MIT Open Access Articles

Network Maximal Correlation

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

As Published	10.1109/TNSE.2017.2716966
Publisher	Institute of Electrical and Electronics Engineers (IEEE)
Version	Original manuscript
Citable link	https://hdl.handle.net/1721.1/135743
Terms of Use	Creative Commons Attribution-Noncommercial-Share Alike
Detailed Terms	http://creativecommons.org/licenses/by-nc-sa/4.0/

Network Maximal Correlation

Soheil Feizi*, Ali Makhdoumi*, Ken Duffy, Manolis Kellis, Muriel Médard

Abstract—We introduce Network Maximal Correlation (NMC) as a multivariate measure of nonlinear association among random variables. NMC is defined via an optimization that infers transformations of variables by maximizing aggregate inner products between transformed variables. For finite discrete and jointly Gaussian random variables, we characterize a solution of the NMC optimization using basis expansion of functions over appropriate basis functions. For finite discrete variables, we propose an algorithm based on alternating conditional expectation to determine NMC. Moreover we propose a distributed algorithm to compute an approximation of NMC for large and dense graphs using graph partitioning. For finite discrete variables, we show that the probability of discrepancy greater than any given level between NMC and NMC computed using empirical distributions decays exponentially fast as the sample size grows. For jointly Gaussian variables, we show that under some conditions the NMC optimization is an instance of the Max-Cut problem. We then illustrate an application of NMC in inference of graphical model for bijective functions of jointly Gaussian variables. Finally, we show NMC’s utility in a data application of learning nonlinear dependencies among genes in a cancer dataset.

Index Terms—Maximum Correlation Problem, Alternating Conditional Expectation (ACE), Hermite-Chebyshev Polynomials, Gaussian Graphical Models, Gene Networks

1 INTRODUCTION

IDENTIFYING relationships among variables in large datasets is an increasingly important task in systems biology [1], social sciences [2], finance [3], and other fields. For independent observations of bivariate data, several measures exist that characterize the strength of the association based on a linear fit (e.g., Pearson’s correlation [4], canonical correlation [5], [6]), rank statistics (e.g., Spearman’s correlation [7]), and information content (e.g., mutual information [8], [9]). Some of these measures have been extended to the multivariate setting. For instance, Chow and Liu [10] have used mutual information in the inference of tree graphical models, while Liu et al. [11], [12] introduced a copula setup based on rank statistics such as Spearman’s [7] correlation coefficient to characterize graphical models for some nonlinear functions of jointly Gaussian variables. Another method to capture a nonlinear association between two variables is the randomized dependence coefficient [13], where it fixes a set of nonlinear functions and then uses randomized rank statistics to compute association between nonlinear transformations of variables.

A classical measure of nonlinear relationships between two random elements X_1 and X_2 is *Maximal Correlation* (MC), introduced by Hirschfeld [14], Gebelein [15], Sarmanov [16] and Rényi [17], having also appeared in the work of Witsenhausen [18], Ahlswede and Gács [19], and Lancaster [20]. MC determines possibly nonlinear transformations of two variables, subject to zero mean and unit variance, to maximize their Pearson’s correlation. MC not only computes an association strength between variables, but it also characterizes a possible functional relationship between them.

Definition 1 (Maximal Correlation). Let $(\Omega, \mathcal{F}, P)$ be a prob-

ability triple and for $k \in \{1, 2\}$ let $(\mathcal{X}_k, \beta_k)$ be a measurable space with $X_k : \Omega \rightarrow \mathcal{X}_k$ a random element. Maximal Correlation (MC) between the two (not necessarily real-valued) random elements X_1 and X_2 is defined as

$$\rho(X_1, X_2) \triangleq \sup_{\phi_1, \phi_2} \mathbb{E}[\phi_1(X_1) \phi_2(X_2)], \quad (1)$$

such that $\phi_k : \mathcal{X}_k \rightarrow \mathbb{R}$ is Borel measurable, $\mathbb{E}[\phi_k(X_k)] = 0$, and $\mathbb{E}[\phi_k(X_k)^2] = 1$, for $k = 1, 2$.

MC can be computed efficiently for both discrete [21] and continuous real-valued [20] random variables (Section 2.1). For discrete random variables, under mild conditions, MC is equal to the second largest singular value of an scaled joint probability distribution matrix and the optimal transformations of the variables can be characterized using right and left singular vectors of the scaled probability distribution matrix. Alternating Conditional Expectation (ACE) was introduced by Breiman and Friedman [21] to compute MC, and was further analyzed by Buja [22]. Recently, MC has been used in many applications. It is related to strong data processing inequalities and contraction coefficients, being recently investigated by Anantharam et al. [23], Polyanskiy [24], and Raginsky [25]. Further studies include applications in information theory [26], [27], information-theoretic security and privacy [28], [29], [30], [31], [32], data processing [23], [25] and other fields [33], [34], [35], [36], [37]. In particular, Beigi and Gohari [33] has introduced multipartite maximal correlation and showed its connection with the multipartite hypercontractivity ribbon. Multipartite maximal correlation aims to find a correlation matrix C with the largest r such that $C - rI$ is positive semidefinite, where I is the identity matrix. This objective function is different than the one of the network maximal correlation (3) discussed in the present paper.

Many modern datasets consist of independent observations of high-dimensional multivariate random elements $X_1, \dots, X_n$ (i.e., $n \gg 2$). One approach to characterize relationships among these elements is to determine the bivariate MC for each

• S. Feizi and A. Makhdoumi contributed equally to this work. S. Feizi is with Stanford University. A. Makhdoumi, M. Kellis and M. Médard are with Massachusetts Institute of Technology (MIT). K. Duffy is with Maynooth University.

• Copyright (c) 2016 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org. An earlier version of this paper is available on MIT DSpace: <https://dspace.mit.edu/handle/1721.1/98878>.

pair of them. In this approach one would solve the following optimization for all pairs (i, i') , $1 \leq i, i' \leq n$:

$$\sup_{\phi_{i,i'}, \phi_{i',i}} \mathbb{E}[\phi_{i,i'}(X_i) \phi_{i',i}(X_{i'})], \quad (2)$$

such that $\phi_{i,i'} : \mathcal{X}_i \rightarrow \mathbb{R}$ is Borel measurable, $\mathbb{E}[\phi_{i,i'}(X_i)] = 0$, and $\mathbb{E}[\phi_{i,i'}(X_i)^2] = 1$, for $1 \leq i, i' \leq n$.

Should they exist, let $\phi_{i,i'}^*$ for $1 \leq i, i' \leq n$ denote the resulting optimizers. By this approach, each element X_i is assigned to $n-1$ transformation functions $\{\phi_{i,i'}^* : 1 \leq i' \leq n, i' \neq i\}$. In some applications, there may be interpretability and overfitting issues. To circumvent these issues, we consider an alternate multivariate extension of the MC optimization (1) with two types of constraints. We seek to formulate an optimization that (i) assigns a single transformation function to each element; and (ii) its objective function can be restricted to a subset of variable pairs. The latter conditioning is described by a graph and is motivated by several distinct considerations: (a) there may be *a priori* known structure that indicates some elements are necessarily unrelated; (b) one may not care about the association of certain elements; or (c) the restriction may serve as a computation reduction technique.

Definition 2 (Network Maximal Correlation). Let $G = (V, E)$ be a graph with vertices $V = \{1, 2, \dots, n\}$ and edges $E \subseteq \{(i, i') : i, i' \in V, i \neq i'\}$. The Network Maximal Correlation (NMC) of $X_1, \dots, X_n$ given G is defined to be

$$\rho_G(X_1, \dots, X_n) \triangleq \sup_{\phi_1, \dots, \phi_n} \sum_{(i, i') \in E} \mathbb{E}[\phi_i(X_i) \phi_{i'}(X_{i'})], \quad (3)$$

such that $\phi_i : \mathcal{X}_i \rightarrow \mathbb{R}$ is Borel measurable, $\mathbb{E}[\phi_i(X_i)] = 0$, and $\mathbb{E}[\phi_i(X_i)^2] = 1$, for all $1 \leq i \leq n$.

When no confusion arises, we use ρ_G to refer to the NMC. The NMC optimization maximizes the aggregate inner products between transformed variables. NMC naturally generalizes the bivariate MC to the multivariate setting. For example, the bivariate MC is equivalent to the NMC when the graph $G = (V, E)$ has two nodes connected by an edge. That is, $V = \{1, 2\}$ and $E = \{(1, 2)\}$. As another example, if $V = \{1, 2, 3\}$ and $E = \{(1, 2), (1, 3)\}$, the objective function of the NMC optimization (3) aims to find transformation functions $\{\phi_i(X_i)\}_{i=1}^3$ which maximize

$$\mathbb{E}[\phi_1(X_1) \phi_2(X_2) + \phi_1(X_1) \phi_3(X_3)].$$

Note that in this example the transformation function $\phi_1(X_1)$ appears in two terms in the objective function, leading to additional coupling constraints in the optimization (3), when compared to multiple bivariate MC optimizations (2). We investigate this optimization for a general graph $G = (V, E)$. In applications, we highlight appropriate selections of G in the NMC optimization¹.

Since $\mathbb{E}[\phi_i(X_i)^2] = 1$ for any $1 \leq i \leq n$, we have

$$\rho_G = |E| - \inf_{\phi_1, \dots, \phi_n} \frac{1}{2} \sum_{(i, i') \in E} \mathbb{E}[(\phi_i(X_i) - \phi_{i'}(X_{i'}))^2], \quad (4)$$

which means that the NMC optimization (3) is equivalent to finding functions of random variables that minimize the Mean Squared Error (MSE) among all neighboring random variables. The form (4) can be useful in different applications such as fitting a nonlinear regression model (e.g., $\phi_i(X_i) = \sum_{i' \in \{1, \dots, n\} \setminus \{i\}} \phi_{i'}(X_{i'}) + Z$).

The techniques described here to characterize local and global optima of the aforementioned NMC optimization can be used in other related formulations as well. For example, we also introduce absolute NMC, which maximizes the total absolute pairwise correlations among variables (Section 8.1). Absolute NMC is appropriate for applications where the strength of an association does not depend on the sign of the correlation coefficient. Moreover, the NMC optimization can be regularized to consider fewer nonlinear transformations or to restrict the set of possible transformations, further avoiding overfitting issues (Section 8.1).

In Section 2, for finite discrete random variables, we characterize the solution of the NMC optimization using a natural basis expansion. For finite discrete random variables, we show that the NMC optimization is an instance of the Maximum Correlation Problem (MCP), which is NP-hard [38], [39], [40], [41].

In Section 3, using the results from the Multivariate Eigenvector Problem (MEP) [38], we characterize necessary conditions that are satisfied at the global optimum of the NMC optimization. We propose an efficient algorithm based on Alternating Conditional Expectation (ACE) [21] that converges to a local optimum of the NMC optimization. We also provide guidelines for choosing appropriate starting points of the ACE algorithm to avoid local maximizers. The proposed algorithm does not require the formation of joint distribution matrices which could be expensive for variables with large number of categories. We develop a distributed version of the ACE algorithm to compute an approximation of the NMC value for large dense graphs using graph partitioning. We characterize a bound on the expected performance of the proposed algorithm for poly-growth graphs which are often used in analyzing distributed graph algorithms (see e.g., [42], [43], [44]).

In Section 4, we prove a finite sample bound and error guarantees for NMC. Under some conditions we prove that NMC of finite discrete random variables is continuous with respect to their joint probability distributions. That is, a small perturbation in the joint distribution results in a small change in the NMC value. Moreover, we show that the probability of discrepancy greater than any given level between NMC and NMC computed using empirical distributions decays exponentially fast as the sample size grows.

In Section 5, for jointly Gaussian variables, we use Hermite-Chebyshev polynomials as the basis for the functions of Gaussian random variables and characterize the solution of the NMC optimization. Under some conditions, we show that the NMC optimization is equivalent to the Max-Cut problem, which is NP-complete [45]. However, there exist algorithms to approximate its solution using Semidefinite Programming (SDP) within certain approximation factors [46]. Moreover, we provide sufficient conditions under which a solution of the NMC optimization can be characterized exactly.

In Section 6, we illustrate some applications of NMC. First we show an application of NMC in inference of graphical model (Definition 8) for bijective functions of jointly Gaussian variables. Graphical models provide a useful framework for characterizing dependencies among variables and for efficient computation of marginal and the mode of distributions [47]. Moreover, Gaussian graphical models play an important role in applications such as linear regression [48], partial correlation [49], maximum likelihood estimation [50], etc. Here we consider a setup in which observed variables are related to

¹We use the terminology *graph* and *network* interchangeably.

latent jointly Gaussian variables through link functions. These link functions are unknown, bijective, and can be linear or nonlinear. We show that, under some conditions, the inverse covariance matrix of the transformed variables computed by the NMC optimization (as well as multiple MC optimizations) characterizes the underlying graphical model. With an example we show that when underlying link functions are monotone the performance of our inference method is comparable to the one of the *copula* model developed by Liu et al. [11], [12]. Moreover, we illustrate that if we violate the model assumption for both our inference framework and for the copula method by considering non-monotone link functions, our inference method appears to outperform the copula one, highlighting the robustness of our proposed inference framework. Finally, we provide an example to demonstrate NMC's utility in a data application where we apply sample NMC to cancer datasets [51] and infer nonlinear gene association networks. Over these inferred networks, we deduce gene modules that are significantly associated with survival times of individuals while these modules are not detected using linear association measures.

2 NETWORK MAXIMAL CORRELATION

2.1 Review of Maximal Correlation

Recall the definition of maximal correlation (Definition 1). For $k = i, i'$, we let $\phi_k^*(X_k)$ denote a solution of optimization (1) should it exist. The existence and uniqueness of a solution ($\phi_i^*(X_i), \phi_{i'}^*(X_{i'})$) to the MC optimization (1) have been investigated in [21]. Maximal correlation $\rho(X_i, X_{i'})$ is always between 0 and 1, where a high MC value indicates a strong association between two variables [15].

Unlike Pearson's linear correlation [4], MC only depends on the joint distribution of the variables and not on their sample spaces $\mathcal{X}_i$. Several works have investigated different aspects of optimization (1) for both discrete [21] and continuous [20] random variables. For discrete variables, under mild conditions, MC is equal to the second largest singular value of the following Q -matrix [21]:

Definition 3. Let $P_{X_i, X_{i'}}$ be the joint distribution of discrete variables X_i and $X_{i'}$ with finite alphabet sizes $|\mathcal{X}_i|$ and $|\mathcal{X}_{i'}|$, respectively. Define a matrix $Q_{i, i'} \in \mathbb{R}^{|\mathcal{X}_i| \times |\mathcal{X}_{i'}|}$ whose (j, j') element is

$$Q_{i, i'}(j, j') \triangleq \frac{P_{X_i, X_{i'}}(j, j')}{\sqrt{P_{X_i}(j)P_{X_{i'}}(j')}}. \quad (5)$$

The matrix is called the Q -matrix of the distribution $P_{X_i, X_{i'}}$.

In this case, optimal transformations of variables can be characterized using right and left singular vectors of the normalized probability distribution matrix.

For Gaussian variables, Lancaster [20] introduced a basis expansion with Hermite-Chebyshev polynomials to compute MC. Interestingly in this case the maximal correlation and Pearson's linear correlation is equivalent.

2.2 NMC Formulation

Recall the definition of the Network Maximal Correlation (Definition 2). NMC infers nonlinear transformation functions assigned to each node variable so that the aggregate pairwise correlations over the graph $G = (V, E)$ is maximized.

Let $\phi_i^*(\cdot)$ be a solution of the NMC optimization (3) should it exist. Then, an edge maximal correlation of variables i and i' is defined as

$$\rho_G^{(i, i')}(X_1, \dots, X_n) \triangleq |\mathbb{E}[\phi_i^*(X_i) \phi_{i'}^*(X_{i'})]|,$$

where $(i, i') \in E$. Unlike the bivariate MC formulation of (1), the edge maximal correlation is a function of the joint distribution of all variables. Thus, an edge maximal correlation of variables X_i and $X_{i'}$ is always smaller than or equal to their bivariate maximal correlation, i.e.,

$$\rho_G^{(i, i')}(X_1, \dots, X_n) \leq \rho(X_i, X_{i'}).$$

We next provide a framework to study NMC and its properties. Following the definition provided in [18, Section 3], for $i = 1, 2, \dots, n$, we define a set of real-valued measurable functions H_{X_i} as

$$H_{X_i} = \{\phi_i \mid \phi_i : \mathcal{X}_i \rightarrow \mathbb{R}, \phi_i \text{ is Borel measurable}, \quad (6)$$

$$\mathbb{E}[\phi_i \circ X_i] < \infty, \mathbb{E}[(\phi_i \circ X_i)^2] < \infty\},$$

where the inner product of two elements of H_{X_i} is defined as $\mathbb{E}[\phi_i(X_i) \phi_{i'}(X_i)]$, and the norm is defined as $\sqrt{\mathbb{E}[(\phi_i \circ X_i)^2]}$. Note that $\phi_i(X_i)$ is a random variable in $\mathbb{R}$ such that for any Borel-measurable set B , we have $\mathbb{P}[\phi_i(X_i) \in B] = \mathbb{P}[X_i \in \phi_i^{-1}(B)]$. We use the notation $\phi_i \circ X_i$ and $\phi_i(X_i)$ interchangeably.

We let $\{\psi_{i, j}\}_{j=1}^\infty$ represent an orthonormal basis for H_{X_i} . We will explicitly construct such basis in the case of discrete random variables and Gaussian random variables, studied in this paper.

Theorem 1. Consider the following optimization:

$$\begin{aligned} \sup_{\{\{a_{i, j}\}_{j=1}^\infty\}_{i=1}^n} \quad & \sum_{(i, i') \in E} \sum_{j, j'} a_{i, j} a_{i', j'} \rho_{i, i'}^{j, j'} \\ & \sum_{j=1}^\infty a_{i, j}^2 = 1, \quad 1 \leq i \leq n, \\ & \sum_{j=1}^\infty a_{i, j} \mathbb{E}[\psi_{i, j}(X_i)] = 0, \quad 1 \leq i \leq n, \end{aligned} \quad (7)$$

where

$$\rho_{i, i'}^{j, j'} \triangleq \mathbb{E}[\psi_{i, j}(X_i) \psi_{i', j'}(X_{i'})]. \quad (8)$$

Let $\phi_i^*(\cdot)$ and $\{a_{i, j}^*\}_{j=1}^\infty$ for $1 \leq i \leq n$ be solutions of optimizations (3) and (7), respectively. Then,

$$\phi_i^*(X_i) = \sum_{j=1}^\infty a_{i, j}^* \psi_{i, j}(X_i). \quad (9)$$

is a solution of the NMC optimization (3). Moreover, $\rho_{i, i'}^{j, j'}$ for $j, j' \geq 1$ are coefficients of the basis expansion of $P_{X_i, X_{i'}}$ with respect to the basis $\{\psi_{i, j} \psi_{i', j'}\}_{j, j'}$, i.e.,

$$P_{X_i, X_{i'}}(x_i, x_{i'}) = \sum_{j, j'} \rho_{i, i'}^{j, j'} \psi_{i, j}(x_i) \psi_{i', j'}(x_{i'}).$$

Proof A proof is presented in Section 7.1. ■

Equation (9) means convergence in the L_2 sense, i.e., $\lim_{n \rightarrow \infty} \|\phi_i^*(X_i) - \sum_{j=1}^n a_{i, j}^* \psi_{i, j}(X_i)\| = 0$. Throughout the paper since we only work with the inner product of functions, without loss of generality, we replace $\phi_i^*(X_i)$ by $\sum_{j=1}^\infty a_{i, j}^* \psi_{i, j}(X_i)$. Also note that for the case of finite discrete random variables, the convergence is in fact pointwise. For further details, see [20], [52].

Selecting appropriate set of functions H_{X_i} is critical to have a tractable optimization problem (7). In the following, we consider the NMC optimization for finite discrete variables, while the Gaussian case is discussed in Section 5.

3 NMC FOR FINITE DISCRETE VARIABLES

In this section, we analyze the NMC optimization (3) for finite discrete random variables, and then introduce an efficient algorithm to compute NMC. We then propose a parallelizable version of the NMC algorithm based on network partitioning and characterize a bound for its expected performance.

First we introduce some notation. For any vector $\mathbf{v} = (v_1, \dots, v_d) \in \mathbb{R}^d$ and $p \geq 1$, we let $\|\mathbf{v}\|_p$ represent the standard p -norm of the vector $\mathbf{v}$ defined as

$$\|\mathbf{v}\|_p = \left(\sum_{i=1}^d v_i^p \right)^{\frac{1}{p}}.$$

For $p = 2$, we drop the subscript, i.e., $\|\mathbf{v}\| = \|\mathbf{v}\|_2$. The infinite norm of a vector is defined as

$$\|\mathbf{v}\|_\infty = \max_{1 \leq i \leq d} v_i.$$

The inner product between two vectors $\mathbf{v}$ and $\mathbf{w}$ is defined as

$$\langle \mathbf{v}, \mathbf{w} \rangle = \sum_{i=1}^d v_i w_i.$$

For a matrix $V \in \mathbb{R}^{d_1} \times \mathbb{R}^{d_1}$, the matrix norm is a vector norm on $\mathbb{R}^{d_1 \times d_1}$. I.e.,

$$\|V\|_p = \sup_{\mathbf{w} \neq 0} \frac{\|V\mathbf{w}\|_p}{\|\mathbf{w}\|_p}.$$

For $p = 2$, we drop the subscript.

3.1 Relationship Between NMC and Maximum Correlation Problem

Let X_i be a discrete random variable with alphabet $\{1, \dots, |\mathcal{X}_i|\}$. Let

$$\psi_{i,j}(x) = \mathbf{1}\{x = j\} \frac{1}{\sqrt{P_{X_i}(j)}} \quad (10)$$

be an orthonormal basis for H_{X_i} , where $\mathbf{1}$ is the indicator function. We assume that all the elements of the alphabet $x_i \in \mathcal{X}_i$ have positive probabilities, as otherwise they can be neglected without loss of generality. We can write

$$\rho_{i,i'}^{j,j'} = \mathbb{E}[\psi_{i,j}(X_i) \psi_{i',j'}(X_{i'})] = \frac{P_{X_i, X_{i'}}(j, j')}{\sqrt{P_{X_i}(j) P_{X_{i'}}(j')}}.$$

Therefore, the optimization (7) is simplified to the following:

$$\begin{aligned} \max_{\{a_{i,j}\}_{i=1}^n} \sum_{(i,i') \in E} \sum_{j,j'} a_{i,j} a_{i',j'} \frac{P_{X_i, X_{i'}}(j, j')}{\sqrt{P_{X_i}(j) P_{X_{i'}}(j')}} \quad (11) \\ \sum_{j=1}^{|\mathcal{X}_i|} (a_{i,j})^2 = 1, \quad 1 \leq i \leq n, \\ \sum_{j=1}^{|\mathcal{X}_i|} a_{i,j} \sqrt{P_{X_i}(j)} = 0, \quad 1 \leq i \leq n. \end{aligned}$$

For $i = 1, \dots, n$, let

$$\begin{aligned} \mathbf{a}_i &\triangleq (a_{i,1}, a_{i,2}, \dots, a_{i,|\mathcal{X}_i|})^T \\ \sqrt{\mathbf{p}_i} &\triangleq (\sqrt{P_{X_i}(1)}, \sqrt{P_{X_i}(2)}, \dots, \sqrt{P_{X_i}(|\mathcal{X}_i|)})^T. \end{aligned} \quad (12)$$

Therefore, the optimization (11) can be re-written as follows:

$$\begin{aligned} \max_{\{\mathbf{a}_i\}_{i=1}^n} \sum_{(i,i') \in E} \mathbf{a}_i^T Q_{i,i'} \mathbf{a}_{i'} \quad (13) \\ \|\mathbf{a}_i\|_2 = 1, \quad 1 \leq i \leq n, \\ \mathbf{a}_i \perp \sqrt{\mathbf{p}_i}, \quad 1 \leq i \leq n, \end{aligned}$$

where $Q_{i,i'}$ is defined in Definition 3 and $\perp$ represents orthogonality between two vectors.

The optimization (13) is not convex nor concave in general. Below, we show that this optimization is an instance of the Maximum Correlation Problem (MCP) proposed by Hotelling [39], [40]. By making this connection, we use established techniques via the Multivariate Eigenvalue Problem (MEP) to solve optimization (13).

For each i , since $I_{|\mathcal{X}_i|} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_i}^T$ is positive semidefinite, we take its square root² and write

$$B_i \triangleq \sqrt{I_{|\mathcal{X}_i|} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_i}^T}, \quad (14)$$

where $I_{|\mathcal{X}_i|}$ is a $|\mathcal{X}_i| \times |\mathcal{X}_i|$ identity matrix. Let $\mathbf{b}_i = B_i \mathbf{a}_i$. Let $U_i \Sigma_i U_i^T$ be the singular value decomposition of B_i where $U_i^{(j)}$ is the j -th column of U_i and $\sigma_i^{(j)}$ is the j -th singular value of B_i . We will show that only one singular value of B_i is zero which corresponds to the singular vector $\sqrt{\mathbf{p}_i}$. Without loss of generality, suppose $\sigma_i^{(1)} = 0$, for all i . Define A_i , a $|\mathcal{X}_i| \times |\mathcal{X}_i|$ matrix, as follows:

$$A_i \triangleq \left(\begin{bmatrix} U_i^{(2)}, \dots, U_i^{(|\mathcal{X}_i|)} \end{bmatrix} \text{diag} \left(\frac{1}{\sigma_i^{(2)}}, \dots, \frac{1}{\sigma_i^{(|\mathcal{X}_i|)}} \right) \begin{bmatrix} U_i^{(2)}, \dots, U_i^{(|\mathcal{X}_i|)} \end{bmatrix}^T \right) \quad (15)$$

As we show in the proof of Theorem 2, $\sigma_i^{(j)} \neq 0$, for all $1 \leq i \leq n$, and $j \geq 2$, A_i is well-defined according to (15).

Theorem 2. The NMC optimization (13) can be re-written as follows:

$$\begin{aligned} \max_{\mathbf{b}_1, \dots, \mathbf{b}_n} \sum_{(i,i') \in E} \mathbf{b}_i^T A_i^T (Q_{i,i'} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_{i'}}^T) A_{i'} \mathbf{b}_{i'} \quad (16) \\ \text{s.t. } \|\mathbf{b}_i\|_2 = 1 \quad 1 \leq i \leq n. \end{aligned}$$

Proof A proof is presented in Section 7.2. ■

Let C be a matrix consisting of submatrices $C_{i,i'} \in \mathbb{R}^{|\mathcal{X}_i| \times |\mathcal{X}_{i'}|}$ where if $(i, i') \in E$,

$$C_{i,i'} \triangleq A_i^T (Q_{i,i'} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_{i'}}^T) A_{i'}, \quad (17)$$

otherwise $C_{i,i'}$ is an all zero matrix of size $|\mathcal{X}_i| \times |\mathcal{X}_{i'}|$. Note that since the graph $G = (V, E)$ does not have self-loops, $C_{i,i}$ is a zero matrix for $1 \leq i \leq n$.

Let $\mathbf{b} \triangleq (\mathbf{b}_1^T, \dots, \mathbf{b}_n^T)^T \in \mathbb{R}^{|\mathcal{X}|}$, where $|\mathcal{X}| = \sum_{i=1}^n |\mathcal{X}_i|$ and $\mathbf{b}_i \in \mathbb{R}^{|\mathcal{X}_i|}$.

Corollary 1. The NMC optimization (16) can be written as follows:

$$\begin{aligned} \max_{\mathbf{b}} \mathbf{b}^T C \mathbf{b} \quad (18) \\ \|\mathbf{b}_i\|_2 = 1, \quad 1 \leq i \leq n. \end{aligned}$$

The optimization (18) is in the standard form of the MCP problem, proposed by Hotelling [39], [40], to find the linear

²The square root of a symmetric positive semidefinite matrix A is defined as $\sqrt{A} = U \Sigma^{1/2} U^T$ where $A = U \Sigma U^T$.

Algorithm 1 Gauss-Seidel Algorithm for MEP

Input: $C \in \mathbb{R}^{|\mathcal{X}|} \times \mathbb{R}^{|\mathcal{X}|}$.
Initialization: $\mathbf{b}^{(0)} \in \mathbb{R}^{|\mathcal{X}|}$.
for $k = 0, 1, \dots$
 for $i = 1, \dots, n$
 $\tilde{\mathbf{b}}_i^{(k)} = \sum_{i'=1}^{i-1} C_{i,i'} \mathbf{b}_{i'}^{(k+1)} + \sum_{i'=i+1}^n C_{i,i'} \mathbf{b}_{i'}^{(k)}$.
 $\mathbf{b}_i^{(k+1)} = \frac{\tilde{\mathbf{b}}_i^{(k)}}{\|\tilde{\mathbf{b}}_i^{(k)}\|_2}$
 end
end

combination of one set of variables that correlates maximally with the linear combination of another set of variables.

Definition 4 (Multivariate Eigenvalue Problem [38]). The first-order optimality condition for optimization (18) is the existence of real-valued scalars, namely, Lagrange multipliers $\lambda_1, \dots, \lambda_n$, such that the following system of equations is satisfied:

$$\sum_{i'=1}^n C_{i,i'} \mathbf{b}_{i'} = \lambda_i \mathbf{b}_i, \quad 1 \leq i \leq n$$

$$\|\mathbf{b}_i\|_2 = 1, \quad 1 \leq i \leq n. \quad (19)$$

This system of equations is called Multivariate Eigenvalue Problem (MEP) ([38], [41]).

3.2 An ACE Approach to Compute NMC

In Section 3.1, we established a connection among the NMC optimization (3), the Maximum Correlation Problem (MCP) and the Multivariate Eigenvalue Problem (MEP) (see e.g., [38], [39], [40], [41]). After showing that the NMC optimization can be reformulated as the MCP, we use existing techniques from the literature to solve it. Several local maximizers exist for cases where finding a global optimum of optimization (18) is computationally difficult [38], [53]. For example, an aggregated power method that iterates on blocks of C was proposed by Horst [54] as a general technique for solving the MEP (Definition 4) numerically.

Below, we summarize general algorithmic ideas to solve MCP:

- (1) First, an efficient algorithm is used to solve MEP which is the necessary first order condition for MCP. This step is studied in references [38], [54].
- (2) Next, a strategy is used to choose starting points of the algorithm or to jump out of the local minima of optimization (18). We adopt this step from [41], [55].

An efficient algorithm to solve MEP: Algorithm 1 proposed by [38] is a Gauss-Seidel algorithm [56] to solve MEP, which is a variant of the classical power iteration method (see e.g. [57]). Let

$$r(\mathbf{b}) = \mathbf{b}^T C \mathbf{b} = \sum_{i,i'=1}^n \mathbf{b}_i^T C_{i,i'} \mathbf{b}_{i'} \quad (20)$$

$$\lambda_i(\mathbf{b}) = \mathbf{b}_i^T [C_{i,1}, \dots, C_{i,n}] \mathbf{b} = \sum_{i'=1}^n \mathbf{b}_i^T C_{i,i'} \mathbf{b}_{i'} \quad (21)$$

Algorithm 1 is an iterative algorithm. Let $\mathbf{b}^{(k)} = (\mathbf{b}_1^{(k)}, \dots, \mathbf{b}_n^{(k)})$ be the update of Algorithm 1 at iteration k .

Algorithm 2 Network ACE to compute NMC

Input: $G, X_1, \dots, X_n$,
Initialization: $\phi_1^{(0)}(X_1), \dots, \phi_n^{(0)}(X_n)$ with mean zero and unit variance.
for $k = 0, 1, \dots$
 for $i = 1, \dots, n$
 $\phi_i^{(k)}(X_i) = \mathbb{E} \left[\sum_{i'=1}^{i-1} \phi_{i'}^{(k+1)}(X_{i'}) | X_i \right] + \mathbb{E} \left[\sum_{i'=i+1}^n \phi_{i'}^{(k)}(X_{i'}) | X_i \right]$, for $i' \in \mathcal{N}(i)$.
 update: $\phi_i^{(k+1)}(X_i) = \frac{\tilde{\phi}_i^{(k)}(X_i)}{\sqrt{\mathbb{E}[\tilde{\phi}_i^{(k)}(X_i)^2]}}$
end
 $\rho_G^{(k+1)} = \sum_{(i,i') \in E} \mathbb{E} [\phi_i^{(k+1)}(X_i) \phi_{i'}^{(k+1)}(X_{i'})]$
end

Theorem 3 (Theorem 5.1 [41]). The sequence $\{r(\mathbf{b}^{(k)})\}$ generated by Algorithm 1 is monotonically increasing and convergent.

Theorem 3 indicates that the Algorithm 1 finds a local optimum of the optimization (18). According to Proposition 1, this solution provides a local optimum for the NMC optimization (3). In the following we introduce a strategy that uses a local optimum of the optimization (18) and constructs another solution for the optimization (18) with strictly higher objective function value.

A strategy for avoiding local optima of MCP:

Proposition 1. Let $\bar{\mathbf{b}}$ be a solution of the MEP (19). Suppose that there exists an $1 \leq i \leq n$ such that $\lambda_i(\bar{\mathbf{b}}) < 0$. Let $\hat{\mathbf{b}} = (\hat{\mathbf{b}}_1, \dots, \hat{\mathbf{b}}_n)$ be defined as: $\hat{\mathbf{b}}_i = -\bar{\mathbf{b}}_i$, for any i such that $\lambda_i(\bar{\mathbf{b}}) < 0$, and $\hat{\mathbf{b}}_{i'} = \bar{\mathbf{b}}_{i'}$, otherwise. Then, we have $r(\hat{\mathbf{b}}) > r(\bar{\mathbf{b}})$.

Proof A proof is presented in Section 7.3. ■

Algorithm 1 finds a local optimum $\bar{\mathbf{b}}$ for the optimization (18). This solution can be translated to a solution of the NMC optimization (3) according to Theorem 1, equations (10), (14), and using $\mathbf{b}_i = B_i \mathbf{a}_i$. This leads to a direct algorithm to find a solution for the NMC optimization (3) based on alternating conditional expectation (Algorithm 2). Similarly to Algorithm 1, Algorithm 2 converges to the local optimum of the NMC optimization (3). We use a strategy similarly to the one of Proposition 1 to avoid remaining at local maximizers. At each iteration of the algorithm, we update transformation functions as follows: Suppose at iteration k , transformation functions are $\{\phi_i^{(k)}\}_{i=1}^n$. If we fix all variables except the transformation function of node i , an optimizer of $\phi_i^{(k+1)}$ can be written as the normalized conditional expectation of functions of its neighbors. In each update, the objective function of the NMC optimization increases or remains the same. Note that, for the bivariate case (i.e., $n = 2$), Algorithm 2 is simplified to the ACE algorithm [21] for the MC computation. If the number of nodes n is large, then computation of NMC may be expensive. In the following, we propose an approach to compute NMC using parallel computation.

3.3 Parallel Computation of NMC

For large and dense networks, exact computation of NMC may become computationally expensive. For those cases, we propose a parallelizable algorithm which approximates NMC using network partitioning. The idea can be described as follow. For a given graph $G = (V, E)$,

- (1) Partition the graph into small disjoint sets.
- (2) Estimate NMC for each partition independently.
- (3) Combine NMC solutions over sub-graphs to form an approximation of NMC for the original graph.

A k -partition $\pi(k, V, E)$ of graph $G = (V, E)$ is defined as a set $\{V_1, \dots, V_M\}$ such that, for any $1 \leq i \neq j \leq M$, $V_i \cap V_j = \emptyset$, $\bigcup_{i=1}^M V_i = V$, and $1 \leq |V_i| \leq k$. We say an edge $e \in E$ belongs to $\pi(k, V, E)$ if $e \in \bigcup_{i=1}^M V_i \times V_i$.

Definition 5. A k -partitioning of graph $G = (V, E)$ is a set $\Pi(k, V, E) = \{\pi_1(k, V, E), \pi_2(k, V, E), \dots\}$ where each $\pi_i(k, V, E)$ is a k -partition of the graph.

Next we define an (ϵ, k) -partitioning of graph $G = (V, E)$.

Definition 6 ([58], Section A). An (ϵ, k) -partitioning of graph $G = (V, E)$ is a distribution over all k -partitioning of the graph such that, for any $e \in E$, $\mathbb{E}[\mathbf{1}\{e \notin \pi_i(k, V, E)\}] \leq \epsilon$.

Definition 7. A graph G is poly-growth if there exists $r > 0$ and $C > 0$, such that for any vertex v in the graph,

$$|N_v(d)| \leq C d^r,$$

where $N_v(d)$ is the number of nodes within distance d of v in G (distance is defined as the shortest path length on the graph).

Reference [58] describes the following procedure for generating an (ϵ, k) -partitioning on a graph:

1. Given $G = (V, E)$, k , and $\epsilon > 0$, we define the truncated geometric distribution as follows:

$$\mathbb{P}[x = l] = \begin{cases} \epsilon(1 - \epsilon)^{l-1}, & l < k, \\ (1 - \epsilon)^{k-1}, & l = k. \end{cases} \quad (22)$$

2. We then order nodes arbitrarily $1, \dots, N$. For node i , we sample R_i according to distribution (22) and assign all nodes within that distance from node i to color i . If a node is already colored, we re-color it with color i .
3. All nodes with the same color form a k -partition of the graph.

Proposition 2 ([58], Lemma 2). If G is a poly-growth graph, then by selecting $k = \Theta(\frac{r}{\epsilon} \log \frac{r}{\epsilon})$, the above procedure results in an (ϵ, Ck^r) -partitioning.

Next, we use an (ϵ, k) -graph partitioning to approximate NMC over large graphs using parallel computations. Consider the following approach:

- (1) We sample a partition $\{V_1, \dots, V_M\}$ of V , given an (ϵ, k) -partitioning of G .
- (2) For each partition $1 \leq m \leq M$, we compute NMC restricted to $G_m = (V_m, E \cap (V_m \times V_m))$, denoted by $\hat{\rho}_{G_m}$.
- (3) Let $\hat{\rho}_G = \sum_{m=1}^M \hat{\rho}_{G_m}$ be an approximation of ρ_G .

In the following, we bound the approximation error by bounding boundary effects:

Theorem 4. Consider an (ϵ, k) -partitioning of the graph G . We have,

$$\mathbb{E}[\hat{\rho}_G] \geq (1 - \epsilon)\rho_G, \quad (23)$$

where the expectation is over (ϵ, k) -partitioning of graph G .

Proof A proof is presented in Section 7.4. ■

4 PROPERTIES OF NETWORK MAXIMAL CORRELATION FOR FINITE DISCRETE RANDOM VARIABLES

In many applications, often only noisy measurements or samples of joint distributions are observed. In this section, we prove a finite sample generalization bound, and error guarantees for Network Maximal Correlation of finite discrete random variables. Specifically, under general conditions, we prove that Network Maximal Correlation is a continuous measure with respect to the joint probability distribution. That is, a small perturbation in the distribution results in a small change in the NMC value. Moreover we prove that the probability of discrepancy between NMC and sample NMC (NMC computed using empirical distributions) greater than any given level decays exponentially fast as the sample size m grows.

Throughout this subsection we only consider finite discrete random variables. Moreover to simplify notation we let P be the matrix representation of probability distribution $P_{X_1, \dots, X_n}$. We assume that all the elements of the alphabet $x_i \in \mathcal{X}_i$ have positive probabilities, as otherwise they can be neglected without loss of generality, and define

$$\delta_{X_i}(P) \triangleq \arg \min_{x_i \in \mathcal{X}_i} P_{X_i}(x_i) > 0, \quad 1 \leq i \leq n. \quad (24)$$

The empirical distribution of these variables using m observed samples is defined as $P^{(m)}(x_1, \dots, x_n) = \frac{1}{m} \sum_{s=1}^m \mathbf{1}\{x_1^{(s)} = x_1, \dots, x_n^{(s)} = x_n\}$, where $\{(x_1^{(s)}, \dots, x_n^{(s)})\}_{s=1}^m$ are i.i.d. samples drawn according to the distribution $P_{X_1, \dots, X_n}$. The vector of observed samples of variable X_i is denoted by $\mathbf{x}_i = (x_i^{(1)}, x_i^{(2)}, \dots, x_i^{(m)})$.

4.1 Continuity of Network Maximal Correlation

Let P and $\tilde{P}$ be two distributions over alphabets $(\mathcal{X}_1, \dots, \mathcal{X}_n)$. Let $K = \max_{1 \leq i \leq n} |\mathcal{X}_i|$. Thus, $\mathcal{K} \triangleq K^n$ is an upper bound on the alphabet size of the joint distribution. In the following, we show that if the infinity norm distance between P and $\tilde{P}$ is small (i.e., $\|P - \tilde{P}\|_\infty = \max_{x_1 \in \mathcal{X}_1, \dots, x_n \in \mathcal{X}_n} |P(x_1, \dots, x_n) - \tilde{P}(x_1, \dots, x_n)| \leq \gamma$), their corresponding NMC values (denoted by ρ_G and $\tilde{\rho}_G$ respectively) are close to each other.

Theorem 5. Network Maximal Correlation is a continuous function of the joint probability distribution P . Let $\|P - \tilde{P}\|_\infty \leq \gamma$, for $\gamma \leq \delta^{3/2} \mathcal{K}^{-1}$. Then, we have

$$|\rho_G - \tilde{\rho}_G| \leq \gamma \mathcal{K} |E| \frac{8}{\delta^2}, \quad (25)$$

where $\delta = \min_{1 \leq i \leq n} (\min\{\delta_{X_i}(P), \delta_{X_i}(\tilde{P})\})$.

Proof A proof is presented in Section 7.5. ■

Corollary 2. Let ρ and $\tilde{\rho}$ be bivariate MCs with respect to distributions P and $\tilde{P}$, respectively. Let $\|P - \tilde{P}\|_\infty \leq \gamma$. For any $\gamma \leq \delta^{3/2} \mathcal{K}^{-1}$, we have

$$|\rho - \tilde{\rho}| \leq \gamma \mathcal{K} \frac{8}{\delta^2}, \quad (26)$$

where $\mathcal{K}$ and δ are defined according to Theorem 5.

4.2 Sample NMC

Let $\{(x_1^{(s)}, \dots, x_n^{(s)})\}_{s=1}^m$ be i.i.d. samples drawn according to a distribution $P_{X_1, \dots, X_n}$. Let $P^{(m)}$ denote the empirical distribution obtained from these samples. Network Maximal Correlation computed using this empirical probability distribution is called *Sample Network Maximal Correlation* and is denoted by $\rho_G^{(m)}(\{(x_1^{(s)}, \dots, x_n^{(s)})\}_{s=1}^m)$. For simplicity, when no confusion arises, we refer to the sample NMC by $\rho_G^{(m)}$. In the following, we show that the probability of discrepancy greater than any given value between $\rho_G^{(m)}$ and ρ_G decays exponentially fast as the sample size m grows.

Theorem 6. For any $\eta, \epsilon > 0$, if

$$m \geq 2 \left(\frac{32|E|\mathcal{K}}{\epsilon\delta'^2} \right)^2 \log \left(\frac{8n}{\eta} \right), \quad (27)$$

then we have

$$\mathbb{P}[|\rho_G^{(m)} - \rho_G| > \epsilon] \leq \eta, \quad (28)$$

where $\delta' = \min_{1 \leq i \leq n} \delta_{X_i}(P)$.

Proof A proof is presented in Section 7.6. ■

Note that the number of required samples to learn the joint probability distribution reliably is a function of the alphabet size. This is reflected in the right hand side of the bound provided in Theorem 6 through the term $\mathcal{K}$.

Corollary 3. Let ρ and $\rho^{(m)}$ be bivariate MC and sample MC, respectively. For

$$m \geq 2 \left(\frac{32\mathcal{K}}{\epsilon\delta'^2} \right)^2 \log \left(\frac{16}{\eta} \right). \quad (29)$$

we have

$$\mathbb{P}[|\rho^{(m)} - \rho| > \epsilon] \leq \eta, \quad (30)$$

where $\mathcal{K}$ and δ' are defined according to Theorem 6.

5 NMC FOR JOINTLY GAUSSIAN VARIABLES

Suppose that $(X_1, \dots, X_n)$ are jointly Gaussian variables with zero means and unit variances. Let $\rho_{i,i'}$ be the correlation coefficient of variables X_i and $X_{i'}$. We assume that $|\rho_{i,i'}| \neq 1$ if $i \neq i'$. Let $G_c = (V_c, E_c)$ be the covariance graph corresponding to these variables where $V_c = \{1, 2, \dots, n\}$, and $(i, i') \in E_c$ iff $\rho_{i,i'} \neq 0$. Moreover for continuous variables $\{\phi_i(\cdot)\}_{i=1}^n$ in optimization (3) are assumed to be continuous and l_2 (equivalent to L_2 for discrete random variables).

The k -th Hermite-Chebyshev polynomial [20] is defined as

$$\Psi_k(x) \triangleq (-1)^k e^{x^2} \frac{d^k}{dx^k} e^{-x^2}. \quad (31)$$

These polynomials form an orthonormal basis with respect to Gaussian distributions [20]. That is,

$$\int_{-\infty}^{\infty} \Psi_j(x_i) \Psi_{j'}(x_{i'}) p(x_i, x_{i'}) dx_i dx_{i'} = (\rho_{i,i'})^j \mathbf{1}\{j = j'\}, \quad (32)$$

where $p(x_i, x_{i'})$ is the joint density function of Gaussian variables X_i and $X_{i'}$ with correlation $\rho_{i,i'}$. Let $\psi_{i,j}$ to be the j -th Hermite-Chebyshev polynomial, for $1 \leq i \leq n$. We have

$$\begin{aligned} \rho_{i,i'}^{j,j'} &= \mathbb{E}[\psi_{i,j}(X_i) \psi_{i',j'}(X_{i'})] \\ &= (\rho_{i,i'})^j \mathbf{1}\{j = j'\}. \end{aligned} \quad (33)$$

Moreover, using the definition of Hermite-Chebyshev polynomials (31), we have

$$\mathbb{E}[\psi_{i,j}(X_i)] = \mathbf{1}\{j = 0\}, \quad 1 \leq i \leq n. \quad (34)$$

because all of these functions for $j \geq 1$ have zero means when integrated against a Gaussian distribution. Therefore, using (33) and (34), the optimization (7) can be written as

$$\begin{aligned} \max_{\{a_{i,j}\}_{i=1}^n, \{j=1}^{\infty}} \sum_{(i,i') \in E} \sum_{j=1}^{\infty} a_{i,j} a_{i',j} (\rho_{i,i'})^j \\ \sum_{j=1}^{\infty} (a_{i,j})^2 = 1, \quad 1 \leq i \leq n. \end{aligned} \quad (35)$$

We establish that solving the optimization (35) is NP-complete by identifying one instance of this problem that reduces to the max-cut problem, which is NP-complete [45].

Theorem 7. Let $s_i \in \{-1, 1\}$ for $1 \leq i \leq n$. Suppose

$$\sum_{i' \neq i} (1 - s_i s_{i'}) \rho_{i,i'} \geq 0, \quad \forall 1 \leq i \leq n,$$

and

$$\sum_{i' \neq i} s_i s_{i'} \rho_{i,i'} \geq \sum_{i' \neq i} (\rho_{i,i'})^2, \quad \forall 1 \leq i \leq n,$$

then, $\mathbf{a}_i^* = (0, s_i, 0, \dots, 0)$, for $1 \leq i \leq n$ is a global maximizer of the NMC optimization (35) over a complete graph without self-loops.

Proof A proof is presented in Section 7.7. ■

Proposition 3. Under assumptions of Theorem 7, the NMC optimization (3) is simplified to the following max-cut optimization:

$$\begin{aligned} \max_{s_i} \sum_{i \neq i'} s_i s_{i'} \rho_{i,i'} \\ s_i \in \{-1, 1\}, \quad 1 \leq i \leq n. \end{aligned} \quad (36)$$

Moreover, for all $1 \leq i \leq n$, we have $\phi_i^*(X_i) = s_i^* X_i$, where ϕ_i^* and s_i^* are solutions of optimizations (2) and (36), respectively.

Proof A proof is presented in Section 7.8. ■

For bivariate jointly Gaussian variables, the conditions of Proposition 3 are always satisfied. For multivariate jointly Gaussian variables however an optimal NMC solution $\phi_i^*(X_i)$ can be different than $\pm X_i$. Proposition 3 provides conditions under which $\phi_i^*(X_i) = \pm X_i$ for the multivariate setup.

In general, the Max-Cut optimization (36) is NP-complete [45]. However, there exist algorithms to approximate its solution using semidefinite programming (SDP) [46].

Assumption 1. Let $(X_1, \dots, X_n)$ be jointly Gaussian variables, with zero means and unit variances such that $\rho_{i,i'} \neq 1$ if $i \neq i'$, and

$$\sum_{i' \neq i} \rho_{i,i'} \geq \sum_{i' \neq i} (\rho_{i,i'})^2, \quad \forall 1 \leq i \leq n, \quad (37)$$

where $\rho_{i,i'}$ is the correlation coefficient of X_i and $X_{i'}$.

If all correlation coefficients are non-negative (i.e., $\rho_{i,i'} \geq 0$ for $1 \leq i, i' \leq n$), Assumption 1 is immediately satisfied. However, Assumption 1 is more general as some correlation coefficients can be negative and the condition will still hold.

Corollary 4. Under Assumption 1, $\phi_i^*(X_i) = X_i$ is a solution of the NMC optimization (35).

6 ILLUSTRATION OF NMC'S USE

In this section, we illustrate some applications of NMC.

6.1 Inference of Graphical Models for Functions of Gaussian Variables

Suppose that $(X_1, \dots, X_n)$ are jointly Gaussian variables with the covariance matrix Λ_X . Without loss of generality, we assume all variables have zero means and unit variances. I.e., $\mathbb{E}[X_i] = 0$ and $\mathbb{E}[X_i^2] = 1$, for all $1 \leq i \leq n$.

Definition 8. The graphical model $G_X = (V_X, E_X)$ is defined such that if $(i, i') \notin E_X$, then

$$X_i \perp\!\!\!\perp X_{i'} | \{X_k, k \neq i, i'\}, \quad (38)$$

where $\perp\!\!\!\perp$ represents independence between variables.

Let J_X be the information (precision) matrix [47] of these variables where $J_X = \Lambda_X^{-1}$.

Theorem 8 ([47], Example 3.3). For jointly Gaussian variables, $(i, i') \in E_X$ if and only if $J_X(i, i') \neq 0$.

This Theorem has been stated in other references as well (e.g. [59]). Theorem 8 represents a way to model explicitly the joint distribution of Gaussian variables using a graphical model $G_X = (V_X, E_X)$. This result is critical in several applications involved with Gaussian variables which requires computation of marginal distributions, or computation of the mode of the distribution. These computations can be performed efficiently over the graphical model using belief propagation approaches [47]. Moreover, Gaussian graphical models play an important role in many applications such as linear regression [48], partial correlation [49], maximum likelihood estimation [50], etc. In many applications, even if variables are not jointly Gaussian, a Gaussian approximation is used often, partially owing to the efficient inference of their graphical models.

In the following, under some conditions, we use the multiple MC (2) and NMC (3) optimizations to characterize graphical models for functions of latent jointly Gaussian variables. These functions are unknown, bijective, and can be linear or nonlinear. More precisely, let $Y_i = f_i(X_i)$, where $f_i : \mathbb{R} \rightarrow \mathbb{R}$ is a bijective function. Our goal is to characterize a graphical model for the variables $(Y_1, Y_2, \dots, Y_n)$ without knowledge of the functions $f_i(\cdot)$.

Consider the following optimizations:

$$\sup_{g_{i,i'}, g_{i',i}} \mathbb{E}[g_{i,i'}(Y_i) g_{i',i}(Y_{i'})], \quad (39)$$

such that $g_{i,i'} : \mathcal{Y}_i \rightarrow \mathbb{R}$ is Borel measurable, $\mathbb{E}[g_{i,i'}(Y_i)] = 0$, and $\mathbb{E}[g_{i,i'}(Y_i)^2] = 1$, for $1 \leq i, i' \leq n$, and

$$\begin{aligned} \max_{g_i : \mathbb{R} \rightarrow \mathbb{R}} \sum_{(i,i')} \mathbb{E}[g_i(Y_i) g_{i'}(Y_{i'})], \quad (40) \\ \mathbb{E}[g_i(Y_i)] = 0, \quad 1 \leq i \leq n, \\ \mathbb{E}[g_i^2(Y_i)] = 1, \quad 1 \leq i \leq n, \end{aligned}$$

such that $g_i : \mathcal{Y}_i \rightarrow \mathbb{R}$ is Borel measurable, $\mathbb{E}[g_i(Y_i)] = 0$, and $\mathbb{E}[g_i(Y_i)^2] = 1$, for $1 \leq i \leq n$. Optimization (39) solves multiple MC between all pairs of variables while optimization (40) solves NMC considering a fully connected graph. Suppose $g_{i,i'}^*$ and g_i^* represent solutions of (39) and (40), respectively.

Theorem 9. $g_{i,i'}^*(Y_i) = \pm X_i$ for all $1 \leq i, i' \leq n$ is a solution of (39). Moreover, if Assumption 1 holds, $g_i^*(Y_i) = X_i$ for all $1 \leq i \leq n$ is a solution of (40).

Fig. 1: Relationships between jointly Gaussian variables X_i and observations Y_i for example 1 (panel a), and for example 2 (panel b).

Proof A proof is presented in Section 7.9. ■

In the multiple MC optimization (39), each variable Y_i is assigned to $n - 1$ transformation functions $\{g_{i,i'}^*(Y_i) : 1 \leq i' \neq i \leq n\}$. However, when variables $X_1, \dots, X_n$ are jointly Gaussian, all these functions are equal to $\pm X_i$. In general this is not true for non-Gaussian distributions. On the other hand, when the Assumption 1 holds, a solution of the NMC optimization (40) recovers the sign of latent variables as well.

We define the matrices $\Lambda_{MultiMC}$ and Λ_{NMC} by

$$\begin{aligned} \Lambda_{MultiMC}(i, i') &= \mathbb{E}[g_{i,i'}^*(Y_i) g_{i',i}^*(Y_{i'})] \\ \Lambda_{NMC}(i, i') &= \mathbb{E}[g_i^*(Y_i) g_{i'}^*(Y_{i'})]. \end{aligned}$$

Moreover, we let $J_{MultiMC} = \Lambda_{MultiMC}^{-1}$ and $J_{NMC} = \Lambda_{NMC}^{-1}$. We define $G_{NMC} = (V_{NMC}, E_{NMC})$ such that $(i, j) \in E_{NMC}$ if and only if $J_{NMC}(i, j) \neq 0$. We also let $\mathcal{G}_{NMC}$ be the set of all possible G_{NMC} since the solution of the optimization (40) may not be unique. Similarly, we define $G_{MultiMC}$ and $\mathcal{G}_{MultiMC}$.

Corollary 5. Let G_Y be a graphical model of variables $Y_i = f_i(X_i)$ according to Definition 8, where $f_i : \mathbb{R} \rightarrow \mathbb{R}$ is a bijective function, for $1 \leq i \leq n$. Then $G_Y \in \mathcal{G}_{MultiMC}$. Moreover, if X_i $1 \leq i \leq n$ satisfy Assumption 1, then we have $G_Y \in \mathcal{G}_{NMC}$.

Corollary 9 characterizes the graphical model of variables $\{Y_i\}$ that are related to latent jointly Gaussian variables $\{X_i\}$ through the unknown bijective functions $\{f_i\}$. The family of distributions considered in this corollary is broad and includes many Gaussian distributions as well as distributions whose variables are bijective functions of Gaussian variables. Graphical models characterized in Corollary 9 can be used in computation of marginal distributions, computation of the mode of the joint distribution, and in other applications of estimation and prediction similarly to the case of Gaussian graphical models.

Next we provide two examples to highlight similarities and differences between our proposed inference framework (Theorem 5) and the one in reference [12]. Liu et al. [12] considers the graphical model inference problem for functions of jointly Gaussian variables when the underlying link functions are monotone. In [12] to characterize the graphical model of variables Y_i according to Definition 8 (or equivalently the covariance matrix of jointly Gaussian variables X_i), rank statistics between variables and connections between Spearman's and Pearson's correlation coefficients are employed.

Consider four zero mean jointly Gaussian variables $X_1, \dots, X_4$ with the covariance matrix

$$\Lambda_X = \begin{bmatrix} 1.0 & 0.4 & 0.2 & 0.3 \\ 0.4 & 1.0 & 0.3 & 0.2 \\ 0.2 & 0.3 & 1.0 & 0.4 \\ 0.3 & 0.2 & 0.4 & 1.0 \end{bmatrix}. \quad (41)$$

In this case we have $J_X(1, 3) \approx 0$ and $J_X(2, 4) \approx 0$ (the underlying graphical model is illustrated in Figure 2). We observe samples from $Y_i = f_i(X_i)$. In Example 1, we have

$$Y_1 = f_1(X_1) = \begin{cases} 10X_1, & \text{if } X_1 \geq 0, \\ \frac{1}{10}X_1, & \text{otherwise,} \end{cases}, \quad Y_2 = f_2(X_2) = e^{20X_2}, \quad (42)$$

$$Y_3 = f_3(X_3) = -X_3, \quad Y_4 = f_4(X_4) = X_4^3.$$

In this example all link functions are continuous and monotone satisfying the model assumption of our proposed inference framework (Corollary 5) and the copula method [12]. In Example 2, we violate the model assumption for both our inference framework and for the copula method by considering non-monotone link functions:

$$\begin{aligned} Y_1 = f_1(X_1) &= \begin{cases} e^{20X_1}, & \text{if } X_1 \geq 0, \\ -e^{-20X_1}, & \text{otherwise,} \end{cases} \\ Y_2 = f_2(X_2) &= \begin{cases} \frac{1}{\max(X_2)}X_2 - 1, & \text{if } X_2 \geq 0, \\ \frac{-1}{\min(X_2)}X_2 + 1, & \text{otherwise,} \end{cases} \\ Y_3 = f_3(X_3) &= -X_3, \quad Y_4 = f_4(X_4) = X_4^3. \end{aligned} \quad (43)$$

Relationships between samples of these variables are illustrated in Figure 1. We have drawn 10,000 samples of these random variables. Note that all variables are normalized to have zero means and unit variances. The functions $f_i(\cdot)$ remain unknown for the inference part.

In the computation of NMC, if variables are continuous and we only observe samples from their joint distributions, the empirical computation of conditional expectations in Algorithm 2 may be challenging owing to the lack of sufficient samples. One approach to compute empirical conditional expectations at point $x_i \in \mathbb{R}$ is to use all samples in its B_i neighborhood. With m samples $\{(x_1^{(s)}, \dots, x_n^{(s)})\}_{s=1}^m$, for any i, i' , this leads to

$$P_{X_i, X_{i'}}^{(m)}(x_i, x_{i'}) = \frac{1}{m} \sum_{s=1}^m \mathbf{1} \left\{ x_i^s \in \left[x_i - \frac{B_i}{2}, x_i + \frac{B_i}{2} \right], \right. \\ \left. x_{i'}^s \in \left[x_{i'} - \frac{B_{i'}}{2}, x_{i'} + \frac{B_{i'}}{2} \right] \right\}. \quad (44)$$

In our ACE implementations to compute NMC for continuous variables we have considered both fixed and variable window sizes (i.e., B_i 's). In our simulations, in the variable window size case, we consider 10 bins and choose B_i 's so that the number of samples in different bins are the same.

Let $\tilde{J}$ be an estimation of the inverse covariance matrix. Since in the underlying graphical model there are no edges between nodes 1 and 3, and nodes 2 and 4 (Figure 2), we use $(|\tilde{J}(1, 3)| + |\tilde{J}(2, 4)|) / \sum_{1 \leq i, j \leq n} |\tilde{J}(i, j)|$ as a measure to evaluate error in inferring the graphical model structure.

Figure 3 shows inference errors for J_{nmc} (i.e., the NMC inference framework), J_{copula} (i.e., the copula method [12]), J_Y (i.e., observed variables) and J_X (i.e., latent variables) for Example 1 and 2. In the case of Example 1, both NMC

Fig. 2: The underlying graphical model considered in Examples 1 and 2 in Section 6.1.

Fig. 3: Inference errors for J_{nmc} (i.e., the NMC inference framework), J_{copula} (i.e., the copula method [12]), J_Y (i.e., observed variables) and J_X (i.e., latent variables) for setups of Example 1 (panel a) and Example 2 (panel b). Experiments have been repeated over 50 random realizations of variables. The red line in the middle of each box is the sample median. The tops and bottoms of each box are the 25-th and 75-th percentiles of the samples, respectively. The performance of NMC and copula inference frameworks are comparable in the setup of Example 1, while in the one of Example 2 NMC outperforms copula.

and copula inference frameworks have small and comparable errors. In the case of Example 2 where we violate the monotonicity assumption of link functions, NMC appears to outperform copula. Experiments have been repeated over 50 random realizations of variables.

In this section we focused on the application of MC and NMC in learning graphical models for functions of jointly Gaussian variables. A similar MC/NMC framework can potentially be useful in learning graphical models in other setups such as tree graphical models with incomplete samples [60]. We leave further exploration of this application for future research.

6.2 Illustration of Sample NMC's Use in a Data Application

Having illustrated applications of the NMC optimization in learning nonlinear dependencies among variables, here we provide an example to demonstrate NMC's utility in a real data application. We validate inference results with second, distinct data set that was not used in the inference part. This validation step provides evidence that inference results are likely to be meaningful.

Data: Cancer is a complex disease involving abnormal cell growth with the potential to invade or spread to other parts

of the body [61]. Different studies have shown associations of micro RNA patterns in different human cancers [51], [62]. In this section, we use normalized RNA sequence counts from the TCGA data portal for the Glioma cancer (GBMLGG) at the gene level [51]. Other cancer types can be considered in our framework as well. We use the processed data provided in [51]. Let the variable X_i denote the RNA sequence counts of the gene i . We have samples from this variable in m patients denoted by $\{x_i^j\}_{j=1}^m$. In the following we explain inference assumptions and steps we have taken in this application.

Inferring nonlinear gene-gene interactions: for each cancer type, first we select the top 500 highly-variant genes based on their normalized variances (i.e., $n = 500$) [63]. Let $\phi_i^*(X_i)$ for $1 \leq i \leq n$ be a solution for the NMC optimization (3) over a complete graph without self-loops. Let $A_{NMC} \in \mathbb{R}^{n \times n}$ be a symmetric adjacency matrix where $A_{NMC}[i, i'] = \mathbb{E}[\phi_i^*(X_i)\phi_{i'}^*(X_{i'})]$ for $1 \leq i, i' \leq n$. Moreover, let A_{Lin} be a symmetric linear covariance matrix where $A_{Lin}[i, i'] = \mathbb{E}[X_i X_{i'}]$ for $1 \leq i, i' \leq n$. Figure 4-a illustrates top 5% of elements of the $|A_{NMC}| - |A_{Lin}|$ matrix for the Glioma Cancer (GBMLGG). The non-zero elements of this matrix represent gene pairs whose RNA sequence counts are strongly and nonlinearly associated to each other. In Figure 5, we consider some other density parameters as well (e.g., 1% and 10%).

It is important to note that this real-data example is based on several heuristic steps and assumptions about the underlying data that we explain below.

Assumptions: we assume that different RNA sequence count samples for a gene (i.e., $\{x_i^j\}_{j=1}^m$) are independent. That is, there is no between-patient correlations among these samples [64]. Also we assume that conditions of Theorem 7 holds. That is, input data comes from, possibly nonlinear, functions of some latent jointly Gaussian variables satisfying conditions of Theorem 7. These functions are unknown and bijective. This assumption is less restrictive than the one of the standard covariance analysis where input variables are assumed to be jointly Gaussian.

Inferring Nonlinear Gene Modules: we partition each network to k groups using a standard spectral clustering algorithm based on the modularity transformation [65]. We consider different values of k (between 1 and 20) to obtain dense and large clusters. Then, in a heuristic manner, we merge small and heavily overlapping modules to form large and dense ones. We define a *gene module* as a group of genes that are densely connected to each other in the network. A gene module is called nonlinear if it is present in the NMC network but not in the linear one. We use a permutation test [66] to compute a p -value for each gene module in the network by permuting the network structure and comparing the density of the module in the original network with the ones in permuted ones. We only consider gene modules with p -values less than 0.05. An example of such a nonlinear gene module is illustrated by a yellow box in Figure 4-a.

Validations using Survival Time Analysis: we partition individuals to two equal-size groups based on their average ranks of normalized RNA sequence counts in that module to evaluate if there are statistically significant differences between survival times of patients in the two groups. In order to do that, we perform a standard survival time analysis for each module based on Kaplan-Meier procedure to estimate the underlying survival function [67]. We compute its associated log-rank p -value to determine its association with individual survival

Fig. 4: (a) A nonlinear gene module of Glioma cancer (GBMLGG) defined in Section 6.2, as a group of genes whose RNA sequence counts are strongly and nonlinearly dependent among each other. (b) Survival time curves for the corresponding nonlinear cancer module. Survival times of cancer patients are significantly associated with average normalized RNA sequence counts of inferred gene module.

times in the considered cancer type [68]. We perform Benjamini and Hochberg multiple hypothesis correction [69] for the computed p -values of different nonlinear modules. We find that this gene module is significantly associated with survival times of cancer patients (Figure 4-b) while it is not detected using linear association measures. Several references [70], [71], [72], [73], [74] have hypothesized that complex nonlinear relationships among genes may play important roles in cancer pathways. The proposed NMC algorithm and inferred nonlinear gene modules can be used in discovering such complex nonlinear relationships in different cancer types. To substantiate these inferences, further experiments should be performed to determine the involvement of these nonlinear gene interactions and modules in different cancers, which is beyond the scope of the present paper.

7 PROOFS

7.1 Proof of Theorem 1

Proof Recall that $\{\psi_{i,j}\}_{j=1}^\infty$ is the orthonormal basis of H_{X_i} for $1 \leq i \leq n$. We can represent functions ϕ_i and $\phi_{i'}$ in terms of the basis functions, i.e., for any $x_i \in \mathcal{X}_i, x_{i'} \in \mathcal{X}_{i'}$,

$$\begin{aligned}\phi_i(x_i) &= \sum_{j=1}^\infty a_{i,j} \psi_{i,j}(x_i), \\ \phi_{i'}(x_{i'}) &= \sum_{j=1}^\infty a_{i',j} \psi_{i',j}(x_{i'}),\end{aligned}$$

for two sequences of coefficients $\{a_{i,j}\}_{j=1}^\infty$ and $\{a_{i',j}\}_{j=1}^\infty$. Thus, the constraint $\mathbb{E}[\phi_k(X_k)^2] = 1$ in optimization (3) would be translated into $\sum_{j=1}^\infty a_{k,j}^2 = 1$ and the constraint $\mathbb{E}[\phi_k(X_k)] = 0$ is simplified to $\sum_{j=1}^\infty a_{k,j} \mathbb{E}[\psi_{k,j}(X_k)] = 0$, for $k = i, i'$. Moreover, we have

$$\mathbb{E}[\phi_i(X_i)\phi_{i'}(X_{i'})] = \sum_{j,j'=1}^\infty a_{i,j}a_{i',j'} \mathbb{E}[\psi_{i,j}(X_i)\psi_{i',j'}(X_{i'})]. \quad (45)$$

Fig. 5: An example illustrating the framework considered in Section 6.2 for Glioma (GBMLGG). Using normalized RNA sequence counts from the TCGA data portal, we construct both linear (A_{Lin}) and nonlinear (A_{NMC}) complete dependency graphs between genes. We then select the top 5% interactions among genes in the NMC network with the largest nonlinear association increases compared to their linear association strengths. Then, using spectral clustering [65] along with a heuristic step to merge small overlapping clusters, we identify dense and large clusters over the thresholded network. In this example, we illustrate the inferred nonlinear gene module over networks with density parameters 1%, 5% and 10%.

Thus, Network Maximal Correlation optimization (3) can be re-written as follows:

$$\begin{aligned} \sup_{\{\{a_{i,j}\}_{i=1}^n\}_{j=1}^\infty} \quad & \sum_{i,i'} \sum_{j,j'} a_{i,j} a_{i',j'} \mathbb{E}[\psi_{i,j}(X_i) \psi_{i',j'}(X_{i'})] \quad (46) \\ & \sum_{j=1}^\infty a_{i,j}^2 = 1, \quad 1 \leq i \leq n, \\ & \sum_{j=1}^\infty a_{i,j} \mathbb{E}[\psi_{i,j}(X_i)] = 0, \quad 1 \leq i \leq n. \end{aligned}$$

Moreover, $\{\psi_{i,j}, \psi_{i',j'}\}_{j,j'}$ form a basis for the following set of functions $\{\phi_{i,i'} : \mathcal{X}_i \times \mathcal{X}_{i'} \rightarrow \mathbb{R} : \mathbb{E}[\phi_{i,i'}^2(X_i, X_{i'})] < \infty\}$. Since $P_{X_i, X_{i'}}(\cdot, \cdot) \in$ belongs to this set of functions, we can write

$$P_{X_i, X_{i'}}(x_i, x_{i'}) = \sum_{j,j'} \rho_{i,i'}^{j,j'} \psi_{i,j}(x_i) \psi_{i',j'}(x_{i'}). \quad (47)$$

This completes the proof. $\blacksquare$

7.2 Proof of Theorem 2

Proof To prove Theorem 2, first we show that the constraints $\mathbb{E}[\phi_i(X_i)] = 0$ and $\mathbb{E}[\phi_i(X_i)^2] = 1$ lead to the same solution as the constraints $\text{var}(\phi_i(X_i)) = 1$. We then show that the constraint $\mathbf{a}_i \perp \sqrt{\mathbf{p}_i}$ can be incorporated into the objective function, without changing the solution.

Lemma 1. The NMC optimization (3) can be written as follows:

$$\begin{aligned} \max_{\phi_1, \dots, \phi_n} \quad & \sum_{(i,i') \in E} \mathbb{E}[(\phi_i(X_i) - \bar{\phi}_i)(\phi_{i'}(X_{i'}) - \bar{\phi}_{i'})] \quad (48) \\ & \text{var}(\phi_i(X_i)) = 1, \quad 1 \leq i \leq n, \end{aligned}$$

where $\bar{\phi}_i$ and $\text{var}(\phi_i(X_i))$ represent the mean and the variance of the random variable $\phi_i(X_i)$.

Proof In this proof, we denote the random variable $\phi_i(X_i)$ by ϕ_i . We let the optimal objective value of optimization (48) be $\tilde{\rho}_G$. We also let ϕ_i^* be an solution of (3). The set of functions ϕ_i^* for $i = 1, \dots, n$ is feasible for optimization (48) and therefore we have $\rho_G \leq \tilde{\rho}_G$. On the other hand, let ϕ_i^{**} be an solution of optimization (48). Let $\tilde{\phi}_i = \phi_i^{**} - \bar{\phi}_i^{**}$. The set of functions $\tilde{\phi}_i$ for $i = 1, \dots, n$ is feasible for optimization (3). Thus, we have $\rho_G \geq \tilde{\rho}_G$. Therefore, we have that $\rho_G = \tilde{\rho}_G$. $\blacksquare$

We have

$$\begin{aligned} 1 &= \mathbb{E}[\phi_i(X_i)^2] - (\mathbb{E}[\phi_i(X_i)])^2 = \|\mathbf{a}_i\|_2^2 - (\mathbf{a}_i \sqrt{\mathbf{p}_i})^2 \quad (49) \\ &= \mathbf{a}_i^T (I - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_i}^T) \mathbf{a}_i. \end{aligned}$$

We next show that the matrix $I - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_i}^T$ is positive semidefinite and the only vectors in its null space are $\mathbf{0}$ and $\sqrt{\mathbf{p}_i}$. This is because

$$\mathbf{x}^T (I - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_i}^T) \mathbf{x} = \|\mathbf{x}\|_2^2 - (\mathbf{x} \sqrt{\mathbf{p}_i})^2 \geq 0, \quad (50)$$

where we use Cauchy-Schwartz and $\|\sqrt{\mathbf{p}_i}\|_2^2 = 1$ to obtain the last inequality (50). This inequality becomes an equality if and only if $\mathbf{x} = \mathbf{0}$ or $\mathbf{x} = \sqrt{\mathbf{p}_i}$.

Now consider the objective function of optimization (48). We have

$$\begin{aligned} \mathbb{E}[(\phi_i(X_i) - \bar{\phi}_i)(\phi_{i'}(X_{i'}) - \bar{\phi}_{i'})] &= \mathbb{E}[\phi_i(X_i) \phi_{i'}(X_{i'})] - \bar{\phi}_i \bar{\phi}_{i'} \\ &= \mathbf{a}_i^T Q_{i,i'} \mathbf{a}_{i'} - (\mathbf{a}_i^T \sqrt{\mathbf{p}_i})(\mathbf{a}_{i'}^T \sqrt{\mathbf{p}_{i'}}) = \mathbf{a}_i^T (Q_{i,i'} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_{i'}}^T) \mathbf{a}_{i'}. \end{aligned}$$

Therefore, optimization (48) (which is equivalent to the NMC optimization (3) according to Lemma 1) can be written as,

$$\begin{aligned} \max_{\mathbf{a}_i} \quad & \sum_{(i,i') \in E} \mathbf{a}_i^T (Q_{i,i'} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_{i'}}^T) \mathbf{a}_{i'} \quad (51) \\ & \mathbf{a}_i^T (I - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_i}^T) \mathbf{a}_i = 1, \quad 1 \leq i \leq n. \end{aligned}$$

For each i , since $I - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_i}^T$ is positive semidefinite. Thus, we can write $I - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_i}^T = B_i B_i^T$. Recall $\mathbf{b}_i = B_i \mathbf{a}_i$. Thus, constraints of optimization (51) can be written as $\mathbf{b}_i^T \mathbf{b}_i = \|\mathbf{b}_i\|_2^2 = 1$. We next write $\mathbf{a}_i$ as a function of $\mathbf{b}_i$. Note that since B_i is not invertible, there are many choices for $\mathbf{a}_i$ as a function of $\mathbf{b}_i$ characterized as follows: Let $U_i \Sigma_i U_i^T$ be the singular value decomposition of B_i . The vector $\sqrt{\mathbf{p}_i}$ is the singular vector corresponding to singular value zero ($\sigma_i^{(1)} = 0$).

$$\begin{aligned} \mathbf{a}_i &= ([U_i^{(2)}, \dots, U_i^{(|\mathcal{X}_i|)}] \text{diag}(1/\sigma_i^{(2)}, \dots, 1/\sigma_i^{(n_i)}) \\ & \quad [U_i^{(2)}, \dots, U_i^{(|\mathcal{X}_i|)}]^T) \mathbf{b}_i + \alpha_i \sqrt{\mathbf{p}_i} = A_i \mathbf{b}_i + \alpha_i \sqrt{\mathbf{p}_i}, \end{aligned} \quad (52)$$

where α_i can be any scalar.³ Below, we show that all choices of $\mathbf{a}_i$ according to (52) lead to the same objective function of

³Since B_i is symmetric it has a set of $|\mathcal{X}_i|$ orthonormal eigenvectors and can be written as

$$B_i = \sum_{j=1}^{|\mathcal{X}_i|} \mathbf{v}_j \sigma_i^{(j)} \mathbf{v}_j^T.$$

We have $\mathbf{b}_i = \sum_{j=2}^{|\mathcal{X}_i|} \beta_j \mathbf{v}_j$ and $\mathbf{a}_i = \sum_{j=1}^{|\mathcal{X}_i|} \alpha_j \mathbf{v}_j$. From $\mathbf{b}_i = B_i \mathbf{a}_i$, we obtain that $\alpha_j = \beta_j / \sigma_i^{(j)}$ for $j \geq 2$, where α_1 can be any scalar.

optimization (51):

$$\begin{aligned}
& \mathbf{a}_i^T (Q_{i,i'} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_{i'}^T}) \mathbf{a}_{i'} = \mathbf{b}_i^T A_i^T (Q_{i,i'} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_{i'}^T}) A_{i'} \mathbf{b}_{i'} \\
& + \mathbf{b}_i^T A_i^T (Q_{i,i'} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_{i'}^T}) \alpha_{i'} \sqrt{\mathbf{p}_{i'}} \\
& + \alpha_i \sqrt{\mathbf{p}_i}^T (Q_{i,i'} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_{i'}^T}) A_{i'} \mathbf{b}_{i'} \\
& + \alpha_i \sqrt{\mathbf{p}_i}^T (Q_{i,i'} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_{i'}^T}) \alpha_{i'} \sqrt{\mathbf{p}_{i'}} \\
& \stackrel{(1)}{=} \mathbf{b}_i^T A_i^T (Q_{i,i'} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_{i'}^T}) A_{i'} \mathbf{b}_{i'} \\
& + \mathbf{b}_i^T A_i^T Q_{i,i'} \alpha_{i'} \sqrt{\mathbf{p}_{i'}} + \alpha_i \sqrt{\mathbf{p}_i}^T Q_{i,i'} A_{i'} \mathbf{b}_{i'} \\
& \stackrel{(2)}{=} \mathbf{b}_i^T A_i^T (Q_{i,i'} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_{i'}^T}) A_{i'} \mathbf{b}_{i'} + \mathbf{b}_i^T A_i^T \sqrt{\mathbf{p}_i} \alpha_{i'} \\
& + \alpha_i \sqrt{\mathbf{p}_i}^T A_{i'} \mathbf{b}_{i'} = \mathbf{b}_i^T A_i^T (Q_{i,i'} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_{i'}^T}) A_{i'} \mathbf{b}_{i'},
\end{aligned}$$

where (1) follows from $\sqrt{\mathbf{p}_{i'}}^T \sqrt{\mathbf{p}_{i'}} = 1$, $\sqrt{\mathbf{p}_i}^T \sqrt{\mathbf{p}_i} = 1$, $\sqrt{\mathbf{p}_i} Q_{i,i'} \sqrt{\mathbf{p}_{i'}} = 1$, $A_i^T \sqrt{\mathbf{p}_i} = 0$, and $\sqrt{\mathbf{p}_{i'}}^T A_{i'} = 0$; and (2) follows from $\sqrt{\mathbf{p}_i}^T Q_{i,i'} = \sqrt{\mathbf{p}_{i'}}^T$ and $Q_{i,i'} \sqrt{\mathbf{p}_{i'}} = \sqrt{\mathbf{p}_i}$. Therefore, the NMC optimization (3) can be written as

$$\max_{\mathbf{b}_1, \dots, \mathbf{b}_n} \sum_{(i,i') \in E} \mathbf{b}_i^T A_i^T (Q_{i,i'} - \sqrt{\mathbf{p}_i} \sqrt{\mathbf{p}_{i'}^T}) A_{i'} \mathbf{b}_{i'}$$

$\|\mathbf{b}_i\|_2 = 1 \quad 1 \leq i \leq n.$

7.3 Proof of Proposition 1

Proof Let $\bar{\mathbf{b}}$ be a solution of the MEP (19). Suppose that there exists an $1 \leq i \leq n$ such that $\lambda_i(\bar{\mathbf{b}}) < 0$. Let $\hat{\mathbf{b}} = (\hat{\mathbf{b}}_1, \dots, \hat{\mathbf{b}}_n)$ be defined as: $\hat{\mathbf{b}}_i = -\bar{\mathbf{b}}_i$, for any i such that $\lambda_i(\bar{\mathbf{b}}) < 0$, and $\hat{\mathbf{b}}_{i'} = \bar{\mathbf{b}}_{i'}$, otherwise. We show that by flipping the sign of $\bar{\mathbf{b}}_i$ while keeping the rest of $\bar{\mathbf{b}}_{i'}$ for $i' \neq i$ the same, $r(\cdot)$ increases. To show this we have

$$\begin{aligned}
r(\hat{\mathbf{b}}) - r(\bar{\mathbf{b}}) & \stackrel{(1)}{=} 2\hat{\mathbf{b}}_i^T \sum_{i'=1}^n C_{i,i'} \hat{\mathbf{b}}_{i'} - 2\bar{\mathbf{b}}_i^T \sum_{i'=1}^n C_{i,i'} \bar{\mathbf{b}}_{i'} \\
& \stackrel{(2)}{=} 2(\hat{\mathbf{b}}_i^T - \bar{\mathbf{b}}_i^T) \sum_{i'=1}^n C_{i,i'} \bar{\mathbf{b}}_{i'} \\
& \stackrel{(3)}{=} -2\bar{\mathbf{b}}_i^T \sum_{i'=1}^n C_{i,i'} \bar{\mathbf{b}}_{i'} \\
& \stackrel{(4)}{=} -2\lambda_i(\bar{\mathbf{b}}) > 0,
\end{aligned}$$

where equations (1) and (4) follow from (20), equation (2) follows from the fact that $\hat{\mathbf{b}}_{i'} = \bar{\mathbf{b}}_{i'}$ for $i' \neq i$, and equation (3) follows from the fact that $\hat{\mathbf{b}}_i = -\bar{\mathbf{b}}_i$. This completes the proof. ■

7.4 Proof of Theorem 4

Proof For any realization of the partitioning, consider NMC over all sub-graphs G_m ($1 \leq m \leq M$) and denote the corresponding functions by ϕ_i^* for $1 \leq i \leq n$. We have

$$\begin{aligned}
\rho_G &= \sum_{(i,i') \in E} \mathbb{E}[\phi_i^*(X_i) \phi_{i'}^*(X_{i'})] \\
&= \sum_{(i,i') \in E \setminus E^c} \mathbb{E}[\phi_i^*(X_i) \phi_{i'}^*(X_{i'})] + \sum_{(i,i') \in E^c} \mathbb{E}[\phi_i^*(X_i) \phi_{i'}^*(X_{i'})] \\
& \stackrel{(1)}{=} \sum_{m=1}^M \sum_{(i,i') \in E_m} \mathbb{E}[\phi_i^*(X_i) \phi_{i'}^*(X_{i'})] + \sum_{(i,i') \in E^c} \mathbb{E}[\phi_i^*(X_i) \phi_{i'}^*(X_{i'})] \\
& \stackrel{(2)}{\leq} \sum_{m=1}^M \hat{\rho}_{G_m} + \sum_{(i,i') \in E^c} \mathbb{E}[\phi_i^*(X_i) \phi_{i'}^*(X_{i'})] \\
&= \hat{\rho}_G + \sum_{(i,i') \in E} \mathbf{1}\{(i,i') \in E^c\} \mathbb{E}[\phi_i^*(X_i) \phi_{i'}^*(X_{i'})],
\end{aligned}$$

where equation (1) comes from the graph partitioning Definition 6, and inequality (2) comes from the fact that $\hat{\rho}_{G_m}$ is the NMC for the partition $G_m = (V_m, E \cap (V_m \times V_m))$. Therefore, by taking expectation over the partitioning, we obtain

$$\rho_G \leq \mathbb{E}[\hat{\rho}_G] + \epsilon \rho_G,$$

which gives us

$$(1 - \epsilon) \rho_G \leq \mathbb{E}[\hat{\rho}_G].$$

7.5 Proof of Theorem 5

Proof We first present a lemma in which we bound the difference between Q matrices, according to Definition 3, by the difference between probability distributions. We then use a sensitivity analysis of the optimization whose optimum is NMC.

Lemma 2. Let P and $\tilde{P}$ be the matrix form of two joint probability distribution on $(\mathcal{X}_i, \mathcal{X}_{i'})$, such that $P_{X_i, X_{i'}}(x_i, x_{i'}) = [P]_{x_i, x_{i'}}$ and $\tilde{P}_{X_i, X_{i'}}(x_i, x_{i'}) = [\tilde{P}]_{x_i, x_{i'}}$ and where $D_{X_i}(P) = \text{diag}(P_{X_i}(x_i) : x_i \in \mathcal{X}_i)$. We can bound the difference between $Q_{i,i'}$ and $\tilde{Q}_{i,i'}$ by the difference between P and $\tilde{P}$, as follows:

$$\begin{aligned}
\|Q_{i,i'} - \tilde{Q}_{i,i'}\|_2 & \leq \frac{1}{2\delta^2} \sqrt{K} \|D_{X_i}(P) - D_{X_i}(\tilde{P})\|_\infty \\
& + \frac{1}{2\delta^2} \sqrt{K} \|D_{X_{i'}}(P) - D_{X_{i'}}(\tilde{P})\|_\infty + \frac{1}{\delta} \sqrt{K} \|P - \tilde{P}\|_\infty, \quad (53)
\end{aligned}$$

where δ is the minimum probability of all elements of $\mathcal{X}_i$ and $\mathcal{X}_{i'}$, under both P and $\tilde{P}$, and K is the maximum alphabet size. In particular, if $\|P_{X_i, X_{i'}} - \tilde{P}_{X_i, X_{i'}}\|_\infty \leq \gamma$ and $K = \max\{|\mathcal{X}_i|, |\mathcal{X}_{i'}|\}$, then we have

$$\|Q_{i,i'} - \tilde{Q}_{i,i'}\|_2 \leq 2\gamma K^n \frac{1}{\delta^2}. \quad (54)$$

Proof We have

$$\begin{aligned}
\|Q_{i,i'} - \tilde{Q}_{i,i'}\|_2 &= \|D_{X_i}(P)^{-\frac{1}{2}} P D_{X_{i'}}(P)^{-\frac{1}{2}} - D_{X_i}(\tilde{P})^{-\frac{1}{2}} \tilde{P} D_{X_{i'}}(\tilde{P})^{-\frac{1}{2}}\|_2.
\end{aligned}$$

We next insert the terms $\pm D_{X_i}(P)^{-\frac{1}{2}} P D_{X_{i'}}(\tilde{P})^{-\frac{1}{2}}$ and $\pm D_{X_i}(P)^{-\frac{1}{2}} \tilde{P} D_{X_{i'}}(\tilde{P})^{-\frac{1}{2}}$, pair the terms, and use triangle inequality to obtain

$$\begin{aligned}
\|Q_{i,i'} - \tilde{Q}_{i,i'}\|_2 & \leq \|D_{X_i}(P)^{-\frac{1}{2}} P\|_2 \|D_{X_{i'}}(P)^{-\frac{1}{2}} - D_{X_{i'}}(\tilde{P})^{-\frac{1}{2}}\|_2 \\
& + \|D_{X_{i'}}(\tilde{P})^{-\frac{1}{2}}\|_2 \|\tilde{P}\|_2 \|D_{X_i}(P)^{-\frac{1}{2}} - D_{X_i}(\tilde{P})^{-\frac{1}{2}}\|_2 \\
& + \|D_{X_{i'}}(\tilde{P})^{-\frac{1}{2}}\|_2 \|D_{X_i}(P)^{-\frac{1}{2}}\|_2 \|P - \tilde{P}\|_2,
\end{aligned} \quad (55)$$

where each term depends on the difference between P and $\tilde{P}$. We have

$$\|D_{X_i}(P)^{-\frac{1}{2}} P\|_2 \stackrel{(1)}{\leq} \frac{1}{\sqrt{\delta}} \sqrt{K} \quad (56)$$

$$\begin{aligned}
\|D_{X_{i'}}(P)^{-\frac{1}{2}} - D_{X_{i'}}(\tilde{P})^{-\frac{1}{2}}\|_2 & \leq \|D_{X_{i'}}(P)^{-\frac{1}{2}} - D_{X_{i'}}(\tilde{P})^{-\frac{1}{2}}\|_\infty \\
& \stackrel{(2)}{\leq} \|D_{X_{i'}}(P) - D_{X_{i'}}(\tilde{P})\|_\infty \frac{1}{2\delta^{3/2}},
\end{aligned}$$

where inequality (1) comes from the Cauchy-Schwartz inequality and the fact that $\|P\|_2 \leq \sqrt{K} \|P\|_1 \leq \sqrt{K}$. Inequality (2) follows from the fact that for $x, y \in \mathbb{R}$, we have

$$\frac{1}{\sqrt{x}} - \frac{1}{\sqrt{y}} = \frac{y - x}{\sqrt{x} \sqrt{y} (\sqrt{x} + \sqrt{y})} \leq |x - y| \frac{1}{2(\min\{x, y\})^{3/2}}. \quad (57)$$

Using (56) in (55) results in (53).

Moreover we have

$$\begin{aligned} & \|D_{X_i}(P) - D_{X_i}(\tilde{P})\|_\infty \\ & \leq \sum_{\{X_1, \dots, X_n\} \setminus \{X_i\}} \|P_{X_1, \dots, X_n} - \tilde{P}_{X_1, \dots, X_n}\|_\infty \leq \gamma K^{n-1}. \end{aligned} \quad (58)$$

Using (58) and the fact that $K > 1$ in (53) results in (54). This completes the proof of this Lemma. ■

Let P and $\tilde{P}$ be two distributions on $\mathcal{X}_1 \times \dots \times \mathcal{X}_n$. We shall compare the solution of the two following optimization problems.

$$\begin{aligned} & \max_{\mathbf{a}_i} \sum_{(i, i') \in E} \mathbf{a}_i^T Q_{i, i'} \mathbf{a}_i \\ & \|\mathbf{a}_i\|_2 = 1, \quad 1 \leq i \leq n, \\ & \mathbf{a}_i \perp \sqrt{\mathbf{p}_i}, \quad 1 \leq i \leq n, \end{aligned} \quad (59)$$

and

$$\begin{aligned} & \max_{\mathbf{a}_i} \sum_{(i, i') \in E} \mathbf{a}_i^T \tilde{Q}_{i, i'} \mathbf{a}_i \\ & \|\mathbf{a}_i\|_2 = 1, \quad 1 \leq i \leq n, \\ & \mathbf{a}_i \perp \sqrt{\tilde{\mathbf{p}}_i}, \quad 1 \leq i \leq n. \end{aligned} \quad (60)$$

Let ρ_G and $\tilde{\rho}_G$ be the optimal values for (59) and (60), respectively.

Recall that

$$\delta = \min_{1 \leq i \leq n} \left(\min\{\delta_{X_i}(P), \delta_{X_i}(\tilde{P})\} \right).$$

Suppose $\mathbf{a}_i^*$ yields the optimum of optimization (59). Based on this solution, we shall construct a feasible solution for optimization (60) and then evaluate its objective function. For any i , let

$$\mathbf{c}_i = \frac{\mathbf{a}_i^* + \nu_i}{\|\mathbf{a}_i^* + \nu_i\|_2},$$

where $\nu_i = \sqrt{\tilde{\mathbf{p}}_i} < \mathbf{a}_i^*, \sqrt{\mathbf{p}_i} - \sqrt{\tilde{\mathbf{p}}_i} >$. Under the assumption of Theorem 5 (i.e., $\gamma \leq \delta^{3/2} K^{-1}$), we show that $\|\nu_i\| \leq 1/2$. We have

$$\|\nu_i\| = \|\sqrt{\tilde{\mathbf{p}}_i} < \mathbf{a}_i^*, \sqrt{\mathbf{p}_i} - \sqrt{\tilde{\mathbf{p}}_i} >\|_2 \stackrel{(3)}{\leq} \|\sqrt{\tilde{\mathbf{p}}_i}\|_2 \|\mathbf{a}_i^*\|_2 \|\sqrt{\mathbf{p}_i} - \sqrt{\tilde{\mathbf{p}}_i}\|_2 \quad (61)$$

$$\begin{aligned} & \stackrel{(4)}{\leq} \|\mathbf{p}_i - \tilde{\mathbf{p}}_i\|_\infty \frac{1}{2\delta^{3/2}} \\ & \stackrel{(5)}{\leq} \frac{\gamma K^{n-1}}{2\delta^{3/2}} \\ & \stackrel{(6)}{\leq} \frac{1}{2}, \end{aligned} \quad (62)$$

where Inequality (3) comes from the Cauchy-Schwartz, Inequality (4) comes from the facts that $\|\sqrt{\tilde{\mathbf{p}}_i}\|_2 = 1$, $\|\mathbf{a}_i^*\|_2 = 1$, (57), Inequality (5) comes from the fact that $\|\mathbf{p}_i(j) - \tilde{\mathbf{p}}_i(j)\|_\infty \leq \gamma K^{n-1}$, and Inequality (6) comes from the fact that $\frac{\gamma K^{n-1}}{\delta^{3/2}} \leq 1$ by the assumption of Theorem 5.

We claim $\mathbf{c}_i$ is feasible for optimization (60). First note that the norm of each $\mathbf{c}_i$ is one. We next show that each $\mathbf{c}_i$ is orthogonal to $\sqrt{\tilde{\mathbf{p}}_i}$:

$$\begin{aligned} \langle \mathbf{c}_i, \sqrt{\tilde{\mathbf{p}}_i} \rangle &= \frac{1}{\|\mathbf{a}_i^* + \sqrt{\tilde{\mathbf{p}}_i} < \mathbf{a}_i^*, \sqrt{\mathbf{p}_i} - \sqrt{\tilde{\mathbf{p}}_i} >\|_2} \\ & \left(\langle \mathbf{a}_i^*, \sqrt{\tilde{\mathbf{p}}_i} \rangle + \langle \mathbf{a}_i^*, \sqrt{\mathbf{p}_i} - \sqrt{\tilde{\mathbf{p}}_i} \rangle \|\sqrt{\tilde{\mathbf{p}}_i}\|_2^2 \right) = 0, \end{aligned}$$

where the last equality follows from $\|\sqrt{\tilde{\mathbf{p}}_i}\|_2 = 1$ and $\mathbf{a}_i \perp \sqrt{\mathbf{p}_i}$. We now plug in the feasible solution $\mathbf{c}_i$ into the objective function of optimization (60).

$$\begin{aligned} \tilde{\rho}_G & \stackrel{(7)}{\geq} \sum_{(i, i') \in E} \mathbf{c}_i^T \tilde{Q}_{i, i'} \mathbf{c}_{i'} \\ & \stackrel{(8)}{=} \sum_{(i, i') \in E} \mathbf{a}_i^{*T} Q_{i, i'} \mathbf{a}_{i'}^* + \mathbf{c}_i^T (\tilde{Q}_{i, i'} - Q_{i, i'}) \mathbf{c}_{i'} + (\mathbf{c}_i^T - \mathbf{a}_i^{*T}) Q_{i, i'} \mathbf{c}_{i'} \\ & \quad + \mathbf{a}_i^{*T} Q_{i, i'} (\mathbf{c}_{i'} - \mathbf{a}_{i'}^*) \\ & \stackrel{(9)}{\geq} \rho_G - \sum_{(i, i') \in E} \left(\underbrace{|\mathbf{c}_i^T (\tilde{Q}_{i, i'} - Q_{i, i'}) \mathbf{c}_{i'}|}_{\text{Term I}} + \underbrace{|(\mathbf{c}_i^T - \mathbf{a}_i^{*T}) Q_{i, i'} \mathbf{c}_{i'}|}_{\text{Term II}} \right. \\ & \quad \left. + \underbrace{|\mathbf{a}_i^{*T} Q_{i, i'} (\mathbf{c}_{i'} - \mathbf{a}_{i'}^*)|}_{\text{Term III}} \right), \end{aligned} \quad (63)$$

where Inequality (7) comes from the fact that $\mathbf{c}_i$ for $1 \leq i \leq n$ is a feasible solution of optimization (60), Equality (8) comes from adding and subtracting $\mathbf{a}_i^{*T} Q_{i, i'} \mathbf{a}_{i'}^*$, $\mathbf{c}_i^T Q_{i, i'} \mathbf{c}_{i'}$ and $\mathbf{a}_i^{*T} Q_{i, i'} \mathbf{c}_{i'}$, and Inequality (9) comes from the triangle inequality and the fact that $\mathbf{a}_i^*$, $1 \leq i \leq n$ is a solution of optimization (59). Next, we bound terms I-III on the righthand side of this relation.

Using Lemma 2, for any i, i' , we have

$$|\mathbf{c}_i^T (\tilde{Q}_{i, i'} - Q_{i, i'}) \mathbf{c}_{i'}| \leq \|\tilde{Q}_{i, i'} - Q_{i, i'}\|_2 \leq 2 \frac{\gamma K^n}{\delta^2}. \quad (64)$$

We also have

$$\begin{aligned} |(\mathbf{c}_i^T - \mathbf{a}_i^{*T}) Q_{i, i'} \mathbf{c}_{i'}| & \leq \|\mathbf{c}_i^T - \mathbf{a}_i^{*T}\|_2 \|Q_{i, i'}\|_2 \|\mathbf{c}_{i'}\|_2 = \|\mathbf{c}_i^T - \mathbf{a}_i^{*T}\|_2 \\ & \leq 2 \frac{\|\nu_i\|_2}{1 - \|\nu_i\|_2}, \end{aligned}$$

where we use the following inequality

$$\begin{aligned} \left\| \frac{\mathbf{a} + \nu_i}{\|\mathbf{a} + \nu_i\|_2} - \mathbf{a} \right\|_2 &= \left\| \frac{\nu_i}{\|\mathbf{a} + \nu_i\|_2} + \mathbf{a} \left(\frac{1}{\|\mathbf{a} + \nu_i\|_2} - 1 \right) \right\|_2 \\ &\leq \left\| \frac{\nu_i}{\|\mathbf{a} + \nu_i\|_2} \right\|_2 + \|\mathbf{a}\|_2 \left| \frac{1}{\|\mathbf{a} + \nu_i\|_2} - 1 \right| \\ &\leq \frac{\|\nu_i\|_2}{1 - \|\nu_i\|_2} + \max\left\{ \left| \frac{1}{1 - \|\nu_i\|_2} - 1 \right|, \left| \frac{1}{1 + \|\nu_i\|_2} - 1 \right| \right\} \\ &\leq 2 \frac{\|\nu_i\|_2}{1 - \|\nu_i\|_2}, \end{aligned} \quad (65)$$

where we used $1 - \|\nu_i\|_2 \leq \|\mathbf{a} + \nu_i\|_2 \leq 1 + \|\nu_i\|_2$ and the fact that $\|\nu_i\| \leq 1/2$ according to (61). Using (61) and (65), we obtain

$$|(\mathbf{c}_i^T - \mathbf{a}_i^{*T}) Q_{i, i'} \mathbf{c}_{i'}| \leq 2 \frac{\gamma K^{n-1}}{\delta^{3/2}}, \quad (66)$$

Using (64) and (66) in (63) leads to

$$\tilde{\rho}_G \geq \rho_G - \sum_{(i, i') \in E} \left(4 \frac{\gamma K^n}{\delta^2} + 4 \gamma K^{n-1} \frac{1}{\delta^{3/2}} \right) \geq \rho_G - \gamma K^n |E| \frac{8}{\delta^2}.$$

Similarly, we have

$$\rho_G \geq \tilde{\rho}_G - \sum_{(i, i') \in E} \left(4 \frac{\gamma K^n}{\delta^2} + 4 \gamma K^{n-1} \frac{1}{\delta^{3/2}} \right) \geq \tilde{\rho}_G - \gamma K^n |E| \frac{8}{\delta^2}.$$

Combining the previous two relations, we obtain

$$|\tilde{\rho}_G - \rho_G| \leq \gamma K^n |E| \frac{8}{\delta^2},$$

which completes the proof. ■

7.6 Proof of Theorem 6

Proof We use the following theorem in the proof.

Theorem 10 ([75], [76]). Let P be a probability distribution on a finite alphabet $\mathcal{X}$. Also, let $P^{(m)}$ denote the empirical probability distribution of X , obtained from m i.i.d. samples, $\{x_i\}_{i=1}^m$, drawn according to P . We have

$$\mathbb{P}[\|P^{(m)} - P\|_\infty > \gamma] \leq 4e^{-m\frac{\gamma^2}{2}}. \quad (67)$$

Theorem 10 establishes that P and $P^{(m)}$ are close to each other for a sufficiently large sample size m^4 . Moreover, one can use Theorem 5 to bound the difference between NMC values of distributions P and $P^{(m)}$ by $\|P^{(m)} - P\|_\infty$. However, to prove Theorem 6, we also need to bound the minimum sample probability of $P^{(m)}$. We explain these steps below.

We have

$$\begin{aligned} \mathbb{P}\left[\delta_{X_i}(P^{(m)}) \geq \frac{\delta'}{2}\right] &\geq \mathbb{P}\left[\|P_{X_i}^{(m)} - P_{X_i}\|_\infty \leq \frac{\delta'}{2}\right] \\ &\geq \mathbb{P}\left[\|P^{(m)} - P\|_\infty \leq \frac{\delta'}{2K^{n-1}}\right] \\ &\geq 1 - \mathbb{P}\left[\|P^{(m)} - P\|_\infty > \frac{\delta'}{2K^{n-1}}\right] \geq 1 - 4e^{-m\frac{1}{2}\left(\frac{\delta'}{2K^{n-1}}\right)^2}, \end{aligned}$$

where the first inequality follows from $\delta_{X_i}(P^{(m)}) \geq -|\delta_{X_i}(P^{(m)}) - \delta_{X_i}(P)| + \delta_{X_i}(P) \geq \delta - \|P_{X_i}^{(m)} - P_{X_i}\|_\infty$ and the second inequality follows from $\|P_{X_i}^{(m)} - P_{X_i}\|_\infty \leq K^{n-1}\|P^{(m)} - P\|_\infty$. Using the union bound, we obtain

$$\begin{aligned} \mathbb{P}\left[\bigcap_{i \in V} \left\{\delta_{X_i}(P^{(m)}) \geq \frac{\delta'}{2}\right\}\right] &= 1 - \mathbb{P}\left[\bigcup_{i \in V} \left\{\delta_{X_i}(P^{(m)}) \leq \frac{\delta'}{2}\right\}\right] \\ &\geq 1 - 4ne^{-m\frac{1}{2}\left(\frac{\delta'}{2K^{n-1}}\right)^2}. \end{aligned}$$

By applying the union bound once more, we have

$$\begin{aligned} \mathbb{P}\left[\left\{\|P^{(m)} - P\|_\infty < \gamma\right\} \cap \left\{\bigcap_{i \in V} \left\{\delta_{X_i}(P^{(m)}) \geq \frac{\delta'}{2}\right\}\right\}\right] \\ \geq 1 - 4ne^{-m\frac{1}{2}\left(\frac{\delta'}{2K^{n-1}}\right)^2} - 4e^{-m\frac{\gamma^2}{2}}. \end{aligned}$$

Thus, with a large probability, i.e., $1 - 4ne^{-m\frac{1}{2}\left(\frac{\delta'}{2K^{n-1}}\right)^2} - 4e^{-m\frac{\gamma^2}{2}}$, the minimum probability of $P^{(m)}$ is bounded by $\delta = \delta'/2$ and $\|P^{(m)} - P\|_\infty$ is bounded by γ .

For given $\eta < 1$ and $\epsilon < 1$, we let $\gamma = \frac{\epsilon(\delta')^2}{32|E|}K^{-n}$ and

$$m_0 \geq 2 \left(\frac{32|E|K^n}{\epsilon(\delta')^2} \right)^2 \log \left(\frac{8n}{\eta} \right).$$

Now we use Theorem 5 with $\delta = \delta'/2$. First note that $\gamma \leq \delta^{3/2}K^{-1}$ which is required as an assumption for invoking Theorem 5. This leads to

$$|\rho_G^{(m)} - \rho_G| \leq \gamma|E|K^n \frac{8}{\delta^2} \leq \epsilon,$$

with probability (at least)

$$1 - \left(4ne^{-m_0\frac{1}{2}\left(\frac{\delta'}{2K^{n-1}}\right)^2} - 4e^{-m_0\frac{\gamma^2}{2}} \right) \geq 1 - \eta. \quad \blacksquare$$

⁴Note that concentration inequalities for empirical average of a discrete random variable has been also studied in several other works including [77], [78], [79], [80]. In particular, [78] exploits properties of the underlying distribution to provide tight bounds on $\mathbb{P}[\|P^{(m)} - P\|_1 \geq \gamma]$. This in turn provides a bound on $\mathbb{P}[\|P^{(m)} - P\|_\infty \geq \gamma]$ as well. However without additional assumptions on the underlying probability distribution this bound is of the same order as the one presented in Theorem 10 (see [78, Section 3]).

7.7 Proof of Theorem 7

Proof

Let $\mathbf{a} = (\mathbf{a}_1, \dots, \mathbf{a}_n)$. Define

$$\begin{aligned} U_L(\mathbf{a}) &\triangleq \sum_{(i,i') \in E} \sum_{j=1}^L a_{i,j} a_{i',j} (\rho_{i,i'})^j \\ U(\mathbf{a}) &\triangleq \sum_{(i,i') \in E} \sum_{j=1}^\infty a_{i,j} a_{i',j} (\rho_{i,i'})^j. \end{aligned} \quad (68)$$

Let $\tilde{\mathbf{a}}^{(L)} = (\tilde{\mathbf{a}}_1^{(L)}, \dots, \tilde{\mathbf{a}}_n^{(L)})$ and $\tilde{\mathbf{a}} = (\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_n)$ be a solution of the following optimizations

$$\begin{aligned} \max_{\mathbf{a}} \quad &U_L(\mathbf{a}) \\ \text{s.t.} \quad &\mathbf{a} \in \mathcal{C}', \end{aligned} \quad (69)$$

and

$$\begin{aligned} \max_{\mathbf{a}} \quad &U(\mathbf{a}) \\ \text{s.t.} \quad &\mathbf{a} \in \mathcal{C}', \end{aligned} \quad (70)$$

respectively, where $\mathcal{C}'$ is the feasible set of the optimization (35). We first show for any L and any $1 \leq i \leq n$, $\tilde{\mathbf{a}}_i^{(L)} = (0, s_i, \dots, 0)$. We then show for any $1 \leq i \leq n$, $\tilde{\mathbf{a}}_i$ must be equal to $(0, s_i, \dots)$, completing the proof.

Claim 1: We first characterize a global maximizer of optimization (69). Note that the constraint set $\mathcal{C}'$ can be truncated w.l.o.g. to consider $1 \leq j \leq L$. Let Λ be the matrix of correlation coefficients where $[\Lambda]_{i,i'} = \rho_{i,i'}$. Diagonal elements of Λ are all zero, as we ignore self-loops. Define

$$\mathbf{x} = (a_{1,1}, a_{2,1}, \dots, a_{n,1}, a_{1,2}, \dots, a_{n,L})^T. \quad (71)$$

Moreover, define A_0 as an $nL \times nL$ matrix composed of L^2 blocks of size $n \times n$ whose m -th diagonal block is equal to $2\Lambda^m$, where $\Lambda^m[i,j] \triangleq (\Lambda[i,j])^m$. Off-diagonal blocks of A_0 are all zeros. Moreover, define A_i for $1 \leq i \leq n$ as an $nL \times nL$ matrix where $A_i[i + (m-1)n, i + (m-1)n] = 1$ for $1 \leq m \leq L$, otherwise it is zero. Therefore, optimization (69) can be re-written as the following standard quadratic optimization:

$$\begin{aligned} \max_{\mathbf{x}} \quad &\frac{1}{2} \mathbf{x}^T A_0 \mathbf{x} \\ &\frac{1}{2} \mathbf{x}^T A_i \mathbf{x} - \frac{1}{2} \leq 0, \quad 1 \leq i \leq n. \end{aligned} \quad (72)$$

Note that equality constraints of optimization (69) are replaced by inequality ones in optimization (72). This is because, since $(\rho_{i,i'})^2 \geq 0$, solutions of optimization (72) belong to the boundary of its feasible set. Optimization (72) is a non-convex quadratic optimization with quadratic constraints. Reference [81] (Proposition 3.2)⁵ characterizes necessary and sufficient conditions for global minimizers of a generalized form of optimization (72). Let

$$\bar{\mathbf{x}} = [s_1, s_2, \dots, s_n, 0, \dots, 0]^T, \quad (73)$$

where $s_i \in \{-1, 1\}$, for $1 \leq i \leq n$. Using Proposition 3.2 and Theorem 3.1 of [81], to have $\bar{\mathbf{x}}$ as a global minimizer of optimization (72), we need to have

$$\left(\sum_{i=1}^n \lambda_i A_i - A_0 \right) \bar{\mathbf{x}} = 0 \quad (74)$$

⁵Optimization (72) can be stated as a minimization by replacing $-A_0$ instead of A_0 to use Proposition 3.2 of Reference [81]

and

$$\sum_{i=1}^n \lambda_i A_i - A_0 \geq 0, \quad (75)$$

where $\lambda_i \geq 0$, and $A \geq 0$ means that A is a positive semi-definite matrix. Using definitions of A_0 , A_i , and $\bar{\mathbf{x}}$, equation (74) is satisfied iff

$$\lambda_i = 2 \sum_{i' \neq i} s_i s_{i'} \rho_{i,i'} \geq 0, \quad 1 \leq i \leq n. \quad (76)$$

Using (76) and Gerschgorin's circle theorem, if

$$\sum_{i' \neq i} (1 - s_i s_{i'}) \rho_{i,i'} \geq 0, \quad \forall 1 \leq i \leq n, \quad (77)$$

$$\sum_{i' \neq i} s_i s_{i'} \rho_{i,i'} \geq \sum_{i' \neq i} \rho_{i,i'}^2, \quad \forall 1 \leq i \leq n, \quad (78)$$

conditions (75) are satisfied. Thus, $\bar{\mathbf{x}}$ is a global minimizer of optimization (72), establishing that for any L and any $1 \leq i \leq n$ we have $\tilde{\mathbf{a}}_i^{(L)} = (0, 1, \dots, 0)$.

Claim 2: Let $\mathbf{a}^* = (\mathbf{a}_1^*, \dots, \mathbf{a}_n^*)$ where $\mathbf{a}_i^* = (0, s_i, \dots)$ for all $1 \leq i \leq n$. We next show that $\tilde{\mathbf{a}} = \mathbf{a}^*$. We proceed by contradiction. Let

$$\Delta \triangleq U(\tilde{\mathbf{a}}) - U(\mathbf{a}^*). \quad (79)$$

By the contradiction assumption $\Delta > 0$. Since $(\rho_{i,i'})^L \rightarrow 0$ as $L \rightarrow \infty$ for all $i \neq i'$, $U_L(\cdot)$ converges uniformly to $U(\cdot)$. Thus there exists L_0 such that for $L \geq L_0$ and any $\mathbf{a}$, we have

$$|U_L(\mathbf{a}) - U(\mathbf{a})| \leq \frac{\Delta}{4}. \quad (80)$$

Therefore we have

$$U_L(\tilde{\mathbf{a}}) \geq U(\tilde{\mathbf{a}}) - \frac{\Delta}{4} \quad (81)$$

$$\begin{aligned} &= U(\mathbf{a}^*) + \Delta - \frac{\Delta}{4} \\ &\geq U_L(\mathbf{a}^*) - \frac{\Delta}{4} + \Delta - \frac{\Delta}{4} \\ &= U_L(\mathbf{a}^*) + \frac{\Delta}{2} > U_L(\mathbf{a}^*). \end{aligned} \quad (82)$$

This is in contradiction with the assumption that $\mathbf{a}^*$ is the maximizer of $U_L(\cdot)$. Putting Claims 1 and 2 together completes the proof. ■

7.8 Proof of Proposition 3

Proof Under assumptions of Theorem 7 and using the definition of Hermite-Chebyshev polynomials (31), we can restrict the feasible set of optimization (3) to the set of functions $\phi_i(X_i) = s_i X_i$ where $s_i \in \{-1, 1\}$ for all $1 \leq i \leq n$. Moreover, we have

$$\mathbb{E}[\phi_i(X_i) \phi_{i'}(X_{i'})] = s_i s_{i'} \rho_{i,i'}.$$

Furthermore, $\mathbb{E}[s_i X_i] = \mathbb{E}[X_i] = 0$, and $\mathbb{E}[(s_i X_i)^2] = \mathbb{E}[X_i^2] = 1$. This completes the proof. ■

7.9 Proof of Theorem 9

Proof The proof of the first part of this Theorem is straightforward. To prove the second part, we re-write optimization (40) as follows:

$$\begin{aligned} &\max_{(i,i')} \sum \mathbb{E}[g_i(f_i(X_i)) g_{i'}(f_{i'}(X_{i'}))], \\ &\mathbb{E}[g_i(f_i(X_i))] = 0, \quad 1 \leq i \leq n, \\ &\mathbb{E}[g_i(f_i(X_i))^2] = 1, \quad 1 \leq i \leq n. \end{aligned} \quad (83)$$

Define $\phi_i(X_i) = g_i(f_i(X_i))$ for $1 \leq i \leq n$. Since f_i 's are bijective, feasible regions of optimizations (83) and (3) are equal. Under the assumptions of Corollary 4, $\phi_i^*(X_i) = X_i$. Thus, $g_i^*(f_i(X_i)) = X_i$. ■

8 DISCUSSION AND CONCLUSION

The techniques we have developed in this paper can be used in other related formulations as well. In the following, we briefly highlight two examples of such formulations, namely absolute NMC and regularized NMC. Considering further properties of these optimizations is beyond the scope of the present paper.

8.1 Other Objective Functions

The optimization (3) maximizes aggregate pairwise correlations over the network. In some applications, the strength of an association does not depend on the sign of the correlation coefficient. In those cases, one can re-write the NMC optimization (3) to maximize the total absolute pairwise correlations over the graph as follows:

Definition 9 (Absolute Network Maximal Correlation). Consider the following optimization:

$$\rho_G^A(X_1, \dots, X_n) \triangleq \sup_{\phi_1, \dots, \phi_n} \sum_{(i,i') \in E} |\mathbb{E}[\phi_i(X_i) \phi_{i'}(X_{i'})]|, \quad (84)$$

such that $\phi_i : \mathcal{X}_i \rightarrow \mathbb{R}$ is Borel measurable, $\mathbb{E}[\phi_i(X_i)] = 0$, and $\mathbb{E}[\phi_i(X_i)^2] = 1$, for all $1 \leq i \leq n$. $G = (V, E)$ is a graph with vertices $V = \{1, 2, \dots, n\}$ and edges $E = \{(i, i') : i, i' \in V, i \neq i'\}$. We refer to this optimization as an absolute NMC optimization.

A similar approach can be used to characterize the absolute NMC optimization (84) by introducing extra variables $s_{i,i'}$ to represent correlation signs of edges (i, i') :

$$\begin{aligned} &\sup_{\phi_1, \dots, \phi_n} \sum_{(i,i') \in E} s_{i,i'} \mathbb{E}[\phi_i(X_i) \phi_{i'}(X_{i'})] \\ &\mathbb{E}[\phi_i(X_i)] = 0, \quad \mathbb{E}[\phi_i^2(X_i)] = 1, \quad 1 \leq i \leq n \\ &s_{i,i'} \in \{-1, 1\}, \quad 1 \leq i, i' \leq n. \end{aligned} \quad (85)$$

The NMC optimization (3) results in n possibly nonlinear transformation functions $\phi_i^*(X_i)$ whose correlation with the original variables can be small. In some applications, one may wish to restrict the set of possible transformations of the NMC optimization to control the correlation between transformed and original variables. This can be done by introducing a regularization term in the definition of NMC as follows.

Definition 10 (Regularized NMC). The regularized NMC of real-valued $X_1, \dots, X_n$ connected by a graph $G = (V, E)$ is defined as the solution of the following optimization:

$$\begin{aligned} \rho_G^R(X_1, \dots, X_n) = \sup_{\phi_1, \dots, \phi_n} (1 - \lambda) \sum_{(i, i') \in E} \mathbb{E}[\phi_i(X_i) \phi_{i'}(X_{i'})] \\ + \lambda \sum_{i \in V} \mathbb{E}[\phi_i(X_i) (X_i - \mathbb{E}[X_i])], \end{aligned} \quad (86)$$

such that $\phi_i : \mathcal{X}_i \rightarrow \mathbb{R}$ is Borel measurable, $\mathbb{E}[\phi_i(X_i)] = 0$, and $\mathbb{E}[\phi_i(X_i)^2] = 1$, for all $1 \leq i \leq n$. $G = (V, E)$ is a graph with vertices $V = \{1, 2, \dots, n\}$ and edges $E = \{(i, i') : i, i' \in V, i \neq i'\}$. $0 \leq \lambda \leq 1$ is the regularization parameter.

Unlike MC and NMC, which only depend on the joint distributions of variables, the regularized NMC depends on both joint distributions and support of variables because of the regularization term. Moreover, one can define regularized absolute NMC similarly to the optimization (84).

Let the optimal transformation functions computed by optimization (86) be $\phi_{i,\lambda}^*(X_i)$. If $\lambda = 0$, $\phi_{i,\lambda}^*(X_i) = \phi_i^*(X_i)$, while if $\lambda = 1$, $\phi_{i,\lambda}^*(X_i) = X_i$. By varying λ between 0 and 1, $\phi_{i,\lambda}^*(X_i)$ vary from $\phi_i^*(X_i)$ to X_i . Define

$$R_{G,\lambda}(X_1, \dots, X_n) \triangleq \sum_{(i, i') \in E} \mathbb{E}[\phi_{i,\lambda}^*(X_i) \phi_{i',\lambda}^*(X_{i'})].$$

Therefore, $R_{G,0} = \rho_G$ while $R_{G,1}$ is the total linear correlations over the network. By the definition of NMC, $R_{G,0} \geq R_{G,1}$. Note that we can use a similar algorithm to Algorithm 2 to compute regularized NMC of Definition 10. The objective function of the regularized NMC optimization (86) can be written as follows:

$$\sum_{i \in V} \mathbb{E} \left[\phi_i(X_i) \left(\frac{1-\lambda}{2} \sum_{j \in \mathcal{N}(i)} \phi_{i'}(X_{i'}) + \lambda(X_i - \mathbb{E}[X_i]) \right) \right], \quad (87)$$

where $\mathcal{N}(i)$ represents the set of neighbors of node i in the graph $G = (V, E)$. To compute the regularized NMC, one can use an algorithm similar to the ACE Algorithm 2 with the following updates for transformation functions:

$$\phi_i^*(X_i) = \mathbb{E} \left[\frac{1-\lambda}{2} \sum_{j \in \mathcal{N}(i)} \phi_{i'}(X_{i'}) + \lambda(X_i - \mathbb{E}[X_i]) | X_i \right]. \quad (88)$$

8.2 Conclusion

In this paper, we propose NMC as a measure to capture nonlinear associations among variables. We show that NMC extends the standard bivariate MC to the case of having large number of variables, by assigning each variable to a single transformation function, thus avoiding over-fitting issues of using multiple MC optimizations over variable pairs. We also introduce a regularized NMC optimization which penalizes total distances of inferred transformed variables from the original ones. One can use other standard regularization techniques to further restrict inferred nonlinear functions in practical applications.

One of the main contributions of this work is providing a unifying framework to compute NMC (and therefore, MC) for both discrete and continuous variables using projections over appropriate Hilbert spaces. Using this framework, we establish a connection between the NMC optimization with the MCP and MEP for discrete random variables, and with the Max-Cut problem for jointly Gaussian variables. Using these relationships, we provide efficient algorithms to compute NMC in different cases. Note that properties of NMC for

finite discrete variables characterized in Theorems 5 and 6 depend on minimum marginal probabilities of variables (i.e., δ). Therefore, extending these properties to the continuous case is not straightforward and requires exploiting techniques tailored for continuous variables and measures. To compute NMC for continuous random variables with general distributions, one can use the proposed optimization framework by choosing appropriate orthonormal basis for Hilbert spaces. For example, we use projections over Hermite-Chebyshev polynomials to characterize a solution of the NMC optimization for jointly Gaussian variables. Finally note that the inferred, possibly nonlinear, functions of the NMC optimization can be used in other applications such as nonlinear regression.

9 ACKNOWLEDGEMENTS

The authors would like to thank Yue Li (MIT) and Gerald Quon (UC Davis) for helpful discussions regarding the cancer application, Yue Li for providing the cancer datasets, and Flavio P. Calmon (MIT) and Vahid Montazerhojati (MIT) for helpful discussions. In addition, the authors would like to thank anonymous reviewers and associate editor for useful comments.

REFERENCES

- [1] R. Albert, "Network inference, analysis, and modeling in systems biology," *The Plant Cell*, vol. 19, no. 11, pp. 3327–3338, 2007.
- [2] M. S. Handcock, *Statistical models for social networks: Inference and Degeneracy*. National Academies Press, 2002.
- [3] A. G. Haldane, "Rethinking the financial network," *Speech delivered at the Financial Student Association, Amsterdam, April*, pp. 1–26, 2009.
- [4] K. Pearson, "Note on regression and inheritance in the case of two parents," *Proceedings of the Royal Society of London*, pp. 240–242, 1895.
- [5] B. Thompson, "Canonical correlation analysis," *Encyclopedia of statistics in behavioral science*, 2005.
- [6] L. L. Scharf and C. T. Mullis, "Canonical coordinates and the geometry of inference, rate, and capacity," 2000.
- [7] C. Spearman, "The proof and measurement of association between two things," *The American Journal of Psychology*, 1904.
- [8] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.
- [9] D. N. Reshef, Y. A. Reshef, H. K. Finucane, S. R. Grossman, G. McVean, P. J. Turnbaugh, E. S. Lander, M. Mitzenmacher, and P. C. Sabeti, "Detecting novel associations in large data sets," *Science*, 2011.
- [10] C. Chow and C. Liu, "Approximating discrete probability distributions with dependence trees," *IEEE Transactions on Information Theory*, vol. 14, no. 3, pp. 462–467, 1968.
- [11] H. Liu, J. Lafferty, and L. Wasserman, "The nonparanormal: Semiparametric estimation of high dimensional undirected graphs," *The Journal of Machine Learning Research*, vol. 10, pp. 2295–2328, 2009.
- [12] H. Liu, F. Han, M. Yuan, J. Lafferty, L. Wasserman *et al.*, "High-dimensional semiparametric gaussian copula graphical models," *The Annals of Statistics*, vol. 40, no. 4, pp. 2293–2326, 2012.
- [13] D. Lopez-Paz, P. Hennig, and B. Schölkopf, "The randomized dependence coefficient," in *Advances in Neural Information Processing Systems*, pp. 1–9, 2013.
- [14] H. O. Hirschfeld, "A connection between correlation and contingency," in *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 31, no. 04, pp. 520–524, 1935.
- [15] H. Gebelein, "Das statistische problem der korrelation als variations- und eigenwertproblem und sein zusammenhang mit der ausgleichsrechnung," *ZAMM-Journal of Applied Mathematics and Mechanics*, vol. 21, no. 6, pp. 364–379, 1941.
- [16] O. Sarmanov, "Maximum correlation coefficient (nonsymmetric case)," *Selected Translations in Mathematical Statistics and Probability*, vol. 2, pp. 207–210, 1962.
- [17] A. Rényi, "On measures of dependence," *Acta Mathematica Hungarica*, vol. 10, no. 3, pp. 441–451, 1959.
- [18] H. S. Witsenhausen, "On sequences of pairs of dependent random variables," *SIAM Journal on Applied Mathematics*, 1975.
- [19] R. Ahlswede and P. Gács, "Spreading of sets in product spaces and hypercontraction of the markov operator," *The Annals of Probability*, pp. 925–939, 1976.

- [20] H. O. Lancaster, "Some properties of the bivariate normal distribution considered in the form of a contingency table," *Biometrika*, vol. 44, no. 1-2, pp. 289–292, 1957.
- [21] L. Breiman and J. H. Friedman, "Estimating optimal transformations for multiple regression and correlation," *Journal of the American Statistical Association*, vol. 80, no. 391, pp. 580–598, 1985.
- [22] A. Buja, "Remarks on functional canonical variates, alternating least squares methods and ace," *The Annals of Statistics*, 1990.
- [23] V. Anantharam, A. Gohari, S. Kamath, and C. Nair, "On maximal correlation, hypercontractivity, and the data processing inequality studied by Erkip and Cover," *arXiv preprint:1304.6133*, 2013.
- [24] Y. Polyanskiy, "Hypothesis testing via a comparator and hypercontractivity," *Prob. Peredachi Inform.*, 2013.
- [25] M. Raginsky, "Strong data processing inequalities and ϕ -sobolev inequalities for discrete channels," *IEEE Transactions on Information Theory*, vol. 62, no. 6, pp. 3355–3389, 2016.
- [26] T. Courtade, "Outer bounds for multiterminal source coding via a strong data processing inequality," in *IEEE International Symposium on Information Theory*, pp. 559–563, 2013.
- [27] F. Calmon, M. Varia, and M. Médard, "An exploration of the role of principal inertia components in information theory," in *Information Theory Workshop*, pp. 252–256, 2014.
- [28] C. T. Li and A. El Gamal, "Maximal correlation secrecy," in *IEEE International Symposium on Information Theory*, pp. 2939–2943, 2015.
- [29] A. Makhdoumi, F. Calmon, and M. Médard, "Forgot your password: Correlation dilution," in *IEEE International Symposium on Information Theory*, pp. 2944–2948, 2015.
- [30] F. Calmon, A. Makhdoumi, and M. Médard, "Fundamental limits of perfect privacy," in *IEEE International Symposium on Information Theory*, pp. 1796–1800, 2015.
- [31] F. Calmon, M. Varia, and M. Médard, "On information-theoretic metrics for symmetric-key encryption and privacy," in *Allerton Conference on Communication, Control, and Computing*, pp. 889–894, 2014.
- [32] F. Calmon, M. Varia, M. Médard, M. M. Christiansen, K. R. Duffy, and S. Tessaro, "Bounds on inference," in *Allerton Conference on Communication, Control, and Computing*, pp. 567–574, 2013.
- [33] S. Beigi and A. Gohari, "On the duality of additivity and tensorization," *arXiv preprint arXiv:1502.00827*, 2015.
- [34] O. Etesami and A. Gohari, "Maximal rank correlation," *arXiv preprint arXiv:1511.04191*, 2015.
- [35] F. Farnia, M. Razaviyayn, S. Kannan, and D. Tse, "Minimum HGR correlation principle: From marginals to joint distribution," *IEEE International Symposium on Information Theory*, pp. 1377–1381, 2015.
- [36] M. Razaviyayn, F. Farnia, and D. Tse, "Discrete Renyi classifiers," *arXiv preprint arXiv:1511.01764*, 2015.
- [37] A. Makur and L. Zheng, "Bounds between contraction coefficients," *arXiv preprint arXiv:1510.01844*, 2015.
- [38] M. T. Chu and J. L. Waterson, "On a multivariate eigenvalue problem, part I: Algebraic theory and a power method," *SIAM Journal on Scientific Computing*, vol. 14, no. 5, pp. 1089–1106, 1993.
- [39] H. Hotelling, "The most predictable criterion," *Journal of Educational Psychology*, 1935.
- [40] —, "Relations between two sets of variates," *Biometrika*, 1936.
- [41] L.-H. Zhang and M. T. Chu, "Computing absolute maximum correlation," *IMA Journal of Numerical Analysis*, vol. 32, no. 1, pp. 163–184, 2012.
- [42] P. Klein, S. A. Plotkin, and S. Rao, "Excluded minors, network decomposition, and multicommodity flow," in *ACM symposium on Theory of computing*, 1993.
- [43] S. Rao, "Small distortion and volume preserving embeddings for planar and Euclidean metrics," in *ACM symposium on Computational geometry*, 1999.
- [44] D. R. Karger and M. Ruhl, "Finding nearest neighbors in growth-restricted metrics," in *ACM symposium on Theory of computing*, 2002.
- [45] M. R. Garey, D. S. Johnson, and L. Stockmeyer, "Some simplified NP-complete graph problems," *Theoretical Computer Science*, 1976.
- [46] M. X. Goemans and D. P. Williamson, "Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming," *Journal of the ACM*, vol. 42, pp. 1115–1145, 1995.
- [47] M. J. Wainwright and M. I. Jordan, "Graphical models, exponential families, and variational inference," *Foundations and Trends in Machine Learning*, vol. 1, no. 1-2, pp. 1–305, 2008.
- [48] J. Neter, M. H. Kutner, C. J. Nachtsheim, and W. Wasserman, *Applied linear statistical models*. Irwin Chicago, 1996.
- [49] K. Baba, R. Shibata, and M. Sibuya, "Partial correlation and conditional correlation as measures of conditional independence," *Australian and New Zealand Journal of Statistics*, vol. 46, pp. 657–664, 2004.
- [50] O. Banerjee, L. El Ghaoui, and A. d'Aspremont, "Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data," *The Journal of Machine Learning Research*, vol. 9, pp. 485–516, 2008.
- [51] Y. Li and Z. Zhang, "Potential microRNA-mediated oncogenic inter-cellular communication revealed by pan-cancer analysis," *Scientific reports*, vol. 4, 2014.
- [52] W. Bryc and A. Dembo, "On the maximum correlation coefficient," *Theory of Probability and its Applications*, vol. 49, no. 1, pp. 132–138, 2005.
- [53] P. Horst, *Factor analysis of data matrices*. Holt, Rinehart and Winston Publisher, 1965.
- [54] —, "Relations among sets of measures," *Psychometrika*, 1961.
- [55] L.-H. Zhang and L.-Z. Liao, "An alternating variable method for the maximal correlation problem," *Journal of Global Optimization*, vol. 54, no. 1, pp. 199–218, 2012.
- [56] D. P. Bertsekas, *Nonlinear programming*. Athena scientific, 1999.
- [57] G. H. Golub and C. F. Van Loan, *Matrix Computations*. JHU Press, 2012.
- [58] K. Jung and D. Shah, "Local algorithms for approximate inference in minor-excluded graphs," in *Advances in Neural Information Processing Systems*, pp. 729–736, 2008.
- [59] J. Whittaker, *Graphical models in applied multivariate statistics*. Wiley Publishing, 2009.
- [60] M. J. Choi, V. Y. Tan, A. Anandkumar, and A. S. Willsky, "Learning latent tree graphical models," *Journal of Machine Learning Research*, vol. 12, pp. 1771–1812, 2011.
- [61] B. W. Stewart, P. Kleihues, and et al., *World cancer report*. IARC press Lyon, vol. 57, 2003.
- [62] G. A. Calin and C. M. Croce, "MicroRNA signatures in human cancers," *Nature Reviews Cancer*, vol. 6, no. 11, pp. 857–866, 2006.
- [63] W. Huber, A. Von Heydebreck, H. Sültmann, A. Poustka, and M. Vingron, "Variance stabilization applied to microarray data calibration and to the quantification of differential expression," *Bioinformatics*, vol. 18, pp. S96–S104, 2002.
- [64] J. K. Pritchard, M. Stephens, and P. Donnelly, "Inference of population structure using multilocus genotype data," *Genetics*, vol. 155, no. 2, pp. 945–959, 2000.
- [65] M. E. Newman, "Modularity and community structure in networks," *Proceedings of the National Academy of Sciences*, 2006.
- [66] P. Good, *Permutation tests: a practical guide to resampling methods for testing hypotheses*. Springer Science & Business Media, 2013.
- [67] E. L. Kaplan and P. Meier, "Nonparametric estimation from incomplete observations," *Journal of the American statistical association*, vol. 53, no. 282, pp. 457–481, 1958.
- [68] J. M. Bland and D. G. Altman, "The logrank test," *BMJ*, vol. 328, no. 7447, p. 1073, 2004.
- [69] Y. Benjamini and Y. Hochberg, "Controlling the false discovery rate: a practical and powerful approach to multiple testing," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 289–300, 1995.
- [70] G. S. Bailey, A. P. Reddy, C. B. Pereira, U. Harttig, W. Baird, J. M. Spitsbergen, J. D. Hendricks, G. A. Orner, D. E. Williams, and J. A. Swenberg, "Nonlinear cancer response at ultralow dose: a 40800-animal ed001 tumor and biomarker study," *Chemical research in toxicology*, vol. 22, no. 7, pp. 1264–1276, 2009.
- [71] X. Huang, W. Qian, I. H. El-Sayed, and M. A. El-Sayed, "The potential use of the enhanced nonlinear properties of gold nanospheres in photothermal cancer therapy," *Lasers in surgery and medicine*, 2007.
- [72] J. S. Lowengrub, H. B. Frieboes, F. Jin, Y. Chuang, X. Li, P. Macklin, S. Wise, and V. Cristini, "Nonlinear modelling of cancer: bridging the gap between cells and tumours," *Nonlinearity*, vol. 23, no. 1, 2010.
- [73] L. Nagel and R. Rohrer, "Computer analysis of nonlinear circuits, excluding radiation (cancer)," *IEEE Journal of Solid-State Circuits*, 1971.
- [74] J. G. Wagner, J. W. Gyves, P. L. Stetson, S. C. Walker-Andrews, I. S. Wollner, M. K. Cochran, and W. D. Ensminger, "Steady-state nonlinear pharmacokinetics of 5-fluorouracil during hepatic arterial and intravenous infusions in cancer patients," *Cancer research*, 1986.
- [75] L. Devroye, "The equivalence of weak, strong and complete convergence in l_1 for kernel density estimates," *The Annals of Statistics*, 1983.
- [76] D. Berend and A. Kontorovich, "On the convergence of the empirical distribution," *arXiv preprint:1205.6711*, 2012.
- [77] G. abor Lugosi and P. Massart, "A sharp concentration inequality with applications," 1999.
- [78] T. Weissman, E. Ordentlich, G. Seroussi, S. Verdu, and M. J. Weinberger, "Inequalities for the l_1 deviation of the empirical distribution," *Hewlett-Packard Labs, Tech. Rep.*, 2003.
- [79] P. Massart, *Concentration inequalities and model selection*. Springer, 2007.
- [80] S. Boucheron, G. Lugosi, and P. Massart, *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- [81] V. Jeyakumar, A. M. Rubinov, and Z. Wu, "Non-convex quadratic minimization problems with quadratic constraints: global optimality conditions," *Mathematical Programming*, vol. 110, no. 3, pp. 521–541, 2007.