielding the same parent set.

One approach to generate candidate approximations involves checking whether modifying a single edge weight in the complete graph, such as setting $\mathbf{X}_{j_1} \rightarrow \mathbf{X}_i$ to have weight $I(\underline{\mathbf{X}}_{\tilde{A}'(i,j)} \rightarrow \mathbf{X}_i)$, does result in a candidate approximation different than the seed. If not, then all subsets of edges inducing the same parent set $\tilde{A}(i, j_1)$ should be modified. Thus if $\tilde{A}(i, j_1) = \tilde{A}(i, j_2) = \tilde{A}(i, j_3)$, then the weights of edges $\{\mathbf{X}_{j_1} \rightarrow \mathbf{X}_i\}$, $\{\mathbf{X}_{j_1} \rightarrow \mathbf{X}_i, \mathbf{X}_{j_2} \rightarrow \mathbf{X}_i\}$, $\{\mathbf{X}_{j_1} \rightarrow \mathbf{X}_i, \mathbf{X}_{j_3} \rightarrow \mathbf{X}_i\}$, and $\{\mathbf{X}_{j_1} \rightarrow \mathbf{X}_i, \mathbf{X}_{j_2} \rightarrow \mathbf{X}_i, \mathbf{X}_{j_3} \rightarrow \mathbf{X}_i\}$ should be modified. Any resulting candidate parent sets that differ from the seed should be retained.

6.3 Near-Optimal Solutions for the Unconstrained Problem

Algorithm 5 generates the r -best approximations, but calls Algorithms 6 which uses an exhaustive search. A greedy search can be used instead to generate r approximations. Consider the first time that Algorithm 6 is called. Let $\{j_1, j_2, \dots, j_K\}$ denote the parent set, in order they were added, for node i . When i 's parent set is changed, in line 5 set

$$\mathcal{B} \leftarrow \{B : \{j_1, \dots, j_{K-1}\} \subseteq B \subseteq [m] \setminus \{i, j_K\}, |B| = K\}.$$

So only the parent added last in the greedy search is changed.

Consider the first time that Algorithm 6 is called where i 's parents in the seed approximation is $\{j_1, \dots, j_{K-1}, j'_K\}$ for some $j'_K \neq j_K$. Then set

$$\mathcal{B} \leftarrow \{B : \{j_1, \dots, j_{K-1}\} \subseteq B \subseteq [m] \setminus \{i, j_K, j'_K\}, |B| = K\}.$$

This can be repeated $m - K - 2$ more times until $|\mathcal{B}| = \emptyset$. When this occurs, the $(K - 1)$ th parent needs to be changed. Set

$$\mathcal{B} \leftarrow \{B : \{j_1, \dots, j_{K-2}, j'_{K-1}\} \subseteq B \subseteq [m] \setminus \{i, j_{K-1}\}, |B| = K\}.$$

where

$$j'_{K-1} = \arg \max_{j \in [m] \setminus \{i, j_{K-1}\}} I(\mathbf{X}_j \rightarrow \mathbf{X}_i \| \underline{\mathbf{X}}_{j_1, \dots, j_{K-2}}),$$

the next $(K - 1)$ th parent selected in a greedy order. Continue in this manner until Algorithm 5 selects r approximations.

Note also that we can combine the modifications discussed here with those in Section 6.2 to identify the top r connected approximations using a greedy search.

7. Complexity of Proposed Algorithms

This section explores the computational complexity of the algorithms and storage complexity of the approximations.

First, calculating $I(\mathbf{X}, \mathbf{Z} \rightarrow \mathbf{Y})$ in general has exponential complexity. Note that

$$I(\mathbf{X}, \mathbf{Z} \rightarrow \mathbf{Y}) = \sum_{t=1}^n I(Y_t; X^{t-1}, Z^{t-1} | Y^{t-1}).$$

The last term in particular, $I(Y_n; X^{n-1}, Z^{n-1} | Y^{n-1})$, involves a sum over all realizations of $3n - 2$ random variables. Thus, with \mathbf{X} denoting the alphabet, the last term has complexity $\mathcal{O}(|\mathbf{X}|^{3n})$. We will assume Markovity of a fixed order l , so

$$I(\mathbf{X}, \mathbf{Z} \rightarrow \mathbf{Y}) = \sum_{t=1}^n I(Y_t; X_{t-l}^{t-1}, Z_{t-l}^{t-1} | Y_{t-l}^{t-1}).$$

The complexity of computing $I(\mathbf{X}, \mathbf{Z} \rightarrow \mathbf{Y})$ then becomes $\mathcal{O}(n|\mathbf{X}|^{3l+1}) = \mathcal{O}(n)$. More generally, computing $I(\underline{\mathbf{X}}_B \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{B'})$, where $|B| + |B'| = K$, has $\mathcal{O}(n|\mathbf{X}|^{(K+1)l+1})$ complexity assuming Markovity.

Assumption 3 *We assume Markovity of order l .*

7.1 Algorithm 1. OPTIMALGENERAL

For each process \mathbf{X}_i , for each of the $\binom{m-1}{K}$ possible subsets B with $|B| = K$, $I(\underline{\mathbf{X}}_B \rightarrow \mathbf{X}_i)$ is computed. Each computation has complexity $\mathcal{O}(n|\mathbf{X}|^{(K+1)l+1})$. Thus, the total complexity for Algorithm 1 under Assumption 3 is $\mathcal{O}(m \binom{m-1}{K} n |\mathbf{X}|^{(K+1)l+1})$, or $\mathcal{O}(m^{K+1} n)$ for fixed K .

7.2 Algorithm 2. OPTIMALCONNECTED

Algorithm 2 computes the same directed information terms as Algorithm 1. It also computes a MWDST, which takes $\mathcal{O}(m^2)$ time (Edmonds, 1967). Under Assumption 3, the total complexity is $\mathcal{O}(m \binom{m-1}{K} n |\mathbf{X}|^{(K+1)l+1} + m^2)$. If K is fixed, the complexity becomes $\mathcal{O}(m^{K+1} n)$.

7.3 Algorithm 3. NEAR-OPTIMALGENERAL

For each process \mathbf{X}_i there are $(m - 1)$ directed information terms computed involving two processes, of the form $I(\mathbf{X}_j \rightarrow \mathbf{X}_i)$. Next there are $(m - 2)$ computed involving three processes, and so on. The complexity is thus

$$\mathcal{O}(m \sum_{i=1}^K (m - 1)n |\mathbf{X}|^{(i+1)l+1}) = \mathcal{O}(m^2 K n |\mathbf{X}|^{(K+1)l+1}).$$

For constant K , this becomes $\mathcal{O}(m^2 n)$.

7.4 Algorithm 4. NEAR-OPTIMALCONNECTED

For each ordered pair of processes $(\mathbf{X}_i, \mathbf{X}_j)$, first there are $(m - 2)$ terms computed involving three processes, such as $I(\mathbf{X}_k \rightarrow \mathbf{X}_i \parallel \mathbf{X}_j)$. Next there are $(m - 3)$ computed involving four processes, and so on. Then a MWDST algorithm is called. The complexity is thus

$$\mathcal{O}(m^2 + m(m - 1) \sum_{i=1}^{K-1} (m - 1 - i)n |\mathbf{X}|^{(i+2)l+1}) = \mathcal{O}(K m^3 n |\mathbf{X}|^{(K+1)l+1}).$$

For constant K , this becomes $\mathcal{O}(m^3 n)$.

7.5 Algorithms 5. TOPRGENERAL

There are three main bottlenecks in generating the top- r solutions. The first is computing the directed information terms, the same as used in Algorithms 1 and 2, $\mathcal{O}(m^{K+1}n)$ for fixed K . The second is sorting those values for each process \mathbf{X}_i . Merge sort, for example, can sort an array of h elements in $\mathcal{O}(h \log h)$ time (Katajainen et al., 1996). For each of the m processes, there are $\binom{m-1}{K}$ values, so sorting takes $\mathcal{O}(m \binom{m-1}{K} \log \binom{m-1}{K}) = \mathcal{O}(m^{K+1} \log m)$ for fixed K . The third bottleneck is the search for candidate solutions. Algorithm 6 generates m new solutions each time it is called, replacing one parent set for each approximation. The branching overall generates $\mathcal{O}(rm)$ candidates. For small r , such as $r = \mathcal{O}(\log \binom{m-1}{K}^m) = \mathcal{O}(m \log m^K) = \mathcal{O}(m \log m)$ for fixed K , the computing and sorting the directed information values dominates. Recall that $\binom{m-1}{K}^m$ is the total number of bounded in-degree approximations. For large r , such as $r = \binom{m-1}{K}^m/c$ for some $c > 1$, then $r = \mathcal{O}(m^{Km})$ and so the branching dominates. The total complexity with fixed K is $\mathcal{O}(m^{K+1}(n + \log m) + rm)$.

7.6 Storage Complexity

An important benefit of using approximations is that they require substantially less storage than the full joint distribution. The full joint distribution has mn random variables, and so requires $\mathcal{O}(|\mathbf{X}|^{mn})$ storage. Under Assumption 3, the storage complexity of the full joint distribution is $\mathcal{O}(n|\mathbf{X}|^{m(l+1)})$ and of approximations of the form (7) is $\mathcal{O}(mn|\mathbf{X}|^{(K+1)l+1}) = \mathcal{O}(mn)$ for fixed K .

8. Simulations

We investigated the performances of Algorithms 1-5 using simulated networks. Comparisons of greedy and optimal search, unconstrained and connected approximations, and the top- r approximations were studied.

8.1 Greedy vs. Optimal Search

We first compare the near-optimal and optimal approximations identified by Algorithm 3 and Algorithm 1 respectively.

8.1.1 SETUP

Markov order-1 autoregressive (AR) networks, of the form $\underline{\mathbf{X}}_t = C\underline{\mathbf{X}}_{t-1} + N_t$, were simulated for a given m by m coefficient matrix C and i.i.d. noise vector N_t . Two network sizes $m \in \{6, 15\}$ were tested. For each m , there were 250 trials. In each trial, the coefficient matrix C was randomly generated. Edges (non-zero off-diagonal entries in C) were selected i.i.d. with probability 1/2. Non-zero AR coefficients were drawn i.i.d. from a standard normal distribution. C was then scaled to be stationary. The noise process $\{N_t\}_{t=1}^n$ had i.i.d. entries drawn from a normal distribution with mean zero and variance 1/4. Data was generated for $n = 1000$ time-steps.

For each network, unconstrained bounded in-degree approximations $\{A_{\text{OPT}}(i)\}_{i=1}^m$ and $\{A_{\text{GRD}}(i)\}_{i=1}^m$ were computed using Algorithms 1 and 3 respectively. For $m = 6$ and $m = 15$,

in-degrees $K = 2$ and $K = 4$ were used respectively. Performance was measured by the ratio

$$\frac{\sum_{i=1}^m I(\underline{\mathbf{X}}_{A_{\text{GRD}}(i)} \rightarrow \mathbf{X}_i)}{\sum_{i=1}^m I(\underline{\mathbf{X}}_{A_{\text{OPT}}(i)} \rightarrow \mathbf{X}_i)}. \quad (14)$$

The value of each sum corresponds to how good that approximation is. For (14), the directed information values were calculated exactly using approximated parent sets.

Both algorithms computed directed information estimates using the simulated data. The estimate for a directed information of the form $I(\mathbf{X} \rightarrow \mathbf{Y} \parallel \mathbf{Z})$ was computed as follows. Least square estimates for the coefficients in two AR models,

$$\begin{aligned} Y_t &= b_1 Y_{t-1} + b_2 Z_{t-1} + b_3 X_{t-1} + N_t \\ Y_t &= b'_1 Y_{t-1} + b'_2 Z_{t-1} + N'_t, \end{aligned}$$

were computed. Let σ and σ' denote $\text{std}(N_t)$ and $\text{std}(N'_t)$ respectively. The entropy $H(Y_t|Y_{t-1}, Z_{t-1}, X_{t-1})$ is $1/2 \log_2(2\pi e \sigma^2)$ (Cover and Thomas, 2006, Theorem 8.4.1), so

$$\begin{aligned} \widehat{I}(\mathbf{X} \rightarrow \mathbf{Y} \parallel \mathbf{Z}) &= \frac{1}{n} \sum_{t=1}^n \widehat{H}(Y_t|Y_{t-1}, Z_{t-1}) - \widehat{H}(Y_t|Y_{t-1}, Z_{t-1}, X_{t-1}) \\ &= \frac{1}{2} \log(2\pi e (\sigma')^2) - \frac{1}{2} \log(2\pi e \sigma^2) \\ &= \log \sigma'/\sigma. \end{aligned}$$

8.1.2 RESULTS

The approximations found by the greedy and optimal search were largely identical. Figure 3 shows histograms of percentages of the ratio (14), normalized by the number of trials. The rightmost column in each histogram corresponds to (14) being one, when the greedy search returned the optimal approximation.

For $m = 6$, the greedy search found the optimal solution in 96.4% of the trials. The average ratio was $99.9\% \pm 0.6\%$. The minimum ratio was 92.0%.

For $m = 15$, the greedy search found the optimal solution in 57.6% of the trials. The average ratio was $99.6\% \pm 0.9\%$. The minimum ratio was 93.3%.

By Theorem 13, the best case lower bound of (14) expected is 63.2%, which corresponds to $\alpha = 1$. On average, the greedy algorithm performed much better than the lower bound, often the same as or close to the optimal.

8.2 Comparison of Top- r Approximations for Algorithms 1,2,3,4

We next investigated how well the approximations from Algorithms 1, 2, 3, and 4 compared to the true parent sets. We used modified versions of the algorithms to produce the top- r approximations for each class, as discussed in Section 6 (such as Algorithm 5). For the MWDST algorithm we used Choudhary (2009).

8.2.1 SETUP

The setup was similar to that described in Section 8.1.1. For the approximations, there were multiple in-degree K values. For $m = 6$, in-degrees $K \in \{1, 2, 4\}$ were used, and for

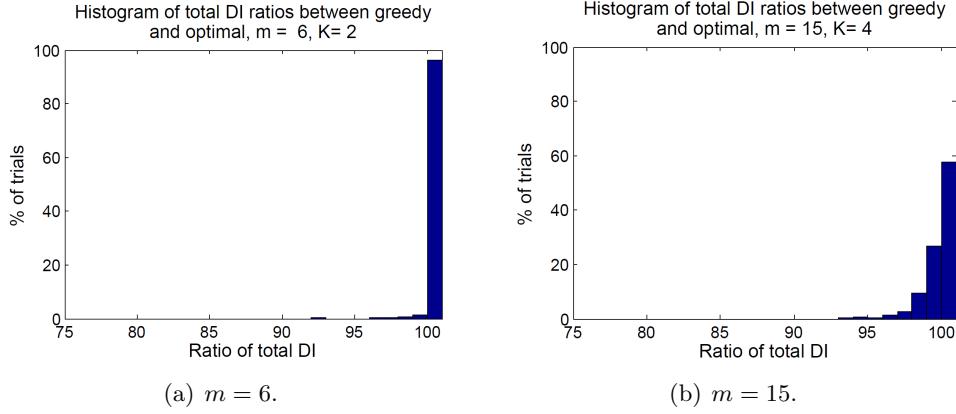


Figure 3: Histograms of the relative performance of Algorithms 1 and 3 using the ratio (14). The right column corresponds to when both algorithms identified the same approximation.

$m = 15, K \in \{2, 4, 8\}$ were used. Performance was measured by the ratio

$$\frac{\sum_{i=1}^m I(\underline{\mathbf{X}}_{A(i)} \rightarrow \mathbf{X}_i)}{\sum_{i=1}^m I(\underline{\mathbf{X}}_{A_{\text{True}}(i)} \rightarrow \mathbf{X}_i)}, \quad (15)$$

where $A(i)$ denotes a parent set induced by an approximation, and $A_{\text{True}}(i)$ is the true parent set. The ratio (15) was calculated exactly. For each type of approximation, the top $r = 10$ approximations were found and performance was averaged across trials.

8.2.2 RESULTS

Overall, the different approximations (unconstrained or connected, using optimal or greedy search) performed comparably for each of the (m, K) pairs. The results are shown in Figure 4. The most noticeable variation in performance was due to different in-degree K values. For both $m = 6$ and $m = 15$, increasing K , especially when K was small, substantially improved performance.

The connected approximations performed only slightly worse than the unconstrained. For larger m and K the difference diminished. For each (m, K) pair, on average the approximations returned by the greedy search performed almost as well as that of the optimal. This is consistent with the results in Section 8.1.2 (see Figure 3).

Performance did not decay appreciably among the top approximations. For $m = 6$ and $K = 1$, the performance difference between the first and tenth approximation is distinguishable, but for others it is not.

Figure 5 shows diagrams of the top $r = 4$ unconstrained and connected approximations for a single trial with $m = 6$ and $K = 2$. Figures 5(a) and 5(e) are the optimal unconstrained and connected approximations respectively. The figures to the right show differences between the optimal and r th best approximations for $r \in \{2, 3, 4\}$. Dashed gray edges are those removed and solid black ones are edges included. For example, the third best unconstrained approximation, Figure 5(c), has all of the same edges as the optimal, Figure 5(a), except it has $\mathbf{X}_2 \rightarrow \mathbf{X}_4$ instead of $\mathbf{X}_1 \rightarrow \mathbf{X}_4$.

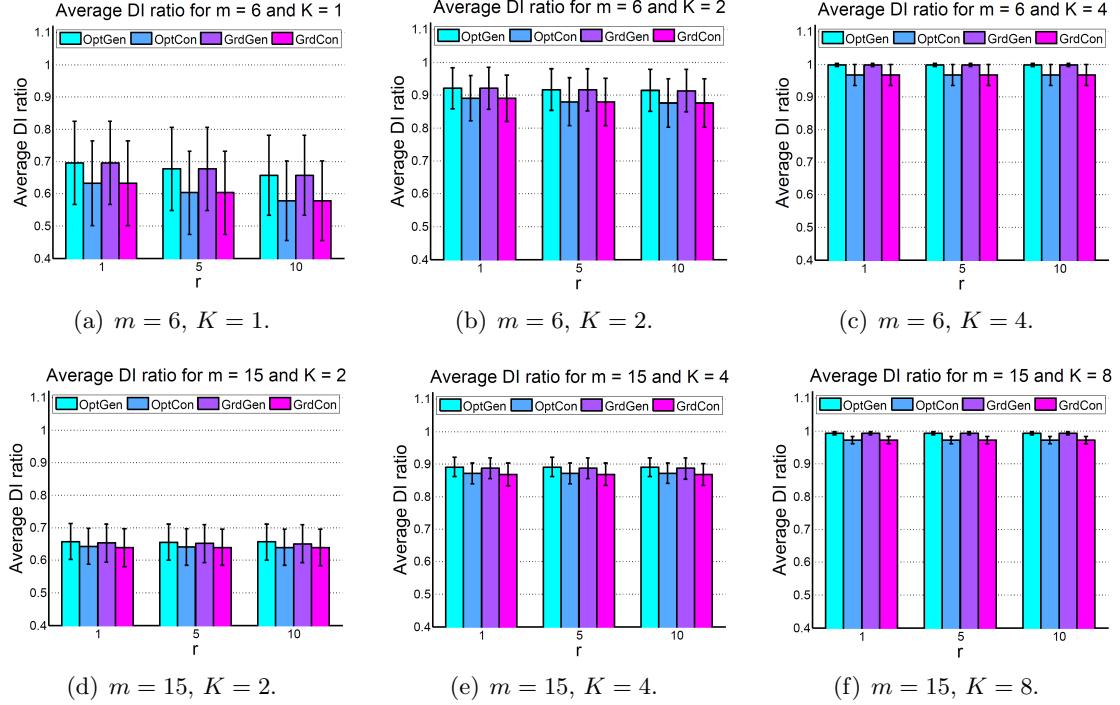


Figure 4: Plots of the average ratio (15) of approximations to the true parent sets. Standard deviation error bars are shown. For each type of search, optimal (“Opt”) and greedy (“Grd”), and each type of approximation, unconstrained (“Gen”) and connected (“Con”), the top- r approximations are shown, with $r \in \{1, 5, 10\}$.

For both the optimal unconstrained and connected approximations, the several next best approximations had only minor changes. Many edges were preserved. For this particular trial, among the top four unconstrained approximations the only differences involved the parent of \mathbf{X}_4 . For the connected approximations, the second, third, and fourth best approximations mostly varied for the parent of \mathbf{X}_5 . The parents of \mathbf{X}_1 , \mathbf{X}_2 , and \mathbf{X}_3 were identical for all of the approximations shown.

8.3 Performance Decay for Large r

We also investigated the decay in performance quality as $r \rightarrow \binom{m-1}{K}^m$.

8.3.1 SETUP

Using a simulation setup similar to Section 8.1.1, 150 trials with $m = 6$ processes and in-degree $K = 2$ were run. Unconstrained bounded in-degree approximations were obtained using Algorithm 5 and a similarly modified Algorithm 3 to obtain optimal and greedy search orderings respectively. All approximations were computed, with $r = \binom{m-1}{K}^m = 10^6$. The ratio (15) was used to measure performance.

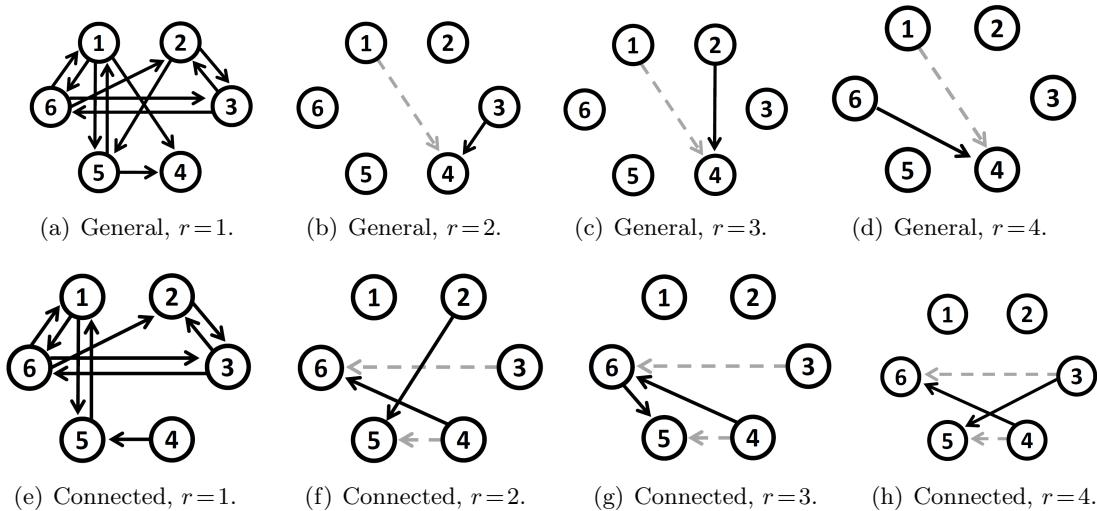


Figure 5: Diagrams of the top $r = 4$ unconstrained and connected approximations using optimal search for a single trial with $m = 6$ and $K = 2$. Figures 5(a) and 5(e) are the optimal unconstrained and connected approximations respectively. The figures to the right show differences between the optimal and r th best approximations. Dashed edges were removed and solid edges were included.

8.3.2 RESULTS

Figure 6(a) depicts the ratio (15) for all approximations, ordered using Algorithm 5, for three trials. Due to estimation error, the actual ratio values did not decay monotonically. Figure 6(b) plots the estimated values.

Figure 7 shows the true ratio (15) value using the ordering returned by Algorithm 5, averaged over 150 trials. The mean and one standard deviation are represented by the black and green curves respectively. Consistent with the curves for individual trials in Figure 6, the shape of the curves in Figure 7 are similar to a logit function. There was a sharp decrease in the quality of approximation in the low and high r regimes, with a nearly linear decay for most r .

The greedy search, using a modified Algorithm 3, performed comparably to the optimal search, Algorithm 5. For $r = 1$, this was shown in Figure 3, and for $r \leq 10$ this was shown in Figure 4. However, the current analysis confirmed that even for large r this was true. The analogous Figure 7 for the greedy ordering was visually indistinguishable and so not shown. However, effects of greedy ordering were clearly seen in some trials. Figure 8 depicts estimate values for approximations in one trial, with the light blue monotonic curve for the optimal ordering and the discontinuous black curve depicting the greedy ordering. The large jumps are characteristic of the depth-first search, as the worst-case along one branch is worse than the best-case of the next, although the former branch was initially more promising. However, such large discontinuities as shown in Figure 8 were rare overall. The ordering returned by the greedy search was largely similar to the optimal ordering.

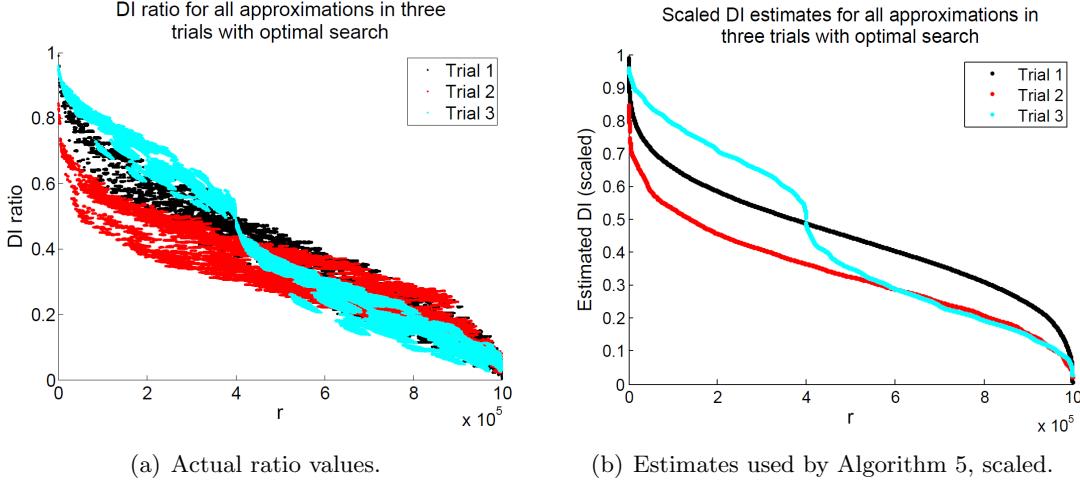


Figure 6: The values of the ratio (15) for all approximations in the order selected by Algorithm 5. Three trials are shown. Figure 6(a) shows the actual ratio values. The spread is due to estimation error. Figure 6(b) shows the estimated values (normalized to the scale of Figure 6(a)).

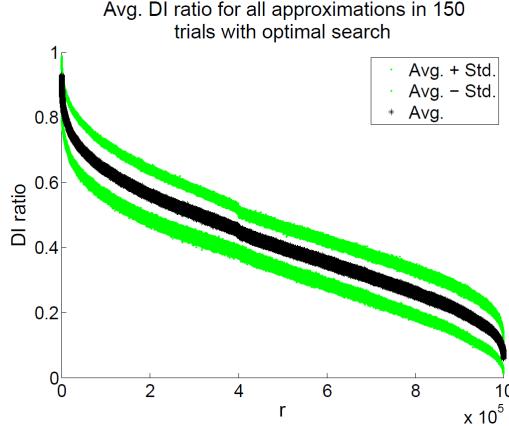


Figure 7: The ratio (15) for all approximations in the order selected by Algorithm 5, averaged over 150 trials. The black and green curves depict the mean standard deviation respectively.

As illustrated in Figure 6, estimation errors led to errors in the optimal search ordering returned by Algorithm 5. It is useful to characterize how large those ordering errors were. Figure 9 shows the error, measured by the difference of the ratio (15) between the r th best approximation and the r th approximation in the ordering returned by Algorithm 5 for $r \leq 1000$. The results were averaged over 150 trials. In terms of percentage points of the sum of directed information of the true parent sets to children, the denominator of (15), Algorithm 5 performed well despite the estimation error. The mean difference was slightly biased above 0 and the standard deviations were within 4% for all $r \leq 1000$. The slight

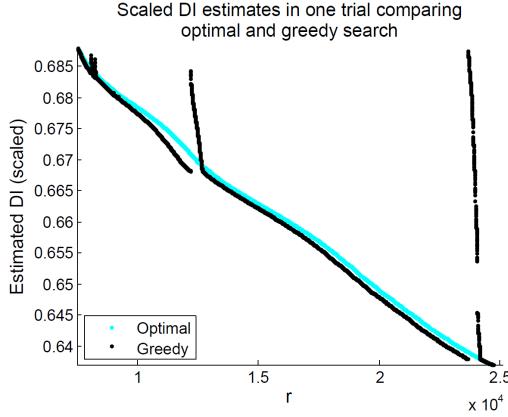


Figure 8: The estimate values used by Algorithm 5 and a modified Algorithm 3 to rank unconstrained approximations using optimal and greedy search. One trial is shown. Estimates are scaled to actual ratio values (15). The smooth light blue curve corresponds to ordering from the optimal search. The black curve corresponds to the greedy search.

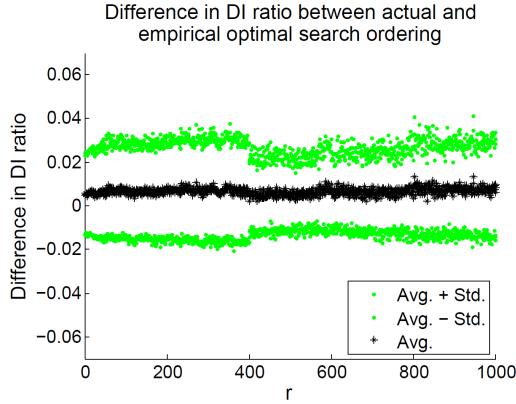


Figure 9: The difference between the ratio (15) for the actual optimal ordering of unconstrained approximations and the ordering returned by Algorithm 3 (using estimates). Results are averaged over 150 trials. The black curve depicts the mean. The green curves depict one standard deviation.

bias in the mean was expected; for small r , the approximations were mostly replaced with worse approximations, if any. Overall, the greedy approximations performed nearly as well as the optimal approximations.

9. Conclusion

In this paper, we presented several novel methods related to approximating directed information graphs. The approximations allowed substantial flexibility for users. Each algorithm took the in-degrees of the nodes as input. With larger in-degrees, the approximations be-

came better, but at the cost of visual simplicity of the graph and computational efficiency. The approximation could be unconstrained or connected, and found using an optimal or a more efficient, near-optimal greedy search. Furthermore, one could generate the top r solutions, not only the best. This enabled evaluation for which edges are most significant as well as finding the best solution of a more constrained class of topologies. Lastly, the empirical results demonstrated the utility of these methods, especially showing that on average the greedy search performed much better than the worst-case lower bound.

Appendix A. Proof of Theorem 8

Proof Let \mathcal{T} be the set of all directed spanning trees on m nodes. For a given tree $T \in \mathcal{T}$, let $\tilde{\mathcal{G}}_K^T \subseteq \tilde{\mathcal{G}}_K$ denote the set of approximations $\hat{P}_{\underline{\mathbf{X}}} \in \tilde{\mathcal{G}}_K$ that contain T as a directed spanning tree subgraph. Every $\hat{P}_{\underline{\mathbf{X}}} \in \tilde{\mathcal{G}}_K$ contains at least one such $T \in \mathcal{T}$ as a subgraph, so $\tilde{\mathcal{G}}_K = \bigcup_{T \in \mathcal{T}} \tilde{\mathcal{G}}_K^T$.

For any tree $T \in \mathcal{T}$, the best approximation $\hat{P}_{\underline{\mathbf{X}}}^T \in \tilde{\mathcal{G}}_K^T$ is the one that for every edge $\mathbf{X}_{a(i)} \rightarrow \mathbf{X}_i$ in T , sets $\tilde{A}(i, a(i))$ as the parent set for node i . This follows from (9), since $a(i) \in \tilde{A}(i, a(i))$, so T will be a subgraph of this approximation, and the sets $\tilde{A}(i, a(i))$ are the best such parent sets. Thus,

$$\max_{\hat{P}_{\underline{\mathbf{X}}} \in \tilde{\mathcal{G}}_K} \sum_{i=1}^m I(\underline{\mathbf{X}}_{A(i)} \rightarrow \mathbf{X}_i) = \max_{T \in \mathcal{T}} \max_{\hat{P}_{\underline{\mathbf{X}}} \in \tilde{\mathcal{G}}_K^T} \sum_{i=1}^m I(\underline{\mathbf{X}}_{A(i)} \rightarrow \mathbf{X}_i) \quad (16)$$

$$= \max_{T \in \mathcal{T}} \sum_{i=1}^m I(\underline{\mathbf{X}}_{\tilde{A}(i, a(i))} \rightarrow \mathbf{X}_i), \quad (17)$$

where (16) follows since $\tilde{\mathcal{G}}_K = \bigcup_{T \in \mathcal{T}} \tilde{\mathcal{G}}_K^T$ and (17) uses that $\hat{P}_{\underline{\mathbf{X}}}^T$ is the best approximation in $\tilde{\mathcal{G}}_K^T$. Algorithm 2 finds the solution to (17) and thus identifies the optimal approximation $\hat{P}_{\underline{\mathbf{X}}}^* \in \tilde{\mathcal{G}}_K$. \blacksquare

Appendix B. Proof of Theorem 13

The proof is based on the proof for a related bound for submodular functions (Nemhauser et al., 1978).

Proof For simplicity, we prove the case $A \cap B = \emptyset$. The other case is almost identical and results in a tighter bound. For that case the greedy algorithm selects each element of B before any element of A . Let $l \leq |B| = L$. Let A_l be the set A but ordered according to how the greedy algorithm would pick elements from A after picking $\{B(1), \dots, B(l)\}$.

We first note two inequalities. For all $l < L$,

$$I(\mathbf{X}_{B(l+1)} \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\{B(1), \dots, B(l)\}}) \geq I(\mathbf{X}_{A_l(1)} \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\{B(1), \dots, B(l)\}}), \quad (18)$$

which holds since the greedy algorithm selects $\mathbf{X}_{B(l+1)}$ after $\{\mathbf{X}_{B(1)}, \dots, \mathbf{X}_{B(l)}\}$, and

$$\alpha^{i-1} I(\mathbf{X}_{A_l(1)} \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\{B(1), \dots, B(l)\}}) \geq I(\mathbf{X}_{A_l(i)} \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\{B(1), \dots, B(l), A_l(1), \dots, A_l(i-1)\}}), \quad (19)$$

which follows from Assumption 1 for the set $A \cup \{B(1), \dots, B(l)\}$.

We now compare an optimal solution A to the first l elements in the greedy solution B .

$$\begin{aligned}
 & I(\underline{\mathbf{X}}_A \rightarrow \mathbf{Y}) - I(\underline{\mathbf{X}}_{\{B(1), \dots, B(l)\}} \rightarrow \mathbf{Y}) \\
 & \leq I(\underline{\mathbf{X}}_{A \cup \{B(1), \dots, B(l)\}} \rightarrow \mathbf{Y}) - I(\underline{\mathbf{X}}_{\{B(1), \dots, B(l)\}} \rightarrow \mathbf{Y}) \\
 & = I(\underline{\mathbf{X}}_{\{B(1), \dots, B(l)\}} \rightarrow \mathbf{Y}) \\
 & \quad + \sum_{i=1}^K I(\mathbf{X}_{A_l(i)} \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\{B(1), \dots, B(l)\} \cup \{A_l(1), \dots, A_l(i-1)\}}) \\
 & \quad - I(\underline{\mathbf{X}}_{\{B(1), \dots, B(l)\}} \rightarrow \mathbf{Y})
 \end{aligned} \tag{20}$$

$$\leq \sum_{i=1}^K \alpha^{i-1} I(\mathbf{X}_{A_l(1)} \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\{B(1), \dots, B(l)\}}) \tag{21}$$

$$\leq \sum_{i=1}^K \alpha^{i-1} I(\mathbf{X}_{B(l+1)} \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\{B(1), \dots, B(l)\}}). \tag{22}$$

Equation (20) follows from the chain rule applied in the order the greedy algorithm would select from $A \cup \{B(1), \dots, B(l)\}$. Equations (21) and (22) follow from (19) and (18) respectively.

Let $\delta_l := I(\underline{\mathbf{X}}_A \rightarrow \mathbf{Y}) - I(\underline{\mathbf{X}}_{\{B(1), \dots, B(l)\}} \rightarrow \mathbf{Y})$. Then $\delta_l - \delta_{l+1} = I(\mathbf{X}_{B(l+1)} \rightarrow \mathbf{Y} \parallel \underline{\mathbf{X}}_{\{B(1), \dots, B(l)\}})$. Also denote $\beta := \sum_{i=1}^K \alpha^{i-1}$. From (22) we have $\delta_l \leq \beta (\delta_l - \delta_{l+1})$, which implies $\delta_{l+1} \leq \left(1 - \frac{1}{\beta}\right) \delta_l$. Thus

$$\delta_l \leq \left(1 - \frac{1}{\beta}\right)^l \delta_0 \leq e^{-\frac{l}{\beta}} \delta_0.$$

The last step uses the bound $(1 - p) \leq e^{-p}$, which holds for all p . For $0 < p < 1$, both sides are positive so the inequality is conserved if powers are taken. Since $\delta_0 = I(\underline{\mathbf{X}}_A \rightarrow \mathbf{Y}) - I(\emptyset \rightarrow \mathbf{Y}) = I(\underline{\mathbf{X}}_A \rightarrow \mathbf{Y})$, this gives

$$I(\underline{\mathbf{X}}_A \rightarrow \mathbf{Y}) - I(\underline{\mathbf{X}}_{\{B(1), \dots, B(l)\}} \rightarrow \mathbf{Y}) \leq e^{-\frac{l}{\beta}} I(\underline{\mathbf{X}}_A \rightarrow \mathbf{Y}),$$

which after rearranging gives the theorem. ■

Appendix C. Proof of Corollary 14

We will prove Corollary 14 by first solving the following optimization problem

$$\max_{\{b_1, \dots, b_K\}} \sum_{i=1}^K b_i \tag{23}$$

$$\text{s.t.} \quad \sum_{i=1}^L b_i \leq c \tag{24}$$

$$0 \leq b_i \leq \alpha b_{i-1}, \quad i = 2, \dots, m, \tag{25}$$

where K and L are integers such that $K > L$ and $\alpha > 1$ and $c > 0$ are real coefficients. Let $\{b_1^*, \dots, b_K^*\}$ denote a solution to (23).

Lemma 19 *For any optimal solution $\{b_1^*, \dots, b_K^*\}$, (24) holds with equality.*

Proof The proof will follow by contradiction. Suppose $c - \sum_{i=1}^L b_i^* > 0$. Let $\gamma := \frac{1}{L}(c - \sum_{i=1}^L b_i^*)$. Define

$$\tilde{b}_i := \begin{cases} b_i^* + \gamma & \text{if } i \leq L, \\ b_i^* & \text{if } i > L. \end{cases}$$

Note that $\sum_{i=1}^L \tilde{b}_i = \sum_{i=1}^L b_i^* + \gamma = c$ so the first constraint is met. Also, for $i \leq L$, $\tilde{b}_i = b_i^* + \gamma \leq \alpha b_{i-1}^* + \alpha \gamma = b_{i-1}^*$, so the second constraint is met. Thus, $\{\tilde{b}_1, \dots, \tilde{b}_K\}$ is feasible and has a larger sum than the optimal solution, $\sum_{i=1}^K \tilde{b}_i > \sum_{i=1}^K b_i^*$, a contradiction. \blacksquare

Lemma 20 *For any optimal solution $\{b_1^*, \dots, b_K^*\}$, (25) holds with equality.*

Proof The proof will follow by contradiction. Suppose there is an index $i > 1$ for which $b_i^* < \alpha b_{i-1}^*$. If $i > L$, then we can set $b_i^* \leftarrow \alpha b_{i-1}^*$ to increase the objective function, which contradicts optimality. If $i \leq L$, replace b_{i-1}^* and b_i^* with \tilde{b}_{i-1} and \tilde{b}_i , where $\tilde{b}_{i-1} = \frac{b_{i-1}^* + b_i^*}{1+\alpha}$ and $\tilde{b}_i = \frac{\alpha(b_{i-1}^* + b_i^*)}{1+\alpha}$. Note that $b_{i-1}^* + b_i^* = \tilde{b}_{i-1} + \tilde{b}_i$, and the constraints are still satisfied. This exchange necessarily results in $b_{i+1}^* \leq \alpha b_i^* < \alpha \tilde{b}_i$. Thus, the exchange can be repeated for larger i until $i = L + 1$. Then set $\tilde{b}_{L+1} \leftarrow \alpha \tilde{b}_L$ and the objective function is necessarily increased, a contradiction. \blacksquare

We can now find the solution to the optimization problem.

Lemma 21 *The optimal solution to (23) is $\sum_{i=1}^K b_i^* = c \frac{1-\alpha^K}{1-\alpha}$.*

Proof By Lemmas 19 and 20, the constraints (24) and (25) hold with equality. We can first solve for b_1^* ,

$$\begin{aligned} c &= \sum_{i=1}^L b_i^* = \sum_{i=1}^L \alpha^{i-1} b_1^* \\ \implies b_1^* &= \frac{c}{\sum_{i=1}^L \alpha^{i-1}}. \end{aligned}$$

Solving for the value of the objective function,

$$\sum_{i=1}^K b_i^* = \sum_{i=1}^K \alpha^{i-1} b_1^* = \sum_{i=1}^K \alpha^{i-1} \frac{c}{\sum_{i=1}^L \alpha^{i-1}}. \quad (26)$$

Using the geometric series formula

$$\sum_{i=1}^K \alpha^{i-1} = \sum_{i=0}^{K-1} \alpha^i = \frac{1-\alpha^K}{1-\alpha},$$

the equation (26) becomes

$$\sum_{i=1}^K b_i^* = c \frac{\sum_{i=1}^K \alpha^{i-1}}{\sum_{i=1}^L \alpha^{i-1}} = c \frac{1 - \alpha^K}{1 - \alpha^L}.$$

■

We can now prove Corollary 14.

Proof Let the elements of A_K and A_L be ordered according to the greedy order. Consider the worst case, with $I(\underline{\mathbf{X}}_{A_K} \rightarrow \mathbf{Y})$ as large as possible, given

$$I(\underline{\mathbf{X}}_{\{A_K(1), \dots, A_K(L)\}} \rightarrow \mathbf{Y}) \leq I(\underline{\mathbf{X}}_{A_L} \rightarrow \mathbf{Y}). \quad (27)$$

The inequality (27) holds by definition of A_L being the optimal parent set of size L . Greedy-submodularity imposes another constraint. For any $0 < i < K$,

$$I(\mathbf{X}_{A_K(i+1)} \rightarrow \mathbf{Y} \mid \underline{\mathbf{X}}_{\{A_K(1), \dots, A_K(i)\}}) \leq \alpha I(\mathbf{X}_{A_K(i)} \rightarrow \mathbf{Y} \mid \underline{\mathbf{X}}_{\{A_K(1), \dots, A_K(i-1)\}}).$$

Corollary 14 follows from Lemma 21, substituting $I(\underline{\mathbf{X}}_{A_L} \rightarrow \mathbf{Y})$ for c and $I(\mathbf{X}_{A_K(i)} \rightarrow \mathbf{Y} \mid \underline{\mathbf{X}}_{\{A_K(1), \dots, A_K(i-1)\}})$ for b_i . ■

Appendix D. Proof for Theorem 17

Proof Let T_2 denote the MWDST picked by Algorithm 2. For an edge $e \in \{\mathbf{X}_j \rightarrow \mathbf{X}_i : 1 \leq j \neq i \leq m\}$ in the complete graph on m nodes, let $w_2(e)$ denote the weight $I(\underline{\mathbf{X}}_{\tilde{A}(i,j)} \rightarrow \mathbf{X}_i)$ assigned by Algorithm 2. Define T_4 and $w_4(e)$ for Algorithm 4 likewise. Also, let $c := (1 - \exp(-L/(\sum_{i=0}^{K-1} \alpha^i)))$. For each edge e in the complete graph,

$$w_4(e) \geq cw_2(e), \quad (28)$$

which follows from Theorem 13. Furthermore,

$$\sum_{e \in T_4} w_4(e) \geq \sum_{e \in T_2} w_4(e) \quad (29)$$

$$\geq c \sum_{e \in T_2} w_2(e). \quad (30)$$

Equation (29) follows since in Algorithm 4, T_4 was selected as the MWDST, and (30) follows from (28). ■

Appendix E. Proof for Theorem 18

To prove Theorem 18, we first show the following lemma.

Lemma 22 For all $1 \leq l \leq r$, the l th best approximation has the same parent sets except one as one of the top $l - 1$ solutions.

Proof The proof follows by induction. The base case, with $l = 1$, holds trivially as it is the only solution. Assume that the statement of the lemma holds for some $1 \leq l < r$. Consider the $(l + 1)$ th best solution, $\{B(i)\}_{i=1}^m$. Let \mathbf{X}_j be a process for which the parent set $B(j)$ is not the same as that of the optimal solution, $A(j)$.

By Corollary 6, parent sets can be identified independently. Also, by Assumption 2, no two parent sets have the same influence. Thus, the optimal parent set is $A(j)$ is better than $B(j)$. Let $A'(j)$ be any parent set for \mathbf{X}_j that is better than $B(j)$. Then, by Corollary 6, the parent sets $\{B(1), \dots, B(j - 1), A'(j), B(j + 1), \dots, B(m)\}$ induce a better approximation than the $(l + 1)$ th best approximation and therefore must be one of the top l approximations. Since this new approximation differs from the $(l + 1)$ th best approximation in precisely one parent set, the lemma holds. \blacksquare

The proof for Theorem 18 follows from Lemma 22, since every approximation selected in Algorithm 5 is used as a seed in Algorithm 6 to generate all of the best solutions that have precisely one parent set different from that of the seed.

Appendix F. Implementation Notes for Algorithm 5

In Algorithm 5, for large m and r , a naive implementation of lines 6–8 can be computationally expensive. Specifically, redundant solutions can appear in \mathcal{S} , and searching to remove redundancies or entries already in Top might be slow. Instead, \mathcal{S} can be kept as a priority queue of value-key pairs, where the value is the sum of the directed information values (8) for the approximation as in line 6 and the key is the index of an approximation. There are $\binom{m-1}{K}^m$ possible bounded in-degree approximations, and a binary vector can track whether an approximation has been seen or not.

We now discuss a method to compute an index for each approximation. First, indices for individual parent sets will be identified, then combined for an index for the whole approximation. Let $\{j_1, j_2, \dots, j_K\}$ denote the elements of parent set $A(i)$, in ascending order. For $k \in [K]$, set $j_k \leftarrow j_k - 1$ if $j_k > i$. Denote the set of these (possibly) modified values by the length K vector idx . Then run Algorithm 7.

Lines 7–9 in Algorithm 7 count how many parent sets of \mathbf{X}_i are lexicographically ordered before $A(i)$. Line 7 counts how many sets have a first element smaller than $idx(1)$. Lines 8–9 use recursion to count how many sets with the same first element $idx(1)$ appear before idx .

Once the index $a_i \leftarrow \text{GetParSetIndex}(m, idx)$ for each parent set $A(i)$ of an approximation is calculated, the index for the approximation can be computed as

$$1 + \sum_{i=1}^m a_i \binom{m-1}{K}^{i-1}.$$

Algorithm 7. GETPARSETINDEX**Input:** m, idx

1. $cnt \leftarrow 0$
2. $K \leftarrow |idx|$
3. **If** $K = 0$
4. **Return** 0
5. **If** $K = 1$
6. **Return** $idx(1) - 1$
7. $cnt \leftarrow \sum_{l=2}^{idx(1)} \binom{m-l}{K-1}$
8. $idx' \leftarrow \{idx(2) - idx(1), \dots, idx(K) - idx(1)\}$
9. $cnt \leftarrow cnt + \text{GetParSetIndex}(m - idx(1), idx')$
10. **Return** cnt

References

Pierre-Olivier Amblard and Olivier J. J. Michel. On directed information theory and Granger causality graphs. *Journal of Computational Neuroscience*, 30(1):7–16, 2011.

Pierre-Olivier Amblard and Olivier J. J. Michel. The relation between Granger causality and directed information theory: A review. *Entropy*, 15(1):113–143, 2012.

Sumanta Basu, Ali Shojaie, and George Michailidis. Network Granger causality with inherent grouping structure. *Journal of Machine Learning Research*, 16:417–453, 2015.

Dhruv Batra, Payman Yadollahpour, Abner Guzman-Rivera, and Gregory Shakhnarovich. Diverse M-best solutions in Markov random fields. In Andrew Fitzgibbon, Svetlana Lazebnik, Pietro Perona, Yoichi Sato, and Cordelia Schmid, editors, *European Conference on Computer Vision (ECCV) 2012*, volume 7576 of *Lecture Notes in Computer Science*, pages 1–16. Springer Berlin-Heidelberg, 2012.

Andrew Bolstad, Barry Van Veen, and Robert Nowak. Causal network inference via group sparse regularization. *IEEE Transactions on Signal Processing*, 59(6):2628–2641, 2011.

Alexandra M. Carvalho and Arlindo L. Oliveira. Learning Bayesian networks consistent with the optimal branching. In *Sixth International Conference on Machine Learning and Applications (ICMLA)*, pages 369 –374, Dec. 2007.

Volkan Cevher and Andreas Krause. Greedy dictionary selection for sparse representation. *IEEE Journal of Selected Topics in Signal Processing*, 5(5):979–988, 2011.

Rui Chang, Jonathan R. Karr, and Eric E. Schadt. Causal inference in biology networks with integrated belief propagation. In *Pacific Symposium on Biocomputing*, volume 20, pages 359–370. World Scientific, 2014.

Camille Charbonnier, Julien Chiquet, and Christophe Ambroise. Weighted-lasso for structured network inference from time course data. *Statistical Applications in Genetics and Molecular Biology*, 9(1):1–29, 2010.

Zhitang Chen, Kun Zhang, and Laiwan Chan. Causal discovery with scale-mixture model for spatiotemporal variance dependencies. In P. Bartlett, F. C. N. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems (NIPS) 25*, pages 1736–1744. Curran Associates Inc., 2012.

David Maxwell Chickering. Learning Bayesian networks is NP-complete. In Doug Fisher and Hans-J. Lenz, editors, *Learning from Data*, volume 112 of *Lecture Notes in Statistics*, pages 121–130. Springer New York, 1996.

Ashish Choudhary. Edmonds algorithm. online, Aug. 2009. URL <http://www.mathworks.com/matlabcentral/fileexchange/24899-edmonds-algorithm>.

C. Chow and C. Liu. Approximating discrete probability distributions with dependence trees. *IEEE Transactions on Information Theory*, 14(3):462–467, 1968.

Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 2006.

Abhimanyu Das and David Kempe. Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection. In *Proceedings of the 28th International Conference on Machine Learning (ICML)*, 2011.

Sanjoy Dasgupta. Learning polytrees. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 134–141. Morgan Kaufmann Publishers Inc., 1999.

Jack Edmonds. Optimum branchings. *Journal of Research of the National Bureau of Standards, Section B*, 71:233–240, 1967.

Michael Eichler. Graphical modelling of multivariate time series. *Probability Theory and Related Fields*, 153(1-2):233–268, 2012.

Natalia Flerova, Emma Rollon, and Rina Dechter. Bucket and mini-bucket schemes for m best solutions over graphical models. In *Graph Structures for Knowledge Representation and Reasoning*, pages 91–118. Springer, 2012.

Menachem Fromer and Amir Globerson. An LP View of the M-best MAP problem. In Y. Bengio, D. Schuurmans, J.D. Lafferty, C.K.I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems (NIPS) 22*, pages 567–575. Curran Associates, Inc., 2009.

Satoru Fujishige. Polymatroidal dependence structure of a set of random variables. *Information and Control*, 39(1):55–72, 1978.

Serge Gaspers, Mikko Koivisto, Mathieu Liedloff, Sebastian Ordyniak, and Stefan Szeider. On finding optimal polytrees. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, pages 750–756, 2012.

Clive W.J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969.

Stefan Haufe, Guido Nolte, Klaus-Robert Müller, and Nicole Krämer. Sparse causal discovery in multivariate time series. In *Advances in Neural Information Processing Systems 2008, Workshop on Causality*, volume 6, pages 97–106, 2010.

Tommi Jaakkola, David Sontag, Amir Globerson, and Marina Meila. Learning Bayesian network structure using LP relaxations. In *International Conference on Artificial Intelligence and Statistics*, pages 358–365, 2010.

Ali Jalali, Christopher C. Johnson, and Pradeep K. Ravikumar. On learning discrete graphical models using greedy methods. In J. Shawe-Taylor, R.S. Zemel, P.L. Bartlett, F. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 1935–1943. Curran Associates, Inc., 2011.

Alexander Jung, Reinhard Heckel, Helmut Bölcskei, and Franz Hlawatsch. Compressive nonparametric graphical model selection for time series. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 769–773, May 2014.

Jyrki Katajainen, Tomi Pasanen, and Jukka Teuhola. Practical in-place mergesort. *Nordic Journal of Computing*, 3(1):27–40, 1996.

Sanggyun Kim, David Putrino, Soumya Ghosh, and Emery N. Brown. A Granger causality measure for point process models of ensemble neural spiking activity. *PLoS Computational Biology*, 7(3), March 2011.

Sanggyun Kim, Christopher J. Quinn, Negar Kiyavash, and Todd P. Coleman. Dynamic and succinct statistical analysis of neuroscience data. *Proceedings of the IEEE*, 102(5):683–698, 2014.

Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press, 2009.

Gerhard Kramer. *Directed information for channels with feedback*. PhD thesis, Swiss Federal Institute of Technology (ETH), Zürich, Switzerland, 1998.

Eugene L. Lawler. A procedure for computing the k best solutions to discrete optimization problems and its application to the shortest path problem. *Management Science*, 18(7):401–405, 1972.

Hans Marko. The bidirectional communication theory—a generalization of information theory. *IEEE Transactions on Communications*, 21(12):1345–1351, Dec 1973.

James Massey. Causality, feedback and directed information. In *Proceedings of the International Symposium on Information Theory and its Applications (ISITA)*, pages 27–30, 1990.

George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. An analysis of approximations for maximizing submodular set functions I. *Mathematical Programming*, 14:265–294, 1978.

Teppo Niinimäki and Mikko Koivisto. Annealed importance sampling for structure learning in Bayesian networks. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1579–1585. AAAI Press, 2013.

Dennis Nilsson. An efficient algorithm for finding the M-most probable configurations in probabilistic expert systems. *Statistics and Computing*, 8(2):159–173, 1998.

Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition, 2009.

Franz Pernkopf and Jeff A. Bilmes. Efficient heuristics for discriminative structure learning of Bayesian network classifiers. *Journal of Machine Learning Research*, 11:2323–2360, Aug. 2010.

Christopher J. Quinn, Todd P. Coleman, Negar Kiyavash, and Nicholas G. Hatsopoulos. Estimating the directed information to infer causal relationships in ensemble neural spike train recordings. *Journal of Computational Neuroscience*, 30(1):17–44, 2011.

Christopher J. Quinn, Negar Kiyavash, and Todd P. Coleman. Directed Information Graphs. *ArXiv e-prints*, April 2012.

Christopher J. Quinn, Negar Kiyavash, and Todd P. Coleman. Efficient methods to compute optimal tree approximations of directed information graphs. *IEEE Transactions on Signal Processing*, 61(12):3173–3182, 2013a.

Christopher J. Quinn, Ali Pinar, and Negar Kiyavash. Optimal bounded-degree approximations of joint distributions of networks of stochastic processes. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, pages 2264–2268, July 2013b.

Arvind Rao, Alfred O. Hero, David J. States, and James D. Engel. Motif discovery in tissue-specific regulatory sequences using directed information. *EURASIP Journal on Bioinformatics and Systems Biology*, 2007:1–13, 2007.

Arvind Rao, Alfred O. Hero, David J. States, and James D. Engel. Using directed information to build biologically relevant influence networks. *Journal of Bioinformatics and Computational Biology*, 6(3):493–520, 2008.

Pradeep Ravikumar, Martin J. Wainwright, and John D. Lafferty. High-dimensional Ising model selection using L1-regularized logistic regression. *The Annals of Statistics*, 38(3):1287–1319, 2010.

Kelvin So, Aaron C. Koralek, Karunesh Ganguly, Michael C. Gastpar, and Jose M. Carmena. Assessing functional connectivity of neural ensembles using directed information. *Journal of Neural Engineering*, 9(026004):1–13, 2012.

Kazutaka Takahashi, Sanggyun Kim, Todd P. Coleman, Kevin A. Brown, Aaron J. Suminski, Matthew D. Best, and Nicholas G. Hatsopoulos. Large-scale spatiotemporal spike patterning consistent with wave propagation in motor cortex. *Nature Communications*, 6(7169):1–11, 2015.

Greg Ver Steeg and Aram Galstyan. 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, 2012. ACM.

Greg Ver Steeg and Aram Galstyan. Information-theoretic measures of influence based on content dynamics. In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, pages 3–12. ACM, 2013.

Chen Yanover and Yair Weiss. Finding the m most probable configurations using loopy belief propagation. In S. Thrun, L.K. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems (NIPS) 16*, pages 289–296. MIT Press, 2004.

Kun Zhang, Zhikun Wang, Jiji Zhang, and Bernhard Schölkopf. On estimation of functional causal models: General results and application to post-nonlinear causal model. *ACM Transactions on Intelligent Systems and Technologies*, 2014.

Kun Zhang, Jiji Zhang, and Bernhard Schölkopf. Distinguishing cause from effect based on exogeneity. In *Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge*, 2015.