
Imputation Estimators for Unnormalized Models with Missing Data

Masatoshi Uehara
Harvard University

Takeru Matsuda
The University of Tokyo
RIKEN Center for Brain Science

Jae Kwang Kim
Iowa State University

Abstract

Several statistical models are given in the form of unnormalized densities, and calculation of the normalization constant is intractable. We propose estimation methods for such unnormalized models with missing data. The key concept is to combine imputation techniques with estimators for unnormalized models including noise contrastive estimation and score matching. In addition, we derive asymptotic distributions of the proposed estimators and construct confidence intervals. Simulation results with truncated Gaussian graphical models and the application to real data of wind direction reveal that the proposed methods effectively enable statistical inference with unnormalized models from missing data.

1 INTRODUCTION

Several statistical models are given in the form of unnormalized densities, and the calculation of the normalization constant (or partition function) is intractable. Namely, a statistical model is defined as

$$p(x; \theta) = \frac{1}{Z(\theta)} \tilde{p}(x; \theta), \quad (1)$$

where $Z(\theta) = \int \tilde{p}(x; \theta) \mu(dx)$ is the normalization constant, μ is a base measure such as the Lebesgue measure or counting measure, and we only have access to $\tilde{p}(x; \theta)$. Such unnormalized models are widely used in many settings: Markov random fields (Besag, 1975), directional statistics (Mardia and Jupp, 1999), Boltzmann machines (Hinton, 2002), overcomplete independent component analysis models (Hyvärinen et al., 2001), and graphical models (Lin et al., 2016; Yu et al., 2016). Several methods for estimating θ without computing the normalizing constant $Z(\theta)$ have been proposed, such as noise con-

trastive estimation (NCE; Gutmann and Hyvärinen, 2012) and score matching (Hyvärinen, 2005).

In practice, we frequently encounter data with missing values, which is called missing data or incomplete data (Tsiatis, 2006; Kim and Shao, 2013). Missing data must be handled properly; otherwise incorrect estimates may be obtained such as nonresponse bias (Little and Rubin, 2002). However, existing estimation methods for unnormalized models are not applicable to missing data, because they assume that the data is fully observed.

In this study, we develop estimation methods for unnormalized models with missing data. The proposed methods utilize NCE and score matching by imputing missing data with importance weights. This method is computationally fast because it does not rely on any sampling techniques. We derive asymptotic distributions of the proposed estimators and construct confidence intervals. On the way, we also discuss how to incorporate multiple imputation (Meng, 1994) and contrastive divergence method (Hinton, 2002).

Note that Rhodes and Gutmann (2019) proposed an estimation method called variational NCE for unnormalized latent variable models, which corresponds to a special case of the current problem (missing completely at random, MCAR). Although variational inference is efficient and useful for large-scale problems, it is not clear how to construct its confidence intervals (Blei et al., 2017). In contrast, the proposed methods are valid under general missing mechanisms, including missing at random (MAR) and missing not at random (MNAR) cases. In addition, the proposed methods allow for the construction of confidence intervals based on asymptotic theory.

Our main contributions are as follows.

- We propose imputation estimators for unnormalized models with missing data. These estimators are \sqrt{n} -consistent under general missing mechanisms, including MNAR, and computationally efficient.
- We derive asymptotic distributions of the proposed estimators and construct confidence intervals.
- We confirm the validity of the proposed methods by

simulation with truncated Gaussian graphical models and apply the proposed methods to analyze real data of wind direction with the bivariate circular model.

2 PRELIMINARY

2.1 Notations

Parameters with a zero in the subscript such as θ_0 and τ_0 denote true parameters. The notation ∇_θ denotes differentiation with respect to θ , and $t(x)^{\otimes 2} = t(x)t(x)^\top$. The notation \xrightarrow{d} denotes a weak convergence. The expectation and variance of $f(x)$ under density $g(x)$ are denoted as $E_g[f(x)]$ and $\text{var}_g[f(x)]$, respectively. The subscript and argument are often omitted when they are clear from the context. A summary of the notation is provided in Appendix A.

2.2 Estimation Methods for Unnormalized Models

Several methods have been developed for estimating the unnormalized model (1) such as noise contrastive estimation (NCE; Gutmann and Hyvärinen, 2012), score matching (Hyvärinen, 2005), and Monte Carlo maximum likelihood estimation (MC-MLE; Geyer, 1994).

2.2.1 Noise Contrastive Estimation (NCE)

In NCE, the unnormalized model (1) is rewritten to a one-parameter extended model $q(x; \tau) = \exp(-c)\tilde{p}(x; \theta)$, where $\tau = (c, \theta)$ and the true value of c is $c = \log Z(\theta)$.

In addition to data samples $\mathbf{x} = \{x_i\}_{i=1}^n$ from the unnormalized model (1), we generate noise samples $\mathbf{y} = \{y_j\}_{j=1}^{n'}$ from a noise distribution with density $a(y)$. Just for simplicity, we set $n' = n$ in the following.

Let $r(z; \tau) = q(z; \tau)/a(z)$ be the density ratio. Then, in NCE, the estimator is defined as the maximizer of the following function with respect to τ ;

$$\sum_{i=1}^n \log \frac{r(x_i; \tau)}{r(x_i; \tau) + 1} + \sum_{j=1}^n \log \frac{1}{r(y_j; \tau) + 1}. \quad (2)$$

This objective function is interpreted as the negative log-likelihood of the naive Bayes classifier. Regarding more intuitive explanations, see Gutmann and Hyvärinen (2012). NCE gives a \sqrt{n} -consistent estimator under mild regularity conditions.

NCE can be generalized from the divergence perspective (Pihlaja et al., 2010; Gutmann and Hirayama, 2011). Let $g(x)$ be the true data distribution and consider the Bregman divergence

$$D_f(g(x), q(x; \tau)) = \int \text{Br}_f \left(\frac{g(x)}{a(x)}, \frac{q(x; \tau)}{a(x)} \right) a(x) \mu(dx),$$

where f is a twice differentiable strictly convex function and $\text{Br}_f(u, v) = f(u) - f(v) - f'(v)(u - v)$. By subtracting a term independent of τ from $D_f(g, q(x; \tau))$, the cross entropy $d_f(g(x), q(x; \tau))$ between $g(x)$ and $q(x; \tau)$ is obtained as

$$E_{g(x)} [m_{nc1}(x; \tau)] + E_{a(y)} [m_{nc2}(y; \tau)],$$

where

$$\begin{aligned} m_{nc1}(x; \tau) &= -f'(r(x; \tau)), \\ m_{nc2}(y; \tau) &= f'(r(y; \tau))r(y; \tau) - f(r(y; \tau)). \end{aligned}$$

Then, the generalized NCE is defined as the minimizer of $M_{nc1}(\mathbf{x}; \tau) + M_{nc2}(\mathbf{y}; \tau)$ with respect to τ , where $M_{nc1}(\mathbf{x}) = n^{-1} \sum_{i=1}^n m_{nc1}(x_i; \tau)$ and $M_{nc2}(\mathbf{y}) = n^{-1} \sum_{j=1}^n m_{nc2}(y_j; \tau)$. By differentiation with respect to τ , the estimator is also given by the solution to $Z_{nc1}(\mathbf{x}; \tau) + Z_{nc2}(\mathbf{y}; \tau) = 0$, where $Z_{nc1}(\mathbf{x}; \tau) = n^{-1} \sum_{i=1}^n z_{nc1}(x_i; \tau)$, $Z_{nc2}(\mathbf{y}; \tau) = n^{-1} \sum_{j=1}^n z_{nc2}(y_j; \tau)$ and

$$\begin{aligned} z_{nc1}(x; \tau) &= -\nabla_\tau \log q(x; \tau) f''(r(x; \tau)) r(x; \tau), \\ z_{nc2}(y; \tau) &= \nabla_\tau \log q(y; \tau) f''(r(y; \tau)) r(y; \tau)^2. \end{aligned}$$

The original NCE (2) corresponds to $f(x) = x \log x - (1+x) \log(1+x)$ and it is optimal in terms of asymptotic variance (Uehara et al., 2018). On the other hand, when $f(x) = x \log x$, the estimator is given by the minimizer of

$$-\frac{1}{n} \sum_{i=1}^n \log q(x_i; \tau) + \frac{1}{n} \sum_{j=1}^n r(y_j; \tau),$$

which is essentially identical to MC-MLE (Geyer, 1994) by profiling-out c .

2.2.2 Score Matching

Score matching was originally developed as a general estimation method for the unnormalized model (1) on \mathbb{R}^d . Let $c(x; \theta) = \nabla_x \log \tilde{p}(x; \theta) \in \mathbb{R}^d$ and denote the s -th coordinate of x by x^s . For data samples $\mathbf{x} = \{x_i\}_{i=1}^n$, the score matching estimator of θ is defined as the minimizer of $M_{sc}(\mathbf{x}; \theta) = n^{-1} \sum_{i=1}^n m_{sc}(x_i; \theta)$, where

$$\begin{aligned} m_{sc}(x; \theta) &= \frac{1}{2} \sum_{s=1}^d c_s(x; \theta)^2 + \sum_{s=1}^d \frac{\partial c_s(x; \theta)}{\partial x^s}, \\ c_s(x; \theta) &= \frac{\partial}{\partial x^s} \log \tilde{p}(x; \theta). \end{aligned}$$

Note that this estimator is also given by the solution to $Z_{sc}(\mathbf{x}; \theta) = 0$, where $Z_{sc}(\mathbf{x}; \theta) = n^{-1} \sum_{i=1}^n z_{sc}(x_i; \theta)$ and $z_{sc}(x; \theta) = \nabla_\theta m_{sc}(x; \theta)$.

Hyvärinen (2007) extended score matching to unnormalized models on $\mathbb{R}_+^d = [0, \infty)^d$. The estimator is defined

as the minimizer of $M_{sc+}(\mathbf{x}; \theta) = n^{-1} \sum_{i=1}^n m_{sc+}(x_i; \theta)$, where $m_{sc+}(x; \theta)$ is given by

$$\sum_{s=1}^d \left\{ 2x^s c_s(x; \theta) + (x^s)^2 c_s(x; \theta)^2 + (x^s)^2 \frac{\partial c_s(x; \theta)}{\partial x^s} \right\}.$$

2.3 Missing Data and Imputation Methods

We briefly review the framework of missing data and imputation methods. For more details, see Kim and Shao (2013).

Suppose that $\{x_i\}_{i=1}^n$ are independently and identically distributed (i.i.d.) samples from a distribution with density $p(x; \theta)$. We consider the situation where some part of x_i may be missing. Let $\{\delta_i\}_{i=1}^n$ be the missing indicators. Accordingly, $x_i = (x_{i,\text{obs}}, x_{i,\text{mis}})$ is fully observed when $\delta_i = 1$, while only $x_{i,\text{obs}}$ is observed and $x_{i,\text{mis}}$ is missing when $\delta_i = 0$. We assume that δ_i follows a Bernoulli distribution with probability $\Pr(\delta_i = 1 | x_i)$. The case with several missing patterns, that is, the case where the dimension of $x_{i,\text{obs}}$ may differ with i , can be easily considered by extending this notation (Seaman et al., 2013). For more details, see Appendix C.

The missing mechanism is called missing at random (MAR) if $\Pr(\delta = 1 | x) = \Pr(\delta = 1 | x_{\text{obs}})$ holds. Importantly, the missing process can be ignored for estimation of θ in MAR cases (Little and Rubin, 2002), because

$$\begin{aligned} p(x_{\text{obs}}; \theta) &= \int p(x_{\text{obs}}, x_{\text{mis}}; \theta) \Pr(\delta | x) \mu(dx_{\text{mis}}) \\ &\propto \int p(x_{\text{obs}}, x_{\text{mis}}; \theta) \mu(dx_{\text{mis}}). \end{aligned}$$

As a special case of MAR, a missing mechanism is referred to as missing completely at random (MCAR) if $\Pr(\delta = 1 | x)$ is independent of x . When MAR does not hold, the missing mechanism is referred to as missing not at random (MNAR).

For estimating θ from missing data, the fundamental algorithm is the expectation maximization (EM) algorithm (Dempster et al., 1977), which maximizes the observed likelihood $p(x_{\text{obs}}; \theta)$. Equivalently, the EM algorithm solves the following observed (mean) score equation with respect to θ (Louis, 1982):

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E} [\nabla_{\theta} \log p(x_i; \theta) | x_{i,\text{obs}}; \theta] = 0. \quad (3)$$

However, the EM algorithm requires a closed-form expression of the conditional expectation in (3), which is often intractable. To overcome this obstacle, a method called fractional imputation (FI) has been proposed (Kim, 2011; Yang and Kim, 2016), which is closely connected with the Monte Carlo EM algorithm (Wei and Tanner, 1990). FI is computationally efficient because it uses importance sampling

and does not rely on MCMC. Another method called multiple imputation (MI) is also commonly used, which utilizes MCMC (Rubin, 1987; Murray, 2018).

3 FINCE and FISCORE

We propose two estimation methods for unnormalized models with missing data: FINCE (fractional imputation noise contrastive estimation) and FISCORE (fractional imputation score matching).

In this section, we focus on the MAR case, that is, $\Pr(\delta = 1 | x) = \Pr(\delta = 1 | x_{\text{obs}})$. In Section 5, we discuss an extension to the MNAR case.

Throughout this section, we assume one missing pattern. For the case of multiple missing patterns, see Appendix C.

3.1 NCE with EM algorithm

First, we incorporate the EM algorithm into NCE. Although the score equation cannot be used as in (3) for unnormalized models, the estimating equation $Z_{nc1}(\mathbf{x}; \tau) + Z_{nc2}(\mathbf{y}; \tau) = 0$ of NCE can be used instead. Thus, the estimator of $\tau = (c, \theta)$ is defined as the solution to the following equation:

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E}[z_{nc1}(x_i; \tau) | x_{i,\text{obs}}; \theta] + \frac{1}{n} \sum_{j=1}^n z_{nc2}(y_j; \tau) = 0, \quad (4)$$

where each conditional expectation in the first term is taken with respect to the posterior

$$p(x_{i,\text{mis}} | x_{i,\text{obs}}; \theta) = \frac{\tilde{p}(x_i; \theta)}{\int \tilde{p}(x_i; \theta) \mu(dx_{i,\text{mis}})}. \quad (5)$$

Note that the first term in the left hand side of (4) formally means

$$\frac{1}{n} \sum_{i=1}^n \{ \delta_i z_{nc1}(x_i; \tau) + (1 - \delta_i) \mathbb{E}[z_{nc1}(x_i; \tau) | x_{i,\text{obs}}; \theta] \}, \quad (6)$$

because the dimension of $x_{i,\text{obs}}$ may vary with i . Throughout this paper, we implicitly assume this conversion following the convention in the literature of missing data (Seaman et al., 2013).

Generally, it is difficult to analytically calculate the conditional expectation in (4). In the subsequent subsection, we develop a method to resolve this problem. Here, assuming that the conditional expectation in (4) can be calculated analytically, we propose the EM algorithm to solve the equation (4), which is given by Algorithm 1.

Note that each update of $\hat{\tau}_t$ in Algorithm 1 can be replaced with M-estimators. For example, when $f(x) =$

Algorithm 1: NCE with EM algorithm

input : $\{x_i\}_{i=1}^n$, $\hat{\tau}_0 = (\hat{c}_0, \hat{\theta}_0)$
output: $\hat{\tau}_T = (\hat{c}_T, \hat{\theta}_T)$
 1 Initialize $t = 0$
 2 Take n samples $\{y_j\}_{j=1}^n$ from $a(y)$
 3 **repeat**
 4 Solve the following equation w.r.t τ and update $\hat{\tau}_{t+1}$:

$$\text{E}[Z_{nc1}(\mathbf{x}; \tau) \mid \mathbf{x}_{\text{obs}}; \theta_t] + Z_{nc2}(\mathbf{y}; \tau) = 0.$$

 5 $t = t + 1$
 6 **until** $\hat{\tau}_t$ converges;

$x \log x - (1 + x) \log(1 + x)$ (original NCE), $\hat{\tau}_{t+1}$ is obtained by maximization with respect to τ of

$$\sum_{i=1}^n \text{E} \left[\log \frac{r(x_i; \tau)}{r(x_i; \tau) + 1} \mid x_{i,\text{obs}}; \hat{\theta}_t \right] + \sum_{j=1}^n \log \frac{1}{r(y_j; \tau) + 1}. \quad (7)$$

When $f(x) = x \log x$ (MC-MLE), $\hat{\tau}_{t+1}$ is obtained by minimization with respect to τ of

$$-\frac{1}{n} \sum_{i=1}^n \text{E}[\log q(x_i; \tau) \mid x_{i,\text{obs}}; \hat{\theta}_t] + \frac{1}{n} \sum_{j=1}^n r(y_j; \tau).$$

Remark 3.1 (Difference from variational NCE) From (7), Algorithm 1 and variational NCE (Rhodes and Gutmann, 2019) are different. See Appendix D for details.

3.2 NCE with Fractional Imputation (FINCE)

It is often infeasible to analytically calculate the conditional expectation in Algorithm 1. Thus, in the same spirit of fractional imputation (FI; Kim, 2011), we incorporate importance sampling using random variables from an auxiliary distribution $b(x_{\text{mis}})$. Namely, we use the formula

$$\begin{aligned} & \int u(x) p(x_{\text{mis}} \mid x_{\text{obs}}; \theta) \mu(dx_{\text{mis}}) \\ &= \frac{\text{E}_{b(x_{\text{mis}})}[u(x) \tilde{p}(x_{\text{mis}}, x_{\text{obs}}; \theta) / b(x_{\text{mis}})]}{\text{E}_{b(x_{\text{mis}})}[\tilde{p}(x_{\text{mis}}, x_{\text{obs}}; \theta) / b(x_{\text{mis}})]} \end{aligned}$$

to calculate $\text{E}[Z_{nc1}(\mathbf{x}; \tau) \mid \mathbf{x}_{\text{obs}}; \theta]$ in (4). The resulting procedure is given in Algorithm 2. Here, \propto in the W-step indicates a normalization so that the summation of w_{ik} over k is equal to 1 for each i .

Note that again the update of $\hat{\tau}_t$ in M-step can be replaced with M-estimators. For example, the conditional expectation in (7) is calculated as

$$\sum_{k=1}^m w_{ik} \log \frac{r(x_i^{*k}; \tau)}{r(x_i^{*k}; \tau) + 1}.$$

Algorithm 2: FINCE

input : $\{x_i\}_{i=1}^n$, $\hat{\tau}_0 = (\hat{c}_0, \hat{\theta}_0)$
output: $\hat{\tau}_T = (\hat{c}_T, \hat{\theta}_T)$
 1 Initialize $t = 0$
 2 Take n samples $\{y_j\}_{j=1}^n$ from $a(y)$.
 3 For i with $\delta_i = 0$, take m samples $\{x_{i,\text{mis}}^{*k}\}_{k=1}^m$ from $b(x_{\text{mis}})$ and set $x_i^{*k} = (x_{i,\text{obs}}, x_{i,\text{mis}}^{*k})$.
 4 For i with $\delta_i = 1$, set m samples $\{x_i^{*k}\}_{k=1}^m$ to $x_i^{*k} = x_i$
 5 **repeat**
 6 W-Step:
 7 For i with $\delta_i = 0$; $w_{ik} \propto q(x_i^{*k}; \hat{\tau}_t) / b(x_{i,\text{mis}}^{*k})$. This means

$$w_{ik} = \frac{q(x_i^{*k}; \hat{\tau}_t) / b(x_{i,\text{mis}}^{*k})}{\sum_{k=1}^m q(x_i^{*k}; \hat{\tau}_t) / b(x_{i,\text{mis}}^{*k})}.$$

 8 For i with $\delta_i = 1$; $w_{ik} = 1/m$
 9 M-step: Solve the following equation w.r.t τ and update $\hat{\tau}_{t+1}$:

$$\left[\frac{1}{n} \sum_{i=1}^n \sum_{k=1}^m w_{ik} z_{nc1}(x_i^{*k}; \tau) \right] + Z_{nc2}(\mathbf{y}; \tau) = 0.$$

 10 $t = t + 1$
 11 **until** $\hat{\tau}_t$ converges;

The choice of the noise and auxiliary distributions is important for improved estimation accuracy. Specifically, the noise distribution $a(x)$ should be generally close to $p(x_{\text{mis}}, x_{\text{obs}}; \theta_0)$, while the auxiliary distribution $b(x_{\text{mis}})$ should be close to $p(x_{\text{mis}} \mid x_{\text{obs}}; \theta_0)$. When there are samples without missing data (complete data) as in Section 6, moment matching with complete data can be used to determine $a(x)$ and $b(x_{\text{mis}})$. We recommend using heavy-tailed distributions for $a(x)$ and $b(x_{\text{mis}})$, following the common strategy of importance sampling (Owen, 2013).

3.3 Score Matching with Fractional Imputation (FISCORE)

Since score matching is defined in the form of Z-estimators like NCE, we can define score matching with the EM algorithm as the solution to $Z_{sc,\text{obs}}(\mathbf{x}_{\text{obs}}; \theta) = 0$, where

$$Z_{sc,\text{obs}}(\mathbf{x}_{\text{obs}}; \theta) = \text{E}[Z_{sc}(\mathbf{x}; \theta) \mid \mathbf{x}_{\text{obs}}; \theta]. \quad (8)$$

Since calculation of the conditional expectation in (8) is often challenging, we again propose using importance sampling with an auxiliary distribution $b(x)$ like FINCE. The resulting procedure of FISCORE is provided in Algorithm 3. A similar algorithm is obtained for the non-negative score matching.

Remark 3.2 (MINCE and MISCORE) We can also

Algorithm 3: FISCORE

input : $\{x_i\}_{i=1}^n, \hat{c}_0, \hat{\theta}_0$
output: $\hat{\theta}_T$

- 1 Initialize $t = 0, \hat{\tau}_0 = (\hat{c}_0, \hat{\theta}_0)$.
- 2 For each i with $\delta_i = 0$, take m samples $\{x_{i,\text{mis}}^{*k}\}_{k=1}^m$ from $b(x)$.
- 3 For i with $\delta_i = 1$, set m samples $\{x_i^{*k}\}_{k=1}^m$ so that $x_i^{*k} = x_i$.
- 4 **repeat**
- 5 **W-Step:**
- 6 If $\delta_i = 0$; $w_{ik} \propto \tilde{p}(x_i^{*k}; \hat{\theta}_t) / b(x_{\text{mis}}^{*k})$.
- 7 If $\delta_i = 1$; $w_{ik} = 1/m$.
- 8 **M-step:** Solve the following equation to obtain for $\hat{\theta}_{t+1}$ w.r.t θ :

$$\frac{1}{n} \sum_{i=1}^n \sum_{k=1}^m w_{ik} z_{sc}(x_i^{*k}; \theta) = 0.$$
- 9 $t = t + 1$
- 10 **until** $\hat{\tau}_t$ converges;

combine MI with NCE and score matching though it is unstable and computationally heavy. See Appendix F.

Remark 3.3 (FICD) We can also extend our approach to the contrastive divergence (CD) method as in Appendix G.

4 ASYMPTOTICS AND CONFIDENCE INTERVALS

We derive the asymptotic distributions of FINCE and FISCORE. Based on the asymptotic distributions, we construct confidence intervals, which enable hypothesis testing. This is an advantage of the proposed methods over variational NCE (Rhodes and Gutmann, 2019). Proofs of the theorems are given in the Appendix.

4.1 FINCE

We start with the analysis of FINCE with the EM algorithm $\hat{\tau}_{nc}$, which is the solution to $Z_{nc,\text{obs}}(\mathbf{x}_{\text{obs}}, \mathbf{y}; \tau) = 0$ where

$$Z_{nc,\text{obs}}(\mathbf{x}_{\text{obs}}, \mathbf{y}; \tau) = \mathbb{E}[Z_{nc1}(\mathbf{x}; \tau) | \mathbf{x}_{\text{obs}}; \tau] + Z_{nc2}(\mathbf{y}; \tau).$$

Based on the theory of Z-estimators (van der Vaart, 1998), its asymptotic distribution is obtained as follows.

Theorem 1 We have

$$\sqrt{n}(\hat{\tau}_{nc} - \tau_0) \xrightarrow{d} N(0, \mathcal{I}_{1,nc}^{-1} \mathcal{J}_{1,nc} (\mathcal{I}_{1,nc}^\top)^{-1}),$$

where

$$\begin{aligned} \mathcal{I}_{1,nc} &= \mathbb{E}[\nabla_{\tau^\top} Z_{nc,\text{obs}}(\mathbf{x}_{\text{obs}}, \mathbf{y}; \tau_0)], \\ \mathcal{J}_{1,nc} &= \text{var}[Z_{nc,\text{obs}}(\mathbf{x}_{\text{obs}}, \mathbf{y}; \tau_0)]. \end{aligned}$$

Next, we investigate FINCE by considering each iteration. Given an initial \sqrt{n} -consistent estimator $\hat{\tau}_p$, we obtain the imputed equation $Z_{nc,m}(\tau | \hat{\tau}_p) = 0$, where $Z_{nc,m}(\tau | \hat{\tau}_p)$ is given by

$$\left\{ \frac{1}{n} \sum_{i=1}^n \sum_{k=1}^m w(x_i^{*k}; \hat{\tau}_p) z_{nc1}(x_i^{*k}; \tau) \right\} + \frac{1}{n} \sum_{j=1}^n z_{nc2}(y_j; \tau),$$

where $x_i^{*k} = (x_{i,\text{obs}}, x_{i,\text{mis}}^{*k})$, $x_{i,\text{mis}}^{*k} \sim b(x_{\text{mis}})$, and $w(x_i^{*k}; \hat{\tau}_p) \propto q(x_i^{*k}; \hat{\tau}_p) / b(x_{i,\text{mis}}^{*k})$. As $m \rightarrow \infty$, $Z_{nc,m}(\tau | \hat{\tau}_p)$ converges to $\bar{Z}_{nc}(\tau | \hat{\tau}_p)$ given by

$$\left\{ \frac{1}{n} \sum_{i=1}^n \mathbb{E}[z_{nc1}(x_i; \tau) | x_{i,\text{obs}}; \hat{\tau}_p] \right\} + \frac{1}{n} \sum_{j=1}^n z_{nc2}(y_j; \tau).$$

Let $\hat{\tau}_{nc,\infty}$ be the solution to $\bar{Z}_{nc}(\tau | \hat{\tau}_p) = 0$. Then, as proved later in Theorem 3, we obtain

$$\hat{\tau}_{nc,\infty} = \hat{\tau}_{nc} + \mathcal{I}_{3,nc}^{-1} \mathcal{I}_{2,nc} (\hat{\tau}_p - \hat{\tau}_{nc}) + o_p(n^{-1/2}),$$

where $\mathcal{I}_{3,nc} = \mathbb{E}[\nabla_{\tau^\top} Z_{nc}(\mathbf{x}, \mathbf{y}; \tau_0)]$ and $\mathcal{I}_{2,nc} = \mathcal{I}_{3,nc} - \mathcal{I}_{1,nc}$.

Let $\hat{\tau}^{(0)} = \hat{\tau}_p$ and define $\hat{\tau}^{(t)}$ to be the solution to $\bar{Z}_{nc}(\tau | \hat{\tau}^{(t-1)}) = 0$ for each t . Then, we obtain the following.

Corollary 1 We have

$$\hat{\tau}^{(t)} = \hat{\tau}_{nc} + (\mathcal{I}_{3,nc}^{-1} \mathcal{I}_{2,nc})^{t-1} (\hat{\tau}^{(0)} - \hat{\tau}_{nc}) + o_p(n^{-1/2}).$$

If the spectral radius of $\mathcal{I}_{3,nc}^{-1} \mathcal{I}_{2,nc}$ is less than 1, then $\hat{\tau}^{(t)}$ converges to $\hat{\tau}_{nc}$ as $t \rightarrow \infty$.

Let $v(x; \tau) = \nabla_{\tau} \log q(x; \tau)$ and $r_0(x) = r(x; \tau_0) = q(x; \tau_0) / a(x)$. For the original NCE, each term in the above is explicitly obtained as follows.

Corollary 2 When $f(x) = x \log x - (1+x) \log(1+x)$,

$$\begin{aligned} \mathcal{I}_{1,nc} &= \mathbb{E} \left[\mathbb{E} \left[\frac{v(x; \tau_0)}{1 + r_0(x)} \middle| x_{\text{obs}} \right] \mathbb{E} [v(x; \tau_0)^\top | x_{\text{obs}}] \right], \\ \mathcal{I}_{3,nc} &= \mathbb{E} \left[\frac{v(x; \tau_0) \otimes 2}{1 + r_0(x)} \right], \\ \mathcal{J}_{1,nc} &= \text{var}_q[\mathbb{E}[z_{nc1}(x; \tau_0) | x_{\text{obs}}]] + \text{var}_a[z_{nc2}(y; \tau_0)], \\ z_{nc1}(x; \tau) &= -\frac{v(x; \tau)}{1 + r(x; \tau)}, \quad z_{nc2}(y; \tau) = \frac{r(y; \tau)v(y; \tau)}{1 + r(y; \tau)}. \end{aligned}$$

For MC-MLE, we can prove the convergence of FINCE as follows.

Corollary 3 When $f(x) = x \log x$,

$$\mathcal{I}_{1,nc} = \mathbb{E} \left[\mathbb{E} [v(x; \tau_0) | x_{\text{obs}}] \otimes 2 \right], \quad \mathcal{I}_{3,nc} = \mathbb{E} [v(x; \tau_0) \otimes 2].$$

Additionally, $\{\mathcal{I}_{3,nc}^{-1} \mathcal{I}_{2,nc}\}^j \rightarrow 0$ as $j \rightarrow \infty$.

4.2 FISCORE

First, we analyze score matching with the EM algorithm $\hat{\theta}_{sc}$, which is the solution of $Z_{sc,obs}(\mathbf{x}_{obs}; \theta) = 0$ where $Z_{sc,obs}(\mathbf{x}_{obs}; \theta)$ is defined by (8). The asymptotic distribution of $\hat{\theta}_{sc}$ is obtained as follows.

Theorem 2 *We have*

$$\sqrt{n}(\hat{\theta}_{sc} - \theta_0) \xrightarrow{d} N(0, \mathcal{I}_{1,sc}^{-1} \mathcal{J}_{1,sc} (\mathcal{I}_{1,sc}^\top)^{-1}),$$

where

$$\begin{aligned} \mathcal{I}_{1,sc} &= E[\nabla_{\theta^\top} Z_{sc,obs}(\mathbf{x}_{obs}; \theta_0)], \\ \mathcal{J}_{1,sc} &= \text{var}[Z_{sc,obs}(\mathbf{x}_{obs}; \theta_0)]. \end{aligned}$$

Next, we study the asymptotic property of FISCORE by focusing on each update. Given an initial \sqrt{n} -consistent estimator $\hat{\theta}_p$ for θ , consider the imputed equation:

$$Z_{sc,m}(\theta | \hat{\theta}_p) := \frac{1}{n} \sum_{i=1}^n \sum_{k=1}^m w(x_i^{*k}; \hat{\theta}_p) z_{sc}(\theta; x_i^{*k}) = 0,$$

where $x_i^{*k} = (x_{i,obs}, x_{i,mis}^{*k})$, $x_{i,mis}^{*k} \sim b(x_{mis})$, and $w(x_i^{*k}; \hat{\theta}_p) \propto \tilde{p}(x_i^{*k}; \hat{\theta}_p) / b(x_{i,mis}^{*k})$.

Here, we consider the case $m \rightarrow \infty$. See Appendix E for the case of finite m . As $m \rightarrow \infty$, the function $Z_{sc,m}(\theta | \hat{\theta}_p)$ converges to

$$\bar{Z}_{sc}(\theta | \hat{\theta}_p) = E[Z_{sc}(\mathbf{x}; \theta) | \mathbf{x}_{obs}; \hat{\theta}_p].$$

Let $\hat{\theta}_{sc,\infty}$ be the solution to $\bar{Z}_{sc}(\theta | \hat{\theta}_p) = 0$. Its asymptotic property is obtained as follows.

Theorem 3 *We have*

$$\hat{\theta}_{sc,\infty} = \hat{\theta}_{sc} + \mathcal{I}_{3,sc}^{-1} \mathcal{I}_{2,sc} (\hat{\theta}_p - \hat{\theta}_{sc}) + o_p(n^{-1/2}),$$

where

$$\begin{aligned} \mathcal{I}_{2,sc} &= -E[\text{cov}[z_{sc}(x; \theta_0), \nabla_{\theta} \log \tilde{p}(x; \theta_0) | x_{obs}]], \\ \mathcal{I}_{3,sc} &= E \left[\sum_{s=1}^d \nabla_{\theta} c_s(x; \theta_0)^{\otimes 2} \right]. \end{aligned}$$

In the proof of Theorem 3, we use the relation $\mathcal{I}_{3,sc} = \mathcal{I}_{1,sc} + \mathcal{I}_{2,sc}$, which corresponds to the missing information principle or Louis' formula for normalized models (Kim and Shao, 2013; Orchard and Woodbury, 1972; Louis, 1982). Specifically, if $Z_{sc}(\mathbf{x}; \theta)$ is replaced by the true score function $S_{sc}(\mathbf{x}; \theta) = \nabla_{\theta} \log p(\mathbf{x}; \theta)$, then Theorem 3 reduces to the result of Wang and Robins (1998). In this case, $\mathcal{I}_{3,sc}$, $\mathcal{I}_{1,sc}$ and $\mathcal{I}_{2,sc}$ are replaced by

$$\begin{aligned} \mathcal{I}_{com} &= E[\nabla_{\theta^\top} S_{sc}(\mathbf{x}; \theta_0)], \quad \mathcal{I}_{obs} = E[\nabla_{\theta^\top} S_{obs}(\mathbf{x}_{obs}; \theta_0)], \\ \mathcal{I}_{mis} &= E[S_{mis}(\mathbf{x}; \theta_0)^{\otimes 2}], \end{aligned}$$

respectively, where

$$\begin{aligned} S_{mis}(\mathbf{x}; \theta) &= S_{sc}(\mathbf{x}; \theta) - E[S_{sc}(\mathbf{x}; \theta) | \mathbf{x}_{obs}; \theta], \\ S_{obs}(\mathbf{x}_{obs}; \theta) &= \int S_{sc}(\mathbf{x}; \theta) \mu(d\mathbf{x}_{mis}). \end{aligned}$$

The relation $\mathcal{I}_{com} = \mathcal{I}_{obs} + \mathcal{I}_{mis}$ holds, and the term $\mathcal{I}_{com}^{-1} \mathcal{I}_{mis}$ is often called the fraction of missing information (Kim and Shao, 2013). For the current problem, $\mathcal{I}_{3,sc}^{-1} \mathcal{I}_{2,sc}$ can be considered as an analog. As seen in Corollary 1, this quantity is important to guarantee the convergence. It is generally difficult to prove that the spectral radius of $\mathcal{I}_{3,sc}^{-1} \mathcal{I}_{2,sc}$ is less than 1 in FISCORE. However, the experimental results presented in Section 6 imply that this algorithm converges in practice.

Remark 4.1 *Similar results hold for the non-negative score matching defined by $M_{sc+}(x; \theta)$. See Appendix I.*

Remark 4.2 *See Appendix I for variance estimators based on Theorems 1, 2 and 3.*

5 EXTENSION TO MNAR CASE

We discuss an extension to the case of missing not at random (MNAR). In general, the nonparametric identification condition does not hold in the MNAR case (Robins and Ritov, 1997). However, assuming the existence of nonresponse instrument and parametric models, the parameter can be identified in some cases (Kim and Kim, 2012; Wang et al., 2014). We hereafter assume the existence of a nonresponse instrument so that the parameter can be identified.

Assumption 1 (Wang et al. (2014)) *There exists nonresponse instrument \mathbf{x}_2 s.t. $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2)$ and $\mathbf{x}_2 \perp\!\!\!\perp \delta | \mathbf{x}_1$.*

To estimate the parameter under MNAR data, FISCORE and FINCE can be still applied. First, we specify a propensity score model $\pi(\delta|x; \phi)$ for $\Pr(\delta|x)$. For the case of FISCORE, we want to solve the equation with respect to η :

$$E \left[\left(\begin{array}{c} Z_{sc}(\mathbf{x}; \theta) \\ \nabla_{\phi} \log \pi(\delta|x; \phi) \end{array} \right) | \mathbf{x}_{obs}, \delta; \eta \right] = 0, \quad (9)$$

where the expectation is taken under $t(\mathbf{x}_{mis} | \mathbf{x}_{obs}, \delta; \eta) \propto p(\mathbf{x}; \theta) \pi(\delta|x; \phi)$, and $\eta = (\theta, \phi)$. Importantly, we must address the selection mechanism unlike in the MAR and MCAR cases, because $p(x_{mis} | x_{obs}) = p(x_{mis} | \delta, x_{obs})$ does not hold. The difference is evident when we compare (9) with (8). Owing to MNAR, the first modification is such that the selection mechanism $\pi(\delta|x)$ appears when calculating the fractional weight: $w_{ik} \propto \tilde{p}(x_i^{*k}; \hat{\theta}_t) \pi(\delta_i | x_i^{*k}; \hat{\phi}_t) / b(x_{i,mis}^{*k})$. The second modification is the score function of the propensity score model which is illustrated in (9).

In the case of FINCE, let $\zeta = (\tau^\top, \phi^\top)^\top$ and $Z_{nc}(\delta, \mathbf{x}, \mathbf{y}; \zeta)$ be defined as an augmented estimating

equation:

$$\begin{pmatrix} Z_{nc1}(\mathbf{x}; \tau) + Z_{nc2}(\mathbf{y}; \tau) \\ \nabla_{\phi} \log \pi(\boldsymbol{\delta} | \mathbf{x}; \phi) \end{pmatrix}.$$

The algorithm is modified to solve the following equation with respect to ζ :

$$\mathbb{E} \left[\begin{pmatrix} Z_{nc1}(\mathbf{x}; \tau) + Z_{nc2}(\mathbf{y}; \tau) \\ \nabla_{\phi} \log \pi(\boldsymbol{\delta} | \mathbf{x}; \phi) \end{pmatrix} \middle| \mathbf{x}_{\text{obs}}, \boldsymbol{\delta}; \zeta \right] = 0.$$

Refer to Appendix H for the details.

6 SIMULATION RESULTS

Here, we confirm the validity of FINCE and FISCORE by simulation. We do not compare them with variational NCE because the latter does not take the MNAR case into account and does not provide a confidence interval.

6.1 Truncated Normal Distribution

First, we consider covariance estimation for the two-dimensional truncated normal distribution:

$$\tilde{p}(x_1, x_2; \Sigma) = \exp \left(-\frac{1}{2} x^{\top} \Sigma^{-1} x \right), \Sigma = \begin{pmatrix} 2.0 & 0.3 \\ 0.3 & 2.0 \end{pmatrix},$$

where $x_1, x_2 \geq 0$. For n samples of $x = (x_1, x_2)$, x_1 was always observed whereas missing data of x_2 was introduced with the following two mechanisms. The binary random variable δ is the missing indicator: x_2 was observed if and only if $\delta = 1$.

- MAR: $\Pr(\delta = 1 | x) = 1/[1 + \exp\{-(x_1 - 0.9)/0.3\}]$
- MNAR: $\Pr(\delta = 1 | x) = 1/[1 + \exp\{-5(x_2 - 0.9)\}]$

For data generation, we used the R-package *mvtnorm* (Genz et al., 2018). In both cases, the overall missing rates were about 30%.

We compared three estimators of Σ : original NCE based on complete data, (original) FINCE, and FISCORE. Note that the noise and auxiliary distributions were set to truncated normal distributions with moment matching and $m = 100$ imputations were used in FINCE and FISCORE.

Table 1 presents the absolute bias and median squared error from 200 simulations. It shows that NCE on complete data has significant bias in both the MAR and MNAR cases, which is expected from the theory of missing data (Little and Rubin, 2002). On the other hand, FINCE and FISCORE achieve better estimation accuracy with reduced bias by appropriately accounting for the missing data.

In addition, we constructed 95% confidence intervals based on the variance estimators in Appendix I. Table 2 shows their coverage probabilities: they are approximately equal to 95% in both FINCE and FISCORE.

Table 1: The absolute bias and median square error

		MAR		
n		NCE	FINCE	FISCORE
500	(bias)	0.29	0.03	0.03
	(mse)	0.040	0.024	0.021
1000	(bias)	0.25	0.02	0.02
	(mse)	0.032	0.015	0.011
		MNAR		
n		NCE	FINCE	FISCORE
500	(bias)	0.33	0.18	0.12
	(mse)	0.041	0.027	0.021
1000	(bias)	0.24	0.14	0.14
	(mse)	0.039	0.020	0.012

Table 2: Coverage probabilities: MAR setting

n	FINCE	FISCORE
500	94%	89%
1000	94%	92%

6.2 Truncated Gaussian Graphical Model

Next, we consider estimation of the truncated Gaussian graphical model (GGM) considered in Lin et al. (2016) with missing data.

Let $G = (V, E)$ be an undirected graph where $V = \{1, \dots, d\}$. Then, a truncated GGM with graph G is defined as the unnormalized model (1) with

$$\tilde{p}(x; \Sigma) = \exp \left(-\frac{1}{2} x^{\top} \Sigma^{-1} x \right) \quad (x \in \mathbb{R}_+^d), \quad (10)$$

where Σ is a $d \times d$ positive definite matrix satisfying $(\Sigma^{-1})_{ij} = 0$ for $(i, j) \notin E$. Similar to the original GGM (Lauritzen, 1996), X_i and X_j are conditionally independent on the other variables X_k ($k \neq i, j$) if $(i, j) \notin E$. Here, we estimate G by using the confidence intervals of the entries of Σ^{-1} .

We generated $n = 1000$ independent samples from a truncated GGM with $d = 10$ and G provided in the top panel of Figure 1. Namely, there are three clusters (x_1, x_2, x_3) , (x_4, x_5, x_6) , and (x_7, x_8, x_9) of three variables and one isolated variable x_{10} . We set all diagonal entries of Σ^{-1} to 1 and all nonzero off-diagonal entries of Σ^{-1} to 0.5. We introduced missing values on x_3 , x_6 and x_9 by using the following MAR mechanism: for $k = 1, 2, 3$, random vector $c_k \in \mathbb{R}^{10}$ was generated by $(c_k)_3 = (c_k)_6 = (c_k)_9 = 0$ and $(c_k)_j \sim N(0, 1)$ ($j \neq 3, 6, 9$) and then x_{3k} was missed with probability $1/(3 + \exp(c_k^{\top} x))$. The proportion of missing data was approximately 60%.

Then, we fitted the truncated GGM by using FINCE and FISCORE with 100 imputations. We used $N(0, 2)$ truncated to the positive orthant as the proposal distribution for missing entries. In FINCE, we generated $n = 1000$ noise

samples from the product of the coordinate-wise exponential distributions with the same mean as the data.

We determined graph G by collecting all edges (i, j) for which the 95 % confidence interval of $(\Sigma^{-1})_{ij}$ did not include zero. Figure 1 presents the result of one realization.

Table 3 shows the proportions of falsely selected edges (false positive) and falsely unselected edges (false negative) in 100 realizations. The coverage probabilities of the confidence intervals are approximately equal to 95% in both FINCE and FISCORE.




Figure 1: True and selected graphs for the truncated GGM

Table 3: Proportions of false positives (FP) and false negatives (FN) in edge selection of truncated GGM

	FINCE	FISCORE
FP	10.5%	6.4%
FN	12.6%	23.3%

7 APPLICATION TO REAL DATA

Many models in directional statistics have intractable normalization constants (Mardia and Jupp, 1999). Here, we consider estimation of the bivariate circular distribution proposed by Singh et al. (2002), which is a probability distribution on two circular variables $x_1, x_2 \in [0, 2\pi)$. It is an unnormalized model (1) with

$$\tilde{p}(x_1, x_2; \theta) = \exp(\kappa_1 \cos(x_1 - \mu_1) + \kappa_2 \cos(x_2 - \mu_2) + \lambda_{12} \sin(x_1 - \mu_1) \sin(x_2 - \mu_2)), \quad (11)$$

where $\theta = (\kappa_1, \kappa_2, \mu_1, \mu_2, \lambda_{12})$. For identifiability, we imposed the parameter constraints $\kappa_1 \geq 0$, $\kappa_2 \geq 0$, $0 \leq \mu_1 < 2\pi$ and $0 \leq \mu_2 < 2\pi$. Although Mardia et al. (2008) developed a method for estimating θ based on pseudo-likelihood, it is not applicable to missing data. We applied FINCE to estimate θ from missing data.

We used the data on wind direction in Tokyo on 00:00 (x_1) and 12:00 (x_2) for each day in 2018¹. Thus, the sample size is $n = 365$. The data were discretized into 16 directions, such as north-northeast. Figure 2 presents a 16×16 histogram of raw data in gray scale.

¹available on Japan Meteorological Agency website




Figure 2: Wind direction data

For each of the 365 samples, we introduced missing for x_2 with probability $1/(1 + \exp(\cos(x_1)))$. Thus, the missing mechanism was MAR. Among 365 samples, x_2 was observed in 212 samples.

Then, we fit the model (11) by FINCE with 100 imputations and 1000 noise samples, where the noise distribution and proposal distribution for missing entries were set to uniform distributions on $[0, 2\pi) \times [0, 2\pi)$ and $[0, 2\pi)$, respectively. For comparison, we also fit the model (11) by NCE, in which 153 samples with x_2 missing were simply discarded, which is called the complete case analysis and known to have bias in MAR cases (Kim and Shao, 2013).

Table 4 presents the confidence intervals obtained by FINCE, NCE on complete cases, and NCE on the full data (365 samples) for reference. Whereas complete case analysis has a large bias, FINCE provides similar estimates to NCE on full data, with wider confidence intervals due to missing. In particular, FINCE succeeds in detecting the 5% significance of $\lambda_{12} \neq 0$, which implies that x_1 and x_2 are not independent. Thus, FINCE enables statistical inference based on unnormalized models by properly handling missing data.

Table 4: Confidence intervals for bivariate circular distribution (11). (cc: complete case)

	FINCE	NCE (cc)	NCE (full)
κ_1	[0.30,1.17]	[0.58,1.99]	[0.26,1.16]
μ_1	[0.59,1.62]	[-0.18,1.31]	[0.69,1.58]
κ_2	[0.11,1.09]	[-0.07,0.96]	[0.16,0.84]
μ_2	[4.08,5.01]	[3.46,5.70]	[4.09,4.85]
λ_{12}	[-2.20,-0.29]	[-1.96,0.94]	[-1.67,-0.48]

8 CONCLUSION

We have proposed estimation methods for unnormalized models with missing data: FINCE and FISCORE. The proposed methods are computationally efficient, valid under general missing mechanisms, and enable statistical inference using the confidence intervals. Extending the recently developed statistically efficient estimators for unnormalized models (Uehara et al., 2020) to missing data is an interesting future research.

Acknowledgements

We would like to thank the anonymous reviewers for their insightful comments and suggestions.

Masatoshi Uehara was supported in part by MASA-SON Foundation. Takeru Matsuda was partially supported by JSPS KAKENHI Grant Numbers 16H06533 and 19K20220.

References

- Besag, J. (1975). Statistical analysis of non-lattice data. *Journal of the Royal Statistical Society. Series D (The Statistician)* 24, 179–195.
- Blei, D. M., A. Kucukelbir, and J. D. McAuliffe (2017). Variational inference: A review for statisticians. *Journal of the American Statistical Association* 112, 859–877.
- Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B* 39, 1–37.
- Genz, A., F. Bretz, T. Miwa, X. Mi, F. Leisch, F. Scheipl, and T. Hothorn (2018). *mvtnorm: Multivariate Normal and t Distributions*. R package version 1.0-8.
- Geyer, C. (1994). On the convergence of monte carlo maximum likelihood calculations. *Journal of the Royal Statistical Society, Series B* 56, 261–274.
- Gutmann, M. and J. Hirayama (2011). Bregman divergence as general framework to estimate unnormalized statistical models. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI 2011)*.
- Gutmann, M. and A. Hyvärinen (2012). Noise-Contrastive Estimation of Unnormalized Statistical Models, with Applications to Natural Image Statistics. *Journal of Machine Learning Research* 13, 307–361.
- Hinton, G. E. (2002). Training products of experts by minimizing contrastive divergence. *Neural Computation* 14, 1771–1800.
- Hyvärinen, A. (2005). Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research* 6, 695–709.
- Hyvärinen, A. (2007). Some extensions of score matching. *Computational Statistics and Data Analysis* 51, 2499–2512.
- Hyvärinen, A., J. Karhunen, and E. Oja (2001). *Independent component analysis*. New York: J. Wiley.
- Kim, J. K. (2011). Parametric fractional imputation for missing data analysis. *Biometrika* 98, 119–132.
- Kim, J. K. and J. Shao (2013). *Statistical Methods for Handling Incomplete Data*. Chapman & Hall / CRC.
- Kim, J. Y. and J. K. Kim (2012). Parametric fractional imputation for nonignorable missing data. *Journal of the Korean Statistical Society* 41, 291–303.
- Lauritzen, S. L. (1996). *Graphical models*. Oxford statistical science series ; 17. Oxford: Clarendon Press.
- Levine, R. A. and G. Casella (2001). Implementations of the monte carlo em algorithm. *Journal of Computational and Graphical Statistics* 10, 422–439.
- Lin, L., M. Drton, and A. Shojaie (2016). Estimation of high-dimensional graphical models using regularized score matching. *Electronic journal of statistics* 10, 806–854.
- Little, R. J. A. and D. B. Rubin (2002). *Statistical analysis with missing data* (2nd ed. ed.). Hoboken, N.J.: Wiley.
- Louis, T. A. (1982). Finding the observed information matrix when using the EM algorithm. *Journal of the Royal Statistical Society, Series B* 44, 226–233.
- Mardia, K. V., G. Hughes, C. C. Taylor, and H. Singh (2008). A multivariate von mises distribution with applications to bioinformatics. *Can. J. Statist.* 89, 99–109.
- Mardia, K. V. and P. E. Jupp (1999). *Directional Statistics*. New York, NY: Wiley.
- Meng, X. L. (1994). Multiple-imputation inferences with uncongenial sources of input (with discussion). *Statistical Science* 9, 538–573.
- Mller, J., A. N. Pettitt, R. Reeves, and K. K. Berthelsen (2006). An efficient markov chain monte carlo method for distributions with intractable normalising constants. *Biometrika* 93, 451–458.
- Murray, I., Z. Ghahramani, and D. MacKay (2006). Mcmc for doubly-intractable distributions. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI 2006)*.
- Murray, J. S. (2018). Multiple imputation: A review of practical and theoretical findings. *Statistical Science* 33, 142–159.
- Orchard, T. and M. Woodbury (1972). A missing information principle: theory and applications. In *Proceedings of the 6th Berkeley Symposium on Mathematical Statistics and Probability*, Volume 1, Berkeley, California, pp. 695–715. University of California Press.
- Owen, A. B. (2013). *Monte Carlo theory, methods and examples*.

- Pihlaja, M., M. Gutmann, and A. Hyvärinen (2010). A family of computationally efficient and simple estimators for unnormalized statistical models. *In Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI 2010)*.
- Rao, C. R. (2008). *Linear Statistical Inference and its Applications: Second Edition*. Hoboken, NJ, USA: John Wiley & Sons, Inc.
- Rhodes, B. and M. Gutmann (2019). Variational noise-contrastive estimation. *In Proceedings of the International Workshop on Artificial Intelligence and Statistics 2019 (AISTATS 2019)*.
- Robins, J. M. and Y. Ritov (1997). Toward a curse of dimensionality appropriate (coda) asymptotic theory for semi-parametric models. *Statistics in medicine* 16, 285–319.
- Robins, J. M. and N. Wang (2000). Inference for imputation estimators. *Biometrika* 87, 113–124.
- Rubin, D. B. (1987). *Multiple Imputation for Nonresponse in Surveys*. New York: Wiley.
- Seaman, S., J. Galati, D. Jackson, and J. Carlin (2013). What is meant by "missing at random"? *Statistical Science* 28, 257–268.
- Singh, H., V. Hnizdo, and E. Demchuk (2002). Probabilistic model for two dependent circular variables. *Biometrika* 89, 719–723.
- Tanner, M. A. and W. H. Wong (1987). The calculation of posterior distribution by data augmentation. *Journal of the American Statistical Association* 82, 528–540.
- Tieleman, T. (2008). Training restricted boltzmann machines using approximations to the likelihood gradient. *In Proceedings of the international conference on machine learning (ICML 2008)*.
- Tsiatis, A. (2006). *Semiparametric theory and missing data*. Springer series in statistics. New York: Springer.
- Uehara, M., T. Kanamori, T. Takenouchi, and T. Matsuda (2020). Unified estimation framework for unnormalized models with statistical efficiency. *In Proceedings of the 23rd International Workshop on Artificial Intelligence and Statistics (AISTATS 2020)*.
- Uehara, M., T. Matsuda, and F. Komaki (2018). Analysis of noise contrastive estimation from the perspective of asymptotic variance. *arXiv preprint arXiv:1808.07983*.
- van der Vaart, A. W. (1998). *Asymptotic statistics*. Cambridge, UK ; New York, NY, USA: Cambridge University Press.
- Wang, N. and J. M. Robins (1998). Large-sample theory for parametric multiple imputation procedures. *Biometrika* 85, 935–948.
- Wang, S., J. Shao, and J. K. Kim (2014). Identifiability and estimation in problems with nonignorable nonresponse. *Statistical Science* 24, 1097 – 1116.
- Wei, G. C. G. and M. A. Tanner (1990). A monte carlo implementation of the em algorithm and the poor man's data augmentation algorithms. *Journal of the American Statistical Association* 85, 699–704.
- Yang, S. and J. K. Kim (2016). Fractional imputation in survey sampling: A comparative review. *Statistical Science* 31, 415–432.
- Younes, L. (1989). Maximum likelihood estimation for Gaussian fields. *Probability Theory and Related Fields* 82, 625–645.
- Yu, M., M. Kolar, and V. Gupta (2016). Statistical inference for pairwise graphical models using score matching. *In Advances in Neural Information Processing Systems* 29, pp. 2829–2837.