

Multiple Imputation with Neural Network Gaussian Process for High-dimensional Incomplete Data

Zongyu Dai

Zhiqi Bu

Graduate Group in Applied Mathematics and Computational Science, University of Pennsylvania

DAIZY@SAS.UPENN.EDU

ZBU@SAS.UPENN.EDU

Qi Long

Division of Biostatistics, University of Pennsylvania

ZBU@SAS.UPENN.EDU

Editors: Emtiyaz Khan and Mehmet Gönen

Abstract

Missing data are ubiquitous in real world applications and, if not adequately handled, may lead to the loss of information and biased findings in downstream analysis. Particularly, high-dimensional incomplete data with a moderate sample size, such as analysis of multi-omics data, present daunting challenges. Imputation is arguably the most popular method for handling missing data, though existing imputation methods have a number of limitations. Single imputation methods such as matrix completion methods do not adequately account for imputation uncertainty and hence would yield improper statistical inference. In contrast, multiple imputation (MI) methods allow for proper inference but existing methods do not perform well in high-dimensional settings. Our work aims to address these significant methodological gaps, leveraging recent advances in neural network Gaussian process (NNGP) from a Bayesian viewpoint. We propose two NNGP-based MI methods, namely MI-NNGP, that can apply multiple imputations for missing values from a joint (posterior predictive) distribution. The MI-NNGP methods are shown to significantly outperform existing state-of-the-art methods on synthetic and real datasets, in terms of imputation error, statistical inference, robustness to missing rates, and computation costs, under three missing data mechanisms, MCAR, MAR, and MNAR.

Keywords: Missing Data, Multiple imputation, Neural Network Gaussian Processes, Statistical Inference

1. Introduction

Missing data are frequently encountered and present significant analytical challenges in many research areas. Inadequate handling of missing data can lead to biased results in subsequent data analysis. For example, complete case analysis that uses only the subset of observations with all variables observed is known to yield biased results and/or loss of information as it does not utilize the information contained in incomplete cases [Little and Rubin \(2019\)](#). Missing value imputation has become increasingly popular for handling incomplete data. Broadly speaking, imputation methods can be categorized as single imputation (SI) or multiple imputation (MI). SI methods impute missing values for a single time which fails to adequately account for imputation uncertainty; in contrast, MI methods impute missing values multiple times by sampling from some (predictive) distribution to account for imputation uncertainty. MI offers another significant advantage over SI in that it can conduct

hypothesis testing or construct confidence intervals using multiply imputed datasets via Rubin’s rule [Little and Rubin \(2019\)](#). Of note, most popular imputation methods in the machine learning literature, such as matrix completion methods, are SI and hence expected to yield invalid statistical inference as shown in our numerical experiments.

When conducting imputation, it is important to know the mechanisms under which missing values originate, namely, missing completely at random (MCAR), missing at random (MAR), or missing not at random (MNAR) [Little and Rubin \(2019\)](#). To be specific, MCAR means that the missingness does not depend on observed or missing data. While most imputation methods are expected to work reasonably well under MCAR, this assumption is typically too strong and unrealistic in practice, particularly for analysis of incomplete data from biomedical studies. MAR and MNAR are more plausible than MCAR. Under MAR, the missingness depends on only the observed values. Under MNAR, the missingness may depend on both observed and missing values, and it is well-known that additional structural assumptions need to be made in order to develop valid imputation methods under MNAR.

Existing state-of-the-art imputation methods can be categorized into discriminative methods and generative methods. The former includes, but not limited to, MICE [Van Buuren \(2007\)](#); [Deng et al. \(2016\)](#); [Zhao and Long \(2016\)](#), MissForest [Stekhoven and Bühlmann \(2012\)](#), KNN [Liao et al. \(2014\)](#) and matrix completion [Mazumder et al. \(2010\)](#); [Hastie et al. \(2015\)](#); the latter includes joint modeling [Schafer \(1997\)](#); [García-Laencina et al. \(2010\)](#), autoencoders [Ivanov et al. \(2018\)](#); [Mattei and Frellsen \(2019\)](#), and generative adversarial networks [Dai et al. \(2021\)](#); [Yoon et al. \(2018\)](#); [Lee et al. \(2019\)](#). However, the existing imputation methods have several drawbacks. MICE imputes missing values through an iterative approach based on conditional distributions and requires to repeating the imputing procedures multiple times till convergence. MICE [Van Buuren \(2007\)](#), known to be computationally expensive, tends to yield poor performance and may become computationally infeasible for high-dimensional data with high missing rates. Joint modeling (JM), another classical imputation method, relies on strong assumptions for the data distribution. Its performance also deteriorates rapidly as the feature dimension increases. SoftImpute [Mazumder et al. \(2010\)](#), a matrix completion method, conducts single imputation based on the low-rank assumption, leading to underestimating the uncertainty of imputed values. In recent years, many deep learning-based imputation methods have been proposed. As the most representative one, GAIN [Yoon et al. \(2018\)](#) can handle mixed data types. However, the appliance for GAIN in practice is limited as it is valid only under MCAR. Most recently, importance-weighted autoencoder based method MIWAE [Mattei and Frellsen \(2019\)](#) and not-MIWAE [Ipsen et al. \(2020\)](#) can deal with MAR and MNAR mechanism, respectively. Plus, optimal transport-based methods [Muzellec et al. \(2020\)](#) including Sinkhorn and Linear RR have been shown to outperform other state-of-the-art imputation methods under MCAR, MAR and MNAR. However, above methods are shown to exhibit appreciable bias in our high-dimensional data experiments. Moreover, Linear RR imputes missing values iteratively like MICE, hence is inherently not suitable for high-dimensional incomplete data.

To address the limitations of existing imputation methods, we leverage recent developments in neural network Gaussian process (NNGP) theory [Williams \(1997\)](#); [Lee et al. \(2017\)](#); [Novak et al. \(2018, 2020\)](#) to develop a new robust multiple imputation approach. The NNGP theory can provide explicit posterior predictive distribution for missing values and allow for Bayesian inference without actually training the neural networks, leading to

substantial savings in training time. Here we take L -layer fully connected neural networks as an example. Suppose a neural network has layer width n^l (for hidden layer l), activation function ϕ and centered normally distributed weights and biases with variance $\frac{\sigma_w^2}{n^l}$ and σ_b^2 at layer l . When each hidden layer width goes to infinity, each output neuron is a Gaussian process $\mathcal{GP}(0, \mathcal{K}^L)$ where \mathcal{K}^L is deterministically computed from $L, \phi, \sigma_w, \sigma_b$. Details can be found in appendix.

Our contribution: Our proposed deep learning imputation method, Multiple Imputation through Neural Network Gaussian Process (MI-NNGP), is designed for the high-dimensional data setting in which the number of variables/features can be large whereas the sample size is moderate. This setting is particularly relevant to analysis of incomplete high-dimensional -omics data in biomedical research where the number of subjects is typically not large. MI-NNGP is the first deep learning based method that yields satisfactory performance in statistical inference for high-dimensional incomplete data under MAR. Empirically speaking, MI-NNGP demonstrates strong performance on imputation error, statistical inference, computational speed, scalability to high dimensional data, and robustness to high missing rate. 1 summarizes the performance of MI-NNGP in comparison with several existing state-of-the-art imputation methods.

Models	MI	Imp Error	Inference	Scalability
MI-NNGP	✓	✓	✓	✓
Sinkhorn	✓	?	?	✗
Linear RR	✓	?	?	✗
MICE	✓	✓	?	✗
SoftImpute	✗	✓	✗	?
MIWAE	✓	✓	?	?
GAIN	✓	✗	✗	✗

Table 1: Summary of imputation methods. Imp Error refers to imputation error. Question mark indicates that the performance depends on specific settings.

2. Problem Setup

To fix ideas, we consider the multivariate K -pattern missing data, meaning that observations can be categorized into K patterns according to which features have missing values. Within each pattern, a feature is either observed in all cases or missing in all cases as visualized in Figure 1 which provides an illustration of 4-pattern missing data. As a motivation, multivariate K -pattern missing data are often encountered in medical research. For example, the Alzheimer’s Disease Neuroimaging Initiative (ADNI) collected high-dimensional multi-omics data and each -omics modality is measured in only a subset of total cases, leading to the multivariate K -pattern missing data. In addition, a general missing data pattern, after some rearranging of rows, can be converted to K -pattern missing data.

Suppose we have a random sample of n observations with p variables. Denote the $n \times p$ observed data matrix by \mathbf{X} , which may include continuous and discrete values and where $\mathbf{X}_{i,j}$ is the value of j -th variable/feature for i -th case/observation. Let $\mathbf{X}_{i,:}$ and $\mathbf{X}_{:,j}$ denote the i -th row vector and j -th column vector, respectively. Since some elements of \mathbf{X} are

missing, the n observations/cases can be grouped into K patterns (i.e. K submatrices) $\mathbf{X}_{P_k,:}$, for $k \in [K]$; see the illustrative example in Figure 1. Here P_k is the index set for the rows in \mathbf{X} which belong to the k -th pattern. Without loss of generality, we let $\mathbf{X}_{P_1,:}$ denote the set of complete cases for which all features are observed. We define $\mathbf{X}_{P_{-k},:} = \mathbf{X} \setminus \mathbf{X}_{P_k,:}$ as the complement data matrix for $\mathbf{X}_{P_k,:}$. We denote by $\mathbf{obs}(k)$ and $\mathbf{mis}(k)$ the index sets for the columns in $\mathbf{X}_{P_k,:}$ that are fully observed and fully missing, respectively.

3. Multiple Imputation via Neural Network Gaussian Process

In this section, we develop two novel MI methods for multivariate K -pattern missing data based on NNGP. Specifically, we first propose MI-NNGP1 which imputes each missing pattern by exploiting information contained in the set of complete cases ($\mathbf{X}_{P_1,:}$). We then propose MI-NNGP2, which imputes each missing pattern iteratively by utilizing the information contained in all observed data. We further improve both methods by incorporating a bootstrap step.

3.1. Imputing Missing Data from an Alternative Viewpoint

MICE is a quite flexible MI method as it learns a separate imputation model for each variable, in an one by one manner. However, MICE is extremely slow in high dimensional setting and incapable of learning the features jointly (hence underestimate the interactions). To overcome these drawbacks, we leverage the NNGP to efficiently impute all missing features of one observation simultaneously. To this end, we propose a ‘transpose’ trick when using NNGP. We regard each column/feature of X , instead of each row/case of X , as a ‘sample’, so that we draw all features jointly instead of drawing all cases jointly as in the conventional NNGP. As demonstrated in our experiments, this appealing property makes our MI-NNGP methods scalable to high-dimensional data where p can be very large.

As a building block, we first consider imputing the k -th pattern of missing data ($k = 1, \dots, K$). We define a training set $\{(\mathbf{X}_{IS,t}, \mathbf{X}_{P_k,t})\}_{t \in \mathbf{obs}(k)}$ and a test set $\{(\mathbf{X}_{IS,t}, \mathbf{X}_{P_k,t})\}_{t \in \mathbf{mis}(k)}$, where $\mathbf{X}_{IS,t}$ is the input data point and $\mathbf{X}_{P_k,t}$ is the output target. Here the index set ‘IS’ represents the cases included as input, which depends on the specific algorithm used. For example, MI-NNGP1 uses P_1 as the IS set; see details in section 3.2. Of note, our goal is to




Figure 1: Multivariate 4-pattern missing data. Orange squares represent observed data and gray squares represent missing data.

predict the test set label $\{\mathbf{X}_{P_k,t}\}_{t \in \text{mis}(k)}$ which is missing. Denote the size of the training and test set by $|\text{obs}(k)|$ and $|\text{mis}(k)|$, respectively.

Given the training and the test sets, we specify a neural network $f_k : \mathbb{R}^{|\text{IS}|} \rightarrow \mathbb{R}^{|P_k|}$ for the k -th pattern. Therefore, each case in the k -th pattern (say the i -th case, $i \in P_k$) corresponds to an output neuron (say the j -th output component): $f_k^j(\mathbf{X}_{\text{IS},t}) = \mathbf{X}_{i,t}$ for $t \in [p]$, if we assume all observed values are noise-free. By considering infinitely-wide layers, each output component/neuron can be regarded as an independent Gaussian process $\mathcal{GP}(0, \mathcal{K})$ in terms of its input. Here the covariance function \mathcal{K} is uniquely determined by two factors, the neural network architecture (including the activation function) and the initialization variance of weight and bias [Lee et al. \(2017\)](#). Hence, for the j -th output component of the network, f_k^j , we know

$$\begin{bmatrix} f_k^j(\mathbf{X}_{\text{IS},1}) \\ \vdots \\ f_k^j(\mathbf{X}_{\text{IS},p}) \end{bmatrix} \sim \mathcal{N}(0, \Sigma) \quad (1)$$

where $\Sigma \in \mathbb{R}^{p \times p}$ whose (u, v) -th element is $\mathcal{K}(\mathbf{X}_{\text{IS},u}, \mathbf{X}_{\text{IS},v})$. Hence, we get:

$$\begin{bmatrix} \mathbf{X}_{i,\text{obs}(k)}^\top \\ \mathbf{X}_{i,\text{mis}(k)}^\top \end{bmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}\right) \quad (2)$$

where the block structure corresponds to the division between the training and the test sets. Specifically, $\Sigma_{11} = \mathcal{K}(\mathbf{X}_{\text{IS},\text{obs}(k)}, \mathbf{X}_{\text{IS},\text{obs}(k)})$, $\Sigma_{22} = \mathcal{K}(\mathbf{X}_{\text{IS},\text{mis}(k)}, \mathbf{X}_{\text{IS},\text{mis}(k)})$, $\Sigma_{12} = \Sigma_{21}^\top = \mathcal{K}(\mathbf{X}_{\text{IS},\text{obs}(k)}, \mathbf{X}_{\text{IS},\text{mis}(k)})$, where $\Sigma_{11} \in \mathbb{R}^{|\text{obs}(k)| \times |\text{obs}(k)|}$ is composed of $\mathcal{K}(\mathbf{X}_{\text{IS},u}, \mathbf{X}_{\text{IS},v})$ for $u, v \in \text{obs}(k)$. Then, (2) indicates that the missing values $\mathbf{X}_{i,\text{mis}(k)}$, conditioned on the known values (either observed or previously imputed), follow a joint posterior distribution,

$$\begin{aligned} & \mathbf{X}_{i,\text{mis}(k)}^\top \Big| \mathbf{X}_{i,\text{obs}(k)}^\top, \mathbf{X}_{\text{IS},\text{obs}(k)}, \mathbf{X}_{\text{IS},\text{mis}(k)} \sim \\ & \mathcal{N}\left(\Sigma_{21}\Sigma_{11}^{-1}\mathbf{X}_{i,\text{obs}(k)}^\top, \Sigma_{22} - \Sigma_{21}\Sigma_{11}^{-1}\Sigma_{12}\right) \end{aligned} \quad (3)$$

Equation (3) allows us to multiply impute $\mathbf{X}_{i,\text{mis}(k)}$. Here we emphasize that NNGP is not a linear method. The imputed values are drawn from a Gaussian distribution.

Note that inverting Σ_{11} is a common computational challenge. The time complexity is cubic in p . We use the efficient implementation in **neural tangents** [Novak et al. \(2020\)](#) to solve this problem. For the setting in our paper, when $p \sim 10000$, inverting Σ_{11} just cost several seconds with GPU P100 and 16G memory.

3.2. MI-NNGP1 — Direct imputation

Our first algorithm is MI-NNGP1 that uses only the complete cases to impute all missing values. More precisely, to impute the missing values in the k -th pattern, we select P_1 as our IS set. Hence, for each k , we essentially divide all features in the first and the k -th pattern into the training set $\{(\mathbf{X}_{P_1,t}, \mathbf{X}_{P_k,t})\}_{t \in \text{obs}(k)}$ and the test set $\{(\mathbf{X}_{P_1,t}, \mathbf{X}_{P_k,t})\}_{t \in \text{mis}(k)}$ as (input,target) pairs. Following the steps described in Section 3.1, the covariance matrices

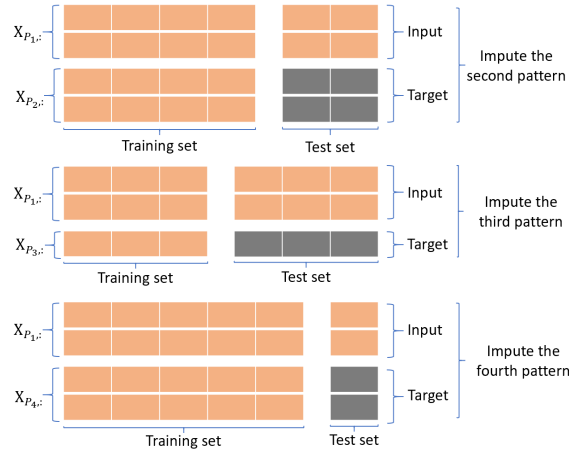


Figure 2: MI-NNGP1 applied to the four-pattern missing data in Figure 1

are

$$\begin{aligned}\Sigma_{11} &= \mathcal{K}(\mathbf{X}_{P_1, \text{obs}(k)}, \mathbf{X}_{P_1, \text{obs}(k)}), \\ \Sigma_{22} &= \mathcal{K}(\mathbf{X}_{P_1, \text{mis}(k)}, \mathbf{X}_{P_1, \text{mis}(k)}), \\ \Sigma_{12} &= \Sigma_{21}^\top = \mathcal{K}(\mathbf{X}_{P_1, \text{obs}(k)}, \mathbf{X}_{P_1, \text{mis}(k)}).\end{aligned}$$

We then draw the imputed values multiple times from the posterior distribution in (3) for all $i \in P_k$. The whole process is summarized in Algorithm 1 and Figure 2. We use the same neural network architecture and the initialization variance of weight and bias for imputing each pattern. Note that the kernel function \mathcal{K} does not depend on the length of the input or that of the output, so the same \mathcal{K} is shared across all patterns. The time complexity of MI-NNGP1 is Kp^3 for imputing a K -pattern missing data.

3.3. MI-NNGP2 — Iterative imputation

In contrast to MI-NNGP1, which imputes each incomplete case basing only on the complete cases, we here propose MI-NNGP2 to impute through an iterative approach that leverages the information contained in incomplete cases. As such, MI-NNGP2 works with a small number of complete cases or even when there is no complete case.

MI-NNGP2 requires an initial imputation $\widehat{\mathbf{X}}$ for the entire data. This can be done by either MI-NNGP1 (if complete cases exist), column mean imputation, or another imputation method. Starting from the initial imputation $\widehat{\mathbf{X}}$, MI-NNGP2 imputes the missing part of each pattern and updates $\widehat{\mathbf{X}}$ iteratively: e.g. the imputed values of k -th pattern is used to impute the missing values of the $(k+1)$ -th pattern. To be more precise, when imputing the k -th pattern, we select P_{-k} as the IS set. Hence, we have the training set $\left\{ \left(\widehat{\mathbf{X}}_{P_{-k}, t}, \widehat{\mathbf{X}}_{P_k, t} \right) \right\}_{t \in \text{obs}(k)}$ and test set $\left\{ \left(\widehat{\mathbf{X}}_{P_{-k}, t}, \widehat{\mathbf{X}}_{P_k, t} \right) \right\}_{t \in \text{mis}(k)}$ of (input, target) pairs.

Then we calculate the covariance matrix

$$\begin{aligned}\Sigma_{11} &= \mathcal{K}(\widehat{\mathbf{X}}_{P_{-k}, \text{obs}(k)}, \widehat{\mathbf{X}}_{P_{-k}, \text{obs}(k)}), \\ \Sigma_{22} &= \mathcal{K}(\widehat{\mathbf{X}}_{P_{-k}, \text{mis}(k)}, \widehat{\mathbf{X}}_{P_{-k}, \text{mis}(k)}), \\ \Sigma_{12} &= \Sigma_{21}^\top = \mathcal{K}(\widehat{\mathbf{X}}_{P_{-k}, \text{obs}(k)}, \widehat{\mathbf{X}}_{P_{-k}, \text{mis}(k)})\end{aligned}$$

and impute the k -th pattern in $\widehat{\mathbf{X}}$ by drawing $\widehat{\mathbf{X}}_{i, \text{mis}(k)}$ from the posterior distribution (3) for each $i \in P_k$. This method is described by Algorithm 2. Similar to MI-NNGP1, \mathcal{K} is shared across all patterns in MI-NNGP2. To conduct multiple imputation, we do not record the imputed values in the first N cycles. After this burn-in period, we choose $\widehat{\mathbf{X}}$ at every T -th iteration.

Algorithm 1 MI-NNGP1: direct imputation

Input: Incomplete matrix \mathbf{X} , imputation times M , neural network architecture, initialization variance of weight and bias

Output: M Imputed matrix

- 1: Calculate the corresponding kernel function \mathcal{K} of the network
 - 2: Create M copies of \mathbf{X} , denoted as $\widehat{\mathbf{X}}^m$, $m \in M$.
 - 3: **for** $k \in \{2, \dots, K\}$ **do**
 - 4: Calculate the matrix $\Sigma_{11}, \Sigma_{21}, \Sigma_{12}$ and Σ_{22} .
 - 5: **for** $i \in P_k$ **do**
 - 6: Draw $\widehat{\mathbf{X}}_{i, \text{mis}(k)}$ for M times from $\mathcal{N}(\Sigma_{21}\Sigma_{11}^{-1}\mathbf{X}_{i, \text{obs}(k)}^\top, \Sigma_{22} - \Sigma_{21}\Sigma_{11}^{-1}\Sigma_{12})$
 - 7: Update $\widehat{\mathbf{X}}^m$ with one $\widehat{\mathbf{X}}_{i, \text{mis}(k)}$ for each $m \in [M]$.
 - 8: **end for**
 - 9: **end for**
 - 10: Output $\widehat{\mathbf{X}}^m$, $m \in [M]$.
-

The time complexity of MI-NNGP2 is $(N + MT)Kp^3$ where M represents imputation times and usually selected as 10. Here N, M, T are bounded by constant and much smaller than K and p . In our experiment, $N = 2$ and $T = 1$ leads to excellent performance. It is important to note that although MI-NNGP2 imputes missing values iteratively, the time cost is expected to increase only modestly compared to MI-NNGP1.

3.4. MI-NNGP with bootstrapping

In the missing data literature, a bootstrap step has been incorporated in nonparametric imputation methods to better account for imputation uncertainty and improve statistical inference. The MI-NNGP methods can also be enhanced by including a bootstrap step. We illustrate this idea for MI-NNGP1. For each incomplete case, MI-NNGP1 essentially draws multiple imputations from the same posterior distribution. However, this may underestimate the uncertainty of imputed values. To overcome this potential drawback, we construct bootstrapping sets of P_1 , denoted as $P_1^{(m)}$ for $m \in [M]$. Each bootstrapping set $P_1^{(m)}$ serves as the IS set for the m -th imputation as visualized in Figure 3. We remark that using the bootstrapping adds negligible additional cost but usually improves the statistical

Algorithm 2 MI-NNGP2: iterative imputation

Input: Initial imputation $\widehat{\mathbf{X}}$, imputation times M , burn-in period N , sampling interval T , neural network architecture, initialization variance of weight and bias

Output: M Imputed matrix

- 1: Calculate the corresponding kernel function \mathcal{K} of the network.
 - 2: **for** $l \in \{1, \dots, N + MT\}$ **do**
 - 3: **for** $k \in \{2, \dots, K\}$ **do**
 - 4: Calculate the matrix $\Sigma_{11}, \Sigma_{21}, \Sigma_{12}$ and Σ_{22} .
 - 5: **for** $i \in P_k$ **do**
 - 6: Draw $\widehat{\mathbf{X}}_{i, \text{mis}(k)}$ from $\mathcal{N}(\Sigma_{21}\Sigma_{11}^{-1}\widehat{\mathbf{X}}_{i, \text{obs}(k)}^\top, \Sigma_{22} - \Sigma_{21}\Sigma_{11}^{-1}\Sigma_{12})$
 - 7: Update $\widehat{\mathbf{X}}$ with $\widehat{\mathbf{X}}_{i, \text{mis}(k)}$
 - 8: **end for**
 - 9: **end for**
 - 10: **if** $l > N$ and $T|(l - N)$ **then**
 - 11: output $\widehat{\mathbf{X}}$
 - 12: **end if**
 - 13: **end for**
-

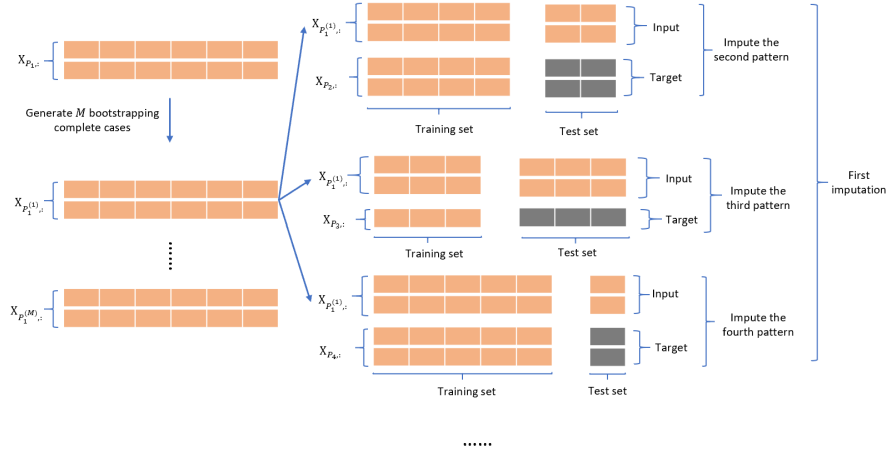


Figure 3: MI-NNGP1 with bootstrapping applied to the four-pattern missing data in Figure 1

coverage. Similarly, a bootstrap step can also be combined with MI-NNGP2. We can first use MI-NNGP1 with bootstrapping to generate multiple initial imputations and then run MI-NNGP2 multiple times from these initial imputations, where we choose $M = 1$ in each track of MI-NNGP2.

4. Experiments

We evaluate the performance of the MI-NNGP methods through extensive synthetic and real data experiments. The details about the experiment setup are provided in Appendix B and C. A brief outline of the synthetic data experiments is as follows. In each synthetic data experiment, we generate the data matrix from a pre-specified data model and then generate missing values under MCAR, MAR or MNAR. We apply an imputation method to each incomplete dataset; SI methods yield one imputed dataset and MI methods yield multiple imputed datasets. To assess the statistical inference performance, each imputed dataset is used to fit a regression model to obtain regression coefficient estimates and Rubin’s rule [Little and Rubin \(2019\)](#) is used to obtain the final regression coefficient estimates $\hat{\beta}$, their standard errors $SE(\hat{\beta})$, and 95% confidence intervals.

4.1. Imputation Methods Compared

Benchmarks. (i) **Complete data** analysis assumes there is no missingness and directly fit a regression on the whole dataset. (ii) **Complete case** analysis does not incorporate imputation and fit a regression using only the complete cases. (iii) Column mean imputation (**ColMean Imp**) is feature-wise mean imputation. Here the complete data analysis serves as a golden standard, representing the best result an imputation method can possibly achieve. The complete case analysis and column mean imputation, two naive methods, are used to benchmark potential bias and loss of information (as represented by larger SE/SD) under MAR and MNAR.

State-of-the-art. (iv) **MICE** (multiple imputation through chained equations) [Van Buuren \(2007\)](#) is a popular and flexible multiple imputation method and has good empirical results and requires little tuning, but it fails to scale to high dimensional settings. (v) **GAIN** [Yoon et al. \(2018\)](#) is a generative neural network (GAN) [Goodfellow et al. \(2014\)](#) based imputation method. (vi) **SoftImpute** [Mazumder et al. \(2010\)](#) is a matrix completion method and uses iterative soft-thresholded SVD to conduct missing data imputation. (vii) **Sinkhorn Muzellec et al. (2020)** is a direct non-parametric imputation method which leverages optimal transport distance. (viii) **Linear RR** [Muzellec et al. \(2020\)](#) is a Round-Robin Sinkhorn Imputation. Similar to MICE, Linear RR iteratively impute missing features using other features in a cyclical manner. (ix) **MIWAE** [Mattei and Frelsen \(2019\)](#) is a importance weighted autoencoder [Burda et al. \(2015\)](#) (IWAE) based imputation method.

Our methods. (x) **MI-NNGP1** uses the complete cases to conduct direct imputation as detailed in Algorithm 1. (xi) **MI-NNGP2** corresponds to Algorithm 2 with burn-in period $N = 10$ and sampling interval $T = 1$. (xii) **MI-NNGP1-BS** is MI-NNGP1 with an added bootstrap step. (xiii) **MI-NNGP2-BS** runs MI-NNGP2 for multiple times with different initial imputations from MI-NNGP1-BS. In each parallel run of MI-NNGP2, we choose $N = 2$ and $M = 1$.

4.2. Performance Metrics

All performance metrics are averaged over 100 Monte Carlo (MC) datasets or repeats unless noted otherwise. To evaluate imputation accuracy and computational costs, we report the imputation mean squared error (Imp MSE) and the computing time in seconds per

imputation (Time(s)). To evaluate statistical inference performance, we report bias of $\hat{\beta}_1$ denoted by $\text{Bias}(\hat{\beta}_1)$, standard error of $\hat{\beta}_1$ denoted by $\text{SE}(\hat{\beta}_1)$, and coverage rate of the 95% confidence interval for β_1 denoted by $\text{CR}(\hat{\beta}_1)$, where $\hat{\beta}_1$ is one of the regression coefficients in the regression model fitted using imputed datasets. Some remarks are in order. A $\text{CR}(\hat{\beta}_1)$ that is well below the nominal level of 95% would lead to inflated false positives, an important factor contributing to lack of reproducibility in research. To benchmark $\text{SE}(\hat{\beta}_1)$, we also report the standard deviation of $\hat{\beta}_1$ across 100 MC datasets denoted by $\text{SD}(\hat{\beta}_1)$, noting that a well-behaved $\text{SE}(\hat{\beta}_1)$ should be close to $\text{SD}(\hat{\beta}_1)$. In addition, while we know the true value of β_1 and can report its bias in the synthetic data experiments, we do not know the true value of β_1 and cannot report its bias in the real data experiment.

4.3. Synthetic data

The synthetic data experiments are conducted for low and high data dimensions, varying missing rates, and continuous and discrete data. In this section, we summarize the results from high dimensional settings (i.e. $p > n$) under MAR. Additional simulation results are included in appendix.

Models	Style	Time(s)	Imp MSE	Bias($\hat{\beta}_1$)	CR($\hat{\beta}_1$)	SE($\hat{\beta}_1$)	SD($\hat{\beta}_1$)
SoftImpute	SI	15.1	0.0200	-0.0913	0.78	0.1195	0.1624
GAIN	SI	39.0	0.8685	0.6257	0.18	0.1463	0.5424
MIWAE	SI	46.3	0.0502	0.0731	0.90	0.1306	0.1379
Linear RR	SI	3134.7	0.0661	0.1486	1.00	0.1782	0.1011
MICE	MI	37.6	0.0234	-0.0061	0.93	0.1167	0.1213
Sinkhorn	MI	31.2	0.0757	0.0205	0.96	0.1864	0.1636
MI-NNGP1	MI	4.9	0.0116	0.0077	0.92	0.1147	0.1223
MI-NNGP1-BS	MI	3.4	0.0149	0.0156	0.96	0.1297	0.1182
MI-NNGP2	MI	5.7	0.0086	0.0012	0.96	0.1179	0.1170
MI-NNGP2-BS	MI	13.9	0.0094	0.0010	0.95	0.1173	0.1206
Complete data	-	-	-	-0.0027	0.90	0.1098	0.1141
Complete case	-	-	-	0.2481	0.88	0.3400	0.3309
ColMean Imp	SI	-	0.1414	0.3498	0.72	0.2212	0.1725

Table 2: Gaussian data with $n = 200$ and $p = 251$ under MAR. Approximately **40%** features and **90%** cases contain missing values. Detailed simulation setup information is in appendix.

Table 2 presents the results for Gaussian data with $n = 200$ and $p = 251$ under MAR. The MI-NNGP methods yield smallest imputation error (Imp MSE) compared to the other methods. In terms of statistical inference, the MI-NNGP methods, MICE and Sinkhorn, all of which are MI methods, lead to small to negligible bias in $\hat{\beta}_1$. The CR for MI-NNGP1-BS, MI-NNGP2, MI-NNGP2-BS, and Sinkhorn is close to the nominal level of 95% and their $\text{SE}(\hat{\beta}_1)$ is close to $\text{SD}(\hat{\beta}_1)$, suggesting that Rubin’s rule works well for these MI methods. Of these methods, our MI-NNGP methods and MICE outperform Sinkhorn in terms of information recovery, as evidenced by their smaller SE compared to Sinkhorn. SoftImpute, Linear RR, MIWAE and GAIN, four SI methods, yield poor performance in statistical inference with considerable bias for $\hat{\beta}_1$ and CR away from the nominal level of 95%. In addition, GAIN yields substantially higher imputation error than the other methods. In

Models	Style	Time(s)	Imp MSE	Bias($\hat{\beta}_1$)	CR($\hat{\beta}_1$)	SE($\hat{\beta}_1$)	SD($\hat{\beta}_1$)
SoftImpute	SI	30.1	0.0442	-0.2862	0.50	0.1583	0.2019
GAIN	SI	111.1	0.7383	0.6897	0.18	0.1697	0.5693
MIWAE	SI	52.5	0.1228	0.5885	0.15	0.1793	0.2162
Sinkhorn	MI	39.3	0.1031	0.6647	0.26	0.2643	0.2195
MI-NNGP1	MI	4.9	0.0119	0.0351	0.89	0.1194	0.1422
MI-NNGP1-BS	MI	4.9	0.0168	0.0383	0.94	0.1424	0.1416
MI-NNGP2	MI	5.8	0.0086	0.0487	0.90	0.1212	0.1343
MI-NNGP2-BS	MI	13.9	0.0092	0.0347	0.93	0.1257	0.1289
Complete data	-	-	-	0.0350	0.94	0.1122	0.1173
Complete case	-	-	-	0.2804	0.76	0.3466	0.4211
ColMean Imp	SI	-	0.1130	0.7024	0.13	0.2574	0.1957

Table 3: Gaussian data with $n = 200$ and $p = 1001$ under MAR. Approximately **40%** features and **90%** cases contain missing values. Linear RR and MICE are not included due to running out of RAM. Detailed simulation setup information is in appendix.

terms of computation, our MI-NNGP methods are the least expensive, whereas Linear RR is the most expensive.

Table 3 presents the results for Gaussian data with $n = 200$ and $p = 1001$ under MAR. As p increases to 1001 from 251 in Table 2, the performance of Sinkhorn deteriorates significantly; Linear RR and MICE run out of RAM; GAIN’s performance remains poor. At the same time, our MI-NNGP methods continue to yield the most satisfactory performance. In particular, MI-NNGPs have smallest imputation error in this setting. In addition, CR($\hat{\beta}_1$) for MI-NNGP with a bootstrap step is closer to the nominal level than MI-NNGP without a bootstrap step, suggesting the bootstrap step indeed improves quantification of uncertainty of imputed values. Also, the computational time for MI-NNGP methods does not increase much as p increases from 251 to 1001, demonstrating that they are scalable to ultra high-dimensional p —a very appealing property. This is because MI-NNGP imputes the set of features with missing values in each pattern jointly, whereas other MI methods such as MICE impute each feature iteratively.




Figure 4: Left: Imputation MSE for varying missing rates. Middle: Bias of $\hat{\beta}_1$ for varying missing rates. Right: Empirical distribution of $\hat{\beta}_1$ from 200 MC datasets when the missing rate is 40%.

Additional results in appendix include synthetic data experiments for small p under MAR (Table 5), for MNAR (Table 7, Table 8, and Table 9), for mix of Gaussian continuous and **discrete** data (Table 11, Table 12), and for **non-Gaussian** continuous data (Table 6 and Table 10). These and other unreported results for MCAR consistently show that

Models	Style	Time(s)	Imp MSE	$\hat{\beta}_1$	SE($\hat{\beta}_1$)
SoftImpute	SI	1008.6	0.0591	0.0213	0.0119
Sinkhorn	MI	843.3	0.0797	0.0223	0.0128
MI-NNGP1	MI	7.1	0.0637	0.0161	0.0104
MI-NNGP1-BS	MI	7.8	0.0685	0.0153	0.0112
MI-NNGP2	MI	11.8	0.0617	0.0171	0.0106
MI-NNGP2-BS	MI	21.3	0.0640	0.0166	0.0110
Complete data	-	-	-	0.0160	0.0085
Complete case	-	-	-	0.0221	0.0185
ColMean Imp	SI	-	0.1534	0.0188	0.0136

Table 4: Real data experiment with $n = 649$ and $p = 10001$ under MAR. Approximately **20%** features and **76%** cases contain missing values. Linear RR, MICE, MIWAE and GAIN are not included due to running out of RAM. Detailed experiment setup information is in appendix.

the MI-NNGP methods outperform the competing state-of-the-art imputation methods, particularly in high-dimensional settings. Of the four MI-NNGP methods, MI-NNGP2-BS offers the best or close to the best performance in all experiments.

To further investigate the impact of varying missing rates on the performance of MI-NNGP2-BS, Figure 4 presents the results from additional experiments under MAR for $n = 200$ and $p = 1001$ in which Sinkhorn and SoftImpute, two closest competitors based on the prior experiments, are also included. As shown in Figure 4, MI-NNGP2-BS always yields the best performance in terms of imputation error and bias of $\hat{\beta}_1$ and is more robust to high missing rates.

4.4. ADNI data

We evaluate the performance of MI-NNGPs using a publicly available, de-identified large-scale dataset from the Alzheimer’s Disease Neuroimaging Initiative (ADNI), containing both image data and gene expression data. This dataset has over 19,000 features and a response variable (y), VBM right hippocampal volume, for 649 patients. The details of the real data experiment are included in appendix. Briefly, we select 10000 centered features and generate missing values under MAR or MNAR. After imputation, we fit a linear regression for y on three features that have highest correlation with response using imputed datasets. Table 4 presents the results under MAR for estimating β_1 , one of the regression coefficients in the linear regression model, as well as the computational time. Again, since we do not know the true value of β_1 , we cannot report its bias and instead we use $\hat{\beta}_1$ from the complete data analysis as a gold standard. The results in Table 4 show that $\hat{\beta}_1$ from the MI-NNGP methods is considerably closer to that from the complete data analysis than the other imputation methods, demonstrating their superior performance. In addition, SE($\hat{\beta}_1$) for MI-NNGP methods is fairly close to that for the complete data analysis and much smaller than that from the complete case analysis, suggesting that our imputation methods results in very limited information loss. In terms of computational costs, SoftImpute and Sinkhorn are much more expensive than MI-NNGP, whereas Linear RR, MICE, MIWAE and GAIN run out of memory. Additional real data experiment results in appendix under MNAR also demonstrate the superiority of our MI-NNGP methods over the existing methods.

5. Discussion

In this work, we develop powerful NNGP-based multiple imputation methods for high dimensional incomplete data with large p and moderate n that are also robust to high missing rates. Our experiments demonstrate that the MI-NNGP methods outperform the current state-of-the-art methods in Table 1 under MCAR, MAR and MNAR. One limitation of the MI-NNGP is that it does not scale well when p becomes extremely large. To overcome this, we can take advantage of recent developments on efficient algorithms for scalable GP computation [Huang et al. \(2015\)](#); [Liu et al. \(2020\)](#). Instead of calculating the GP, we can approximate the GP and keep a balance between performance and computational complexity. This is our future research interest.

References

- Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. *arXiv preprint arXiv:1509.00519*, 2015.
- Youngmin Cho and Lawrence K Saul. Kernel methods for deep learning. In *Advances in neural information processing systems*, pages 342–350, 2009.
- Zongyu Dai, Zhiqi Bu, and Qi Long. Multiple imputation via generative adversarial network for high-dimensional blockwise missing value problems. In *2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 791–798, 2021. doi: 10.1109/ICMLA52953.2021.00131.
- Yi Deng, Changgee Chang, Moges Seyoum Ido, and Qi Long. Multiple imputation for general missing data patterns in the presence of high-dimensional data. *Scientific reports*, 6(1):1–10, 2016.
- Pedro J García-Laencina, José-Luis Sancho-Gómez, and Aníbal R Figueiras-Vidal. Pattern classification with missing data: a review. *Neural Computing and Applications*, 19(2): 263–282, 2010.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- Trevor Hastie, Rahul Mazumder, Jason D Lee, and Reza Zadeh. Matrix completion and low-rank svd via fast alternating least squares. *The Journal of Machine Learning Research*, 16(1):3367–3402, 2015.
- Wenbing Huang, Deli Zhao, Fuchun Sun, Huaping Liu, and Edward Chang. Scalable gaussian process regression using deep neural networks. In *Twenty-fourth international joint conference on artificial intelligence*, 2015.
- Niels Bruun Ipsen, Pierre-Alexandre Mattei, and Jes Frellsen. not-miwa: Deep generative modelling with missing not at random data. *arXiv preprint arXiv:2006.12871*, 2020.

- Oleg Ivanov, Michael Figurnov, and Dmitry Vetrov. Variational autoencoder with arbitrary conditioning. *arXiv preprint arXiv:1806.02382*, 2018.
- Dongwook Lee, Junyoung Kim, Won-Jin Moon, and Jong Chul Ye. Collagan: Collaborative gan for missing image data imputation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2487–2496, 2019.
- Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. Deep neural networks as gaussian processes. *arXiv preprint arXiv:1711.00165*, 2017.
- Serena G Liao, Yan Lin, Dongwan D Kang, Divay Chandra, Jessica Bon, Naftali Kaminski, Frank C Scieurba, and George C Tseng. Missing value imputation in high-dimensional phenomic data: imputable or not, and how? *BMC bioinformatics*, 15(1):346, 2014.
- Roderick JA Little and Donald B Rubin. *Statistical analysis with missing data*, volume 793. John Wiley & Sons, 2019.
- Haitao Liu, Yew-Soon Ong, Xiaobo Shen, and Jianfei Cai. When gaussian process meets big data: A review of scalable gps. *IEEE transactions on neural networks and learning systems*, 31(11):4405–4423, 2020.
- Pierre-Alexandre Mattei and Jes Frellesen. Miwae: Deep generative modelling and imputation of incomplete data sets. In *International Conference on Machine Learning*, pages 4413–4423. PMLR, 2019.
- Rahul Mazumder, Trevor Hastie, and Robert Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. *The Journal of Machine Learning Research*, 11: 2287–2322, 2010.
- Boris Muzellec, Julie Josse, Claire Boyer, and Marco Cuturi. Missing data imputation using optimal transport. *arXiv preprint arXiv:2002.03860*, 2020.
- Roman Novak, Lechao Xiao, Jaehoon Lee, Yasaman Bahri, Greg Yang, Jiri Hron, Daniel A Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein. Bayesian deep convolutional networks with many channels are gaussian processes. *arXiv preprint arXiv:1810.05148*, 2018.
- Roman Novak, Lechao Xiao, Jiri Hron, Jaehoon Lee, Alexander A. Alemi, Jascha Sohl-Dickstein, and Samuel S. Schoenholz. Neural tangents: Fast and easy infinite neural networks in python. In *International Conference on Learning Representations*, 2020. URL <https://github.com/google/neural-tangents>.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- Joseph L Schafer. *Analysis of incomplete multivariate data*. CRC press, 1997.

- Daniel J Stekhoven and Peter Bühlmann. Missforest—non-parametric missing value imputation for mixed-type data. *Bioinformatics*, 28(1):112–118, 2012.
- Stef Van Buuren. Multiple imputation of discrete and continuous data by fully conditional specification. *Statistical methods in medical research*, 16(3):219–242, 2007.
- Christopher KI Williams. Computing with infinite networks. In *Advances in neural information processing systems*, pages 295–301, 1997.
- Jinsung Yoon, James Jordon, and Mihaela Van Der Schaar. Gain: Missing data imputation using generative adversarial nets. *arXiv preprint arXiv:1806.02920*, 2018.
- Yize Zhao and Qi Long. Multiple imputation in the presence of high-dimensional data. *Statistical Methods in Medical Research*, 25(5):2021–2035, 2016.