

SUPPORT CONSISTENCY OF DIRECT SPARSE-CHANGE LEARNING IN MARKOV NETWORKS¹

BY SONG LIU*, TAIJI SUZUKI^{†,‡}, RAISSA RELATOR[§], JUN SESE[§],
MASASHI SUGIYAMA[¶] AND KENJI FUKUMIZU*

Institute of Statistical Mathematics, Tokyo Institute of Technology,[†]
PRESTO, Japan Science and Technology Agency,[‡] National Institute of
Advanced Industrial Science and Technology[§] and University of Tokyo[¶]*

We study the problem of learning sparse structure changes between two Markov networks P and Q . Rather than fitting two Markov networks separately to two sets of data and figuring out their differences, a recent work proposed to learn changes *directly* via estimating the ratio between two Markov network models. In this paper, we give sufficient conditions for *successful change detection* with respect to the sample size n_p, n_q , the dimension of data m and the number of changed edges d . When using an unbounded density ratio model, we prove that the true sparse changes can be consistently identified for $n_p = \Omega(d^2 \log \frac{m^2+m}{2})$ and $n_q = \Omega(n_p^2)$, with an exponentially decaying upper-bound on learning error. Such sample complexity can be improved to $\min(n_p, n_q) = \Omega(d^2 \log \frac{m^2+m}{2})$ when the boundedness of the density ratio model is assumed. Our theoretical guarantee can be applied to a wide range of discrete/continuous Markov networks.

1. Introduction. Learning changes in interactions between random variables play an important role in many real-world applications. For example, genes may regulate each other in different ways when external conditions are changed. The number of daily flu-like symptom reports in nearby hospitals may become correlated when a major epidemic disease breaks out. EEG signals from different regions of the brain may be synchronized/desynchronized when the patient is performing different activities. Identifying such changes in interactions helps us expand our knowledge on these real-world phenomena.

In this paper, we consider the problem of learning changes between two undirected graphical models. Such a model, also known as a Markov network (MN) [7], expresses interactions via the conditional independence between random variables. The Hammersley–Clifford theorem [5] states that the joint distribution of an MN can be factorized over subsets of interacted random variables and general

Received December 2015; revised April 2016.

¹Supported in part by KAKENHI (15H06823, 25730013, 25120012, 26280009, 15H01678, 15H05707, 15H01717, 15H05713, 24240044 and 25120012), JST-PRESTO No. JPMJPR14E4, and JST-CREST No. JPMJCR1304 and JPMJCR14D7.

MSC2010 subject classifications. Primary 62F12; Secondary 68T99.

Key words and phrases. Markov networks, change detection, density ratio estimation.

MNs may have factors over arbitrary numbers of random variables. For simplicity, we focus on a special case, namely *pairwise MNs*, whose joint distribution can be factorized over only single or pairwise random variables.

The problem of learning structure of MN itself has been thoroughly investigated in the last decade. The graphical lasso method [1, 4] learns a sparse precision (inverse covariance) matrix from data by using the ℓ_1 -norm, while the neighborhood regression methods [8, 11, 15] solve a nodewise lasso program to identify the neighborhood of each single node.

One naive approach to learning changes in MNs is to apply these methods to two MNs separately and compare the learned models. However, such a two-step approach does not work well when the MNs themselves are dense (this can happen even when the change in MNs is sparse). A recent study [28] adopts a neighborhood selection procedure to learn sparse changes between Gaussian MNs via a fused-lasso type regularizer [20]. However, no theoretical guarantee was given on identifying changes. Furthermore, extension of the above mentioned methods to general non-Gaussian MNs is hard due to the computational intractability of the normalization term.

To cope with these problems, an novel algorithm has been proposed recently [9]. Its basic idea is to model the changes between two MNs P and Q as the ratio between two MN density functions $p(\mathbf{x})$ and $q(\mathbf{x})$, and the ratio $p(\mathbf{x})/q(\mathbf{x})$ is directly estimated in one-shot without estimating $p(\mathbf{x})$ and $q(\mathbf{x})$ themselves [19]. Since parameters in the density ratio model represent the parametric difference between P and Q , sparsity constraints can be directly imposed for sparse change learning. Thus, the density-ratio approach can work well even when each MN is dense as long as the change is sparse. Furthermore, the normalization term in the density-ratio approach can be approximately computed by the straightforward sample average, and thus there is no computational bottleneck in using non-Gaussian MNs. Experimentally, the density-ratio approach was demonstrated to perform well. However, its theoretical properties have not been explored yet.

The ability of recovering a sparsity pattern via a sparse learning algorithm has been studied under the name of *support consistency* or sparsistency [23], that is, the support of the estimated parameter converges to the true support. Previous works for *successful structure recovery* are available for ℓ_1 -regularized maximum (pseudo-)likelihood estimators [15, 26]. However, density ratio estimator in [9] brought us a new question: What is the sparsistency of identifying correct sparse changes *without* learning individual MNs? Such a concern is very practical since in applications such as learning changes in gene expression between stimuli conditions, we only care about changes rather than individual structures before or after changes. We illustrate such an example in the problem of gene regulatory networks in Section 7.

In this paper, we theoretically investigate the success of the density-ratio approach and provide sufficient conditions for *successful change detection* with respect to the number of samples n_p, n_q , data dimension m and the number of

changed edges d . More specifically, we prove that if $n_p = \Omega(d^2 \log \frac{m^2+m}{2})$ and $n_q = \Omega(n_p^2)$, changes between two MNs can be consistently learned under mild assumptions, regardless the sparsity of individual MNs. Such sample complexity can be further improved to $\min(n_p, n_q) = \Omega(d^2 \log \frac{m^2+m}{2})$ when the boundedness of the density ratio model is assumed. Technically, our contribution can be regarded as an extension of support consistency of lasso-type programs [23] to the *ratio* of MNs. The convergence rate does not rely on the individual sparsity of each MN, thus structures like hub-nodes can exist. Such hub structure is common in many applications, such as gene expression data where one gene regulates many other genes. Our theorem holds for the most general log-linear MN models, and does *not* assume any special type of individual MNs (such as Gaussian or Ising).

Note that the theoretical results presented in this paper are fundamentally different from previous works on learning a “jumping MN” [6], where the focuses are learning the partition boundaries between jumps, and the successful recovery of graphical structure within each partition, rather than learning sparse changes between partitions.

In previous works [15, 26], the (upper/lower) boundedness of the Fisher information matrix, or log-partition function derivatives of a density model, are often assumed. In this work, similar assumptions are imposed on the true density ratio model. Moreover, we show that such assumptions have profound links with the smoothness of our model, which implies the magnitude of change should not be too drastic for keeping the density ratio model well behaved. These assumptions are also automatically satisfied under some special cases.

The target of [9] coincides with another recently proposed method where a *differential network* is learned directly using a different technique [29]. Without learning a precision matrix for each MN, this approach estimates a differential network utilizing a special equality obtained for Gaussian MNs. However, such an objective function does not generalize to ordinary pairwise MNs. The theorems obtained in this paper and the ones in [29] both rely on one similar assumption: The changes are sparse. However, theorems in this paper manage to achieve the same sample complexity of recovering the correct structure changes without explicitly assuming Gaussianity over datasets.

This paper is organized as follows: First, we introduce the problem formulation of learning changes between two MNs in Section 2. Second, we review the density ratio estimation method proposed in [9]. Then, as the main focus of this paper, we analyze the sufficient conditions for successful change detection, that is, the support consistency of such algorithm in Sections 3 and 4. Moreover, in Section 5, we study the key assumptions in this paper, and discuss their consequences. Through experiments in Section 6, we demonstrate the validity of our theorems and compare the performance of the density ratio approach with a state of the art method. Finally, in Section 7, we show the density ratio method successfully identifies key changes in a gene network between two stimuli conditions.

2. Direct change learning between Markov networks. In this section, we review a direct structural change detection method [9].

2.1. *Problem formulation.* Consider two sets of independent samples drawn separately from two probability distributions P and Q on \mathbb{R}^m :

$$\{\mathbf{x}_p^{(i)}\}_{i=1}^{n_p} \stackrel{\text{i.i.d.}}{\sim} P \quad \text{and} \quad \{\mathbf{x}_q^{(i)}\}_{i=1}^{n_q} \stackrel{\text{i.i.d.}}{\sim} Q.$$

We assume that P and Q belong to the family of *Markov networks* (MNs) consisting of univariate and bivariate factors, that is, their respective probability densities p and q are expressed as

$$(1) \quad p(\mathbf{x}; \boldsymbol{\theta}^{(p)}) = \frac{1}{Z(\boldsymbol{\theta}^{(p)})} \exp\left(\sum_{u,v=1, u \geq v}^m \boldsymbol{\theta}_{u,v}^{(p)\top} \boldsymbol{\psi}_{u,v}(x_u, x_v)\right),$$

where $\mathbf{x} = (x_1, \dots, x_m)^\top$ is the m -dimensional random variable, \top denotes the transpose, $\boldsymbol{\theta}_{u,v}^{(p)}$ is the b -dimensional parameter vector for the elements x_u and x_v , and

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

is the entire parameter vector. $\boldsymbol{\psi}_{u,v}(x_u, x_v)$ is a bivariate vector-valued basis function, and $Z(\boldsymbol{\theta}^{(p)})$ is the normalization factor defined as

$$Z(\boldsymbol{\theta}^{(p)}) = \int \exp\left(\sum_{u,v=1, u \geq v}^m \boldsymbol{\theta}_{u,v}^{(p)\top} \boldsymbol{\psi}_{u,v}(x_u, x_v)\right) d\mathbf{x}.$$

$q(\mathbf{x}; \boldsymbol{\theta}^{(q)})$ is defined in the same way. Using this notation, we can define two well-known MNs as examples:

Ising-model (see, e.g., [7]). One of the earliest and widely known graphical models is the Ising model, where $\boldsymbol{\psi}_{u,v}(x_u, x_v) = x_u x_v$, and $x_u, x_v \in \{-1, 1\}$. For all pairs $(u, v) \in E$, where E is the edge set of the graphical model, $\theta_{u,v}$ is a scalar and has a nonzero value.

Gaussian MN. Gaussian MN is a representative of continuous MN. $\boldsymbol{\psi}_{u,v}(x_u, x_v) = x_u x_v$ and $x_u, x_v \in \mathbb{R}$. For all pairs $(u, v) \in E$ or $u = v$, $\theta_{u,v}$ is a scalar and has a nonzero value.

The research problem now becomes clear: Given two parametric models $p(\mathbf{x}; \boldsymbol{\theta}^{(p)})$ and $q(\mathbf{x}; \boldsymbol{\theta}^{(q)})$, we hope to discover *changes in parameters* from P to Q , that is, $\boldsymbol{\theta}^{(p)} - \boldsymbol{\theta}^{(q)}$.

2.2. *Density ratio formulation for structural change detection.* The key idea in [9] is to consider the ratio of p and q :

$$\frac{p(\mathbf{x}; \boldsymbol{\theta}^{(p)})}{q(\mathbf{x}; \boldsymbol{\theta}^{(q)})} \propto \exp\left(\sum_{u,v=1, u \geq v}^m (\boldsymbol{\theta}_{u,v}^{(p)} - \boldsymbol{\theta}_{u,v}^{(q)})^\top \boldsymbol{\psi}_{u,v}(x_u, x_v)\right),$$

where $\theta_{u,v}^{(p)} - \theta_{u,v}^{(q)}$ encodes the difference between P and Q for factor $\psi_{u,v}(x_u, x_v)$, that is, $\theta_{u,v}^{(p)} - \theta_{u,v}^{(q)}$ is zero if there is no change in the factor $\psi_{u,v}(x_u, x_v)$.

Once the ratio of p and q is considered, each parameter $\theta_{u,v}^{(p)}$ and $\theta_{u,v}^{(q)}$ does not have to be estimated, but only their difference $\theta_{u,v} = \theta_{u,v}^{(p)} - \theta_{u,v}^{(q)}$ is sufficient to be estimated for change detection. Thus, in this density-ratio formulation, p and q are no longer modeled separately, but it models the changes from p to q directly as

$$(2) \quad r(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{N(\boldsymbol{\theta})} \exp\left(\sum_{u,v=1, u \geq v}^m \boldsymbol{\theta}_{u,v}^\top \boldsymbol{\psi}_{u,v}(x_u, x_v)\right),$$

where $N(\boldsymbol{\theta})$ is the normalization term. This direct formulation also halves the number of parameters from both $\boldsymbol{\theta}^{(p)}$ and $\boldsymbol{\theta}^{(q)}$ to only $\boldsymbol{\theta}$.

The normalization term $N(\boldsymbol{\theta})$ is chosen to fulfill $\int q(\mathbf{x})r(\mathbf{x}; \boldsymbol{\theta}) d\mathbf{x} = 1$:

$$N(\boldsymbol{\theta}) = \int q(\mathbf{x}) \exp\left(\sum_{u,v=1, u \geq v}^m \boldsymbol{\theta}_{u,v}^\top \boldsymbol{\psi}_{u,v}(x_u, x_v)\right) d\mathbf{x},$$

which is the expectation over $q(\mathbf{x})$. This expectation form of the normalization term is another notable advantage of the density-ratio formulation because it can be easily approximated by the sample average over $\{\mathbf{x}_q^{(i)}\}_{i=1}^{n_q} \stackrel{\text{i.i.d.}}{\sim} q(\mathbf{x})$:

$$\hat{N}(\boldsymbol{\theta}; \mathbf{x}_q^{(1)}, \dots, \mathbf{x}_q^{(n_q)}) := \frac{1}{n_q} \sum_{i=1}^{n_q} \exp\left(\sum_{u,v=1, u \geq v}^m \boldsymbol{\theta}_{u,v}^\top \boldsymbol{\psi}_{u,v}(x_{q,u}^{(i)}, x_{q,v}^{(i)})\right).$$

Thus, one can always use this empirical normalization term for any (non-Gaussian) models $p(\mathbf{x}; \boldsymbol{\theta}^{(p)})$ and $q(\mathbf{x}; \boldsymbol{\theta}^{(q)})$.

An important observation can be made from this formulation: Although two MNs may have sophisticated models individually, their changes might be “simple” since many terms may be canceled while taking the ratio, that is, $\theta_{u,v}^{(p)} - \theta_{u,v}^{(q)}$ might be zero. Thus, if we use $\psi_{u,v}(x_u, x_v) = x_u x_v$ in our ratio model, it does not mean we assume two individual MNs are Gaussian or Ising. It simply means we assume the changes of interactions are linear while other nonlinear interactions remain unchanged. This formulation allows us to consider highly complicated MNs as long as their changes are “simple.” We will give a concrete example later.

Throughout the rest of the paper, we simplify the notation from $\psi_{u,v}$ to $\boldsymbol{\psi}$ by assuming the feature functions are the same for all pairs of random variables. However, our analysis still holds if this assumption is violated. Next, we study the density ratio formulation in the case of Gaussian MNs.

Gaussian MN. Given two m -dimensional zero-mean Gaussian MNs $p(\mathbf{x}; \boldsymbol{\Theta}^{(p)})$ and $q(\mathbf{x}; \boldsymbol{\Theta}^{(q)})$ parameterized by the precision matrix $\boldsymbol{\Theta}^{(p)}$ and $\boldsymbol{\Theta}^{(q)}$, respectively, it is reasonable to parametrize a density ratio model

$$(3) \quad r(\mathbf{x}; \boldsymbol{\Delta}) = \frac{1}{N(\boldsymbol{\Delta})} \exp\left(-\frac{1}{2} \mathbf{x}^\top \boldsymbol{\Delta} \mathbf{x}\right),$$

where Δ is a symmetric real-valued matrix and

$$N(\Delta) = \int q(\mathbf{x}; \Theta^{(q)}) \exp\left(-\frac{1}{2} \mathbf{x}^\top \Delta \mathbf{x}\right) d\mathbf{x} \\ = \frac{\det(\Theta^{(q)})^{1/2}}{\det(\Delta + \Theta^{(q)})^{1/2}}.$$

However, this formulation brings a problem: It still contains an unknown parameter $\Theta^{(q)}$, meaning that we will have to learn $\Theta^{(q)}$ first before we can model the difference between two MNs. To solve this problem, one may use the empirical version of the normalization term instead:

$$(4) \quad \hat{N}(\Delta) = \frac{1}{n_q} \sum_{i=1}^{n_q} \exp\left(-\frac{1}{2} \mathbf{x}_q^{(i)\top} \Delta \mathbf{x}_q^{(i)}\right).$$

Interestingly, by using this model, the Gaussianity assumption has been loosened: the normalization term is obtained by an empirical average and did not use the analytical form offered by the Gaussianity of $q(\mathbf{x}; \Theta^{(q)})$. Thus, it can actually model the density ratio for any p and q as long as their changes are limited to the quadratic components.

2.3. *Direct density-ratio estimation.* Density ratio estimation has been recently introduced to the machine learning community and is proven to be useful in a wide range of applications [19]. In [9], a density ratio estimator called the *Kullback–Leibler importance estimation procedure* (KLIEP) for log-linear models [18, 22] was employed in learning structural changes.

For a density ratio model $r(\mathbf{x}; \theta)$, the KLIEP method minimizes the Kullback–Leibler divergence from $p(\mathbf{x})$ to $\hat{p}(\mathbf{x}; \theta) = q(\mathbf{x})r(\mathbf{x}; \theta)$:

$$(5) \quad \text{KL}[p \parallel \hat{p}_\theta] = \int p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})r(\mathbf{x}; \theta)} d\mathbf{x} \\ = \text{Const.} - \int p(\mathbf{x}) \log r(\mathbf{x}; \theta) d\mathbf{x}.$$

Note that the density-ratio model (2) automatically satisfies the nonnegativity and normalization constraints:

$$r(\mathbf{x}; \theta) \geq 0 \quad \text{and} \quad \int q(\mathbf{x})r(\mathbf{x}; \theta) d\mathbf{x} = 1.$$

Here, we define

$$\hat{r}(\mathbf{x}; \theta) = \frac{\exp(\sum_{u,v=1, u \geq v}^m \theta_{u,v}^\top \psi(x_{q,u}, x_{q,v}))}{\hat{N}(\theta; \mathbf{x}_q^{(1)}, \dots, \mathbf{x}_q^{(n_q)})}$$

as the *empirical density ratio model*. In practice, one minimizes the negative empirical approximation of the rightmost term in equation (5):

$$\begin{aligned} \ell_{\text{KLIEP}}(\boldsymbol{\theta}) &= -\frac{1}{n_p} \sum_{i=1}^{n_p} \log \hat{r}(\mathbf{x}_p^{(i)}; \boldsymbol{\theta}) \\ &= -\frac{1}{n_p} \sum_{i=1}^{n_p} \sum_{u,v=1, u \geq v}^m \boldsymbol{\theta}_{u,v}^\top \boldsymbol{\psi}(x_{p,u}^{(i)}, x_{p,v}^{(i)}) \\ &\quad + \log \left(\frac{1}{n_q} \sum_{i=1}^{n_q} \exp \left(\sum_{u,v=1, u \geq v}^m \boldsymbol{\theta}_{u,v}^\top \boldsymbol{\psi}(x_{q,u}^{(i)}, x_{q,v}^{(i)}) \right) \right). \end{aligned}$$

Because $\ell_{\text{KLIEP}}(\boldsymbol{\theta})$ is convex with respect to $\boldsymbol{\theta}$, its global minimizer can be numerically found by standard optimization techniques such as gradient descent or quasi-Newton methods. The gradient of ℓ_{KLIEP} with respect to $\boldsymbol{\theta}_{u,v}$ is given by

$$\begin{aligned} \nabla_{\boldsymbol{\theta}_{u,v}} \ell_{\text{KLIEP}}(\boldsymbol{\theta}) &= -\frac{1}{n_p} \sum_{i=1}^{n_p} \boldsymbol{\psi}(x_{p,u}^{(i)}, x_{p,v}^{(i)}) \\ (6) \quad &\quad + \frac{1}{n_q} \sum_{i=1}^{n_q} \hat{r}(\mathbf{x}^{(i)}; \boldsymbol{\theta}) \boldsymbol{\psi}(x_{q,u'}^{(i)}, x_{q,v'}^{(i)}), \end{aligned}$$

that can be computed in a straightforward manner for *any* feature vector $\boldsymbol{\psi}(x_u, x_v)$.

Importance sampling. From the gradient of KLIEP (6), we can observe a clear link between KLIEP and Importance Sampling (see, e.g., [16]). The second term on the right-hand side is an “importance sampled” approximation of $\mathbb{E}_p[\boldsymbol{\psi}(x_u, x_v)]$ using our density ratio model while the first term is a straightforward sample average. The population version of (6) equals zero if and only if $p(\mathbf{x}) = \hat{r}(\mathbf{x}; \boldsymbol{\theta})q(\mathbf{x})$. Therefore, one should aware that the assignment of p and q may affect the performance of such an approximation as importance sampling can be easily affected by the choice of the instrumental distribution (in this case, q). To reduce the estimation variance, q is usually picked as the one with a thicker tail [24]. This observation reveals a fundamental asymmetry of KLIEP which will be discussed in Section 8.

Gaussian MN. By using the density ratio model of (3) and the normalization term (4), we can write the objective function and its gradient as

$$\begin{aligned} \ell_{\text{KLIEP}}(\boldsymbol{\Delta}) &= \frac{1}{2n_p} \sum_{i=1}^{n_p} \mathbf{x}_p^{(i)\top} \boldsymbol{\Delta} \mathbf{x}_p^{(i)} \\ (7) \quad &\quad + \log \frac{1}{n_q} \sum_{i=1}^{n_q} \exp \left(-\frac{1}{2} \mathbf{x}_q^{(i)\top} \boldsymbol{\Delta} \mathbf{x}_q^{(i)} \right) \end{aligned}$$

and

$$\begin{aligned} &\nabla_{\Delta_{u,v}} \ell_{\text{KLIEP}}(\mathbf{\Delta}) \\ &= \frac{1}{2n_p} \sum_{i=1}^{n_q} x_{p,u}^{(i)} \Delta_{u,v} x_{p,v}^{(i)} \\ &\quad - \frac{\frac{1}{n_q} \sum_{i=1}^{n_q} \exp(-\frac{1}{2} \mathbf{x}_q^{(i)\top} \mathbf{\Delta} \mathbf{x}_q^{(i)}) x_{q,u}^{(i)} \Delta_{u,v} x_{q,v}^{(i)}}{\frac{2}{n_q} \sum_{j=1}^{n_q} \exp(-\frac{1}{2} \mathbf{x}_q^{(j)\top} \mathbf{\Delta} \mathbf{x}_q^{(j)})}. \end{aligned}$$

2.4. *Sparsity-inducing norm.* To find a sparse change between P and Q , one may regularize the KLIEP solution with a sparsity-inducing norm $\sum_{u \geq v} \|\theta_{u,v}\|$, that is, the *group-lasso* penalty [27] where we use $\|\cdot\|$ to denote the ℓ_2 norm. Note that the separate density estimation approaches sparsify both θ_p and θ_q so that the difference $\theta_p - \theta_q$ is also sparsified. On the other hand, the density-ratio approach [9] directly sparsifies the difference $\theta_p - \theta_q$, and thus intuitively this method can still work well even if θ_p and θ_q are dense as long as $\theta_p - \theta_q$ is sparse.

Now we have reached our final objective:

$$(8) \quad \hat{\theta} = \underset{\theta}{\operatorname{argmin}} \ell_{\text{KLIEP}}(\theta) + \lambda_{n_p} \sum_{u,v=1, u \geq v}^m \|\theta_{u,v}\|.$$

3. Support consistency of direct sparse-change detection. The above density-ratio approach to change detection was demonstrated to be promising in empirical studies [9]. However, its theoretical properties have not yet been investigated. In this section, we give theoretical guarantees of the convex program (8) on sparse structural change learning. More specifically, we give *sufficient conditions* for detecting correct changes in terms of the sample size n_p and n_q , data dimensions m and the number of changed edges d , followed by the discussion of the insights we can gain from such theoretical analysis.

3.1. *Notation.* In the previous section, a sub-vector of θ indexed by (u, v) corresponds to a specific edge of an MN. From now on, we use new indices with respect to the ‘‘oracle’’ sparsity pattern of the true parameter for notational simplicity. We introduce the ‘‘true parameter’’ notation θ^* , $p(\mathbf{x}) = q(\mathbf{x})r(\mathbf{x}; \theta^*)$, and the pairwise index set $E = \{(u, v) | u \geq v\}$. Two sets of *sub-vector indices* regarding to θ^* and E are defined as $S = \{t' \in E | \|\theta_{t'}^*\| \neq 0\}$, $S^c = \{t'' \in E | \|\theta_{t''}^*\| = 0\}$. We rewrite the objective (8) as

$$(9) \quad \hat{\theta} = \underset{\theta}{\operatorname{argmin}} \ell_{\text{KLIEP}}(\theta) + \lambda_{n_p} \sum_{t' \in S} \|\theta_{t'}\| + \lambda_{n_p} \sum_{t'' \in S^c} \|\theta_{t''}\|.$$

Similarly, we can define $\hat{S} = \{t' \in E \mid \|\hat{\theta}_{t'}\| \neq 0\}$ and \hat{S}^c accordingly. Sample Fisher information matrix $\mathcal{I} \in \mathbb{R}^{\frac{b(m^2+m)}{2} \times \frac{b(m^2+m)}{2}}$ denotes the Hessian of the log-likelihood: $\mathcal{I} = \nabla^2 \ell_{\text{KLIEP}}(\theta^*) = \nabla^2 \log \hat{N}(\theta^*)$ where we simplify $\hat{N}(\theta; \mathbf{x}_q^{(1)}, \dots, \mathbf{x}_q^{(n_q)})$ as $\hat{N}(\theta)$. \mathcal{I}_{AB} is a sub-matrix of \mathcal{I} indexed by two sets of indices $A, B \subseteq E$ on rows and columns.

We also concatenate $\psi \in \mathbb{R}^2 \mapsto \mathbb{R}^b$ to get a “linearized” version of the feature function $f : \mathbb{R}^m \mapsto \mathbb{R}^{\frac{b(m^2+m)}{2}}$ as

$$f(\mathbf{x}) = (\psi^\top(x_1, x_1), \dots, \psi^\top(x_m, x_1), \psi^\top(x_2, x_2), \dots, \psi^\top(x_m, x_2), \dots, \psi^\top(x_m, x_m))^\top$$

and $f_A(\mathbf{x})$ is the partial output of $f(\mathbf{x})$ indexed by a set of indices $A, A \subseteq E$.

Gaussian MN. Here, we derive the Fisher information matrix for the Gaussian MN ratio model. Define an auxiliary matrix $H(\Delta) \in \mathbb{R}^{n_q \times n_q}$:

$$H(\Delta) := \frac{1}{\hat{N}^2(\Delta)} (\hat{N}(\Delta) I_{n_q} - \mathbf{e}^\top(\Delta) \mathbf{e}(\Delta)),$$

where $\mathbf{e}(\Delta) := [\exp(-\frac{1}{2} \mathbf{x}_q^{(1)\top} \Delta \mathbf{x}_q^{(1)}), \dots, \exp(-\frac{1}{2} \mathbf{x}_q^{(n_q)\top} \Delta \mathbf{x}_q^{(n_q)})]$. The Fisher information matrix $\mathcal{I} = \nabla_{\Delta}^2 \ell_{\text{KLIEP}}(\Delta^*)$ of the likelihood function using the Gaussian density ratio model described in (7) has the form

$$\mathcal{I}_{(u,v),(u',v')} := \sum_{i=1}^{n_q} \sum_{j=1}^{n_q} x_{q,u}^{(i)} x_{q,v}^{(i)} x_{q,u'}^{(j)} x_{q,v'}^{(j)} H_{i,j}(\Delta^*), \quad \mathcal{I} \in \mathbb{R}^{m^2 \times m^2}.$$

3.2. Assumptions. There is an important guideline for imposing assumptions in this paper: We try *not* to put any explicit constraints on the types of individual MN P or Q nor their structures, but only on the changes between them. This is crucial since KLIEP is a *direct* and *flexible* change learning method and have no restrictions on the types of individual MNs on which it works. Therefore, we hope to obtain the most generic theorem for this method.

Similarly to previous researches on sparsity recovery analysis [15, 23], the first two assumptions are made on the Fisher information matrix.

ASSUMPTION 1 (Dependency assumption). The sample Fisher information submatrix \mathcal{I}_{SS} has bounded eigenvalues:

$$\Lambda_{\min}(\mathcal{I}_{SS}) \geq \lambda_{\min} > 0,$$

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

This assumption on the *submatrix* of \mathcal{I} is to ensure that the model is identifiable (see B.1 in the supplementary article [10] for details). Note “ λ ” denotes either eigenvalue or regularization parameter depending on its subscript.

ASSUMPTION 2 (Incoherence assumption).

$$\max_{t'' \in S^c} \|\mathcal{I}_{t''S} \mathcal{I}_{SS}^{-1}\|_1 \leq 1 - \alpha, \quad 0 < \alpha \leq 1$$

with probability 1, where $\|Y\|_1 = \sum_{i,j} \|Y_{i,j}\|_1$.

This assumption says the unchanged edges cannot exert overly strong effects on changed edges and is a common assumption can be found in previous literatures on support consistency analysis such as [15, 23].

ASSUMPTION 3 (Smoothness assumption on likelihood ratio). The log-likelihood ratio $\ell_{\text{KLIEP}}(\boldsymbol{\theta})$ is smooth around its optimal value, that is, it has bounded derivatives

$$\begin{aligned} (10) \quad & \max_{\boldsymbol{\delta}, \|\boldsymbol{\delta}\| \leq \|\boldsymbol{\theta}^*\|} \|\nabla^2 \ell_{\text{KLIEP}}(\boldsymbol{\theta}^* + \boldsymbol{\delta})\| \\ & = \max_{\boldsymbol{\delta}, \|\boldsymbol{\delta}\| \leq \|\boldsymbol{\theta}^*\|} \|\nabla^2 \log \hat{N}(\boldsymbol{\theta}^* + \boldsymbol{\delta})\| \leq \lambda_{\max} < \infty, \end{aligned}$$

$$\begin{aligned} (11) \quad & \max_{t \in S \cup S^c} \max_{\boldsymbol{\delta}, \|\boldsymbol{\delta}\| \leq \|\boldsymbol{\theta}^*\|} \|\|\nabla_{\boldsymbol{\theta}_t} \nabla^2 \ell_{\text{KLIEP}}(\boldsymbol{\theta}^* + \boldsymbol{\delta})\|\| \\ & = \max_{t \in S \cup S^c} \max_{\boldsymbol{\delta}, \|\boldsymbol{\delta}\| \leq \|\boldsymbol{\theta}^*\|} \|\|\nabla_{\boldsymbol{\theta}_t} \nabla^2 \log \hat{N}(\boldsymbol{\theta}^* + \boldsymbol{\delta})\|\| \leq \lambda_{3,\max} < \infty, \end{aligned}$$

with probability 1.

$\|\cdot\|, \|\|\cdot\|\|$ are the spectral norms of a matrix and a tensor, respectively (see e.g., [21] for the definition of spectral norm of a tensor). Note that (10) also implies the bounded largest eigenvalue of \mathcal{I} . Assumption 3 can be regarded as an analogy of assumptions on the log-normalization function in [26]. As we set no explicit restrictions on the type of distribution P and Q , this assumption guarantees the log-likelihood function is well behaved.

Now, we make the following assumptions on the density ratio.

ASSUMPTION 4 (The correct model assumption). The density ratio model is correct, that is, there exists $\boldsymbol{\theta}^*$ such that

$$p(\mathbf{x}) = r(\mathbf{x}; \boldsymbol{\theta}^*)q(\mathbf{x}).$$

Assumptions 1, 2 and 3 are in fact related to distribution Q . However, the density ratio estimation objective is an M-estimator summed up over samples from P .

Assumption 4 provides a transform between P and Q and allows us to perform analysis on such an M-estimator in an “importance sampling” fashion.

Next, we impose assumptions on the “smoothness” of the density ratio model. Generally speaking, if we expect good performance from the density ratio estimator, the density ratio model should be “well behaved.” The following assumption quantifies such an intuition.

ASSUMPTION 5 (Smooth density ratio model assumption). For any vector $\delta \in \mathbb{R}^{\dim(\theta^*)}$ such that $\|\delta\| \leq \|\theta^*\|$ and every $a \in \mathbb{R}$, the following inequality holds:

$$\mathbb{E}_q[\exp(a(r(\mathbf{x}, \theta^* + \delta) - 1))] \leq \exp(10a^2).$$

We list a few consequences of the Assumption 5.

PROPOSITION 1. For all $\varepsilon > 0$ and for any vector $\delta \in \mathbb{R}^{\dim(\theta^*)}$ such that $\|\delta\| \leq \|\theta^*\|$, $P(r(\mathbf{x}, \theta^* + \delta) - 1 \geq \varepsilon) \leq 2 \exp(-\frac{\varepsilon^2}{40})$.

Using Assumption 5, we get Proposition 1 that provides a tail probability bound of the density ratio model on Q , which is further used to obtain an exponentially decaying upper-bound of empirical approximation error of the log-normalization term (see Proposition 1 in the supplementary article [10] for details).

PROPOSITION 2. For any vector $\delta \in \mathbb{R}^{\dim(\theta^*)}$ such that $\|\delta\| \leq \|\theta^*\|$, $\text{Var}_q[r(\mathbf{x}; \theta^* + \delta) - 1] \leq 20$.

PROOF. Noting $\mathbb{E}_q[r(\mathbf{x}; \theta^* + \delta) - 1] = 0$, the above inequalities is the consequence of sub-Gaussianity. \square

Since the density ratio can be thought as the magnitude of change between two MNs, Proposition 2 tells the fact that the change should not be too drastic in order to keep our ratio-model well behaved.

We are now ready to state the main theorem.

3.3. *Sufficient conditions for successful change detection.* The following theorem establishes sufficient conditions of change detection in terms of parameter sparsity. Its proof is provided in Section 4.1. First, let us define $g(m) = \frac{\log(m^2+m)}{(\log \frac{m^2+m}{2})^2}$ (see Figure 3 in the supplementary article [10] for its plot) which is smaller than 1 when $m > 4$.

THEOREM 1. Suppose that Assumptions 1, 2, 3, 4 and 5 as well as $\min_{t \in S} \|\theta_t^*\| \geq \frac{10}{\lambda_{\min}} \sqrt{d} \lambda_{n_p}$ are satisfied, where d is the number of changed edges

defined as $d = |S|$, that is, the cardinality of the set of nonzero parameter groups. Suppose also that the regularization parameter is chosen so that

$$(12) \quad \frac{8(2 - \alpha)}{\alpha} \sqrt{\frac{M_1 \log \frac{m^2+m}{2}}{n_p}} \leq \lambda_{n_p} \leq \frac{4(2 - \alpha)M_1}{\alpha} \min\left(\frac{\|\theta^*\|}{\sqrt{b}}, 1\right),$$

where $M_1 = \lambda_{\max}b + 2$, $n_q \geq M_2 n_p^2 g(m)$ and M_2 is a positive constant. Then there exist some constants L_1, K_1 and K_2 such that if $n_p \geq L_1 d^2 \log \frac{m^2+m}{2}$, with the probability at least

$$(13) \quad 1 - \exp(-K_1 \lambda_{n_p}^2 n_p) - 4 \exp(-K_2 d n_q \lambda_{n_p}^4),$$

the following properties hold:

- *Unique solution:* The solution of (9) is unique.
- *Successful change detection:* $\hat{S} = S$ and $\hat{S}^c = S^c$, where \hat{S} and \hat{S}^c are estimated sparse/nonsparse indices.

First, it is interesting to analyze the sample complexity of n_q , which is a novel element in this research. Intuitively, one should obtain a sufficient number of samples from Q to accurately approximate the normalization term. Theorem 1 states n_q should grow at least quadratically with respect to n_p , which is undesirable if n_p is large. In the next corollary, we discuss a relaxed coupling between n_p and n_q with some extra but mild cost.

Second, Assumption 5 together with Proposition 2 shows the variation allowed for the density ratio model is irrelevant to the number of changed edges d . This implies that, if d is large, we are only able to detect weak changes that do not cause huge fluctuations in the density ratio model, which is rather restrictive and unrealistic in some occasions, since the magnitude of change usually increases when the number of changed edges d increases. Below, we consider another more relaxed scenario, where the assumption on the smoothness of the density ratio model is allowed to grow with d .

ASSUMPTION 6. For any vector $\delta \in \mathbb{R}^{\dim(\theta^*)}$ such that $\|\delta\| \leq \|\theta^*\|$ and every $a \in \mathbb{R}$, the following inequality holds:

$$\mathbb{E}_q[\exp(a(r(\mathbf{x}, \theta^* + \delta) - 1))] \leq \exp(10da^2),$$

where d is the number of changed edges.

PROPOSITION 3. For some small constants ε and any vector $\delta \in \mathbb{R}^{\dim(\theta^*)}$ such that $\|\delta\| \leq \|\theta^*\|$, then $P(r(\mathbf{x}, \theta^* + \delta) - 1 \geq \varepsilon) \leq 2 \exp(-\frac{\varepsilon^2}{40d})$.

PROPOSITION 4. For any vector $\delta \in \mathbb{R}^{\dim(\theta^*)}$ such that $\|\delta\| \leq \|\theta^*\|$, $\text{Var}_q[r(\mathbf{x}; \theta^* + \delta) - 1] \leq 20d$.

From Proposition 4, we can see the magnitude of changes between MNs are allowed to grow at most linearly with d . Now we see how much this will bring changes to our sufficient conditions:

COROLLARY 1. *Suppose that Assumptions 1, 2, 3, 4 and 6 are satisfied, $\min_{t \in S} \|\theta_t^*\|$ satisfies the condition in Theorem 1, and the regularization parameter is chosen so that*

$$\frac{2 - \alpha}{\alpha} \sqrt{\frac{M_1 \log \frac{m^2+m}{2}}{n_p^{\frac{3}{4}}}} \leq \lambda_{n_p} \leq \frac{4(2 - \alpha)M_1}{\alpha} \min\left(\frac{\|\theta^*\|}{\sqrt{b}}, \frac{1}{n_p^{1/8}}\right),$$

where $M_1 = \lambda_{\max} b + 2, n_q \geq M_2 d n_p g(m)$, and M_2 is some positive constant. Then there exist some constant L_1 such that if $n_p \geq L_1 d^{\frac{8}{3}} (\log \frac{m^2+m}{2})^{\frac{4}{3}}$, KLIEP has the same properties as those stated in Theorem 1.

See Section B.6 in the supplementary article [10] for the proof. Corollary 1 states that it is possible to drop the growth rate of n_q on n_p from 2 to 1 with the cost that n_p has to grow with $d^{\frac{8}{3}}$ (rather than just d^2 in the previous case). This is an encouraging result, since with slight changes on growth rate with respect to d and $\log(\frac{m^2+m}{2})$, we are able to consider a milder coupling between n_p and n_q .

Moreover, under the weaker Assumption 6, n_q now grows linearly with d . It shows the prices we need to pay when consider the magnitude of changes increasing with d .

So far, we have only considered the scaling quadruple (n_p, n_q, d, m) . However, it is also interesting to consider that the scalability of our theorem relative to b , the dimension of the pairwise feature vector. This is a realistic scenario: It may be difficult to know the true underlying model of MN in practice, and thus we may adopt a model that contains many features to be “flexible enough” to describe the interactions among data. In the following corollary, we restate Theorem 1 with b and a new scalar s , which is the maximum number of nonzero elements in a pairwise feature vector ψ . We assume that the positions of nonzero elements are independent of each sample x .

COROLLARY 2. *Suppose that Assumptions 1, 2, 3, 4 and 5 are satisfied, $\min_{t \in S} \|\theta_t^*\|$ satisfies the condition in Theorem 1, and the regularization parameter is chosen so that*

$$\frac{8(2 - \alpha)}{\alpha} \sqrt{\frac{M_1 s \log \frac{m^2+m}{2}}{n_p}} \leq \lambda_{n_p} \leq \frac{4(2 - \alpha)M_1}{\alpha} \min\left(\frac{\|\theta^*\|}{\sqrt{b}}, 1\right),$$

where $M_1 = \lambda_{\max} b + 2$, $n_q \geq M_2 s n_p^2 g'(m)$ and M_2 is some positive constant, and $g'(m) = \frac{\log((m^2+m)\binom{b}{s})}{(\log \frac{m^2+m}{2})^2}$. Then there exists some constant L_1 such that if $n_p \geq L_1 s d^2 \log \frac{m^2+m}{2}$, KLIEP has the same properties as those stated in Theorem 1.

See Section B.6 in the supplementary article [10] for the proof. From Corollary 2, we can see that required n_p and n_q for change detection grows only linearly with respect to s , and n_q grows mildly with respect to $\binom{b}{s}$. Therefore, it is possible for one to consider a highly flexible model in practice.

3.4. Discussions. From the above theorem, one may gather some interesting insights into change detection based on density ratio estimation.

First, the required number of samples depends solely on d and m and is irrelevant to the number of edges of each MN. In contrast, separate graphical structural learning methods require more samples when each MN gets denser in terms of the number of edges or neighborhood [11, 14, 15]. This establishes the superiority of the density-ratio approach in sparse change detection between dense MNs. In other words, in order to detect sparse changes, the density-ratio approach does not require the individual MN to be sparse.

Second, the growth of n_q is also lower-bounded and grows quadratically with respect to n_p . This result illustrates the consequence of introducing a sample approximated normalization term $\hat{N}(\theta)$. An insufficient number of samples from Q would lead to poor approximation of the normalization term, and makes change detection more difficult. Fortunately, such growth rate can be further relaxed, and with slightly increased sample complexity of n_p .

Finally, our theorem also points out the limits of the density-ratio approach. Our analysis shows that the density ratio model may not deviate too much from its mean 1 over distribution Q . A previous study on another density ratio estimator also has a similar observation [25]. Since the density ratio indicates how much P differs from Q , this analysis generally implies that to make KLIEP work, the discrepancy between P and Q should be mild. This is a reasonable assumption since we have already assumed that the changes in the MN structure are sparse. In a high-dimensional setting, it implies P and Q are similar. Similar assumption can be found in [29] where the ℓ -1 norm of the differences between two precision matrices is bounded.

4. Proof of support consistency.

4.1. The proof outline of the main theorem. The procedure of the main proof partially follows the steps of previous support consistency proofs using the *primal-dual witness* method [23]; however, the problem settings are quite different: First,

ℓ_{KLIEP} is a *likelihood ratio* between two densities which means that two sets of samples are involved in this proof and we have to consider the sparsity recovery conditions not only on one dataset, but with respect to two different MNs. Second, we did not explicitly limit the types of distribution for P and Q , and the parameter of each factor $\theta_t, t \in S \cup S^c$ is a vector rather than a scalar, which gives enough freedom for modeling highly complicated distributions. To the best of our knowledge, this is the first sparsity recovery analysis on learning changes between two type-free MNs. From now on, ℓ_{KLIEP} is shortened as ℓ .

First, define a dual variable \hat{z} associated with $\hat{\theta}$ using the following equality:

$$(14) \quad \nabla \ell(\hat{\theta}) + \lambda_{n_p} \hat{z} = \mathbf{0}$$

and if \hat{z}_t is the subgradient of $\|\hat{\theta}_t\|$, that is, $\hat{z}_t \in \nabla_{\theta_t} \|\hat{\theta}_t\|, t \in S \cup S^c$, (14) is the optimality condition of (9) and $\hat{\theta}$ is an optimal solution to (9). Moreover, the next lemma tells the relationship between dual variable \hat{z} and sparsity patterns of any other optimal solution of (9).

LEMMA 1. *If there exists an optimal $\hat{\theta}$ of (9) with associated \hat{z} in (14) such that $\|\hat{z}_{t''}\| < 1$, for all $t'' \in S^c$. Then any optimal $\tilde{\theta}$ of (9) should have $\tilde{\theta}_{t''} = \mathbf{0}$ for all $t'' \in S^c$.*

See Section B.1 in the supplementary article [10] for the proof.

Now we illustrate the proof procedure of Theorem 1:

- Solve the constrained optimization problem

$$(15) \quad \hat{\theta}_S = \underset{\theta_S}{\operatorname{argmin}} \ell \left(\begin{bmatrix} \theta_S \\ \mathbf{0} \end{bmatrix} \right) + \lambda_{n_p} \sum_{t' \in S} \|\theta_{t'}\|;$$

- For all $t' \in S$, set $\hat{z}_{t'} = \nabla \|\hat{\theta}_{t'}\|$, and let $\hat{\theta} = [\hat{\theta}_S, \mathbf{0}]$;
- Obtain $\hat{z}_{t''}$ for all $t'' \in S^c$ using equality (14);
- Show $\max_{t'' \in S^c} \|\hat{z}_{t''}\| < 1$ with high probability under certain conditions. According to Lemma 1, we conclude that for any optimal $\tilde{\theta}$ from (9), the correct sparsity pattern is recovered.

Bounding $\max_{t'' \in S^c} \|\hat{z}_{t''}\|$ requires obtaining $z_{t''}$ from (14). More specifically, from (14) we have

$$\nabla \ell(\hat{\theta}) + \lambda_{n_p} \hat{z} = \mathbf{0} \quad \Rightarrow \quad \nabla \ell(\hat{\theta}) + \lambda_{n_p} \hat{z} - \nabla \ell(\theta^*) = -\nabla \ell(\theta^*).$$

Applying the mean-value theorem,

$$(16) \quad \underbrace{\nabla^2 \ell(\theta^*)}_{\mathcal{I}} [\hat{\theta} - \theta^*]^\top + \lambda_{n_p} \hat{z} = \underbrace{-\nabla \ell(\theta^*)}_{\mathbf{w}} + \underbrace{[\nabla^2 \ell(\theta^*) - \nabla^2 \ell(\bar{\theta})]}_{\mathbf{g}} [\hat{\theta} - \theta^*]^\top,$$

where $\bar{\theta}$ is between θ^* and $\hat{\theta}$ in a coordinate fashion. We can then rewrite (16) in blockwise fashion:

$$(17) \quad \begin{aligned} \mathcal{I}_{S,S}[\hat{\theta}_S - \theta_S^*] + \lambda_{n_p} \hat{z}_S &= \mathbf{w}_S + \mathbf{g}_S, \\ \mathcal{I}_{t'',S}[\hat{\theta}_S - \theta_S^*] + \lambda_{n_p} \hat{z}_{t''} &= \mathbf{w}_{t''} + \mathbf{g}_{t''}, \quad t'' \in S^c. \end{aligned}$$

Substitute $\hat{\theta}_S - \theta_S^* = \mathcal{I}_{SS}^{-1}[\mathbf{w}_S + \mathbf{g}_S - \lambda_{n_p} \hat{z}_S]$ into (17), we have

$$\mathcal{I}_{t''S} \mathcal{I}_{SS}^{-1}[\mathbf{w}_S + \mathbf{g}_S - \lambda_{n_p} \hat{z}_S] + \lambda_{n_p} \hat{z}_{t''} = \mathbf{w}_{t''} + \mathbf{g}_{t''}.$$

Rearranging terms, we have

$$\lambda_{n_p} \hat{z}_{t''} = \mathbf{w}_{t''} + \mathbf{g}_{t''} - \mathcal{I}_{t''S} \mathcal{I}_{SS}^{-1}[\mathbf{w}_S + \mathbf{g}_S - \lambda_{n_p} \hat{z}_S].$$

According to the triangle inequality,

$$\begin{aligned} \lambda_{n_p} \max_{t'' \in S^c} \|\hat{z}_{t''}\| &\leq \max_{t'' \in S^c} \|\mathbf{w}_{t''}\| + \max_{t'' \in S^c} \|\mathbf{g}_{t''}\| \\ &\quad + \max_{t'' \in S^c} \|\mathcal{I}_{t''S} \mathcal{I}_{SS}^{-1}\|_1 \left(\max_{t' \in S} \|\mathbf{w}_{t'}\| + \max_{t' \in S} \|\mathbf{g}_{t'}\| + \lambda_{n_p} \right). \end{aligned}$$

By assumption, $\max_{t'' \in S^c} \|\mathcal{I}_{t''S} \mathcal{I}_{SS}^{-1}\|_1 \leq (1 - \alpha)$, and we obtain

$$\max_{t'' \in S^c} \|\hat{z}_{t''}\| \leq \frac{(2 - \alpha)}{\lambda_{n_p}} \left(\max_{t \in \text{SUS}^c} \|\mathbf{w}_t\| + \max_{t \in \text{SUS}^c} \|\mathbf{g}_t\| \right) + (1 - \alpha).$$

Now we need to show the boundedness of \mathbf{w} and \mathbf{g} .

The boundedness of \mathbf{w} , which is the gradient of log-likelihood function on θ^* , is guaranteed by the following lemma.

LEMMA 2. *There exist constants $c = \lambda_{\max} b + 2$, c' and c'' and if the regularization parameter λ_{n_p} satisfies*

$$\frac{8(2 - \alpha)}{\alpha} \sqrt{\frac{c \log \frac{m^2+m}{2}}{n_p}} \leq \lambda_{n_p} \leq \frac{4(2 - \alpha)c}{\alpha} \min\left(\frac{\|\theta^*\|}{\sqrt{b}}, 1\right),$$

then

$$(18) \quad \begin{aligned} P\left(\max_{t \in S^c \cup S} \|\mathbf{w}_t\| \geq \frac{\alpha}{4(2 - \alpha)} \lambda_{n_p}\right) \\ \leq \exp(-c' n_p) + 4 \exp\left(-c'' n_q \left(\frac{\log \frac{m^2+m}{2}}{n_p}\right)^2 + b \log(m^2 + m)\right). \end{aligned}$$

Quite different from similar lemmas in previous works (such as Lemma 2 in [15]), Lemma 2 is not a simple concentration of sample mean converging to its population mean, since the gradient of the likelihood contains two sets of data

from different distributions P and Q . Moreover, the sub-Gaussianity or the boundedness of the ψ of P or Q were not assumed, so the concentration inequality cannot be applied here. Instead, via the smoothness behavior of the likelihood ratio function (Assumption 3), we are able to derive such a boundedness of \mathbf{w} without using any explicit properties of two distributions.

The consequence of such differences is important: this analysis allows us to consider a very wide range of distributions which may not be well behaved (e.g., heavy-tailed), as long as the change between two distributions are minor. After all, all assumptions are imposed on the density ratio model r only, rather than P or Q . This analysis preserves the flexibility of the density ratio estimation methodology. The proof of Lemma 2 can be found in Section B.2 of the supplementary article [10].

The next lemma bounds the difference between the estimated parameter and the true parameter over the nonsparse indices, which is further used to bound \mathbf{g} and derive the sample complexity.

LEMMA 3. *If $d\lambda_{n_p} \leq \frac{\lambda_{\min}^2}{20\lambda_{3,\max}}$ and $\max_{t \in S^c \cup S} \|\mathbf{w}_t\| \leq \frac{\lambda_{n_p}}{4}$, then $\|\boldsymbol{\theta}_S^* - \hat{\boldsymbol{\theta}}_S\| \leq \frac{10}{\lambda_{\min}} \sqrt{d}\lambda_{n_p}$.*

The boundedness of \mathbf{g} is finally given by the following.

LEMMA 4. *If $\lambda_{n_p} d \leq \frac{\lambda_{\min}^2}{100\lambda_{3,\max}} \frac{\alpha}{4(2-\alpha)}$, and $\max_{t \in S \cup S^c} \|\mathbf{w}_t\| \leq \frac{\lambda_{n_p}}{4}$, then $\max_{t \in S \cup S^c} \|\mathbf{g}_t\| \leq \frac{\alpha\lambda_{n_p}}{4(2-\alpha)}$.*

See Sections B.3 and B.4 in the supplementary article [10] for proofs. Using Lemmas 2, 3 and 4, we have $\max_{t'' \in S^c} \|\hat{\mathbf{z}}_{t''}\| \leq 1 - \frac{\alpha}{2} < 1$.

To show the correct nonzero pattern recovery, it suffices to show

$$\max_{t \in S^c \cup S} \|\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}_t^*\| < \frac{1}{2} \min_{t \in S} \|\boldsymbol{\theta}_t^*\|.$$

Since Lemma 3 shows $\max_{t \in S \cup S^c} \|\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}_t^*\| \leq \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\| < \frac{10}{\lambda_{\min}} \sqrt{d}\lambda_{n_p}$, we just need $\min_{t \in S} \|\boldsymbol{\theta}_t^*\| > \frac{20}{\lambda_{\min}} \sqrt{d}\lambda_{n_p}$ to ensure such recovery.

4.2. *Sample complexity.* The sample complexity for n_p and n_q are derived from the conditions of the lemmas. To make Lemma 2 hold, we may set $\lambda_{n_p} = C \sqrt{\frac{\log \frac{m^2+m}{2}}{n_p}}$, where C is chosen so that the lower bound of λ_{n_p} in Lemma 2 is satisfied. Since the upper bound is a constant while such setting of λ_{n_p} is always decaying as n_p grows, it is automatically satisfied at some point.

Moreover, λ_{n_p} should also satisfy the upper-bound condition in Lemma 4: $\lambda_{n_p} d \leq \frac{\lambda_{\min}^2}{100\lambda_{3,\max}} \frac{\alpha}{4(2-\alpha)}$, and this inequality can be satisfied when $n_p = \Omega(d^2 \log \frac{m^2+m}{2})$.

The upper bound of λ_{n_p} in Lemma 4 is tighter than it is in Lemma 3, so the condition of Lemma 3 is automatically satisfied. However, one still needs to make sure that the tail probability term that involves n_q in Lemma 2 decays, that is,

$$4 \exp\left(-c'' n_q \left(\frac{\log \frac{m^2+m}{2}}{n_p}\right)^2 + b \log(m^2 + m)\right) \rightarrow 0, \quad n_p \rightarrow \infty.$$

This can be guaranteed by setting $n_q = \Omega(n_p^2 g(m))$.

5. Analysis of assumptions. In this section, we investigate the conditions under which the maximum (minimum) eigenvalues of likelihood ratio derivatives are bounded. We show that under mild regularity conditions of sample statistics of distribution Q and tightened smoothness conditions of the density ratio model, Assumptions 1 and 3 holds automatically.

5.1. Bounded density ratio model. Since the derivatives of the log-likelihood ratio expresses the curvature of our objective function, we expect the smoothness of the density ratio model $r(\mathbf{x}; \boldsymbol{\theta})$ may play an important role in such analysis. To begin with, consider a simple bounded-ratio model by replacing the smoothness Assumption 5 (or 6) with a tightened Assumption 7.

ASSUMPTION 7 (Smooth density ratio model assumption). For any vector $\boldsymbol{\delta} \in \mathbb{R}^{\dim(\boldsymbol{\theta}^*)}$ such that $\|\boldsymbol{\delta}\| \leq \|\boldsymbol{\theta}^*\|$, the following inequality holds:

$$0 < C_{\min} \leq r(\mathbf{x}, \boldsymbol{\theta}^* + \boldsymbol{\delta}) \leq C_{\max} < \infty.$$

As consequences, $\frac{1}{C_{\text{ratio}}} \leq \hat{r}(\mathbf{x}; \boldsymbol{\theta}^* + \boldsymbol{\delta}) \leq C_{\text{ratio}}$ and $\|\mathbf{f}_t(\mathbf{x})\| \leq C_{f_t, \max}$, where $C_{f_t, \max}$ and C_{ratio} are all constants.

Since Assumption 7 is stronger than Assumption 5 or 6 (for appropriately chosen C_{\max} and C_{\min}), the proof of Theorem 1 still holds if one uses the above assumptions to substitute Assumptions 5 or 6. However, as we will demonstrate later, an improved sample complexity for $\min(n_p, n_q)$ can be derived.

As Assumptions 1 and 3 are constructed using samples from distribution Q , it is also natural to assume some basic regularity conditions on sample statistics.

ASSUMPTION 8 (Bounded moments). The feature transform $\mathbf{f}(\mathbf{x}) \in \mathbb{R}^{b(m^2+m)/2}$, where \mathbf{x} is drawn from Q , has upper-bounded moments with proba-

bility one, that is,

$$\begin{aligned} \max_{t \in \text{SUS}^c} \widehat{\mathbb{E}}_q[\|\mathbf{f}_t(\mathbf{x})\|] &\leq D_{\max,1} < \infty, \\ \|\widehat{\mathbb{E}}_q[\mathbf{f}(\mathbf{x})\mathbf{f}(\mathbf{x})^\top]\| \quad \text{and} \quad \|\widehat{\text{Cov}}_q[\mathbf{f}(\mathbf{x})]\| &\leq D_{\max,2} < \infty \end{aligned}$$

and

$$(19) \quad \Lambda_{\min}\{\widehat{\text{Cov}}_q[\mathbf{f}_S(\mathbf{x})]\} \geq D_{\min,2} > 0,$$

is bounded with probability $1 - \delta_{n_q}$.

$\widehat{\text{Cov}}$ is the sample covariance estimator and $\widehat{\mathbb{E}}_q[\mathbf{g}(\mathbf{x})] = \frac{1}{n_q} \sum_{i=1}^{n_q} \mathbf{g}(\mathbf{x}^{(i)})$ is the empirical expectation over samples drawn from Q . From now on, we remove subscript p or q from a random sample \mathbf{x} when summing up, as long as the indices give enough context for telling in which distribution the sample is drawn. For example, $\sum_{i=1}^{n_q} \mathbf{x}^{(i)}$ is a summation over samples drawn from distribution Q .

Of course, one may impose similar bounded moments constraints on the corresponding population quantities, then the above assumption automatically holds with high probability under certain regularity conditions. To avoid lengthy proofs, we stick to assumptions using sample quantities in this paper.

The following propositions show Assumptions 7 and 8 guarantee the boundedness of the derivatives of the likelihood ratio function.

PROPOSITION 5 (Bounded Hessian). *For any vector $\boldsymbol{\delta} \in \mathbb{R}^{\dim(\boldsymbol{\theta}^*)}$ such that $\|\boldsymbol{\delta}\| \leq \|\boldsymbol{\theta}^*\|$, if Assumptions 7 and 8 hold then $\|\nabla_{\boldsymbol{\theta}}^2 \ell(\boldsymbol{\theta}^* + \boldsymbol{\delta})\| \leq 2C_{\text{ratio}} D_{\max,2}$ with probability $1 - \delta_{n_q}$.*

PROPOSITION 6 (Bounded third-order derivative). *For any vector $\boldsymbol{\delta} \in \mathbb{R}^{\dim(\boldsymbol{\theta}^*)}$ such that $\|\boldsymbol{\delta}\| \leq \|\boldsymbol{\theta}^*\|$, if Assumptions 7 and 8 hold, then*

$$\max_{t \in \text{SUS}^c} \|\|\nabla_{\boldsymbol{\theta}_t} \nabla^2 \ell(\boldsymbol{\theta}^* + \boldsymbol{\delta})\|\| \leq 6C_{\text{ratio}}^2 D_{\max,1} D_{\max,2}$$

with probability $1 - \delta_{n_q}$.

PROPOSITION 7 (Eigenvalue lower bound of invertible Hessian submatrix). *For any vector $\boldsymbol{\delta} \in \mathbb{R}^{\dim(\boldsymbol{\theta}^*)}$ such that $\|\boldsymbol{\delta}\| \leq \|\boldsymbol{\theta}^*\|$, if Assumptions 7 and 8 hold, then*

$$\Lambda_{\min}[\nabla_{\boldsymbol{\theta}_S}^2 \ell(\boldsymbol{\theta}^* + \boldsymbol{\delta})] \geq \frac{D_{\min,2}}{C_{\text{ratio}}^2}$$

with probability $1 - \delta_{n_q}$.

Proofs of the above propositions are listed in Sections B.7, B.8 and B.9 in the supplementary article [10].

5.2. *Sufficient conditions under bounded density ratio.* Now, we give a variation of Theorem 1 based on the totally bounded density ratio model. Consider the objective

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \ell(\theta) + \lambda_{n_p, n_q} \sum_{t' \in S} \|\theta_{t'}\| + \lambda_{n_p, n_q} \sum_{t'' \in S^c} \|\theta_{t''}\|,$$

which is identical to (9) but the regularization parameter is now determined with respect to both n_p and n_q .

COROLLARY 3. *Suppose that Assumptions 2, 4, 7 and 8 as well as $\min_{t \in S} \|\theta_t^*\| \geq \frac{10}{\lambda_{\min}} \sqrt{d} \lambda_{n_p, n_q}$ are satisfied. Suppose also that the regularization parameter is chosen so that*

$$\frac{24(2 - \alpha)}{\alpha} \sqrt{\frac{M \log \frac{m^2 + m}{2}}{\min(n_p, n_q)}} \leq \lambda_{n_p, n_q},$$

where M is a positive constant. Then there exist some constants L and K such that if $\min(n_p, n_q) \geq Ld^2 \log \frac{m^2 + m}{2}$, with the probability at least $1 - 4 \exp(-K \lambda_{n_p, n_q}^2 \min(n_p, n_q)) - \delta_{n_q}$, KLIEP has the same properties as those stated in Theorem 1.

The proof of this corollary is done by replacing Lemma 2 with Lemma 5.

LEMMA 5. *If $\lambda_{n_p, n_q} \geq \frac{24(2 - \alpha)}{\alpha} \cdot \sqrt{\frac{c \log \frac{(m^2 + m)}{2}}{\min(n_p, n_q)}}$, then*

$$P\left(\max_{t \in S \cup S^c} \|\mathbf{w}_t\| \geq \frac{\alpha \lambda_{n_p, n_q}}{4(2 - \alpha)}\right) \leq 4 \exp(-c'' \min(n_p, n_q)),$$

where c and c'' are some constants.

See Section B.5 in the supplementary article [10] for the proof. Note that we have ditched Assumptions 1 and 3 since we have already shown that they are automatically satisfied with probability $1 - \delta_{n_q}$ given Assumptions 7 and 8.

5.3. *Smoothness assumption relaxed.* In the previous derivation, the assumption of boundedness of density ratio model guarantees its empirical counterpart $\hat{r}(\mathbf{x}; \theta)$ is always upper-bounded by C_{ratio} and is lower bounded by $\frac{1}{C_{\text{ratio}}}$, but this was somewhat restrictive. In this section, we discuss a relaxation of Assumption 7 as follows.

ASSUMPTION 9 (Smooth density ratio model assumption). For any vector $\delta \in \mathbb{R}^{\dim(\theta^*)}$ such that $\|\delta\| \leq \|\theta^*\|$, the following inequality holds:

$$(20) \quad \begin{aligned} &0 < r(\mathbf{x}, \theta^* + \delta) \leq C_{\max}, \\ &\mathbb{E}_q \left[\inf_{\delta \in \mathbb{R}^{\dim(\theta^*)}: \|\delta\| \leq \|\theta^*\|} r(\mathbf{x}, \theta^* + \delta) \right] \geq 1 - c, \quad 0 < c < 1, \end{aligned}$$

where c is a constant.

Now the strictly positive lower bound of the density ratio model is removed, and we add a new uniform lower bound on the expectation of density ratio model around the true model. Such condition allows us to control the tail of the empirical density ratio model so that the ratio model of a specific sample \mathbf{x} from Q , that is, $\inf_{\delta \in \mathbb{R}^{\dim(\theta^*)}: \|\delta\| \leq \|\theta^*\|} \hat{r}(\mathbf{x}; \theta + \delta)$, would not deviate “too much” from 1.

Note that since smoothness Assumption 9 is still stronger than 5 or 6, the proof of Theorem 1 can be used without modification to show the support consistency when this assumption is substituted. However, the proof of Corollary 3 cannot be used when Assumption 9 is imposed, since the proof requires the boundedness of $\|f_t\|$ which is not implied by this assumption. From now on, we show that such relaxed regime together with moment-bounding Assumption 8 also allows us to bound eigenvalues of derivatives of the likelihood function.

First, we give an example showing that if the expectation of the density ratio derivative is bounded over Q , the above assumption holds.

PROPOSITION 8. *If $\|\theta^*\| \mathbb{E}_q[\sup_{\delta \in \mathbb{R}^{\dim(\theta^*)}: \|\delta\| \leq \|\theta^*\|} \|\nabla r(x, \theta^* + \delta)\|] \leq c, 0 < c < 1$ then (20) holds.*

The proof is listed in Section B.10 in the supplementary article [10]. This proposition intuitively shows that as long as the density ratio model is smooth in the first order, and the changes in parameter is not too drastic, our new assumption holds.

Bounding $\|\nabla_{\theta}^2 \ell(\theta)\|$ and $\|\|\nabla_{\theta_t} \nabla^2 \ell(\theta)\|\|$. It can be seen from Propositions 5 and 6 that upper-bounding the second or the third-order derivative relies on the upper-bound of the empirical density ratio model $\hat{r} \leq C_{\text{ratio}}$. However, under this new assumption, the empirical density ratio model is no longer explicitly bounded. Now we derive the upper-boundedness of \hat{r} using the new assumption.

PROPOSITION 9 (Uniformly upper-bounded \hat{r}). *If Assumption 9 holds,*

$$\sup_{\delta \in \mathbb{R}^{\dim(\theta^*)}: \|\delta\| \leq \|\theta^*\|} \hat{r}(\mathbf{x}; \theta^* + \delta) \leq C'_{\text{ratio}} < \infty,$$

holds with probability at least $1 - \exp(-\frac{2n_q \varepsilon^2}{C_{\max}^2})$.

The proof is in Section B.11 in the supplementary article [10]. Therefore, replace C_{ratio} in (20), (21) and (23) with C'_{ratio} , and we have that the Assumption 3 holds with high probability.

Bounding minimum eigenvalue of $\nabla_{\theta_S}^2 \ell(\theta^)$.* Under Assumption 9, the lower bound of empirical density ratio model $\min_j \hat{r}(\mathbf{x}^{(j)}; \theta^*)$ is no longer valid since the density ratio can approach to 0. Here, we illustrate another proof showing the boundedness of the minimum eigenvalue using concentration inequalities.

PROPOSITION 10. *If Assumption 9 holds, then*

$$(21) \quad 1 - \varepsilon \leq \frac{1}{n_q} \sum_{i=1}^{n_q} r(\mathbf{x}^{(i)}; \theta^*) = \frac{\hat{N}(\theta^*)}{N(\theta^*)} \leq 1 + \varepsilon$$

holds with probability at least $1 - 2 \exp(-\frac{2n_q \varepsilon^2}{C_{\max}^2})$.

PROPOSITION 11. *If Assumption 9 holds and the importance-sampled covariance using true density ratio model $r(\mathbf{x}; \theta^*)$ satisfies*

$$\Lambda_{\min}(\widehat{\text{Cov}}_{qr_{\theta^*}}[f_S(\mathbf{x})]) \geq D'_{\min,2},$$

with probability δ_{n_q} , then $\Lambda_{\min}(\nabla_{\theta_S}^2 \ell(\theta^)) \geq D'_{\min,2}/(1 + \varepsilon)^2, \forall \varepsilon < \infty$ holds with probability at least $1 - 2 \exp(-\frac{2n_q \varepsilon^2}{C_{\max}^2}) - \delta_{n_q}$.*

Proofs are listed in Sections B.12 and B.13 in the supplementary article [10]. In fact, following the derivations used in above proofs, we can rewrite derivatives of $\ell(\theta)$ as higher-order sample statistics importance-sampled by samples from Q with the empirical density ratio model $\hat{r}(\mathbf{x}; \theta^*)$. See Proposition 2 in the supplementary article [10] for a precise statement.

5.4. *Bounded density ratio assumption: How strong is it?* In this section, we have considered a few stronger alternative assumptions to Assumptions 5 and 6 in order to derive the boundedness of derivatives of the likelihood function, which are crucial to the proof of Theorem 1. However, it is natural to ask, how strong are these assumptions?

The main advantage of the density ratio based change detection described in [9] is that such a method does not limit itself to certain distributions. Therefore, limiting the differences between two distributions help us avoid making assumptions on individual MNs.

In fact, the totally bounded density ratio assumption (Assumption 7) is very well justified through our “interest”: learning the changes between patterns (MNs). The power of density ratio is the magnitude of the changes between two density functions. If the change itself is “insanely” big, such a change detection task would

not make any sense in the first place. Unfortunately, such a restriction will rule out some common distributions for change detection, such as Gaussian distribution whose density ratio value is not necessarily upper bounded. Nonetheless, it does not forbid us to consider *truncated Gaussian distributions* where we focus on a “confined area” as our interested region of learning changes.

To loosen this restriction, we utilize another fact that density ratio, like density functions, are naturally lower bounded by 0. Therefore, Section 5.3 is dedicated to the case where the density ratio can decay unbounded toward 0, and thus it allows us to consider the sufficient statistics f_t with unbounded ℓ_2 norm. Illustrative figures of the applicability of our smoothness assumptions are given in Section A.1 in the supplementary article [10].

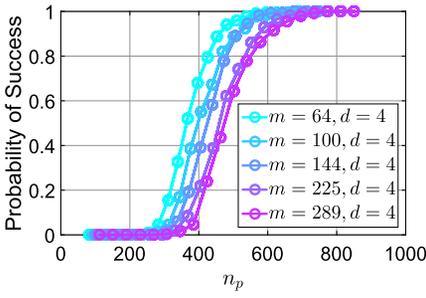
6. Synthetic experiments. In this section, we validate our theorem and compare KLIEP with a state of the art method on synthetic datasets. For a practical usage of KLIEP, see Section 7 for details. The MATLAB code skeleton that is used for our experiments can be found at <https://github.com/lamfeeling/LearningSparseMNChange>.

If all the *sufficient conditions* in Theorem 1 are satisfied, the solution of our optimization problem in (9) should successfully recover the sparsity pattern in θ^* with high probability. Therefore, we can validate our theorem by examining the probability of successful detection of changed edges, that is, the proportion of the simulation where the method exactly recovers the support of the changed edges. We set the regularization parameter as a scaling variable: $\lambda_{n_p} = C \sqrt{\frac{\log m}{n_p}}$, where C is a chosen constant, so the right-hand side inequality of (12) may be satisfied at some point as n_p grows. As $\log \frac{m^2+m}{2}$ is upper bounded by $2 \log m$ if $m > 1$, the left-hand side of (12) is also satisfied if C is appropriately chosen. Note that this is *not* how the hyper-parameter is chosen in practice.

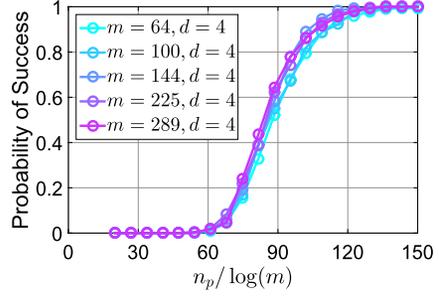
Now using the same reasoning illustrated in Section 4.2, we can deduce that when fixing d , the number of samples n_p required for detecting the correct sparse changes grows linearly with $\log m$, so the success rate versus $n_p/\log m$ plot should align well for MNs with different number of nodes (dimensions) m .

Moreover, our theorem does not have “a preference” on any specific graph structure (such as trees or stars), nor the connectivity of each *individual* MN. Therefore, as long as the number of changed edges d is the same, the success rate plot should have similar behaviors for MNs with different structures. This is a unique feature of *direct change detection* comparing to methods involving learning two separate MNs. See Section E in the supplementary article [10] for detailed experimental settings.

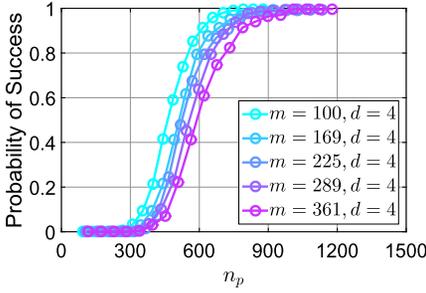
6.1. n_p versus $\log m$. We now illustrate above effects via experiments. Since the density ratio estimator involves two sets of data with size n_p and n_q , to avoid complication, we first set n_q to a sufficiently large value ($n_q = 1000$), and examine



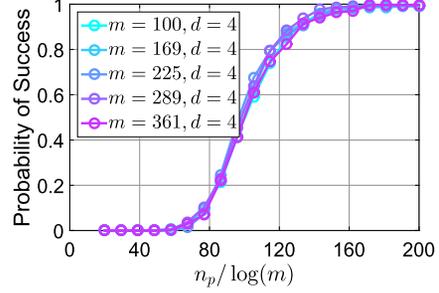
(a) 4-neighbor Lattice, Gaussian, without scaling with $\log m$



(b) Lattice, Gaussian, rescaled by $\log m$



(c) Random structure ($\sim 5\%$ connectivity), Gaussian, without scaling with $\log m$



(d) Random structure, Gaussian, rescaled by $\log m$

FIG. 1. The relationship between n_p and $\log m$, while $n_q = 1000$ is kept fixed. Success rates are computed over 300 runs, same below.

the relationship between m and n_p with $d = 4$ fixed. The results in Figure 1 show all success rate curves align well over MNs of different sizes for both “lattice” or randomly shaped structures.

6.2. Changing n_q . Our theorem also states that n_q should also satisfy a certain relationship with n_p . In this experiment, we vary n_q to observe the change of success rate pattern using the “random” and “lattice” dataset in the previous experiment. As we can see from Figure 2(a), if we set $n_q = 1000$, all successful rate curves reach 100% but are not aligned. Figure 2(b) shows, when $n_q = 500$, we cannot reach 100% success rate even for an ever growing n_p and the probability of success even decays in the final stage. This can be explained by (13) in Theorem 1. If n_q is large enough, the second term in (13) can be safely ignored. However, as λ_{n_p} decays when n_p grows, a small n_q may not be able to suppress the second term and the overall probability of success starts to decay eventually. By setting $n_q = 0.01n_p^2$, we obtain a perfectly aligned result [Figure 2(c)], as our theorem indicated. It also shows that though n_q is required to grow quadratically with n_p , it can be rescaled by a small constant (in this case, 0.01). Moreover, Corol-

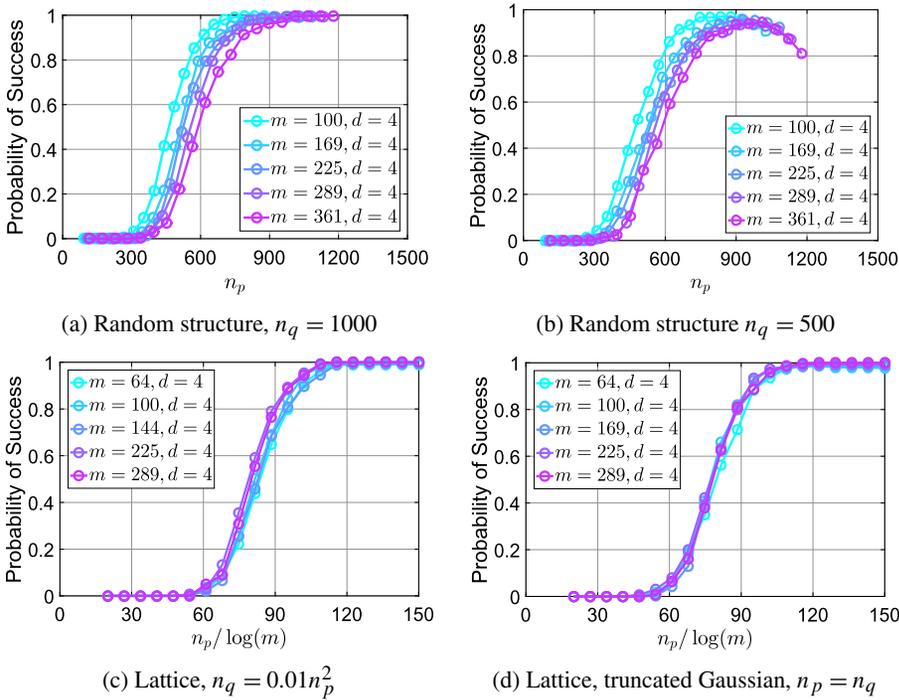


FIG. 2. The relationship between n_p and $\log m$, when varying n_q . (d) Gaussian is truncated within a ball centred at the origin with the radius of 15.

lary 3 points out if the density ratio model is bounded, we may relax the coupling condition between n_p and n_q . To verify this, we truncate a Gaussian distribution by rejecting samples fall out of a ball centered at origin with radius 15, then let $n_q = n_p$. From Figure 2(d), we can see a similar patten of success rates alignment.

6.3. *Changing d.* As Theorem 1 indicates, n_p should grow at least quadratically with d , the number of *changed edges*. However, in our experiments, it shows such condition is overly conservative. Figure 3(a) shows the success rate depends on d only very mildly (see $n_p / \log m$ needed for success rates passing 80%), which is a good news. This indicates that the bound can be tightened under certain regimes.

6.4. *Non-Gaussian distribution.* We next perform experiments on the truncated “8”-shaped distribution (see Figure 2 in the supplementary article [10] for details). The MNs are constructed as lattices, and the samples are generated via slice sampling [13]. Figure 3(c) shows, for the lattice grids with dimensions $m = 60 \sim 256$, the curves of success rates are well aligned with the setting $n_q = 5n_p$.

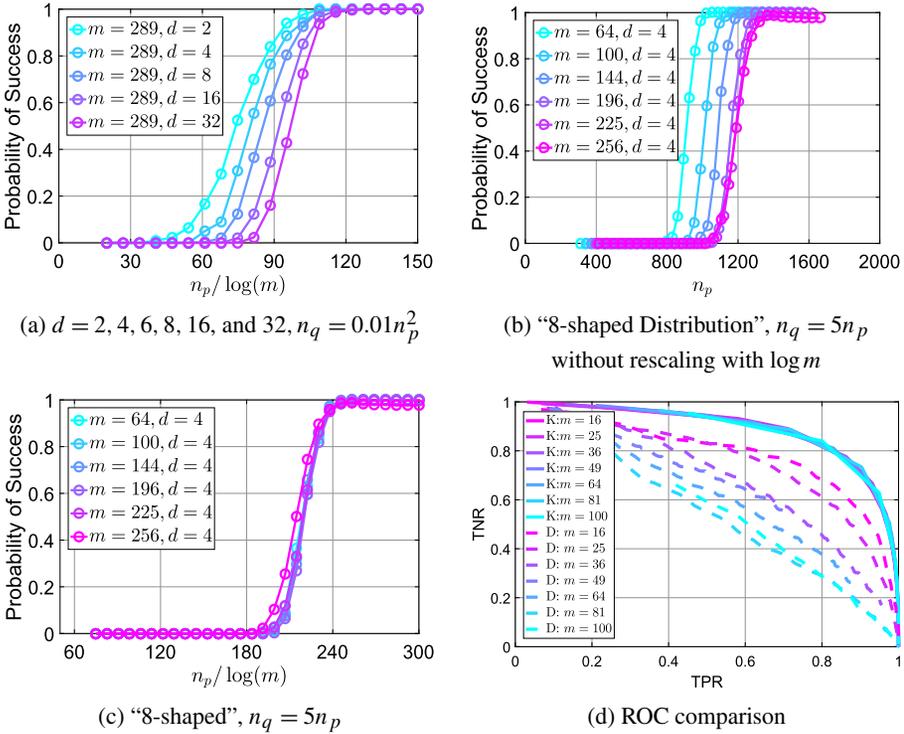


FIG. 3. The relationship between n_p and $\log m$, when varying d (a) or under a non-Gaussian distribution [(b) and (c)]. (d) is the ROC comparison between (K)LIEP and (D)ifferential learning.

6.5. Comparison with differential network learning. In this section, we conduct experimental comparison between KLIEP and the differential network learning proposed in [29], where the following constrained objective is minimized:

$$\hat{\Delta} = \underset{\Delta}{\operatorname{argmin}} \|\Delta\|_1 \quad \text{subject to} \quad \|\hat{\Sigma}^{(p)} \Delta \hat{\Sigma}^{(q)} + \hat{\Sigma}^{(p)} - \hat{\Sigma}^{(q)}\|_\infty < \varepsilon,$$

where $\hat{\Sigma}^{(p)}$ and $\hat{\Sigma}^{(q)}$ are the sample covariance matrices, ε is a small constant and $\hat{\Delta}$ is the estimated differential network. To obtain a sparse solution, Zhao et al. [29] thresholds the solution at a certain level τ , that is, for all u, v that $\Theta_{u,v} < \tau$ is thresholded to 0.

We compare the performance between KLIEP and differential network learning by Receiver Operating Characteristic (ROC) plot of True Positive and True Negative rate described in [29]. See Section F in the supplementary article [10] for more detailed settings. Experimental data are constructed using the synthetic Gaussian 4-neighbor lattice MN grid described in Section 6.1. However, instead of fixing some key variables, we adopt a more practical setting: $n_p = n_q = 50$ and $d = \sqrt{m}$ (i.e., the number of changed edges increases with m). The ROC curves of KLIEP (solid) and the differential method (dashed) are reported in Figure 3(d).

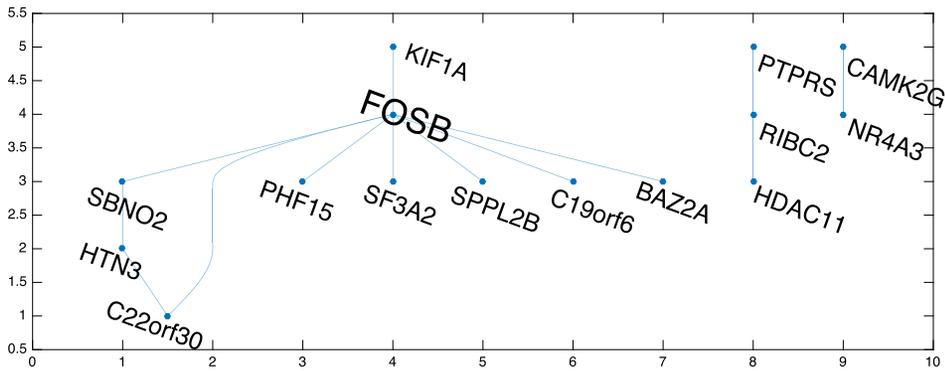


FIG. 4. Gene expression change graph highlighting a hub node.

As it can be seen from Figure 3(d), KLIEP performs consistently better than the differential network learning method. Particularly, when m and d increases, the performance almost remains the same for KLIEP while for the differential method, it decays significantly.

Figure 3(d) is plotted up to $m = 100$, due to the fact that the differential method fails to return an output within a reasonable amount of time (10 hours) when m is scaled up to 121, using either the authors' or our implementation. This is consistent with the authors' claim in [29], where the differential method requires more than 14 hours to obtain the result for a MN sized 120. As for KLIEP, we can compute a full ROC curve within 3 hours even for $m = 625$ when $n_p = n_q = 50$.

7. Gene expression analysis. We applied KLIEP to gene expression profiles for estimating changes in gene networks activated by two different stimuli: epidermal growth factor (EGF) and heregulin (HRG). EGF is known to induce proliferation in MCF7 human breast cancer cells, while HRG induces differentiation. We used the gene expression data from [12]. The expression profiles were from cells stimulated with two controls, resulting in 29 EGF and 28 HRG sample conditions ($n_p = 29, n_q = 28$). We extracted 1835 genes ($m = 1835$) from the gene set in [12] by selecting genes with high expression variance (at least three times the mean of all variances). The values were log₂-transformed and normalized using the 2% trimmed mean before finally getting the respective ratios with the controls. The change graph is obtained by reducing the regularization parameter until $|\hat{S}| > 10$.

Figure 4 shows a learned change network. Each node represents a gene, and each edge indicates that the regulation between them is different from EGF stimuli to HRG stimuli. The leftmost large component includes 10 genes and 10 interactions, and a hub node of the component is the FOSB gene, which is a member of the Fos family of transcription factors, regulating expressions of other genes. This indicates that KLIEP successfully found that FOSB regulates other genes without any prior knowledge, and suggests that the regulation has been changed between

stimuli. Moreover, FOSB is known as a regulator of cell proliferation and differentiation [12], showing that the detected change network agrees with biological knowledge. The result confirms that KLIEP can detect known biologically significant changes from the expression data. We also swapped P and Q , and confirmed that detected changes are similar.

To evaluate the reliability of this experiment, we also conducted bootstrap experiments which are detailed in Section G in the supplementary article [10].

8. Discussion. Let us further discuss some properties of KLIEP.

8.1. *Asymmetry of KLIEP.* It can be noticed that KLIEP is asymmetric, that is, the structural change learning performance may differ when swapping P and Q . This is caused by directly parametrizing the density ratio which is naturally asymmetric. Such asymmetry is also reflected in Theorem 1, where the sample complexity is not the same as $n_p = \Omega(d \log m(m+1)/2)$ and $n_q = \Omega(n_p^2)$, though they are “symmetrized” in Corollary 3 as $\min(n_p, n_q) = \Omega(d \log(m(m+1)/2))$ under the stronger bounded ratio assumption (Assumption 7). These analyses shed some light on choosing P and Q in practice when two datasets are given: Q should be “wide” and more “spread-out” compared to P , so that the boundedness of the density ratio can be guaranteed, since a “sharp” Q would lead to very sharp and unbounded density ratio. This insight implies that KLIEP is in fact a *directional* method, and achieves better performance when the change itself shows a tendency of evolving from one “general” state to another more “specialized” state.

Symmetric measures of differences between densities are also available without using the density ratio, such as ℓ_2 distance [17]. In general, such difference cannot be parameterized using the difference between parameters of individual MNs, thus cannot serve as a difference measure for direct structure change learning. An alternative is the differential network learning [29] where the objective function is symmetric for both P and Q . However, such an objective is only sensible when both P and Q are Gaussian and it cannot be easily generalized to non-Gaussian cases.

Moreover, inspired by Corollary 3, we may consider a “symmetrized” version of KLIEP by learning p/q and q/p independently, after which we mark changing edges by taking the *union* of two sparsity patterns in both models. For the support recovery probability, we simply need a union bound applied on the current results. Such a union-support algorithm is similar to the nodewise regression that is discussed in [15].

8.2. *Comparison with differential network learning [29].* Other than the asymmetry issues mentioned above, the theoretical analysis in this work and the one in [29] share some key similarities and both have good guarantees in a high-dimensional setting. First, they both set assumptions on the true difference/changes

between two MNs. In [29], such a constrain is explicitly expressed: The true differential network $\|\Delta^*\|_1$ is bounded by a constant that does not grow with m (Condition 1, [29]). In our analysis, the assumption is implicitly made via limiting the “smoothness” of the density ratio (Assumptions 5, 6, 7, 9). Second, differential network learning prohibits strong “connections” among covariates. It assumes the magnitude of off-diagonal values in the covariance matrices decays as the number of changed edges increases (Condition 2, [29]), while in our work, we assume feature vectors on unchanged edges should not have strong correlations with those on changed edges (Assumption 2). Both assumptions are imposed to satisfy the incoherence condition [3] needed for Lasso-type model selection. The convergence results are also similar: The required sample size scales with $\log m$ for differential network learning (Theorem 2, [29]) and KLIEP [$\log(m^2 + m) \approx 2 \log m$]; the difference between the true parameter and the estimated one vanishes at the speed of $\sqrt{\frac{\log m}{\min(n_p, n_q)}}$ in the ℓ_2 norm. However, we manage to achieve this rate without any assumptions of Gaussianity, while differential network learning described in [29] cannot be directly applied on non-Gaussian MNs.

8.3. *Joint structural change learning.* Another emerging trend in graphical model structural learning is to learn multiple *similar* MNs simultaneously [2]. For example, one may use the following fused-lasso penalized objective function to learn K -Gaussian MNs at the same time:

$$\begin{aligned}
 (22) \quad & \min_{\Theta^{(1)}, \dots, \Theta^{(K)}} \sum_{a=1}^K \ell_{\text{MLE}}(\Theta^{(a)}) + \lambda_1 \sum_{\substack{a < b \\ a, b \in \{1, \dots, K\}}} \|\Theta^{(a)} - \Theta^{(b)}\|_1 \\
 & + \lambda_2 \sum_{a \in \{1, \dots, K\}} \|\Theta^{(a)}\|_1,
 \end{aligned}$$

where $\ell_{\text{MLE}}(\Theta)$ is the negative Gaussian MN log-likelihood parametrized by precision matrix Θ and the assumption is that all $\Theta^{(a)}$ share a similar structure. Though the final outputs are estimated sparse precision matrices, one may still obtain a differential graph by taking the differences. We refer readers to [9] for more empirical comparisons between KLIEP and Fused-lasso differential network learning.

Following the same spirit, the KLIEP based change detection can also be utilized to learn *changes* from K -MNs at the same time by assuming all MNs sharing a similar structure, and one of the options is

$$\begin{aligned}
 (23) \quad & \min_{\theta^{a,b}, a, b \in \{1, \dots, K\}, a \neq b} \sum_{\substack{a, b \in \{1, \dots, K\} \\ a \neq b}}^K \ell_{\text{KLIEP}}^{a,b}(\theta^{a,b}) \\
 & + \lambda \sum_{a, b, c \in \{1, \dots, K\}} \sum_{u, v=1, u \leq v}^m \|\theta_{u,v}^{a,b} + \theta_{u,v}^{b,c}\|_2,
 \end{aligned}$$

where $\ell_{\text{KLIEP}}^{a,b}(\theta^{a,b})$ is the KLIEP log-likelihood of density ratio p^a/p^b , a, b, c are the triples from $1, \dots, K$ and by definition we have

$$\theta_{u,v}^{a,b} + \theta_{u,v}^{b,c} \equiv \theta_{u,v}^a - \theta_{u,v}^b + \theta_{u,v}^b - \theta_{u,v}^c \equiv \theta_{u,v}^a - \theta_{u,v}^c,$$

and let $\theta_{u,v}^{a,a} \equiv \mathbf{0}$.

The objective described in (22) still requires a tractable likelihood and cannot be easily generalized to non-Gaussian models. It also assumes each individual MN is sparse. However, the variation of KLIEP given in (23) does not impose any assumptions on each individual MNs and can be computed even in non-Gaussian cases. We will explore this algorithm and other possible alternatives in future works.

8.4. Uncertainty of estimation. In this work, we have only focused on the successful change detection (Theorem 1) which states $P(S = \hat{S} \text{ and } S^c = \hat{S}^c)$ converges to one eventually. However, under very a low sample regime, the learned change structure may have high uncertainty, that is, the results may be sensitive to some minor modifications or the randomness of the dataset. This would degrade the reliability of the obtained results. To evaluate such an uncertainty, we performed bootstrap experiments on the gene dataset used in Section 7, and the results are presented in Section G in the supplementary article [10]. Such a method is useful for practitioners to measure the reliability of an estimated network. However, the rigours quantification of such an uncertainty is still an open question and would be an important future direction to pursue.

Acknowledgements. The authors thank two anonymous reviewers who provided valuable suggestions on this paper.

SUPPLEMENTARY MATERIAL

Supplement to “Support consistency of direct sparse-change learning in Markov networks” (DOI: [10.1214/16-AOS1470SUPP](https://doi.org/10.1214/16-AOS1470SUPP); .pdf). Due to the page limit, we present the proofs of Lemmas 1–5, corollaries and propositions in this supplementary article. We also use this supplementary material to show some detailed experimental settings, extended simulation results and illustrations of a few concepts in the paper.

REFERENCES

- [1] BANERJEE, O., EL GHAOU, L. and D’ASPROMONT, A. (2008). Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *J. Mach. Learn. Res.* **9** 485–516. [MR2417243](https://arxiv.org/abs/0806.2644)
- [2] DANAHER, P., WANG, P. and WITTEN, D. M. (2014). The joint graphical lasso for inverse covariance estimation across multiple classes. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **76** 373–397. [MR3164871](https://arxiv.org/abs/1308.4033)

- [3] DONOHO, D. L. and HUO, X. (2001). Uncertainty principles and ideal atomic decomposition. *IEEE Trans. Inform. Theory* **47** 2845–2862. [MR1872845](#)
- [4] FRIEDMAN, J., HASTIE, T. and TIBSHIRANI, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9** 432–441.
- [5] HAMMERSLEY, J. M. and CLIFFORD, P. (1971). *Markov Fields on Finite Graphs and Lattices*. Available at <http://www.statslab.cam.ac.uk/~grg/books/hammfest/hamm-cliff.pdf>.
- [6] KOLAR, M. and XING, E. P. (2012). Estimating networks with jumps. *Electron. J. Stat.* **6** 2069–2106. [MR3020257](#)
- [7] KOLLER, D. and FRIEDMAN, N. (2009). *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge, MA. [MR2778120](#)
- [8] LEE, S.-I., GANAPATHI, V. and KOLLER, D. (2007). Efficient structure learning of Markov networks using l_1 -regularization. In *Advances in Neural Information Processing Systems* (B. Schölkopf, J. Platt and T. Hoffman, eds.) **19** 817–824. MIT Press, Cambridge, MA.
- [9] LIU, S., QUINN, J. A., GUTMANN, M. U., SUZUKI, T. and SUGIYAMA, M. (2014). Direct learning of sparse changes in Markov networks by density ratio estimation. *Neural Comput.* **26** 1169–1197. [MR3234638](#)
- [10] LIU, S., SUZUKI, T., RELATOR, R., SESE, J., SUGIYAMA, M. and FUKUMIZU, K. (2016). Supplement to “Support consistency of direct sparse-change learning in Markov networks.” DOI:[10.1214/16-AOS1470SUPP](https://doi.org/10.1214/16-AOS1470SUPP).
- [11] MEINSHAUSEN, N. and BÜHLMANN, P. (2006). High-dimensional graphs and variable selection with the lasso. *Ann. Statist.* **34** 1436–1462. [MR2278363](#)
- [12] NAGASHIMA, T., SHIMODAIRA, H., IDE, K., NAKAKUKI, T., TANI, Y., TAKAHASHI, K., YUMOTO, N. and HATAKEYAMA, M. (2007). Quantitative transcriptional control of erbB receptor signaling undergoes graded to biphasic response for cell differentiation. *Journal of Biological Chemistry* **282** 4045–4056.
- [13] NEAL, R. M. (2003). Slice sampling. *Ann. Statist.* **31** 705–767. [MR1994729](#)
- [14] RASKUTTI, G., YU, B., WAINWRIGHT, M. J. and RAVIKUMAR, P. (2009). Model selection in Gaussian graphical models: High-dimensional consistency of l_1 -regularized mle. In *Advances in Neural Information Processing Systems* (D. Koller, D. Schuurmans, Y. Bengio and L. Bottou, eds.) **21** 1329–1336. Curran Associates, Red Hook, NY.
- [15] RAVIKUMAR, P., WAINWRIGHT, M. J. and LAFFERTY, J. D. (2010). High-dimensional Ising model selection using l_1 -regularized logistic regression. *Ann. Statist.* **38** 1287–1319. [MR2662343](#)
- [16] ROBERT, C. P. and CASELLA, G. (2004). *Monte Carlo Statistical Methods*. Springer, New York. [MR1707311](#)
- [17] SUGIYAMA, M., KANAMORI, T., SUZUKI, T., DU PLESSIS, M. C., LIU, S. and TAKEUCHI, I. (2013). Density-difference estimation. *Neural Comput.* **25** 2734–2775. [MR3114585](#)
- [18] SUGIYAMA, M., NAKAJIMA, S., KASHIMA, H., VON BÜNAU, P. and KAWANABE, M. (2008). Direct importance estimation with model selection and its application to covariate shift adaptation. In *Advances in Neural Information Processing Systems* 20 (J. C. Platt, D. Koller, Y. Singer and S. T. Roweis, eds.). Curran Associates, Red Hook, NY.
- [19] SUGIYAMA, M., SUZUKI, T. and KANAMORI, T. (2012). *Density Ratio Estimation in Machine Learning*. Cambridge Univ. Press, Cambridge. [MR2895762](#)
- [20] TIBSHIRANI, R., SAUNDERS, M., ROSSET, S., ZHU, J. and KNIGHT, K. (2005). Sparsity and smoothness via the fused lasso. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **67** 91–108. [MR2136641](#)
- [21] TOMIOKA, R. and SUZUKI, T. (2014). Spectral norm of random tensors. Preprint. Available at [arXiv:1407.1870](https://arxiv.org/abs/1407.1870).
- [22] TSUBOI, Y., KASHIMA, H., HIDO, S., BICKEL, S. and SUGIYAMA, M. (2009). Direct density ratio estimation for large-scale covariate shift adaptation. *Journal of Information Processing* **17** 138–155.

- [23] WAINWRIGHT, M. J. (2009). Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (Lasso). *IEEE Trans. Inform. Theory* **55** 2183–2202. [MR2729873](#)
- [24] WASSERMAN, L. (2004). *All of Statistics: A Concise Course in Statistical Inference*. Springer, New York. [MR2055670](#)
- [25] YAMADA, M., SUZUKI, T., KANAMORI, T., HACHIYA, H. and SUGIYAMA, M. (2013). Relative density-ratio estimation for robust distribution comparison. *Neural Comput.* **25** 1324–1370. [MR3075782](#)
- [26] YANG, E., GENEVERA, A., LIU, Z. and RAVIKUMAR, P. (2012). Graphical models via generalized linear models. In *Advances in Neural Information Processing Systems* (F. Pereira, C. J. C. Burges, L. Bottou and K. Q. Weinberger, eds.) **25** 1358–1366. Curran Associates, Red Hook, NY.
- [27] YUAN, M. and LIN, Y. (2006). Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **68** 49–67. [MR2212574](#)
- [28] ZHANG, B. and WANG, Y. J. (2010). Learning structural changes of Gaussian graphical models in controlled experiments. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence (UAI2010)* 701–708. AUAI Press, Corvallis, OR.
- [29] ZHAO, S. D., CAI, T. T. and LI, H. (2014). Direct estimation of differential networks. *Biometrika* **101** 253–268. [MR3215346](#)

S. LIU
 K. FUKUMIZU
 RESEARCH CENTER FOR STATISTICAL
 MACHINE LEARNING
 INSTITUTE OF STATISTICAL MATHEMATICS
 10-3 MIDORI-CHO, TACHIKAWA
 TOKYO 190-8562
 JAPAN
 E-MAIL: liu@ism.ac.jp
fukumizu@ism.ac.jp

M. SUGIYAMA
 DEPARTMENT OF COMPLEXITY SCIENCE
 AND ENGINEERING
 GRADUATE SCHOOL OF FRONTIER SCIENCES
 UNIVERSITY OF TOKYO
 7-3-1 HONGO, BUNKYO-KU
 TOKYO 113-0033
 JAPAN
 E-MAIL: sugi@k.u-tokyo.ac.jp

T. SUZUKI
 DEPARTMENT OF MATHEMATICAL
 AND COMPUTING SCIENCE
 SCHOOL OF COMPUTING
 TOKYO INSTITUTE OF TECHNOLOGY
 2-12-1 O-OKAYAMA, MEGURO
 TOKYO 152-8552
 JAPAN
 E-MAIL: suzuki.t.ct@m.titech.ac.jp

R. RELATOR
 J. SESE
 ARTIFICIAL INTELLIGENCE RESEARCH CENTER
 NATIONAL INSTITUTE OF ADVANCED-
 INDUSTRIAL SCIENCE AND TECHNOLOGY
 2-4-7 AOMI, KOTO-KU
 TOKYO 135-0064
 JAPAN
 E-MAIL: rai.relator@aist.go.jp
sese.jun@aist.go.jp