
Structure Learning of Partitioned Markov Networks

Song Liu[†]
Taiji Suzuki[‡]
Masashi Sugiyama[★]
Kenji Fukumizu[†]

LIU@ISM.AC.JP
SUZUKI.T.CT@M.TITECH.AC.JP
SUGI@K.U-TOKYO.AC.JP
FUKUMIZU@ISM.AC.JP

[†]The Institute of Statistical Mathematics, Tokyo 190-8562, Japan.

[‡]Tokyo Institute of Technology, Tokyo 152-8552, Japan; PRESTO, Japan Science and Technological Agency (JST).

[★]University of Tokyo, Chiba 277-8561, Japan.

Abstract

We learn the structure of a Markov Network between two groups of random variables from joint observations. Since modelling and learning the full MN structure may be hard, learning the links between two groups directly may be a preferable option. We introduce a novel concept called the *partitioned ratio* whose factorization directly associates with the Markovian properties of random variables across two groups. A simple one-shot convex optimization procedure is proposed for learning the *sparse* factorizations of the partitioned ratio and it is theoretically guaranteed to recover the correct inter-group structure under mild conditions. The performance of the proposed method is experimentally compared with the state of the art MN structure learning methods using ROC curves. Real applications on analyzing bipartisanship in US congress and pairwise DNA/time-series alignments are also reported.

1. Introduction

An undirected graphical model, or a Markov Network (MN) (Koller & Friedman, 2009; Wainwright & Jordan, 2008) has a wide range of applications in real world, such as natural language processing, computer vision, and computational biology. The structure of MN, which encodes the interactions among random variables, is one of the key interests of MN learning tasks. However, on a high-dimensional dataset, learning the full MN structure can be cumbersome since we may not have enough knowledge to model the entire MN, or our application only concerns a specific portion of the MN structure.




Figure 1. An illustration of a full MN (left) and PMN (right). Full MN models all the connections among random variables, while PMN only models the interactions between groups (red edges) and does not care connections within groups.

Rather than considering the full MN structure over the complete set of random variables, we focus on learning a portion of the MN structure that links *two groups of random variables*, namely the Partitioned Markov Network (PMN). PMN is suitable for describing the “inter-group relations”. For example, politicians in US Congress are naturally grouped into two parties (Democrats and Republicans). Learning a PMN on congresspersons via their voting records will reveal bipartisan collaborations among them. A full gene network may have complicated structure. However if genes can be clustered into a few homologous groups, PMN can help us understand how genes in different functioning groups interact with each other. An illustration of a full MN and a PMN is shown in Figure 1.

Since a PMN can be regarded as a “sub-structure” of a full MN, a naive approach may be learning a full MN over the complete set of random variables and figuring out its PMN. In fact, the machine learning community has seen huge progresses on learning the *sparse* structures of MNs, thanks to the pioneer works on sparsity inducing norms (Tibshirani, 1996; Zhao & Yu, 2006; Wainwright, 2009).

A majority of the previous works fall into the category of the regularized maximum likelihood approach which maximizes the likelihood function of a probabilistic model under sparsity constraints. Graphical lasso (Friedman et al., 2008; Banerjee et al., 2008) considers a joint Gaussian model parameterized by the inverse covariance matrix,

where zero elements indicate the conditional independence among random variables, while others have developed useful variations of graphical lasso in order to loosen the Gaussianity assumed on data (Liu et al., 2009; Loh & Wainwright, 2012). SKEPTIC (Liu et al., 2012) is a semi-parametric approach that replaces the covariance matrix with the correlation matrix, such as Kendall’s Tau in MN learning.

The latest advances along this line of research has been made by considering a node-wise conditional probabilistic model. Instead of learning all the structures in one shot, such a method focuses on learning the neighborhood structure of a single random variable at a time. Maximizing the conditional likelihood leads to simple logistic regression (in the case of the Ising model) (Ravikumar et al., 2010) or linear regression (in the case of the Gaussian model) (Meinshausen & Bühlmann, 2006).

Unfortunately, the maximum (conditional) likelihood method can be difficult to compute for general non-Gaussian graphical models, since computing the normalization term is in general intractable. Though one may use sampling such as Monte-carlo methods (Robert & Casella, 2005) to approximate the normalization term, there is no universal guideline telling how to choose sampling parameters so that the approximation error is minimized.

A more severe problem is that sparsity approaches may have difficulties when learning a dense MN. Specifically, the samples size required for a successful structure recovery grows quadratically with the number of connected neighbors (Raskutti et al., 2009; Ravikumar et al., 2010). However, it is quite reasonable to assume that in some applications, one node may have many neighbors within its own group while connections to the other group are sparse: a congressperson is very well connected to other members inside his/her party but has only a few links with the opposition party. Genes in a homologous group may have dense structure but they only interact with another group of genes via a few ties.

Is there a way to *directly* obtain the PMN structure? Neither maximizing a joint nor conditional likelihood take the “partition information” into account and interactions are modelled *globally*. However PMN encodes only the *local* conditional independence between groups, and the requirement for obtaining a good estimator should be much milder.

The above intuition leads us to a novel concept of the Partitioned Ratio (PR). Given a set of partitioned random variables $X = (X_1, X_2)$, PR is the ratio between the joint probability $P(X)$ and the product between its marginals $P(X_1)P(X_2)$, i.e. $\frac{P(X)}{P(X_1)P(X_2)}$. In the same way that the joint distribution can be decomposed into clique potentials

of MN, we prove PR also factorizes over subgraph structures called *passages*, which indicate the connectivity between two groups of random variables X_1 and X_2 in a PMN.

Conventionally, PR is a measure of the independence between two sets of random variables. In this paper, we show that the factorization of this quantity indicates the linkage between two groups of random variables, which is a natural extension of the regular usage of PR.

Most importantly, we show the sparse factorization of this quantity may be learned via a one shot convex optimization procedure, which can be solved efficiently even for the general, non-Gaussian distributions. The correct recovery of sparse passage structure is theoretically guaranteed under the assumption that the sample size increases with the number of passages which is not related to the structure density of the entire MN.

This paper is organized as follows. In Section 2, we review the Hammersley and Clifford theorem (Section 2.1) and define some notations as preliminaries (Section 2.2). The factorization theorems of PMN are introduced in Section 3 with a few simplifications. We give an estimator to obtain the sparse factorization of PR in Section 4 and prove its recovered structure is consistent in Section 5. Finally, experimental results on both artificial and real-world datasets are reported in Section 6.

2. Background and Preliminaries

In this section, we review the factorization theorems of MN. We limit our discussions on strictly positive distributions from now on. A graph is always assumed to be finite, simple, and undirected.

2.1. Background and Motivation

Definition 1 (MN). For a joint probability $P(X)$ of random variables $X = \{X_1, X_2, \dots, X_m\}$, if for all i , $P(X_i | \setminus X_i) = P(X_i | X_{N(i)})$, where $X_{N(i)}$ is the neighbors of node X_i in graph G , then P is an MN with respect to G .

Definition 2 (Gibbs Distribution). For a joint distribution P on a set of random variables X , if the joint density can be factorized as

$$P(X) = \frac{1}{Z} \prod_{C \in \mathbf{C}(G)} \phi_C(X_C),$$

where Z is the normalization term, $\mathbf{C}(G)$ is the set of complete subgraphs of G and each factor ϕ_C is defined only on a subset of random variables X_C , then P is called a Gibbs distribution that factorizes over G .

Theorem 1 (See e.g., Hammersley & Clifford (1971)). If




Figure 2. If (I) is an MN over X , then (I), (II), (III) are all PMNs over X . If (I) is a PMN over X , (I), (II), (III) are not necessarily the MN over X (but still PMNs over X).

P is an MN with respect to G (Definition 1), then P is a Gibbs distribution that factorizes over G (Definition 2)

Theorem 2 (See e.g., Koller & Friedman (2009)). *If P is a Gibbs distribution that factorizes over G then P is an MN with respect to G .*

Theorems 1 and 2 are the keystones of many MN structure learning methods. It states, by learning a sparse factorization of a joint distribution, we are able to spot the structure of a graphical model. However, learning a joint distribution has never been an easy task due to the normalization issue and if the task is to learn a PMN that only concerns conditional independence across two groups, such an approach seems to “solve a more general task as an intermediate step” (Vapnik, 1998).

Does there exist an alternative to the joint distribution, whose factorization relates to the structure of PMN? Ideally, such factorization should be efficiently estimated from samples with a tractable normalization term and the estimation procedure should provide good statistical guarantees.

In the rest of the paper, we show PR has the desired properties to indicate the structure of a PMN: It is factorized over the structure of a PMN (Section 3) and easy to estimate from joint samples (Section 4) with good statistical properties (Section 5).

2.2. Definitions

Notations. Sets are denoted by upper-case letters, e.g., A, B . An upper-case with a lower-case subscript A_i means the i -th element in A . Set operator $A \setminus B$ means excluding set B from set A . $\setminus B$ means the whole set excluding the set B . $A = (A_1, A_2)$ is a partition of set A and an upper-case followed by an integer number, e.g. A_1, A_2 means groups divided by such a partition. Given a graph $L = \langle N, E \rangle$ and a subgraph $K \subseteq L$, N_K or E_K denotes the subset of N or E whose elements are indexed topologically by K . Upper-case with bold font, e.g. \mathbf{K} , is a set of sets.




Figure 3. (Left) ABCD and (Right) AB...Z are two passages.

PMN and Gibbs Partitioned Ratio. Now, we formally define a graph $G = \langle X, E \rangle$, where X is a set of random variables and $X = (X_1, X_2)$, i.e. $X_1 \cap X_2 = \emptyset$, $X_1 \cup X_2 = X$ and $X_1, X_2 \neq \emptyset$. The concept of PMN can now be defined.

Definition 3 (PMN). *For a joint probability $P(X)$, $X = (X_1, X_2)$, if*

$$P(X_i | X_1 \cup X_{N(i)} \setminus X_i) = P(X_i | \setminus X_i), \forall X_i \in X_1, \quad (1)$$

$$P(X_i | X_2 \cup X_{N(i)} \setminus X_i) = P(X_i | \setminus X_i), \forall X_i \in X_2, \quad (2)$$

then P is a PMN with respect to G .

The following proposition is a consequence of Definition 3, and an example is visualized in Figure 2.

Proposition 1. *If P is an MN with respect to G , then P is a PMN with respect to G , but not vice versa.*

Proposition 2. *If P is a PMN with respect to G , $X_u \in X_1, X_v \in X_2$, and $v \notin N(u)$, then $X_u \perp\!\!\!\perp X_v | \setminus \{X_u, X_v\}$.*

See Appendix A for the proof.

The concept of *Passage* is defined as follows:

Definition 4 (Passage). *Let $X = (X_1, X_2)$. We define a **passage** B of G as a subgraph of G , such that $X_B \cap X_1 \neq \emptyset$, $X_B \cap X_2 \neq \emptyset$, and $\forall X_u \in (X_1 \cap X_B), \forall X_v \in (X_2 \cap X_B)$, we have $\text{edge}(X_u, X_v) \in E_B$.*

Here we highlight two of the passage structures of two graphs in Figure 3.

From definition, we can see all cliques that go across two groups are passages, but not all passages are cliques:

Proposition 3. *Let $X = (X_1, X_2)$. Given a passage B of G , B is a complete subgraph if and only if $\forall X_u, X_v \in X_B \cap X_1$, $\text{edge}(X_u, X_v) \in E_B$ and $\forall X_u, X_v \in X_B \cap X_2$, $\text{edge}(X_u, X_v) \in E_B$.*

As an analogy to a Gibbs distribution used in the Hammersley-Clifford Theorem, we define the Gibbs partitioned ratio.

Definition 5 (Gibbs Partitioned Ratio). *For a joint distribution P over $X = (X1, X2)$, if the partitioned ratio has the form*

$$\frac{P(X1, X2)}{P(X1)P(X2)} = \frac{1}{Z} \prod_{B \in \mathbf{B}(G)} \phi_B(X_B),$$

where $\mathbf{B}(G)$ is the set of all passages in G , then $\frac{P(X1, X2)}{P(X1)P(X2)}$ is called the Gibbs partitioned ratio (GPR) over G .

3. Factorization over Passages

In this section, we will investigate the question: can we have a similar factorization theorem like Theorems 1 and 2 for PMN? If so, learning the sparse factorization of PR may reveal the Markovian properties among random variables.

3.1. Fundamental Properties

There are two steps for introducing our factorization theorems. The first step is establishing the Markovian property of random variables using the factorization of PR.

Theorem 3. *Given $X = (X1, X2)$, if PR $\frac{P(X1, X2)}{P(X1)P(X2)}$ is a GPR over a graph G then P is a PMN with respect to G .*

See Appendix B for the proof.

Next, let us prove the other direction: From the Markovian property to the factorization.

Theorem 4. *Given $X = (X1, X2)$, if P is a PMN with respect to a graph G , then $\frac{P(X1, X2)}{P(X1)P(X2)}$ is a GPR over G .*

See Appendix C for the proof.

Simply, the factorization of a GPR is only related to the “linkage” (or rigorously, passages) between two groups. Interestingly, if we have an MN whose groups are linked via a few “bottleneck” passages, then the factorization is simply over those sparse passages, no matter how densely the graph are connected within each group. This gives PMN a significant advantage over traditional MN in terms of modelling: If the interactions between groups are simple (e.g. linear), we do not need to care the interactions within groups, even if they are highly complicated (e.g. non-linear). For example, in the bipartisan analysis problem, a PR over congresspersons can be represented only via a few cross-party links, and a large chunk of connections between congresspersons within their own party can be ignored, no matter how complicated they are.

Theorems 3 and 4 point out a promising direction for structural learning of a PMN: Once the sparse factorization of a GPR is learned, we are able to recover the *sparse passages* of a PMN partitioned into two groups.

3.2. Simplification of Passage Factorization

The Hammersley-Clifford theorem (Theorem 1) shows P factorizes over cliques of G , given P is an MN with respect to G . However, if one does not know the maximum size of cliques, the model of a probability function has to consider factors on all potential cliques, i.e., all subsets of X . It is unrealistic to construct a model with $2^{|X|}$ factors under the high-dimensional setting.

Therefore, a popular assumption called “pairwise MN” (Koller & Friedman, 2009; Murphy, 2012) has been widely used to lower the computational burden of MN structure learning. It assumes that in P , all clique factors can be further recovered using only bivariate and univariate components which give rise to a pairwise model with only $(|X|^2 + |X|)/2$ factors. Some well known MNs, such as Gaussian MN and Ising model are all examples of pairwise MNs.

Similar issues also happen when modelling GPR. There are $(2^{|X1|} - 1)(2^{|X2|} - 1)$ possible passage potentials for the set of random variables $X = (X1, X2)$. Following the same spirit, we can consider a simplified model of PR by assuming that all passage potentials of the GPR must factorize in a pairwise fashion, i.e.:

Definition 6 (Pairwise PR). *For a joint distribution P over $X = (X1, X2)$, if the partitioned ratio has the form*

$$\begin{aligned} \frac{P(X1, X2)}{P(X1)P(X2)} &= \frac{1}{Z} \prod_{B \in \mathbf{B}(G)} \phi_B(X_B) \\ &= \frac{1}{Z} \prod_{B \in \mathbf{B}(G)} \prod_{X_u, X_v \in X_B, u \leq v} h_{u,v}(X_u, X_v), \end{aligned}$$

then $\frac{P(X1, X2)}{P(X1)P(X2)}$ is called the pairwise Gibbs partitioned ratio (pairwise PR) over G .

If we can assume the GPR we hope to learn is also a pairwise PR, the model may only contain $(|X|^2 + |X|)/2$ pairwise factors, and is much easier to construct.

In fact, pairwise PR does not have straightforward relationship with pairwise MN, i.e., a PR of a pairwise MN may not be a pairwise PR, meanwhile the joint distribution corresponding to a pairwise PR may not be a pairwise MN, since the pairwise MN and the pairwise PR apply the same assumption on the parameterizations of two fundamentally different quantities, the joint probability and the PR respectively.

Whether one should impose such an assumption on joint probability or PR is totally up to the application, as neither parameterization is always superior to the other. If the application focuses on learning the connections between two groups, we believe imposing such an assumption on PR directly is more sensible.

However, as a special case, a joint Gaussian distribution is a pairwise MN, and its PR is also a pairwise PR.

Proposition 4. *If P over $X = (X1, X2)$ is a zero-mean Gaussian distribution, then the PR $\frac{P(X1, X2)}{P(X1)P(X2)}$ is a pairwise PR.*

Since the Gaussian distribution factorizes over pairwise potentials, and the marginal distribution $P(X1)$ and $P(X2)$ are still Gaussian distributions. From the construction of the potential function (6) in the proof of Theorem 4, we can verify this statement. Moreover, one can show it has the pairwise factor $h_{u,v}(X_u, X_v) = \exp(\theta_{u,v} \cdot X_u X_v)$, where $\theta_{u,v}$ is the parameter.

This pairwise assumption together with factorization theorems motivate us to recover the structure of PMN by learning a sparse pairwise PR model: For any $X_u \in X1, X_v \in X2$, if X_u, X_v appear in the same pairwise factor of a PR model, they must be at least involved in one of the passage potentials.

4. Estimating PR from Samples

To estimate PR using such a model, we require a set of samples

$$\{\mathbf{x}^{(i)}\}_{i=1}^n \stackrel{\text{i.i.d.}}{\sim} P, \quad \mathbf{x} \in \mathbb{R}^m,$$

and each sample vector $\mathbf{x}^{(i)}$ is a joint sample, i.e. $\mathbf{x}^{(i)} = (\mathbf{x}_1^{(i)}, \mathbf{x}_2^{(i)})$ where $\mathbf{x}_1, \mathbf{x}_2$ are subvectors corresponding to two groups.

We define a log-linear pairwise PR model $g(\mathbf{x}; \boldsymbol{\theta})$:

$$g(\mathbf{x}; \boldsymbol{\theta}) := \frac{1}{N(\boldsymbol{\theta})} \exp\left[\sum_{u \leq v} \boldsymbol{\theta}_{u,v}^\top \boldsymbol{\psi}(\mathbf{x}_{u,v})\right],$$

where $\boldsymbol{\theta}_{u,v} \in \mathbb{R}^b$ is a column vector,

$$\boldsymbol{\theta} = (\boldsymbol{\theta}_{1,2}^\top, \dots, \boldsymbol{\theta}_{1,m}^\top, \boldsymbol{\theta}_{2,3}^\top, \dots, \boldsymbol{\theta}_{2,m}^\top, \dots, \boldsymbol{\theta}_{m-1,m}^\top)^\top,$$

and $\boldsymbol{\psi}$ is a vector valued feature function $\boldsymbol{\psi} : \mathbb{R}^2 \rightarrow \mathbb{R}^b$. Notice that we still have to model all pairwise features in \mathbf{x} , but the vast majority of these pairs are going to be nullified due to Theorem 4 if links between two groups are sparse.

$N(\boldsymbol{\theta})$ is defined as a normalization function of $g(\mathbf{x}; \boldsymbol{\theta})$:

$$N(\boldsymbol{\theta}) := \int p(\mathbf{x}_1)p(\mathbf{x}_2) \exp\left[\sum_{u \leq v} \boldsymbol{\theta}_{u,v}^\top \boldsymbol{\psi}(\mathbf{x}_{u,v})\right] d\mathbf{x}, \quad (3)$$

where $p(\mathbf{x}_1)$ and $p(\mathbf{x}_2)$ are the marginal distributions of $p(\mathbf{x})$, so it is guaranteed that $\int p(\mathbf{x}_1)p(\mathbf{x}_2)g(\mathbf{x}; \boldsymbol{\theta})d\mathbf{x} = 1$.

$N(\boldsymbol{\theta})$ in (3) can be approximated via two-sample U-statistics (Hoeffding, 1963) using the dataset,

$$N(\boldsymbol{\theta}) \approx \hat{N}(\boldsymbol{\theta}) := \frac{1}{\binom{n}{2}} \sum_{j \neq k} \exp\left[\sum_{u \leq v} \boldsymbol{\theta}_{u,v}^\top \boldsymbol{\psi}(\mathbf{x}_{u,v}^{[j,k]})\right],$$

where $\mathbf{x}^{[j,k]}$ is a *permuted sample*: $\mathbf{x}^{[j,k]} = (\mathbf{x}_1^{(j)}, \mathbf{x}_2^{(k)})$.

Notice that the normalization term $N(\boldsymbol{\theta})$ in (3) is an integral with respect to a probability distribution $p(\mathbf{x}_1)p(\mathbf{x}_2)$. Though we do not have samples directly from such a distribution, U-statistics help us “simulate” such an expectation using joint samples. In Maximum Likelihood Estimation, density models are in general hard to compute since their normalization term is not with respect to a sample distribution. In comparison, $N(\boldsymbol{\theta})$ can always be easily approximated for *any choice* of $\boldsymbol{\psi}$. This gives us the flexibility to consider complicated PR models beyond the conventional Gaussian or Ising models.

This model can be learned via the algorithm of maximum likelihood mutual information (MLMI) (Suzuki et al., 2009), by simply minimizing the Kullback-leibler divergence between $p(\mathbf{x})$ and $p_\theta(\mathbf{x}) = p(\mathbf{x}_1)p(\mathbf{x}_2)g(\mathbf{x}; \boldsymbol{\theta})$:

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} \operatorname{KL}[p||p_\theta].$$

Substitute the model of $g(\mathbf{x}; \boldsymbol{\theta})$ into the above objective and approximate $N(\boldsymbol{\theta})$ by $\hat{N}(\boldsymbol{\theta})$, then the estimated parameter $\hat{\boldsymbol{\theta}}$ is obtained as

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} - \underbrace{\sum_{i=1}^n \sum_{u \leq v} \boldsymbol{\theta}_{u,v}^\top \boldsymbol{\psi}(\mathbf{x}_{u,v}^{(i)})}_{\ell_{\text{MLMI}}(\boldsymbol{\theta})} + \log \hat{N}(\boldsymbol{\theta}) + C,$$

where C is some constant. From now on, we denote $\ell_{\text{MLMI}}(\boldsymbol{\theta})$ as the *negative* likelihood function. Due to Theorem 4 and our parametrization, if the passages between two groups are rare, then $\boldsymbol{\theta}$ is very sparse. Therefore, we may use sparsity inducing group-lasso penalties (Yuan & Lin, 2006) to encourage the sparsity on each subvector $\boldsymbol{\theta}_{u,v}$:

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} \ell_{\text{MLMI}}(\boldsymbol{\theta}) + \lambda \sum_{u \leq v} \|\boldsymbol{\theta}_{u,v}\|. \quad (4)$$

This objective is convex, unconstrained, and can be easily solved by standard sub-gradient methods. λ is a regularization parameter that can be tuned via cross-validation.

Now let us define the “true parameter” $\boldsymbol{\theta}^*$, such that $p(\mathbf{x}) = q(\mathbf{x})g(\mathbf{x}; \boldsymbol{\theta}^*)$. The learned parameter $\hat{\boldsymbol{\theta}}$ is an estimate of $\boldsymbol{\theta}^*$, where $\boldsymbol{\theta}_{u,v}^*$ is non-zero on pairwise features that are at least involved in one of the passage potentials. Moreover, as Theorem 3 and Proposition 2 show, if $X_u \in X1$ and $X_v \in X2$ are not in any of the passage structures, i.e., $\boldsymbol{\theta}_{u,v}^* = \mathbf{0}$, then $X_u \perp\!\!\!\perp X_v | \setminus \{X_u, X_v\}$.

Given the optimization problem (4), it is natural to consider the structure recovery consistency, i.e., under what conditions, the sparsity pattern of $\hat{\boldsymbol{\theta}}$ is the same as that of $\boldsymbol{\theta}^*$?

5. High-dimensional Structure Recovery Consistency

To better state the structure recovery consistency theorem, we use new indexing system with respect to the sparsity pattern of the parameter. Denoting the pairwise index set as $H = \{(u, v) | u \geq v\}$, two sets of *subvector indices* can be defined as $S = \{t' \in H \mid \|\theta_{t'}^*\| \neq 0\}$, $S^c = \{t'' \in H \mid \|\theta_{t''}^*\| = 0\}$. We rewrite the objective (4) as

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \ell(\theta) + \lambda_n \sum_{t' \in S} \|\theta_{t'}\| + \lambda_n \sum_{t'' \in S^c} \|\theta_{t''}\|. \quad (5)$$

Similarly we can define \hat{S} and \hat{S}^c . From now on, we simplify $\ell_{\text{MLMI}}(\theta^*)$ as $\ell(\theta^*)$.

Now we state our assumptions.

Assumption 1 (Dependency). *The minimum eigenvalue of the submatrix of the log-likelihood Hessian is lower-bounded:*

$$\Lambda_{\min}(\nabla_{\theta_S} \nabla_{\theta_S} \ell(\theta^*)) \geq \lambda_{\min} > 0,$$

with probability 1, where Λ_{\min} is the minimum-eigenvalue operator of a symmetric matrix

Assumption 2 (Incoherence).

$$\max_{t'' \in S^c} \left\| \left[\nabla_{\theta_{t''}} \nabla_{\theta_S} \ell(\theta^*) \right] \left[\nabla_{\theta_S} \nabla_{\theta_S} \ell(\theta^*) \right]^{-1} \right\|_1 \leq 1 - \alpha,$$

with probability 1, where $0 < \alpha \leq 1$, and $\|Y\|_1 = \sum_{i,j} \|Y_{i,j}\|_1$.

The first two assumptions are common in the literatures of support consistency. The first assumption guarantees the identifiability of the problem. The second assumption ensures the pairwise factors in passages are not too easily affected by those are not in any passages. The third assumption states the likelihood function is ‘‘well-behaved’’.

Assumption 3 (Smoothness on Likelihood Objective). *The log-likelihood ratio $\ell(\theta)$ is smooth around its optimal value, i.e., it has bounded derivatives*

$$\begin{aligned} \max_{\delta, \|\delta\| \leq \|\theta^*\|} \left\| \nabla^2 \ell(\theta^* + \delta) \right\| &\leq \lambda_{\max} < \infty, \\ \max_{t \in \{S \cup S^c\}} \max_{\delta, \|\delta\| \leq \|\theta^*\|} \left\| \nabla_{\theta_t} \nabla^2 \ell(\theta^* + \delta) \right\| &\leq \lambda_{3,\max} < \infty, \end{aligned}$$

with probability 1.

$\|\cdot\|$, $\|\cdot\|$ are the spectral norms of a matrix and a tensor respectively (See e.g., Tomioka & Suzuki (2014) for the definition of the spectral norm of a tensor).

Assumption 4 (Bounded PR Model). *For any vector $\delta \in \mathbb{R}^{\dim(\theta^*)}$ such that $\|\delta\| \leq \|\theta^*\|$, the following inequality holds:*

$$\begin{aligned} 0 < C_{\min} \leq g(\mathbf{x}; \theta) &\leq C_{\max} < \infty, \\ \|\mathbf{f}_t\|_{\infty} \leq \frac{C_{\mathbf{f}_t, \max}}{\sqrt{b}} \text{ and } \|\mathbf{f}_t\| &\leq C'_{\mathbf{f}_t, \max}, \forall t \in (S \cup S^c). \end{aligned}$$

This assumption simply indicates our PR model is bounded from above and below around the optimal value. Though it rules out the Gaussian distribution whose PR is not necessarily upper/lower-bounded, as a theory of generic pairwise models, we think it is acceptable.

Theorem 5. *Suppose that Assumptions 1, 2, 3, and 4 are satisfied as well as $\min_{t \in S} \|\theta_t^*\| \geq \frac{10}{\lambda_{\min}} \sqrt{|S|} \lambda_n$. Suppose also that the regularization parameter is chosen so that*

$$\frac{24(2 - \alpha)}{\alpha} \sqrt{\frac{M \log \frac{m^2 + m}{2}}{n}} \leq \lambda_n,$$

where M is a positive constant. Then there exist some constants L , K_1 and K_2 such that if $n \geq L|S|^2 \log \frac{m^2 + m}{2}$, with the probability at least $1 - K_1 \exp(-K_2 \lambda_n^2 n)$, MLMI in (5) has the following properties:

- *Unique Solution: The solution of (5) is unique.*
- *Successful Passage Recovery: $\hat{S} = S$ and $\hat{S}^c = S^c$.*
- $\|\hat{\theta} - \theta^*\| = \mathcal{O}\left(\sqrt{\log \frac{m^2 + m}{2}}\right)$.

The proof of Theorem 5 is detailed in Appendix D. Since the PR function is a *density ratio function* between $p(\mathbf{x})$ and $p(\mathbf{x}_1)p(\mathbf{x}_2)$, and (5) is also a sparsity inducing Kullback-Leibler Importance Estimation Procedure (KLIEP) (Sugiyama et al., 2008), the previously developed support consistency theorem (Liu et al., 2015; 2016) can be applied here as long as we can verify a few assumptions and lemmas.

The sample size required for the proposed method increases with $\log m$ (since $\log \frac{m^2 + m}{2} \leq 2 \log m$ if $m > 2$) and the estimation error on θ vanishes at the speed of $\sqrt{\frac{\log m}{n}}$. They are the same as the optimal rates obtained in previous researches for Gaussian graphical model structure learning (Ravikumar et al., 2010; Raskutti et al., 2009).

This theorem also indicates that the sample size required is not influenced by the structural density of the entire MN structure, but by the number of pairwise factors in the passage potentials. This is encouraging since we are allowed to explore PMNs with dense groups which would be hard to learn using conventional methods.

6. Experiments

Unless specified otherwise, we use pairwise feature function $\psi(x_u, x_v) = x_u x_v$. Note this does *not* mean we assume the Gaussianity over the joint distribution, since this is a parameterization of a PR rather than a joint distribution.

6.1. Synthetic Datasets

We are interested in comparing the proposed method with a few possible alternatives: **LL** (Meinshausen & Bühlmann,

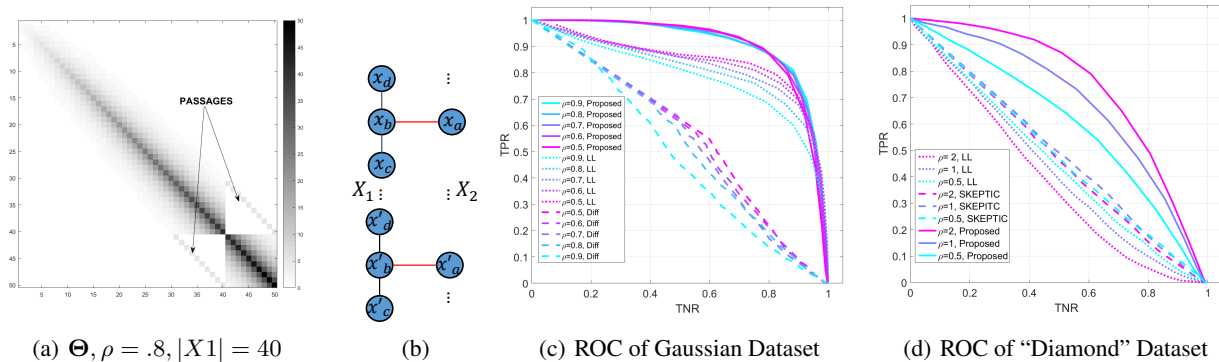


Figure 4. Synthetic experiments

2006; Ravikumar et al., 2010), **SKEPTIC** (Liu et al., 2012) and **Diff** (Zhao et al., 2014): A *direct* difference estimation method that learns the differences between two MNs without learning each individual precision matrix separately. In this paper, we employed this method to learn the differences between two Gaussian densities: $p(\mathbf{x})$ and $p(\mathbf{x}_1)p(\mathbf{x}_2)$.

We first generate a set of joint samples $\{\mathbf{x}^{(i)}\}_{i=1}^{50} \sim \mathcal{N}(\mathbf{0}, \Theta^{-1})$, where $\Theta \in \mathbb{R}^{50 \times 50}$ and is constructed in two steps. First, create

$$\Theta_{i,j} = \begin{cases} \rho^{|i-j|} \sqrt{ij}, & i, j < 40 \text{ or } i, j > 40, \\ 0, & \text{Otherwise,} \end{cases}$$

where $0 < \rho < 1$ is a coefficient controlling the dominance of the diagonal entries. Second, let Λ be the 15th smallest eigenvalue of Θ , and fill the submatrices $\Theta_{\{41, \dots, 50\}, \{31, \dots, 40\}}$ and $\Theta_{\{31, \dots, 40\}, \{41, \dots, 50\}}$ with $\Lambda \mathbf{I}_{10}$, where \mathbf{I}_{10} is a 10×10 identity matrix. By such a construction, we have created two groups over X : $X = (X_{\{1, \dots, 40\}}, X_{\{41, \dots, 50\}})$ and 10 passages between them. Notably, within two groups, the precision matrix is *dense*, and random variables interact with each other via powerful links when ρ is large. An example of Θ when $\rho = 0.8$ is plotted in Figure 4(a). We measure the performance of three methods using the True Postive Rate (TPR) and True Negative Rate (TNR). The detailed definition of TPR and TNR is deferred to Appendix, E.

The ROC curve in Figure 4(c) can be plotted by adjusting the sensitivity of each method: Tuning the regularization parameter of the proposed method and LL, or the threshold parameter of Diff.

As we can see, the proposed method has the best overall performance on all ρ choices, comparing to both LL and Diff. Also, as the links within each group get more and more powerful (by increasing ρ), the performance of LL and Diff decay significantly, while the proposed method almost remain unchanged.

As the proposed method is capable of handling complex models, we draw 50 samples from a 52-dimensional ‘‘diamond’’ distribution used in (Liu et al., 2014) where the correlation among random variables are non-linear. To speed-up the sampling procedure, the graphical model of this distribution is constructed by concatenating 13 simple 4-variable MNs whose density functions are defined as

$$p(x_a, x_b, x_c, x_d) \propto \exp(-\rho x_a^2 x_b^2 - .5x_b x_c - .5x_b x_d) \cdot \mathcal{N},$$

where \mathcal{N} is short for a normal density $\mathcal{N}(\mathbf{0}, .5\mathbf{I}_4)$ over x_a, x_b, x_c and x_d . Notice this distribution does not have a closed form normalization term. The graphical model of such a distribution is illustrated in Figure 4(b). In this experiment, the coefficient ρ is used to control the strength of inter-group interactions ($x_a \leftrightarrow x_b$), and we set $\psi(x_u, x_v) = x_u^2 x_v^2$. Other than LL, we include SKEPTIC due to the non-Gaussian nature of this dataset. The performance is compared in Figure 4(d) using ROC curves.

The correlation among random variables are completely non-linear. As the power of interactions on passages increases, LL performs worse and worse since it still relies on the Gaussian model assumption. Thanks to the correct PR model, the proposed method performs reasonably well and gets better when ρ increases. As the density model does not fit into the Gaussian copula model, SKEPTIC also performs poorly.

6.2. Bipartisanship in 109th US Senate

We use the proposed method to study the bipartisanship between Democrats and Republicans in the 109th US Senate via the recorded votes. There were totally 100 senators (45 Democrats and 55 Republicans) casting votes on 645 questions with ‘‘yea’’, ‘‘nay’’ or ‘‘not voting’’. The task is to discover the cross-party links between senators. We construct a dataset $\{\mathbf{x}^{(i)}\}_{i=1}^{645} \sim X$ using all 645 questions as observations, where each observation $\mathbf{x} \in \{1, -1, 0\}^{100}$ corresponds to the votes on a single question by 100 senators, and random variables $X = (X_{\{1, \dots, 45\}}, X_{\{46, \dots, 100\}})$




Figure 5. Bipartisanship in 109th US Senate. Prefix “(D)” or “(R)” indicates the party membership of a senator. Red: positive influence, Blue: negative influence. Edge widths are proportional to $|\theta_{u,v}|$.

are senators partitioned according to party memberships.

We run the proposed method directly on this dataset, and decrease λ from 10 until $|\hat{S}| > 15$. To avoid complication, we only plot edges that contain nodes from different groups in Figure 5.

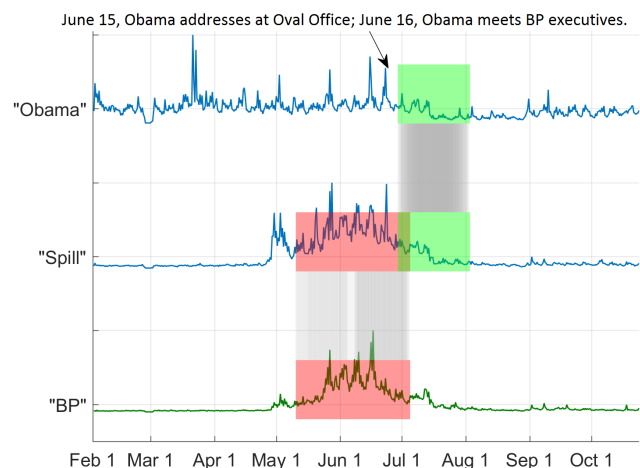
It can be seen that Ben Nelson, a conservative Democrat, who “frequently voting against his party” (Wikipedia, 2016a), has multiple links with the other side. On the right, Democrat Tom Carper tends to agree with Republican Lincoln Chafee. Carper collaborated with Chafee on multiple bipartisan proposals (Press-Release, a;b) while Chafee, who “support for fiscal and social policies that often opposed those promoted by the Republican Party” (Wikipedia, 2016b) finally switched his affiliation to Democratic in 2013. Interestingly, we have also observed a cluster of senators who tend to disagree with each other.

6.3. Pairwise Sequences Alignment


PMN can also be used to “align” sequences. Given a pair of sequences where points are collected from the domain \mathcal{X} , we pick sequence 1 and construct the dataset by sliding a window sized n toward future, until reaching the end. Suppose there are m_1 windows generated, then we can create a dataset $\{\mathbf{x}_1^{(i)}\}_{i=1}^n, \mathbf{x} \in \mathcal{X}^{m_1}$. Similarly, we construct another dataset $\{\mathbf{x}_2^{(i)}\}_{i=1}^n, \mathbf{x} \in \mathcal{X}^{m_2}$ on sequence 2, and make joint samples by letting $\mathbf{x}^{(i)} = (\mathbf{x}_1^{(i)}, \mathbf{x}_2^{(i)})$. After learning a PMN over two groups, if X_u and X_v are connected, then we regard the elements in the u -th window and the elements in the v -th window are “aligned”. See Figure 7 in Appendix for an illustration.

We run the proposed method to learn PMNs over two datasets: Twitter keyword count sequences (Liu et al., 2013) and Amino acid sequences with Genebank ID: AAD01939 and AAQ67266. The results were obtained by decreasing λ from 10 so $|\hat{S}| > 15$.

For the Twitter dataset, we collect normalized frequencies of keywords as time-series over 8 months, during the



(a) Twitter keyword frequency time-series alignments. $n = 50, m = 962$ and $\mathcal{X} = \mathbb{R}$.



(b) Amino acid sequence alignments between AAD01939 (human) and AAQ67266 (fly). $n = 10, m = 592, \phi(x_i, x_j) = \delta(x_i, x_j)$ and $\mathcal{X} = \{\text{amino acid dictionary}\}$.

Figure 6. Sequence alignment. For two aligned windows with size n , we plot n gray lines between two windows linking each pair of elements. Since lines are so close to each other, they look like “gray shades” on the plot. The color box contains the region of consecutively aligned windows.

event “Deepwater Horizon oil spill” in 2010. We learn alignments between two pairs of keywords: “Obama” vs. “Spill” and “Spill” vs. “BP”. The results are plotted in Figure 6(a) where we can see the sequences of two pairs are aligned well in chronological order. The two popular keywords, “BP” and “Spill” are synchronized throughout almost the entire event while “Spill” and “Obama” are only synchronized later on after he delivered his speech in Oval Office on this crisis on June 15th, 2010.

The next experiment uses two amino acid string sequences, consisting codes such as ‘V’, ‘I’, ‘L’ and ‘F’, etc. Figure 6(b) shows that the proposed method has successfully identified the aligned segment between eyeless gene of *Drosophila melanogaster* (a fruitfly) and human aniridia genes. The same segment is also spotted by widely used Needleman-Wunsch (NW) algorithm (Needleman & Wunsch, 1970) with statistical significance.

Acknowledgments

SL acknowledges JSPS Grant-in-Aid for Scientific Research Activity Start-up 15H06823. TS acknowledges JST-PRESTO. MS acknowledges the JST CREST program. KF acknowledges JSPS Grant-in-Aid for Scientific Research on Innovative Areas 25120012. Authors would like to thank four anonymous reviewers for their helpful comments.

References

- Banerjee, O., El Ghaoui, L., and d’Aspremont, A. Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *Journal of Machine Learning Research*, 9:485–516, March 2008.
- Friedman, J., Hastie, T., and Tibshirani, R. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.
- Hammersley, J. M. and Clifford, P. Markov fields on finite graphs and lattices. 1971.
- Hoeffding, W. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
- Koller, D. and Friedman, N. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge, MA, USA, 2009.
- Liu, H., Lafferty, J., and Wasserman, L. The nonparametric: Semiparametric estimation of high dimensional undirected graphs. *Journal of Machine Learning Research*, 10:2295–2328, 2009.
- Liu, H., Han, F., Yuan, M., Lafferty, J., and Wasserman, L. High-dimensional semiparametric Gaussian copula graphical models. *The Annals of Statistics*, pp. 2293–2326, 2012.
- Liu, S., Yamada, M., Collier, N., and Sugiyama, M. Change-point detection in time-series data by relative density-ratio estimation. *Neural Networks*, 43:72–83, 2013.
- Liu, S., Quinn, J. A., Gutmann, M. U., Suzuki, T., and Sugiyama, M. Direct learning of sparse changes in Markov networks by density ratio estimation. *Neural Computation*, 26(6):1169–1197, 2014.
- Liu, S., Suzuki, T., and Sugiyama, M. Support consistency of direct sparse-change learning in Markov networks. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI2015)*, 2015.
- Liu, S., Relator, R., Sese, J., Suzuki, T., and Sugiyama, M. Support consistency of direct sparse-change learning in Markov networks. *Annals of Statistics*, 2016. to appear.
- Loh, P.-L. and Wainwright, M. J. Structure estimation for discrete graphical models: Generalized covariance matrices and their inverses. In Pereira, F., Burges, C.J.C., Bottou, L., and Weinberger, K.Q. (eds.), *Advances in Neural Information Processing Systems 25*, pp. 2087–2095. 2012.
- Meinshausen, N. and Bühlmann, P. High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, 34(3):1436–1462, 06 2006.
- Murphy, K. P. *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012.
- Needleman, S. B. and Wunsch, C. D. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *Journal of molecular biology*, 48(3):443–453, 1970.
- Press-Release. Carper urges bipartisan compromise on clean air, a. URL <http://www.epw.senate.gov/pressitem.cfm?party=rep&id=230919>.
- Press-Release. Carper-chafee-feinstein will offer bipartisan budget plan similar to “blue dog” house proposal, b. URL <http://www.carper.senate.gov/public/index.cfm/pressreleases?ID=e5603ed2-80ae-483b-a98d-4aa9932edeaf>.
- Raskutti, G., Yu, B., Wainwright, M. J., and Ravikumar, P. Model selection in Gaussian graphical models: High-dimensional consistency of ℓ_1 -regularized mle. In Koller, D., Schuurmans, D., Bengio, Y., and Bottou, L. (eds.), *Advances in Neural Information Processing Systems 21*, pp. 1329–1336. Curran Associates, Inc., 2009.
- Ravikumar, P., Wainwright, M. J., and Lafferty, J. D. High-dimensional Ising model selection using ℓ_1 -regularized logistic regression. *The Annals of Statistics*, 38(3):1287–1319, 2010.
- Robert, C. P. and Casella, G. *Monte Carlo Statistical Methods*. Springer-Verlag, Secaucus, NJ, USA, 2005.
- Steinwart, I. and Christmann, A. *Support vector machines*. Springer Science & Business Media, 2008.
- Sugiyama, M., Suzuki, T., Nakajima, S., Kashima, H., von Bünau, P., and Kawanabe, M. Direct importance estimation for covariate shift adaptation. *Annals of the Institute of Statistical Mathematics*, 60(4):699–746, 2008.
- Suzuki, T., Sugiyama, M., and Tanaka, T. Mutual information approximation via maximum likelihood estimation of density ratio. In *Information Theory, 2009. ISIT*

2009. *IEEE International Symposium on*, pp. 463–467, June 2009.

Tibshirani, R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.

Tomioka, R. and Suzuki, T. Spectral norm of random tensors. *arXiv preprint arXiv:1407.1870 [math.ST]*, 2014.

Vapnik, V. N. *Statistical Learning Theory*. Wiley, New York, NY, USA, 1998.

Wainwright, M. J. Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (lasso). *IEEE Trans. Inf. Theor.*, 55(5):2183–2202, May 2009.

Wainwright, M. J. and Jordan, M. I. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1-2):1–305, 2008.

Wikipedia. Ben Nelson — Wikipedia, the free encyclopedia, 2016a. URL https://en.wikipedia.org/wiki/Ben_Nelson. [Online; accessed 30-Jan-2016].

Wikipedia. Lincoln Chafee — Wikipedia, the free encyclopedia, 2016b. URL https://en.wikipedia.org/wiki/Lincoln_Chafee. [Online; accessed 30-Jan-2016].

Yuan, M. and Lin, Y. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.

Zhao, P. and Yu, B. On model selection consistency of lasso. *The Journal of Machine Learning Research*, 7: 2541–2563, 2006.

Zhao, S., Cai, T., and Li, H. Direct estimation of differential networks. *Biometrika*, 101(2):253–268, 2014.