Learning sparse structural changes in high-dimensional Markov networks

A review on methodologies and theories

Song Liu¹ · Kenji Fukumizu¹ · Taiji Suzuki^{2,3}

Received: 15 September 2016 / Accepted: 9 January 2017 / Published online: 13 February 2017
© The Author(s) 2017. This article is published with open access at Springerlink.com

Abstract Recent years have seen an increasing popularity of learning the sparse *changes* in Markov Networks. Changes in the structure of Markov Networks reflect alternations of interactions between random variables under different regimes and provide insights into the underlying system. While each individual network structure can be complicated and difficult to learn, the overall change from one network to another can be simple. This intuition gave birth to an approach that *directly* learns the sparse changes without modelling and learning the individual (possibly dense) networks. In this paper, we review such a direct learning method with some latest developments along this line of research.

Keywords Markov network · Density ratio estimation · Change detection

1 Introduction

The problem of learning the *changes of interactions* between random variables can be useful in many applications. For example, genes may regulate each other in different ways when external conditions are changed; the number of daily flu-like

Communicated by Brandon Malone.

✉ Song Liu
liu@ism.ac.jp

Kenji Fukumizu
fukumizu@ism.ac.jp

Taiji Suzuki
suzuki.t.ct@m.titech.ac.jp

¹ The Institute of Statistical Mathematics, 10-3 Midori-cho, Tachikawa, Tokyo 190-8562, Japan

² Tokyo Institute of Technology, 2 Chome-12-1 Ookayama, Meguro, Tokyo 152-8550, Japan

³ PRESTO, Japan Science and Technological Agency (JST), 4-1-8 Honcho, Kawaguchi-shi, Saitama 332-0012, Japan

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 subject is performing different activities. Spotting such changes in interactions may provide key insights into the underlying system.

The interactions among random variables can be formulated as undirected probabilistic graphical models, or Markov Networks (MNs) (Koller and Friedman 2009), expressing the interactions via the *conditional independence*. We consider a simple model: the *pairwise MNs* where the links are only encoded for single or pairs of random variables. Due to the Hammersley–Clifford theorem (Hammersley and Clifford 1971), the underlying joint probability density function can be represented as the product of univariate and bivariate factors.

As an important challenge, structure learning of MNs has also attracted a significant amount of attention. Earlier methods (Spirtes et al. 2000) use hypothesis testing to learn the conditional independence among random variables, which reflects the absences of edges. It is proved that such a problem is generally NP-hard (Chickering 1996). Methods restricted to a sub-class of graphical models (such as trees or forests) (Chow and Liu 1968; Geman and Geman 1984; Liu et al. 2011) also suffer from growing computational cost.

However, the Hammersley–Clifford theorem together with the recent breakthrough on sparsity-inducing methods (Tibshirani 1996; Zhao and Yu 2006; Wainwright 2009) gave birth to many sparse structure learning ideas where the sparse factorization of the joint/conditional density function was estimated to infer the underlying structure of the MN (Friedman et al. 2008; Banerjee et al. 2008; Meinshausen and Bühlmann 2006; Ravikumar et al. 2010). Although most works focused on parametric models, the structure learning has been conducted on semi-parametric ones in recent years (Liu et al. 2009, 2012).

There is also a trend of learning the *changes* between MNs (Zhang and Wang 2010; Liu et al. 2014; Zhao et al. 2014). Comparing to standard structure learning, the learning of changes views the problem in a more *dynamic* fashion: Instead of estimating a static pattern, we hope to obtain a dynamic one, namely “the change” by comparing two sets of data since in some applications, the static pattern may not be computationally tractable, or simply too hard to comprehend. However, the difference between two patterns may be represented by some simple incremental effects involving only a small number of nodes or bonds. Thus it takes much less effort to learn and understand.

One of the main uses of structural change learning is to spot responding variables in “controlled experiments” (Zhang and Wang 2010) where some key external factors of the experiments are altered, and two sets of samples are obtained. By discovering the changes in the MNs, we can see how random variables have responded to the change of the external stimuli.

In this paper, we first review a recently proposed method of structural change learning between MNs (Liu et al. 2014). This follows a simple idea: if the MNs are products of the pairwise factors, the *ratio* of two MNs must also be proportional to the ratios of those factors. Moreover, factors that do not change between two MNs will have no contribution to the ratio. This naturally suggests the idea of modelling

the changes between two MNs P and Q as the ratio between two MN density functions $p(\mathbf{x})$ and $q(\mathbf{x})$. The ratio $p(\mathbf{x})/q(\mathbf{x})$ is directly estimated from a one-shot estimation (Sugiyama et al. 2012). This density-ratio approach can work well even when each MN is dense (as long as the change is sparse).

We also present some very recent theoretical results along this line of research. These works prove the *consistency* of the density ratio method in the *high-dimensional setting*. The *support consistency* indicates the support of the estimated parameter converges to the support of the true parameter in probability. This is an important property for sparsity inducing methods. It is shown that under certain conditions the density ratio method recovers the correct parameter sparsity with high probability (Liu et al. 2017b). Moreover, Fazayeli and Banerjee introduced a theorem for the regularized density ratio estimator showing the estimation error, i.e., the ℓ_2 distance between the estimated parameter and the true parameter converges to zero under milder conditions.

As comparisons, we will also show a few alternative approaches to the change detection problem between MNs. The differential graphical model learning approach (Zhao et al. 2014) uses a covariance-precision matrix equality to learn changes without going through the learning of the individual MNs. The “jumping” MNs (Kolar and Xing 2012) setting considers a scenario where the observations are received as a sequence and multiple sub-sequences are generated via different parametrizations of MN.

We organize this paper as follows: Firstly, we introduce the problem formulation of learning changes between MNs in Sect. 2. Secondly, the density ratio approach and two other alternatives are explained in Sect. 3. Section 4 reviews the theoretical results of these approaches. Synthetic and real-world experiments are conducted in Sect. 5 to compare the performance of methods. Finally, in Sects. 6 and 7, we give a few possible future directions and conclude the current developments along this line of research.

2 Formulating changes

In this section, we focus on formulating the change of MNs using density ratio. At the end of this section, a few alternatives are also introduced.

2.1 Structural changes by parametric differences

Detecting changes naturally involves two sets of data. Consider independent samples drawn separately from two probability distributions P and Q on $\mathbb{R}^m$:

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

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

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

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 pair (x_u, x_v) , and

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

is the entire parameter vector. The feature function $\boldsymbol{\psi}_{u,v}(x_u, x_v)$ 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. Such a parametrization is generic when representing pairwise graphical models.

Directly estimating an MN in this generic form is challenging since $Z(\boldsymbol{\theta}^{(p)})$ usually does not have a closed form except for a few special cases (e.g. Gaussian distribution). Markov Chain Monte Carlo (Robert and Casella 2005) is used to approximate such an integral. However, this would bring extra approximation errors.

Nonetheless, we can define *changes* between two MNs as the difference between their parameters. Therefore, 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 , i.e.,

$$\boldsymbol{\delta} = \boldsymbol{\theta}^{(p)} - \boldsymbol{\theta}^{(q)}.$$

Note that by its definition, the *changes* are continuous. This is more advantageous than only considering discrete changes of the MN structure, since a weak change of interaction does not necessarily shatter or flip the bond between two random variables.

2.2 Density ratio modelling

An important observation is that although two MNs may be complex individually, their changes might be “simple” since many terms may be cancelled while taking the difference, i.e. $\boldsymbol{\theta}_{u,v}^{(p)} - \boldsymbol{\theta}_{u,v}^{(q)}$ might be zero. The key idea in (Liu et al. 2014) is to consider the *ratio* of p and q :

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

where $\boldsymbol{\theta}_{u,v}^{(p)} - \boldsymbol{\theta}_{u,v}^{(q)}$ encodes the difference between P and Q for factor $\boldsymbol{\psi}_{u,v}(x_u, x_v)$, i.e., $\boldsymbol{\theta}_{u,v}^{(p)} - \boldsymbol{\theta}_{u,v}^{(q)}$ is zero if there is no change in the factor $\boldsymbol{\psi}_{u,v}(x_u, x_v)$.

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. Their difference $\delta_{u,v} = \theta_{u,v}^{(p)} - \theta_{u,v}^{(q)}$ is sufficient for change detection, as $\mathbf{x}$ only interacts with such a parametric difference in the ratio model. Thus, in this density-ratio formulation, p and q are no longer modelled separately, but *directly* as

$$r(\mathbf{x}; \delta) = \frac{1}{N(\delta)} \exp \left(\sum_{u,v=1, u \geq v}^m \delta_{u,v}^\top \psi_{u,v}(x_u, x_v) \right), \quad (3)$$

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

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

$$N(\delta) = \int q(\mathbf{x}) \exp \left(\sum_{u,v=1, u \geq v}^m \delta_{u,v}^\top \psi_{u,v}(x_u, x_v) \right) d\mathbf{x},$$

which is the expectation over $q(\mathbf{x})$.¹ Note this integral is with respect to a true distribution where our samples are generated.² 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 $\{\mathbf{x}_q^{(i)}\}_{i=1}^{n_q} \stackrel{\text{i.i.d.}}{\sim} Q$:

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

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

Interestingly, if one uses $\psi_{u,v}(x_u, x_v) = x_u x_v$ in the ratio model, it does not mean one assumes two individual MNs are Gaussian or Ising, it simply means we assume the changes of interactions are linear while other non-linear interactions remain unchanged. This formulation allows us to consider highly complicated MNs as long as their changes are “simple”.

Throughout the rest of the paper, we simplify the notation from $\psi_{u,v}$ to ψ by assuming the feature functions are the same for all pairs of random variables.

2.3 Quasi log-likelihood equality

Density ratio is not the only direct modelling approach. Particularly for Gaussian MNs, where two distributions are parametrized as $p(\mathbf{x}; \theta^{(p)})$ and $p(\mathbf{x}; \theta^{(q)})$ with the

¹ If one models the ratio $\frac{q(\mathbf{x})}{p(\mathbf{x})}$, the normalization

$$N(\delta) = \int p(\mathbf{x}) \exp \left(\sum_{u,v=1, u \geq v}^m \delta_{u,v}^\top \psi_{u,v}(x_u, x_v) \right) d\mathbf{x}$$

should be used.

² $q(\mathbf{x})$ should not be confused with $q(\mathbf{x}; \theta)$.

precision matrix Θ , one alternative was proposed using the following equality (Zhao et al. 2014):

$$\Sigma^{(p)}(\Theta^{(p)} - \Theta^{(q)})\Sigma^{(q)} + \Sigma^{(p)} - \Sigma^{(q)} = \mathbf{0}, \quad (4)$$

where $\Sigma^{(p)}$ is the covariance matrix of the Gaussian distribution p . As we replace the covariance matrices $\Sigma^{(p)}$ and $\Sigma^{(q)}$ with their sample versions $\widehat{\Sigma}^{(p)}$ and $\widehat{\Sigma}^{(q)}$, it can be seen that $\Theta^{(p)} - \Theta^{(q)}$ is the only variable interacting with the data. Therefore, one may replace it with a single parameter Δ and later minimize the sample version of (4) (see Sect. 3.3 for details).

This direct formulation specifically uses a property of Gaussian MN that the covariance matrix computed from the data and the precision matrix that encodes the MN structure should approximately cancel each other when multiplied. However, such a relationship does not hold for other distributions in general. Studies on the generality of this equality is an interesting open question (see Sect. 6).

Remark In fact, it is not necessary to combine $\theta^{(p)} - \theta^{(q)}$ in (2) (or $\Theta^{(p)} - \Theta^{(q)}$ in (4)) into one parameter. However, such a model will be unidentifiable since there are too many combinations of $\Theta^{(p)}$ and $\Theta^{(q)}$ can produce the same difference. Nonetheless, such an indirect modelling may still be useful when the individual structures of the MNs are also our interests. We review an example of such indirect modelling in Sect. 3.4.

3 Learning sparse changes in Markov networks

3.1 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 (Sugiyama et al. 2012). In Liu et al. (2014), a density ratio estimator called the *Kullback–Leibler importance estimation procedure* (KLIEP) for log-linear models (Sugiyama et al. 2008; Tsuboi et al. 2009) was employed in learning structural changes.

For a density ratio model $r(\mathbf{x}; \delta)$ (as introduced in (3)), the KLIEP method minimizes the Kullback–Leibler divergence from $p(\mathbf{x})$ to $\hat{p}(\mathbf{x}; \delta) = q(\mathbf{x})r(\mathbf{x}; \delta)$:

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

Note that the density ratio model (3) automatically satisfies the non-negativity and normalization constraints:

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

Here we define

$$\hat{r}(\mathbf{x}; \boldsymbol{\delta}) = \frac{\exp\left(\sum_{u,v=1, u \geq v}^m \boldsymbol{\delta}_{u,v}^\top \boldsymbol{\psi}(x_u, x_v)\right)}{\hat{N}(\boldsymbol{\delta}; \mathbf{x}_q^{(1)}, \dots, \mathbf{x}_q^{(n_q)})}$$

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

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

Optimization Since $\ell_{\text{KLIEP}}(\boldsymbol{\delta})$ consists of a linear part and a *log-sum-exp* function (Boyd and Vandenberghe 2004), it is convex with respect to $\boldsymbol{\delta}$, and its global minimizer can be numerically found by standard optimization techniques such as gradient descent. The gradient of ℓ_{KLIEP} with respect to $\boldsymbol{\delta}_{u,v}$ is given by

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

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

3.2 Sparsity inducing and regularizations

In the search for *sparse* changes, one may regularize the KLIEP solution with a sparsity-inducing norm $\sum_{u \geq v} \|\boldsymbol{\delta}_{u,v}\|$, i.e., the *group-lasso* penalty (Yuan and Lin 2006) where we use $\|\cdot\|$ to denote the ℓ_2 norm.

Note that the density-ratio approach (Liu et al. 2014) directly sparsifies the difference $\boldsymbol{\theta}^{(p)} - \boldsymbol{\theta}^{(q)}$, and thus intuitively this method can still work well even if $\boldsymbol{\theta}^{(p)}$ and $\boldsymbol{\theta}^{(q)}$ are dense as long as $\boldsymbol{\theta}^{(p)} - \boldsymbol{\theta}^{(q)}$ is sparse. The following is the objective function used in (Liu et al. 2014):

$$\hat{\boldsymbol{\delta}} = \underset{\boldsymbol{\delta}}{\operatorname{argmin}} \ell_{\text{KLIEP}}(\boldsymbol{\delta}) + \lambda \sum_{u,v=1, u \geq v}^m \|\boldsymbol{\delta}_{u,v}\|. \quad (7)$$

In a recent work (Fazayeli and Banerjee 2016), authors considered *structured* changes, such as sparse, block sparse, node-perturbed sparse and so on. These structured changes can be represented via suitable atomic norms (Chandrasekaran et al. 2012; Mohan et al. 2014). For example, a KLIEP objective with a node-perturbation regularizer is

$$\begin{aligned} \hat{\mathcal{A}} = & \underset{\mathcal{A} \in \mathbb{R}^{m \times m}, \mathcal{L} \in \mathbb{R}^{m \times m}}{\operatorname{argmin}} \ell_{\text{KLIEP}}(\mathcal{A}) + \lambda_1 \|\mathcal{A}\|_1 + \lambda_2 \sum_{v=1}^m \left(\sum_{u=1}^m L_{u,v}^k \right)^{\frac{1}{k}} \\ & \text{subject to: } \mathcal{A} = \mathcal{L} + \mathcal{L}^\top, \end{aligned} \quad (8)$$

Such a regularization can be used to discover perturbed nodes i.e., nodes that have a completely different connectivity pattern to other nodes among two networks.

Optimization Although the original objective of KLIEP was smooth and convex, the sparsity inducing norms are in general non-smooth. Proximal gradient methods, such as Fast Iterative Shrinkage Thresholding Algorithms (FISTA) (Beck and Teboulle 2009) can be utilized to solve regularized KLIEP objectives. A FISTA-like algorithm was proposed in (Fazayeli and Banerjee 2016) with a faster rate of convergence.

3.3 Covariance-precision matching

As mentioned above, the density ratio formulation is not the only way that may motivate the direct modelling. For the formulation using the equality (4), we can solve the following sparsity inducing objective which was introduced in Zhao et al. (2014).

$$\hat{\mathcal{A}} = \underset{\mathcal{A}}{\operatorname{argmin}} \|\mathcal{A}\|_1 \text{ subject to } \|\hat{\Sigma}^{(p)} \mathcal{A} \hat{\Sigma}^{(q)} + \hat{\Sigma}^{(p)} - \hat{\Sigma}^{(q)}\|_\infty \leq \epsilon, \quad (9)$$

where ϵ is a hyper-parameter. To obtain a sparse solution, we set a threshold for the solution at a certain level τ , i.e. the value for $|\hat{\mathcal{A}}_{u,v}| < \tau$ is rounded to 0. The constraint enforces the equality (4) and we used single parameter $\mathcal{A}$ replacing $\Theta^{(p)} - \Theta^{(q)}$.

Optimization This method is quite computationally demanding as the dimension m grows. The Alternating Direction Method of Multipliers (ADMM) procedure (Boyd et al. 2011) was implemented based on an augmented version of (9) (see Section 3.3, Zhao et al. 2014 for details).

3.4 Maximizing joint likelihood

As it was mentioned in Sect. 2.3, one does not have to use the direct modelling to learn sparse changes between MNs. In fact, separated modelling may not only discover changes, but also can recover the individual MN themselves. Recently, a method based on *fused-lasso* (Tibshirani et al. 2005) has been developed (Zhang and Wang 2010). This method also sparsifies $\theta^{(p)} - \theta^{(q)}$ directly.

The original method conducts *feature-wise neighborhood regression* (Meinshausen and Bühlmann 2006) jointly for P and Q , which can be conceptually understood as maximizing the local conditional Gaussian likelihood jointly on each random variable t . A slightly more general form of the learning criterion may be summarized as

$$\begin{aligned} \min_{\theta_t^{(p)} \in \mathbb{R}^{m-1}, \theta_t^{(q)} \in \mathbb{R}^{m-1}} & \ell_t(\theta_t^{(p)}; \mathcal{X}_p) + \ell_t(\theta_t^{(q)}; \mathcal{X}_q) \\ & + \lambda_1(\|\theta_t^{(p)}\|_1 + \|\theta_t^{(q)}\|_1) + \lambda_2\|\theta_t^{(p)} - \theta_t^{(q)}\|_1, \end{aligned} \quad (10)$$

where $\ell_t(\theta; \mathcal{X}_p)$ is the *negative* log conditional likelihood for the t -th random variable $x_t \in \mathbb{R}$ given the rest $\mathbf{x}_{\setminus t} \in \mathbb{R}^{m-1}$:

$$\ell_t(\theta; \mathcal{X}_p) = -\frac{1}{n_p} \sum_{i=1}^{n_p} \log p(x_{p,t}^{(i)} | \mathbf{x}_{p,\setminus t}^{(i)}; \theta),$$

where each dimension of θ corresponds to one of its potential neighborhood. $\ell_t(\theta; \mathcal{X}_q)$ is defined in the same way as $\ell_t(\theta; \mathcal{X}_p)$.

Since the Fused-lasso-based method directly sparsifies the changes in MN structure, it can work well even when each MN is not sparse (when λ_1 is set to 0). *Learning Changes in Sequence* Another recent development (Kolar and Xing 2012) along this line of research assumes the data points are received *sequentially*, i.e., we observe $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(T)}$ over time points $\mathcal{T} = \{1, 2, \dots, T\}$. Suppose $\mathcal{T}$ can be segmented into K disjoint unknown subsets: $\mathcal{T} = \cup_{k \in \{1, \dots, K\}} \mathcal{T}_k$ and $\mathbf{x}_{\mathcal{T}_k} \sim p(\mathbf{x}, \theta^{(\mathcal{T}_k)})$. The task is to segment such a sequence and learn an estimate $\hat{\theta}^{(\mathcal{T}_k)}$ for each segment. We can extend the idea of fused-lasso in (10), and maximize the joint likelihood over each single observation:

$$\operatorname{argmin}_{\theta^{(i)}, i \in \{1, \dots, T\}} \sum_{i=1}^T \ell(\theta^{(i)}; \mathbf{x}^{(i)}) + \lambda_1 \sum_{i=1}^T \|\theta^{(i)}\|_1 + \lambda_2 \sum_{i=1}^{T-1} \|\theta^{(i+1)} - \theta^{(i)}\|_1,$$

where the fused lasso term sparsifies the changes between MNs at adjacency time points, thus the learned $\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(T)}$ is “piecewise-constant” and the segments are automatically determined from it. A block-coordinate descent procedure was proposed to solve this problem efficiently (Kolar et al. 2010).

4 Theoretical analysis

The KLIEP algorithm does not only perform well in practice, but is also justified theoretically. In this section, we first introduce the support recovery theorem of KLIEP and then review some recent theoretical developments of direct change learning.

4.1 Preliminaries

In the previous section, a sub-vector of δ indexed by a pair (u, v) corresponds to a specific edge of an MN. From now on, we switch to a “unitary” index system as our analysis is not dependent on the edge nor the structure setting of the graph.

We introduce the “true parameter” notation $\delta^*, p(\mathbf{x}) = q(\mathbf{x})r(\mathbf{x}; \delta^*)$, and the pairwise index set $E = \{(u, v) | u \geq v\}$. Two sets of *sub-vector indices* regarding δ^*

and E are defined as $S = \{t' \in E \mid \|\delta_{t'}^*\| \neq 0\}$, $S^c = \{t'' \in E \mid \|\delta_{t''}^*\| = 0\}$. We rewrite the objective (7) as

$$\hat{\delta} = \underset{\delta}{\operatorname{argmin}} \ell_{\text{KLIEP}}(\delta) + \lambda_{n_p} \sum_{t \in S \cup S^c} \|\delta_t\|. \quad (11)$$

Similarly we can define $\hat{S} = \{t' \in E \mid \|\hat{\delta}_{t'}\| \neq 0\}$ 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: $\mathcal{I} = \nabla^2 \ell_{\text{KLIEP}}(\delta^*)$. $\mathcal{I}_{AB}$ is a sub-matrix of $\mathcal{I}$ indexed by two sets of indices $A, B \subseteq E$ are indices on rows and columns.

In this section, we prove the support consistency, i.e. with high probability that $S = \hat{S}, S^c = \hat{S}^c$ (see e.g., Chapter 11 in Hastie et al. 2015 for an introduction of support consistency).

4.2 Assumptions

We try *not* to impose assumptions directly on each individual MNs, as the essence of KLIEP method is that it can handle various changes regardless the types of individual MNs.

The first two assumptions are essential to many support consistency theorems (e.g. Eqs. (15) and (16) in Wainwright 2009, Assumption A1 and A2 in Ravikumar et al. 2010). These 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 $1 - \xi_q$, where $\lambda_{\min}$ is the minimum-eigenvalue operator of a symmetric matrix.

This assumption on the *submatrix* of $\mathcal{I}$ is to ensure that the density ratio model is identifiable and the objective function is “reasonably convex”.

Assumption 2 (*Incoherence assumption*) $\max_{t'' \in S^c} \|\mathcal{I}_{t''S} \mathcal{I}_{SS}^{-1}\|_1 \leq 1 - \alpha, 0 < \alpha \leq 1$, 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. Note this assumption is sometimes called “*irrepresentability*” condition.

Assumption 3 (*Smoothness assumption on likelihood ratio*) The log-likelihood ratio $\ell_{\text{KLIEP}}(\delta)$ is smooth around its optimal value, i.e., it has bounded derivatives

$$\begin{aligned} \max_{u, \|u\| \leq \|\delta^*\|} \|\nabla^2 \ell_{\text{KLIEP}}(\delta^* + u)\| &\leq \lambda_{\max} < \infty, \\ \max_{t \in S \cup S^c} \max_{u, \|u\| \leq \|\delta^*\|} \|\nabla_{\delta_t} \nabla^2 \ell_{\text{KLIEP}}(\delta^* + u)\| &\leq \lambda_{3,\max} < \infty, \end{aligned}$$

with probability 1.

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

assumption guarantees the log-likelihood function is well-behaved. Now, we state the following assumptions on the density ratio:

Assumption 4 (*Correct model assumption*) The density ratio model is correct, i.e. there exists δ^* such that

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

Although analysing the mis-specified ratio model (Kanamori et al. 2010) is certainly an interesting open question, we focus on correctly specified models in this section.

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

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

where $M > 0$ is a constant independent from m . This assumption states that the density ratio model, around its optimal parameter, should not often obtain large values over samples from Q .

4.3 Successful support recovery of KLIEP (Liu et al. 2017a, b)

Theorem 1 Suppose that Assumptions 1–5 as well as

$$\min_{t' \in S} \|\delta_{t'}^*\| \geq \frac{10}{\lambda_{\min}} \sqrt{d} \lambda_{n_p} \quad (12)$$

are satisfied, where d is the number of changed edges defined as $d = |S|$, i.e., the cardinality of the set of non-zero parameter groups. Suppose also that the regularization parameter is chosen so that

$$M_1 \sqrt{\frac{\log \frac{m^2+m}{2}}{n_p}} \leq \lambda_{n_p} \leq M_2 \min\left(\frac{\|\delta^*\|}{\sqrt{b}}, 1\right), \quad (13)$$

and $n_q \geq M_3 n_p^2$, where M_1, M_2 and M_3 are constants. 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

$$1 - \exp\left(-K_1 \lambda_{n_p}^2 n_p\right) - 4 \exp\left(-K_2 d n_q \lambda_{n_p}^4\right) - \xi_q,$$

the following properties hold:

- *Unique Solution:* The solution of (11) is unique.
- *Successful Change Detection:* $\hat{S} = S$ and $\hat{S}^c = S^c$.

The proof of this theorem follows the Primal-dual witness construction (see, e.g., Section 11.4.2 in Hastie et al. 2015).

Remark The main conclusion of this theorem states that if the regularization parameter is *reasonably chosen* (13) and the true non-zero groups $\|\delta_{t'}^*\|$, $t' \in S$ is *large enough* (12), with high probability, we are guaranteed to have the correct support of parameters. The samples needed for n_p only grow linearly with $\log m$ and n_q grows quadratically with n_p .

4.4 ℓ_2 Consistency of KLIEP with atomic norm (Fazayeli and Banerjee 2016)

As it was introduced in Sect. 3.2, atomic norms can be used to learn changes with special topological structures. Instead of support recovery, we focus on the ℓ_2 loss between the estimated parameter $\hat{\delta}$ and the true parameter δ^* , i.e., $\|\delta^* - \hat{\delta}\|$.

First, we generalize our objective function as

$$\hat{\delta} = \underset{\delta \in \mathbb{R}^{\frac{m^2+m}{2}}}{\operatorname{argmin}} \ell_{\text{KLIEP}}(\delta) + \lambda_{n_p, n_q} R(\delta), \quad (14)$$

where R is an atomic norm function.

Such a theorem relies on the *Restricted Strong Convex (RSC)* property on the *Error Set* of the objective function. Intuitively, if $\ell_{\text{KLIEP}}(\delta)$ is “highly curved”, small $|\ell_{\text{KLIEP}}(\hat{\delta}) - \ell_{\text{KLIEP}}(\delta^*)|$ ensures small $\|\hat{\delta} - \delta^*\|$. Thus we only need to figure out how $|\ell_{\text{KLIEP}}(\hat{\delta}) - \ell_{\text{KLIEP}}(\delta^*)|$ reaches zero as number of samples goes to infinity and this is a more accessible target.

To make sure our objective has such a “strongly convex” curvature, one needs to impose a uniform lower-bound on the eigenvalues of the objective Hessian (a.k.a., sample Fisher information matrix $\mathcal{I}$). However, this is not realistic for the high-dimensional setting, as $\mathcal{I}$ is certainly rank-deficient. As an alternative, we impose an assumption on the convexity of ℓ_{KLIEP} over a constrained set:

Restricted strong convex condition The function ℓ is Restricted Strong Convex (RSC) at a cone C if there exists a constant κ such that $\forall \mathbf{u} \in C$

$$\ell(\delta^* + \mathbf{u}) - \ell(\delta^*) - \langle \mathbf{u}, \nabla_{\ell}(\delta^*) \rangle \geq \kappa \|\mathbf{u}\|^2.$$

If $\delta^* - \hat{\delta} \in C$, it is possible to obtain a *deterministic* bound (Theorem 2 in Banerjee et al. 2014) on the ℓ_2 estimation error

$$\|\delta^* - \hat{\delta}\|_2 = O\left(\frac{\lambda_{n_p, n_q}}{\kappa} \Psi(C)\right),$$

where $\Psi(C)$ is the norm compatibility constant (Negahban et al. 2009) and can be easily bounded. Note that although this bound itself is not probabilistic, the parameter λ_{n_p, n_q} is random and the RSC may hold with a probability. One can infer the sample complexity from these bounds.

Two things remain to be shown. First, we need to find such a cone which contains $\hat{\delta} - \delta^*$. Second, we need to prove ℓ_{KLIEP} is RSC on this cone. We start with the first problem.

Error set (Lemma 1 in Banerjee et al. 2014) For any convex loss $\ell(\delta)$, if λ_{n_p, n_q} is large enough, i.e.,

$$\lambda_{n_p, n_q} \geq \beta R^*(\nabla \ell(\delta^*)), \quad \beta > 1$$

where R^* is the dual norm of R , it can be proven that the estimation error $\mathbf{u} = \delta^* - \hat{\delta}$ lies in an Error Set:

$$E_r = \left\{ \mathbf{u}, \mathbf{u} \in \text{dom}(\delta) \mid R(\delta^* + \mathbf{u}) \leq R(\delta^*) + \frac{1}{\beta} R(\mathbf{u}) \right\},$$

where $\text{dom}(\mathbf{y})$ is the domain of $\mathbf{y}$. Let us define $C_r = \text{cone}(E_r)$.

In fact, it can be shown that if

$$\lambda_{n_p, n_q} \geq \frac{c \cdot (w(\Omega_R) + \epsilon)}{\sqrt{\min(n_p, n_q)}},$$

where $w(A)$ is the Gaussian width of a set A (Ledoux and Talagrand 2013) and $\Omega_R = \{\mathbf{u} \mid R(\mathbf{u}) \leq 1\}$, then $\lambda_{n_p, n_q} \geq \beta R^*(\nabla \ell(\delta^*))$ holds automatically with high probability (Theorem 1 in Fazayeli and Banerjee 2016). Now we have a cone C_r where $\hat{\delta} - \delta^*$ resides.

As to the second problem, it can be proven that ℓ_{KLIEP} is RSC at C_r with high probability once $n_q \geq n_0$, $n_0 = w^2(C_r \cap S)$, where S is a unit hypersphere (Theorem 2 in Fazayeli and Banerjee 2016). Thus n_0 is the minimum number of samples required from Q to be able to apply this theorem.

Putting everything together, we have the main theorem proved in Fazayeli and Banerjee (2016):

Theorem 2 (ℓ_2 Consistency of atomic norms) *If Assumption 5 holds, and $\hat{\delta}$ is the minimizer of (14), then with probability at least $1 - M_1 \exp(-\epsilon^2)$ the followings hold::*

$$\lambda_{n_p, n_q} \geq \frac{M_2}{\sqrt{\min(n_p, n_q)}} (w(\Omega_R) + \epsilon)$$

and for $n_q \geq c_1 w^2(C_r \cap S)$, with high probability, the estimate $\hat{\delta}$ satisfies

$$\|\hat{\delta} - \delta^*\|_2 = O\left(\frac{w(\Omega_R)}{\sqrt{\min(n_p, n_q)}}\right) \Psi(C_r)$$

Note the constants M_1 and M_2 listed in this theorem are not the same as the ones in Theorem 1. To apply this theorem, we need to know the bounds of $w(\Omega_R)$ and $\Psi(C_r)$ for specific R norms. These bounds have been proven in previous literatures (see, e.g. Banerjee et al. 2014). For example, if R is ℓ_1 norm, then $\Psi(C_r) \leq 4\sqrt{d}$ and $w(\Omega_R) \leq c \log m$ so applying the above theorem, we have

$$\|\hat{\delta} - \delta^*\|_2 = O\left(\sqrt{\frac{d \log m}{\min(n_p, n_q)}}\right).$$

Remark Although this bound does not directly prove the support consistency, we can learn that sample complexity $\min(n_p, n_q) = \Omega(d \log m)$ guarantees the convergence of estimation error in ℓ_2 norm. As to n_q , it should also satisfy $n_q \geq c_1 w^2 (C_r \cap \mathcal{S})$, which is again $n_q = \Omega(d \log m)$ in the case of ℓ_1 norm. This sample complexity is milder than what Liu et al. have obtained in the previous section $\Omega(d^2 \log(m^2 + m)/2)$ and $n_q = \Omega(n_p^2)$. Nonetheless, both theories can be applied to high-dimensional regime $m \gg \min(n_p, n_q)$.

4.5 Support consistency of covariance-precision matching (Zhao et al. 2014)

In this section, we introduce the support recovery theorem of the Covariance-Precision Matching method (9) in terms of support consistency on Gaussian MNs. Specifically for Gaussian MNs, we need a slightly different set of notations, as they are parametrized in matrix forms. $\Sigma_{j,k}^{(p)}$ is the j, k -th elements of matrix $\Sigma_{j,k}^{(p)}$ and $\Sigma_{\max}^{(p)}$ is $\max_j \Sigma_{j,j}^{(p)}$.

The first assumption is to ensure that the “amount of change” is fixed and the change is always sparse, and does not grow with the number of dimension m .

Assumption 6 The difference matrix $\mathbf{A}$ has $d \leq m$ non-zero elements in its upper triangular sub-matrix. $|\mathbf{A}|_1 \leq M_0$, and both d and M_0 do not depend on dimension m .

The second assumption assures that the covariates are not strongly dependent if there are many changes in the precision matrix. This is similar to the incoherence assumption used in Assumption 2.

Assumption 7 The constants $\mu^{(p)} = \max_{j \neq k} |\Sigma_{j,k}^{(p)}|$ and $\mu^{(q)} = \max_{j \neq k} |\Sigma_{j,k}^{(q)}|$ must satisfy $\mu = 4 \max(\mu^{(p)} \Sigma_{\max}^{(q)}, \mu^{(q)} \Sigma_{\max}^{(p)}) \leq \frac{\Sigma_{\min}^S}{2d}$, where

$$\Sigma_{\min}^S = \min_{j,k} \left(\Sigma_{jj}^{(q)} \Sigma_{jj}^{(p)}, \Sigma_{kk}^{(q)} \Sigma_{jj}^{(p)} + 2 \Sigma_{kj}^{(q)} \Sigma_{jk}^{(p)} + \Sigma_{jj}^{(q)} \Sigma_{kk}^{(p)} \right).$$

We first intuitively explain how the proof works. The proof of the support consistency can be thought as controlling $\|\hat{\mathbf{A}} - \mathbf{A}^*\|_\infty$. Clearly, for the population covariance matrices $\Sigma^{(p)}$ and $\Sigma^{(q)}$, $\Sigma^{(p)} \mathbf{A}^* \Sigma^{(q)} + \Sigma^{(p)} - \Sigma^{(q)} = \mathbf{0}$. If we replace the above population covariances with their sample versions, we can expect $\|\hat{\Sigma}^{(p)} \mathbf{A}^* \hat{\Sigma}^{(q)} + \hat{\Sigma}^{(p)} - \hat{\Sigma}^{(q)}\|_\infty \leq \epsilon$, if the number of samples is large enough. Furthermore, ϵ can be a function decreasing with $\min(n_p, n_q)$ as the estimated covariances are getting closer and closer to the population ones.

Therefore, if we set the ϵ to a decreasing function, we can still “contain” the optimal parameter Δ^* in the feasible zone with high probability. By definition, the estimated difference $\hat{\Delta}$ should also be in the feasible zone; thus they should not be far off, if the zone is small enough. The rigorous proof of the above statements is given in the Appendix of Zhao et al. (2014).

Now, we give the support recovery theorem³ as follows (see Section 4 in Zhao et al. 2014 for details):

Theorem 3 (Support consistency of covariance-precision matching) *Suppose P and Q are Gaussian, Assumptions 6 and 7 hold, $\min(n_p, n_q) \geq \log m$ and*

$$\tau_{n_p, n_q} = \Omega\left(\sqrt{\frac{\log m}{\min(n_p, n_q)}}\right), \epsilon_{n_p, n_q} = M_1 \cdot \sqrt{\frac{\log m}{\min(n_p, n_q)}}$$

and $\min_{j,k} |\Delta_{j,k}^| \geq 2\tau_{n_p, n_q}$,⁴ then with high probability, (9) can recover the correct support of Δ^* .*

This support consistency theorem, although only applies to Gaussian MNs, has similar structure to the one derived for KLIEP (Sect. 4.3). First, they both assume the true non-zero parameter should be large enough. Second, they both assume the sparsity inducing factor (λ_{n_p, n_q} and τ_{n_p, n_q}) should decay as the sample size $\min(n_p, n_q)$ increases, while increase as the log-dimension $\log m$ increases.

4.6 Summary and discussion

Now, we summarize and compare these theoretical results. First we discuss the similarities of these theorems.

- None of the above proofs require the sparsity assumption on each individual MN. Thus in theory, all methods should work well when individual MNs are dense.
- The efficiency of all methods are affected by the sparsity of changes (i.e. d). This make sense since the sparsity assumption is made on the changes between two MNs.
- All theorems apply to the high-dimensional regime ($m \gg \min(n_p, n_q)$). None requires n_p or n_q to be comparable to the dimensionality m .

However, there is one important difference among these theorems. The sample complexities introduced in Sects. 4.3 and 4.4 are *not* symmetric; the sample complexity of n_q is more restrictive comparing to that of n_p . This is understandable since KLIEP itself is an asymmetric method (KL divergence is asymmetric). In comparison, the sample complexity of Covariance-Precision Matching is symmetric, i.e., the theorem does not show the “bias” toward either of the datasets. Thus, if

³ In fact, the support recovery theorem was proved for a slightly augmented version of (9).

⁴ $\tau_{n_p, n_q}, \epsilon_{n_p, n_q}$ is the sample-dependent version of τ, ϵ introduced in Sect. 3.3.

one has perfectly balanced Gaussian datasets, it might be natural to use Covariance-Precision Matching to learn the differences.

5 Experiments

In this section, we compare the performance of two direct change detection methods: KLIEP and Covariance-Precision (CP) Matching using synthetic and real-world examples.

5.1 Implementations

Sparsity-inducing KLIEP can be implemented using sub-gradient descent approach. The MATLAB[®] code can be found at http://www.ism.ac.jp/~liu/kliep_sparse/demo_sparse.html.

The R (R Core Team 2016) implementation of CP matching using ADMM can be obtained at <https://github.com/sdzhao/dpm>.

5.2 Synthetic examples

Illustrative example Now we illustrate the performance of both KLIEP and CP matching using two 50-dimensional multivariate zero-mean Gaussian distributions. First, we randomly generate a 50×50 symmetric adjacency matrix $A^{(P)}$ with 10% connectivity and draw 500 samples from a Gaussian distribution with the following precision matrix:

$$\Theta_{ij}^{(P)} = \begin{cases} 2 & i = j \\ 0.4 & A_{ij}^{(P)} \neq 0, i \neq j \end{cases} \quad (15)$$

Then we remove six edges randomly from it, resulting a change pattern shown in Fig. 1a and use it as $A^{(Q)}$. Following the same step above we construct $\Theta^{(q)}$ and generate 500 samples again. As it was suggested by Theorem 1, we set $\lambda = \frac{\alpha \log 50}{500}$, and the learned $\hat{A}^5$ are shown in Fig. 1c–e using different α .

The same experiments are repeated using the CP matching method. However, since the sparsity control of CP matching is via the selection of the threshold τ , we set $\epsilon = 0.2$ which shows good performance empirically and plot the learned $\hat{A}$ using different thresholds. Results are shown in Fig. 1f–h.

As we can see, both approaches recover the change pattern well as we increase the sparsity control parameter.

ROC-curves In this experiment, we compare two methods quantitatively using ROC curves. We adopt the True Positive (TP) and True Negative (TN) rate as described in Zhao et al. (2014):

⁵ We convert $\hat{\delta}$ into its corresponding matrix form.

Fig. 1 Illustrative experiments

$$\text{TPR} = \frac{\sum_{t' \in S} \delta(\hat{\delta}_{t'} \neq \mathbf{0})}{\sum_{t' \in S} \delta(\delta_{t'}^* \neq \mathbf{0})}, \quad \text{TNR} = \frac{\sum_{t'' \in S^c} \delta(\hat{\delta}_{t''} = \mathbf{0})}{\sum_{t'' \in S^c} \delta(\delta_{t''}^* = \mathbf{0})}.$$

We generate an adjacency matrix $A^{(P)}$ with four-neighbour lattice structure and randomly remove $d = \sqrt{m}$ edges producing $A^{(Q)}$. Two sets of $n_p = n_q = 50$ samples are generated using the same criteria mentioned in (15). The ROC curves averaged over 50 trials with different dimensions are shown in Fig. 1b, and the AUCs are reported in Table 1.

It can be seen that as the number of both dimension and changed edges increases, KLIEP method can retain stable performance while the performance of CP approach decays rapidly.

Table 1 The area under the curve (AUC) of ROC plot in Fig. 1b (“K” for KLIEP and “CP” for CP matching)

	$m = 9$	$m = 16$	$m = 25$	$m = 36$	$m = 49$	$m = 64$	$m = 81$	$m = 100$
K	0.8746	0.8865	0.8899	0.8890	0.8902	0.8878	0.8903	0.8866
CP	0.8165	0.7917	0.7627	0.6829	0.5574	0.5914	0.5475	0.5656

5.3 Running time

Although the rigorous timing comparison is difficult due to the different implementations of KLIEP and CP matching, from our experience, KLIEP is faster but more memory-consuming as our implementation stores the entire parameter vector into the memory. On a server with 16 Xeon cores, it takes KLIEP about 15 min to run experiments needed for plotting Fig. 1b, while it takes CP matching roughly 1 h.

As to KLIEP, we also observe that “early stopping” heuristics (e.g., stopping at 100 iterations) can provide an accurate non-zero pattern within a short period of time.

5.4 Image difference detection

Two photos were taken in a rainy afternoon using a camera pointing at the parking lot of The Institute of Statistical Mathematics (ISM). In this task, we are interested in learning the changes of the parking patterns marked by green boxes in Fig. 2(b). As we can see from Fig. 2a, b, the light conditions and positions of raindrops vary in two pictures.

To construct samples, we use windows of pixels (Fig. 2c). Each window is a dimension of a dataset, and the samples are the pixel RGB values within this window. By sliding the window across the entire picture, we may obtain samples of different dimensions. Two sets of data can be obtained by using this sample generating mechanism over two images.

Assuming an image can be represented by an MN of windows, changes of pixels values within a window may cause changes of “interactions” between neighbouring windows. In other words, we can discover a change by looking at the change of the dependency of pixel values between a certain window and its neighbours. This is more advantageous than simply looking at the pixel values since changing the brightness of a picture may increase the pixel values in many windows simultaneously, even if the “contrast” between two windows does not change by much.

By applying KLIEP on such two sets of data and highlighting adjacent window pairs that are involved in the changes of pairwise interactions, we may spot changes between two images. In our experiment, we use sliding windows of size 16×16 on a 200×150 image, generating two sets of samples with $m = 999$ and $n_p = n_q = 256$. We reduce λ until $|\hat{S}| > 40$. The spotted changes were plotted in

Fig. 2 Detecting changes of parking patterns from two photos

Fig. 2d. It can be seen that KLIEP has correctly labelled almost all changed parkings between two images except one missing on the left.

Note that here we set $\psi(\mathbf{x}_u, \mathbf{x}_v) = \exp\left(-\frac{\|\mathbf{x}_u - \mathbf{x}_v\|^2}{0.5}\right)$, and the underlying MN is highly non-Gaussian so CP matching cannot be applied here.

6 Open problems

Although pioneering works have been conducted in this area, there are still important unsolved open problems. In this section, we list a few examples.

Generalized covariance-precision matching In Sect. 3.3, we introduced an equality between Gaussian covariance and precision matrix (4). This leads to a *direct* sparse change learning approach. However, it does not apply to more general pairwise MNs. A natural question is, can we extend this relationship between covariance and precision matrices to a more general principle? Particularly, in a recent work (Loh and Wainwright 2013), the *generalized covariance matrix* was used to learn a non-

Gaussian graphical model structure. Would a generalized equality (4) provide us with a universal framework of learning changes between MNs?

Learning changes from multiple MNs In this paper, we have only reviewed the algorithms that learn changes between two MNs. In fact, in some applications, datasets may be obtained as multiple “snapshots”. For example, gene activities may be measured at a few different time points. Under the same assumption that changes between adjacent time points are “mild” and “sparse”, can we perform multiple change detections in one shot?

Asymmetry versus symmetry As we have pointed out in Sect. 4.6, there exists an asymmetry in KLIEP while Covariance-Precision matching has a symmetric formulation. An interesting future direction is to systematically investigate how such an asymmetry affects the change detection results, and more importantly, how can we automatically determine which density to be Q and which one to be P in the ratio formulation.

We believe thorough investigations in these three directions will significantly expand our knowledge over the domain of learning changes between MNs in the future.

7 Conclusion

In this paper, we have reviewed an MN change learning method based on density ratio estimation and other alternative approaches. Statistical guarantees regarding the support recovery and ℓ_2 consistency were also reported and compared. Through their direct modelling and theoretical results, we can see an interesting common pattern in all these methods: they work well regardless of the difficulty of learning individual MNs.

These results are inspiring as they shed lights on a new family of methods that only learn the *incremental* patterns. They show that if the change itself is simple enough, even with limited amount of information, we can have good learning performance. Compared to classic, *static* pattern recognition, such methods are well-suited for analysing *dynamic* datasets, where the “absolute” pattern is not the main interest, but learning the change itself is more valuable.

These works have offered a new vision of research on learning changes between patterns. We believe these methods and theorems may have many potential applications in the years to come.

Acknowledgements We would like to thank Masashi Sugiyama, John Quinn and Michael Gutmann for their tremendous help during the development of the density ratio modelling idea. This work was partially supported by JSPS Grant-in-Aid for Scientific Research Activity Start-up 15H06823, MEXT kakenhi (25730013, 25120012, 26280009, 15H01678 and 15H05707), and JST-PRESTO. Authors would like to thank anonymous reviewers for their helpful comments. We would like to thank anonymous reviewers and Matthew Ames for their helpful comments and suggestions on this paper.

Compliance with ethical standards

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Banerjee O, El Ghaoui L, d’Aspremont A (2008) Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *J Mach Learn Res* 9:485–516
- Banerjee A, Chen S, Fazayeli F, Sivakumar V (2014) Estimation with norm regularization. *Adv Neural Inf Process Syst* 26:1556–1564
- Beck A, Teboulle M (2009) A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J Imaging Sci* 2(1):183–202
- Boyd S, Vandenberghe L (2004) *Convex optimization*. Cambridge University Press, Cambridge
- Boyd S, Parikh N, Chu E, Peleato B, Eckstein J (2011) Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found Trends® Mach Learn* 3(1):1–122
- Chandrasekaran V, Recht B, Parrilo PA, Willsky AS (2012) The convex geometry of linear inverse problems. *Found Comput Math* 12(6):805–849
- Chickering DM (1996) Learning Bayesian networks is NP-complete. In: *Learning from data*. Springer, Berlin, pp 121–130
- Chow C, Liu C (1968) Approximating discrete probability distributions with dependence trees. *IEEE Trans Inf Theory* 14(3):462–467
- Fazayeli F, Banerjee A (2016) Generalized direct change estimation in ising model structure. In: *Proceedings of the 33rd international conference on machine learning*, pp 2281–2290. <http://jmlr.org/proceedings/papers/v48/fazayeli16.html>
- Friedman J, Hastie T, Tibshirani R (2008) Sparse inverse covariance estimation with the graphical Lasso. *Biostatistics* 9(3):432–441
- Geman S, Geman D (1984) Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans Pattern Anal Mach Intell* 6:721–741
- Hammersley JM, Clifford P (1971) Markov fields on finite graphs and lattices (unpublished)
- Hastie T, Tibshirani R, Wainwright M (2015) *Statistical learning with sparsity: the Lasso and generalizations*. CRC Press, Boca Raton
- Kanamori T, Suzuki T, Sugiyama M (2010) Theoretical analysis of density ratio estimation. *IEICE Trans Fundam Electron Commun Comput Sci* E93-A(4):787–798
- Kolar M, Xing EP (2012) Estimating networks with jumps. *Electron J Stat* 6:2069–2106
- Kolar M, Song L, Ahmed A, Xing EP (2010) Estimating time-varying networks. *Ann Appl Stat* 4(1):94–123
- Koller D, Friedman N (2009) *Probabilistic graphical models: principles and techniques*. MIT Press, Cambridge
- Ledoux M, Talagrand M (2013) *Probability in Banach spaces: isoperimetry and processes*. Springer Science & Business Media, Berlin
- Liu H, Lafferty J, Wasserman L (2009) The nonparanormal: semiparametric estimation of high dimensional undirected graphs. *J Mach Learn Res* 10:2295–2328
- Liu H, Xu M, Gu H, Gupta A, Lafferty J, Wasserman L (2011) Forest density estimation. *J Mach Learn Res* 12(Mar):907–951
- Liu H, Han F, Yuan M, Lafferty J, Wasserman L (2012) The nonparanormal skeptic. In: *Proceedings of the 29th international conference on machine learning (ICML2012)* (accepted)
- Liu S, Quinn JA, Gutmann MU, Suzuki T, Sugiyama M (2014) Direct learning of sparse changes in Markov networks by density ratio estimation. *Neural Comput* 26(6):1169–1197
- Liu S, Suzuki T, Relator R, Sese J, Sugiyama M, Fukumizu K (2017a) Supplement to “support consistency of direct sparse-change learning in Markov networks” (accepted)
- Liu S, Suzuki T, Relator R, Sese J, Sugiyama M, Fukumizu K (2017b) Support consistency of direct sparse-change learning in Markov networks. *Ann Stat* (accepted)
- Loh P-L, Wainwright MJ (2013) Structure estimation for discrete graphical models: generalized covariance matrices and their inverses. *Ann Stat* 41(6):3022–3049

- Meinshausen N, Bühlmann P (2006) High-dimensional graphs and variable selection with the Lasso. *Ann Stat* 34(3):1436–1462
- Mohan K, London P, Fazel M, Witten DM, Lee S (2014) Node-based learning of multiple gaussian graphical models. *J Mach Learn Res* 15(1):445–488
- Negahban S, Yu B, Wainwright MJ, Ravikumar PK (2009) A unified framework for high-dimensional analysis of m -estimators with decomposable regularizers. *Adv Neural Inf Process Syst* 21:1348–1356
- R Core Team (2016) R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna. <https://www.R-project.org/>
- Ravikumar P, Wainwright MJ, Lafferty JD (2010) High-dimensional Ising model selection using ℓ_1 -regularized logistic regression. *Ann Stat* 38(3):1287–1319
- Robert CP, Casella G (2005) Monte Carlo statistical methods. Springer, Berlin
- Spirtes P, Glymour CN, Scheines R (2000) Causation, prediction, and search. MIT Press, Cambridge
- Sugiyama M, Nakajima S, Kashima H, von Bünau P, Kawanabe M (2008) Direct importance estimation with model selection and its application to covariate shift adaptation. In: *Advances in neural information processing systems*, vol 20, pp 1433–1440
- Sugiyama M, Suzuki T, Kanamori T (2012) Density ratio estimation in machine learning. Cambridge University Press, Cambridge
- Tibshirani R (1996) Regression shrinkage and selection via the Lasso. *J R Stat Soc Ser B (Methodological)* 58(1):267–288
- Tibshirani R, Saunders M, Rosset S, Zhu J, Knight K (2005) Sparsity and smoothness via the fused Lasso. *J R Stat Soc Ser B (Stat Methodol)* 67(1):91–108
- Tomioka R, Suzuki T (2014) Spectral norm of random tensors. arXiv preprint [arXiv:1407.1870 \[math.ST\]](https://arxiv.org/abs/1407.1870)
- Tsuboi Y, Kashima H, Hido S, Bickel S, Sugiyama M (2009) Direct density ratio estimation for large-scale covariate shift adaptation. *J Inf Process* 17:138–155
- Wainwright MJ (2009) Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (Lasso). *IEEE Trans Inf Theory* 55(5):2183–2202
- Yuan M, Lin Y (2006) Model selection and estimation in regression with grouped variables. *J R Stat Soc Ser B (Stat Methodol)* 68(1):49–67
- Zhang B, Wang YJ (2010) Learning structural changes of Gaussian graphical models in controlled experiments. In: *Proceedings of the twenty-sixth conference on uncertainty in artificial intelligence (UAI2010)*, pp 701–708
- Zhao P, Yu B (2006) On model selection consistency of Lasso. *J Mach Learn Res* 7:2541–2563
- Zhao S, Cai T, Li H (2014) Direct estimation of differential networks. *Biometrika* 101(2):253–268