

# UNLEASHING THE POTENTIAL OF DIFFUSION MODELS FOR INCOMPLETE DATA IMPUTATION

Anonymous authors

Paper under double-blind review

## ABSTRACT

Generative models play an important role in missing data imputation in that they aim to learn the joint distribution of full data. However, applying advanced deep generative models (such as Diffusion models) to missing data imputation is challenging due to 1) the inherent incompleteness of the training data and 2) the difficulty in performing conditional inference from unconditional generative models. To deal with these challenges, this paper introduces DIFFPUTER, a tailored diffusion model combined with the Expectation-Maximization (EM) algorithm for missing data imputation. DIFFPUTER iteratively trains a diffusion model to learn the joint distribution of missing and observed data and performs an accurate conditional sampling to update the missing values using a tailored reversed sampling strategy. Our theoretical analysis shows that DIFFPUTER’s training step corresponds to the maximum likelihood estimation of data density (M-step), and its sampling step represents the Expected A Posteriori estimation of missing values (E-step). Extensive experiments across ten diverse datasets and comparisons with 17 different imputation methods demonstrate DIFFPUTER’s superior performance. Notably, DIFFPUTER achieves an average improvement of 8.10% in MAE and 5.64% in RMSE compared to the most competitive existing method.

## 1 INTRODUCTION

In the field of data science and machine learning, missing data in tabular datasets is a common issue that can severely impair the performance of predictive models and the reliability of statistical analysis. Missing data can result from various factors, including data entry errors, non-responses in surveys, and system errors during data collection (Barnard & Meng, 1999; Lillard et al., 1986; Eckert et al., 2020). Properly handling missing data is essential, as improper treatment can lead to biased estimates, reduced statistical power, and invalid conclusions.

A plethora of work proposed over the past decades has propelled the development of missing data imputation research. Early classical methods often relied on partially observed statistical features to impute missing values or were based on conventional machine learning techniques, such as KNN (Pujianto et al., 2019), or simple parametric models, such as Bayesian models (Ma & Chen, 2018) or Gaussian Mixture Models (García-Laencina et al., 2010). With the advent of deep learning, recent research has primarily focused on predictive (Stekhoven & Bühlmann, 2012; Van Buuren & Karin, 2011; Kyono et al., 2021) or generative deep models (Yoon et al., 2018; Mattei & Frellsen, 2019; Richardson et al., 2020) for missing data imputation. Predictive models learn to predict the target entries conditioned on other observed entries, guided by masking mechanisms (Du et al., 2024) or graph regularization techniques (You et al., 2020; Zhong et al., 2023). By contrast, generative methods learn the joint distribution of missing entries and observed entries and aim to impute the missing data via conditional sampling (Richardson et al., 2020; Nazabal et al., 2020; Mattei & Frellsen, 2019; Yoon et al., 2018; Ouyang et al., 2023; Zheng & Charoenphakdee, 2022). However, generative imputation methods still fall short compared to predictive methods even when employing state-of-the-art generative models (Kingma & Welling, 2013; Rezende & Mohamed, 2015; Goodfellow et al., 2014; Ho et al., 2020). We think this is primarily due to the **incomplete likelihood** nature issue in missing data imputation: generative models need to estimate the joint distribution of missing data and observed data. However, since the missing data itself is unknown, there is an inherent error in the estimated data density. The classical Expectation-Maximization (EM) algorithm (Dempster

et al., 1977) offers an elegant route to handle this issue, being capable of addressing the incomplete likelihood issue by iteratively refining the values of the missing data.

However, integrating the EM algorithm with deep generative models (DGMs) has been less explored, probably because existing deep generative models (Xu et al., 2019; Liu et al., 2023; Zhang et al., 2024) struggle to perform both the M-step and E-step of EM effectively simultaneously. 1) In the M-step, we need deep generative models to faithfully **reconstruct the ground-truth distribution**. 2) While in the **E-step**, we want deep generative models to **perform accurate conditional inference**. Although existing DGM-based imputation methods (Mattei & Frellsen, 2019) can conveniently perform conditional inference, they struggle to learn accurate tabular data distributions (Xu et al., 2019; Liu et al., 2023). Meanwhile, existing SOTA tabular generation models, such as diffusion-based models (Kotelnikov et al., 2023; Zhang et al., 2024), model the joint distribution of the entire table, making it difficult to infer the distribution of remaining columns based on partial columns.

This paper introduces DIFFPUTER, a principled generative method for missing data imputation. DIFFPUTER explores a path that makes diffusion models compatible with the EM framework, allowing both E-step and M-step to be effective. Specifically: 1) In the M-step, DIFFPUTER employs a diffusion model to learn the joint distribution of the missing and observed data. With the powerful ability of diffusion models to learn tabular data distributions, in the M-step, DIFFPUTER learns a parameterized distribution that faithfully recovers the joint distribution of observed and missing entries. 2) In the E-step, DIFFPUTER uses the learned diffusion model to perform flexible and accurate conditional sampling by mixing the forward process for observed entries with the reverse process for missing entries. Theoretically, we show that DIFFPUTER’s M-step corresponds to the maximum likelihood estimation of the data density, while its E-step represents the *Expected A Posteriori* (EAP) estimation of the missing values, conditioned on the observed values.

We conduct experiments on ten benchmark tabular datasets containing both continuous and discrete features under various missing data scenarios. We compare the performance of DIFFPUTER with 17 competitive imputation methods from different categories. Experimental results demonstrate the superior performance of DIFFPUTER across all settings and on almost all datasets. In addition, experiments demonstrate that DIFFPUTER’s iterative training can effectively and gradually reduce the error in density estimation, thereby improving the performance of imputation. Ablation studies also demonstrate DIFFPUTER’s robustness to different missing mechanisms. We have released the code to reproduce the results of the proposed DIFFPUTER and all the baseline methods. The code is available at <https://anonymous.4open.science/r/ICLR-DiffPutter>.

## 2 RELATED WORKS

**Iterative Methods for Missing Data Imputation.** Iterative imputation is a widely used approach due to its ability to continuously refine predictions of missing data, resulting in more accurate imputation outcomes. This iterative process is especially crucial for methods requiring an initial estimation of the missing data. The Expectation-Maximization (EM) algorithm (Dempster et al., 1977), a classical method, can be employed for missing data imputation. However, earlier applications often assume simple data distributions, such as mixtures of Gaussians for continuous data or Bernoulli and multinomial densities for discrete data (García-Laencina et al., 2010). These assumptions limit the imputation capabilities of these methods due to the restricted density estimation of simple distributions. The integration of EM with deep generative models remains underexplored. A closely related approach is MCFlow (Richardson et al., 2020), which iteratively imputes missing data using normalizing flows (Rezende & Mohamed, 2015). However, MCFlow focuses on recovering missing data through maximum likelihood rather than expectation, and its conditional imputation is achieved through soft regularization instead of precise sampling based on the conditional distribution. Beyond EM, the concept of iterative training is prevalent in state-of-the-art deep learning-based imputation methods. For instance, IGRM (Zhong et al., 2023) constructs a graph from all dataset samples and introduces the concept of friend networks, which are iteratively updated during the imputation learning process. HyperImpute (Jarrett et al., 2022) proposes an AutoML imputation method that iteratively refines both model selection and imputed values.

**Diffusion Models for Missing Data Imputation.** We are not the first to utilize diffusion models for missing data imputation. In computer vision, diffusion models have been widely applied in image

inpainting, either in the data space (Lugmayr et al., 2022) or the latent space (Corneanu et al., 2024). In the tabular data area, TabCSDI (Zheng & Charoenphakdee, 2022) employs a conditional diffusion model to learn the distribution of masked observed entries conditioned on the unmasked observed entries. MissDiff (Ouyang et al., 2023) uses a diffusion model to learn the density of tabular data with missing values by masking the observed entries. Although MissDiff was not originally intended for imputation, it can be easily adapted for this task. Other methods, despite claiming applicability for imputation, are trained on complete data and evaluated on incomplete testing data (Zhang et al., 2024; Jolicoeur-Martineau et al., 2024a). This approach contradicts the focus of this study, where the training data itself contains missing values. Additionally, all the aforementioned methods use one-step imputations, which overlook the issue that missing data in the training set can lead to inaccurate data density estimation. **By contrast, the proposed DIFFPUTER is the first to integrate a diffusion-based generative model into the EM framework. Additionally, we achieved accurate conditional sampling by mixing the forward and reverse processes of diffusion and demonstrated the effectiveness of this approach through theoretical analysis.**

### 3 PRELIMINARIES

#### 3.1 MISSING VALUE IMPUTATION FOR INCOMPLETE DATA

This paper addresses the missing value imputation task for incomplete data, where only partial data entries are observable during the training process. Formally, let the complete  $d$ -dimensional data be denoted as  $\mathbf{x} \sim p_{\text{data}}(\mathbf{x}) \in \mathbb{R}^d$ . For each data sample,  $\mathbf{x}$ , there is a binary mask  $\mathbf{m} \in \{0, 1\}^d$  indicating the location of missing entries for  $\mathbf{x}$ . Let the subscript  $k$  denote the  $k$ -th entry of a vector, then  $m_k = 1$  stands for missing entries while  $m_k = 0$  stands for observable entries.

We further use  $\mathbf{x}^{\text{obs}}$  and  $\mathbf{x}^{\text{mis}}$  to denote the observed data and missing data, respectively (i.e.,  $\mathbf{x} = (\mathbf{x}^{\text{obs}}, \mathbf{x}^{\text{mis}})$ ). Note that  $\mathbf{x}^{\text{obs}}$  is the fixed ground-truth observation, while  $\mathbf{x}^{\text{mis}}$  is conceptual and unknown, and we aim to estimate it. The missing value imputation task aims to predict the missing entries  $\mathbf{x}^{\text{mis}}$  based on the observed entries  $\mathbf{x}^{\text{obs}}$ .

**In-sample vs. Out-of-sample imputation.** The missing data imputation task can be categorized into two types: in-sample and out-of-sample. In-sample imputation means that the model only has to impute the missing entries in the training set, while out-of-sample imputation requires the model’s capacity to generalize to the unseen data records without fitting its parameters again. Not all imputation methods can generalize to out-of-sample imputation tasks. For example, methods that directly treat the missing values as learnable parameters (Muzellec et al., 2020; Zhao et al., 2023) are hard to apply to unseen records. A desirable imputation method is expected to perform well on both in-sample and out-of-sample imputation tasks, and this paper studies both of the two settings.

#### 3.2 FORMULATING MISSING DATA IMPUTATION WITH EXPECTATION-MAXIMIZATION

Treating  $\mathbf{x}^{\text{obs}}$  as observed variables and  $\mathbf{x}^{\text{mis}}$  as latent variables, with the estimated density of the complete data distribution parameterized as  $p_{\theta}(\mathbf{x})$ , we can formulate the missing value imputation problem using the Expectation-Maximization (EM) algorithm. Specifically, when the complete data distribution  $p_{\theta}(\mathbf{x})$  is available, the optimal estimation of missing values is given by  $\mathbf{x}^{\text{mis}*} = \mathbb{E}_{\mathbf{x}^{\text{mis}}} p(\mathbf{x}^{\text{mis}} | \mathbf{x}^{\text{obs}}, \theta)$ . Conversely, when the missing entries  $\mathbf{x}^{\text{mis}}$  are known, the density parameters can be optimized via maximum likelihood estimation:  $\theta^* = \arg \max_{\theta} p(\mathbf{x} | \theta)$ . Consequently, with the initial estimation of missing values  $\mathbf{x}^{\text{mis}}$ , the model parameters  $\theta$  and the missing value  $\mathbf{x}^{\text{mis}}$  can be optimized by iteratively applying M-step and E-step:

- **Maximization-step:** Fix  $\mathbf{x}^{\text{mis}}$ , update  $\theta^* = \arg \max_{\theta} p(\mathbf{x} | \theta) = \arg \max_{\theta} p_{\theta}(\mathbf{x}^{\text{obs}}, \mathbf{x}^{\text{mis}})$ .
- **Expectation-step:** Fix  $\theta$ , update  $\mathbf{x}^{\text{mis}*} = \mathbb{E}_{\mathbf{x}^{\text{mis}}} p(\mathbf{x}^{\text{mis}} | \mathbf{x}^{\text{obs}}, \theta)$ .




Figure 1: An overview of the architecture of the proposed DIFFPUTER. DIFFPUTER utilizes one-hot encoding to transform discrete variables into continuous ones and use the mean of observed values to initialize the missing entries. The EM algorithm alternates the process of 1) fixing  $x^{mis}$  and estimate diffusion model parameter  $\theta$ , 2) fixing  $\theta$  and estimate  $x^{mis}$ , for  $K$  iterations. The final imputation result  $x^*$  is returned from the E-step of the last iteration.

## 4 METHODOLOGY

In this section, we introduce DIFFPUTER- Iterative Missing Data **Im**putation with **Diff**usion. Based on the Expectation-Maximization (EM) algorithm, DIFFPUTER updates the density parameter  $\theta$  and hidden variables  $x^{mis}$  in an iterative manner. Fig. 1 shows the overall architecture and training process of DIFFPUTER: 1) The M-step fixes the missing entries  $x^{mis}$ , then a diffusion model is trained to estimate the density of the complete data distribution  $p_\theta(x) = p_\theta(x^{obs}, x^{mis})$ ; 2) The E-step fixes the model parameters  $\theta$ , then we update the missing entries  $x^{mis}$  via the reverse process of the learned diffusion model  $p_\theta(x)$ . The above two steps are executed iteratively until convergence. The following sections introduce the M-step and E-step of DIFFPUTER, respectively. To avoid confusion, we use  $x, x_t, x^{obs}$ , etc. to denote samples from real data,  $\tilde{x}, \tilde{x}_t, \tilde{x}^{obs}$ , etc. to denote samples obtained by the model  $\theta$ , while  $\hat{x}, \hat{x}_t, \hat{x}^{obs}$ , etc. to denote the specific values of variables.

### 4.1 M-STEP: DENSITY ESTIMATION WITH DIFFUSION MODELS

Given an estimation of complete data  $x = (x^{obs}, x^{mis})$ , M-step aims to learn the density of  $x$ , parameterized by model  $\theta$ , i.e.,  $p_\theta(x)$ . Inspired by the impressive generative modeling capacity of diffusion models (Song et al., 2021b; Karras et al., 2022), DIFFPUTER learns  $p_\theta(x)$  through a diffusion process, which consists of a forward process that gradually adds Gaussian noises of increasing scales to  $x$ , and a reverse process that recovers the clean data from the noisy one:

$$x_t = x_0 + \sigma(t)\epsilon, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (\text{Forward Process}) \quad (1)$$

$$dx_t = -2\dot{\sigma}(t)\sigma(t)\nabla_{x_t} \log p(x_t)dt + \sqrt{2\dot{\sigma}(t)\sigma(t)}d\omega_t, \quad (\text{Reverse Process}) \quad (2)$$

In the forward process,  $x_0 = x = (x^{obs}, x^{mis})$  is the currently estimated data at time 0, and  $x_t$  is the diffused data at time  $t$ .  $\sigma(t) = t$  is the noise level (and  $\dot{\sigma}(t)$  is its derivative w.r.t.  $t$ ), i.e., the standard deviation of Gaussian noise, at time  $t$ . The forward process has defined a series of data distribution  $p(x_t) = \int_{x_0} p(x_t|x_0)p(x_0)dx_0$ , and  $p(x_0) = p(x)$ . Note that when restricting the mean of  $x_0$  to 0 and keeping the variance of  $x_0$  small (e.g., via standardization),  $p(x_t)$  approaches a tractable prior distribution  $\pi(x)$  at  $t = T$  when  $\sigma(T)$  is large enough, meaning  $p(x_T) \approx \pi(x)$  (Song et al., 2021a). In our formulation in Eq. 1,  $p(x_T) \approx \pi(x) = \mathcal{N}(\mathbf{0}, \sigma^2(T)\mathbf{I})$ .

In the reverse process,  $\nabla_{x_t} \log p(x_t)$  is the gradient of  $x_t$ 's log-probability w.r.t., to  $x_t$ , and is also known as the *score function*.  $\omega_t$  is the standard Wiener process. The model is trained by (conditional) score matching (Song et al., 2021b), which utilizes a neural network  $\epsilon_\theta(x_t, t)$  (called denoising/score network) to approximate the conditional score-function  $\nabla_{x_t} \log p(x_t|x_0)$ , which in expectation approximates  $\nabla_{x_t} \log p(x_t)$ :

$$\mathcal{L}_{\text{SM}} = \mathbb{E}_{\mathbf{x}_0 \sim p(\mathbf{x}_0)} \mathbb{E}_{t \sim p(t)} \mathbb{E}_{\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \|\boldsymbol{\varepsilon}_{\boldsymbol{\theta}}(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t | \mathbf{x}_0)\|_2^2, \text{ where } \mathbf{x}_t = \mathbf{x}_0 + \sigma(t)\boldsymbol{\varepsilon}. \quad (3)$$

Since the score of conditional distribution has analytical solutions, i.e.,  $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t | \mathbf{x}_0) = \frac{\nabla_{\mathbf{x}_t} p(\mathbf{x}_t | \mathbf{x}_0)}{p(\mathbf{x}_t | \mathbf{x}_0)} = -\frac{\mathbf{x}_t - \mathbf{x}_0}{\sigma^2(t)} = -\frac{\boldsymbol{\varepsilon}}{\sigma(t)}$ , Eq. 3 can be interpreted as training a neural network  $\boldsymbol{\varepsilon}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)$  to approximate the scaled noise. Therefore,  $\boldsymbol{\varepsilon}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)$  is also known as the denoising function, and in this paper, it is implemented as a five-layer MLP (see Appendix D.4).

**Remark 1.** Starting from the prior distribution  $\mathbf{x}_T \sim \pi(\mathbf{x})$ , and apply the reverse process in Eq. 2 with  $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)$  replaced with  $\boldsymbol{\varepsilon}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)$ , we obtain a series of distributions  $p_{\boldsymbol{\theta}}(\mathbf{x}_t)$ ,  $\forall t \in [0, T]$

**Remark 2.** (Corollary 1 in (Song et al., 2021a)) Let  $p(\mathbf{x}) = p(\mathbf{x}_0)$  be the data distribution, and  $p_{\boldsymbol{\theta}}(\mathbf{x}) = p_{\boldsymbol{\theta}}(\mathbf{x}_0)$  be the marginal data distribution obtained from the reverse process, then the score-matching loss function in Eq. 3 is an upper bound of the negative-log-likelihood of real data  $\mathbf{x} \sim p(\mathbf{x})$  over  $\boldsymbol{\theta}$ . Formally,

$$-\mathbb{E}_{p(\mathbf{x})}[\log p_{\boldsymbol{\theta}}(\mathbf{x})] \leq \mathcal{L}_{\text{SM}}(\boldsymbol{\theta}) + \text{Const}, \quad (4)$$

where Const is a constant independent of  $\boldsymbol{\theta}$ .

Remark 2 indicates that the  $\boldsymbol{\theta}$  optimized by Eq. 3 approximates the maximum likelihood estimation of data distribution  $p(\mathbf{x})$ . Consequently,  $p_{\boldsymbol{\theta}}(\mathbf{x})$  approximates  $p(\mathbf{x})$  accurately with enough capacity. The algorithmic illustration of DIFFPUTER’s M-step is summarized in Algorithm 1.

#### 4.2 E-STEP: MISSING DATA IMPUTATION WITH A LEARNED DIFFUSION MODEL

Given the current estimation of data distribution  $p_{\boldsymbol{\theta}}(\mathbf{x})$ , the E-step aims to obtain the distribution of complete data, conditional on the observed values, i.e.,  $p_{\boldsymbol{\theta}}(\mathbf{x} | \mathbf{x}^{\text{obs}})$ , such that the estimated complete data  $\mathbf{x}^*$  can be updated by taking the expectation, i.e.,  $\mathbf{x}^* = \mathbb{E}_{\mathbf{x}} p_{\boldsymbol{\theta}}(\mathbf{x} | \mathbf{x}^{\text{obs}})$ .

When there is an explicit density function for  $p_{\boldsymbol{\theta}}(\mathbf{x})$ , or when the conditional distribution  $p_{\boldsymbol{\theta}}(\mathbf{x} | \mathbf{x}^{\text{obs}})$  is tractable (e.g., can be sampled), computing  $\mathbb{E}_{\mathbf{x}} p(\mathbf{x} | \mathbf{x}^{\text{obs}}, \boldsymbol{\theta})$  becomes feasible. While most deep generative models such as VAEs and GANs support convenient unconditional sampling from  $p_{\boldsymbol{\theta}}(\mathbf{x})$ , they do not naturally support conditional sampling, e.g.,  $\tilde{\mathbf{x}} \sim p_{\boldsymbol{\theta}}(\mathbf{x} | \mathbf{x}^{\text{obs}})$ . Luckily, since the diffusion model preserves the size and location of features in both the forward diffusion process and the reverse denoising process, it offers a convenient and accurate solution to perform conditional sampling  $p_{\boldsymbol{\theta}}(\mathbf{x} | \mathbf{x}^{\text{obs}})$  from an unconditional model  $p_{\boldsymbol{\theta}}(\mathbf{x}) = p_{\boldsymbol{\theta}}(\mathbf{x}^{\text{obs}}, \mathbf{x}^{\text{mis}})$ .

Specifically, let  $\mathbf{x}$  be the data to impute,  $\mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}}$  be the values of observed entries,  $\mathbf{m}$  be the 0/1 indicators of the location of missing entries, and  $\tilde{\mathbf{x}}_t$  be the imputed data at time  $t$ . Then, we can obtain the imputed data at time  $t - \Delta t$  via combining the observed entries from the forward process on  $\mathbf{x}_0 = \mathbf{x}$ , and the missing entries from the reverse process on  $\tilde{\mathbf{x}}_t$  from the prior step  $t$  (Lugmayr et al., 2022; Song et al., 2021b):

$$\mathbf{x}_{t-\Delta t}^{\text{forward}} = \mathbf{x} + \sigma(t - \Delta t) \cdot \boldsymbol{\varepsilon}, \text{ where } \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (\text{Forward for observed entries}) \quad (5)$$

$$\mathbf{x}_{t-\Delta t}^{\text{reverse}} = \tilde{\mathbf{x}}_t + \int_t^{t-\Delta t} d\tilde{\mathbf{x}}_t, \text{ where } d\tilde{\mathbf{x}}_t \text{ is defined in Eq. 2} \quad (\text{Reverse for missing entries}) \quad (6)$$

$$\tilde{\mathbf{x}}_{t-\Delta t} = (1 - \mathbf{m}) \odot \mathbf{x}_{t-\Delta t}^{\text{forward}} + \mathbf{m} \odot \mathbf{x}_{t-\Delta t}^{\text{reverse}}. \quad (\text{Merging}) \quad (7)$$

Based on the above process, starting at a random noise from the maximum time  $T$ , i.e.,  $\tilde{\mathbf{x}}_T \sim \mathcal{N}(\mathbf{0}, \sigma^2(T)\mathbf{I})$ , we can obtain a reconstructed  $\tilde{\mathbf{x}}_0$ , such that the observed entries of  $\tilde{\mathbf{x}}_0$  are the same as those of  $\mathbf{x}_0$ , i.e.,  $\tilde{\mathbf{x}}_0^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}}$ . In Theorem 1, we prove that via the algorithm above, the obtained  $\tilde{\mathbf{x}}_0$  is sampled from the desired conditional distribution, i.e.,  $\tilde{\mathbf{x}}_0 \sim p_{\boldsymbol{\theta}}(\mathbf{x} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$ .

**Theorem 1.** Let  $\tilde{\mathbf{x}}_T$  be a sample from the prior distribution  $\pi(\mathbf{x}) = \mathcal{N}(\mathbf{0}, \sigma^2(T)\mathbf{I})$ ,  $\mathbf{x}$  be the data to impute, and the known entries of  $\mathbf{x}$  are denoted by  $\mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}}$ . The score function  $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)$  is approximated by neural network  $\boldsymbol{\varepsilon}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)$ . Applying Eq. 5, Eq. 6, and Eq. 7 iteratively from  $t = T \gg 0$  until  $t = 0$  with  $\Delta t \rightarrow 0$ , then  $\tilde{\mathbf{x}}_0$  is a sample from  $p_{\boldsymbol{\theta}}(\mathbf{x})$ , under the condition that its observed entries  $\tilde{\mathbf{x}}_0^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}}$ . Formally,

$$\tilde{\mathbf{x}}_0 \sim p_{\boldsymbol{\theta}}(\mathbf{x} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}}) \quad (8)$$



See proof in Appendix B.1. Theorem 1 demonstrates that with a learned diffusion model  $\theta$ , we are able to obtain samples exactly from the conditional distribution  $p_\theta(\mathbf{x}|\mathbf{x}^{\text{obs}})$  through the aforementioned imputation process. For inference, we use Monte Carlo estimation to compute the expectation of the missing values, i.e.,  $\mathbf{x}^* = \mathbb{E}_\mathbf{x} p_\theta(\mathbf{x}|\mathbf{x}^{\text{obs}}) \approx \sum_{j=1}^N \tilde{\mathbf{x}}_0^{(j)} / N$ .

**Discretization.** The above imputing process involves recovering  $\tilde{\mathbf{x}}_t$  continuously from  $t = T$  to  $t = 0$ , which is infeasible in practice. In implementation, we discretize the process via  $M + 1$  discrete descending timesteps  $t_M, t_{M-1}, \dots, t_i, \dots, t_0$ , where  $t_M = T, t_0 = 0$ . Therefore, starting from  $\tilde{\mathbf{x}}_{t_M} \sim p_\theta(\mathbf{x}_{t_M} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$ , we obtain  $p_\theta(\mathbf{x}_{t_i} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$  from  $i = M - 1$  to  $i = 0$ .

**Algorithm 1:** M-step: Density Estimation using Diffusion Model

**Input:** Data samples from  $p(\mathbf{x})$ .

**Output:** Score network  $\epsilon_\theta$

```

1 while not converging do
2   Sample  $\mathbf{x} \sim p(\mathbf{x}) \in \mathbb{R}^d$ 
3   Sample  $t \sim p(t)$ 
4   Sample  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \in \mathbb{R}^d$ 
5    $\mathbf{x}_0 = \mathbf{x}$ 
6    $\mathbf{x}_t = \mathbf{x}_0 + \sigma(t) \cdot \epsilon$ 
7    $\ell(\theta) = \|\epsilon_\theta(\mathbf{x}_t, t) - \frac{-\epsilon}{\sigma(t)}\|_2^2$ 
8   Update  $\theta$  via Adam optimizer
```

**Algorithm 2:** E-step: Missing Data Imputation

**Input:** Score network  $\epsilon_\theta(\mathbf{x}_t, t)$ . Data with missing values  $\mathbf{x} \in \mathbb{R}^d$ , mask  $\mathbf{m} \in \{0, 1\}^d$ . Number of samples  $N$ . Number of sampling steps  $M$ .

**Output:** Imputed data  $\mathbf{x}^*$

```

1 for  $j \leftarrow 1$  to  $N$  do
2   Sample  $\tilde{\mathbf{x}}_{t_M}^{(j)} \sim \mathcal{N}(\mathbf{0}, \sigma^2(t_M)\mathbf{I})$ 
3   for  $i \leftarrow M$  to 1 do
4      $\mathbf{x}_{t_{i-1}}^{\text{forward},(j)} = \mathbf{x} + \sigma(t_{i-1}) \cdot \epsilon$ 
5      $\mathbf{x}_{t_{i-1}}^{\text{reverse},(j)} = \tilde{\mathbf{x}}_{t_i}^{(j)} + \int_{t_i}^{t_{i-1}} d\tilde{\mathbf{x}}_{t_i}^{(j)}$ 
6      $\tilde{\mathbf{x}}_{t_{i-1}}^{(j)} = \mathbf{m} \odot \mathbf{x}_{t_{i-1}}^{\text{forward},(j)} + (1 - \mathbf{m}) \odot \mathbf{x}_{t_{i-1}}^{\text{reverse},(j)}$ 
7  $\mathbf{x}^* = \sum_{j=1}^N \tilde{\mathbf{x}}_0^{(j)} / N = \sum_{j=1}^N \tilde{\mathbf{x}}_0^{(j)} / N$ 
```

Since the desired imputed data  $\mathbf{x}^{\text{mis}*}$  is the expectation, i.e.,  $\mathbf{x}^{\text{mis}*} = \mathbb{E}_{\mathbf{x}^{\text{mis}}} p(\mathbf{x}^{\text{mis}} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$ , we sample  $\tilde{\mathbf{x}}_0^{\text{mis}}$  for  $N$  times, and take the average value as the imputed  $\mathbf{x}^{\text{mis}*}$ . The algorithmic illustration of DIFFPUTER's E-step is summarized in Algorithm 2.

### 4.3 IMPLEMENTATIONS

**Data Processing.** Diffusion models are inherently suited for continuous data but not for discrete data. Given that real-world tabular data often contains both types of data, we use one-hot encoding to convert each dimension of discrete data into multi-dimensional 0/1 encoding, thereby treating it as continuous data<sup>1</sup>. After that, we perform standardization on each column of data, ensuring that it has a mean of 0 and a variance of 1. The final imputation performance is measured by the difference between the predicted value and the ground-truth missing value **after standardization**.

**Initialization of Missing Data.** The execution of the EM algorithm requires the initialized values of missing entries  $\mathbf{x}^{\text{mis}(0)}$ , which might have a huge impact on the model's convergence. For simplicity, we initialize the missing entries of each column to be the mean of the column's observed values, equivalent to setting  $\mathbf{x}^{\text{mis}(0)} = 0$  everywhere (since the data has been standardized).

**Training.** To obtain a more accurate estimation of complete data  $\mathbf{x}$ , DIFFPUTER iteratively executes the M-step and E-step. To be specific, let  $\mathbf{x}^{(k)} = (\mathbf{x}^{\text{obs}(k)}, \mathbf{x}^{\text{obs}(k)}) = (\mathbf{x}^{\text{obs}}, \mathbf{x}^{\text{obs}(k)})$  be the estimation of complete data at  $k$ -th iteration,  $\theta^{(k)}$  be the diffusion model's parameters at  $k$ -th iteration, we write  $\theta^{(k+1)}$  as a function of  $\mathbf{x}^{(k)}$ , i.e.,  $\theta^{(k+1)} = \text{M-step}(\mathbf{x}^{(k)})$ , and  $\mathbf{x}^{(k+1)}$  as a function of  $\theta^{(k+1)}$  and  $\mathbf{x}^{(k)}$ , i.e.,  $\mathbf{x}^{(k+1)} = \text{E-step}(\mathbf{x}^{(k)}, \theta^{(k+1)})$ . Therefore, with the initialized  $\mathbf{x}^{(0)}$  and the maximum iteration  $K$ , we are able to obtain  $\mathbf{x}^{(k)}$  from  $k = 1$  to  $K$ .

**Inference.** For in-sample imputation, the imputed values are obtained iteratively during the training process. For out-of-sample imputation, with the observed incomplete data  $\mathbf{x}$ , the mask  $\mathbf{m}$ , and the trained score network  $\epsilon_\theta$ , we directly apply the E-step once in Algorithm 2 to obtain the imputation values  $\mathbf{x}^*$ .

<sup>1</sup>We've tried encoding techniques such as analog-bit encoding, but observe little difference.

## 5 EXPERIMENTS

In this section, we conduct experiments to study the efficacy of the proposed DIFFPUTER in missing data imputation tasks. The code and documents for reproducing all the experimental results of the proposed DIFFPUTER and all the baselines are at <https://anonymous.4open.science/r/ICLR-DiffPuter>.

### 5.1 EXPERIMENTAL SETTINGS

**Datasets.** We evaluate the proposed DIFFPUTER on ten public real-world datasets of varying scales. We consider five datasets of only continuous features: California, Letter, Gesture, Magic, and Bean, and five datasets of both continuous and discrete features: Adult, Default, Shoppers, Beijing, and News. The detailed information of these datasets is presented in Appendix D.2. Following previous works (Muzellec et al., 2020; Zhao et al., 2023), we study three missing mechanisms: MCAR, MAR, and MNAR. Differences between the three settings are in Appendix D.3. In this section, we only report the performance in the MCAR setting, while the results of the other two settings are in Appendix E. In the main experiments, we set the missing rate as  $r = 30\%$ . For each dataset, we generate 10 masks according to the missing mechanism and report the mean and standard deviation of the imputing performance.

**Baselines.** We compare DIFFPUTER with 16 powerful imputation methods from different categories: 1) Distribution-matching methods based on optimal transport, including TDM (Zhao et al., 2023) and MOT (Muzellec et al., 2020). 2) Graph-based imputation methods, including GRAPE (You et al., 2020): a pure bipartite graph-based framework for data imputation, and IGRM (Zhong et al., 2023): a graph-based imputation method that iteratively reconstructs the friendship network. 3) Iterative methods, including EM with multivariate Gaussian priors (García-Laencina et al., 2010), MICE (Van Buuren & Karin, 2011), MIRACLE (Kyono et al., 2021), SoftImpute (Hastie et al., 2015), and MissForest (Stekhoven & Bühlmann, 2012). 4) Deep generative models, including MIWAE (Mattei & Frellsen, 2019), GAIN (Yoon et al., 2018), MCFlow (Richardson et al., 2020), MissDiff (Ouyang et al., 2023) and TabCSDI (Zheng & Charoenphakdee, 2022). It is worth noting that MissDiff and TabCSDI are also based on diffusion. We also compare with two recent SOTA imputing methods ReMasker (Du et al., 2024) and HyperImpute (Jarrett et al., 2022).

**Evaluation Protocols.** For each dataset, we use 70% as the training set, and the remaining 30% as the testing set. All methods are trained on the training set. The imputation is applied to both the missing values in the training set and the testing set. Consequently, the imputation of the training set is the 'in-sample' setting, while imputing the testing set is the 'out-of-sample' setting. The performance of DIFFPUTER is evaluated by the divergence between the predicted values and ground-truth values of missing entries. For continuous features, we use Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE), while for discrete features, we use Accuracy. Note that both the MAE and RMSE are calculated based on the input data after standardization (zero mean and unit variance). The implementation details and hyperparameter settings are presented in Appendix D.4.

### 5.2 MAIN RESULTS (MASK COMPLETELY AT RANDOM)

We first evaluate DIFFPUTER's performance in the in-sample imputation task. Figure 2 compares the performance of different imputation methods regarding continuous columns using MAE and RMSE. Table 1 compares the performance of different methods regarding discrete columns using accuracy. Our observations on these experimental results are summarized as follows.

**Superior performance of DIFFPUTER.** Across all datasets, DIFFPUTER provides high-quality imputation results, matching the best methods on some datasets and significantly outperforming the second-best methods on the remaining datasets. Compared with the other two diffusion-based methods, MissDiff and TabCSDI, the substantial performance advantage of DIFFPUTER convincingly demonstrates the correctness and superiority of combining the EM algorithm with Diffusion.

**Traditional machine learning methods are still powerful imputors.** A well-recognized claim in tabular data machine learning is that *traditional machine learning methods can sometimes be better than deep learning methods* (Lalande & Doya, 2022; Suh & Song, 2023; Jolicoeur-Martineau et al., 2024b), and we have similar observations in missing data imputation. For example, the simple EM




Figure 2: MCAR, In-sample imputation performance on **continuous columns**: Comparing DIFFPUTER with **17 baselines** on imputing continuous data on all the **ten datasets**. A blank column indicates that the method fails or gets out-of-memory for that dataset. DIFFPUTER outperforms the most competitive baseline method by 8.10% (MAE score) and 5.64% (RMSE score) by average. The circled number after the model name denotes its ranking among all methods.

algorithm under the mixture-Gaussian assumption and the KNN algorithm give decent imputation performance and outperform a lot of early deep generative imputation methods.

**The preference of different types of algorithms for heterogeneous tabular data.** Another interesting finding is that discriminative methods (e.g., Remasker, GRAPE, MOT) and generative methods (e.g., DIFFPUTER) have different preferences for continuous data and discrete data. Overall, generative models appear to be more effective at imputing continuous columns. As shown in Figure 2, our model significantly outperforms these discriminative models. Conversely, discriminative models are more proficient at imputing discrete data: in Table 1, the best discriminative model, Remasker, can achieve results comparable to our method on almost all datasets.

Next, we compare the performance in the out-of-sample imputation tasks. Noting that some imputation methods are specifically designed for in-sample imputation and cannot be applied to the out-of-sample setting, the number of baselines in this task is significantly reduced. Table 5 in Appendix E.1 compares the MAEs and RMSEs in the OOS imputation task. Comparing it with the results of in-sample imputation, we can easily observe that some methods exhibit significant performance differences between the two settings. For example, graph-based methods GRAPE and IGRM perform well in in-sample imputation, but their performance degrades significantly in out-of-sample imputation. IGRM even fails on all datasets in the out-of-sample imputation setting. In contrast, our DIFFPUTER demonstrates excellent performance in both in-sample and out-of-sample imputation. This highlights DIFFPUTER’s superior performance and robust generalization capabilities.

### 5.3 ABLATION STUDIES

In this section, we conduct ablation studies to evaluate the efficacy of the individual designs of DIFFPUTER, as well as its robustness w.r.t. different factors.

**Impacts of the number of EM iterations.** In Fig. 3, we present the performance of DIFFPUTER’s imputation results from increasing EM iterations. Note that  $k = 0$  represents the imputation result of a randomly initialized denoising network, and  $k = 1$  represents the performance of a pure diffusion model without iterative refinement. It is clearly observed that a single diffusion imputation achieves only suboptimal performance, while DIFFPUTER’s performance steadily improves as the number of



Table 1: MCAR, In-sample imputation performance on **discrete columns**. Comparison of the imputation accuracy on discrete columns on five datasets. DIFFPUTER ranks the first among **19** imputation methods. Mean/Median/MF means using the mean/median/most frequent value as the imputation, which will give the same imputation result in this case.

Method	Adult	Default	Shoppers	News	Beijing	Average	Rank
<i>Statistical</i>							
Mean/Median/MF	55.20 $\pm$ 0.02	53.37 $\pm$ 0.11	50.60 $\pm$ 0.29	18.60 $\pm$ 0.14	14.06 $\pm$ 0.11	38.37	(15)
<i>Traditional iterative</i>							
EM	61.27 $\pm$ 0.12	57.80 $\pm$ 0.22	50.94 $\pm$ 0.27	39.37 $\pm$ 0.06	17.75 $\pm$ 0.18	45.42	(10)
MICE	50.31 $\pm$ 0.21	51.88 $\pm$ 0.30	43.43 $\pm$ 0.19	30.05 $\pm$ 1.17	14.89 $\pm$ 0.90	38.11	(16)
MIRACLE	62.28 $\pm$ 0.21	55.79 $\pm$ 1.52	45.74 $\pm$ 0.69	39.33 $\pm$ 0.27	20.88 $\pm$ 0.58	44.61	(12)
SoftImpute	56.18 $\pm$ 0.08	56.01 $\pm$ 0.07	50.94 $\pm$ 0.27	18.70 $\pm$ 0.05	17.19 $\pm$ 0.07	39.80	(14)
MissForest	63.51 $\pm$ 0.31	56.97 $\pm$ 0.23	51.10 $\pm$ 0.61	37.55 $\pm$ 0.84	19.13 $\pm$ 0.21	45.65	(9)
<i>Dist. match</i>							
MOT	63.85 $\pm$ 0.16	71.95 $\pm$ 0.06	56.91 $\pm$ 0.11	38.62 $\pm$ 0.20	19.51 $\pm$ 0.08	50.16	(5)
TDM	62.64 $\pm$ 1.36	63.41 $\pm$ 7.16	55.00 $\pm$ 0.46	30.62 $\pm$ 1.22	19.38 $\pm$ 0.47	46.21	(8)
<i>GNN</i>							
GRAPE	69.84 $\pm$ 0.25	73.85 $\pm$ 0.01	57.32 $\pm$ 0.15	40.83 $\pm$ 0.44	19.92 $\pm$ 0.40	52.35	(4)
IGRM	69.21 $\pm$ 0.29	OOM	57.45 $\pm$ 0.16	OOM	OOM	-	-
<i>Generative models</i>							
MIWAE	51.33 $\pm$ 0.08	48.11 $\pm$ 0.02	49.23 $\pm$ 0.61	17.32 $\pm$ 0.01	13.31 $\pm$ 0.28	35.86	(18)
GAIN	53.73 $\pm$ 0.38	54.49 $\pm$ 1.80	47.25 $\pm$ 0.95	34.20 $\pm$ 0.06	14.71 $\pm$ 0.22	40.88	(13)
MCFlow	60.31 $\pm$ 0.18	68.36 $\pm$ 0.9	52.73 $\pm$ 0.29	30.08 $\pm$ 3.01	15.84 $\pm$ 0.19	45.46	(11)
TabCSDI	58.72 $\pm$ 0.25	50.57 $\pm$ 0.02	43.24 $\pm$ 0.05	17.49 $\pm$ 0.14	13.03 $\pm$ 0.08	36.61	(17)
MissDiff	57.39 $\pm$ 7.34	66.01 $\pm$ 1.88	48.10 $\pm$ 1.86	41.39 $\pm$ 0.59	19.02 $\pm$ 0.05	46.38	(7)
<i>Other</i>							
KNN	63.28 $\pm$ 0.52	71.64 $\pm$ 0.34	53.93 $\pm$ 0.28	38.59 $\pm$ 0.17	18.40 $\pm$ 0.39	49.17	(6)
Remasker	69.18 $\pm$ 0.08	76.88 $\pm$ 0.62	57.86 $\pm$ 0.09	44.33 $\pm$ 0.08	<b>27.32</b> $\pm$ 0.16	55.11	(2)
HyperImpute	64.99 $\pm$ 0.06	74.39 $\pm$ 2.07	<b>59.19</b> $\pm$ 0.04	40.29 $\pm$ 0.96	26.91 $\pm$ 0.03	53.15	(3)
DIFFPUTER	<b>70.12</b> $\pm$ 0.17	<b>77.64</b> $\pm$ 0.32	58.82 $\pm$ 0.09	<b>44.69</b> $\pm$ 0.13	27.17 $\pm$ 0.11	<b>55.69</b>	(1)

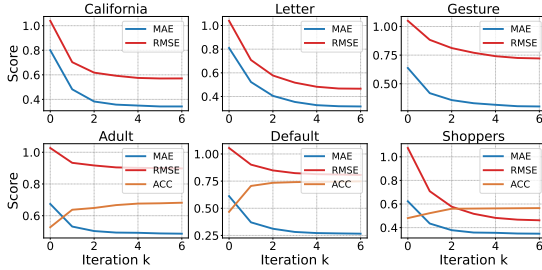


Figure 3: Impacts of the number of EM iterations. DIFFPUTER’s performance steadily improves as the number of EM interactions increases.




Figure 4: Impacts of the number of sampled imputations per iteration. A very small  $N$  leads to poor performance and large variance.

iterations increases. Additionally, we observe that DIFFPUTER does not require a large number of iterations to converge. In fact, 4 to 5 iterations are sufficient for DIFFPUTER to converge to a stable and satisfying state.

**Impacts of the number of sampled imputations per iteration.** To obtain the expected value of missing entries, we have to sample a number of examples from the conditional distribution (via the reverse diffusion process). Therefore, the number  $N$  should have a huge impact on DIFFPUTER’s performance. In Figure 4, we study the impacts of the number of samples on the imputation results. As expected, a very small  $N$  leads to poor performance and large variance, while increasing the sampling number consistently improves the performance and reduces the variance. An optimal and stable performance can be achieved at  $N \geq 10$ . Since the sampling time is linear w.r.t. to the sample number, we use  $N = 10$  by default.




Figure 5: Impacts of sampling steps on the training time and imputation performance (MAE). Reducing diffusion sampling steps greatly reduces the training time at the cost of a slight performance drop.




Figure 6: Impacts of extremely large missing ratios. The performance of DIFFPUTER is upper-bounded by the initialized missing values via mean imputation

Table 2: Comparison of the real training time on an Nvidia RTX 4090 GPU. DIFFPUTER has a similar training cost to SOTA methods, yet brings from 8% to 25% performance improvement.

Datasets	MOT	TDM	GRAPE	IGRM	HyperImpute	Remasker	DIFFPUTER
Time: <b>California</b>	446.47s	216.91s	635.7s	1267.5s	1277.3s	1320.1s	1927.2s
Time: <b>Adult</b>	396.68s	514.56s	2347.1s	3865.1s	1806.9s	1902.4s	2142.9s
DIFFPUTER’s perf. improve.	21.47%	21.37%	25.94%	20.97%	8.44%	10.65%	-

**Impacts of the number of sampling steps in diffusion.** Another hyperparameter impacting the performance and training speed is the number of sampling steps  $M$  in the reverse process of diffusion. By default, DIFFPUTER set  $M = 50$ , and we study the impact of reducing  $M$  in Figure 5. As demonstrated, reducing the diffusion sampling steps can greatly reduce the training time at the cost of only a slight performance drop. More detailed, reducing the sampling step from 50 to 20 reduces about 25% of training time cost, with only 3% performance drop. Therefore, when time costs are high, and precision requirements are not stringent, the training speed can be further increased by reducing the sampling steps.

**Impacts of large missing ratios.** In Figure 6, we study the performance of DIFFPUTER with missing ratio from 30% to 99% (observed ratio from 70% to 1%). Consistent with the intuition, when the observed ratio decreases (i.e., missing ratio increases), there is a significant increase in the imputation MAE score, indicating that the imputation performance becomes increasingly worse. However, we can observe a clear performance lower bound when the missing ratio approaches 99%, which is close to the performance of directly imputing using the mean of observed values. This is because, in DIFFPUTER’s first E-step, each column’s missing values are always initialized as the mean of the observed ones. Therefore, DIFFPUTER can still give a reasonable imputation result even if the missing ratio is extremely large.

**Comparison of Training Time.** Intuitively, the training of DIFFPUTER would be very time-consuming, as it requires iteratively alternating between the training and sampling processes of the diffusion model. However, in practice, DIFFPUTER’s training efficiency is significantly improved through the following design: 1) light-weighted MLP-architected denoising function; 2) reduced number of sampling steps ( $M = 50 \ll 1000$  in traditional diffusion methods). To demonstrate this, we compare the training time of DIFFPUTER with other SOTA imputation methods. As demonstrated in Table 2, DIFFPUTER has similar training cost (similar scale) to these representative SOTA methods, yet brings from 8% to 25% performance improvement, which is acceptable.

## 6 CONCLUSIONS

In this paper, we have proposed DIFFPUTER for missing data imputation. DIFFPUTER is an iterative method that combines the Expectation-Maximization algorithm and diffusion models, where the diffusion model serves as both the density estimator and missing data imputer. We demonstrate theoretically that the training and sampling process of a diffusion model precisely corresponds to the M-step and E-step of the EM algorithm. Therefore, we can iteratively update the density of the complete data and the values of the missing data. Extensive experiments have demonstrated the efficacy of the proposed method.

## REPRODUCIBILITY STATEMENT

We describe the algorithm of DIFFPUTER in Algorithm 1 and Algorithm 2. The proofs are in Appendix B. The detailed implementations are provided in Appendix D. The codes for implementing baselines and the proposed DIFFPUTER are provided in the anonymous github repo <https://anonymous.4open.science/r/ICLR-DiffPuter>.

## REFERENCES

- John Barnard and Xiao-Li Meng. Applications of multiple imputation in medical studies: from aids to nhanes. *Statistical methods in medical research*, 8(1):17–36, 1999.
- Ciprian Corneanu, Raghudeep Gadde, and Aleix M Martinez. Latentpaint: Image inpainting in latent space with diffusion models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 4334–4343, 2024.
- Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society: series B (methodological)*, 39(1): 1–22, 1977.
- Tianyu Du, Luca Melis, and Ting Wang. Remasker: Imputing tabular data with masked autoencoding. In *International Conference on Learning Representations*, 2024.
- Fabian Eckert, Teresa C Fort, Peter K Schott, and Natalie J Yang. Imputing missing values in the us census bureau’s county business patterns. Technical report, National Bureau of Economic Research, 2020.
- Pedro J García-Laencina, José-Luis Sancho-Gómez, and Aníbal R Figueiras-Vidal. Pattern classification with missing data: a review. *Neural Computing and Applications*, 19:263–282, 2010.
- Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems*, pp. 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.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pp. 6840–6851, 2020.
- Daniel Jarrett, Bogdan C Cebere, Tennison Liu, Alicia Curth, and Mihaela van der Schaar. Hyperimpute: Generalized iterative imputation with automatic model selection. In *International Conference on Machine Learning*, pp. 9916–9937. PMLR, 2022.
- Alexia Jolicoeur-Martineau, Kilian Fatras, and Tal Kachman. Generating and imputing tabular data via diffusion and flow-based gradient-boosted trees. In *International Conference on Artificial Intelligence and Statistics*, pp. 1288–1296. PMLR, 2024a.
- Alexia Jolicoeur-Martineau, Kilian Fatras, and Tal Kachman. Generating and imputing tabular data via diffusion and flow-based gradient-boosted trees. In *International Conference on Artificial Intelligence and Statistics*, pp. 1288–1296. PMLR, 2024b.
- Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, pp. 26565–26577, 2022.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Akim Kotelnikov, Dmitry Baranchuk, Ivan Rubachev, and Artem Babenko. Tabddpm: Modelling tabular data with diffusion models. In *International Conference on Machine Learning*, pp. 17564–17579. PMLR, 2023.
- Trent Kyono, Yao Zhang, Alexis Bellot, and Mihaela van der Schaar. Miracle: Causally-aware imputation via learning missing data mechanisms. *Advances in Neural Information Processing Systems*, 34:23806–23817, 2021.

- Florian Lalande and Kenji Doya. Numerical data imputation: Choose knn over deep learning. In *International Conference on Similarity Search and Applications*, pp. 3–10. Springer, 2022.
- Lee Lillard, James P Smith, and Finis Welch. What do we really know about wages? the importance of nonreporting and census imputation. *Journal of Political Economy*, 94(3, Part 1):489–506, 1986.
- Tennison Liu, Zhaozhi Qian, Jeroen Berrevoets, and Mihaela van der Schaar. Goggle: Generative modelling for tabular data by learning relational structure. In *International Conference on Learning Representations*, 2023.
- Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11461–11471, 2022.
- Zhihua Ma and Guanghui Chen. Bayesian methods for dealing with missing data problems. *Journal of the Korean Statistical Society*, 47:297–313, 2018.
- Pierre-Alexandre Mattei and Jes Frellsen. Miwae: Deep generative modelling and imputation of incomplete data sets. In *International conference on machine learning*, pp. 4413–4423. PMLR, 2019.
- Boris Muzellec, Julie Josse, Claire Boyer, and Marco Cuturi. Missing data imputation using optimal transport. In *International Conference on Machine Learning*, pp. 7130–7140. PMLR, 2020.
- Alfredo Nazabal, Pablo M Olmos, Zoubin Ghahramani, and Isabel Valera. Handling incomplete heterogeneous data using vaes. *Pattern Recognition*, 107:107501, 2020.
- Yidong Ouyang, Liyan Xie, Chongxuan Li, and Guang Cheng. Missdiff: Training diffusion models on tabular data with missing values. *arXiv preprint arXiv:2307.00467*, 2023.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: an imperative style, high-performance deep learning library. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pp. 8026–8037, 2019.
- Utomo Pujiyanto, Aji Prasetya Wibawa, Muhammad Iqbal Akbar, et al. K-nearest neighbor (k-nn) based missing data imputation. In *2019 5th International Conference on Science in Information Technology (ICSITech)*, pp. 83–88. IEEE, 2019.
- Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pp. 1530–1538. PMLR, 2015.
- Trevor W Richardson, Wencheng Wu, Lei Lin, Beilei Xu, and Edgar A Bernal. Mcflow: Monte carlo flow models for data imputation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14205–14214, 2020.
- Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. In *Advances in Neural Information Processing Systems*, 2021a.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *The Ninth International Conference on Learning Representations*, 2021b.
- Daniel J Stekhoven and Peter Bühlmann. Missforest—non-parametric missing value imputation for mixed-type data. *Bioinformatics*, 28(1):112–118, 2012.
- Heajung Suh and Jongwoo Song. A comparison of imputation methods using machine learning models. *Communications for Statistical Applications and Methods*, 30(3):331–341, 2023.
- Stef Van Buuren and Karin. mice: Multivariate imputation by chained equations in r. *Journal of statistical software*, 45:1–67, 2011.
- Lei Xu, Maria Skoularidou, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. Modeling tabular data using conditional gan. In *Advances in Neural Information Processing Systems*, pp. 7335–7345, 2019.



- Jinsung Yoon, James Jordon, and Mihaela Schaar. Gain: Missing data imputation using generative adversarial nets. In *International Conference on Machine Learning*, pp. 5689–5698. PMLR, 2018.
- Jiaxuan You, Xiaobai Ma, Yi Ding, Mykel J Kochenderfer, and Jure Leskovec. Handling missing data with graph representation learning. *Advances in Neural Information Processing Systems*, 33: 19075–19087, 2020.
- Hengrui Zhang, Jiani Zhang, Balasubramaniam Srinivasan, Zhengyuan Shen, Xiao Qin, Christos Faloutsos, Huzefa Rangwala, and George Karypis. Mixed-type tabular data synthesis with score-based diffusion in latent space. In *The twelfth International Conference on Learning Representations*, 2024.
- He Zhao, Ke Sun, Amir Dezfouli, and Edwin V Bonilla. Transformed distribution matching for missing value imputation. In *International Conference on Machine Learning*, pp. 42159–42186. PMLR, 2023.
- Shuhan Zheng and Nontawat Charoenphakdee. Diffusion models for missing value imputation in tabular data. In *NeurIPS 2022 First Table Representation Workshop*, 2022.
- Jiajun Zhong, Ning Gui, and Weiwei Ye. Data imputation with iterative graph reconstruction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 11399–11407, 2023.

## A DEFINITION OF SYMBOLS

In Table 3, we explain the definition of every symbol used in this paper.

Table 3: Explanations of the symbols used in this paper.

Symbol	Explanation
$d$	the number of columns of tabular data
$\mathbf{x} \in \mathbb{R}^d$	random variable, representing a row of the tabular data
$p(\mathbf{x}) = p_{\text{data}}(\mathbf{x})$	the distribution of the complete data
$\mathbf{m} \in \{0, 1\}^d$	binary mask vector
$\mathbf{x}^{\text{obs}}$	the observed part of $\mathbf{x}$ , indicated by $m_k = 0$
$\mathbf{x}^{\text{mis}}$	the missing part of $\mathbf{x}$ , indicated by $m_k = 1$
$t$	the timestep of diffusion process
$\sigma(t) = t$	the noise level at timestep $t$
$T$	the maximum timestep
$\varepsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$	standard Gaussian noise
$\mathbf{x}_0 = \mathbf{x}$	unforwarded $\mathbf{x}$ (at timestep 0)
$\mathbf{x}_0^{\text{obs}} = \mathbf{x}^{\text{obs}}$	observed part of $\mathbf{x}$
$\mathbf{x}_0^{\text{mis}} = \mathbf{x}^{\text{mis}}$	missing part of $\mathbf{x}$
$\mathbf{x}_t = \mathbf{x}_0 + \sigma(t)\varepsilon$	forwarded $\mathbf{x}$ at timestep $t$
$\mathbf{x}_t^{\text{obs}} = \mathbf{x}_0^{\text{obs}} + \sigma(t)\varepsilon$	observed part of $\mathbf{x}_t$
$\mathbf{x}_t^{\text{mis}} = \mathbf{x}_0^{\text{mis}} + \sigma(t)\varepsilon$	missing part of $\mathbf{x}_t$
$\theta$	diffusion model’s parameters
$p_\theta(\mathbf{x})$	density of data induced by the diffusion model
$\tilde{\mathbf{x}}_T \sim \pi(\mathbf{x}) = \mathcal{N}(\mathbf{0}, \sigma^2(T)\mathbf{I})$	sampled random noise for the reverse process
$\tilde{\mathbf{x}}_T^{\text{obs}}$	observed part of $\tilde{\mathbf{x}}_T$
$\tilde{\mathbf{x}}_T^{\text{mis}}$	missing part of $\tilde{\mathbf{x}}_T$
$\tilde{\mathbf{x}}_t$	sampled data at timestep $t$ in the reverse process
$\tilde{\mathbf{x}}_t^{\text{obs}}$	observed part of $\tilde{\mathbf{x}}_t$
$\tilde{\mathbf{x}}_t^{\text{mis}}$	missing part of $\tilde{\mathbf{x}}_t$
$\hat{\mathbf{x}}_0^{\text{obs}}$	The exact ground-truth value(s) of the observed part

## B PROOFS

### B.1 PROOF FOR THEOREM 1

*Proof.* First, it is obvious that the observed entries from Eq. 7, i.e.,  $\tilde{\mathbf{x}}_t^{\text{obs}}$ , satisfy

$$\tilde{\mathbf{x}}_t^{\text{obs}} \sim p(\mathbf{x}_t^{\text{obs}} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}}), \forall t. \quad (9)$$

Then, we introduce the following Lemma.

**Lemma 1.** Let  $\tilde{\mathbf{x}}_t \sim p_\theta(\mathbf{x}_t | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$ . If  $\tilde{\mathbf{x}}_{t-\Delta t}$  is a sample obtained from Eq. 5, Eq. 6, and Eq. 7, then  $\tilde{\mathbf{x}}_{t-\Delta t} \sim p_\theta(\mathbf{x}_{t-\Delta t} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$ , when  $\Delta t \rightarrow 0^+$ .

*Proof.* First, note that  $\mathbf{x}_t$  ( $\mathbf{x}_{t-\Delta t}$ ) is obtained via adding **dimensionally-independent** Gaussian noises to  $\mathbf{x}_0 = \mathbf{x}$  (see the forward process in Eq. 1), we have

$$\begin{aligned} p_\theta(\mathbf{x}_{t-\Delta t} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}}) &= p_\theta(\mathbf{x}_{t-\Delta t}^{\text{obs}}, \mathbf{x}_{t-\Delta t}^{\text{mis}} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}}) \\ &= \underbrace{p(\mathbf{x}_{t-\Delta t}^{\text{obs}} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})}_{\text{Correspond to } \tilde{\mathbf{x}}_{t-\Delta t}^{\text{obs}} \text{ in Eq. 5}} \cdot p_\theta(\mathbf{x}_{t-\Delta t}^{\text{mis}} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}}). \end{aligned} \quad (10)$$

Therefore,  $\tilde{\mathbf{x}}_{t-\Delta t}$ 's observed entries  $\tilde{\mathbf{x}}_{t-\Delta t}^{\text{obs}}$  is sampled from distribution  $p(\mathbf{x}_{t-\Delta t}^{\text{obs}} | \mathbf{x}_0^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$ . Then, we turn to the missing entries, where we have the following derivation:

$$\begin{aligned} p_{\theta}(\mathbf{x}_{t-\Delta t}^{\text{mis}} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}}) &= \int p_{\theta}(\mathbf{x}_{t-\Delta t}^{\text{mis}} | \mathbf{x}_t, \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}}) p_{\theta}(\mathbf{x}_t | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}}) d\mathbf{x}_t \\ &= \mathbb{E}_{p_{\theta}(\mathbf{x}_t | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})} p_{\theta}(\mathbf{x}_{t-\Delta t}^{\text{mis}} | \mathbf{x}_t, \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}}) \\ &\approx \mathbb{E}_{p_{\theta}(\mathbf{x}_t | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})} p_{\theta}(\mathbf{x}_{t-\Delta t}^{\text{mis}} | \mathbf{x}_t) \\ &= p_{\theta}(\mathbf{x}_{t-\Delta t}^{\text{mis}} | \tilde{\mathbf{x}}_t), \end{aligned} \quad (11)$$

where  $\tilde{\mathbf{x}}_t$  is a random sample from  $p_{\theta}(\mathbf{x}_t | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$ . The first ' $\approx$ ' holds because when  $\Delta t \rightarrow 0$ ,  $\mathbf{x}_{t-\Delta t}^{\text{mis}}$  is almost predictable via  $\mathbf{x}_t$  without  $\mathbf{x}_0^{\text{obs}}$ . Therefore, when  $\tilde{\mathbf{x}}_t \sim p_{\theta}(\mathbf{x}_t | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$ ,  $p_{\theta}(\mathbf{x}_{t-\Delta t}^{\text{mis}} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$  is approximately tractable via  $p_{\theta}(\mathbf{x}_{t-\Delta t}^{\text{mis}} | \tilde{\mathbf{x}}_t)$ .

Note that  $\mathbf{x}_{t-\Delta t}^{\text{mis}}$  denotes the missing entries of  $\mathbf{x}_{t-\Delta t}$ . According to the reverse denoising process in Eq. 2, given  $\mathbf{x}_t = \tilde{\mathbf{x}}_t$ ,  $\mathbf{x}_{t-\Delta t}$  is obtained via integrating  $d\mathbf{x}_t$  from  $t$  to  $t - \Delta t$ , i.e.,

$$\mathbf{x}_{t-\Delta t} = \tilde{\mathbf{x}}_t + \int_t^{t-\Delta t} d\mathbf{x}_t. \quad (12)$$

Therefore,  $\tilde{\mathbf{x}}_{t-\Delta t}^{\text{mis}} = \mathbf{x}_{t-\Delta t}^{\text{mis}}$  obtained from Eq. 6 is a sample from  $p_{\theta}(\mathbf{x}_{t-\Delta t}^{\text{mis}} | \tilde{\mathbf{x}}_t)$ . And, approximately,  $\tilde{\mathbf{x}}_{t-\Delta t}^{\text{mis}} \sim p_{\theta}(\mathbf{x}_{t-\Delta t}^{\text{mis}} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$ .

Since  $\tilde{\mathbf{x}}_{t-\Delta t}^{\text{obs}} \sim p(\mathbf{x}_{t-\Delta t}^{\text{obs}} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$ ,  $\tilde{\mathbf{x}}_{t-\Delta t}^{\text{mis}} \sim p_{\theta}(\mathbf{x}_{t-\Delta t}^{\text{mis}} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$ , we have  $\tilde{\mathbf{x}}_{t-\Delta t} \sim p_{\theta}(\mathbf{x}_{t-\Delta t} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$ , according to Eq. 10. Therefore, the proof for Lemma 1 is completed.  $\square$

With Lemma 1, we are able to prove Theorem 1 via induction, as long as  $\tilde{\mathbf{x}}_T$  is also sampled from  $p(\mathbf{x}_T | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$ . This holds because  $p(\mathbf{x}_T | \mathbf{x}_0) \approx p(\mathbf{x}_T)$ , given the condition that  $\mathbf{x}_0 = \mathbf{x}$  has zero mean and unit variance, and  $\sigma(T) \gg 1$ .

Note that the score function  $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)$  in Eq. 2 is intractable, it is replaced with the output of the score neural network  $\epsilon_{\theta}(\mathbf{x}_t, t)$ . Therefore, the finally obtained distribution can be rewritten as  $p_{\theta}(\mathbf{x} | \mathbf{x}^{\text{obs}} = \hat{\mathbf{x}}_0^{\text{obs}})$ . Therefore, the proof for Theorem 1 is complete.  $\square$

## C FURTHER EXPLANATION OF THE DIFFUSION MODEL

The diffusion model we adopt in Section 4.1 is actually a simplified version of the Variance-Exploding SDE proposed in (Song et al., 2021b).

Note that (Song et al., 2021b) has provided a unified formulation via the Stochastic Differential Equation (SDE) and defines the forward process of Diffusion as

$$d\mathbf{x} = \mathbf{f}(\mathbf{x}, t)dt + g(t) d\mathbf{w}_t, \quad (13)$$

where  $\mathbf{f}(\cdot)$  and  $g(\cdot)$  are the drift and diffusion coefficients and are selected differently for different diffusion processes, e.g., the variance preserving (VP) and variance exploding (VE) formulations.  $\omega_t$  is the standard Wiener process. Usually,  $\mathbf{f}(\cdot)$  is of the form  $\mathbf{f}(\mathbf{x}, t) = f(t) \mathbf{x}$ . Thus, the SDE can be equivalently written as

$$d\mathbf{x} = f(t) \mathbf{x} dt + g(t) d\mathbf{w}_t. \quad (14)$$

Let  $\mathbf{x}$  be a function of the time  $t$ , i.e.,  $\mathbf{x}_t = \mathbf{x}(t)$ , then the conditional distribution of  $\mathbf{x}_t$  given  $\mathbf{x}_0$  (named as the perturbation kernel of the SDE) could be formulated as:

$$p(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; s(t)\mathbf{x}_0, s^2(t)\sigma^2(t)\mathbf{I}), \quad (15)$$

where

$$s(t) = \exp\left(\int_0^t f(\xi) d\xi\right), \text{ and } \sigma(t) = \sqrt{\int_0^t \frac{g^2(\xi)}{s^2(\xi)} d\xi}. \quad (16)$$

Therefore, the forward diffusion process could be equivalently formulated by defining the perturbation kernels (via defining appropriate  $s(t)$  and  $\sigma(t)$ ).

Variance Exploding (VE) implements the perturbation kernel Eq. 15 by setting  $s(t) = 1$ , indicating that the noise is directly added to the data rather than weighted mixing. Therefore, The noise variance (the noise level) is totally decided by  $\sigma(t)$ . The diffusion model used in DIFFPUTER belongs to VE-SDE, but we use linear noise level (i.e.,  $\sigma(t) = t$ ) rather than  $\sigma(t) = \sqrt{t}$  in the vanilla VE-SDE (Song et al., 2021b). When  $s(t) = 1$ , the perturbation kernels become:

$$p(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \mathbf{0}, \sigma^2(t)\mathbf{I}) \Rightarrow \mathbf{x}_t = \mathbf{x}_0 + \sigma(t)\boldsymbol{\varepsilon}, \quad (17)$$

which aligns with the forward diffusion process in Eq. 1.

The sampling process of diffusion SDE is given by:

$$d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t) - g^2(t)\nabla_{\mathbf{x}} \log p_t(\mathbf{x})]dt + g(t)d\mathbf{w}_t. \quad (18)$$

For VE-SDE,  $s(t) = 1 \Leftrightarrow \mathbf{f}(\mathbf{x}, t) = \mathbf{f}(t) \cdot \mathbf{x} = \mathbf{0}$ , and

$$\begin{aligned} \sigma(t) &= \sqrt{\int_0^t g^2(\xi)d\xi} \Rightarrow \int_0^t g^2(\xi)d\xi = \sigma^2(t), \\ g^2(t) &= \frac{d\sigma^2(t)}{dt} = 2\sigma(t)\dot{\sigma}(t), \\ g(t) &= \sqrt{2\sigma(t)\dot{\sigma}(t)}. \end{aligned} \quad (19)$$

Plugging  $g(t)$  into Eq. 18, the reverse process in Eq. 2 is recovered:

$$d\mathbf{x}_t = -2\sigma(t)\dot{\sigma}(t)\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)dt + \sqrt{2\sigma(t)\dot{\sigma}(t)}d\mathbf{w}_t. \quad (20)$$

## D EXPERIMENTAL DETAILS

### D.1 CONFIGURATIONS.

We conduct all experiments with:

- Operating System: Ubuntu 22.04.3 LTS
- CPU: Intel 13th Gen Intel(R) Core(TM) i9-13900K
- GPU: NVIDIA GeForce RTX 4090 with 24 GB of Memory
- Software: CUDA 12.2, Python 3.9.16, PyTorch (Paszke et al., 2019) 1.12.1

### D.2 DATASETS

We use ten real-world datasets of varying scales, and all of them are available at Kaggle<sup>2</sup> or the UCI Machine Learning repository<sup>3</sup>. We consider five datasets of only continuous features: California<sup>4</sup>, Letter<sup>5</sup>, Gestur<sup>6</sup>, Magic<sup>7</sup>, and Bean<sup>8</sup>, and five datasets of both continuous and discrete features: Adult<sup>9</sup>, Default<sup>10</sup>, Shoppers<sup>11</sup>, Beijing<sup>12</sup>, and News<sup>13</sup>. The statistics of these datasets are presented in Table 4.

<sup>2</sup><https://www.kaggle.com/>

<sup>3</sup><https://archive.ics.uci.edu/>

<sup>4</sup><https://www.kaggle.com/datasets/camnugent/california-housing-prices>

<sup>5</sup><https://archive.ics.uci.edu/dataset/59/letter+recognition>

<sup>6</sup><https://archive.ics.uci.edu/dataset/302/gesture+phase+segmentation>

<sup>7</sup><https://archive.ics.uci.edu/dataset/159/magic+gamma+telescope>

<sup>8</sup><https://archive.ics.uci.edu/dataset/602/dry+bean+dataset>

<sup>9</sup><https://archive.ics.uci.edu/dataset/2/adult>

<sup>10</sup><https://archive.ics.uci.edu/dataset/350/default+of+credit+card+clients>

<sup>11</sup><https://archive.ics.uci.edu/dataset/468/online+shoppers+purchasing+intention+dataset>

<sup>12</sup><https://archive.ics.uci.edu/dataset/381/beijing+pm2+5+data>

<sup>13</sup><https://archive.ics.uci.edu/dataset/332/online+news+popularity>

Table 4: Statistics of datasets. # Num stands for the number of numerical columns, and # Cat stands for the number of categorical columns.

Dataset	# Rows	# Num	# Cat	# Train (In-sample)	# Test (Out-of-Sample)
<b>California</b> Housing	20,640	9	-	14,303	6,337
<b>Letter</b> Recognition	20,000	16	-	14,000	6,000
<b>Gesture</b> Phase Segmentation	9,522	32	-	6,665	2,857
<b>Magic</b> Gamma Telescope	19,020	10	-	13,314	5,706
<b>Dry Bean</b>	13,610	17	-	9,527	4,083
<b>Adult</b> Income	32,561	6	8	22,792	9,769
<b>Default</b> of Credit Card Clients	30,000	14	10	21,000	9,000
Online <b>Shoppers</b> Purchase	12,330	10	7	8,631	3,699
<b>Beijing</b> PM2.5 data	43,824	6	5	29,229	12,528
Online <b>News</b> Popularity	39,644	45	2	27,790	11,894

### D.3 MISSING MECHANISMS

According to how the masks  $\mathbf{m}$  are generated, there are three mainstream mechanisms of missingness, namely missing patterns: 1) Missing completely at random (MCAR) refers to the case that the probability of an entry being missing is independent of the data, i.e.,  $p(\mathbf{m}|\mathbf{x}) = p(\mathbf{m})$ . 2) In missing at random (MAR), the probability of missingness depends only on the observed values, i.e.,  $p(\mathbf{m}|\mathbf{x}) = p(\mathbf{m}|\mathbf{x}^{\text{obs}})$ . 3) All other cases are classified as missing not at random (MNAR), where the probability of missingness might also depend on other missing entries.

We follow the methods proposed in Zhao et al. (2023) to implement MAR and MNAR.

- "For MAR, we first sample a subset of features (columns in  $X$ ) that will not contain missing values, and then we use a logistic model with these non-missing columns as input to determine the missing values of the remaining columns, and we employ line search of the bias term to get the desired proportion of missing values."
- For MNAR, we use the first approach proposed in [4], "Using a logistic model with the input masked by MCAR". Specifically, similar to the MAR setting, we first divide the data columns into two groups, with one group serving as input for the logistic model, outputting the missing probability for the other set of columns. The difference is that after determining the missing probability of the second set, we apply MCAR to the input columns (the first set). Hence, missing values from the second set will depend on the masked values of the first set.

The code for generating masks according to the three missing mechanisms is also provided at <https://anonymous.4open.science/r/ICLR-DiffPuter>.

### D.4 IMPLEMENTATIONS AND HYPERPARAMETERS

We use a fixed set of hyperparameters, which will save significant efforts in hyperparameter-tuning when applying DIFFPUTER to more datasets. For the diffusion model, we set the maximum time  $T = 80$ , the noise level  $\sigma(t) = t$ , which is linear to  $t$ . The score/denoising neural network  $\epsilon(\mathbf{x}_t, t)$  is implemented as a 5-layer MLP with hidden dimension 1024.  $t$  is transformed to sinusoidal timestep embeddings and then added to  $\mathbf{x}_t$ , which is subsequently passed to the denoising function. When using the learned diffusion model for imputation, we set the number of discrete steps  $M = 50$  and the number of sampling times per data sample  $N = 10$ . DIFFPUTER is implemented with Pytorch, and optimized using Adam (Kingma & Ba, 2015) optimizer with a learning rate of  $1 \times 10^{-4}$ .

**Architecture of denoising neural network.** We use the same architecture of denoising neural network as in two recent tabular diffusion models for tabular data synthesis (Kotelnikov et al., 2023; Zhang et al., 2024). The denoising MLP takes the current time step  $t$  and the feature vector  $\mathbf{x}_t \in \mathbb{R}^{1 \times d}$  as input. First,  $\mathbf{x}_t$  is fed into a linear projection layer that converts the vector dimension to be  $d_{\text{hidden}} = 1024$ :

$$\mathbf{h}_0 = \text{FC}_{\text{in}}(\mathbf{x}_t) \in \mathbb{R}^{1 \times d_{\text{hidden}}}, \quad (21)$$



where  $\mathbf{h}_0$  is the transformed vector, and  $d_{\text{hidden}}$  is the output dimension of the input layer.

Then, following the practice in TabDDPM (Kotelnikov et al., 2023), the sinusoidal timestep embeddings  $\mathbf{t}_{\text{emb}} \in \mathbb{R}^{1 \times d_{\text{hidden}}}$  is added to  $\mathbf{h}_0$  to obtain the input vector  $\mathbf{h}_{\text{hidden}}$ :

$$\mathbf{h}_{\text{in}} = \mathbf{h}_0 + \mathbf{t}_{\text{emb}}. \quad (22)$$

The hidden layers are three fully connected layers of the size  $d_{\text{hidden}} - 2 * d_{\text{hidden}} - 2 * d_{\text{hidden}} - d_{\text{hidden}}$ , with SiLU activation functions:

$$\begin{aligned} \mathbf{h}_1 &= \text{SiLU}(\text{FC}_1(\mathbf{h}_0) \in \mathbb{R}^{1 \times 2 * d_{\text{hidden}}}), \\ \mathbf{h}_2 &= \text{SiLU}(\text{FC}_2(\mathbf{h}_1) \in \mathbb{R}^{1 \times 2 * d_{\text{hidden}}}), \\ \mathbf{h}_3 &= \text{SiLU}(\text{FC}_3(\mathbf{h}_2) \in \mathbb{R}^{1 \times d_{\text{hidden}}}). \end{aligned} \quad (23)$$

The estimated score is obtained via the last linear layer:

$$\epsilon_{\theta}(\mathbf{x}_t, t) = \mathbf{h}_{\text{out}} = \text{FC}_{\text{out}}(\mathbf{h}_3) \in \mathbb{R}^{1 \times d}. \quad (24)$$

Finally,  $\epsilon_{\theta}(\mathbf{x}_t, t)$  is applied to Eq. 3 for model training.

## D.5 IMPLEMENTATIONS OF BASELINES

We implement most of the baseline methods according to the publicly available codebases:

- Remasker (Du et al., 2024): <https://github.com/tydusky/remasker>.
- TDM (Zhao et al., 2023): <https://github.com/hezgit/TDM>
- MOT (Muzellec et al., 2020): <https://github.com/BorisMuzellec/MissingDataOT>
- GRAPE (You et al., 2020): <https://github.com/maxiaoba/GRAPE>
- IGRM (Zhong et al., 2023): <https://github.com/G-AILab/IGRM>
- TabCSDI (Zheng & Charoenphakdee, 2022): <https://github.com/pfnet-research/TabCSDI>
- MCFlow (Richardson et al., 2020): <https://github.com/trevor-richardson/MCFlow>.
- For HyperImputer (Jarrett et al., 2022), MissForest (Stekhoven & Bühlmann, 2012), MICE (Van Buuren & Karin, 2011), SoftImpute (Hastie et al., 2015), EM (García-Laencina et al., 2010), GAIN (Yoon et al., 2018), MIRACLE (Kyono et al., 2021), and MIWAE (Mattei & Frellsen, 2019), we use the implementations at: <https://github.com/vanderschaarlab/hyperimpute>.

MissDiff (Ouyang et al., 2023) does not provide its official implementations. Therefore, we obtain its results based on our own implementation.

The codes for all the methods are available at <https://anonymous.4open.science/r/ICLR-DiffPutter>

## D.6 HYPERPARAMETER SETTINGS OF BASELINES

Most of the deep learning baselines recommend the use of one set of hyperparameters for all datasets. For these methods, we directly follow their guidelines and use the default hyperparameters:

- ReMasker (Du et al., 2024): we use the recommended hyperparameters provided in Appendix A.2 in the original paper (Du et al., 2024). This set of hyperparameters is searched by tuning on the Letter dataset and is deployed for all the datasets in the original paper; hence, we follow this setting.
- HyperImpute (Jarrett et al., 2022): since HyperImpute works by searching over the space of classifiers/regressors and their hyperparameters, it does not have hyperparameters itself except parameters related to the AutoML search budget. We adopt the default budget

parameters of HyperImpute’s official implementation for all datasets. The default budget parameters and AutoML search space are provided<sup>14</sup> and Table 5 in the original paper (Jarrett et al., 2022).

- **MOT** (Muzellec et al., 2020) and **TDM** (Zhao et al., 2023): There is a main hyperparameter representing the number of subset pairs sampled from the dataset for computing optimal transport loss. Sinkhorn algorithm and TDM are controlled by hyperparameter `n_iter`. While the default value is 3000, we set it as 12000 for all datasets to ensure the algorithm converges sufficiently. For the round-robin version of the algorithm, the number of sampled pairs is controlled by `max_iter` and `rr_iter`; we adopt the default value 15, which is enough for the algorithm to converge. For the remaining hyperparameters related to network architectures, we use the default ones for all datasets.
- **kNN**: we follow the common practice of selecting the number of nearest neighbors as  $\sqrt{n}$ , where  $n$  is the number of samples in the dataset.
- **GRAPE** (You et al., 2020) and **IGRM** (Zhong et al., 2023): we adopt the recommended set of hyperparameters used in the original paper for all datasets. For a detailed explanation of the meaning of the parameters, please see the github repos: GRAPE<sup>15</sup> and IGRM<sup>16</sup>.
- **MissDiff**: since the original implementation is not available and it is based on the diffusion model, for a fair comparison, we simply use the same set of hyperparameters with our DiffPutter.
- **TabCSDI** (Zheng & Charoenphakdee, 2022): we follow the guide for selecting hyperparameters in the original paper (Appendix B in [3]). Specifically, we use a large version of the TabCSDI model with a number of layers set to 4 (see more detailed hyperparameters about the large TabCSDI model<sup>17</sup>. For batch size, we take the official choice of batch size (8) for the breast dataset (700 samples) as a base and scale the batch size accordingly with the sample size of our datasets: since most of the datasets we used have the number of samples between 20000 to 40000, we scale the batch size to 256 and use it for all datasets.
- **MCFlow** (Richardson et al., 2020): we adopt the recommended hyperparameters provided in the official implementation for all datasets<sup>18</sup>.
- For the remaining classical machine learning methods, including EM, GAIN, MICE, Miracle, MissForest, and Softimpute where hyperparameters might be important. Since we use the implementations from the ‘hyperimpute’ package, we tune the hyperparameters within the hyperparameter space provided in the package<sup>19</sup>. To be specific, we set the maximum budget as 50, then we sample 50 different hyperparameter combinations according to the hyperparameter space. Finally, we report the optimal performance over the 50 trials.

## E ADDITIONAL RESULTS

### E.1 MCAR: OUT-OF-SAMPLE IMPUTATION

We present the out-of-sample imputation performance comparison under MCAR setting in Table 5.

### E.2 MISSING AT RANDOM (MAR)

We present the performance comparison under MAR setting in Table 6, Table 7, and Table 8.

<sup>14</sup>[https://github.com/vanderschaarlab/hyperimpute/blob/main/src/hyperimpute/plugins/imputers/plugin\\_hyperimpute.py](https://github.com/vanderschaarlab/hyperimpute/blob/main/src/hyperimpute/plugins/imputers/plugin_hyperimpute.py)

<sup>15</sup>[https://github.com/maxiaoba/GRAPE/blob/master/train\\_mdi.py](https://github.com/maxiaoba/GRAPE/blob/master/train_mdi.py)

<sup>16</sup><https://github.com/G-AILab/IGRM/blob/main/main.py>

<sup>17</sup>[https://github.com/pfnet-research/TabCSDI/blob/main/config/census\\_onehot\\_analog.yaml](https://github.com/pfnet-research/TabCSDI/blob/main/config/census_onehot_analog.yaml)

<sup>18</sup><https://github.com/trevor-richardson/MCFlow/blob/master/main.py>

<sup>19</sup>[https://github.com/vanderschaarlab/hyperimpute/tree/main/src/hyperimpute/plugins/imputers/plugin\\_missforest.py](https://github.com/vanderschaarlab/hyperimpute/tree/main/src/hyperimpute/plugins/imputers/plugin_missforest.py)

Table 5: MCAR, Out-of-sample imputation performance on MAE and RMSE metrics (using base  $10^{-2}$  for better presentation). DIFFPUTER outperforms the most competitive baseline methods by 13.37% on MAE, and by 4.43% on RMSE.

	Method	California	Letter	Gesture	Adult	Default	Shoppers	Beijing	Average
MAE	MCFlow	60.74 $\pm$ 5.99	66.35 $\pm$ 0.62	82.72 $\pm$ 2.43	77.54 $\pm$ 0.98	64.12 $\pm$ 2.16	73.90 $\pm$ 3.34	45.56 $\pm$ 2.76	67.27
	IGRM	138.22 $\pm$ 4.25	121.59 $\pm$ 1.16	149.88 $\pm$ 3.92	149.64 $\pm$ 0.18	OOM	132.28 $\pm$ 0.01	OOM	-
	GRAPE	36.54 $\pm$ 0.12	42.41 $\pm$ 0.19	92.50 $\pm$ 0.07	59.25 $\pm$ 0.58	57.51 $\pm$ 0.09	51.94 $\pm$ 1.20	45.11 $\pm$ 0.95	55.04
	MOT	44.20 $\pm$ 0.81	46.75 $\pm$ 0.13	36.74 $\pm$ 0.01	51.17 $\pm$ 0.32	31.42 $\pm$ 0.50	41.55 $\pm$ 0.42	32.32 $\pm$ 0.12	40.59
	Remasker	33.97 $\pm$ 0.40	34.03 $\pm$ 0.33	35.29 $\pm$ 4.13	49.27 $\pm$ 0.53	37.02 $\pm$ 3.38	41.59 $\pm$ 1.39	30.53 $\pm$ 0.37	37.39
	DIFFPUTER	<b>33.47<math>\pm</math>0.25</b>	<b>30.69<math>\pm</math>0.61</b>	<b>28.53<math>\pm</math>0.18</b>	<b>49.14<math>\pm</math>0.62</b>	<b>24.63<math>\pm</math>0.24</b>	<b>34.46<math>\pm</math>0.48</b>	<b>27.96<math>\pm</math>0.51</b>	<b>32.70</b>
	Improv.	<b>1.47%</b>	<b>9.81%</b>	<b>19.16%</b>	<b>0.26%</b>	<b>21.61%</b>	<b>17.06%</b>	<b>8.42%</b>	<b>13.37%</b>
RMSE	MCFlow	84.57 $\pm$ 7.57	87.73 $\pm$ 0.58	119.73 $\pm$ 1.10	120.68 $\pm$ 0.66	101.33 $\pm$ 1.82	115.75 $\pm$ 0.82	90.75 $\pm$ 3.37	102.93
	IGRM	170.00 $\pm$ 4.72	151.16 $\pm$ 2.05	172.25 $\pm$ 2.99	179.17 $\pm$ 0.88	OOM	174.42 $\pm$ 1.73	OOM	-
	GRAPE	56.05 $\pm$ 0.03	58.11 $\pm$ 0.29	120.31 $\pm$ 0.11	99.19 $\pm$ 0.50	91.94 $\pm$ 0.42	92.37 $\pm$ 0.41	83.61 $\pm$ 3.17	72.88
	MOT	71.02 $\pm$ 1.19	64.69 $\pm$ 0.44	71.64 $\pm$ 0.84	94.67 $\pm$ 1.50	71.35 $\pm$ 1.08	82.63 $\pm$ 2.19	77.23 $\pm$ 2.54	76.18
	Remasker	<b>52.96<math>\pm</math>0.09</b>	48.60 $\pm$ 0.45	68.78 $\pm$ 3.60	<b>90.38<math>\pm</math>1.41</b>	76.94 $\pm$ 1.72	80.00 $\pm$ 0.88	73.80 $\pm$ 2.63	70.20
	DIFFPUTER	54.18 $\pm$ 0.25	<b>47.74<math>\pm</math>0.39</b>	<b>60.90<math>\pm</math>0.52</b>	91.31 $\pm$ 0.61	<b>70.50<math>\pm</math>0.83</b>	<b>74.01<math>\pm</math>0.40</b>	<b>71.03<math>\pm</math>1.95</b>	<b>67.09</b>
	Improv.	-	<b>1.77%</b>	<b>11.46%</b>	-	<b>1.19%</b>	<b>7.49%</b>	<b>3.75%</b>	<b>4.43%</b>

Table 6: MAR, In-sample imputation MAE.

	Method	California	Letter	Gesture	Magic	Bean	Adult	Shoppers	News	Default	Beijing
<i>Traditional iterative</i>	EM	39.17	58.05	35.43	48.13	15.41	62.54	43.01	39.52	33.78	42.35
	MICE	58.56	82.32	65.14	72.81	22.17	98.44	75.70	61.12	63.92	71.21
	MIRACLE	47.13	71.10	62.56	53.87	29.06	73.06	56.89	38.21	43.72	53.24
	SoftImpute	59.47	67.01	53.95	58.55	31.55	70.39	61.91	64.85	42.66	51.28
	MissForest	45.90	62.51	37.69	47.70	29.21	73.25	41.97	41.18	38.09	53.42
<i>Dist. match</i>	MOT	44.52	50.86	44.38	46.95	30.41	58.00	40.53	49.92	36.48	41.57
	TDM	39.27	46.86	42.69	45.31	29.11	59.59	45.00	63.25	37.93	40.32
<i>GNN</i>	GRAPE	35.14	42.45	38.08	39.87	18.81	60.73	53.73	50.85	48.43	51.61
	IGRM	35.23	41.31	36.45	40.09	18.97	61.35	59.86	-	-	-
<i>Generative models</i>	MIWAE	76.27	72.30	57.30	77.31	62.64	57.99	46.74	63.26	51.10	55.62
	GAIN	69.68	67.50	73.70	66.11	53.46	92.81	46.69	57.49	60.04	65.23
	MCFlow	60.00	68.23	65.59	55.27	30.40	84.36	72.99	56.36	67.87	63.75
	TabCSDI	83.17	77.42	54.64	76.93	75.34	57.11	58.53	66.41	46.48	69.89
	MissDiff	119.49	56.69	229.19	55.82	45.62	55.01	48.19	54.41	50.62	111.03
<i>Other</i>	KNN	49.25	49.41	45.52	49.70	30.81	50.83	46.81	48.53	35.15	39.31
	ReMasker	34.26	33.93	36.52	50.21	16.85	52.60	38.24	30.37	40.58	35.69
	HyperImpute	36.22	39.74	36.44	42.23	16.93	49.91	36.93	26.85	29.66	35.37
	DIFFPUTER	32.66	32.22	30.39	29.66	15.92	49.51	35.19	28.22	26.61	32.53

### E.3 MISSING NOT AT RANDOM (MNAR)

We present the performance comparison under MNAR setting in Table 9, Table 10, and Table 11.

Table 7: MAR, In-sample imputation RMSE.

Method	California	Letter	Gesture	Magic	Bean	Adult	Default	Shoppers	News	Beijing
<i>Traditional iterative</i>										
EM	60.28	77.77	67.68	74.17	39.37	88.74	79.48	65.76	82.32	78.25
MICE	85.13	106.80	92.85	102.74	49.96	126.50	109.18	89.10	105.35	105.60
MIRACLE	105.54	96.64	104.10	105.48	100.19	102.18	107.70	105.52	105.53	103.06
SoftImpute	83.23	88.33	99.13	87.11	56.32	95.44	100.26	94.04	90.61	89.30
MissForest	72.08	84.20	75.70	74.75	51.24	106.97	83.22	70.15	93.03	104.47
<i>Dist. match</i>										
MOT	74.86	69.71	86.57	76.53	57.92	85.96	88.74	81.01	92.78	82.33
TDM	67.88	64.66	80.44	80.77	62.96	86.98	91.34	95.22	95.02	80.50
<i>GNN</i>										
GRAPE	55.73	57.71	72.88	66.88	42.33	86.67	96.19	80.85	95.50	90.65
IGRM	56.13	56.44	69.51	67.45	41.61	85.55	97.30	0.00	0.00	0.00
<i>Generative models</i>										
MIWAE	105.54	96.64	104.10	105.48	100.19	102.18	107.70	105.52	105.53	103.06
GAIN	98.01	88.89	112.56	93.54	79.02	123.26	103.13	90.26	117.09	108.51
MCFlow	82.84	90.07	109.49	89.02	52.83	112.48	117.18	86.55	112.21	99.69
TabCSDI	111.64	100.27	101.72	101.65	101.15	87.37	101.74	100.07	104.73	109.09
MissDiff	256.09	79.60	1183.01	90.38	97.26	98.17	121.62	676.10	1464.76	423.92
<i>Other</i>										
KNN	80.05	68.59	92.64	75.52	62.92	90.24	85.50	74.04	89.12	70.74
ReMasker	56.08	47.97	70.55	77.09	36.87	78.08	78.43	58.94	93.74	75.19
HyperImpute	61.94	57.50	71.98	69.71	39.01	88.73	79.84	58.96	71.46	81.93
DIFFPUTER	57.12	48.12	68.59	67.29	35.19	79.42	77.52	60.11	65.53	71.38

Table 8: MAR, Out-of-sample imputation performance on MAE and RMSE metrics.

	Method	California	Letter	Gesttrue	Magic	Bean	Adult	Default	Shoppers	News	Beijing
MAE	MCFlow	54.34	69.83	83.55	64.87	46.55	78.31	77.20	91.12	61.01	61.46
	IGRM	86.61	114.96	125.91	97.17	92.10	85.58	64.32	-	-	-
	GRAPE	32.54	45.97	81.08	58.99	14.74	56.12	51.71	89.43	54.06	44.81
	MOT	47.13	61.02	42.69	61.32	47.64	54.36	45.08	53.76	43.98	34.51
	KNN	58.60	53.29	42.87	50.71	19.36	59.84	48.43	46.70	29.51	35.83
	MissDiff	76.35	58.62	137.93	52.24	29.56	65.61	46.68	51.78	164.06	126.31
	ReMasker	34.45	36.71	37.10	59.09	11.64	47.15	40.37	38.75	36.57	29.43
	DIFFPUTER	34.18	31.19	29.83	36.16	11.08	48.09	26.81	36.25	32.09	28.91
RMSE	MCFlow	76.15	91.91	117.75	99.19	77.20	115.98	118.85	811.81	95.71	98.14
	IGRM	122.73	140.69	156.16	135.74	119.44	129.87	121.19	-	-	-
	GRAPE	51.42	62.36	109.57	90.57	31.73	94.06	92.37	805.00	82.69	80.16
	MOT	71.10	81.90	79.62	92.84	68.51	92.28	84.57	83.39	79.65	73.90
	KNN	105.10	72.92	85.13	78.14	45.59	102.01	80.49	74.06	74.88	80.73
	MissDiff	152.68	81.17	495.66	90.04	56.94	113.84	81.80	1646.18	13653.41	438.17
	ReMasker	47.68	51.98	70.74	88.51	31.18	86.30	77.54	776.25	70.65	70.94
	DIFFPUTER	50.84	47.91	61.29	65.84	29.65	87.66	72.51	75.86	67.84	71.19

Table 9: MNAR, In-sample imputation MAE.

Method	California	Letter	Gesture	Magic	Bean	Adult	Shoppers	News	Default	Beijing
<i>Traditional iterative</i>										
EM	38.71	56.10	37.49	50.75	10.48	56.17	42.57	38.93	30.06	37.77
MICE	57.63	81.26	66.33	77.06	17.12	98.34	76.30	62.27	59.55	73.40
MIRACLE	42.42	71.68	71.83	45.59	15.27	59.32	78.56	41.05	37.41	43.59
SoftImpute	64.13	65.58	55.31	60.76	28.86	61.11	60.12	63.48	43.39	46.02
MissForest	43.11	58.61	39.74	51.15	25.37	57.40	42.58	39.57	34.49	45.18
<i>Dist. match</i>										
MOT	42.48	47.70	44.78	48.05	25.21	49.39	42.20	48.14	35.84	33.60
TDM	35.93	45.22	43.43	46.24	20.91	50.29	43.48	53.17	41.07	49.88
<i>GNN</i>										
GRAPE	35.82	41.00	40.46	40.68	13.24	57.97	55.06	49.36	44.70	51.55
IGRM	34.89	40.18	40.82	40.92	13.57	54.52	58.14	0.00	0.00	0.00
<i>Generative models</i>										
MIWAE	75.91	83.18	60.93	79.34	62.09	59.03	49.34	64.40	56.06	58.00
GAIN	79.60	67.04	83.73	64.75	39.68	123.87	55.89	57.33	85.44	76.43
MCFlow	59.23	65.53	67.03	60.37	26.30	81.16	77.87	54.96	67.03	60.15
TabCSDI	72.26	77.71	58.16	75.46	75.69	62.09	60.29	65.88	58.88	60.58
MissDiff	52.37	49.11	325.12	113.14	38.52	98.34	168.17	75.23	63.11	34.98
<i>Other</i>										
KNN	54.23	53.21	48.79	49.26	24.00	62.99	46.27	47.83	40.17	45.24
ReMasker	33.22	33.65	38.98	75.05	12.94	47.66	39.89	29.24	39.44	31.38
HyperImpute	33.96	40.20	38.30	42.88	11.96	50.19	39.80	31.74	28.36	35.16
DIFFPUTER	34.08	33.13	31.54	40.82	11.42	48.59	37.25	28.51	26.94	32.13

Table 10: MNAR, In-sample imputation RMSE.

Method	California	Letter	Gesture	Magic	Bean	Adult	Default	Shoppers	News	Beijing
<i>Traditional iterative</i>										
EM	59.75	75.78	74.88	76.01	28.90	92.30	83.26	66.16	71.03	80.53
MICE	84.36	105.91	98.00	106.59	41.91	132.01	112.55	91.16	96.87	110.40
MIRACLE	73.98	109.85	125.48	74.93	50.36	106.10	150.95	82.04	100.59	127.62
SoftImpute	102.74	87.02	103.04	90.21	48.27	96.65	99.56	95.59	83.84	88.51
MissForest	65.80	80.53	81.17	76.66	42.72	99.58	91.87	70.15	85.34	103.21
<i>Dist. match</i>										
MOT	70.10	66.95	88.63	76.35	49.43	91.16	93.92	83.08	86.44	83.23
TDM	59.81	63.17	86.36	78.16	47.59	91.81	96.36	88.88	94.51	92.89
<i>GNN</i>										
GRAPE	55.78	56.43	78.33	63.91	32.66	94.56	99.99	81.17	84.97	92.97
IGRM	55.34	55.32	78.74	64.30	32.57	93.64	97.58	—	—	—
<i>Generative models</i>										
MIWAE	102.77	107.54	111.03	107.10	96.50	102.70	113.33	107.99	114.09	107.78
GAIN	106.38	88.99	128.97	89.24	58.92	184.39	116.56	93.92	145.92	133.40
MCFlow	82.16	86.87	114.77	91.89	47.98	117.01	123.68	88.87	105.53	101.86
TabCSDI	93.59	101.11	107.65	100.20	101.99	98.54	106.69	102.92	107.07	104.25
MissDiff	137.09	68.98	6116.33	227.21	61.13	229.50	893.48	449.46	829.78	87.14
<i>Other</i>										
KNN	83.15	73.01	101.66	75.44	53.33	103.03	86.35	85.52	98.56	92.82
ReMasker	54.34	48.02	75.66	102.23	32.15	85.60	82.65	58.61	80.17	77.93
HyperImpute	58.93	57.84	75.84	69.87	31.24	89.04	86.96	61.48	76.42	83.08
DIFFPUTER	57.48	48.54	69.54	67.66	32.59	85.26	78.82	59.83	66.74	72.88



Table 11: MNAR, Out-of-sample imputation performance on MAE and RMSE metrics.

	Method	California	Letter	Gestrue	Magic	Bean	Adult	Default	Shoppers	News	Beijing
MAE	MCFLOW	64.23	65.18	85.36	60.98	29.30	83.03	76.07	87.28	72.31	57.94
	IGRM	117.60	129.00	158.38	102.21	122.01	99.16	76.46	-	-	-
	GRAPE	37.73	42.90	96.20	55.52	17.72	62.80	48.99	85.83	63.88	54.00
	MOT	46.90	47.66	39.16	49.99	25.47	52.79	40.86	46.32	34.14	34.44
	KNN	54.85	50.38	47.11	57.39	22.64	64.49	43.72	45.90	30.16	43.41
	MissDiff	73.40	48.97	205.74	171.93	44.65	104.14	159.56	88.56	43.58	33.29
	ReMasker	33.40	33.86	39.40	78.60	12.94	54.12	38.70	37.86	42.33	32.96
	DIFFPUTER	33.81	31.51	30.84	38.53	11.42	49.14	27.83	36.94	32.49	29.18
RMSE	MCFLOW	90.73	86.13	126.45	92.99	48.42	125.68	115.22	119.79	114.70	100.65
	IGRM	158.26	155.77	184.34	139.09	151.60	146.65	132.44	-	-	-
	GRAPE	56.95	58.69	124.21	83.88	35.08	102.13	88.56	121.00	103.72	98.57
	MOT	77.44	65.50	75.48	76.99	44.94	96.83	80.02	314.35	78.72	83.92
	KNN	86.58	68.61	93.60	90.29	47.98	105.68	81.98	76.92	80.41	82.96
	MissDiff	1651.42	69.37	3552.97	437.17	69.27	236.68	781.38	2778.77	476.74	78.07
	ReMasker	58.18	48.11	74.44	107.53	30.11	95.48	75.20	72.17	83.17	80.18
	DIFFPUTER	50.99	48.01	61.84	66.08	29.77	88.91	73.05	75.55	69.91	72.41