

Learning Non-Gaussian Graphical Models via Hessian Scores and Triangular Transport

Ricardo Baptista*

RSB@CALTECH.EDU

*Department of Computing and Mathematical Sciences
California Institute of Technology
Pasadena, CA 91125-2100, USA*

Youssef Marzouk

YMARZ@MIT.EDU

*Center for Computational Science and Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139-4301, USA*

Rebecca Morrison*

REBECCAM@COLORADO.EDU

*Department of Computer Science
University of Colorado Boulder
Boulder, CO 80309-0430, USA*

Olivier Zahm

OLIVIER.ZAHM@INRIA.FR

*Université Grenoble Alpes
Inria, CNRS, Grenoble INP, LJK
38000 Grenoble, France*

Editor: Gal Elidan

Abstract

Undirected probabilistic graphical models represent the conditional dependencies, or Markov properties, of a collection of random variables. Knowing the sparsity of such a graphical model is valuable for modeling multivariate distributions and for efficiently performing inference. While the problem of learning graph structure from data has been studied extensively for certain parametric families of distributions, most existing methods fail to consistently recover the graph structure for non-Gaussian data. Here we propose an algorithm for learning the Markov structure of continuous and non-Gaussian distributions. To characterize conditional independence, we introduce a score based on integrated Hessian information from the joint log-density, and we prove that this score upper bounds the conditional mutual information for a general class of distributions. To compute the score, our algorithm SING estimates the density using a deterministic coupling, induced by a triangular transport map, and iteratively exploits sparse structure in the map to reveal sparsity in the graph. For certain non-Gaussian datasets, we show that our algorithm recovers the graph structure even with a biased approximation to the density. Among other examples, we apply SING to learn the dependencies between the states of a chaotic dynamical system with local interactions.

Keywords: Undirected graphical models, structure learning, non-Gaussian distributions, conditional mutual information, transport map, sparsity

*. Authors contributed equally.

1. Introduction

Consider a collection of random variables $\mathbf{Z} = (Z_1, \dots, Z_d)$ with probability measure ν_π and Lebesgue density π . An undirected graphical model is a representation of the conditional independence structure, or Markov properties, satisfied by \mathbf{Z} . In particular, it is a graph $\mathcal{G} = (V, E)$ where the set of vertices $V = \{1, \dots, d\}$ contains the indices of the random variables and where the set of edges E encodes the statistical dependence structure of \mathbf{Z} in the following way: for any disjoint subsets A , B , and C of V , \mathbf{Z}_A and \mathbf{Z}_B are conditionally independent given \mathbf{Z}_C if C separates A and B in the graph \mathcal{G} . That is, if after removing nodes C from \mathcal{G} , there is no path between sets A and B in the resulting graph, then the conditional independence property

$$\mathbf{Z}_A \perp\!\!\!\perp \mathbf{Z}_B \mid \mathbf{Z}_C \tag{1}$$

holds. In this case, we say that \mathcal{G} is an independence map (I-map), or Markov network, for ν_π . In this work, we focus on minimal I-maps, which are the sparsest I-maps for ν_π . Our objective is to identify the minimal I-map of ν_π using samples from ν_π .

Learning the Markov properties of a distribution, given a set of data drawn from it, is useful for various reasons. Undirected graphs can reveal hierarchical and cyclic interactions between the variables (Epskamp et al., 2018; Shen et al., 2009). These graph structures have been used to interpret datasets from many application areas such as handwriting recognition (Fotak et al., 2011; Fischer et al., 2013), natural language processing (Biemann, 2006; Mihalcea and Radev, 2011), protein folding (Thomas et al., 2008), and many more (Bloom and Golomb, 1977; Saaty and Busacker, 1965). The graphical model also defines efficient factorizations of high-dimensional distributions by representing the density as a product of potential functions that each depend on low-dimensional subsets of variables (Wainwright et al., 2007). In some fields, such as image processing and spatial statistics, the Markov structure of the problem may be immediately recognizable (Koller and Friedman, 2009), and exploiting the structure of the undirected graphical model greatly simplifies algorithms for inference and prediction.

Previous work in learning probabilistic graphical models has mainly focused on certain parametric families of distributions (Drton and Maathuis, 2017). For Gaussian random variables, the conditional independence properties are encoded by the sparsity of the inverse covariance, or *precision*, matrix. That is, the (i, j) entry of the precision matrix is zero if and only if variables Z_i and Z_j are conditionally independent given the rest. Thus, learning the graph reduces to identifying the support of the non-zero entries in the precision matrix. An active area of research considers how to learn the graph with (very) sparse data relative to the dimension d of the variables. One of the best-known methods to tackle this problem is the graphical lasso (GLASSO), introduced by Banerjee et al. (2008); Yuan and Lin (2007), which solves an ℓ_1 -penalized maximum likelihood estimation problem for the precision matrix. One popular method for finding the GLASSO estimator is the coordinate descent algorithm of Friedman et al. (2008). In the case of discrete random variables, approaches for finding sparse graphs include ℓ_1 -penalized logistic regression (Wainwright et al., 2007), and the estimation of a generalized covariance matrix whose inverse (via its support) encodes conditional independence properties (Loh and Wainwright, 2012).

In the case of continuous and non-Gaussian data, the connection between the inverse covariance matrix and conditional independence is lost. Outside of the Gaussian setting, a

regularized score matching method was proposed to learn sparse graphs for distributions within the exponential family (Lin et al., 2016). Recently, a large class of multivariate graphical models was considered by combining node-wise conditional distributions in the exponential family (Yang et al., 2015; Suggala et al., 2017). Another area of research has proposed semiparametric methods based on Gaussian copulas to model non-Gaussian data (Liu et al., 2009). In this case, the observations are assumed to come from marginal nonlinear transformations of a multivariate Gaussian random vector with known Markov properties. The marginal transformations yield potentially non-Gaussian marginal distributions while at the same time preserving the I-map of the original multivariate Gaussian distribution. Per Sklar’s theorem, any multivariate distribution can be written in terms of a copula and the univariate marginals (Nelsen, 2007). Recently, these copulas were extended to fit models based on elliptical distributions (Liu et al., 2012b). However, the specific families of copulas that can be easily or explicitly described remain rather limited (Asmussen and Glynn, 2007). Moreover, data generated via Gaussian copula transformations may fail to fully test a structure learning algorithm’s ability to handle non-Gaussianity.

So far, there has been relatively little work for general non-Gaussian distributions outside of the exponential family. The current paper is concerned with providing a mathematical and algorithmic framework to describe and identify the Markov properties of a continuous and non-Gaussian distribution. As an expansion of our NeurIPS paper, Morrison et al. (2017), the main contributions are as follows.

First, we establish a framework which allows for the description and computation of conditional independence properties in the non-Gaussian setting. This is represented by a new *conditional independence score* matrix Ω : each entry Ω_{ij} is a score for the independence of Z_i and Z_j conditioned on the remaining variables. Each score Ω_{ij} is defined by the expected magnitude of certain mixed derivatives—in other words, integrated *Hessian* information—from the joint density π . For strictly positive and continuously differentiable π , this score is zero if and only if Z_i and Z_j are conditionally independent. We also show that, under certain assumptions, the score provides an upper bound for the *conditional mutual information*, a widely used measure of conditional dependence. To compute Ω given only samples from π , an estimate of the joint density is needed. This is achieved via a transport map—a transformation that deterministically couples one probability measure to another.

Second, we expand the use and analysis of an algorithm called SING (Sparsity Identification in Non-Gaussian distributions). A key element of the algorithm is a thresholding scheme, and thus we propose a class of threshold estimators that are *consistent* for graph recovery. The success of the algorithm also depends on a strong and explicit connection between the sparsity of the graph and the sparsity of triangular transport maps associated with π . This connection is exploited to iteratively reveal sparsity in the graph. We show numerically that this iterative algorithm provides an improved estimator for the conditional independence score, compared to a non-iterative approach that does not account for the structure in the map.

Third, we explore the relationship between the distribution in question and how to learn its corresponding graph. Because there exist an infinite number of probability distributions with the same minimal I-map, intuitively it should be easier to identify the graph than to estimate the entire joint distribution. We refer to this notion as the *information gap*. Through several empirical and analytical studies, we explore the extent to which this notion

is in fact true and discuss its consequences for graph identification. Moreover, we see how this notion can be directly exploited by the algorithm. That is, in some cases, we still recover the correct graph even with a biased approximation to the density resulting from a constrained parameterization for the transport map. On the other hand, it is also possible that constraining the form of the map yields an incorrect graph. We show examples of both cases in Section 6.

The remainder of the paper is organized as follows. Section 2 introduces the conditional independence score and its connections to conditional mutual information. Section 3 describes a sample-based estimator for the score based on a transport map approximation to the density. Section 4 introduces our first algorithm for learning the Markov structure, and proposes a class of threshold estimators that are consistent for this structure. To take advantage of the connection between the sparsity of the graph and the transport map, Section 5 presents the iterative algorithm SING. Section 6 demonstrates the performance of the algorithm on a variety of numerical examples and explores the notion of the information gap. Section 7 provides a discussion and outlook on future research directions. Finally, Appendix A contains the proof of our main result on the conditional independence score, Appendix B collects the proofs of results related to transport maps, and Appendix C contains the proof of consistency for the graph estimator.

2. Measures of conditional independence

An alternative to the global Markov properties (1) for characterizing conditional independence are the pairwise Markov properties. A distribution for $\mathbf{Z} = (Z_1, \dots, Z_d)$ satisfies a pairwise Markov property between two variables when the associated nodes are not connected in the graph $\mathcal{G} = (V, E)$. That is, given two variables Z_i and Z_j for $i \neq j$, the lack of an edge between nodes i and j means the two variables are conditionally independent given the remaining variables:

$$(i, j) \notin E \iff Z_i \perp\!\!\!\perp Z_j \mid \mathbf{Z}_{-ij}.$$

Here \mathbf{Z}_{-ij} denotes the random vector obtained by removing the i th and j th components from \mathbf{Z} . The pairwise Markov property is, in general, weaker than the global, but when the density is strictly positive, i.e., $\pi(\mathbf{z}) > 0 \forall \mathbf{z} \in \mathbb{R}^d$, the global and the pairwise Markov properties are equivalent (see Lauritzen, 1996). In this paper, we restrict our attention to this setting.

By definition, $Z_i \perp\!\!\!\perp Z_j \mid \mathbf{Z}_{-ij}$ means that the joint conditional density $\pi(z_i, z_j \mid \mathbf{z}_{-ij}) = \pi(z_i, z_j, \mathbf{z}_{-ij}) / \int \pi(z'_i, z'_j, \mathbf{z}_{-ij}) dz'_i dz'_j$ factorizes as the product of the conditional marginals

$$\pi(z_i, z_j \mid \mathbf{z}_{-ij}) = \pi(z_i \mid \mathbf{z}_{-ij}) \pi(z_j \mid \mathbf{z}_{-ij}), \tag{2}$$

where $\pi(z_i \mid \mathbf{z}_{-ij}) = \int \pi(z_i, z_j \mid \mathbf{z}_{-ij}) dz_j$ and $\pi(z_j \mid \mathbf{z}_{-ij}) = \int \pi(z_i, z_j \mid \mathbf{z}_{-ij}) dz_i$. Thus, $Z_i \perp\!\!\!\perp Z_j \mid \mathbf{Z}_{-ij}$ allows the joint density to factorize as $\pi(\mathbf{z}) = \pi(z_i \mid \mathbf{z}_{-ij}) \pi(z_j \mid \mathbf{z}_{-ij}) \pi(\mathbf{z}_{-ij})$. If we further assume that π is continuously differentiable, this yields

$$\partial_i \partial_j \log \pi(\mathbf{z}) = 0, \tag{3}$$

for any $\mathbf{z} \in \mathbb{R}^d$. Conversely, if a strictly positive continuously differentiable $\pi(\mathbf{z})$ satisfies (3) for any $\mathbf{z} \in \mathbb{R}^d$, then $\pi(z_i, z_j \mid \mathbf{z}_{-ij})$ necessarily factors as in (2) so that $Z_i \perp\!\!\!\perp Z_j \mid \mathbf{Z}_{-ij}$. The

characterization (3) of conditional independence has been already observed in Lemma 4.1 of Spantini et al. (2018). Based on (3), we propose to measure the conditional independence of Z_i and Z_j by the score $\Omega_{ij} \geq 0$ defined as

$$\Omega_{ij} := \int |\partial_i \partial_j \log \pi(\mathbf{z})|^2 \pi(\mathbf{z}) d\mathbf{z}. \quad (4)$$

A similar measure also appears in Morrison et al. (2017) under the name of “generalized precision,” the difference being the square inside the integral.¹ For a strictly positive and continuously differentiable density π , the condition $\Omega_{ij} = 0$ yields (3) so that Z_i and Z_j are conditionally independent. That is, the sparsity pattern of Ω gives the Markov structure of π . The entries of Ω also provide a natural score for conditional independence: a value of Ω_{ij} near zero means that Z_i and Z_j are nearly conditionally independent, whereas a large value of Ω_{ij} means that Z_i and Z_j are strongly conditionally dependent. In Section 4.2, we will use this interpretation to estimate the Markov structure of π by *thresholding* the entries of an estimator for Ω .

We can relate the magnitude of score Ω_{ij} to another popular measure of conditional independence, the conditional mutual information (CMI). The CMI $I(Z_i; Z_j | \mathbf{Z}_{-ij})$ is defined as the expected (with respect to \mathbf{Z}_{-ij}) Kullback–Leibler divergence from the product of the marginal conditionals $\pi(z_i | \mathbf{z}_{-ij})$ and $\pi(z_j | \mathbf{z}_{-ij})$ to the joint conditional $\pi(z_i, z_j | \mathbf{z}_{-ij})$; that is,

$$I(Z_i; Z_j | \mathbf{Z}_{-ij}) = \int \left(\int \pi(z_i, z_j | \mathbf{z}_{-ij}) \log \frac{\pi(z_i, z_j | \mathbf{z}_{-ij})}{\pi(z_i | \mathbf{z}_{-ij}) \pi(z_j | \mathbf{z}_{-ij})} dz_i dz_j \right) \pi(\mathbf{z}_{-ij}) d\mathbf{z}_{-ij} \quad (5)$$

$$= \int \log \left(\frac{\pi(z_i, z_j | \mathbf{z}_{-ij})}{\pi(z_i | \mathbf{z}_{-ij}) \pi(z_j | \mathbf{z}_{-ij})} \right) \pi(\mathbf{z}) d\mathbf{z}. \quad (6)$$

The CMI is widely adopted in part due to its information theoretic interpretations. The following theorem shows that $I(Z_i; Z_j | \mathbf{Z}_{-ij})$ can in fact be *bounded above* by Ω_{ij} . This result relies on logarithmic Sobolev inequalities.

Definition 1 *A probability density function π on \mathbb{R}^d satisfies the logarithmic Sobolev inequality if there exists a constant $C < \infty$ such that*

$$\int h \log \frac{h}{\int h d\pi} d\pi \leq \frac{C}{2} \int \|\nabla \log h\|_2^2 h d\pi, \quad (7)$$

holds for any continuously differentiable function $h : \mathbb{R}^d \rightarrow \mathbb{R}_{>0}$. Here $\|\cdot\|_2$ denotes the canonical Euclidean norm on \mathbb{R}^d . The smallest constant $C = C(\pi)$ such that (7) holds is called the logarithmic Sobolev constant of π .

Theorem 2 *Let π be a strictly positive continuously differentiable probability density function on \mathbb{R}^d . Assume that all the conditional densities of the form $\pi(\cdot | \mathbf{z}_{-i}) : z_i \mapsto \pi(z_i | \mathbf{z}_{-i})$ and $\pi(\cdot, \cdot | \mathbf{z}_{-ij}) : (z_i, z_j) \mapsto \pi(z_i, z_j | \mathbf{z}_{-ij})$ satisfy the logarithmic Sobolev inequality with constants uniformly bounded by some constant $C_0 < \infty$, meaning*

$$C(\pi(\cdot | \mathbf{z}_{-i})) \leq C_0 \quad \text{and} \quad C(\pi(\cdot, \cdot | \mathbf{z}_{-ij})) \leq C_0, \quad (8)$$

1. In this paper, we no longer use the term “generalized precision,” but keep the same notation Ω to denote the new conditional independence score.

for all $\mathbf{z}_{-i} \in \mathbb{R}^{d-1}$ and $\mathbf{z}_{-ij} \in \mathbb{R}^{d-2}$ and for all $1 \leq i \neq j \leq d$. Then

$$I(Z_i; Z_j | \mathbf{Z}_{-ij}) \leq C_0^2 \Omega_{ij}, \tag{9}$$

holds for any $i \neq j$.

The proof of Theorem 2 is provided in Appendix A.

Let us now comment on assumption (8) of Theorem 2. In general, there is no simple way to compute exactly the logarithmic Sobolev constant of a density. The Holley–Stroock perturbation criterion (Holley and Stroock, 1976) and the Bakry–Émery criterion (Bakry and Émery, 1985) are commonly used to bound the logarithmic Sobolev constant. Following Zahm et al. (2022), one can combine these two criteria to show that if there exists a log-concave probability density π_0 and two scalars $\alpha > 0$ and $\beta < \infty$ such that $\alpha\pi_0(\mathbf{z}) \leq \pi(\mathbf{z}) \leq \beta\pi_0(\mathbf{z})$ for any $\mathbf{z} \in \mathbb{R}^d$, then π satisfies (8) with constant

$$C_0 = \frac{\beta}{\alpha\lambda}, \tag{10}$$

where $\lambda > 0$ is a lower bound for the smallest eigenvalue of $-\nabla^2 \log \pi_0(\mathbf{z})$ for any $\mathbf{z} \in \mathbb{R}^d$. We refer the reader to Zahm et al. (2022, Section 2.2) for more details and examples of probability densities that satisfy the above conditions.

Remark 3 (Gaussian case) Suppose that $\mathbf{Z} \sim \mathcal{N}(\mathbf{m}, \Sigma)$ is a Gaussian vector with mean $\mathbf{m} \in \mathbb{R}^d$ and non-singular covariance $\Sigma \in \mathbb{R}^{d \times d}$. Because $\pi(\mathbf{z}) \propto \exp(-\frac{1}{2}(\mathbf{z} - \mathbf{m})^\top \Sigma^{-1}(\mathbf{z} - \mathbf{m}))$ we have

$$\Omega_{ij} = (\Sigma^{-1})_{ij}^2, \tag{11}$$

and so the sparsity pattern of the precision matrix Σ^{-1} gives the Markov structure of \mathbf{Z} . This is a well known property of Gaussian vectors; see for instance Proposition 3.3.6. in Drton et al. (2008). This also explains the name “generalized precision” in Morrison et al. (2017).

Next we show how sharp the inequality (9) is for a $d = 2$ dimensional Gaussian vector $\mathbf{Z} = (Z_1, Z_2)$ with covariance given $\Sigma = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$ for some correlation $-1 \leq \rho \leq 1$. One can compute the (conditional) mutual information analytically $I(Z_1; Z_2 | \mathbf{Z}_{-12}) = I(Z_1; Z_2) = -\frac{1}{2} \log(1 - \rho^2)$ and the score $\Omega_{12} = (\frac{\rho}{1 - \rho^2})^2$. It remains to compute C_0 . Using the formula (10) with $\beta = \alpha = 1$ (since π is log-concave we can chose $\pi_0 = \pi$), we obtain $C_0 = \lambda_{\min}(-\nabla^2 \log \pi(\mathbf{z}))^{-1} = \lambda_{\max}(\Sigma)$. Because $\lambda_{\max}(\Sigma) \leq 2$, we then have

$$I(Z_1; Z_2) = -\frac{1}{2} \log(1 - \rho^2) \stackrel{(9)}{\leq} C_0^2 \Omega_{12} = \lambda_{\max}(\Sigma)^2 \left(\frac{\rho}{1 - \rho^2} \right)^2 \leq 4 \left(\frac{\rho}{1 - \rho^2} \right)^2.$$

We complete this section by comparing the *computational cost* of evaluating the CMI with that of evaluating the conditional independence score Ω_{ij} , for one variable pair (Z_i, Z_j) . Given the joint density $\pi(\mathbf{z})$, computing the CMI requires the ability to evaluate the *normalized* conditional density $\pi(z_i, z_j | \mathbf{z}_{-ij})$ and the two normalized marginal conditionals, $\pi(z_i | \mathbf{z}_{-ij})$ and $\pi(z_j | \mathbf{z}_{-ij})$, for any value of $z_i, z_j \in \mathbb{R}$, $\mathbf{z}_{-ij} \in \mathbb{R}^{d-2}$. Breaking this down further, evaluating $\pi(z_i, z_j | \mathbf{z}_{-ij})$ requires integrating the joint density with respect to (z_i, z_j) , while the marginal densities involve further integration with respect to the z_i and z_j variables individually. Then one must take an outer expectation over \mathbf{z} . Evaluating the CMI therefore requires computing

nested integrals, which is known to be challenging without additional structure (Rainforth et al., 2018). In practice, information theoretic quantities that involve nested integration, such as CMI, are approximated using nested Monte Carlo (NMC) estimators. Given a budget of n samples, the best-case root-mean-squared error of NMC converges at a slow asymptotic rate of $\mathcal{O}(n^{-1/3})$. Moreover, as these Monte Carlo estimators involve the repeated computation of normalizing constants, controlling variance requires the construction of suitable biasing distributions for importance sampling (Gelman and Meng, 1998), often in a way that depends on the conditioning variables (Feng and Marzouk, 2019).

On the other hand, computing the score entry Ω_{ij} requires evaluating a single mixed derivative of the log density $\log \pi(\mathbf{z})$ for any $\mathbf{z} \in \mathbb{R}^d$. In computing derivatives of the log density, we only require access to the unnormalized $\pi(\mathbf{z})$. Moreover, the outer expectation of the score can be approximated with standard non-nested Monte Carlo estimators, converging at the usual $\mathcal{O}(n^{-1/2})$ rate. Samples can be drawn directly from the target π , and no importance sampling of any kind is required. For these reasons, in this work we focus on efficiently estimating the score matrix rather than computing the CMI.

In either case, evaluating the CMI or the conditional independence score requires an explicit functional form of the joint density π . In a setting where we only have access to samples from π , many structure learning algorithms rely on density estimation techniques as a step to learn the Markov structure. In general, density estimation in high dimensions, especially for non-Gaussian data, can become computationally expensive. Thus, a major question becomes: What is an efficient or advantageous way to represent the density given the goal of learning the Markov structure of a distribution? To answer this, we rely on a particular method of density estimation, using transport maps.

3. Transport maps

Given a target probability measure ν_π and a reference probability measure ν_η , both on \mathbb{R}^d , a transport map $S: \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a measurable function that defines a deterministic coupling between these two measures (Villani, 2008). For random variables $\mathbf{Z} \sim \nu_\pi$ and $\mathbf{X} \sim \nu_\eta$, this map ensures $S(\mathbf{Z}) = \mathbf{X}$ in distribution. If this map is invertible, we can sample from π by generating i.i.d. samples $\{\mathbf{x}^l\}_{l>0}$ from η and evaluating the inverse map at these samples to obtain i.i.d. samples $S^{-1}(\mathbf{x}^l)$ from π . The measure of $S(\mathbf{Z})$ is known as the pushforward measure of ν_π through S and is denoted by $S_\# \nu_\pi$. If the measure ν_π has density π , the density of the pushforward measure is denoted by $S_\# \pi$. Similarly, the measure of $S^{-1}(\mathbf{X})$ is known as the pullback measure of ν_η through S and is denoted by $S^\# \nu_\eta$. If the measure ν_η has density η , the density of the pullback measure is denoted by $S^\# \eta$. For a diffeomorphism S (i.e., a differentiable bijective map with differentiable inverse), the pullback density $S^\# \eta$ is given by:

$$S^\# \eta(\mathbf{z}) = \eta \circ S(\mathbf{z}) |\det \nabla S(\mathbf{z})|, \tag{12}$$

where $\det \nabla S(\mathbf{z})$ denotes the determinant of the Jacobian of the map S at \mathbf{z} , and \circ represents the composition operator.

In general, there may exist many different transport maps that couple two arbitrary probability distributions. Optimal transport identifies one such map by minimizing an integrated transportation cost that represents the effort of “moving” samples from one distribution to another (Peyré et al., 2019). We will instead consider transport maps with

triangular structure, and show how to learn such maps given only samples from the target density π . Triangular structure will confer several advantages, which we detail below. Besides the structure of the map, another important degree of freedom is the choice of reference distribution. For the remainder of this paper, we choose a standard Gaussian reference, $\nu_\eta = \mathcal{N}(\mathbf{0}, I_d)$. As we will see below, this choice, coupled with triangular structure, leads to useful simplicity in the optimization problem used to build estimators of the map. Moreover, the fact that ν_η is a product measure guarantees that the triangular transport map inherits sparsity from the Markov structure of ν_π (Spantini et al., 2018). This connection will be exploited when we discuss our iterative algorithm for learning the graphical model in Section 5.

3.1 Lower triangular transport maps

We consider triangular transports between smooth and strictly positive densities π and η on \mathbb{R}^d . In this setting,² Rosenblatt (1952); Knothe (1957); Bogachev et al. (2005) showed that there exists a unique monotone lower triangular map S —known as the Knothe-Rosenblatt (KR) rearrangement—such that $S^\# \eta = \pi$. A monotone lower triangular map S is a multivariate function of the form

$$S(\mathbf{z}) = \begin{bmatrix} S^1(z_1) \\ S^2(z_1, z_2) \\ S^3(z_1, z_2, z_3) \\ \vdots \\ S^d(z_1, \dots, z_d) \end{bmatrix},$$

where the k th component S^k only depends on the first k input variables $\mathbf{z}_{1:k} := (z_1, \dots, z_k)$, and the map $\xi \mapsto S^k(z_1, \dots, z_{k-1}, \xi)$ —the restriction of the component onto its first $k - 1$ inputs—is monotone increasing for all $(z_1, \dots, z_{k-1}) \in \mathbb{R}^{k-1}$. We denote the space of such lower triangular maps as \mathcal{S}_Δ .

For differentiable triangular maps, the Jacobian determinant of S can be easily evaluated as the product of its partial derivatives, i.e., $\det \nabla S = \prod_{k=1}^d \partial_k S^k$. This enables the pullback density in (12) to be efficiently evaluated. For more discussion on the advantages of a triangular structure for S we refer the reader to Marzouk et al. (2016).

Remark 4 (Gaussian case) *Suppose that $\mathbf{Z} \sim \pi = \mathcal{N}(\mathbf{0}, \Sigma)$ is a Gaussian vector with non-singular covariance and $\mathbf{X} \sim \eta = \mathcal{N}(\mathbf{0}, I_d)$. Let $L^T L = \Sigma^{-1}$ be the Cholesky decomposition of the inverse covariance matrix of \mathbf{Z} . Then, L is a linear operator that maps samples $\mathbf{z} \sim \pi$ from the target density to samples $\mathbf{x} \sim \eta$ from the reference density, and similarly, L^{-1} maps \mathbf{x} to \mathbf{z} . That is,*

$$L\mathbf{z} = \mathbf{x}, \quad L^{-1}\mathbf{x} = \mathbf{z}.$$

Thus, $L\mathbf{z}$ is an example of a linear lower triangular transport map $S(\mathbf{z})$. We note that other non-triangular transformations also map \mathbf{z} to \mathbf{x} (e.g., any inverse square root of Σ^{-1}), but L is the unique lower triangular map.

More generally, an affine transport map is sufficient to represent Gaussian target distributions. On the other hand, nonlinear maps S can represent non-Gaussian target distributions.

2. More generally, we only need the measures ν_π, ν_η to be absolutely continuous with respect to the Lebesgue measure on \mathbb{R}^d (see Santambrogio, 2015).

3.2 Diagonal transport maps

Diagonal maps $D: \mathbb{R}^d \rightarrow \mathbb{R}^d$ are one class of lower triangular transport maps where each map component D^k depends only on the k th input variable z_k .

Let ν_ρ be an arbitrary measure with density ρ . While the pullback measure of ν_ρ through a general lower triangular transport map S has a different Markov structure than ν_ρ , the pullback measure of ν_ρ through a diagonal map has the same Markov structure as ν_ρ . This is formalized in the following proposition, which we prove in Appendix B.

Proposition 5 *Let ν_ρ be a measure with strictly positive density ρ that is Markov with respect to \mathcal{G} , and let D be a differentiable diagonal transport map. Then, the pullback measure $D^\# \nu_\rho$ is also Markov with respect to \mathcal{G} .*

If the pullback density $\pi = D^\# \rho$ satisfies the hypotheses in Theorem 2, we can also observe this property of diagonal transformations in the conditional independence score of π . The following proposition shows that the entries of the score matrix are related to the mixed partial derivatives of the log-density of ρ and the derivatives of the diagonal map components.

Proposition 6 *Let ρ be a strictly positive continuously differentiable density and let D and its inverse be differentiable diagonal transport maps. Then, the conditional independence score of the pullback density $D^\# \rho$ is:*

$$\Omega_{ij} = \int |\partial_i \partial_j \log \rho(\mathbf{x}) \partial_i (D^i)^{-1}(x_i) \partial_j (D^j)^{-1}(x_j)|^2 \rho(\mathbf{x}) d\mathbf{x}. \quad (13)$$

This result is proved in Appendix B.

Remark 7 (Gaussian case) *Suppose that ρ is a multivariate Gaussian density with non-singular covariance $\Sigma \in \mathbb{R}^{d \times d}$. Then the conditional independence score of $D^\# \rho$ has the form:*

$$\Omega_{ij} = (\Sigma^{-1})_{ij}^2 \int |\partial_i (D^i)^{-1}(x_i) \partial_j (D^j)^{-1}(x_j)|^2 \rho(\mathbf{x}) d\mathbf{x}.$$

Thus, if the transformation D is strictly increasing so that $D^i(z_i) \neq 0$ for all i , then the support of Ω is identical to the support of the inverse covariance matrix, i.e., $\Omega_{ij} = 0$ if and only if $(\Sigma^{-1})_{ij} = 0$. Furthermore, each nonzero entry of the conditional independence score Ω_{ij} is proportional to $(\Sigma^{-1})_{ij}^2$.

The pullback of a multivariate Gaussian density ρ through a diagonal map D is an example of a Gaussian copula. It is well known that Gaussian copulas preserve the Markov properties of ρ , while introducing non-Gaussianity via nonlinear diagonal transformations. This class of distributions was considered in the context of structure learning in Liu et al. (2009) and will be studied numerically in Section 6.

3.3 Monotone transport map parameterizations

To approximate the KR rearrangement between a pair of densities on \mathbb{R}^d , we consider a differentiable and monotone increasing representation for the map. Monotonicity is enforced component-wise by ensuring the derivative of the k th component S^k with respect to the k th variable is a strictly positive function, i.e., $\partial_k S^k(z_1, \dots, z_k) > 0$ for all $(z_1, \dots, z_k) \in \mathbb{R}^k$. A general parameterization that guarantees monotonicity of S^k is:

$$S^k(z_1, \dots, z_k) = c_k(z_1, \dots, z_{k-1}) + \int_0^{z_k} g(h_k(z_1, \dots, z_{k-1}, t)) dt, \quad (14)$$

for some positive function $g: \mathbb{R} \rightarrow \mathbb{R}_+$ and functions $c_k: \mathbb{R}^{k-1} \rightarrow \mathbb{R}$ and $h_k: \mathbb{R}^k \rightarrow \mathbb{R}$ (Ramsay, 1998). Two common choices for g in (14) are $g(x) = \exp(x)$ and $g(x) = x^2$, which result in the “integrated exponential” and the “integrated squared” parameterizations, respectively. In this work, we use the integrated squared parameterization because of its computational advantage of closed-form integration under certain choices for h_k (Spantini et al., 2018; Brennan et al., 2020). Following Baptista et al. (2023), we parameterize the functions c_k and h_k using linear expansions

$$c_k(\mathbf{z}) = \sum_j c_{k,j} \psi_j(\mathbf{z}), \quad h_k(\mathbf{z}) = \sum_j h_{k,j} \phi_j(\mathbf{z}), \quad (15)$$

in terms of tensor product Hermite functions (ψ_j, ϕ_j) and unknown coefficients $\boldsymbol{\alpha} = (c_{k,j}, h_{k,j})$. We note that recent methods for autoregressive density estimation alternatively represent the functions c_k and h_k using neural networks (Papamakarios et al., 2017; Jaini et al., 2019).

In practice, we truncate the expansions in (15) by prescribing a maximum total degree β for the multivariate Hermite functions in c_k and h_k . We denote the space of lower-triangular maps with total degree β as \mathcal{S}_Δ^β . As expected, a higher degree provides a richer basis for density estimation, but requires more computational effort to optimize and more samples to accurately estimate the coefficients. Here we follow the convention in Spantini et al. (2018) where a map of degree β uses basis functions up to degree β for c_k and $\beta - 1$ for h_k , since the latter is then integrated once. To include affine maps within the space \mathcal{S}_Δ^β , we also include constant and linear functions with respect to each variable in the expansions (15). Computations in Section 6 using the transport map parameterizations above are performed using the publicly-available software TRANSPORTMAPS.³

3.4 Optimization of the transport map

We complete the discussion of transport maps by describing the optimization procedure for finding the map. After defining a maximum polynomial degree for the basis functions within each map component, the resulting map $S_\alpha \in \mathcal{S}_\Delta^\beta$ is parameterized by a finite number of coefficients $\boldsymbol{\alpha} \in \mathbb{R}^p$. In this subsection we include the subscript $\boldsymbol{\alpha}$ on S to emphasize the map’s parametric dependence.

A computational approach to find the KR rearrangement that was explored in Marzouk et al. (2016) is to minimize the Kullback–Leibler divergence $D_{\text{KL}}(\pi || S^\# \eta) = \mathbb{E}_\pi[\log(\pi/S^\# \eta)]$

3. <http://transportmaps.mit.edu>

from the pullback density $S^\sharp \eta$ to π over the space of monotone increasing triangular maps \mathcal{S}_Δ . After parameterizing the maps with coefficients α , this is equivalent to solving

$$\begin{aligned} \alpha^* &= \operatorname{argmin}_{\alpha} D_{\text{KL}} \left(\pi \parallel S_{\alpha}^{\sharp} \eta \right) \\ &= \operatorname{argmax}_{\alpha} \mathbb{E}_{\pi} \left[\log S_{\alpha}^{\sharp} \eta(\mathbf{Z}) \right] \end{aligned} \quad (16)$$

As shown in Marzouk et al. (2016); Parno and Marzouk (2018), for standard Gaussian η and lower triangular S , this optimization problem is separable across the map components S^1, \dots, S^d . In addition, when using the parameterizations in Section 3.3 for each map component, the optimization problem is unconstrained and differentiable with respect to the coefficients α . Therefore, in practice we can use an iterative method such as BFGS to find the optimal solution (Nocedal and Wright, 2006). While the problem is not in general convex in α , suitable choices of g in (14) can ensure that the problem has a unique global minimizer (Baptista et al., 2023).

Given i.i.d. data $\{\mathbf{z}^l\}_{l=1}^n$ from π , we can approximate the expectation in (16) and maximize the likelihood associated with this data set. That is,

$$\begin{aligned} \hat{\alpha} &= \operatorname{argmax}_{\alpha} \frac{1}{n} \sum_{l=1}^n \log S_{\alpha}^{\sharp} \eta(\mathbf{z}^l), \\ &= \operatorname{argmax}_{\alpha} \frac{1}{n} \sum_{l=1}^n \sum_{k=1}^d \left[-\frac{1}{2} S_{\alpha}^k(\mathbf{z}_{1:k}^l)^2 + \log \partial_k S_{\alpha}^k(\mathbf{z}_{1:k}^l) \right]. \end{aligned} \quad (17)$$

where the last equality follows from the form of the standard Gaussian reference density η . The optimal solution $\hat{\alpha}$ in (17) is the maximum likelihood estimate (MLE) of α^* . Here we assume that the solutions of (16) and (17) are unique.

Under suitable regularity conditions on the log-likelihood in (17), $\hat{\alpha}$ is a consistent estimator of α^* . Furthermore, it is a random variable that converges in distribution as $n \rightarrow \infty$ to a normal random vector given by

$$\sqrt{n}(\hat{\alpha} - \alpha^*) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \Gamma(\alpha^*)^{-1}), \quad (18)$$

where $\Gamma(\alpha) \in \mathbb{R}^{p \times p}$ is the non-singular Fisher information matrix (Casella and Berger, 2002) for the transport map representation of the density π with coefficients α . Entry (i, j) of the Fisher information matrix is given by

$$\Gamma(\alpha)_{ij} := -\mathbb{E}_{\pi} \left[\partial_{\alpha_i} \partial_{\alpha_j} \log S_{\alpha}^{\sharp} \eta(\mathbf{Z}) \right]. \quad (19)$$

We conclude this section with a description of the closed-form solution to (17) when the coefficients parameterize transport maps $S_{\alpha}(\mathbf{z})$ that are affine in \mathbf{z} . A proof of this result is presented in Appendix B.

Proposition 8 (Affine map optimization) *Suppose π is an arbitrary continuous density on \mathbb{R}^d , and let $\{\mathbf{z}^l\}_{l=1}^n$ be a sample drawn from π with $n \geq d$. If the map components are restricted to be affine functions of the input variables (i.e., polynomial degree $\beta = 1$), maximizing the log-likelihood function that follows from the pullback density in (17) yields a Gaussian approximation to π given by $S_{\hat{\alpha}}^{\sharp} \eta = \mathcal{N}(\hat{\mathbf{m}}, \hat{\Sigma})$ with empirical mean $\hat{\mathbf{m}} = \frac{1}{n} \sum_{l=1}^n \mathbf{z}^l$ and empirical covariance matrix $\hat{\Sigma} = \frac{1}{n} \sum_{l=1}^n (\mathbf{z}^l - \hat{\mathbf{m}})(\mathbf{z}^l - \hat{\mathbf{m}})^T$.*

4. Learning the Markov structure

In this section, we present our first algorithm for learning the edge set of a minimal I-map for ν_π . This algorithm uses the transport map representation of the target density π to compute the conditional independence score Ω . In Section 4.1 we define a sample-based estimator for Ω and in Section 4.2 we present a thresholding procedure to identify the sparsity of Ω . Section 4.3 presents the complete algorithm for learning the edges in the graph, and Section 4.4 shows that this procedure is consistent.

4.1 Computation of Ω

After optimizing the coefficients, the pullback density $S_{\hat{\alpha}}^\# \eta$ defines an approximation to the target density π . The conditional independence score Ω_{ij} in (4) is then estimated as

$$\widehat{\Omega}_{ij} = \mathbb{E}_\pi |\partial_i \partial_j \log S_{\hat{\alpha}}^\# \eta(\mathbf{Z})|^2. \quad (20)$$

We can approximate the expectation above using n i.i.d. samples $\{\mathbf{z}^l\}_{l=1}^n$ from π , yielding a sample-based estimator of $\widehat{\Omega}$:

$$\widetilde{\Omega}_{ij} = \frac{1}{n} \sum_{l=1}^n |\partial_i \partial_j \log S_{\hat{\alpha}}^\# \eta(\mathbf{z}^l)|^2. \quad (21)$$

In this work we estimate $\widehat{\Omega}$ using the same n samples from π as those used in (17) to estimate the map coefficients $\hat{\alpha}$. Reusing the samples produces a biased approximation of the non-zero values of $\widehat{\Omega}$ at finite n . We observe in our numerical experiments, however, that this bias has no significant impact on the estimation of the sparsity pattern of Ω .

4.2 Threshold estimator of Ω

As a result of the sample-based approximations of S and $\widehat{\Omega}$, the sparsity pattern of $\widetilde{\Omega}$ will not exactly match that of Ω . For instance, a zero entry in Ω may result in a small but numerically non-zero entry in $\widetilde{\Omega}$. To account for this mismatch, we introduce a threshold estimator $\overline{\Omega}$ defined as

$$\overline{\Omega}_{ij} = \begin{cases} \widetilde{\Omega}_{ij}, & \widetilde{\Omega}_{ij} \geq \tau_{ij} \\ 0, & \text{otherwise} \end{cases}, \quad (22)$$

for some $\tau_{ij} > 0$. Threshold estimators are commonly used for sparse covariance matrix estimation (see Cai and Liu (2011)) and τ_{ij} is usually chosen proportional to the standard deviation of $\widetilde{\Omega}_{ij}$. The rationale behind this choice is to threshold the entries of $\widetilde{\Omega}$ whose standard deviation makes them indistinguishable from zero.

To compute the standard deviation or variance of $\widetilde{\Omega}_{ij}$, empirical estimation is not feasible since we only have a unique realization of $\widetilde{\Omega}_{ij}$. Instead we approximate its variance with

$$\mathbb{V}(\widetilde{\Omega}_{ij}) \approx \frac{1}{n} \left(\nabla_{\alpha} \widetilde{\Omega}_{ij} \right)^T \Gamma(\alpha)^{-1} \left(\nabla_{\alpha} \widetilde{\Omega}_{ij} \right) \Big|_{\alpha=\hat{\alpha}} := \frac{\widetilde{v}_{ij}^2}{n}, \quad (23)$$

where $\Gamma(\alpha)$ is the Fisher information in (19) and $\nabla_{\alpha} \widetilde{\Omega}_{ij}|_{\alpha=\hat{\alpha}}$ denotes the gradient of $\alpha \mapsto \widetilde{\Omega}_{ij}(\alpha) = \frac{1}{n} \sum_{l=1}^n |\partial_i \partial_j \log S_{\alpha}^\# \eta(\mathbf{z}^l)|^2$ evaluated at $\hat{\alpha}$. We assume here that $\alpha \mapsto \widetilde{\Omega}_{ij}(\alpha)$

is a continuously differentiable function of the parameters $\boldsymbol{\alpha}$ for each entry (i, j) . The variance approximation in (23) is inspired by the delta method (Oehlert, 1992), which exploits the fact that, given a sequence of random variables $\theta_n \in \mathbb{R}^p$ satisfying $\sqrt{n}(\theta_n - \theta) \xrightarrow{D} \mathcal{N}(\mathbf{0}, \Lambda)$ and a continuously differentiable function $g: \mathbb{R}^p \rightarrow \mathbb{R}$ such that $\nabla g(\theta) \neq 0$, we have $\sqrt{n}(g(\theta_n) - g(\theta)) \xrightarrow{D} \mathcal{N}(0, \nabla g(\theta)^T \Lambda \nabla g(\theta))$. Here we consider θ_n to be the map coefficients $\widehat{\boldsymbol{\alpha}}$ and g to be the sample-based score estimator $\widetilde{\Omega}_{ij}$. The main difference between (23) and the variance predicted by the delta method is that we evaluate the gradients and the Fisher information at the estimated map coefficients $\widehat{\boldsymbol{\alpha}}$ because the limiting coefficients $\boldsymbol{\alpha}^*$ are unknown in practice.

For the threshold in (22), we then use

$$\tau_{ij} = f(n) \frac{\widetilde{v}_{ij}}{\sqrt{n}}, \quad (24)$$

where $f(n)$ is a function that increases slowly with n . In our numerical experiments, we will choose $f(n) \propto \sqrt{\log n}$. In Section 4.4, we will justify this choice by identifying a class of functions f that make the resulting threshold estimator $\overline{\Omega}$ a *consistent* estimator of the sparsity pattern of Ω as $n \rightarrow \infty$.

4.3 Non-iterative SING algorithm

The tools above suffice to build an algorithm to learn the Markov structure of a continuous, non-Gaussian distribution. Algorithm 1 learns the graph structure from the support of the threshold estimator $\overline{\Omega}$ for the conditional independence score. We refer to this sequence of steps as the non-iterative Sparsity Identification in Non-Gaussian distributions (N-SING) algorithm. In Section 5 we propose an iterative version of this algorithm that uses the sparsity of $\overline{\Omega}$ to define improved estimators for the map and the resulting score matrix Ω .

Algorithm 1: Non-iterative Sparsity Identification in Non-Gaussian distributions

(N-SING)

-
- Input** : i.i.d. sample $\{\mathbf{z}^l\}_{l=1}^n \sim \pi$, maximum polynomial degree β , threshold scaling function f
- Output** : Edge set \widehat{E}_n of minimal I-map for $\boldsymbol{\nu}_\pi$
- 1 Compute transport map: $S_{\widehat{\boldsymbol{\alpha}}} = \operatorname{argmax}_{S_{\boldsymbol{\alpha}} \in \mathcal{S}_{\Delta}^{\beta}} \sum_{l=1}^n \log S_{\boldsymbol{\alpha}}^{\sharp} \eta(\mathbf{z}^l)$
 - 2 Estimate Ω : $\widetilde{\Omega}_{ij} = \frac{1}{n} \sum_{l=1}^n |\partial_i \partial_j \log S_{\boldsymbol{\alpha}}^{\sharp} \eta(\mathbf{z}^l)|^2$
 - 3 Threshold $\widetilde{\Omega}$, with $\tau_{ij} = f(n) \widetilde{v}_{ij} / \sqrt{n}$, to yield $\overline{\Omega}$
 - 4 Compute \widehat{E}_n : $(i, j) \in \widehat{E}_n$ if $\overline{\Omega}_{ij} \neq 0$
-

4.4 Analysis of consistency

Here we establish conditions under which the proposed threshold estimator is consistent for recovering the edge set E of the minimal I-map for $\boldsymbol{\nu}_\pi$. For simplicity, we consider a variant

of N-SING that uses an exact expectation $\widehat{\Omega}_{ij} = \mathbb{E}_\pi |\partial_i \partial_j \log S_{\widehat{\alpha}}^\# \eta(\mathbf{Z})|^2$ in step 2 of Algorithm 1. Using the finite-sample approximation $\widehat{\Omega}_{ij}$ would complicate but not fundamentally change the following analysis. We also assume that the map parameterization is sufficiently rich to recover the target density exactly, i.e., that there exists a map $S_{\alpha^*} \in \mathcal{S}_\Delta^\beta$ with coefficients $\alpha^* \in \mathbb{R}^p$, for some polynomial degree β , such that $\pi = (S_{\alpha^*})^\# \eta$, where α^* is obtained from the solution of (16). Intuitively, since the maximum likelihood estimate $\widehat{\alpha}$ converges to α^* as $n \rightarrow \infty$ (recall (18)), each entry $\widehat{\Omega}_{ij}$ of the estimated score matrix should then converge to the true score Ω_{ij} :

$$\widehat{\Omega}_{ij} = \mathbb{E}_\pi |\partial_i \partial_j \log S_{\widehat{\alpha}}^\# \eta(\mathbf{Z})|^2 \xrightarrow[n \rightarrow \infty]{} \mathbb{E}_\pi |\partial_i \partial_j \log S_{\alpha^*}^\# \eta(\mathbf{Z})|^2 = \Omega_{ij}.$$

We will formalize this notion within the analysis below.

To recover the support of Ω from $\widehat{\Omega}$, we consider the threshold estimator

$$\overline{\Omega}_{ij} = \widehat{\Omega}_{ij} \mathbb{1}(\widehat{\Omega}_{ij} > \tau_{ij}), \quad \text{with } \tau_{ij} = f(n) \widehat{v}_{ij} / \sqrt{n}, \quad (25)$$

where $\widehat{v}_{ij}^2 := (\nabla_{\alpha} \widehat{\Omega}_{ij})^T \Gamma(\alpha)^{-1} (\nabla_{\alpha} \widehat{\Omega}_{ij})|_{\alpha=\widehat{\alpha}}$. A proper choice of f is critical to guaranteeing that the support of the threshold estimator $\overline{\Omega}$ converges to the support of Ω with increasing n . For instance, when $\Omega_{ij} = 0$, both $\widehat{\Omega}_{ij}$ and $\widehat{v}_{ij} / \sqrt{n}$ go to zero at the same rate as $n \rightarrow \infty$. As a result, the event $\{\widehat{\Omega}_{ij} > \widehat{v}_{ij} / \sqrt{n}\}$ asymptotically occurs with a constant non-zero probability, resulting in false positive edges. The role of $f(n)$ in this case is to adjust the rate of convergence of the threshold to ensure that $\widehat{\Omega}_{ij} < \tau_{ij}$ asymptotically, i.e., that there are no false positives. A similar argument holds for false negatives. The following proposition gives sufficient conditions on f to guarantee the recovery of the edge set in the minimal I-map for ν_π . The proof is provided in Appendix C.

Proposition 9 *For the threshold estimator (25), let f be a function such that $f(n) \rightarrow \infty$ and $f(n)/\sqrt{n} \rightarrow 0$ as $n \rightarrow \infty$. Then the edge set \widehat{E}_n returned by the associated N-SING algorithm is a consistent estimator of E , i.e., $\mathbb{P}(\widehat{E}_n = E) \rightarrow 1$ as $n \rightarrow \infty$.*

The main idea behind this result is to show that the probability of false positives (i.e., type 1 errors) or false negatives (i.e., type 2 errors) converges to zero with the prescribed threshold. As demonstrated above, the occurrence of these events is determined by the magnitude of the estimated score $\widehat{\Omega}_{ij}$ relative to the threshold τ_{ij} , or, equivalently, by the magnitude of the ratio $\sqrt{n} \widehat{\Omega}_{ij} / \widehat{v}_{ij}$ in comparison to $f(n)$. For each pair (i, j) , we consider the function $g(\alpha) = \mathbb{E}_\pi |\partial_i \partial_j \log S_{\alpha}^\# \eta(\mathbf{z})|^2$ and analyze the asymptotic statistics of the ratio $\sqrt{n} \widehat{\Omega}_{ij} / \widehat{v}_{ij} = \sqrt{n} g(\widehat{\alpha}) / (\nabla_{\alpha} g(\widehat{\alpha})^T \Gamma(\widehat{\alpha})^{-1} \nabla_{\alpha} g(\widehat{\alpha}))^{1/2}$ as a function of the estimated map coefficients $\widehat{\alpha}$. When the variables Z_i and Z_j are conditionally dependent, we have $g(\alpha^*) = \Omega_{ij} \neq 0$, and we assume that the gradient $\nabla_{\alpha} g(\alpha^*) \neq 0$. Then, the limiting distribution of the ratio is Gaussian by an application of the delta method. On the other hand, when Z_i and Z_j are conditionally independent, we have not only $g(\alpha^*) = \Omega_{ij} = 0$ but also $\nabla_{\alpha} g(\alpha^*) = 0$, because both of these terms depend on $\partial_i \partial_j \log \pi(\mathbf{z})$, which is zero for all $\mathbf{z} \in \mathbb{R}^d$. Thus both $\widehat{\Omega}_{ij}$ and \widehat{v}_{ij} approach zero, and their ratio becomes singular as $n \rightarrow \infty$. In this case, the delta method is no longer valid; instead, we must consider the asymptotic distribution of singular Wald statistics, as analyzed in Drton et al. (2016); Pillai and Meng

(2016). Here, to characterize the limiting distribution, we assume that $\nabla_{\alpha}^2 g(\alpha^*) \neq 0$, but higher-order derivatives could also be considered if this condition does not hold. Under the asymptotic distributions for these two scenarios, we show that the probabilities of false positives and false negatives converge to zero given the conditions above on the function f in (25).

While Proposition 9 guarantees the threshold estimator is consistent for *any* f satisfying the criteria above, the selected f may affect the algorithm’s finite-sample performance. A function f that grows more quickly with n will produce higher thresholds and reduce the probability of false positive edges, while a more slowly growing function will produce lower thresholds and reduce the probability of false negative edges. Future work will investigate the impact of these choices on finite-sample bounds, i.e., the number of samples required by the algorithm to recover the graph with high probability.

5. Improved estimator for the Markov structure

In Section 4 we showed how to estimate the conditional independence score using a transport map representation of the target density. This transport map is a lower-triangular function and thus is a sparse map where each component does not depend on all of its input variables. However, when the target measure ν_{π} satisfies conditional independence properties, the transport map we seek to approximate can inherit additional sparse structure (see Spantini et al., 2018, Section 5). In this Section, we take advantage of this connection between the sparsity of the graph and the sparsity of the transport map to present an iterative algorithm for learning the graph.

5.1 Sparsity of the transport map

For a lower triangular function S , the sparsity pattern of the map, \mathcal{I}_S , is defined in Spantini et al. (2018) as:

$$\mathcal{I}_S := \{(j, k) : j < k, \partial_j S^k = 0\}. \quad (26)$$

That is, the sparsity pattern is the set of all integer pairs (j, k) with $j < k$, such that the k th component of the map does not depend on the j th input variable.⁴ The complement of this set, i.e.,

$$\mathcal{I}_S^c := \{(j, k) : j < k, \partial_j S^k \neq 0\}, \quad (27)$$

determines the *active variables* of the map. That is, if $(j, k) \in \mathcal{I}_S^c$, then the k th component of the map must depend on the j th input variable. We denote the set of lower triangular maps that respect the sparsity pattern given by \mathcal{I}_S as $\mathcal{S}_{\mathcal{I}_S} \subset \mathcal{S}_{\Delta}$.

Given a target density π , Spantini et al. (2018) showed that the Markov structure of ν_{π} yields a *tight lower bound* on the sparsity pattern \mathcal{I}_S of the KR rearrangement that pulls back η to π . Knowledge of this sparsity can be used when solving the variational problem in (16) by restricting the feasible domain to transport maps with a reduced set of active variables. To determine this sparsity pattern, we perform a series of graph operations on the minimal I-map \mathcal{G} of the target measure ν_{π} . These operations define the active variables for each map component based on a sequence of intermediate graphs (\mathcal{G}^k) . The graph \mathcal{G}^k is

4. The lower triangular function also satisfies $\partial_j S^k = 0$ for all $j > k$ by construction.

identical to the graph obtained in the variable elimination algorithm before marginalizing node k according to the elimination ordering $(d, d-1, \dots, 1)$. However, we emphasize that this sparsity pattern is identified only by inspecting the graph, without actually performing variable elimination or additional computation (e.g., marginalization) on the joint density. We restate the relevant part of this result below.⁵

Theorem 10 (Spantini et al. (2018), Theorem 3 (Part 1)) *Let $\mathbf{X} \sim \nu_\eta$, $\mathbf{Z} \sim \nu_\pi$ with Lebesgue absolutely continuous measures ν_η , ν_π , and let ν_η be a product measure on \mathbb{R}^d . Moreover, assume that ν_π is globally Markov with respect to \mathcal{G} , and define, recursively, the sequence of graphs $(\mathcal{G}^k)_{k=1}^d$ as: (1) $\mathcal{G}^d := \mathcal{G}$ and (2) for all $1 \leq k < d$, \mathcal{G}^{k-1} is obtained from \mathcal{G}^k by removing node k and by turning its neighborhood $Nb(k, \mathcal{G}^k)$ into a clique. Then the following holds:*

1. If \mathcal{I}_S is the sparsity pattern of the transport map S , then

$$\widehat{\mathcal{I}}_S \subset \mathcal{I}_S, \tag{28}$$

where $\widehat{\mathcal{I}}_S$ is the set of integer pairs (j, k) such that $j \notin Nb(k, \mathcal{G}^k)$.

Remark 11 *The sequence of elimination steps performed on the graph in Theorem 10 are equivalent to the steps taken by the symbolic Cholesky factorization of a sparse symmetric matrix to identify the nonzero structure of its triangular Cholesky factor; see Lülkesmann et al. (2010) for more details on this equivalence. Thus, the sparsity bound $\widehat{\mathcal{I}}_S$ identified by Theorem 10 corresponds to the symbolically predicted sparsity of the Cholesky factor of a symmetric matrix whose (j, k) th entry is zero if $Z_j \perp\!\!\!\perp Z_k | \mathbf{Z}_{-jk}$.*

5.2 Ordering variables in the map

In the process above, the sparsity pattern of the map decreases (relative to that of the original I-map \mathcal{G}) when adding edges to the intermediate graphs \mathcal{G}^k to create cliques. These edges produce *fill-in*. Fill-in will occur unless \mathcal{G} is chordal *and* the variable ordering corresponds to the perfect elimination ordering (Rose et al., 1976). Whether or not \mathcal{G} is chordal, the amount of fill-in is dependent on the ordering of the input variables. For example, consider the graph and variable ordering in Figure 1a. The corresponding lower bound for the sparsity pattern of the map is

$$\widehat{\mathcal{I}}_S = \{(1, 4), (2, 4), (1, 5), (2, 5), (3, 5)\}, \tag{29}$$

and the dependence of each map component on the input variables is

$$S(\mathbf{z}) = \begin{bmatrix} S^1(z_1) \\ S^2(z_1, z_2) \\ S^3(z_1, z_2, z_3) \\ S^4(\quad z_3, z_4) \\ S^5(\quad z_4, z_5) \end{bmatrix}. \tag{30}$$

5. Parts 2 and 3 of Theorem 3 in Spantini et al. (2018) provide sparsity bounds for the transport map S^{-1} , which are related to the marginal independence of ν_π and do not concern the current work.




Figure 1: (a) A sparse graph with an optimal node ordering; (b) Suboptimal ordering induces extra edges.

With the variable ordering shown in Figure 1a, no edges are added during the process to identify $\widehat{\mathcal{L}}_S$ and the resulting transport map is more sparse than a dense lower-triangular map. In contrast, using the suboptimal ordering shown in Figure 1b, edges must be added to the induced graph, shown in dashed lines. The associated sparsity pattern is now $\widehat{\mathcal{L}}_S = \{(1, 5), (2, 5)\}$, and hence the transport map is predicted to be less sparse.

With a larger set $\widehat{\mathcal{L}}_S$, we can simplify the parameterization of the map and reduce the size of the coefficient vector α ; this in turn reduces the variance of the estimator $\widehat{\alpha}$ (17), for any given sample size n (see demonstrations in Morrison et al., 2017, Section 3.2). Thus, it is worthwhile to find a variable ordering that maximizes the sparsity of the map. A variable ordering is equivalent to a permutation $\varphi: [d] \rightarrow [d]$ of the nodes in the graph. Spantini et al. (2018) presented an optimization problem for identifying the permutation that induces the least fill-in—specifically, the largest cardinality of the resulting sparsity bound $|\widehat{\mathcal{L}}_S|$. This optimization problem is in general NP-complete (Yannakakis, 1981). Nevertheless, several heuristics have been proposed to construct permutations based only on the graph \mathcal{G} , including weighted min-fill and min-degree orderings (Koller and Friedman, 2009).

Given a matrix whose sparsity defines the Markov structure of ν_π , an optimal variable ordering can be identified from its sparse Cholesky factorization. From Remark 11, we have that the variable dependence in the transport map identified by Theorem 10 corresponds to the fill-in added to the Cholesky factor L of a sparse matrix $\Sigma^{-1} = L^T L$ encoding the Markov structure. As a result, finding an ordering that minimizes fill-in in the graph can be cast as seeking a permutation matrix P so that $P\Sigma^{-1}P^T = L^T L$ where L is as sparse as possible; see Raskutti and Uhler (2018) for consistency guarantees when solving this ideal problem with estimators of Σ^{-1} .

Remark 12 *In the setting of Gaussian Markov random fields (Rue and Held, 2005), we can take the matrix Σ^{-1} to be the (sparse) precision matrix of the Gaussian distribution $\mathcal{N}(\mathbf{0}, \Sigma)$. Following Remark 4, the lower-triangular Cholesky factor L of Σ^{-1} defines a linear transport map $S(\mathbf{z}) = L\mathbf{z}$ that pulls back the reference distribution $\mathcal{N}(\mathbf{0}, I_d)$ to the target $\mathcal{N}(\mathbf{0}, \Sigma)$. Our goal is to seek a variable ordering that yields the map with sparsest variable dependence, i.e., the Cholesky factor L with the largest number of zero entries.*

For arbitrary non-Gaussian distributions, the sparsity of the conditional independence score after thresholding, $\overline{\Omega}$, provides an estimate of the Markov structure of ν_π . We therefore can find a good ordering by seeking a permutation that induces minimum fill-in in the

symbolic Cholesky factorization of $\bar{\Omega}$. We use standard heuristics to solve this problem, which are implemented in software packages such as CHOLMOD (Davis, 2008). The CHOLMOD package uses both the approximate minimum degree algorithm and the nested dissection algorithm implemented in the METIS graph partitioning library (Karypis and Kumar, 1998) to identify candidate orderings, and then chooses the variable ordering that induces the fewest non-zero entries in the Cholesky factor of $\bar{\Omega}$. We use this ordering scheme for the numerical examples in Section 6.

5.3 The iterative SING algorithm

In Section 5.1 we saw how sparsity in the graph implies sparsity in the transport map. To take advantage of this sparsity, we need to know the Markov structure of the target density π . This Markov structure is not available *a priori* given only samples from π . However, the N-SING algorithm from Section 4 can provide an initial estimate of this Markov structure, based on a dense lower triangular map S (i.e., with $\mathcal{I}_S = \emptyset$). This Markov structure can be used to identify a new variable ordering, and a corresponding bound $\hat{\mathcal{I}}_S$ on the sparsity pattern of the transport map via Theorem 10. We can then enforce this sparsity bound to compute a new estimate of the transport map, and in turn obtain a new threshold estimate of the Markov structure of π . We repeat this procedure until the sparsity of the estimated graph no longer changes. This defines an iterative application of N-SING that we call the SING algorithm.⁶ The formal steps of SING are presented in Algorithm 2.

The first steps of SING are identical to the N-SING algorithm. The rationale for the remaining steps, and for subsequent iterations, is essentially to exploit sparsity for variance reduction. Sparsity in the graph, coupled with a good variable ordering, leads to sparsity in the map, i.e., a smaller number of active variables. As noted above, such a map can be described by fewer coefficients α ; a maximum likelihood estimate of these coefficients—based on the same finite sample from π as in previous iterations—in turn has a smaller variance, as observed in Morrison et al. (2017). Thus, the iterations of the SING algorithm provide improved estimators of the target density and of the conditional independence score for the goal of learning the graph.

Remark 13 *For a multivariate Gaussian target density $\pi = \mathcal{N}(\mathbf{0}, \Sigma)$, the SING algorithm with affine maps (i.e., polynomial degree $\beta = 1$) alternates between estimating a sparse Cholesky factor of the inverse covariance matrix (see Proposition 8 and Remark 4) and defining a threshold estimator for Σ^{-1} to learn the Markov structure of π . The sparsity of the Cholesky factor is dependent on the sparsity of Σ^{-1} and the ordering of the input variables. Recently, several methods have also been proposed to learn sparse Cholesky factors of a sparse inverse covariance matrix for the goal of density estimation in the Gaussian setting. These methods are based on ℓ_1 -penalized maximum likelihood estimation (Huang et al., 2006), banded sparsity patterns for the Cholesky factor (Bickel et al., 2008; Levina et al., 2008), and combinations of multiple variable orderings (Kang and Deng, 2020).*

To conclude, let us comment on the stopping criterion used in Algorithm 2. During the procedure, the edges in the estimated graph can change freely. For instance, we do not impose any constraint on the edge set \hat{E}^t at iteration t (e.g., $\hat{E}^t \subseteq \hat{E}^{t-1}$) or on the sparsity pattern

6. The SING algorithm first appeared with slight modifications in Morrison et al. (2017).

of the map. Nevertheless, we have observed in most of our numerical experiments that the cardinality of estimated edges $|\widehat{E}^t|$ decreases until the algorithm finds a good estimate for E . Thus, checking that $|\widehat{E}^t|$ is non-decreasing works well as a practical stopping criterion.

Algorithm 2: Sparsity Identification in Non-Gaussian distributions (SING)

Input : i.i.d. sample $\{\mathbf{z}^l\}_{l=1}^n \sim \pi$, maximum polynomial degree β , threshold scaling function f

Output : Edge set \widehat{E}_n of minimal I-map for ν_π

Define : $\mathcal{S}_{\widehat{\mathcal{I}}^1}^\beta = \mathcal{S}_\Delta^\beta$, $|\widehat{E}^0| = d(d-1)/2$, $t = 1$

- 1 **while** $|\widehat{E}^t|$ *is decreasing* **do**
- 2 Compute transport map: $S_{\widehat{\alpha}} = \operatorname{argmax}_{S_\alpha \in \mathcal{S}_{\widehat{\mathcal{I}}^t}^\beta} \sum_{l=1}^n \log S_{\widehat{\alpha}}^\# \eta(\mathbf{z}^l)$
- 3 Estimate Ω : $(\widetilde{\Omega}^t)_{ij} = \frac{1}{n} \sum_{l=1}^n |\partial_i \partial_j \log S_{\widehat{\alpha}}^\# \eta(\mathbf{z}^l)|^2$
- 4 Threshold $\widetilde{\Omega}^t$ with $\tau_{ij} = f(n)(\widetilde{v}^t)_{ij}/\sqrt{n}$, to yield $\overline{\Omega}^t$
- 5 Compute $|\widehat{E}^t|$ where $(i, j) \in \widehat{E}^t$ if $(\overline{\Omega}^t)_{ij} \neq 0$
- 6 Find permutation of the variables φ^{t+1} (using, e.g., CHOLMOD)
- 7 Identify sparsity pattern of subsequent map $\widehat{\mathcal{I}}^{t+1}$
- 8 $t \leftarrow t + 1$
- 9 Set $\widehat{E}_n = \widehat{E}^t$

6. Examples

In this section we apply the SING algorithm to learn the Markov structure of several non-Gaussian datasets. Section 6.1 presents results for the butterfly distribution and demonstrates the value of an iterative algorithm for recovering the Markov structure. Section 6.2 applies SING to data from nonparanormal distributions that were considered in Liu et al. (2009). Surprisingly, affine transport map approximations to the target density (i.e., with polynomial degree $\beta = 1$) work well for these nonparanormal examples, even with highly non-Gaussian marginals. In Section 6.3 we use an analytical example to further investigate why a linear map might still work for non-Gaussian data. We then examine a generalization of the nonparanormal setting in Section 6.4, where the target distribution is given by a diagonal transformation of a non-Gaussian base distribution. In these examples, we use the approximation class of the transport map that represents the base density also to approximate the target density, and still recover the correct graph. Finally, in Sections 6.5 and 6.6 we consider two application-oriented examples: a higher-dimensional physics-based dataset arising from the Lorenz-96 dynamical system and a data set that captures interactions in a protein-signaling network. The code to reproduce the results is available at github.com/baptistar/SING.

In all of our numerical experiments, we parameterize the transport maps using the integrated squared representation introduced in Section 3.3, thereby enforcing the monotonicity of each map component by construction. Within the SING algorithm, we use the ordering

heuristics described in Section 5.2 to identify a good permutation at each iteration. Following the analysis in Section 4.4, we set the threshold to $f(n) = c\sqrt{\log n}$ where $c \in \mathbb{R}$ is a constant. While c can be chosen via cross-validation to improve empirical performance, any value is consistent for learning the graph, and we set $c = 1$ in our numerical investigations. Before running SING, we standardize each variable in the dataset by subtracting the empirical mean and dividing by the empirical standard deviation. This normalization step ensures all of the variables are on a similar scale and improves empirical performance.

To quantitatively evaluate the results of the SING algorithm for recovering the Markov structure using samples from π , we measure the errors in the estimated edge sets \widehat{E} (in examples where the true edge set E is available). For each graph, we measure the number of false positives: edges that are in \widehat{E} and not in E (Type 1 errors) and the number of false negatives: edges in E that are not in \widehat{E} (Type 2 errors). In the figures below, we report the mean type 1 and 2 errors across 25 runs of the algorithm with independent batches of samples, as well as the 95% confidence interval for the mean.

6.1 Butterfly distribution

The first example consists of r i.i.d. pairs of random variables (P_i, Q_i) , where:

$$P_i \sim \mathcal{N}(0, 1) \tag{31}$$

$$Q_i = W_i P_i, \quad \text{with } W_i \sim \mathcal{N}(0, 1), \quad W_i \perp\!\!\!\perp P_i, \quad i = 1, \dots, r. \tag{32}$$

One such pair of random variables, or a variation of the above, is a commonly used example to illustrate how two random variables can be uncorrelated but not independent.

Figures 2a–2b show the minimal I-map and corresponding adjacency matrix of the graph for $r = 5$ pairs, with the variables ordered as $P_1, Q_1, \dots, P_5, Q_5$. Figure 2c shows the one- and two-dimensional marginal densities for one pair (P_i, Q_i) . Each one-dimensional marginal is symmetric and unimodal, but the two-dimensional marginal (shown as samples) displays strongly non-Gaussian behavior.

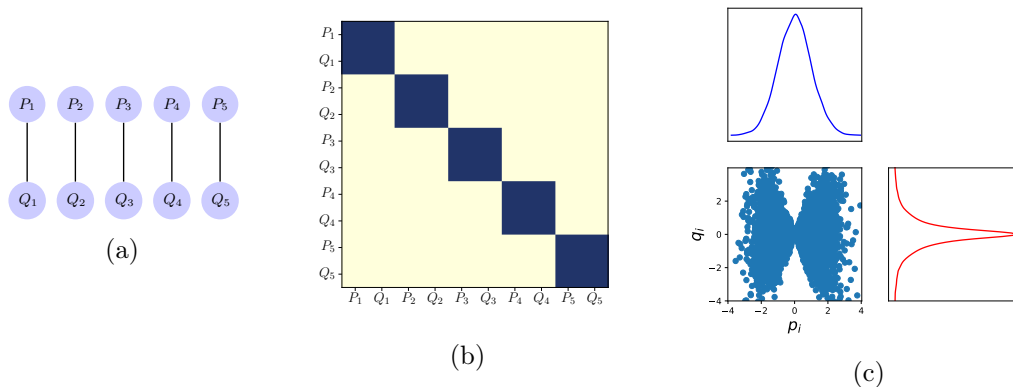


Figure 2: (a) The undirected graphical model; (b) Adjacency matrix of true graph (dark blue corresponds to an edge, off-white to no edge); (c) One- and two-dimensional marginal densities for one pair (P_i, Q_i) .

6.1.1 GRAPH IDENTIFICATION

Figure 3 shows the progression of the identified graph (based on the sparsity of the estimated conditional independence score) over the iterations of SING, with $n = 3000$ samples and a polynomial degree $\beta = 3$. The variables in the data set are initially permuted, to verify that SING identifies a good ordering. After the first iteration of the algorithm (the output of the N-SING algorithm), the estimator for the conditional independence score has the block diagonal pattern in Figure 3b, but the off-diagonals of $\bar{\Omega}$ are not yet zero, resulting in many extra edges. In the next iterations, the algorithm leverages the sparsity of the graph estimated thus far to reveal sparsity in the transport map and improve the estimator for Ω in Figures 3c and 3d, thereby removing all erroneous edges. After the fifth iteration, the sparsity of the graph (and thus the size of the edge set) has not changed and the algorithm returns the the correct graph in Figure 3e.




Figure 3: (a) Adjacency matrix of original graph under random permutation; (b) Entrywise logarithm of the thresholded score matrix $\bar{\Omega}$ (white indicates a zero entry for $\bar{\Omega}$) after iteration 1; (c) Iteration 3; (d) Iteration 5; (e) Adjacency matrix of the graph after iteration 5. SING returns the correct graph with $\beta = 3$.

In contrast, assuming that the data follow a normal or nonparanormal distribution returns the incorrect graph, as illustrated by the experiments in Figure 4. If the data were normal, then a linear map ($\beta = 1$) would suffice. However, SING with $\beta = 1$ (Figure 4(a)) yields not only an incorrect graph, but in fact fails to detect any edges at all. This result does not improve with increasing n , and incorrectly implies that all ten variables are marginally independent. Furthermore, this result indicates that methods based on precision matrix estimation that assume the data to be Gaussian (e.g., GLASSO) will also fail to recover the correct graph in this example, as demonstrated in Morrison et al. (2017). Similarly, algorithms that assume the data to be Gaussian after a (possibly nonlinear) diagonal transformation of the marginals do not recover the correct graph for this dataset; see Figures 4b-4d. These algorithms transform the data marginal using either the empirical cumulative distribution functions (CDFs) in Figure 4b, or rank-based quantities such as the Spearman’s rho and Kendall’s tau statistics (that are only based on the sample ordering and do not require computing CDFs) in Figures 4c and 4d, respectively. After computing these transformations, Liu et al. (2009, 2012a) showed that a GLASSO estimator applied with a correlation coefficient matrix of the transformed data will consistently estimate the graph for nonparanormal data. For the distribution considered in Figure 4 which is outside of this class, however, these algorithms consistently miss the edges between each pair of non-Gaussian variables.

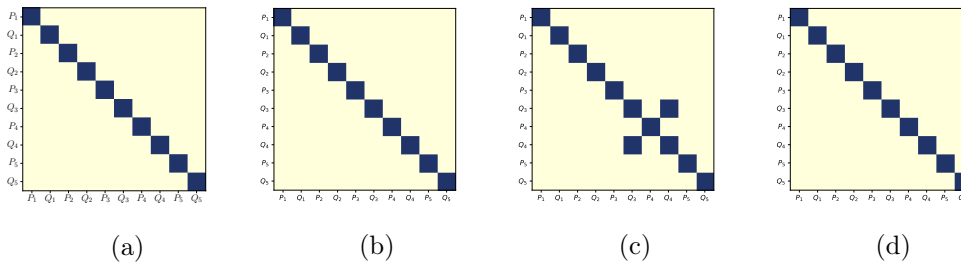


Figure 4: When the data are assumed (a) normal (i.e., using $\beta = 1$) or nonparanormal using (b) empirical CDF marginal transformations, (c) Spearman correlations or (d) Kendall-Tau correlations, all ten variables are incorrectly found to be independent or to have spurious dependencies.

Another important comment on this example is that the underlying density estimation problem is quite challenging: the exact transport map that characterizes π in this case would in fact require an infinite expansion of polynomials. That is, no transport map with finite polynomial degree β can perfectly represent π . However, the correct graph can be identified with $\beta = 3$. Thus, in this case, allowing for some nonlinearity in the map is sufficient. A linear map does not work at all, but β also need not be excessively high.⁷

6.1.2 COMPUTATIONAL COST

Here we provide details on the runtime and memory requirements of applying SING to the butterfly dataset. We have observed that memory usage is relatively consistent across different data sets and computing setups. Runtime can vary more widely, although it tends to be on the order of seconds to just a few minutes.

The tables in Appendix D provide specific memory usage with $n = 1000$ samples for different dataset dimensions d and maximum polynomial degrees β . The tables break down the memory required for: (i) maximum likelihood estimation of the map coefficients, (ii) estimating the score, and (iii) estimating the variance of the score, all for the first two iterations of SING. We observe that variance estimation (see equation (23)) is the main contributor to memory usage. This is not surprising because the variance estimate relies on computing derivatives of $\tilde{\Omega}$ with respect to map parameters, and $\tilde{\Omega}$ already contains second derivatives of the log density with respect its arguments. Evaluations of the resulting derivatives rely on a tensor whose size has quartic growth with respect to the number of variables d . We also note that the first iteration of SING uses the most memory by far. After each iteration, SING may remove active variables from the map, making subsequent iterations less memory intensive. During the first iteration, we do not assume any sparsity of the map or the graph, and hence this iteration is the most costly.

Though our current implementation of the algorithm can be memory intensive, we emphasize that there are no theoretical limits to recovering the exact graph with SING (assuming the transport map class is appropriate); indeed, even the one-step algorithm will recover the correct graph in the limit of infinite samples. (This is in contrast with many

⁷ A tutorial on this example is provided online by Transport Maps Team (2018).

existing methods that make distributional assumptions, usually related to Gaussianity.) Though our present focus is on establishing essential features of the provably consistent algorithm, it is natural to then seek approximations that could lessen computational cost in practice. Some options are to use a lower β than absolutely necessary, to devise further approximations to the variance of $\tilde{\Omega}$, or to exploit prior knowledge that imposes some sparsity on the initial transport map. Another natural extension, which would significantly reduce memory usage, is to devise a *local* version of SING that employs neighborhood selection. We comment further on this possibility in Section 7.

6.2 Nonparanormal data: Gaussian CDF and power transformations

Now we test SING on data from nonparanormal distributions with arbitrary Markov structures. Let $D^k : \mathbb{R} \rightarrow \mathbb{R}$ be a monotone and differentiable transformation for $k = 1, \dots, d$. Following Liu et al. (2009), we say that $\mathbf{Z} = (Z_1, Z_2, \dots, Z_d)$ has a nonparanormal distribution with measure ν_π and density π when $D(\mathbf{Z}) = (D^1(Z_1), D^2(Z_2), \dots, D^d(Z_d))$ follows a multivariate Gaussian distribution with measure ν_ρ and density $\rho = \mathcal{N}(\mathbf{0}, \Sigma)$. We refer to this Gaussian as the “base” distribution in the following subsections. The transformation D acts component-wise on each variable, and thus defines a diagonal transport map that pushes forward π to ρ (or equivalently D pulls back ρ to π). From the result in Proposition 5 for diagonal maps, the minimal I-maps of ν_π and of ν_ρ are equivalent. Thus the Markov structure of the target distribution is prescribed by the sparsity of the precision matrix Σ^{-1} of ρ , but the data can have very non-Gaussian features (see Figure 5b for one example). Within the field of structure learning from non-Gaussian data, this type of nonparanormal distribution (also known as a Gaussian copula) is a test case for algorithms that handle non-Gaussianity.

To generate data from a nonparanormal distribution with an arbitrary graph structure, we follow the steps in Liu et al. (2009). First, a random sparse graph $\mathcal{G} = (V, E)$ is generated. For each node $i \in (1, \dots, d)$, we associate a pair of random variables $(Y_i^{(1)}, Y_i^{(2)}) \in [0, 1]^2$ to i where

$$Y_1^{(l)}, Y_2^{(l)}, \dots, Y_d^{(l)} \sim \mathcal{U}[0, 1] \quad (33)$$

for $l = 1, 2$. Then, each pair of nodes (i, j) is included in the edge set E with probability

$$\mathbb{P}((i, j) \in E) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\|y_i - y_j\|_2^2}{2s}\right), \quad (34)$$

where s is a parameter that controls the sparsity of the graph, $y_k \equiv (y_k^{(1)}, y_k^{(2)})$ is a sample of $(Y_k^{(1)}, Y_k^{(2)})$, and $\|\cdot\|_2$ represents the Euclidean norm. In our numerical experiments we set $s = 3$ and limit the maximum degree of the graph, i.e., the number of the edges incident to each node, to be four. A realization of a graph generated according to this procedure is shown in Figure 5a. After defining the graph, the entries of the inverse covariance Σ^{-1} are given by:

$$\Sigma_{ij}^{-1} = \begin{cases} 1 & i = j \\ 0.245 & (i, j) \in E \\ 0 & \text{otherwise.} \end{cases} \quad (35)$$

We note that with maximum degree four, the value 0.245 ensures that the inverse covariance matrix is positive definite by the Gershgorin circle theorem.

To sample from ν_π , we generate i.i.d. samples \mathbf{x}^l from ρ and apply the inverse diagonal transformation to generate i.i.d. samples $\mathbf{z}^l = D^{-1}(\mathbf{x}^l)$ from π . In this work we consider two possibilities for the function D as in Liu et al. (2009): a scaled Gaussian CDF and a power transformation. We now detail how we construct the two functions D^k .

Gaussian CDF transformation. Let $f_0: \mathbb{R} \rightarrow [0, 1]$ be the univariate Gaussian CDF with mean μ_{f_0} and standard deviation σ_{f_0} , i.e.,

$$f_0(t) = \Phi\left(\frac{t - \mu_{f_0}}{\sigma_{f_0}}\right), \quad (36)$$

where Φ is the univariate standard Gaussian CDF. The inverse CDF transformation $F^k = (D^k)^{-1}$ applied to the k th variable is defined as:

$$F^k(x_k) = \sigma_k \left(\frac{f_0(x_k) - \int f_0(t) \Phi'\left(\frac{t - \mu_k}{\sigma_k}\right) dt}{\sqrt{\int \left(f_0(y) - \int f_0(t) \Phi'\left(\frac{t - \mu_k}{\sigma_k}\right) dt\right)^2 \Phi'\left(\frac{y - \mu_k}{\sigma_k}\right) dy}} \right) + \mu_k. \quad (37)$$

In our experiments we apply the same transformation to each marginal by setting $\mu_{f_0} = 0.05$, $\sigma_{f_0} = 0.4$, $\mu_k = 0$ and $\sigma_k = \sqrt{\Sigma_{kk}}$. A representative marginal PDF of π is shown (as a histogram) in Figure 5b. Each marginal displays very non-Gaussian behavior as a result of the nonlinearity in the function F^k .




Figure 5: (a) The true minimal I-map of ν_π . (b) A histogram of one of the marginals from the nonparanormal dataset. In this case, the data has bimodal, non-Gaussian behavior.

Power transformation. Let $f_0(t) = \text{sign}(t)|t|^a$, for $a > 0$. The inverse power transformation $F^k = (D^k)^{-1}$ applied to the k th variable is defined as

$$F^k(x_k) = \sigma_k \left(\frac{f_0(x_k - \mu_k)}{\sqrt{\int (f_0(t - \mu_k))^2 \Phi\left(\frac{t - \mu_k}{\sigma_k}\right) dt}} \right) + \mu_k. \quad (38)$$

We set μ_k and σ_k as above in the CDF transformation, and set $a = 3$.

Surprisingly, the conditional independence properties of these nonparanormal distributions are recovered with a linear map. As seen in Figure 6a for the Gaussian CDF transformation, we recover the correct graph with $n = 3000$ samples and a polynomial degree $\beta = 1$. Figure 6b shows that the correct graph is also returned with $\beta = 2$. While a transport map with higher polynomial degree could be used, a biased approximation to π based on a $\beta = 1$ map in this case is sufficient for learning the graph. We note that the dominant entries of the inverse of the empirical covariance matrix of the data, shown in Figure 6c, for this example also reveal the true graph—just as the sparsity of the precision matrix would in the Gaussian case. The next subsection will investigate this connection further. But the inverse of the empirical covariance contains many noisy and non-zero entries as compared to the final score matrix $\bar{\Omega}$, as shown in Figure 6d; the latter benefits from the thresholding process in SING and the resulting sparsification of the transport map used to estimate the target density.




Figure 6: (a) Recovered graph with $\beta = 1$. The correct graph is returned. (b) The recovered graph with $\beta = 2$ is also correct, but Hermite functions of polynomial degree two in the transport map are unnecessary. (c) Entrywise logarithm of the inverse of the empirical covariance matrix of the data. (d) Entrywise logarithm of $\bar{\Omega}$ for $\beta = 1$.

Figure 7 shows the errors made by the SING algorithm versus the sample size n for both the Gaussian CDF and the power transformation, with $\beta = 1$. For both transformations, the number of type 1 errors (i.e., erroneous edges that are not present in the true graph) is always zero for all n , and the number of type 2 errors (i.e., undetected edges that are present in the true graph) decreases to zero as n increases.

6.3 Nonparanormal data: cubic transformation

In this section we investigate the effect of biased approximations to the target density when learning the graph. We consider a simplified form of the power transformation applied to a 3-dimensional Gaussian vector $\mathbf{X} \in \mathbb{R}^3$. Let $\mathbf{X} \sim \rho = \mathcal{N}(\mathbf{0}, \Sigma_\rho)$ where the precision and the covariance matrix are given by

$$\Sigma_\rho^{-1} = \begin{bmatrix} 1 & 0.2 & 0 \\ 0.2 & 1 & 0.2 \\ 0 & 0.2 & 1 \end{bmatrix}, \quad \Sigma_\rho = 0.92 \begin{bmatrix} 0.96 & -0.2 & 0.04 \\ -0.2 & 1 & -0.2 \\ 0.04 & -0.2 & 0.96 \end{bmatrix} \quad (39)$$

From the sparsity pattern of the precision matrix, the random variables satisfy $X_1 \perp\!\!\!\perp X_3 | X_2$. The corresponding graph is a chain connecting nodes X_1 to X_2 and X_2 to X_3 .

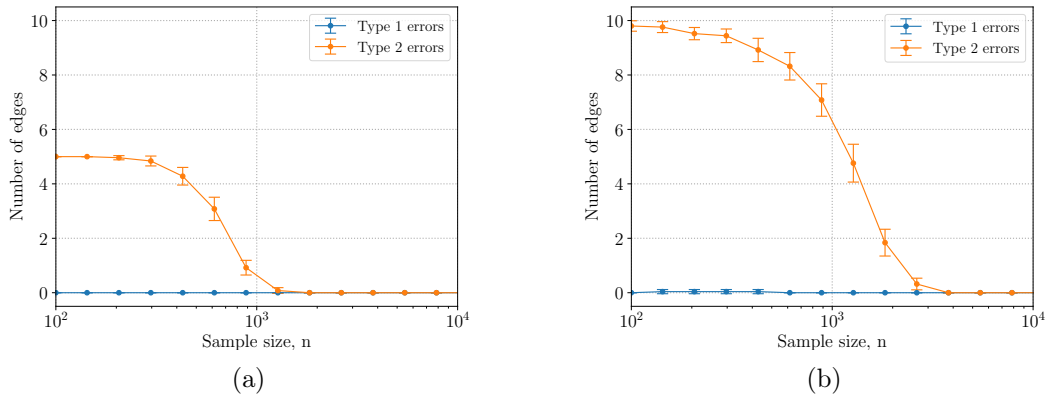


Figure 7: Type 1 and 2 errors for recovering the graph of a nonparanormal distribution versus sample size n for the (a) Gaussian CDF, and (b) power transformation using $\beta = 1$.

We now consider a component-wise monotonic transformation (i.e., a diagonal transport map) given by $F^k(x_k) = x_k^3$ for $k = 1, 2, 3$. The transformed random variable $\mathbf{Z} = F(\mathbf{X}) = (F^1(X_1), F^2(X_2), F^3(X_3))$ is non-Gaussian with density $\pi = F_{\#}\rho$. By the result of Proposition 6 for diagonal transformations, π satisfies the conditional independence property $Z_1 \perp\!\!\!\perp Z_3 | Z_2$, and the Markov structure of π is equivalent to the Markov structure of ρ .

To characterize the target density, suppose we approximate π by the pullback of a standard Gaussian density through a triangular transport map with polynomial degree $\beta = 1$. Using Proposition 8 with $n \rightarrow \infty$, the approximate density is a multivariate Gaussian distribution $\mathcal{N}(\mu_\pi, \Sigma_\pi)$ with mean $\mu_\pi = \mathbb{E}_\pi[\mathbf{Z}] = \mathbb{E}_\rho[F(\mathbf{X})] = \mathbf{0}$ and covariance $\Sigma_\pi = \mathbb{E}_\pi[\mathbf{Z}\mathbf{Z}^T] = \mathbb{E}_\rho[F(\mathbf{X})F(\mathbf{X})^T]$. Each entry of the covariance matrix can be expressed analytically in terms of $\sigma_{ij} = (\Sigma_\rho)_{ij}$ as

$$(\Sigma_\pi)_{ij} = 9\sigma_{ii}\sigma_{jj}\sigma_{ij} + 6\sigma_{ij}^3, \quad (40)$$

since each entry is in fact a higher-order moment of the multivariate Gaussian distribution ν_ρ . As a result, we can also analytically compute the inverse of Σ_π , which is given by

$$\Sigma_\pi^{-1} = \begin{bmatrix} 5.96 \times 10^{-2} & 6.99 \times 10^{-3} & -5.56 \times 10^{-4} \\ 6.99 \times 10^{-3} & 5.36 \times 10^{-2} & 6.99 \times 10^{-3} \\ -5.56 \times 10^{-4} & 6.99 \times 10^{-3} & 5.96 \times 10^{-2} \end{bmatrix}, \quad (41)$$

up to three significant digits. The (1,3) entry of Σ_π^{-1} is not zero, so in principle a Gaussian approximation to π will *not* recover the correct graph. The (1,3) entry is still quite small, however—and, importantly, the relative magnitudes of entries in Σ_π^{-1} are similar to those in Σ_ρ^{-1} (see Morrison et al. (2022) for a more theoretical explanation for why these entries are still small). Thus, when using a numerical approximation, the small (1,3) entry can easily be thresholded and set to zero. And after this thresholding, the correct graph is in fact returned.

Figures 8a and 8b investigate this phenomenon, by showing the errors made by SING for sample sizes $n \in [10^2, 10^6]$ using $\beta = 1$ and $\beta = 2$, respectively. The correct graph is

returned for a broad range of sample sizes. This is a case where the finite-sample estimators of S and Ω and an under-parameterized map (i.e., a biased approximation to π) interact in a surprisingly beneficial way to correctly learn the graph. We believe the same phenomenon explains the success of the linear $\beta = 1$ transport maps in the previous subsection. When the sample size becomes large enough to resolve the smallest entries in the precision matrix (41) with sufficiently high confidence, however, we observe in Figure 8 that the SING algorithm with $\beta = 1$ also includes $(1, 3)$ in the estimated edge set of the graph. A similar trend is observed for $\beta = 2$. Note that the class of $\beta = 2$ transport maps is still insufficient to exactly capture π in this example (the exact transport map would be the composition of a linear map with a diagonal map that applies component-wise cubic root transformations). For sufficiently large n , the bias for $\beta = 2$ yields the type 1 errors seen at the right of Figure 8b.




Figure 8: Type 1 and 2 errors for recovering the graph of the non-paranormal distribution with the cubic transformation versus sample size n using (a) $\beta = 1$, and (b) $\beta = 2$.

6.4 Diagonal transformations of a non-Gaussian base distribution

In this example, we consider the class of probability distributions with measure ν_π defined as pullbacks of a *non-Gaussian* base measure ν_ρ through a nonlinear diagonal transport map. When the base measure ν_ρ is represented as the pullback of a standard Gaussian measure through a triangular transport map that uses polynomials of maximum degree β , we refer to ν_ρ as a distribution in base class β . The nonparanormal distributions in Sections 6.2–6.3 are a subset of this class, where ρ is a multivariate Gaussian density (i.e., ρ is in base class $\beta = 1$). For $\beta > 1$, the class considered here is a *generalization* of the nonparanormal family of distributions.

As an example, we consider a distribution ν_ρ in base class $\beta = 2$. Its density ρ is given by the pullback of a standard Gaussian density η through a sparse transport map S of the form

$$S^1(x_1) = ax_1, \quad S^k(x_1, x_k) = (x_1^2 + b) + ax_k \quad \text{for } k = 2, \dots, d, \quad (42)$$

where we set the parameters $a = 1$ and $b = 1$ to adjust the moments of the distribution $S^\# \nu_\eta$. The transformed random variable $\mathbf{X} = S^{-1}(\mathbf{Y})$ for $\mathbf{Y} \sim \eta$ has the density $S^\# \eta$ and its Markov structure is displayed in Figure 9 for a $d = 5$ dimensional problem. The star

graph associated with the conditional independence structure of ρ is a commonly used graph benchmark for structure learning algorithms (Jalali et al., 2011).



Figure 9: Markov structure and sparsity pattern \mathcal{I}_S of the transport map S for the $d = 5$ dimensional pullback density $S^\# \eta$.

After defining the base density ρ , we apply the nonlinear inverse CDF transformation in (37) to each component of \mathbf{X} in order to define the random variable $\mathbf{Z} = D^{-1}(\mathbf{X})$. Our target density is that of \mathbf{Z} : the pullback of a standard Gaussian density η through the composition of S and the diagonal map D , which we denote as $\pi = (S \circ D)^\# \eta$. To sample from π , we generate i.i.d. samples \mathbf{y}^l from the standard Gaussian reference density η and apply the composition of the inverse maps $D^{-1} \circ S^{-1}: \mathbb{R}^d \rightarrow \mathbb{R}^d$ to each sample, thus generating i.i.d. samples $\mathbf{z}^l = D^{-1} \circ S^{-1}(\mathbf{y}^l)$.

To learn the graph structure of π , we run the SING algorithm with $\beta = 2$ using $n = 10^4$ samples. The true graph and the recovered graph are displayed in Figures 10a and 10b, respectively. The graph structure is learned correctly. On the other hand, Figure 10c shows that using $\beta = 1$ does not recover the true graph of π .




Figure 10: (a) The true graph structure. (b) The recovered graph with $\beta = 2$ is correct. (c) Recovered graph with $\beta = 1$ is incorrect.

To demonstrate the sample sizes needed to learn the true graph in this example, we run the SING algorithm for sample sizes $n \in [10^2, 10^5]$. Figures 11a and 11b display the average type 1 and type 2 errors in the estimated graphs for $\beta = 1$ and $\beta = 2$, respectively, at each sample size n . For $n \geq 10^4$ samples roughly, the number of type 1 and type 2 errors using $\beta = 2$ remains close to zero and represents successful graph recovery. In contrast, Figure 11a shows that using $\beta = 1$, there is no sample size within the tested range where one can recover the exact graph. The conclusion here is analogous to the nonparanormal case of subsection 6.2: the SING algorithm can return the correct Markov structure when the β

parameter matches the polynomial degree necessary to represent the base distribution and not the target.




Figure 11: Type 1 and 2 errors for recovering the graph of a nonlinear marginal transformation of a *non-Gaussian* base distribution, versus sample size n , using (a) $\beta = 1$, and (b) $\beta = 2$.

6.5 Lorenz-96 dynamical system

We now apply SING to a physics-based dataset generated by a chaotic dynamical system. In particular, we consider the Lorenz-96 model, which is commonly used to represent features of the atmosphere (e.g., temperature or vorticity) on a mid-latitude circle of the Earth (Lorenz, 1996). The state at d discrete locations represents the discretization of a spatially *periodic* domain, described by the state vector $\mathbf{Z}(t) = (Z_1(t), \dots, Z_d(t)) \in \mathbb{R}^d$ at time t . The evolution of this state in time is defined by the set of coupled nonlinear ODEs

$$\frac{dZ_j}{dt} = (Z_{j+1} + Z_{j-2})Z_{j-1} - Z_j + F, \quad j = 1, \dots, d, \quad (43)$$

where $Z_{-1} \equiv Z_{d-1}$, $Z_0 \equiv Z_d$, $Z_1 \equiv Z_{d+1}$. In our experiments we use $d = 15$ and $F = 8$, which leads to chaotic dynamics (Reich and Cotter, 2015).

Our goal is to characterize the conditional independence properties of the invariant distribution of the state; thus we collect data from long-time trajectories of the system. Figure 12a shows a sample trajectory of three consecutive variables, and Figure 12b displays a heat map of all 15 variables. While there appears to be some dependence among neighboring nodes, the dependence structure is certainly not revealed by data visualization alone. Indeed, this example does not have a known graph representing its “true” conditional dependence, which is also the case with other data sets representative of real atmospheric dynamics. Nonetheless, approximating the Markov properties of the state of such systems is important to localization schemes for data assimilation (see Spantini et al. (2022)) and to learning the dynamics directly (see Ott et al. (2002)).

To generate a trajectory, we sample a random initial condition $\mathbf{Z}(0) \sim \mathcal{N}(\mathbf{0}, I_d)$ and use a 4th order Runge-Kutta method with a time step of $\Delta t = 0.01$ for $t \in [0, 1600]$ to approximate the state vector $\mathbf{Z}(k\Delta t)$ of the ODE in (43) as \mathbf{Z}^k , for $k \in \mathbb{N}_0$. To reduce

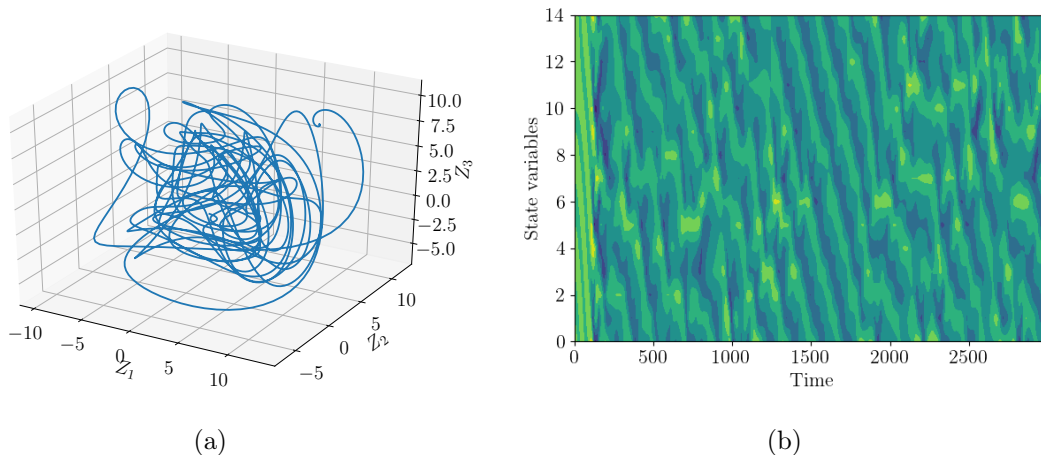


Figure 12: 15-dimensional Lorenz-96 dataset: (a) A sample trajectory of three neighboring state variables (Z_1, Z_2, Z_3). (b) Magnitude of all 15 state variables over time.

correlation between the samples, we sub-sample the trajectory by collecting samples only for $k = 40m$, $m \in \mathbb{N}_0$. Furthermore, to reduce the effect of the initial condition, we discard the first 1000 samples of the trajectory.

Given this dataset, we then learn the Markov structure of the invariant distribution of the Lorenz-96 system. We run the SING algorithm with $\beta = 2$ using $n = 3000$ samples from this dataset. In this problem, the dynamics at each time step introduce local interactions between each state variable Z_j and its neighboring variables on the discretized periodic domain, as seen in the structure of the ODE in (43). As a result, the repeated application of the dynamics results in full dependence amongst the variables. However, the invariant distribution of the system can be well approximated by a Markov random field where each variable conditioned on its closest few neighbors in the physical grid is independent of the others. To account for the weak conditional independence between distant states in the grid, we use a threshold for each entry of the conditional independence score given by $\tau_{ij} = \tau_0 + f(n)\tilde{v}_{ij}/\sqrt{n}$, where $\tau_0 \geq 0$ is a constant offset for all (i, j) . This can be seen as a generalization of the threshold proposed in subsection 4.2 and applied in the previous numerical examples, where we simply used $\tau_0 = 0$. Here we set $\tau_0 = 0.1$.

The thresholded conditional independence score $\bar{\Omega}$ and the corresponding adjacency matrix of the graph found with SING are shown in Figures 13a-13d, for both $\beta = 2$ and $\beta = 1$. To emphasize the decay of the entries in $\bar{\Omega}$ away from the diagonal, we plot the entrywise logarithm of $\bar{\Omega}$. Figure 13a demonstrates the banded dependence of each variable on neighboring variables separated by at most 3 nodes in either direction, as well as the periodic structure of the graph. On the other hand, the SING algorithm with $\beta = 1$ (i.e., a Gaussian approximation to the target density π) entirely misses the conditional dependence of each variable on its immediate neighbors, as seen in Figure 13d.




Figure 13: 15-dimensional Lorenz-96 dataset: Top row: $\beta = 2$; Bottom row: $\beta = 1$. Left column: Logarithm of nonzero entries in $\bar{\Omega}$, at the final iteration of SING; Right column: Adjacency matrix of the graph learned using SING. (d) In particular, the adjacency matrix obtained with $\beta = 1$ incorrectly suggests that each variable Z_j is conditionally independent of its immediate neighbors, Z_{j-1} and Z_{j+1} .

6.6 Cell signaling data

Now we apply SING to learn conditional dependence relationships among the components of a cellular signaling network, based on single-cell data obtained via flow cytometry (Sachs et al., 2005). The data comprise simultaneous measurements of $d = 11$ phosphorylated proteins and phospholipids in $n = 7446$ individual primary human immune system cells. This dataset was used in Sachs et al. (2005) to learn a Bayesian network (i.e., a directed acyclic graph) and in Friedman et al. (2008) to fit a Gaussian graphical model using the GLASSO algorithm. Here we investigate the effect of accounting for non-Gaussianity, via a nonlinear transport map, on the learned dependencies.

As the measurements are all non-negative, we first apply a marginal log transformation to each variable. This pre-processing step transforms the support of the samples from \mathbb{R}_+^d to \mathbb{R}^d . Hence, it makes the transport map pulling back the standard Gaussian reference to the target distribution better behaved and easier to approximate. Note that this diagonal transformation does not change the graph structure; see Section 3.2.

Figures 14a and 14b display the log of the estimated conditional independence score $\bar{\Omega}$ and the adjacency matrix, respectively, from running SING with $\beta = 2$ and a variance threshold based on $f(n) = \delta \log(n)/\sqrt{n}$. Figures 14c and 14d display the corresponding conditional independence score and adjacency matrix with $\beta = 1$, i.e., a linear transport

map. For both map parameterizations, we select $\delta \in \{1.0, 1.25, 1.5, 1.75, 2.0\}$ by performing a 10-fold cross-validation procedure to identify the value which yields, after thresholding, an approximate density $S_{\hat{\alpha}}^{\sharp} \eta$ that maximizes the average log-likelihood on the test sets.

In Figure 14a, and in both Figures 14b and 14d, we observe strong conditional dependence (i.e., larger values for the estimated score $\bar{\Omega}$), and hence the presence of edges, between variable pairs (Z_1, Z_2) , (Z_3, Z_4) , (Z_6, Z_7) and (Z_9, Z_{10}) . These pairs are also connected by edges in the directed graph of Figure 15, whose relationships were verified experimentally by Sachs et al. (2005) and thus can be considered as a *true* reference. For variable pairs with weaker dependence, however, the SING algorithm finds different edge sets when accounting for non-Gaussian structure (with $\beta = 2$) than in the Gaussian case ($\beta = 1$). For example, the non-Gaussian approach identifies edges between Z_8 and the set (Z_6, Z_7) , which are also found in the directed graph by Sachs et al. (2005). In contrast, the Gaussian graphical model fails to identify edges between Z_8 and (Z_6, Z_7, Z_9) . These edges are present when accounting for non-Gaussianity with $\beta = 2$, and they follow from the true directed graph after moralization.




Figure 14: 11-dimensional cell signaling dataset. Top row: $\beta = 2$; Bottom row: $\beta = 1$. Left column: Logarithm of nonzero entries in $\bar{\Omega}$, at the final iteration of SING; Right column: Adjacency matrix of the graph learned using SING.

7. Discussion

This paper develops a framework for learning the Markov structure of continuous non-Gaussian probability distributions from data. The framework is built on two key elements. First, we introduce a computationally tractable score for conditional independence, based on




Figure 15: Directed acyclic graph identified in Sachs et al. (2005) from cell-signaling data. To simplify the presentation, we assign a numeric label to each protein as follows: 1=Raf, 2=Mek, 3=Plc γ , 4=PIP2, 5=PIP3, 6=Erk, 7=Akt, 8=PKA, 9=PKC, 10=P38, and 11=Jnk.

averaged Hessian information from the target log-density. Though this score is immediately useful in its own right, we also show that it bounds the conditional mutual information under appropriate assumptions. Second, we use a transport-based density estimation method, based on parametric approximations of the triangular Knothe–Rosenblatt rearrangement, that explicitly and iteratively exploits sparsity. In particular, the algorithm uses sparsity bounds for triangular maps, which follow directly from the Markov structure; variance-based thresholding for the score estimator; and variable ordering schemes designed to preserve sparsity. Putting these elements together yields the SING algorithm for structure learning. Our analysis shows consistent graph recovery for a single iteration of the algorithm in an asymptotic regime, while numerical results demonstrate the benefits of an iterative algorithm with finite samples.

We demonstrate in multiple examples that this framework can recover the correct conditional independence structure in cases where a Gaussian approximation would yield entirely incorrect results. At the same time, our results show that the map parameterization need not be chosen rich enough to exactly capture the target density. The latter property points to an “information gap,” which reflects the idea that learning the graph structure should require capturing less information than the target distribution itself. Indeed, any undirected graph encodes the Markov properties of an infinite number of distributions.

Put more specifically: there are infinitely many different distributions that will have the same conditional independence score matrix, and even more that will have score matrices with the same sparsity pattern. Moreover, as shown in the analytical example of Section 6.3, entries of the score matrix resulting from a biased approximation to the map—i.e., entries that would be zero in expectation if the density approximation were unbiased—can be an order of magnitude smaller than the non-spurious entries. It would be very useful to understand why this is the case, and under what conditions one can expect a favorable interaction between finite sample sizes, threshold estimators, and small spurious entries in the score matrix. These insights could guide the choice of transport map parameterization for any given data set.

Indeed, the parameterization of the triangular transport is a useful degree of freedom of our overall framework. Polynomials (or the closely related Hermite functions) are convenient as they include the Gaussian setting as a special case, but future work can certainly explore other options, e.g., the nonlinear separable representations in Spantini et al. (2022), or single

layers of autoregressive flows. Further analysis of the information gap described above might make it possible to develop transport map estimators with the *minimal* representation that is needed to learn the Markov structure of the target distribution. Relatedly, it would be valuable to develop information theoretic (*representation-independent*) lower bounds on the number of samples needed to identify the graph in the non-Gaussian setting. To our knowledge, these bounds only exist for Gaussian and discrete Markov random fields (Wang et al., 2010; Santhanam and Wainwright, 2012).

We outline a few additional avenues for future work as follows:

Different data sources. In practical settings, it is of interest to learn the graph from heterogeneous data that may not be independent and identically distributed. Examples include data collected from several related populations that have common structure, and data collected over time. There has been some work to address these problems in the Gaussian setting: Guo et al. (2011); Danaher et al. (2014) propose to combine optimization problems for separate precision matrices with shared ℓ_1 -penalties to identify common sparsity, while Zhou et al. (2010) exploits smooth changes in the precision matrix to estimate the evolution of the graph over time.

Latent structure. We also envision extending the SING algorithm to learn the structure of graphical models with latent variables and partial observations. For Gaussian distributions, Chandrasekaran et al. (2012) proposes a penalized maximum likelihood approach to identify a precision matrix with sparse and low-rank structure. A similar approach may explore other properties of the conditional independence score matrix Ω , besides sparsity, to reveal hidden variables and multiscale structure in the target density.

High-dimensional models. To learn the graph of a d -dimensional distribution, the SING algorithm computes transport maps with d components and stores $\mathcal{O}(d^2)$ entries for the estimated conditional independence score in memory. For high-dimensional datasets, it may not be feasible to jointly estimate the associated graph. Instead, neighborhood selection methods identify local dependencies by independently estimating the neighborhood of each node in \mathcal{G} , i.e., $\text{Nb}(k, \mathcal{G})$ for $k = 1, \dots, d$ such that $\pi(z_k | \mathbf{z}_{-k}) = \pi(z_k | \mathbf{z}_{\text{Nb}(k, \mathcal{G})})$. Methods for finding the neighborhood include greedy selection strategies (Bresler, 2015) and penalized maximum likelihood estimators (Meinshausen and Bühlmann, 2006). Future work will consider a neighborhood selection version of the SING algorithm that avoids estimating the global transport map for the joint density; this would reduce the run time and memory required to learn the graph.

Acknowledgments

We would like to acknowledge support for this project from the National Science Foundation (NSF grant IIS-9988642), the Multidisciplinary University Research Program of the Department of Defense (MURI N00014-00-1-0637), and the INRIA associate team Unquestionable. We also thank Daniele Bigoni and Teo Price-Broncucia for programming support. RB and YM acknowledge support from the Department of Energy, Office of Advanced Scientific Computing Research, AEOLUS (Advances in Experimental design, Optimal control, and Learning for Uncertain complex Systems) center. RB acknowledges support from an NSERC

PGSD fellowship. RM acknowledges Dr. Dusa McDuff and the Johnson&Johnson Foundation for support—financial and otherwise—through the Women in STEM2D Scholars Program.

Appendix A. Proof of Theorem 2

Proof The proof contains two steps: we first show the result when $d = 2$, then extend it to $d > 2$.

Assume $d = 2$. The conditional mutual information $I(Z_i; Z_j | \mathbf{Z}_{-ij})$ becomes the mutual information $I(Z_i; Z_j)$, given by

$$\begin{aligned} I(Z_i; Z_j) &= \int \pi(z_i, z_j) \log \left(\frac{\pi(z_i, z_j)}{\pi(z_i)\pi(z_j)} \right) dz_i dz_j \\ &= \int h(\mathbf{z}) \log \left(\frac{h(\mathbf{z})}{\int h(\mathbf{z}') \pi(z'_i) \pi(z'_j) dz'_i dz'_j} \right) \pi(z_i) \pi(z_j) dz_i dz_j, \end{aligned}$$

where $h(\mathbf{z}) = h(z_i, z_j) = \pi(z_i, z_j) / (\pi(z_i)\pi(z_j))$. Next we apply the logarithmic Sobolev inequality to bound $I(Z_i; Z_j)$. To do that, we need to show that the product density $(z_i, z_j) \mapsto \pi(z_i)\pi(z_j)$ satisfies the logarithmic Sobolev inequality. By Theorem 4.4. in Guionnet and Zegarlinski (2003), $(z_i, z_j) \mapsto \pi(z_i)\pi(z_j)$ has a logarithmic Sobolev constant $C(\pi(z_i)\pi(z_j))$ bounded by $\max\{C(\pi(z_i)); C(\pi(z_j))\}$. By Assumption (8) the joint density $\pi(z_i, z_j)$ has a logarithmic Sobolev inequality bounded by C_0 and therefore the marginals $\pi(z_i)$ and $\pi(z_j)$ satisfy $C(\pi(z_i)) \leq C_0$ and $C(\pi(z_j)) \leq C_0$ (this can be easily shown by restricting (8) to univariate functions $h : (z_i, z_j) \mapsto h(z_i)$ or $h : (z_i, z_j) \mapsto h(z_j)$). We deduce that the product of marginals satisfies the logarithmic Sobolev inequality with constant $C(\pi(z_i)\pi(z_j)) \leq C_0$. We can write

$$\begin{aligned} I(Z_i; Z_j) &\leq \frac{C_0}{2} \int \left\| \begin{matrix} \partial_i \log h(z_i, z_j) \\ \partial_j \log h(z_i, z_j) \end{matrix} \right\|_2^2 h(z_i, z_j) \pi(z_i) \pi(z_j) dz_i dz_j \\ &= \frac{C_0}{2} \int \left(\int (\partial_i \log \pi(z_i, z_j) - \partial_i \log \pi(z_i))^2 \pi(z_j | z_i) dz_j \right) \pi(z_i) dz_i \\ &\quad + \frac{C_0}{2} \int \left(\int (\partial_j \log \pi(z_i, z_j) - \partial_j \log \pi(z_j))^2 \pi(z_i | z_j) dz_i \right) \pi(z_j) dz_j. \end{aligned} \quad (44)$$

Next we apply the Poincaré inequality to bound the two integrands in the above expression. By assumption (8), the density $z_i \mapsto \pi(z_i | z_j)$ has a logarithmic Sobolev constant bounded by C_0 so that it satisfies the Poincaré inequality

$$\int \left(f(z_i) - \int f(z'_i) \pi(z'_i | z_j) dz'_i \right)^2 \pi(z_i | z_j) dz_i \leq C_0 \int f'(z_i)^2 \pi(z_i | z_j) dz_i, \quad (45)$$

for any continuously differentiable function $f : \mathbb{R} \rightarrow \mathbb{R}$. The Poincaré inequality (45) is classically obtained from a logarithmic Sobolev inequality (7) by letting $h = 1 + \varepsilon f$ and by taking a Taylor expansion as $\varepsilon \rightarrow 0$. Notice that

$$\int (\partial_j \log \pi(z'_i, z_j)) \pi(z'_i | z_j) dz'_i = \int \frac{\partial_j \pi(z'_i, z_j)}{\pi(z'_i, z_j)} \frac{\pi(z'_i, z_j)}{\pi(z_j)} dz'_i = \frac{\partial_j \int \pi(z'_i, z_j) dz'_i}{\pi(z_j)} = \partial_j \log \pi(z_j),$$

so that the Poincaré inequality (45) with $f(z_i) = \partial_j \log \pi(z_i, z_j)$ yields

$$\int (\partial_j \log \pi(z_i, z_j) - \partial_j \log \pi(z_j))^2 \pi(z_i|z_j) dz_i \leq C_0 \int (\partial_i \partial_j \log \pi(z_i, z_j))^2 \pi(z_i|z_j) dz_i.$$

In the same way, by permuting i and j , we obtain

$$\int (\partial_i \log \pi(z_i, z_j) - \partial_i \log \pi(z_i))^2 \pi(z_j|z_i) dz_j \leq C_0 \int (\partial_i \partial_j \log \pi(z_i, z_j))^2 \pi(z_j|z_i) dz_j.$$

Using the above two inequalities in (44) yields

$$\begin{aligned} I(Z_i; Z_j) &\leq \frac{C_0}{2} \int \left(C_0 \int (\partial_i \partial_j \log \pi(z_i, z_j))^2 \pi(z_j|z_i) dz_j \right) \pi(z_i) dz_i \\ &\quad + \frac{C_0}{2} \int \left(C_0 \int (\partial_i \partial_j \log \pi(z_i, z_j))^2 \pi(z_i|z_j) dz_i \right) \pi(z_j) dz_j \\ &= C_0^2 \int (\partial_i \partial_j \log \pi(z_i, z_j))^2 \pi(z_i, z_j) dz_i dz_j. \end{aligned}$$

This shows that $I(Z_i; Z_j) \leq C_0^2 \Omega_{i,j}$ when $d = 2$.

Assume now that $d > 2$. For any $\mathbf{z}_{-ij} \in \mathbb{R}^{d-2}$, replacing $\pi(z_i, z_j)$ by $\pi(z_i, z_j|\mathbf{z}_{-ij})$ in the previous analysis allows us to write

$$\begin{aligned} &\int \pi(z_i, z_j|\mathbf{z}_{-ij}) \log \left(\frac{\pi(z_i, z_j|\mathbf{z}_{-ij})}{\pi(z_i|\mathbf{z}_{-ij})\pi(z_j|\mathbf{z}_{-ij})} \right) dz_i dz_j \\ &\leq C_0^2 \int (\partial_i \partial_j \log \pi(z_i, z_j|\mathbf{z}_{-ij}))^2 \pi(z_i, z_j|\mathbf{z}_{-ij}) dz_i dz_j \\ &= C_0^2 \int (\partial_i \partial_j \log \pi(\mathbf{z}))^2 \pi(z_i, z_j|\mathbf{z}_{-ij}) dz_i dz_j, \end{aligned}$$

where the last equality is obtained by $\log \pi(z_i, z_j|\mathbf{z}_{-ij}) = \log \pi(\mathbf{z}) - \log \pi(\mathbf{z}_{-ij})$. Multiplying by the marginal $\pi(\mathbf{z}_{-ij})$ and integrating over $\mathbf{z}_{-ij} \in \mathbb{R}^{d-2}$ we obtain (9), which concludes the proof. \blacksquare

Appendix B. Proofs of transport map results

Proof of Proposition 5 Let ν_ρ be a measure on \mathbb{R}^d that is Markov with respect to a graph \mathcal{G} and has a strictly positive density ρ . By the Hammersley-Clifford theorem, the density ρ factorizes as

$$\rho(\mathbf{x}) = \frac{1}{\mathcal{Z}} \prod_{c \in \mathcal{C}} \varphi_c(\mathbf{x}_c), \tag{46}$$

where φ_c are nonnegative potential functions, \mathcal{Z} is a normalizing constant and \mathcal{C} is the set of maximal cliques of \mathcal{G} (Lauritzen, 1996). A clique is a fully connected subset of nodes and a maximal clique is a clique that is not a strict subset of another clique.

The pullback density of ρ through a differentiable diagonal transport map D is given by

$$\begin{aligned} D^\sharp \rho(\mathbf{z}) &= \rho \circ D(\mathbf{z}) |\det \nabla D(\mathbf{z})| \\ &= \rho \circ D(\mathbf{z}) \prod_{i=1}^d \partial_i D^i(z_i) \\ &= \frac{1}{\mathcal{Z}} \prod_{c \in \mathcal{C}} \varphi_c(D^c(\mathbf{z}_c)) \prod_{i=1}^d \partial_i D^i(z_i), \end{aligned}$$

where D^c represents the subset of components of D corresponding to the nodes in clique c . Collecting the derivatives of the map components for the nodes in each clique, we have

$$\begin{aligned} D^\sharp \rho(\mathbf{z}) &= \frac{1}{\mathcal{Z}} \left(\varphi_{c_1}(D^{c_1}(\mathbf{z}_{c_1})) \prod_{i \in c_1} \partial_i D^i(z_i) \right) \left(\varphi_{c_2}(D^{c_2}(\mathbf{z}_{c_2})) \prod_{i \in c_2; i \notin c_1} \partial_i D^i(z_i) \right) \dots \\ &\quad \left(\varphi_{c_M}(D^{c_M}(\mathbf{z}_{c_M})) \prod_{i \in c_M; i \notin c_1, \dots, c_{M-1}} \partial_i D^i(z_i) \right) \\ &\equiv \frac{1}{\mathcal{Z}} \psi_{c_1}(\mathbf{z}_{c_1}) \psi_{c_2}(\mathbf{z}_{c_2}) \dots \psi_{c_M}(\mathbf{z}_{c_M}), \end{aligned} \quad (47)$$

where $\psi_{c_j}(z_{c_j})$ defines new potential functions of the variables \mathbf{Z}_{c_j} in the maximal clique c_j , and $M = |\mathcal{C}|$ represents the cardinality of \mathcal{C} . From (47), the density of $D^\sharp \rho$ also factorizes according to \mathcal{G} . Thus, by Proposition 3.8 in Lauritzen (1996), $D^\sharp \nu_\rho$ is Markov with respect to \mathcal{G} . \blacksquare

We note that the contrapositive of Proposition 5 immediately follows: If the minimal I-map of $S^\sharp \nu_\rho$ is not equivalent to that of ν_ρ , then the map S is not diagonal.

Proof of Proposition 6 Let ρ be a strictly positive density on \mathbb{R}^d and let D be a differentiable diagonal transport map. The log of the pullback density $\pi = D^\sharp \rho$ is given by

$$\begin{aligned} \log \pi(\mathbf{z}) &= \log D^\sharp \rho(\mathbf{z}) = \log \rho \circ D(\mathbf{z}) + \log |\det \nabla D(\mathbf{z})| \\ &= \log \rho \circ D(\mathbf{z}) + \sum_{k=1}^d \log \partial_k D^k(z^k). \end{aligned} \quad (48)$$

The partial derivatives of the log-density with respect to z_i, z_j are given by

$$\partial_i \partial_j \log \pi(\mathbf{z}) = \partial_i \partial_j \log \rho \circ D(\mathbf{z}) \partial_i D^i(z_i) \partial_j D^j(z_j),$$

by using that each term in the log-determinant of the map's Jacobian only depends on a single variable. Thus, entry (i, j) of the conditional independence score for π is given by

$$\Omega_{ij} = \int |\partial_i \partial_j \log \rho \circ D(\mathbf{z}) \partial_i D^i(z_i) \partial_j D^j(z_j)|^2 \pi(\mathbf{z}) d\mathbf{z}. \quad (49)$$

Using the measure transformation $\nu_\pi = D^\sharp \nu_\rho$, we rewrite the conditional independence score of π in (49) as the following expectation over ρ

$$\Omega_{ij} = \int \left| \partial_i \partial_j \log \rho(\mathbf{x}) \partial_i D^i((D^i)^{-1}(x_i)) \partial_j D^j((D^j)^{-1}(x_j)) \right|^2 \rho(\mathbf{x}) d\mathbf{x}. \quad (50)$$

Finally, by applying the inverse function theorem to each continuously differentiable map component to get $\partial_i D^i((D^i)^{-1}(x_i)) = \partial_i (D^i)^{-1}(x_i)$, we arrive at the score in (13). \blacksquare

Proof of Proposition 8 For the polynomial degree $\beta = 1$, the transport map S is an affine function $S(\mathbf{z}) = L(\mathbf{z} - \mathbf{c})$ with an invertible lower triangular matrix $L \in \mathbb{R}^{d \times d}$ and vector $\mathbf{c} \in \mathbb{R}^d$. For this class of transport maps, the minimization problem in (17) becomes

$$\begin{aligned} \operatorname{argmax}_S \frac{1}{n} \sum_{l=1}^n \log S^\sharp \eta(\mathbf{z}^l) &= \operatorname{argmax}_S \frac{1}{n} \sum_{l=1}^n \log \eta \circ S(\mathbf{z}^l) + \log |\det \nabla S(\mathbf{z}^l)| \\ &= \operatorname{argmax}_{L, \mathbf{c}} \frac{1}{n} \sum_{l=1}^n \left(-\frac{1}{2} (\mathbf{z}^l - \mathbf{c})^T L^T L (\mathbf{z}^l - \mathbf{c}) + \log \det L \right). \end{aligned} \quad (51)$$

For any invertible matrix $L^T L$, the optimal \mathbf{c} is the empirical mean $\hat{\mathbf{m}} := \frac{1}{n} \sum_{l=1}^n \mathbf{z}^l$. Substituting this value for \mathbf{c} in (51), the objective for L is given by

$$\operatorname{argmin}_L \left\{ \operatorname{Tr} \left(L^T L \hat{\Sigma} \right) - \log \det (L^T L) \right\}, \quad (52)$$

where $\hat{\Sigma} := \frac{1}{n} \sum_{i=1}^n (\mathbf{z}^i - \hat{\mathbf{m}})(\mathbf{z}^i - \hat{\mathbf{m}})^T$ denotes the empirical covariance matrix and Tr is the matrix trace operator. Setting the gradient of (52) with respect to $L^T L$ to zero yields the optimal L to be the inverse of the Cholesky factor of $\hat{\Sigma}$ (which exists for $n \geq d$). Thus, the pullback density $S^\sharp \eta$ for a standard Gaussian density η yields a multivariate Gaussian approximation to π with mean $\hat{\mathbf{m}}$ and covariance matrix $\hat{\Sigma}$. \blacksquare

Appendix C. Proof of Proposition 9

Proof The proof considers two types of errors: a false positive occurs when the true score Ω_{ij} is zero but the threshold estimate $\bar{\Omega}_{ij} = \hat{\Omega}_{ij} \mathbb{1}(\hat{\Omega}_{ij} > \tau_{ij})$ is nonzero; and a false negative occurs when the true score Ω_{ij} is nonzero but the threshold estimate $\bar{\Omega}_{ij}$ is zero.

For each pair of variables (i, j) , let $g: \boldsymbol{\alpha} \mapsto \mathbb{E}_\pi |\partial_i \partial_j \log S_{\boldsymbol{\alpha}}^\sharp \eta(\mathbf{z})|^2$ be a continuous function of the coefficients $\boldsymbol{\alpha}$. Then, we have $\hat{\Omega}_{ij} = g(\hat{\boldsymbol{\alpha}})$ (as defined in (20)) and $\hat{\nu}_{ij} = (\nabla_{\boldsymbol{\alpha}} g(\hat{\boldsymbol{\alpha}})^T \Gamma(\hat{\boldsymbol{\alpha}})^{-1} \nabla_{\boldsymbol{\alpha}} g(\hat{\boldsymbol{\alpha}}))^{1/2}$ (defined below (25)). Assuming that g is twice differentiable, a Taylor expansion of the score estimator around $\boldsymbol{\alpha}^*$ yields

$$g(\hat{\boldsymbol{\alpha}}) = g(\boldsymbol{\alpha}^*) + \nabla_{\boldsymbol{\alpha}} g(\boldsymbol{\alpha}^*)^T (\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}^*) + \frac{1}{2} (\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}^*)^T \nabla_{\boldsymbol{\alpha}}^2 g(\boldsymbol{\alpha}^*) (\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}^*) + o_p(\|\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}^*\|^2), \quad (53)$$

where the remainder is a term that tends to zero in probability.

For conditionally dependent variables, we have $g(\boldsymbol{\alpha}^*) = \Omega_{ij} \neq 0$ and, by assumption, $\nabla_{\boldsymbol{\alpha}} g(\boldsymbol{\alpha}^*) \neq 0$. Therefore, we can truncate the expansion in (53) to first-order

terms and from the delta method we have $\sqrt{n}(\widehat{\Omega}_{ij} - \Omega_{ij}) = \sqrt{n}(g(\widehat{\boldsymbol{\alpha}}) - g(\boldsymbol{\alpha}^*)) \xrightarrow{d} \mathcal{N}(0, \nabla_{\boldsymbol{\alpha}} g(\boldsymbol{\alpha}^*)^T \Gamma(\boldsymbol{\alpha}^*)^{-1} \nabla_{\boldsymbol{\alpha}} g(\boldsymbol{\alpha}^*))$ as $n \rightarrow \infty$. From the continuous mapping theorem, \widehat{v}_{ij} tends in probability to the constant $v_{ij} = (\nabla_{\boldsymbol{\alpha}} g(\boldsymbol{\alpha}^*)^T \Gamma(\boldsymbol{\alpha}^*)^{-1} \nabla_{\boldsymbol{\alpha}} g(\boldsymbol{\alpha}^*))^{1/2} \neq 0$ (nonzero for positive-definite Fisher information matrix $\Gamma(\boldsymbol{\alpha}^*)$ by the assumption $\nabla_{\boldsymbol{\alpha}} g(\boldsymbol{\alpha}^*) \neq 0$) as $\widehat{\boldsymbol{\alpha}} \rightarrow \boldsymbol{\alpha}^*$. Then from Slutsky's theorem (Corollary 2.3.2 in Lehmann (2004)), the ratio $\sqrt{n}(\widehat{\Omega}_{ij}/\widehat{v}_{ij} - \Omega_{ij}/v_{ij}) \xrightarrow{d} R$ as $n \rightarrow \infty$ where $R \sim \mathcal{N}\left(0, \nabla_{\boldsymbol{\alpha}} g(\boldsymbol{\alpha}^*)^T \Gamma(\boldsymbol{\alpha}^*)^{-1} \nabla_{\boldsymbol{\alpha}} g(\boldsymbol{\alpha}^*)/v_{ij}^2\right) = \mathcal{N}(0, 1)$. Thus, the probability of false negatives given by

$$\begin{aligned} \mathbb{P}(\overline{\Omega}_{ij} = 0; \Omega_{ij} \neq 0) &= \mathbb{P}\left(\widehat{\Omega}_{ij} < \frac{f(n)\widehat{v}_{ij}}{\sqrt{n}}; \Omega_{ij} \neq 0\right) \\ &= \mathbb{P}\left(\sqrt{n}\left(\frac{\widehat{\Omega}_{ij}}{\widehat{v}_{ij}} - \frac{\Omega_{ij}}{v_{ij}}\right) < f(n) - \frac{\sqrt{n}\Omega_{ij}}{v_{ij}}; \Omega_{ij} \neq 0\right), \end{aligned}$$

asymptotically converges to

$$\mathbb{P}\left(R < f(n) - \frac{\sqrt{n}\Omega_{ij}}{v_{ij}}\right) \leq e^{-\frac{1}{2}\left(f(n) - \frac{\sqrt{n}\Omega_{ij}}{v_{ij}}\right)^2}, \quad (54)$$

where the last inequality provides an explicit bound and follows from the Gaussian tail bound $\mathbb{P}(R < r) \leq e^{-r^2/2}$ for $r \leq -1$. For $f(n)/\sqrt{n} \rightarrow 0$, the term $f(n) - \frac{\sqrt{n}\Omega_{ij}}{v_{ij}}$ tends to negative infinity as $n \rightarrow \infty$, and therefore, the probability of a false negative in (54) converges to zero for any f that grows more slowly than \sqrt{n} . Conversely, if $f(n) = c\sqrt{n}$ for some $c \geq \Omega_{ij}/v_{ij}$, the left-hand side in (54) does not go to zero.

For conditionally independent variables, we have $\partial_i \partial_j \log S_{\boldsymbol{\alpha}^*}^{\sharp} \eta(\mathbf{z}) = 0$ for all $\mathbf{z} \in \mathbb{R}^d$; thus $g(\boldsymbol{\alpha}^*) = \mathbb{E}|\partial_i \partial_j \log S_{\boldsymbol{\alpha}^*}^{\sharp} \eta(\mathbf{Z})|^2 = 0$ and $\nabla_{\boldsymbol{\alpha}} g(\boldsymbol{\alpha}^*) = 2\mathbb{E}[\partial_i \partial_j \log S_{\boldsymbol{\alpha}^*}^{\sharp} \eta(\mathbf{Z}) \nabla_{\boldsymbol{\alpha}} \partial_i \partial_j \log S_{\boldsymbol{\alpha}^*}^{\sharp} \eta(\mathbf{Z})] = 0$. Under the assumption that $\nabla_{\boldsymbol{\alpha}}^2 g(\boldsymbol{\alpha}^*) \neq 0$, we can truncate the expansion in (53) to second-order terms and approximate $\widehat{\Omega}_{ij}$ asymptotically using a quadratic form (i.e., a degree two polynomial). By Proposition 2.1 in Drton et al. (2016), the Wald statistic $n\widehat{\Omega}_{ij}^2/\widehat{v}_{ij}^2 = \frac{ng(\widehat{\boldsymbol{\alpha}})^2}{\nabla_{\boldsymbol{\alpha}} g(\widehat{\boldsymbol{\alpha}})^T \Gamma(\widehat{\boldsymbol{\alpha}})^{-1} \nabla_{\boldsymbol{\alpha}} g(\widehat{\boldsymbol{\alpha}})} \xrightarrow{d} W$, where the random variable W satisfies, by Proposition 3.4 in Drton et al. (2016), $\mathbb{P}(W \geq x) \leq \mathbb{P}(\frac{1}{4}X \geq x)$ for $X \sim \chi_p^2$ and all $x > 0$. Here, χ_p^2 denotes a chi-squared variable with $p \geq 1$ degrees of freedom where p corresponds to the dimension of $\boldsymbol{\alpha}$. Then, the probability of false positives, which can be written as

$$\mathbb{P}(\overline{\Omega}_{ij} \neq 0; \Omega_{ij} = 0) = \mathbb{P}\left(\widehat{\Omega}_{ij} > \frac{f(n)\widehat{v}_{ij}}{\sqrt{n}}; \Omega_{ij} = 0\right) = \mathbb{P}\left(\frac{n\widehat{\Omega}_{ij}^2}{\widehat{v}_{ij}^2} > f(n)^2; \Omega_{ij} = 0\right)$$

since $\widehat{\Omega}_{ij}$ and \widehat{v}_{ij} are non-negative, is asymptotically given by

$$\mathbb{P}(W > f(n)^2) \leq \mathbb{P}\left(X - p > 4p\left(\frac{f(n)^2}{p} - \frac{1}{4}\right)\right) \leq e^{-f(n)^2/p+1/4}, \quad (55)$$

where the last inequality uses the χ_p^2 squared tail bound $\mathbb{P}(X - p \geq 4px) \leq \mathbb{P}(X - p \geq 2\sqrt{px} + 2x) \leq e^{-x}$ for $x \geq 1$ (Lemma 1 in Laurent and Massart (2000)). For $f(n) \rightarrow \infty$, the right-hand side in (55) and hence the left-hand side converge to zero as $n \rightarrow \infty$.

To complete the proof we use a union bound over all pairs (i, j) to bound the probability of failing to recover the true edge set

$$\mathbb{P}(\widehat{E}_n \neq E) \leq \sum_{(i,j) \in E} \mathbb{P}(\bar{\Omega}_{ij} = 0; \Omega_{ij} \neq 0) + \sum_{(i,j) \notin E} \mathbb{P}(\bar{\Omega}_{ij} \neq 0; \Omega_{ij} = 0). \quad (56)$$

Using the asymptotic results in (54) and (55) to bound each term in (56) we have $\mathbb{P}(\widehat{E}_n \neq E) \rightarrow 0$ as $n \rightarrow \infty$. ■

Appendix D. Data on memory usage

Following the discussion in Section 6.1.2, we report the memory usage of the first two iterations of SING for different data set dimensions d and polynomial degrees β , applied to the butterfly dataset of Section 6.1.

| Polynomial degree β | Dimension d | Map estimation | Estimating $\tilde{\Omega}$ | Estimating $\text{Var}(\tilde{\Omega})$ |
|---------------------------|---------------|--------------------|-----------------------------|-----------------------------------------|
| 1 | 4 | 51.8×10^3 | 19.3×10^6 | 10.2×10^7 |
| 2 | 4 | 96.2×10^5 | 55.3×10^6 | 29.4×10^7 |
| 1 | 6 | 18.3×10^6 | 44.9×10^6 | 33.3×10^7 |
| 2 | 6 | 21.9×10^6 | 18.0×10^7 | 12.2×10^8 |
| 1 | 8 | 57.6×10^6 | 13.2×10^7 | 84.5×10^7 |
| 2 | 8 | 73.1×10^6 | 49.6×10^7 | 35.9×10^8 |
| 1 | 10 | 36.0×10^6 | 13.4×10^7 | 15.9×10^8 |
| 2 | 10 | 21.6×10^6 | 99.6×10^7 | 94.4×10^8 |
| 1 | 12 | 49.4×10^6 | 13.3×10^7 | 27.8×10^8 |
| 2 | 12 | 35.7×10^6 | 19.0×10^8 | 22.0×10^9 |

Table 1: Breakdown of memory usage (in bytes) for the first iteration of SING applied to the butterfly data set with $n = 1000$ samples.

References

Søren Asmussen and Peter W Glynn. *Stochastic simulation: algorithms and analysis*, volume 57. Springer Science & Business Media, 2007.

Dominique Bakry and Michel Émery. Diffusions hypercontractives. In *Séminaire de Probabilités XIX 1983/84*, pages 177–206. Springer, 1985.

Onureena Banerjee, Laurent El Ghaoui, and Alexandre d’Aspremont. Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *Journal of Machine Learning Research*, 9(Mar):485–516, 2008.

Ricardo Baptista, Youssef Marzouk, and Olivier Zahm. On the representation and learning of monotone triangular transport maps. *Foundations of Computational Mathematics*, pages 1–46, 2023.

| Polynomial degree β | Dimension d | Map estimation | Estimating $\tilde{\Omega}$ | Estimating $\text{Var}(\tilde{\Omega})$ |
|---------------------------|---------------|--------------------|-----------------------------|-----------------------------------------|
| 1 | 4 | 36.8×10^5 | 13.4×10^6 | 33.0×10^6 |
| 2 | 4 | 21.5×10^6 | 34.4×10^6 | 63.7×10^6 |
| 1 | 6 | 13.1×10^6 | 28.2×10^6 | 97.4×10^6 |
| 2 | 6 | 51.9×10^6 | 34.3×10^6 | 13.8×10^7 |
| 1 | 8 | 78.1×10^6 | 84.1×10^6 | 24.8×10^7 |
| 2 | 8 | 12.9×10^7 | 96.0×10^6 | 34.2×10^7 |
| 1 | 10 | 52.8×10^6 | 38.5×10^6 | 35.9×10^7 |
| 2 | 10 | 13.9×10^7 | 48.1×10^6 | 52.8×10^7 |
| 1 | 12 | 18.1×10^6 | 59.7×10^6 | 61.3×10^7 |
| 2 | 12 | 16.2×10^7 | 66.4×10^6 | 89.6×10^7 |

Table 2: Breakdown of memory usage (in bytes) for the second iteration of SING applied to the butterfly data set with $n = 1000$ samples

Peter J Bickel, Elizaveta Levina, et al. Regularized estimation of large covariance matrices. *The Annals of Statistics*, 36(1):199–227, 2008.

Chris Biemann. Chinese whispers: an efficient graph clustering algorithm and its application to natural language processing problems. In *Proceedings of the first workshop on graph based methods for natural language processing*, pages 73–80. Association for Computational Linguistics, 2006.

Gary S Bloom and Solomon W Golomb. Applications of numbered undirected graphs. *Proceedings of the IEEE*, 65(4):562–570, 1977.

Vladimir Igorevich Bogachev, Aleksandr Viktorovich Kolesnikov, and Kirill Vladimirovich Medvedev. Triangular transformations of measures. *Sbornik: Mathematics*, 196(3):309, 2005.

Michael Brennan, Daniele Bigoni, Olivier Zahm, Alessio Spantini, and Youssef Marzouk. Greedy inference with structure-exploiting lazy maps. *Advances in Neural Information Processing Systems*, 33:8330–8342, 2020.

Guy Bresler. Efficiently learning Ising models on arbitrary graphs. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 771–782, 2015.

Tony Cai and Weidong Liu. Adaptive thresholding for sparse covariance matrix estimation. *Journal of the American Statistical Association*, 106(494):672–684, 2011.

George Casella and Roger L Berger. *Statistical inference*, volume 2. Duxbury Pacific Grove, CA, 2002.

Venkat Chandrasekaran, Pablo A. Parrilo, and Alan S. Willsky. Latent variable graphical model selection via convex optimization. *Ann. Statist.*, 40(4):1935–1967, 08 2012.

Patrick Danaher, Pei Wang, and Daniela M Witten. The joint graphical lasso for inverse covariance estimation across multiple classes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(2):373–397, 2014.

- Timothy A Davis. User guide for CHOLMOD: a sparse Cholesky factorization and modification package. *Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL, USA*, 2008.
- Mathias Drton and Marloes H Maathuis. Structure learning in graphical modeling. *Annual Review of Statistics and Its Application*, 4:365–393, 2017.
- Mathias Drton, Bernd Sturmfels, and Seth Sullivant. *Lectures on algebraic statistics*, volume 39. Springer Science & Business Media, 2008.
- Mathias Drton, Han Xiao, et al. Wald tests of singular hypotheses. *Bernoulli*, 22(1):38–59, 2016.
- Sacha Epskamp, Lourens J Waldorp, René Møttus, and Denny Borsboom. The Gaussian graphical model in cross-sectional and time-series data. *Multivariate Behavioral Research*, 53(4):453–480, 2018.
- Chi Feng and Youssef M Marzouk. A layered multiple importance sampling scheme for focused optimal Bayesian experimental design. *arXiv preprint arXiv:1903.11187*, 2019.
- Andreas Fischer, Ching Y Suen, Volkmar Frinken, Kaspar Riesen, and Horst Bunke. A fast matching algorithm for graph-based handwriting recognition. In *International Workshop on Graph-Based Representations in Pattern Recognition*, pages 194–203. Springer, 2013.
- Tomislav Fotak, Miroslav Bača, and Petra Koruga. Handwritten signature identification using basic concepts of graph theory. *WSEAS Transactions on Signal Processing*, 7:117–129, 2011.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.
- Andrew Gelman and Xiao-Li Meng. Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical science*, pages 163–185, 1998.
- Alice Guionnet and B Zegarliński. Lectures on logarithmic Sobolev inequalities. In *Séminaire de probabilités XXXVI*, pages 1–134. Springer, 2003.
- Jian Guo, Elizaveta Levina, George Michailidis, and Ji Zhu. Joint estimation of multiple graphical models. *Biometrika*, 98(1):1–15, 2011.
- RA Holley and DW Stroock. L_2 theory for the stochastic Ising model. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 35(2):87–101, 1976.
- Jianhua Z Huang, Naiping Liu, Mohsen Pourahmadi, and Linxu Liu. Covariance matrix selection and estimation via penalised normal likelihood. *Biometrika*, 93(1):85–98, 2006.
- Priyank Jaini, Kira A Selby, and Yaoliang Yu. Sum-of-squares polynomial flow. In *International Conference on Machine Learning*, pages 3009–3018, 2019.

- Ali Jalali, Christopher C Johnson, and Pradeep K Ravikumar. On learning discrete graphical models using greedy methods. In *Advances in Neural Information Processing Systems*, pages 1935–1943, 2011.
- Xiaoning Kang and Xinwei Deng. An improved modified Cholesky decomposition approach for precision matrix estimation. *Journal of Statistical Computation and Simulation*, 90(3): 443–464, 2020.
- George Karypis and Vipin Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on scientific Computing*, 20(1):359–392, 1998.
- Herbert Knothe. Contributions to the theory of convex bodies. *The Michigan Mathematical Journal*, 1957(1028990175), 1957.
- Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.
- Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.
- Steffen L Lauritzen. *Graphical models*, volume 17. Clarendon Press, 1996.
- Erich Leo Lehmann. *Elements of large-sample theory*. Springer Science & Business Media, 2004.
- Elizaveta Levina, Adam Rothman, Ji Zhu, et al. Sparse estimation of large covariance matrices via a nested Lasso penalty. *The Annals of Applied Statistics*, 2(1):245–263, 2008.
- Lina Lin, Mathias Drton, and Ali Shojaie. Estimation of high-dimensional graphical models using regularized score matching. *Electronic journal of statistics*, 10(1):806, 2016.
- Han Liu, John Lafferty, and Larry Wasserman. The nonparanormal: Semiparametric estimation of high dimensional undirected graphs. *Journal of Machine Learning Research*, 10:2295–2328, 2009.
- Han Liu, Fang Han, Ming Yuan, John Lafferty, and Larry Wasserman. The nonparanormal skeptic. In *29th International Conference on Machine Learning, ICML 2012*, pages 1415–1422, 2012a.
- Han Liu, Fang Han, and Cun-hui Zhang. Transelliptical graphical models. In *Advances in Neural Information Processing Systems*, pages 800–808, 2012b.
- Po-Ling Loh and Martin J Wainwright. Structure estimation for discrete graphical models: Generalized covariance matrices and their inverses. In *NIPS*, pages 2096–2104, 2012.
- Edward N Lorenz. Predictability: A problem partly solved. In *Proc. Seminar on predictability*, 1996.
- Michael Lüllesmann, Simon R Leßenich, and H Martin Bückner. Interactively exploring elimination orderings in symbolic sparse Cholesky factorization. *Procedia Computer Science*, 1(1):867–874, 2010.

- Youssef Marzouk, Tarek Moselhy, Matthew Parno, and Alessio Spantini. Sampling via measure transport: An introduction. In Roger Ghanem, David Higdon, and Houman Owhadi, editors, *Handbook of Uncertainty Quantification*. Springer, 2016. ISBN 978-3-319-11259-6. doi: 10.1007/978-3-319-11259-6_23-1.
- Nicolai Meinshausen and Peter Bühlmann. High-dimensional graphs and variable selection with the Lasso. *The Annals of Statistics*, pages 1436–1462, 2006.
- Rada Mihalcea and Dragomir Radev. *Graph-based natural language processing and information retrieval*. Cambridge university press, 2011.
- Rebecca Morrison, Ricardo Baptista, and Youssef Marzouk. Beyond normality: Learning sparse probabilistic graphical models in the non-Gaussian setting. In *Advances in Neural Information Processing Systems 30*, pages 2359–2369, 2017.
- Rebecca Morrison, Ricardo Baptista, and Estelle Basor. Diagonal nonlinear transformations preserve structure in covariance and precision matrices. *Journal of Multivariate Analysis*, 190:104983, 2022. ISSN 0047-259X.
- Roger B Nelsen. *An introduction to copulas*. Springer Science & Business Media, 2007.
- Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Science & Business Media, 2006.
- Gary W Oehlert. A note on the delta method. *The American Statistician*, 46(1):27–29, 1992.
- Edward Ott, Brian R. Hunt, Istvan Szunyogh, Aleksey V. Zimin, Eric J. Kostelich, Matteo Corazza, Eugenia Kalnay, D. J. Patil, and James A. Yorke. A local ensemble kalman filter for atmospheric data assimilation, 2002. URL <https://arxiv.org/abs/physics/0203058>.
- George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, pages 2338–2347, 2017.
- Matthew D Parno and Youssef M Marzouk. Transport map accelerated Markov chain Monte Carlo. *SIAM/ASA Journal on Uncertainty Quantification*, 6(2):645–682, 2018.
- Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- Natesh S Pillai and Xiao-Li Meng. An unexpected encounter with Cauchy and Lévy. *The Annals of Statistics*, pages 2089–2097, 2016.
- Tom Rainforth, Rob Cornish, Hongseok Yang, Andrew Warrington, and Frank Wood. On nesting Monte Carlo estimators. In *International Conference on Machine Learning*, pages 4267–4276. PMLR, 2018.
- James O Ramsay. Estimating smooth monotone functions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(2):365–375, 1998.

- Garvesh Raskutti and Caroline Uhler. Learning directed acyclic graph models based on sparsest permutations. *Stat*, 7(1):e183, 2018.
- S. Reich and C. Cotter. *Probabilistic Forecasting and Bayesian Data Assimilation*. Cambridge Texts in Applied Mathematics. Cambridge University Press, 2015. ISBN 9781107069398.
- Donald J Rose, R Endre Tarjan, and George S Lueker. Algorithmic aspects of vertex elimination on graphs. *SIAM Journal on computing*, 5(2):266–283, 1976.
- Murray Rosenblatt. Remarks on a multivariate transformation. *The Annals of Mathematical Statistics*, 23(3):470–472, 1952.
- H. Rue and L. Held. *Gaussian Markov Random Fields: Theory and Applications*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. CRC Press, 2005. ISBN 9780203492024.
- Thomas L Saaty and Robert G Busacker. *Finite Graphs and Networks: an introduction with applications*. McGraw-Hill Book Company, 1965.
- Karen Sachs, Omar Perez, Dana Pe’er, Douglas A Lauffenburger, and Garry P Nolan. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308(5721):523–529, 2005.
- Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, 55(58-63):94, 2015.
- Narayana P Santhanam and Martin J Wainwright. Information-theoretic limits of selecting binary graphical models in high dimensions. *IEEE Transactions on Information Theory*, 58(7):4117–4134, 2012.
- Huawei Shen, Xueqi Cheng, Kai Cai, and Mao-Bin Hu. Detect overlapping and hierarchical community structure in networks. *Physica A: Statistical Mechanics and its Applications*, 388(8):1706–1712, 2009.
- Alessio Spantini, Daniele Bigoni, and Youssef Marzouk. Inference via low-dimensional couplings. *The Journal of Machine Learning Research*, 19(1):2639–2709, 2018.
- Alessio Spantini, Ricardo Baptista, and Youssef Marzouk. Coupling techniques for nonlinear ensemble filtering. *SIAM Review*, 64(4):921–953, 2022.
- Arun Suggala, Mladen Kolar, and Pradeep K Ravikumar. The expxorclist: nonparametric graphical models via conditional exponential densities. In *Advances in neural information processing systems*, pages 4446–4456, 2017.
- John Thomas, Naren Ramakrishnan, and Chris Bailey-Kellogg. Protein design by sampling an undirected graphical model of residue constraints. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 6(3):506–516, 2008.
- Transport Maps Team. TransportMaps v2.0. <http://transportmaps.mit.edu/docs/example-inverse-sparsity-identification.html>, 2018.

- C. Villani. *Optimal Transport: Old and New*, volume 338. Springer, 2008.
- Martin J Wainwright, John D. Lafferty, and Pradeep K. Ravikumar. High-dimensional graphical model selection using ℓ_1 -regularized logistic regression. In B. Schölkopf, J. C. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 1465–1472. MIT Press, 2007.
- Wei Wang, Martin J Wainwright, and Kannan Ramchandran. Information-theoretic bounds on model selection for Gaussian Markov random fields. In *2010 IEEE International Symposium on Information Theory*, pages 1373–1377. IEEE, 2010.
- Eunho Yang, Pradeep Ravikumar, Genevera I Allen, and Zhandong Liu. Graphical models via univariate exponential family distributions. *The Journal of Machine Learning Research*, 16(1):3813–3847, 2015.
- Mihalis Yannakakis. Computing the minimum fill-in is NP-complete. *SIAM Journal on Algebraic Discrete Methods*, 2(1):77–79, 1981.
- Ming Yuan and Yi Lin. Model selection and estimation in the Gaussian graphical model. *Biometrika*, 94(1):19–35, 2007.
- Olivier Zahm, Tiangang Cui, Kody Law, Alessio Spantini, and Youssef Marzouk. Certified dimension reduction in nonlinear bayesian inverse problems. *Mathematics of Computation*, 91(336):1789–1835, 2022.
- Shuheng Zhou, John Lafferty, and Larry Wasserman. Time varying undirected graphs. *Machine Learning*, 80(2-3):295–319, 2010.