
Identifiable Generative Models for Missing Not at Random Data Imputation

Chao Ma^{1,2} * Cheng Zhang²

¹University of Cambridge ²Microsoft Research Cambridge
cm905@cam.ac.uk
cheng.zhang@microsoft.com

Abstract

Real-world datasets often have missing values associated with complex generative processes, where the cause of the missingness may not be fully observed. This is known as missing not at random (MNAR) data. However, many imputation methods do not take into account the missingness mechanism, resulting in biased imputation values when MNAR data is present. Although there are a few methods that have considered the MNAR scenario, their model’s identifiability under MNAR is generally not guaranteed. That is, model parameters can not be uniquely determined even with infinite data samples, hence the imputation results given by such models can still be biased. This issue is especially overlooked by many modern deep generative models. In this work, we fill in this gap by systematically analyzing the identifiability of generative models under MNAR. Furthermore, we propose a practical deep generative model which can provide identifiability guarantees under mild assumptions, for a wide range of MNAR mechanisms. Our method demonstrates a clear advantage for tasks on both synthetic data and multiple real-world scenarios with MNAR data.

1 Introduction

Missing data is an obstacle in many data analysis problems, which may seriously compromise the performance of machine learning models, as well as downstream tasks based on these models. Being able to successfully recover/impute missing data in an unbiased way is the key to understanding the structure of real-world data. This requires us to identify the underlying data-generating process, as well as the probabilistic mechanism that decides which data is missing.

In general, there are three types of missing mechanisms [44]. The first type is missing completely at random (MCAR), where the probability of a data entry being missing is independent of both the observed and unobserved data (Figure 1 (a)). In this case, no statistical bias is introduced by MCAR. The second type is missing at random (MAR), which assumes that the missing data mechanism is independent of the value of unobserved data (Figure 1 (b)). Under this assumption, maximum likelihood learning methods without explicit modeling of the missingness mechanism can be applied by marginalizing out the missing variables [3, 25, 28]. However, both MCAR and MAR do not hold in many real-world applications, such as recommender systems [8, 14], healthcare [13], and surveys [51]. For example, in a survey, participants with financial difficulties are more likely to refuse to complete the survey about financial incomes. This is an example of missing not at random (MNAR), where the cause of the missingness (financial income) can be unobserved. In this case, ignoring the missingness mechanism will result in biased imputation, which will jeopardize down-stream tasks.

*This work was performed when the authors were (part-time) associated with Microsoft Research, Cambridge

Figure 1: Exemplar missing data situations: (a): MCAR; (b): MAR; (c)-(i): MNAR.

There are few works considering the MNAR setting in scalable missing value imputation. On the one hand, many practical methods for MNAR does not have identifiability guarantees [24]. That is, the parameters can not be uniquely determined, even with access to infinite samples [38]. As a result, missing value imputation based on such parameter estimation could be biased. On the other hand, there are theoretical analyses on the identifiability in certain scenarios [35, 38, 53, 55, 57], but without associated practical algorithms for flexible and scalable settings (such as deep generative models). Moreover, MNAR data have many possible cases (Figure 1) based on different independence assumptions [38], making the discussion of identifiability difficult. This motivates us to fill this gap by extending identifiability results of deep generative models to different missing mechanisms, and provide a scalable practical solution. Our contributions are threefold:

- We provide a theoretical analysis of identifiability for generative models under different MNAR scenarios (Section 3). More specifically, we provide sufficient conditions, under which the ground truth parameters can be uniquely identified via maximum likelihood (ML) learning using observed information [24]. We also demonstrate how the assumptions can be relaxed in the face of real-world datasets. This provides foundation for practical solutions using deep generative models.
- Based on our analysis, we propose a practical algorithm model based on VAEs (Section 4), named GINA (deep generative imputation model for missing not at random). This enables us to apply flexible deep generative models in a principled way, even in the presence of MNAR data.
- We demonstrate the effectiveness and validity of our approach by experimental evaluations (Section 6) on both synthetic data modeling, missing data imputation in real-world datasets, as well as downstream tasks such as active feature selection under missing data.

2 Backgrounds

2.1 Problem Setting

A critical component to develop model to impute MNAR data is the model identifiability [43]. We give the definition below:

Definition 2.1 (Model identifiability). Assume $p(X)$ is a distribution of some random variable X , θ is its parameter that takes values in some parameter space. Then, if $p(X)$ satisfies $p_1(X) \neq p_2(X) \implies \theta_1 \neq \theta_2$, we say that $p(X)$ is identifiable w.r.t. on θ .

In other words, a model $p(X)$ is identifiable, if different parameter configurations implies a different probabilistic distributions over the observed variables. With identifiability guarantee, if the model assumption is correct, the true generation process can be recovered. Next, we introduce necessary notations and of missing data, and set up a concrete problem setting.

Basic Notation. Similar to the notations introduced by [1, 44], let X be the complete set of variables in the system of interest. We call X_o the set of observable variables. Let $I = \{1, \dots, D\}$ be the index set of all observable variables, i.e., $X_o = \{X_i | i \in I\}$. Let X_u denote the set of unobservable variables here. $O \subseteq I$ is a index set such that $X_o = \{X_i | i \in O\}$. We call O the observable pattern. Similarly, X_u denotes the set of missing/unobserved variables and $X = X_o \cup X_u$. Additionally, we use R to denote the missing mask indicator variable, such that $R_i = 1$ indicates X_i is observed, and $R_i = 0$ indicates otherwise. We call a probabilistic distribution $p(X)$ on X the reference distribution that is, the distribution that we would have observed if no missing mechanism is present; and we call the conditional distribution $p(R|X)$ the missing mechanism, which decides the probability of each X_i being missing. Then, we can define the marginal distribution of partially observed variables which is given by $\log p(X_o; R) = \log \int_{X_u} p(X_o; X_u; R) dX_u$. Finally, we will use lowercase letters to denote the realized values of the corresponding random variable. For example, $(x_o; r) \sim p(X_o; R)$ is the realization/samples of X_o and R , and the dimensionality of x_o may vary for each realizations.

Problem setting. Suppose that we have a ground truth data generating process, denoted by $p_D(X_o; R)$, from which we can obtain (partially observed) samples $(x_o; r) \sim p_D(X_o; R)$. We also have a model to be optimized, denoted by $p(\cdot | \theta)(X_o; X_u; R)$, where θ is the parameter of reference distribution $p(X)$, and ψ is the parameter of missing mechanism $p(R|X)$. Our goal can then be described as follows:

- To establish the identifiability of the model $p(\cdot | \hat{\theta})(X_o; R)$. That is, we wish to uniquely and correctly identify $\hat{\theta}$, such that $p_{\hat{\theta}}(X) = p_D(X)$, given in finite amount of partially observed data samples from ground truth $(x_o; r) \sim p_D(X_o; R)$.
- Then, given the identified parameter, we will be able to perform missing data imputation, using $p_{\hat{\theta}}(X_u | X_o)$. If our parameter estimate is unbiased, then our imputation is also unbiased, that is, $p_{\hat{\theta}}(X_u | X_o) = p_D(X_u | X_o)$ for all possible configurations of X_o .

2.2 Challenges in MNAR imputation

Recall the three types of missing mechanisms: if data is MCAR, $p(R|X) = p(R)$; if data is MAR, $p(R|X) = p(R|X_o)$; otherwise, we call it MNAR. When missing data is MCAR or MAR, missing mechanism can be ignored when performing maximum likelihood (ML) inference based only on the observed data [44], as:

$$\arg \max_{(\theta, \psi)} E_{(x_o; r) \sim p_D(X_o; R)} \log p(X_o = x_o) = \arg \max_{(\theta, \psi)} E_{(x_o; r) \sim p_D(X_o; R)} \log p(X_o = x_o; R = r)$$

where $\log p(X_o) = \log \int_{X_u} p(X_o; X_u) dX_u$. In practice, ML learning of X_o can be done by EM algorithm [3, 24]. However, when missing data is MNAR, the above argument does not hold, and the missing data mechanism cannot be ignored during learning. Consider the representative graphical model example in Figure 1 (d), which has appeared in many context of machine learning. In this graphical model X is the cause of R , and the connections between X and R are fully connected, i.e., each single node in R are caused by the entire set X . All nodes in R are conditionally independent from each other given X .

Clearly, this is an example of a data generating process with MNAR mechanism. In this case, Rubin proposed to jointly optimize both the reference distribution $p(X)$ and the missing data mechanism $p(R|X)$, by maximizing:

$$\arg \max_{(\theta, \psi)} E_{(x_o; r) \sim p_D(X_o; R)} [\log p(X_o = x_o) + \log p(R = r | X_o = x_o)] \quad (1)$$

This factorization is referred as selection modeling [12, 24]. There are multiple challenges if we want to Eq. 1 to obtain a practical model that provide unbiased imputation. First, we need model assumption to be consistent with the real-world data generation process $p_D(X_o; R)$. Given a wide range of possible MNAR scenarios, it is a challenge to design a general model. Secondly, the model need to be identifiable to enable the possibility to learn the underlying process which leads to unbiased imputation.

2.3 Variational Autoencoders and its identifiability

Variational auto-encoders [17, 41, 64] is a flexible deep generative model that is commonly used for estimating densities of $p_D(X)$. It takes the following form:

$$\log p(X) = \log \int_Z p(X|Z)p(Z); \quad (2)$$

where Z is some latent variable model with prior $p(Z)$, and $p(X|Z)$ is given by $p(X|Z) = N(f(Z); \Sigma)$, with $f(\cdot)$ being a neural network parameterized by θ . Generally, VAEs do not have identifiability guarantees w.r.t. [16]. Nevertheless, inspired by the identifiability of nonlinear ICA, [16] shows that the identifiability of VAE can be established up to equivalence permutation under mild assumptions, if the unconditional prior $p(Z)$ of VAE is replaced by the following the conditionally factorial exponentially family prior,

$$p_{T_i}(Z|U) \propto \prod_{i=1}^Y Q(Z_i) \exp\left[-\sum_{j=1}^K T_{ij}(Z_i) \eta_{ij}(U)\right]; \quad (3)$$

where U is some additional observations (called auxiliary variables), $Q(Z_i)$ is some base measure, $T_i(U) = (T_{i,1}, \dots, T_{i,K})$ the sufficient statistics, and $\eta_i(U) = (\eta_{i,1}, \dots, \eta_{i,K})$ the corresponding natural parameters. Then, the new VAE model given by

$$\log p(X|U) = \log \int_Z p(X|Z)p_{T_i}(Z|U) \quad (4)$$

is identifiable (Theorem 1 and 2 of [16], see Appendix G). We call the model (4) the identifiable VAE. Unfortunately, this identifiability results for VAE only hold when all variables are fully observed; thus, it cannot be immediately applied to address the challenges of dealing with MNAR data stated in Section 2.2. Next, we will analyze the identifiability of generative models under general MNAR settings (Section 3), and propose a practical method that can be used in MNAR (Section 4).

3 Establishing model identifiability under MNAR

One key issue of training probabilistic models under MNAR missing data is its identifiability. Recall that (Definition 2.1) model identifiability characterizes the property that the mapping from parameter to the distribution $p(X)$ is one-to-one. This is often closely related to maximum likelihood learning. In fact, it is not hard to show that Definition 2.1 is equivalent to the following Definition 3.1:

Definition 3.1 (Equivalent definition of identifiability). We say a model $p(X)$ is identifiable, if:

$$\arg \max_{\theta} E_{x \sim p(X)} \log p(X=x) = \theta; \quad (5)$$

In other words, the "correct" model parameter can be identified via maximum likelihood learning (under complete data), and the ML solution is unbiased. Similarly, when MNAR missing mechanism is present, we perform maximum likelihood learning on both X_0 and R using Eq. 1. Thus, we need $\log p_{\theta}(X_0; R)$ to be identifiable under MNAR, so that we can correctly identify the ground truth data generating process, and achieve unbiased imputation. The identifiability of $(X_0; R)$ under MNAR is usually not guaranteed, even in some simplistic settings. In this section, we will give the sufficient conditions for model identifiability under MNAR, and study how these can be relaxed for real-world applications.

3.1 Sufficient conditions for identifiability under MNAR

In this section, we give sufficient conditions where the model parameters can be uniquely identified by Rubin's objective, Eq. 1. Our aim is to i), find a set of model assumptions, so that it can cover many common scenarios and be flexible for practical interests; and ii), under those conditions, we want to show that its parameters can be uniquely determined by the partial ML solution Eq. 1. As shown in Figure 1, MNAR have many possible difference cases depending on its graphical structures. We want our results to cover every situation.

Instead of doing case by case analysis, we will start our identifiability analysis with one fairly general case as the example shown in Figure 1 (h) where the missingness can be caused by other partially

observed variable, by itself (self-masking) or by latent variables. Then, we will discuss how these analysis can be applied to other MNAR scenarios in Section 3.2.

Data setting D1. Suppose the ground truth data generation process satisfies the following conditions: all variables X are generated from a shared latent confounder Z and there are no connections among X ; and the missingness indicator variable can not be the parent of other variables. A typical example of such distribution is depicted in Figure 1(h). We further assume that $p_D(X_o; X_u; R)$ has the following parametric form $p_D(X_o; X_u; R) = \int_Z p_d(X_d|Z)p(Z)p(R|X; Z)dZ$, where $p(R|X; Z) = \prod_d p_d(R_d|X; Z)$, for some θ .

Then, consider the following model:

Model assumption A1. We assume that our model has the same graphical representation, as well as parametric form as data setting D1, that is, our model can be written as:

$$p_{\theta}; (X_o; R) = \int_{X_u; Z} \prod_d p_d(X_d|Z) \prod_d p_d(R_d|X; Z) p(Z) \quad (6)$$

Here, $(\theta; \gamma)$ are learnable parameters that belong to some parameter space Θ . Each θ is the parameter that parameterizes the conditional distribution that connects Z , $p_d(X_d|Z)$. Assume that the ground truth parameter θ^* belongs to the model parameter space; $\theta^* \in \Theta$.

Given such a model, our goal is to correctly identify the ground truth parameter settings given partially observed samples from $p_D(X_o; X_u; R)$. That is, let $(\hat{\theta}, \hat{\gamma}) = \arg \max_{(\theta; \gamma) \in \Theta} E_{(x_o; r) \sim p_D(X_o; R)} \log p_{(\theta; \gamma)}(X_o = x_o; R = r)$, we would like to achieve $\hat{\theta} = \theta^*$. In order to achieve this, we must make additional assumptions.

Assumption A2. Subset identifiability: There exist a partition of Θ , denoted by $A_1 = \{O_s\}_{s \in S}$, such that: for all $O_s \in A_1$, $p(X_{O_s})$ is identifiable on a subset of parameters $\theta_s \in O_s$.

This assumption basically formalizes the idea of divide and conquer: we partition the whole index set into several smaller subsets O_s , on which each reference distribution $p(X_{O_s})$ is only responsible for the identifiability on a subset of parameters.

Assumption A3. There exists a collection of observable patterns, denoted by O_1^0, \dots, O_L^0 , such that: 1) A_1 is a cover² of Θ ; 2) $p_D(X; R_{O_1^0} = 1; R_{\setminus O_1^0}) > 0$ for all $1 \leq l \leq L$; and 3), for all index $c \in O_1^0$, there exists $O_s \in A_1$ defined in A2, such that $c \in O_s \cap O_1^0$.

This assumption is about the strict positivity of the ground truth data generating process, $p_D(X_o; X_u; R)$. Instead of assuming that complete case data are available as in (1), here we assume we should at least have some observations $p(X_o; R_o = 1; R_u = 0) > 0$ for $O \in A_1$, on which $p(X_o)$ is identifiable.

To summarize, A1 ensures that our model has the same graphical representation/parametric forms as the ground truth, A2 $p(X_o) = \int_{X_u} p(X_o; X_u)dX_u$ should be at least identifiable for a collection of observable patterns that forms a partition of Θ and Assumption A3 ensures that $p_D(X_o; X_u; R)$ should be positive for certain important patterns (i.e., those on which $p(X_o)$ is identifiable). In Appendix C, we will provide a practical example that satisfies those assumptions. Given these assumptions, we have the following proposition (See Appendix C for proof.):

Proposition 1 (Sufficient conditions for identifiability under MNAR) Let $p_{\theta}; (X_o; X_u; R)$ be a model on the observable variables, and missing pattern R , and $p_D(X_o; X_u; R)$ be the ground truth distribution. Assume that they satisfy Data setting D1 Assumptions A1, A2 and A3.

Let $\hat{\theta} = \arg \max_{(\theta; \gamma) \in \Theta} E_{(x_o; r) \sim p_D(X_o; R)} \log p_{(\theta; \gamma)}(X_o = x_o; R = r)$ be the set of ML solutions of Equation 1. Then, we have $\hat{\theta} = \theta^*$. That is, the ground truth model parameter can be uniquely identified via (partial) maximum likelihood learning.

Missing value imputation as inference. Given a model $p_{\theta}(X_o; X_u)$, the missing data imputation problem can be then formulated by the Bayesian inference problem $\int p(X_u|X_o) / p(X_u; X_o)$. If the assumptions of Proposition 1 are satisfied, it enables us to correctly identify the ground truth

²It can be arbitrary partition in the set theory sense.

reference model parameter. Therefore, the imputed values sampled from the posterior $p(\mathbf{X}_u | \mathbf{X}_o)$ will be unbiased, and can be used for down stream decision making tasks.

Remark: Note that Proposition 1 can be extended to the case where model identifiability is defined by equivalence classes [16, 53]. See Appendix F for details.

3.2 Relaxing “correctness of parametric model” assumption (A1)

In this section, we further extend our previous results to the general MNAR cases including all different examples in Figure 1. In particular, we would like to see if the same model setting in Section 3.1 can be applied to scenarios where $p_D(\mathbf{X}_o; \mathbf{X}_u; R)$ and $p;(\mathbf{X}_o; \mathbf{X}_u; R)$ might have different parametric forms, or even different graphical representations.

To start with, we would like to point out that the mismatch between $p_D(\mathbf{X}_o; \mathbf{X}_u; R)$ and the model $p;(\mathbf{X}_o; \mathbf{X}_u; R)$ can be, to a certain extent, modeled by mappings between spaces of parameters. Let \mathcal{R}^I denote the parameter domain of our model, $(\mathbf{X}_o; \mathbf{X}_u; R)$. Suppose we have a mapping $\phi: \mathcal{R}^I \rightarrow \mathcal{R}^J$, such that $(\theta; \gamma) \in \mathcal{R}^I$ is mapped to another parameter space $(\theta; \gamma) = (\theta; \gamma) \in \mathcal{R}^J$ via ϕ . Here, \mathcal{R}^I is a subset of \mathcal{R}^J on which ϕ is defined. Then, the re-parameterized $p;(\mathbf{X}_o; \mathbf{X}_u; R)$ on parameter space can be rewritten as:

$$p;(\mathbf{X}_o; \mathbf{X}_u; R) := p^{-1}(\phi; \cdot)(\mathbf{X}_o; \mathbf{X}_u; R)$$

Assuming that the inverse mapping ϕ^{-1} exists. Then trivially, if $p;(\mathbf{X}_o; R)$ is identifiable with respect to θ and γ , then $p;(\mathbf{X}_o; R)$ should be also identifiable with respect to θ and γ :

Proposition 2. Let \mathcal{R}^I be the parameter domain of the model $(\mathbf{X}_o; \mathbf{X}_u; R)$. Assume that the mapping $\phi: (\theta; \gamma) \in \mathcal{R}^I \rightarrow (\theta; \gamma) \in \mathcal{R}^J$ is one-to-one on \mathcal{R}^I (equivalently, the inverse mapping $\phi^{-1}: \mathcal{R}^J \rightarrow \mathcal{R}^I$ is injective, and \mathcal{R}^I is its image set). Consider the induced distribution with parameter space, defined as $p;(\mathbf{X}_o; R) := p^{-1}(\phi; \cdot)(\mathbf{X}_o; R)$. Then, $p;(\mathbf{X}_o; R)$ is identifiable w.r.t. $(\theta; \gamma)$, if $p;(\mathbf{X}_o; R)$ is identifiable w.r.t. θ and γ .

Proposition 2 basically shows that if two distributions $p;(\mathbf{X}_o; R)$ and $p;(\mathbf{X}_o; R)$ are related by a mapping ϕ with nice properties, then the identifiability will translate between them. This already covers many scenarios of the data-model mismatch. For example, consider the case where ground truth data generation process satisfies the following assumption:

Data setting D2 Suppose the ground truth $p_D(\mathbf{X}_o; \mathbf{X}_u; R)$ satisfies: \mathbf{X} are all generated by shared latent confounders Z (as in D1), and R cannot be the cause of any other variables as in D1. Typical examples are given by any of the cases in Fig 1 (excluding (j) where R is the cause of R_2). Furthermore, the ground truth data generating process is given by the parametric form $p_D(\mathbf{X}_o; \mathbf{X}_u; R) = p;(\mathbf{X}_o; \mathbf{X}_u; R)$, where \mathcal{R} denotes its parameter space.

Then, for such ground truth data generating process, we can show that we can always find a model in the form of Equation 6, such that there exists some mapping ϕ that can model their relationship:

Lemma 1. Suppose the ground truth data generating process $p_D(\mathbf{X}_o; \mathbf{X}_u; R)$ satisfies setting D2. Then, there exists a model $p;(\mathbf{X}_o; \mathbf{X}_u; R)$, such that: 1) $p;(\mathbf{X}_o; \mathbf{X}_u; R)$ can be written in the form of Equation 6 (i.e. Assumption A1 and 2), there exists a mapping as described in Proposition 2, such that $p;(\mathbf{X}_o; R) = p^{-1}(\phi; \cdot)(\mathbf{X}_o; R)$, for all $(\theta; \gamma) \in \mathcal{R}$.

Model identification under data-model mismatch. Since we showed the identifiability can be preserved under the parameter space mapping (Proposition 2), next we will prove that if the model $p;(\mathbf{X}_o; \mathbf{X}_u; R)$ is trained on partially observed data points sampled from $p_D(\mathbf{X}_o; \mathbf{X}_u; R)$ that satisfies data setting D2 then the ML solution is still unbiased. For this purpose, inspired by Lemma 1, we work with the following additional assumption:

Model Assumption A4 Let $p;(\mathbf{X}_o; \mathbf{X}_u; R)$ denote our ground truth data generating process that satisfies data setting D2. Then, we assume our model $p;(\mathbf{X}_o; \mathbf{X}_u; R)$ is the one that satisfies the description given by Lemma 1. That is, its parametric form is given by Equation 6, and there exists a mapping ϕ as described in Proposition 2, such that $p;(\mathbf{X}_o; R) = p^{-1}(\phi; \cdot)(\mathbf{X}_o; R)$.

Then, we have the following proposition:

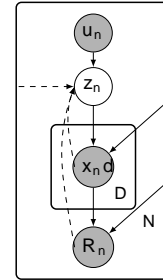
Proposition 3 (Sufficient conditions for identifiability under MNAR and data-model mismatch)
 Let $p_\theta(X_o; X_u; R)$ be a model on the observable variables X_o and missing pattern R , and $p_D(X_o; X_u; R)$ be the ground truth distribution. Assume that they satisfy Assumption A2, A3, and A4. Let $\hat{\theta} = \arg \max_{\theta} \int \int p_D(x_o, r) \log p_\theta(x_o; x_u; r) dx_o dr$ be the set of ML solutions of Equation 1. Then, we have $\hat{\theta} = \theta^*$. Namely, the ground truth model parameter θ^* of p_D can be uniquely identified (as $\hat{\theta}$) via ML learning.

Remark: practical implications Proposition 3 allows us to deal with the cases where the parameterization of ground truth data generating process and model distribution are related through a set of mappings f . In general, the graphical structure of $p_D(X_o; X_u; R)$ can be any cases in Figure 1 excluding (j). Then, in those cases, we are still able to use a model that corresponds to Equation 6 (Fig 1 (h)) to perform ML learning, provided that our model is flexible enough (Assumption A4). This greatly improves the applicability of our identifiability results, and we can build a practical algorithm based on Equation 6 to handle many practical MNAR cases.

4 GINA: A Practical Imputation Algorithm for MNAR

In the previous section, we have established the identifiability conditions for models in the form of Equation (6). However, in order to derive a practically useful algorithm, we still need to specify a parametric form of the model, that is both flexible and compatible with our assumptions. In this section, by utilizing the results in Section 3, we propose GINA, a deep generative imputation model for MNAR data (Figure 2). GINA fulfills identifiability assumptions above, and can handle general MNAR case as discussed in section 3.2. The code is released at <https://github.com/microsoft/project-azua>.

The parametric form of GINA We use utilize the flexibility of deep generative models to model the data generating process. We assume that the reference model $p(X)$ is parameterized by an identifiable VAE (see Section 2.3) to satisfy Assumption A2. That is, $p(X|U) = \int_Z dZ p(X|f(Z))p(Z|U)$, where U is some fully observed auxiliary inputs. The decoder $p(X|f(Z))$ is parameterized by a neural network $f: \mathbb{R}^H \rightarrow \mathbb{R}^D$. For convenience, we will drop the input U to $p(X|U)$, and simply use $p(X)$ to denote $p(X|U)$. Finally, for the missing model $p(R|X; Z)$, we use a Bernoulli likelihood model $p(R|X; Z) := \prod_d p_d(X; Z)^{R_d} (1 - p_d(X; Z))^{1 - R_d}$, where $p_d(X; Z)$ is parameterized by a neural network.



In Appendix G, we show that GINA fulfills the required assumptions of Proposition 1 and 3. Thus, we can use GINA to identify the ground truth data generating process, and perform missing value imputation under MNAR. The consistency of estimation result is also given in Appendix H.

Learning and imputation In practice, the joint likelihood in Equation 1 is intractable. Similar to the approach proposed in [2], we introduce a variational inference network $q(Z|X_o)$, which enable us to derive a importance weighted lower bound of $\log p_\theta(X_o; R)$:

$$\log p_\theta(X_o; R) \geq \mathbb{E}_{Z \sim q} \log \frac{1}{K} \sum_k w_k p(X_u|Z) q(Z|X_o)$$

where $w_k = \frac{p(R|X_o; X_u = x_u^k; Z = z^k) p(X_o; Z = z^k)}{q(Z = z^k|X_o)}$ is the importance weights. Note that we did not notate the missing pattern R as additional input to q , as this information is already contained in X_o . Then, we can optimize the parameters θ by solving the following optimization problem

$$\hat{\theta} = \arg \max_{\theta} \mathbb{E}_{(x_o, r) \sim p_D(X_o; R)} L_K(\theta; X_o = x_o; R = r)$$

Given $\hat{\theta}$, we can impute missing data by solving the approximate inference problem:

$$p(X_u|X_o) = \int_Z p(X_u|Z) p(Z|X_o) dZ \approx \int_Z p(X_u|Z) q(Z|X_o) dZ$$

Thus, GINA can be used to predict missing data even when the data are MNAR.

5 Related works

We mainly review recent works for handling MNAR data. In Appendix A, we provide a brief review of traditional methods that deal with MCAR and MAR.

When the missing data is MNAR, a general framework is to learn a joint model on both observable variables and missing patterns, in which a model of missing data is usually assumed [1]. This approach is also widely adopted in imputation tasks. For example, in the field of recommender systems, different probabilistic models are used within such a framework [58, 22, 21]. A similar approach has also been taken in the context of causal approach to imputation [20]. Similar to the use of the missing model, they have used an explicit model of exposure and adopted a causal view, where MNAR is treated as a confounding bias. Apart from these, inverse probability weighting methods are also used to debias the effect of MNAR [48, 58, 29] for imputation.

One issue that is often ignored by many MNAR methods is the model identifiability. Both parametric and non-parametric identifiability under MNAR has been discussed for certain cases [35, 57, 55, 53]). For example, [7] proposed the instrumental variable approach to help the identification of MNAR data. [3] investigated the identifiability of normal and normal mixture models, and showed that identifiability for parametric models is highly non-trivial under MNAR. [4] studied conditions for nonparametric identification using shadow variable technique. Despite the resemblance to the auxiliary variable in our approach, [3, 34] mainly considers the supervised learning (multivariate regression) scenario. [38, 37, 50] also discussed a similar topic based on a graphical and causal approach in a non-parametric setting. Although the notion of recoverability has been extensively discussed, their methods do not directly lead to practical imputation algorithms in a scalable setting. On the contrary, our work takes a different approach, in which we handle MNAR with a parametric setting, by dealing with learning and inference in latent variable models. We step aside from the computational burden with the help of recent advances in deep generative models for scalable imputation.

There has been a growing interest in applying deep generative models to missing data imputation. In [28, 25, 40], scalable methods for training VAEs under MAR have been proposed. Similar methods have also been advocated in the context of importance weighted VAEs, multiple imputation, and heterogeneous tabular data imputation [40, 26, 27]. Generative adversarial networks (GANs) have also been applied to MCAR data [63, 19]. More recently, deep generative models under MNAR have been studied [2, 5, 6], where different approaches such as selection models [57] and pattern-set mixture models [23] has been combined with partial variational inference for training VAEs. However, without additional assumptions, the model identifiability remains unclear in these approaches, and the posterior distribution of missing data conditioned on observed data might be biased.

6 Experiments

We study the empirical performance of the proposed algorithm of Section 4 with both synthetic data (Section 6.1) and two real-world datasets with music recommendation (Section 6.2) and personalized education (Section 6.3). The experimental setting details can be found in Appendix B.

6.1 Synthetic MNAR dataset

We first consider 3D synthetic MNAR datasets. We generate three synthetic datasets with nonlinear data generation process (shown in Appendix B.1). For all datasets, $X_2; X_3$ are generated via the latent variables $Z_1; Z_2; Z_3$, where X_1 are fully observed and X_2 and X_3 are MNAR. For dataset A, we apply self-masking (similar to Figure 1(c)) X_i will be missing if $X_i > 0$. For datasets B and C, we apply latent-dependent self-masking, X_i will be missing, if $g(X_i; Z_1; Z_2; Z_3) > 0$, where g is a linear mapping whose coefficients are randomly chosen.

We train GINA and baseline models with partially observed data. Then, we use the trained models to generate random samples. By comparing the generated samples with the ground truth data density, we can evaluate whether $p_\theta(X)$ is correctly identified. Results are visualized in Figure 3. In addition, we show the imputation results in Appendix J. Across three datasets, PVAE performs poorly, as it does not account for the MNAR mechanism. Not-MIWAE performs better than PVAE, as it is able to generate samples that are closer to the mode. However, it is still biased more towards the observed

Figure 3: Visualization of generated X_2 and X_3 from synthetic experiment. Row-wise (A-C) plots for dataset A, B, and C, respectively. Column-wise (i-iv): training set (only displays fully observed samples), PVAE samples, Not-MIWAE samples, and GINA samples, respectively. Our plot: kernel density estimate of ground truth density of complete data;

values. On the other hand, GINA is much more aligned to ground truth, and is able to recover the ground truth from partially observed data. This experiment showed the clear advantage of our method under different MNAR situations.

6.2 Recommender dataset imputation with random test set

Method	Test MSE
Matrix Factorization Methods	
PMF	1.401
IPW-PMF	1.375
Deconfounded-PMF	1.329
PMF-MNAR	1.483
PMF-MAR	1.480
VAE-based models	
PVAE	1.259 0.003
PVAE w/o IW	1.261 0.004
Not-MIWAE	1.078 0.000
GINA	1.052 0.002
Others	
CPTv-MNAR	1.056
Logitvd-MNAR	1.141
AutoRec	1.199
Oracle-test	1.057

Table 1: Test MSE on Yahoo! R3

Next, we apply our models to recommendation systems on Yahoo! R3 dataset [40, 60] for user-song ratings which is designed to evaluate MNAR imputation. It contains an MNAR training set of more than 300K self-selected ratings from 15,400 users on 1,000 songs, and an MCAR test set of randomly selected ratings from 5,400 users on 10 random songs. We train all models on the MNAR training set, and evaluate on MCAR test set. This is repeated 10 times with different random seeds. Both the missing model for GINA ($p(R_j|X; Z)$) and Not-MIWAE ($p(R_j|X)$) are parameterized by linear neural nets with Bernoulli likelihood functions. The decoders for GINA, PVAE, and Not-MIWAE uses Gaussian likelihoods with the same network structure. See Appendix B for implementation details and network structures.

We compare to the following baselines: 1), probabilistic matrix factorization (PMF) [2]; 2), inverse probability weighted PMF [48]; 3), Deconfounded PMF [60]; 4), PMF with MNAR/MAR data [8]; 5), CPTv and Logitv models for MNAR ratings [30]; 6), Oracle [8]: predicts ratings based on their marginal distribution in the test set; and 7) AutoRec [49]: Autoencoders that ignores missing data.

Results are shown in Table 1. Our method (GINA) gives the best performance among all methods. Also, VAE-based methods are consistently better than PMF-based methods, and MNAR-based models consistently outperform their MAR versions. More importantly, among VAE-based models, our GINA outperforms its non-identifiable counterpart (Not-MIWAE), and MAR counterpart (PVAE), where both models can not generate unbiased imputation.

6.3 Missing data imputation and active question selection on Eedi education dataset

Finally, we apply our methods to the Eedi education dataset, one of the largest real-world education response datasets. We consider the Eedi competition task 3 dataset, which contains over 1 million responses from 4918 students to 948 multiple-choice diagnostic questions. Each diagnostic question is a multiple-choice question. We consider predicting whether a student answers a question correctly or not. Over 70% of the entries are missing. The dataset also contains student metadata which we use as the auxiliary variables. In this experiment, we randomly split the data in a 90% train/10% test/10% validation ratio, and train our models on the response outcome data.

We evaluate our model on two tasks. Firstly, we perform missing data imputation, where different methods perform imputation over the test set. As opposed to Yahoo! R3 dataset, now the test set is MNAR, thus we use the evaluation method suggested by [6] where we evaluate the average per-question MSE. For each question, over all students with non-empty response. Then, the MSEs of all questions are averaged. We call this metric the debiased MSE. While regular MSE might be biased toward questions with more responses, the debiased MSE treats all questions equally, and can avoid selection bias to a certain degree. We report results for 10 repeats in the first column in Table 2. We can see that our proposed GINA achieves significantly improved results comparing to the baselines.

Secondly, we evaluate personalized education through active question selection on the test set which is task 4 from this competition dataset. The procedure is as follows: for each student in the test set, at each step, the trained generative models are used to decide which is the most informative missing question to collect next. This is done by maximizing the information reward [28] (see Appendix I for details). Since at each step, different students might collect different questions, there isn't a simple way to debias the predictive MSE as in the imputation task. Alternatively, we evaluate each method with the help of question meta data (difficulty level, which is a scalar). Intuitively, when the student response to the previously collected question is correct, we expect the next diagnostic question which has higher difficulty levels, and vice versa. Thus, we can evaluate the mean level change after correct/incorrect responses, and expect them to have significant differences. We also perform t-test between the level changes after incorrect/correct responses and report the p-value. We

Table 2: Performance on Eedi education dataset (with standard errors)

Method	Debiased MSE		Level change (correct)		Level change (incorrect)		p-value
PVAE	0.194	0.001	0.131	0.138	-0.101	0.160	0.514
Not-MIWAE	0.192	0.000	0.062	0.142	-0.073	0.179	0.561
GINA	0.188	0.001	0.945	0.151	-0.353	0.189	1.01 $\cdot 10^{-7}$

can see in Table 2, GINA is the only method that reports a significant p-value (<0.05) between the level changes of next collected questions after incorrect/correct responses which are desired. This further indicates that our proposed GINA predicts the unobserved answer with the desired behavior.

7 Conclusion

In this paper, we provide a analysis of identifiability for generative models under MNAR, and studies sufficient conditions of identifiability under different scenarios. We provide sufficient conditions under which the model parameters can be uniquely identified, via joint maximum likelihood learning on X_o and R . Therefore, the learned model can be used to perform unbiased missing data imputation. We proposed a practical algorithm based on VAEs, which enables us to apply flexible generative models that is able to handle missing data in a principled way. The main limitation of our proposed practical algorithm is the need for auxiliary variables, which is inherited from identifiable VAE models [16]. In practice, they may not be always available. For future work, we will investigate how to address such limitation, and how to extend to more complicated scenarios.

Acknowledgements and Disclosure of Funding

We thank Martin Kukla, Angus Lamb, Yingzhen Li, Dawen Liang, Chang Liu, Ruibo Tu, Yordan Zaykov (in alphabetical order) for helpful discussions and implementation support.

References

- [1] Paul D Allison. Missing data Sage publications, 2001.
- [2] James L Arbuckle, George A Marcoulides, and Randall E Schumacker. Full information estimation in the presence of incomplete data. *Advanced structural equation modeling: Issues and techniques* 243:277, 1996.
- [3] 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.
- [4] Craig K Enders and Deborah L Bandalos. The relative performance of full information maximum likelihood estimation for missing data in structural equation models. *Structural equation modeling* 8(3):430–457, 2001.
- [5] Sahra Ghalebikesabi, Rob Cornish, Chris Holmes, and Luke Kelly. Deep generative missingness pattern-set mixture models. *International Conference on Artificial Intelligence and Statistics* pages 3727–3735. PMLR, 2021.
- [6] Yu Gong, Hossein Hajimirsadeghi, Jiawei He, Thibaut Durand, and Greg Mori. Variational selective autoencoder: Learning from partially-observed heterogeneous data. *International Conference on Artificial Intelligence and Statistics* pages 2377–2385. PMLR, 2021.
- [7] James J Heckman. Sample selection bias as a specification error. *Econometrica: Journal of the Econometric Society* pages 153–161, 1979.
- [8] José Miguel Hernández-Lobato, Neil Houlsby, and Zoubin Ghahramani. Probabilistic matrix factorization with non-random missing data. *International Conference on Machine Learning* pages 1512–1520. PMLR, 2014.
- [9] Nicholas J Horton and Stuart R Lipsitz. Multiple imputation in practice: comparison of software packages for regression models with missing variables. *American Statistician* 55(3):244–254, 2001.
- [10] Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association* 47(260):663–685, 1952.
- [11] Joseph G Ibrahim, Stuart R Lipsitz, and M-H Chen. Missing covariates in generalized linear models when the missing data mechanism is non-ignorable. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 61(1):173–190, 1999.
- [12] Niels Bruun Ipsen, Pierre-Alexandre Mattei, and Jes Frellsen. not-mi-wae: Deep generative modelling with missing not at random data. *arXiv preprint arXiv:2006.12871*, 2020.
- [13] Janus Christian Jakobsen, Christian Gluud, Jørn Wetterslev, and Per Winkel. When and how should multiple imputation be used for handling missing data in randomised clinical trials—a practical guide with flowcharts. *BMC medical research methodology* 17(1):1–10, 2017.
- [14] Dietmar Jannach, Markus Zanker, Alexander Felfernig, and Gerhard Friesinger. *Recommender systems: an introduction*. Cambridge University Press, 2010.
- [15] Phimmarin Keerin, Werasak Kurutach, and Tossapon Boongoen. Cluster-based knn missing value imputation for dna microarray data. *2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC)* pages 445–450. IEEE, 2012.
- [16] Ilyes Khemakhem, Diederik Kingma, Ricardo Monti, and Aapo Hyvarinen. Variational autoencoders and nonlinear ica: A unifying framework. *International Conference on Artificial Intelligence and Statistics* pages 2207–2217. PMLR, 2020.

- [17] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114, 2013.
- [18] Daphne Koller and Nir Friedman. Probabilistic graphical models: principles and techniques. MIT press, 2009.
- [19] Steven Cheng-Xian Li, Bo Jiang, and Benjamin Marlin. Misgan: Learning from incomplete data with generative adversarial networks. arXiv preprint arXiv:1902.09599, 2019.
- [20] Dawen Liang, Laurent Charlin, and David M Blei. Causal inference for recommendation.
- [21] Dawen Liang, Laurent Charlin, James McInerney, and David M Blei. Modeling user exposure in recommendation. In Proceedings of the 25th international conference on World Wide Web, pages 951–961, 2016.
- [22] Guang Ling, Haiqin Yang, Michael R Lyu, and Irwin King. Response aware model-based collaborative filtering. arXiv preprint arXiv:1210.4869, 2012.
- [23] Roderick JA Little. Pattern-mixture models for multivariate incomplete data. Journal of the American Statistical Association, 88(421):125–134, 1993.
- [24] Roderick JA Little and Donald B Rubin. Statistical analysis with missing data, volume 793. John Wiley & Sons, 2019.
- [25] Chao Ma, Wenbo Gong, José Miguel Hernández-Lobato, Noam Koenigstein, Sebastian Nowozin, and Cheng Zhang. Partial vae for hybrid recommender system.
- [26] Chao Ma, Sebastian Tschiatschek, José Miguel Hernández-Lobato, Richard Turner, and Cheng Zhang. Vaem: a deep generative model for heterogeneous mixed type data. arXiv preprint arXiv:2006.11941, 2020.
- [27] Chao Ma, Sebastian Tschiatschek, Yingzhen Li, Richard Turner, Jose Miguel Hernandez-Lobato, and Cheng Zhang. Hm-vaes: a deep generative model for real-valued data with heterogeneous marginals. In Symposium on Advances in Approximate Bayesian Inference, pages 1–8. PMLR, 2020.
- [28] Chao Ma, Sebastian Tschiatschek, Konstantina Palla, José Miguel Hernández-Lobato, Sebastian Nowozin, and Cheng Zhang. Eddi: Efficient dynamic discovery of high-value information with partial vae. arXiv preprint arXiv:1809.11142, 2018.
- [29] Wei Ma and George H Chen. Missing not at random in matrix completion: The effectiveness of estimating missingness probabilities under a low nuclear norm assumption. arXiv preprint arXiv:1910.12774, 2019.
- [30] Benjamin M Marlin and Richard S Zemel. Collaborative prediction and ranking with non-random missing data. In Proceedings of the third ACM conference on Recommender systems, pages 5–12, 2009.
- [31] Herbert W Marsh. Pairwise deletion for missing data in structural equation models: Nonpositive definite matrices, parameter estimates, goodness of fit, and adjusted sample sizes. Structural Equation Modeling: A Multidisciplinary Journal, 5(1):22–36, 1998.
- [32] Pierre-Alexandre Mattei and Jes Frellsen. Miwae: Deep generative modelling and imputation of incomplete data sets. International Conference on Machine Learning, pages 4413–4423. PMLR, 2019.
- [33] Wang Miao, Peng Ding, and Zhi Geng. Identifiability of normal and normal mixture models with nonignorable missing data. Journal of the American Statistical Association, 111(516):1673–1683, 2016.
- [34] Wang Miao, Lan Liu, Eric Tchetgen Tchetgen, and Zhi Geng. Identification, doubly robust estimation, and semiparametric efficiency theory of nonignorable missing data with a shadow variable. arXiv preprint arXiv:1509.02556, 2015.

- [35] Wang Miao and Eric Tchetgen Tchetgen. Identification and inference with nonignorable missing covariate data. *Statistica Sinica* 28(4):2049, 2018.
- [36] Andriy Mnih and Russ R Salakhutdinov. Probabilistic matrix factorization. *Advances in neural information processing systems* 20:1257–1264, 2007.
- [37] Karthika Mohan and Judea Pearl. Graphical models for recovering probabilistic and causal queries from missing data. Technical report, CALIFORNIA UNIV LOS ANGELES COGNITIVE SYSTEMS LAB, 2014.
- [38] Karthika Mohan, Judea Pearl, and Jin Tian. Graphical models for inference with missing data. *Advances in neural information processing systems* 26:1277–1285, 2013.
- [39] Jared S Murray et al. Multiple imputation: a review of practical and theoretical findings. *Statistical Science* 33(2):142–159, 2018.
- [40] Alfredo Nazabal, Pablo M Olmos, Zoubin Ghahramani, and Isabel Valera. Handling incomplete heterogeneous data using variational autoencoders. *Pattern Recognition* 107:107501, 2020.
- [41] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *International conference on machine learning*, pages 1278–1286. PMLR, 2014.
- [42] James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association* 89(427):846–866, 1994.
- [43] Thomas J Rothenberg. Identification in parametric models. *Econometrica: Journal of the Econometric Society*, pages 577–591, 1971.
- [44] Donald B Rubin. Inference and missing data. *Biometrika* 63(3):581–592, 1976.
- [45] Donald B Rubin. Formalizing subjective notions about the effect of nonrespondents in sample surveys. *Journal of the American Statistical Association* 72(359):538–543, 1977.
- [46] Donald B Rubin. An overview of multiple imputation. *Proceedings of the survey research methods section of the American statistical association*, pages 79–84. Citeseer, 1988.
- [47] Donald B Rubin. *Multiple imputation for nonresponse in surveys*, volume 81. John Wiley & Sons, 2004.
- [48] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. Recommendations as treatments: Debiasing learning and evaluation. *arXiv preprint arXiv:1602.05352*, 2016.
- [49] Suvash Sedhain, Aditya Krishna Menon, Scott Sanner, and Lexing Xie. Autorec: Autoencoders meet collaborative filtering. In *Proceedings of the 24th international conference on World Wide Web*, pages 111–112, 2015.
- [50] Ilya Shpitser, Karthika Mohan, and Judea Pearl. Missing data as a causal and probabilistic problem. Technical report, CALIFORNIA UNIV LOS ANGELES DEPT OF COMPUTER SCIENCE, 2015.
- [51] Fiona M Shrive, Heather Stuart, Hude Quan, and William A Ghali. Dealing with missing data in a multi-question depression scale: a comparison of imputation methods. *Medical research methodology* 6(1):1–10, 2006.
- [52] Aude Sportisse, Claire Boyer, and Julie Josse. Imputation and low-rank estimation with missing not at random data. *Statistics and Computing* 30(6):1629–1643, 2020.
- [53] Aude Sportisse, Claire Boyer, and Julie Josse. Estimation and imputation in probabilistic principal component analysis with missing not at random data. *Advances in Neural Information Processing Systems* 33, 2020.

- [54] Daniel J Stekhoven and Peter Bühlmann. Missforest—non-parametric missing value imputation for mixed-type data. *Bioinformatics* 28(1):112–118, 2012.
- [55] Gong Tang, Roderick JA Little, and Trivellore E Raghunathan. Analysis of multivariate missing data with nonignorable nonresponse. *Biometrika* 90(4):747–764, 2003.
- [56] Ruibo Tu, Cheng Zhang, Paul Ackermann, Karthika Mohan, Hedvig Kjellström, and Kun Zhang. Causal discovery in the presence of missing data. *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1762–1770. PMLR, 2019.
- [57] Sheng Wang, Jun Shao, and Jae Kwang Kim. An instrumental variable approach for identification and estimation with nonignorable nonresponse. *Statistica Sinica*, pages 1097–1116, 2014.
- [58] Xiaojie Wang, Rui Zhang, Yu Sun, and Jianzhong Qi. Doubly robust joint learning for recommendation on data missing not at random. *International Conference on Machine Learning*, pages 6638–6647. PMLR, 2019.
- [59] Yixin Wang and David M. Blei. The blessings of multiple causes, 2019.
- [60] Yixin Wang, Dawen Liang, Laurent Charlin, and David M Blei. The deconfounded recommender: A causal inference approach to recommendation. *arXiv preprint arXiv:1808.06581*, 2018.
- [61] Zichao Wang, Angus Lamb, Evgeny Saveliev, Pashmina Cameron, Yordan Zaykov, José Miguel Hernández-Lobato, Richard E Turner, Richard G Baraniuk, Craig Barton, Simon Peyton Jones, Simon Woodhead, and Cheng Zhang. Diagnostic questions: The neurips 2020 education challenge. *arXiv preprint arXiv:2007.12061*, 2020.
- [62] Ian R White, Patrick Royston, and Angela M Wood. Multiple imputation using chained equations: issues and guidance for practice. *Statistics in medicine* 30(4):377–399, 2011.
- [63] Jinsung Yoon, James Jordon, and Mihaela Schaar. Gain: Missing data imputation using generative adversarial nets. *International Conference on Machine Learning*, pages 5689–5698. PMLR, 2018.
- [64] Cheng Zhang, Judith Bütepage, Hedvig Kjellström, and Stephan Mandt. Advances in variational inference. *IEEE transactions on pattern analysis and machine intelligence* 40(8):2008–2026, 2018.

A Traditional methods for handling missing data

Methods for handling missing data has been extensively studied in the past few decades. Those methods can be roughly classified into two categories: complete case analysis (CCA) based, and imputation based methods. CCA based methods, such as listwise deletion and pairwise deletion [31] focuses on deleting data instances that contains missing entries, and keeping those that are complete. Listwise/pairwise deletion methods are known to be unbiased under MCAR, and will be biased under MAR/MNAR. On the contrary, imputation based methods tries to replace missing values by imputed/predicted values. One popular imputation technique is called single imputation, where only produce one single set of imputed values for each data instance. Standard techniques of single imputation include mean/zero imputation, regression-based imputation of parametric methods [5, 54]. As opposed to single imputation, the multiple imputation (MI) methods such as MICE [62], was first proposed by Rubin [45, 46, 9, 39] is essentially a simulation-based methods that returns multiple imputation values for subsequent statistical analysis. Unlike single imputation, the standard errors of estimated parameters produced with MI is known to be unbiased apart from MI, there exists other methods such as full information maximum likelihood [20] and inverse probability weighting [2, 10], which can be directly applied to MAR without introducing additional bias. However, these methods assumes a MAR missing data mechanism, and cannot be directly applied to MNAR without introducing bias.

B Implementation details

We first introduce the general settings of GINA and other baselines. Our model (GINA) is based on the practical algorithm in Section 4. By default, we will set the auxiliary variable to be some fully observed meta feature (if there's any) or the missing mask pattern (if the dataset does not have a fully observed meta feature). The most important baselines are as follows: 1), Partial VAE (PVAE) [27], a VAE model with slightly modified ELBO objective, specially designed for MAR data; and 2), Not-MIWAE [12], a VAE model for MNAR data trained by jointly maximizing the likelihood on both the partially observed data and the missing pattern. As opposed to our model, the latent priors $p(Z)$ for both PVAE and Not-MIWAE are parameterized by a standard normal distribution, hence no auxiliary variables are used. Also, note that the graphical model of Not-MIWAE is described by Fig 1 (d), and does not handle the scenarios where the ground truth data distribution follows other graphs like Fig 1 (g). Finally, the inference model $q(Z|X)$ for the underlying VAEs is set to be diagonal Gaussian distributions whose mean and variance are parameterized by neural nets as in standard VAEs [17] (with missing values replaced by zero [12, 32]), or a permutation invariant set function proposed in [28]. See Appendix B for more implementation details for each tasks.

B.1 Synthetic dataset implementation details

Data generation The ground truth data generating process is given by $Z_1, Z_2, Z_3 \sim N(0, 1); X_1 = h_w(Z_1; Z_2; Z_3) + \epsilon_1; X_2 = f_1(X_1; Z_1; Z_2; Z_3) + \epsilon_2; X_3 = f_2(X_1; X_2; Z_1; Z_2; Z_3) + \epsilon_3$ where h_w is a linear mapping with coefficients w , f is some non-linear mapping whose functional form is given by Appendix B, ϵ_1 & ϵ_2 are two different sets of parameters for and $\epsilon_1, \epsilon_2, \epsilon_3$ are observational noise variables with mean 0, variance 0.01. We randomly sample three different sets of parameters, and generate the corresponding datasets (Figure 3), namely dataset A, B, and C. Each dataset consists of 2000 samples. Then, we apply different missing mechanisms for each dataset. For all datasets, we assume that X_1 are fully observed and X_2 and X_3 are MNAR. , and missing mechanism will be only applied to X_2 and X_3 . Finally, all observable variables are standardized.

Remark Note that in Dataset A, the ground truth missing mechanism does not depend on the latent variable model. Therefore, in this case, the not-MIWAE model does not have model-misspecification problem, hence any less satisfying performance is due to non-identifiability.

Network structure and training We use 5 dimensional latent space with fully factorized standard normal priors. The decoder part $p(X|Z)$ uses a 5-10 structure, where $D = 3$ in our case. For inference net, we use a zero imputation [28] with structure $D-10-10-5$, that maps the concatenation of observed data (with missing data filled with zero) and mask variables to distributional parameters of the latent space. For the factorized prior $p(Z|J)$ of the i-VAE component of GINA, we used

a linear network with one auxiliary input (which is set to be fully observed dimension). The missing mode $p(R_j|X)$ for GINA and i-NotMIWAE is a single layer neural network with 10 hidden units. All neural networks use Tanh activations (except for output layer, where no activation function is used). All baselines uses importance weighted VAE objective with 5 importance samples. The observational noise for continuous variables are $\text{std} = 2$. All methods are trained with Adam optimizer with batchsize with 100, and learning rate 0.001 for 20k epochs.

B.2 Yahoo!R3 experiment implementation details

Before training, all user ratings are scaled to be between 0 and 1 (such scaling will be reverted during evaluation). For all baselines, we use Gaussian likelihood with variance of 0.02. We use 20 dimensional latent space, and the decoder $p(X|Z)$ uses a 20-10 structure. We use Tanh activation function for the decoder (except for output layer, where no activation function is used). For inference net, we uses the point net structure proposed in [28], we use 20 dimensional feature mapping parameterized by a single layer neural network and 20 dimensional ID vectors for each variable. The symmetric operator is set to be the summation operator. The missing mode $p(R_j|X)$ for GINA and i-NotMIWAE is parameterized by linear neural network. All methods are trained with 400 epochs with batchsize 100.

B.3 Eedi dataset experiment implementation details

Since Eedi dataset is a binary matrix with 1/0 indicating that the student response is correct/incorrect, we use Bernoulli likelihood for decoder $p(X|Z)$. We use 50 dimensional latent space, and the decoder $p(X|Z)$ uses a 50-20-50 structure. Such structure is chosen using the validation set using grid search. We use ReLU activation function for the decoder (except for output layer, where no activation function is used). For inference net, we uses the point net structure that were used in Yahoo!R3 dataset. Here, the difference is that we use 50 dimensional feature mapping parameterized by a single layer neural network and 10 dimensional ID vectors for each variable. All methods are trained with 1k epochs with batchsize 100. A trick that we used for both not-MIWAE and GINA to improve the imputation performance, is to turn down the weight of the likelihood term for $p(R_j|X)$, by multiplying a factor of $w = 0.5$. This is due to that majority of the student response matrix is missing, the $p(R_j|X)$ will most likely dominate the training, hence the learning algorithm will prefer more about learning the models that explains the missing mechanism better, over the models that explains the observable variables better.

C Proof for Proposition 1

Proof : First, we show that $(X_{O_1}; R)$ is partially identifiable (i.e., identifiable on subset of parameters) on $\mathcal{G}_{d_2 O_1}$ for $\mathcal{O}_1 \subseteq \mathcal{A}_1$. We prove this by contradiction. Suppose there exists two different set of parameters $(\theta_1; \gamma_1)$ and $(\theta_2; \gamma_2)$, such that there exists at least one index O_1^0 for some, such that $\theta_1 \in \mathcal{G}_{d_2 O_1}$, and $p_{(\theta_1; \gamma_1)}(X_{O_1}; R) = p_{(\theta_2; \gamma_2)}(X_{O_1}; R)$. That is, $p(X_{O_1}; R)$ is not identifiable on $\mathcal{G}_{d_2 O_1}$.

According to Assumption A3, there exists $O_s \subseteq \mathcal{A}_1$, such that $O_s \cap O_1^0 = \emptyset$. Then, consider the marginal

$$p(X_{O_s}) = \int_{Z; R; X_{O_1}} \int_{d_2 O_s} p_d(X_d|Z) p(R_j|X; Z) p(Z) = p_{d_2 O_s}(X_{O_s})$$

. Since $p_{(\theta_1; \gamma_1)}(X_{O_1}; R) = p_{(\theta_2; \gamma_2)}(X_{O_1}; R)$, we have $p_{(\theta_1; \gamma_1)}(X_{O_s}) = p_{(\theta_2; \gamma_2)}(X_{O_s})$ (the joint uniquely determines marginals). However, this contradicts with Assumption A2 that $p_{O_s}(X_{O_s})$ is identifiable: this identifiability assumption implies that we should have $p_{(\theta_1; \gamma_1)}(X_{O_s}) \neq p_{(\theta_2; \gamma_2)}(X_{O_s})$. Therefore, by contradiction, we have $(X_{O_1}; R)$ is partially identifiable on $\mathcal{G}_{d_2 O_1}$ for $\mathcal{O}_1 \subseteq \mathcal{A}_1$.

Then, we proceed to prove that the ground truth parameter can be uniquely identified via ML learning. Based on our Assumption A1, upon optimal ML solution,

$$ML = \arg \max_{(\theta; \gamma)} E_{(x_o; r) \sim p_D(X_o; R)} \log p_{(\theta; \gamma)}(X_o = x_o; R = r)$$

, we have the following identity:

$$p_{(\theta; \phi)}(X_o; R) = p_{(\theta; \phi)}(X_o; R)$$

holds for all $(\theta; \phi) \in \Theta \times \Phi$, and all $\theta \in \Theta$ that satisfies $p_{\theta}(X_o; X_u; R_o = 1; R_u = 0) > 0$.

Note also that:

$$p_{(\theta; \phi)}(X_o; R) = \int_{Z: X_{in} = o} \int_{d \in \mathcal{D}} p_{\theta}^{ML}(X_d | Z) p_{\phi}^{ML}(R | X) p(Z)$$

, which depends on both θ and ϕ . Since we have already shown that $p_{\theta}(X_o; R)$ are partially identifiable on $\mathcal{D}_{\theta} \subseteq \mathcal{D}$ for $\theta \in \Theta$ and according to Assumption A3, $p_{\theta}(X_o; X_u; R_o = 1; R_{in} = 0) > 0$. Therefore, upon optimal solution

, we have that

$$f_{\theta} = g_{\theta}$$

holds for all $\theta \in \Theta$. Since we have assumed that Θ is a cover of \mathcal{D} , this guarantees that

$$f_{\theta} = g_{\theta}$$

for all d . In other words, we are able to uniquely identify from observed data, therefore

$$f_{\theta} = g_{\theta}$$

□

Remark (examples). To better illustrate the implication of the proposition, we provide an example that satisfies the assumptions of Proposition 1. One example is the self-masking one-way ANOVA [24], which contains l observable variables $X_i = f(X_1; \dots; X_l)g$, generated according to

$$X_i | Z_i \sim N(X_i; Z_i; \sigma^2); \quad i = 1; \dots; l;$$

where σ^2 is some known observational noise variance, $Z_i = f(Z_1; \dots; Z_l)g$ are latent variables generated by

$$Z_i \sim N(Z_i; \mu_i; \sigma^2);$$

Hence, the learnable parameters given by $\theta = f(\mu_1; \dots; \mu_l)g$. Then, assume that each observable variable X_i is missing MNAR according to the following mechanism:

$$p(R_i = 1 | X; Z) = \text{Sigmoid}(\mu_1 Z_i + \nu_i);$$

where $\nu_i = f(\mu_1; \dots; \mu_l)g$ is the set of learnable parameters for missing mechanism. Now, we verify that $p_{\theta}(X; Z)$ satisfies the assumptions of proposition 1:

- Assumption A2. By taking $A_1 = \{O_s | \mu_s \neq 0\}$, where $O_s = \{s | \mu_s \neq 0\}$, we have $p_{\theta}(X_s) = N(X_s; \mu_s; \sigma^2 + \sigma^2)$, which is identifiable over the subset of parameters $\theta = f(\mu_s; \dots; \mu_l)g$. Note that the partition $A_1 = \{O_s | \mu_s \neq 0\}$ is not unique; in fact, since $p_{\theta}(X_{O_s}) = \prod_{i \in O_s} N(X_i; \mu_i; \sigma^2 + \sigma^2)$ is identifiable on $\{s | \mu_s \neq 0\}$ for all non-empty $O_s \subseteq \{1, \dots, l\}$, any partition $A_1 = \{O_s | \mu_s \neq 0\}$ will satisfy Assumption A2.
- Assumption A3. Since $p_{\theta}(X)$ and $p_{\theta}(R | X; Z)$ are strictly positive for all possible settings of $X, R,$ and Z , Assumption A3 is trivially satisfied.

D Proof of Proposition 2

Proof Let $(\theta_1; \phi_1)$ and $(\theta_2; \phi_2)$ be two different parameters in $\Theta \times \Phi$. Then, we have

$$\begin{aligned} & p_{(\theta_1; \phi_1)}(X_o; R) \\ &= p_{(\theta_1; \phi_1)}(X_o; R) \\ & \neq p_{(\theta_2; \phi_2)}(X_o; R) \\ &= p_{(\theta_2; \phi_2)}(X_o; R) \end{aligned}$$

where the third line is due to the fact that θ_1 is injective and $p_{\theta_1}(X_o; R)$ is identifiable with respect to θ_1 and ϕ_1 . □

E Relaxing Assumption A1

E.1 Proof of Lemma 1

Lemma 1. Suppose the ground truth data generating process $(X_o; X_u; R)$ satisfies Assumption D2. Then, there exists a model $(X_o; X_u; R)$, such that: 1) $p_{D^*}(X_o; X_u; R)$ can be written in the form of Equation 6 (i.e. Assumption A1 and 2), there exists a mapping as described in Proposition 2, such that $p_{D^*}(X_o; R) = p_{D^*}(\cdot; \cdot)(X_o; R)$, for all $(\cdot; \cdot) \in \mathcal{D}$. Additionally, such is decoupled, i.e. $(\cdot; \cdot) = (\cdot; \cdot)$.

Proof: ³

Case 1 (connections among X): Suppose the ground truth data generating process $p_D(X; R)$ is given by Figure 1 (i). That is, $p_D(X; R) = \int p_i(X_i; Z; pa(X_i)) p(Z) dZ$. Without loss of generality, assume that probabilistic distribution $p_i(X_i; Z; pa(X_i))$ takes the form $p_i(X_i; Z; pa(X_i)) = f_i(X_i; pa(X_i); X; Z) p(i)$. Therefore, we have

$$\begin{aligned}
 & \int_{\mathcal{Z}} p_i(X_i; Z; pa(X_i)) p(Z) dZ \\
 &= \int_{\mathcal{Z}} \int_{\mathcal{Y}} p_i(X_i; Z; pa(X_i)) p(Z) dZ \\
 &= \int_{\mathcal{Z}} \int_{\mathcal{Y}} f_i(X_i; pa(X_i); X; Z) p(i) p(Z) dZ \\
 &= \int_{\mathcal{Z}} \int_{\mathcal{Y}} f_i(X_i; pa(X_i); X; Z) p(i) p(Z) dZ \\
 &= \int_{\mathcal{Z}} \int_{\mathcal{Y}} f_i(X_i; pa(X_i); X; Z) p(i) p(Z) dZ
 \end{aligned}$$

Apparently, there exists a set of functions $f_i(X_i; pa(X_i); X; Z)$, such that:

$$\begin{aligned}
 & \int_{\mathcal{Z}} \int_{\mathcal{Y}} f_i(X_i; pa(X_i); X; Z) p(i) p(Z) dZ \\
 &= \int_{\mathcal{Z}} \int_{\mathcal{Y}} f_i(X_i; pa(X_i); X; Z) p(i) p(Z) dZ \\
 &= \int_{\mathcal{Z}} \int_{\mathcal{Y}} f_i(X_i; pa(X_i); X; Z) p(i) p(Z) dZ
 \end{aligned}$$

³We mainly consider the case where all variables are continuous. Discrete variables will complicate the discussion, but will not change the conclusion.

Where $\text{anc}(i)$ is the shorthand for

$$\{k \mid X_k \text{ is an ancestor of } X_i\}$$

Note that, the graphical model of the new parameterization,

$$p(X) = \prod_{Z \in \mathcal{Z}} p(Z) \prod_{X_i \in \mathcal{X}} p(X_i \mid \text{anc}(i); Z) \prod_{X_j \in \mathcal{X}} p(X_j \mid Z)$$

has a new aggregated latent space \mathcal{Z} . That is, for each X_i that has non empty neighbour in \mathcal{X} , a new latent variable will be created. With this new latent space, the connections among \mathcal{X} can be decoupled, and the new graphical structure $\mathcal{G}(R)$ corresponds to Figure 1 (h).

The mapping that connects $p_i(X; R)$ and $p(X; R)$ can now be defined as identity mapping, since no new parameters are introduced/removed when reparameterizing $p_i(X; R)$ into $p(X; R)$. Hence, the two requirements of Lemma 1 are fulfilled.

Case 2(subgraph): Next, consider the case that the ground truth data generating process $p_D(X; R) = p_i(X; R)$ is given by one of the Figure 1 (a)-(g). That is, it is a subgraph of Figure 1 (h). Without loss of generality, assume that $\text{pa}(R_i) = \{X_j \mid p_i(X_j; Z) = \logit^{-1}(f_j(\text{pa}(R_i)))\}$, and $\text{pa}(R_i) \cap \mathcal{X} = \emptyset$; in other words, certain connections from \mathcal{X} to R_i is missing. Consider the model distribution parameterized by $p_i(X; Z) = \logit^{-1}(f_j(\text{pa}(R_i)) + g_j(X; Z; \text{pa}(R_i)))$, satisfying $g_j(\emptyset) = 0$. Therefore, the mapping ϕ^{-1} is given as $\phi^{-1}(i) := (i; i = 0)$. Apparently, ϕ^{-1} is injective, hence satisfying the requirement of Proposition 2.

□

E.2 Proof for Proposition 3

Proposition 3 (Sufficient conditions for identifiability under MNAR and data-model mismatch)

Let $p_i(X_o; X_u; R)$ be a model on the observable variables X and missing pattern R , and $p_D(X_o; X_u; R)$ be the ground truth distribution. Assume that they satisfy Data setting D2 Assumption A2, A3, and A4. Let $\hat{\theta} = \arg \max_{(\theta; \gamma)} \mathbb{E}_{(X_o; R)} p_D(X_o; R) \log p_i(X_o = x_o; R = r)$ be the set of ML solutions of Equation 1. Then, we have $\hat{\theta} = \theta^*$. Namely, the ground truth model parameter θ^* of p_D can be uniquely identified (as $\hat{\theta}$) via ML learning.

Proof: First, it is not hard to show that $p_i(X_o; R)$ is partially identifiable on \mathcal{O}_1^0 for \mathcal{A}_1 . This has been shown in the proof of Proposition 1, and we will not repeat this proof again.

Next, given data setting D2 and Assumption A4, define

$$\hat{\theta}_{ML} = \arg \max_{(\theta; \gamma)} \mathbb{E}_{(X_o; R)} p_D(X_o; R) \log p_i(X_o = x_o; R = r)$$

, then we have:

$$p_{(\hat{\theta}_{ML}; \hat{\gamma}_{ML})}(X_o; R) = p_{(\theta^*; \gamma^*)}(X_o; R)$$

holds for all $(\theta_{ML}; \gamma_{ML}) \in \hat{\theta}_{ML}$, and all \mathcal{O}_1^0 that satisfies $p(X_o; X_u; R_o = 1; R_u = 0) > 0$.

Since $p_i(X_o; R)$ are partially identifiable on \mathcal{O}_1^0 for \mathcal{A}_1 and according to Assumption A3, $p_D(X_o; X_u; R_o = 1; R_u = 0) > 0$. Therefore,

$$f_d = \phi^{-1}(\hat{\theta}_{ML}; \hat{\gamma}_{ML})$$

must be true for all $\mathcal{O}_1^0 \in \mathcal{A}_1$, where $\phi^{-1}(\hat{\theta}_{ML}; \hat{\gamma}_{ML})$ denotes the components of $\phi^{-1}(\hat{\theta}_{ML}; \hat{\gamma}_{ML})$ that corresponds to the entries of \mathcal{O}_1^0 . Since we have assumed that $X_o = 1$ in Assumption 3 (i.e. \mathcal{A}_1 is a cover of \mathcal{O}_1^0), this guarantees that

$$d^{ML} = f^{-1}(g; \theta)_d$$

for all d . In other words, we are able to uniquely identify from observed data, therefore

$$= f^{-1}(g; \theta)$$

Finally, according to Assumption 4 and the proof of Lemma 1, is decoupled as $(\theta; \theta)$. Therefore, we can write $= f^{-1}(g; \theta)$. That is, the ground truth model parameter of p_D can be uniquely identified (as (θ)).

□

F Identifiability based on equivalence classes

In this section, we introduce the notion of identifiability based on equivalence classes. Let \sim be an equivalence relation on a parameter space Θ . That is, it satisfies reflexivity ($\theta \sim \theta$), symmetry ($\theta \sim \theta' \iff \theta' \sim \theta$), and transitivity (if $\theta \sim \theta'$ and $\theta' \sim \theta''$, then $\theta \sim \theta''$). Then, an equivalence class of $\theta \in \Theta$ is defined as $[\theta] = \{\theta' \in \Theta : \theta' \sim \theta\}$. We denote this by $[\theta]$. Then, we are able to give the definition of model identifiability based on equivalence classes:

Definition F.1 (Model identifiability based on equivalence relation). Let $p(X; \theta)$ be a distribution of some random variable X , θ is its parameter that takes values in some parameter space Θ and \sim is some equivalence relation on Θ . Then, if $p(X; \theta) \neq p(X; \theta')$ for $\theta \in [\theta_1]$ and $\theta' \in [\theta_2]$, we say that p is identifiable w.r.t. \sim on Θ .

Apparently, Definition 2.1 is a special case of Definition F.1, where \sim is given by the equality operator, $\theta \sim \theta' \iff \theta = \theta'$. When the discussion is based on the identifiability under equivalence relation, then it is obvious that all the arguments of Proposition 1, 2, and 3 still holds. Also, the statement of the results needs to be adjusted accordingly. For example, in Proposition 1, instead of "the ground truth model parameter can be uniquely identified", we now have "the ground truth model parameter can be uniquely identified up to an equivalence relation,".

G Subset identifiability (A2) for identifiable VAEs

The GINA model needs to satisfy the requirement on model of Proposition 1 or 3, if we wish to use it to fit to the partially observed data and then perform (unbiased) missing data imputation. In order to show that the identifiability result of Proposition 1/3 can be applied to GINA, the key assumption that we need to verify is the local identifiability (Assumption A2).

To begin with, in [16], the following theorem on VAE identifiability has been proved:

Theorem 1. Assume we sample data from the model given by $p(X; Z; U) = p(X; f(Z; U))p_T(Z; U)$, where f is a multivariate function $f: \mathbb{R}^H \times \mathbb{R}^D \rightarrow \mathbb{R}^D$. $p_T(Z; U)$ is parameterized by exponential family of the form $p_T(Z; U) \propto \prod_{i=1}^M Q(Z_i) \exp[\sum_{j=1}^K T_{ij}(Z_i) \theta_{ij}(U)]$, where $Q(Z_i)$ is some base measure M is the dimensionality of the latent variable Z , $T_{ij}(U) = (T_{i,1}; \dots; T_{i,K})$ are the sufficient statistics, and $\theta_{ij}(U) = (\theta_{i,1}; \dots; \theta_{i,K})$ are the corresponding parameters, depending on U . Assume the following holds:

1. The set $X \times Z \times U \rightarrow \mathbb{R}^D$ has zero measure, where $\mathbb{1}_A$ is the characteristic function of A ;
2. The multivariate function f is injective;
3. T_{ij} are differentiable a.e., and $(T_{ij})_{1 \leq j \leq K}$ are linearly independent on any subset of U of measure greater than zero;
4. There exist $n_k + 1$ distinct points U^0, \dots, U^{n_k} , such that the matrix $L = (U^1; \dots; U^{n_k}; U^0)$ of size $(n_k + 1) \times n_k$ is invertible.

Then, the parameters $(\theta; T; \theta)$ are A -identifiable, where A is the equivalence class defined as (see also Appendix F):

$$(\theta; T; \theta) \sim (\theta'; T'; \theta') \iff \exists A; c; j; T(f^{-1}(X)) = AT(f^{-1}(X)) + c$$

. Here, A is a n_k by n_k matrix, and c is a vector.

Note that under additional mild assumptions, then the A equivalence relation can be further reduced to a permutation matrix. That is, the model parameters can be identified, such that the latent variables differs up to a permutation. This is inconsequential in many applications. We refer to [16] for more discussions on permutation equivalence.

So far, Theorem 1 only discussed the identifiability on the full variables, $X = X_o \oplus X_u$. However, in Assumption A2, we need the reference model to be (partially) identifiable on a partition $O_s \subseteq A_1, p(X_{O_s})$. Naturally, we need additional assumptions on the injective functions stated below:

Assumption A5 There exists an integer D_0 , such that $f_O : \mathbb{R}^H \rightarrow \mathbb{R}^{jO}$ is injective for all O that $|jO| \leq D_0$. Here, f_O is the entries from the output of that corresponds to the index set

Remark Note that, under assumption A5, the Assumption A3 in Section 3 becomes more intuitive: it means that in order to uniquely recover the ground truth parameters, our training data must contain training examples that have more than observed features. This is different from some previous works ([38] for example), where complete case data must be available.

Finally, given these new assumptions, it is easy to show that:

Corollary 1 (Local identifiability). Assume that $p(X; Z; jU) = p(X; f(Z))_{PT}$; $(Z; jU)$ is the model parameterized according to Theorem 1. Assume that the assumptions in Theorem 1 holds. Additionally, assume that satisfies assumption A5.

Then, consider the subset of variables. Then, $p(X_{OjU})$ is A -identifiable on $(f_O; T;)$ for all O that satisfies $|jO| \leq D_0$, where f_O is the entries from the output of that corresponds to the index set O .

Proof : it is trivial to see that the assumptions 1, 3, and 4 in Theorem 1 automatically holds regarding $p(X_{OjU})$. f_O is injective according to Assumption A5. Hence, $p(X_{OjU})$ satisfies all the assumptions in Theorem 1, and $p(X_{OjU})$ is A -identifiable on $(f_O; T;)$ for all O that satisfies $|jO| \leq D_0$. \square

Remark In practice, Assumption A5 is often satisfied. For example, consider that is parameterized by the following MLP composite function:

$$f(Z) = h(W \cdot g(Z)) \tag{7}$$

, where g is a D_0 dimensional, injective multivariate function: $\mathbb{R}^H \rightarrow \mathbb{R}^{D_0}$, h is some activation function $h : \mathbb{R} \rightarrow \mathbb{R}$, and W is a injective linear mapping $W : \mathbb{R}^{D_0} \rightarrow \mathbb{R}^D$ represented by the matrix $W_{D \times D_0}$, whose submatrices that consists of D_0 arbitrary selected columns are also injective. Note that this assumption is not hard to fulfill: a randomly generated matrix (e.g., with element-wise i.i.d. Gaussian prior) satisfies this condition with probability 1. To verify injective for all $|jO| \leq D_0$, notice that $f_O(Z) = h(W_O \cdot g(Z))$, where W_O is the output dimensions of W that corresponds to the index set. Since W is injective and $|jO| \leq D_0$, we have that W_O is also injective, hence f_O is also injective.

H Consistency of estimation for GINA

In [16], a result regarding the consistency of estimation for identifiable VAE. Similarly, it is trivial to show that similar result holds for GINA:

Theorem 2 (Consistency of estimation) Assuming that

1. $q(Z = z^k | X_o)$ is expressive enough to contain the true posterior $p(Z | X_o)$, for all X_o , and .
2. The model in Section 4 is correctly specified, and its parameters are estimated by maximizing $L_K(\hat{\theta}; X_o; R)$ w.r.t. θ, μ , and σ .

Then, under perfect information (infinite samples from data) and θ is recovered up to A .

Proof Since $(Z = z^k | X_o)$ is expressive enough to contain the true posterior $q(z; X_o; R)$, it recovers the true likelihood function $p(z; X_o; R)$ by simply maximizing $\log p(z; X_o; R)$. Therefore, the problem of maximizing $\log p(z; X_o; R)$ is equivalent to maximum likelihood estimation problem. Therefore, since we assumed that the model is correctly specified, the consistency of MLE trivially implies the consistency of GINA model trained via maximizing $\log p(z; X_o; R)$. \square

I Active question selection

Suppose X_o be the set of observed variables, that represents the correctness of student's response to questions that are presented to them. Then, in the problem of active question selection, we start with $O = \emptyset$, and we would like to decide which variable X_i from X_U to observe/query next, so that it will most likely provide the most valuable information for some target variable of interest, X_j ; Meanwhile, we should while make as few queries as possible. Once we have decided which X_i to observed next, we will make query and add O . This process is done by maximizing the information reward proposed by [28]:

$$i = \arg \max_{i \in U} R(i | X_o) := E_{X_i | p(X_i | X_o)} \text{KL} [p(X_j | X_i; X_o) | p(X_j | X_o)]:$$

In the Eedi dataset, as we do not have a specific target variable of interest, it is defined as $X_j = X_U$. In this case, X_j could be very high-dimensional, and direct estimation of $\text{KL} [p(X_j | X_i; X_o) | p(X_j | X_o)]:$ could be inefficient. In [28], a fast approximation has been proposed:

$$\begin{aligned} R(i | X_o) &= E_{X_i | p(X_i | X_o)} D_{\text{KL}} [p(Z | X_i; X_o) | p(Z | X_o)] \\ &= E_{X_i | p(X_i | X_o)} D_{\text{KL}} [p(Z | X_i; X_o) | p(Z | X_o)]: \\ &= E_{X_i | p(X_i | X_o)} D_{\text{KL}} [q(Z | X_i; X_o) | q(Z | X_o)] \\ &= E_{X_i | p(X_i | X_o)} D_{\text{KL}} [q(Z | X_i; X_o) | q(Z | X_o)]: \end{aligned}$$

In this approximation, all calculation happens in the latent space of the model, hence we can make use of the learned inference net to efficiently estimate $R(i | X_o)$.

J Additional results

J.1 Imputation results for synthetic datasets

In addition to the data generation samples visualized in Figure 3, we present the imputation results for synthetic datasets in Figure 4. The procedure of generating the imputed samples are as follows. First, each model are trained on the randomly generated, partially observed synthetic dataset described in Section 6.1. Once the models are trained, they are used to impute the missing data in the training set. For each training data, we draw exactly one sample from the (approximate) conditional distribution $p_\theta(X_U | X_o)$. Thus, we have "complete" version of the training set, one for each different model. Finally, we draw the scatter plot for each imputed training set, per dataset and per model. If the model is doing a good job recovering the ground truth distribution $p_\theta(X)$ from training set, then its scatter plot should be close to the KDE estimate of the ground truth density of complete data. According to Figure 4, the imputed distribution is similar to the generated distribution in Figure 3.

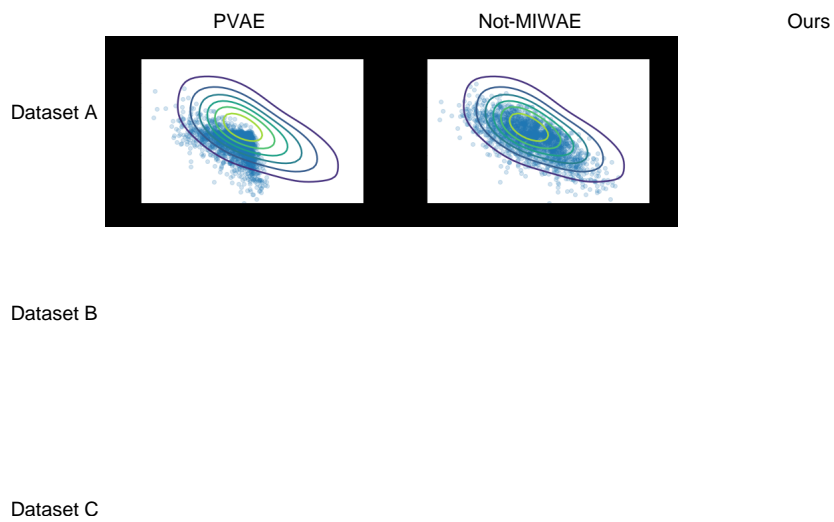


Figure 4: Visualization of imputed X_2 and X_3 from synthetic experiment. **Row-wise (A-C)** plots for dataset A, B, and C, respectively; **Column-wise:** PVAE imputed samples, Not-MIWAE imputed samples, and GINA imputed samples, respectively. **Contour plot:** kernel density estimate of ground truth density of complete data;