

Variational Gibbs Inference for Statistical Model Estimation from Incomplete Data

Vaidotas Simkus
Benjamin Rhodes
Michael U. Gutmann
*School of Informatics
University of Edinburgh*

VAIDOTAS.SIMKUS@ED.AC.UK
BEN.RHODES@ED.AC.UK
MICHAEL.GUTMANN@ED.AC.UK

Editor: Mohammad Emtiyaz Khan

Abstract

Statistical models are central to machine learning with broad applicability across a range of downstream tasks. The models are controlled by free parameters that are typically estimated from data by maximum-likelihood estimation or approximations thereof. However, when faced with real-world data sets many of the models run into a critical issue: they are formulated in terms of fully-observed data, whereas in practice the data sets are plagued with missing data. The theory of statistical model estimation from incomplete data is conceptually similar to the estimation of latent-variable models, where powerful tools such as variational inference (VI) exist. However, in contrast to standard latent-variable models, parameter estimation with incomplete data often requires estimating exponentially-many conditional distributions of the missing variables, hence making standard VI methods intractable. We address this gap by introducing variational Gibbs inference (VGI), a new general-purpose method to estimate the parameters of statistical models from incomplete data. We validate VGI on a set of synthetic and real-world estimation tasks, estimating important machine learning models such as variational autoencoders and normalising flows from incomplete data. The proposed method, whilst general-purpose, achieves competitive or better performance than existing model-specific estimation methods.

Keywords: statistical model estimation, variational inference, Gibbs sampling, missing data, amortised inference

1. Introduction

This paper introduces a new general-purpose method to estimate statistical models from incomplete data that is well-suited for modern (deep) statistical models. Estimating statistical models is one of the core tasks in machine learning because the fitted model can be used in many practical downstream tasks, such as, classification, prediction, anomaly detection, data augmentation, and missing data imputation (e.g. Goodfellow et al., 2016, Chapter 5.1.1). However, most of the current methods require large amounts of fully-observed data at training time, and hence they remain largely impractical in many real-world domains that are overwhelmed with incomplete data. For example, the vast amounts of data gathered by online systems is sparse, with ratings data used in recommender systems often missing 95-99% of the total data (Marlin et al., 2011). Similarly, a review of medical trial studies has identified that 95% of studies contained missing data, with as much as 70%

of the data values missing in some studies (Bell et al., 2014). This prevalence of missing data in real-world scenarios warrants a need for principled approaches to efficiently handle missing data in machine learning.

A principled classical approach to estimating statistical models from incomplete data is expectation-maximisation (EM, Dempster et al., 1977) which aims at maximising the likelihood, however, it is mostly limited to simple models. Monte Carlo EM (MCEM, Wei and Tanner, 1990) is a less limited version of classical EM that can be understood as an iterative method that fits the statistical model on imputations derived from itself. However, exact conditional sampling for modern statistical models is typically impossible. One way to (approximately) sample imputations from a joint statistical model is via Markov chain Monte Carlo methods (MCMC, e.g. Barber, 2017, Chapter 27.4), however they tend to be computationally intensive and hence scale poorly to larger data sets (Blei et al., 2017). Variational inference (VI, Jordan et al., 1999) is often a computationally more performant alternative, while sometimes lacking the asymptotic exactness guarantees of MCMC. In case of missing data, however, as we will elaborate in the paper (Sections 2.1-2.2), existing (amortised) VI methods would require $2^d - 1$ variational distributions, one for each non-trivial pattern of missingness, and thus scale poorly with the dimension of the data d . Their applicability to estimating statistical models from incomplete data has thus been strongly limited and our paper addresses this gap in the literature.

1.1 Main Contributions

Our main contribution is a novel general-purpose method for estimating statistical models from incomplete data. The method combines the computational performance of VI and the expressiveness of Markov chains, which makes it well-suited for modern (deep) statistical models. Crucially, the method only requires d rather than $2^d - 1$ variational distributions and thereby overcomes the limitations of existing (amortised) VI methods that prevented their use for model estimation from incomplete data. We achieve this reduction from exponential to linear growth by leveraging techniques that are related to those used by popular imputation methods (see Section 2.3). As the proposed method is based on variational inference (VI) and the Gibbs sampler, we call it variational Gibbs inference (VGI).

Figure 1 illustrates VGI.¹ The method starts with K initial random imputations of the missing values of each incomplete data-point. Each iteration of our algorithm has two steps—a learning and an imputation step. In the learning step, the statistical model and the variational distributions are updated by maximising a variational lower-bound on the log-likelihood (Ⓐ and Ⓒ in the figure). In the imputation step, the missing values are imputed via pseudo-Gibbs sampling using the learnt variational distributions (Ⓑ in the figure). These imputations are *persistent* and updated in subsequent iterations. In this way, the initial imputations iteratively adapt to the target statistical model, such that they follow the joint distribution of the missing variables conditional on the observed data. Moreover, our method facilitates parameter sharing across missingness patterns by using an amortised inference model and is amenable to parallelisation, which further increases computational efficiency.

1. An interactive demo is available at github.com/vsimkus/variational-gibbs-inference.




Figure 1: A schematic illustration of our method on a 3D toy model $p_{\theta}(x_1, x_2, x_3)$, with missingness only in the first two dimensions for ease of exposition.

The top row along the arrow shows contours of the model $p_{\theta^{(t)}}$ with parameters $\theta^{(t)}$ learnt iteratively as it approaches the complete-data estimate p_* .

The zoomed-in pane details an iteration of the algorithm, where:

(a) univariate variational conditionals $q(x_j | \mathbf{x}_{\setminus j})$ are learnt from $p_{\theta}(\mathbf{x})$ and

(b) used to update the imputations via pseudo-Gibbs sampling;

(c) $p_{\theta}(\mathbf{x})$ is learnt using the imputed data via a variational objective.

Compared to standard VI, our approach reduces the number of variational distributions that need to be learnt from $2^d - 1$ to d .

1.2 Overview of the Paper

The paper is organised as follows. Section 2 provides background on statistical model estimation and standard variational inference for incomplete data, discusses the technical gap in the literature on amortised variational inference of missing data, and describes a classical approach for missing data imputation that is related to our work.

In Section 3, we derive the VGI optimisation objective (Section 3.1), present the VGI algorithm (Section 3.2), and discuss practical considerations for modelling the variational Gibbs conditionals (Section 3.3). We further introduce a block-Gibbs version of VGI (Section 3.4), which we later use in Section 5.3 to adapt our method specifically to the variational autoencoder (VAE, Kingma and Welling, 2013; Rezende et al., 2014). The details of evaluating models fitted with VGI for model selection using incomplete held-out data are presented next (Section 3.5). A discussion in Section 3.6 on the similarities and differences between VGI and related methods closes Section 3.

In Section 4, we validate and analyse our method on low- and high-dimensional toy problems where analytical solutions exist. In Section 5 we compare VGI against model-specific estimation methods on a VAE model and show that VGI produces competitive results in terms of model accuracy. In Section 6 we apply our general-purpose method to normalising flow estimation (Rezende and Mohamed, 2015) and show that it can outperform a flow-specific estimation method.

In Section 7 we summarise our findings and discuss possible future research directions.

2. Background

In this section we provide the background on statistical model estimation and variational inference with missing data, amortised variational inference, and a popular missing data imputation method based on conditional modelling. We highlight the shortcomings of those methodologies that the proposed method addresses.

2.1 Model Estimation and Variational Inference with Incomplete Data

Statistical model estimation from observed data is typically solved via maximum-likelihood estimation (MLE). Following Little and Rubin (2002) the marginal observed likelihood for incomplete data is defined via a joint model $p_{\theta}(\mathbf{x}_{\text{obs}}, \mathbf{m})$ of the observed variables of interest \mathbf{x}_{obs} and the binary missingness mask \mathbf{m} ,

$$p_{\theta}(\mathbf{x}_{\text{obs}}, \mathbf{m}) = \int p_{\theta}(\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}}) p(\mathbf{m} \mid \mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}}) d\mathbf{x}_{\text{mis}}, \quad (1)$$

where the complete variable of interest \mathbf{x} is defined by its observed and missing components $\mathbf{x} = (\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})$, $p_{\theta}(\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})$ is the statistical model of the data, and $p(\mathbf{m} \mid \mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})$ is the missingness model representing the missing data mechanism. In this paper we assume that data are missing at random (MAR) or missing completely at random (MCAR) so that the missingness model can be ignored when estimating the statistical model of the data $p_{\theta}(\mathbf{x})$ (Rubin, 1976). The maximum likelihood estimate of the parameters θ is then given

\mathbf{x}^1	x_1^1	?	x_3^1	x_4^1	\mathbf{m}^1	1	0	1	1	$q(\mathbf{x}_{\text{mis}}^i \mathbf{x}_{\text{obs}}^i)$
\mathbf{x}^2	?	x_2^2	x_3^2	?	\mathbf{m}^2	0	1	1	0	$q(x_1^2, x_4^2 x_2^2, x_3^2)$
\mathbf{x}^3	?	?	?	x_4^3	\mathbf{m}^3	0	0	0	1	$q(x_1^3, x_2^3, x_3^3 x_4^3)$
\vdots	\vdots			\vdots	\vdots				\vdots	

Figure 2: Left: Observed incomplete data. Middle: missingness mask. There are potentially as many as 2^d different missingness patterns. Right: corresponding variational posterior distributions.

by

$$\arg \max_{\boldsymbol{\theta}} \prod_{i=1}^N p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{obs}}^i, \mathbf{m}^i) = \arg \max_{\boldsymbol{\theta}} \prod_{i=1}^N \int p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{obs}}^i, \mathbf{x}_{\text{mis}}^i) d\mathbf{x}_{\text{mis}}^i, \quad (2)$$

where N is the number of data-points and i denotes the index of a data-point. A brief review of the different missingness mechanisms and a proof of the above is provided in Appendix A.

The integral over the missing components in (2) is typically intractable, which renders the likelihood and hence standard maximum-likelihood estimation intractable. Exactly the same problem occurs when estimating latent-variable models. Indeed, we can consider the missing components to be latent variables and hence obtain a tractable lower bound on the log-likelihood as done in variational inference (VI, Jordan et al., 1999). Following the standard derivation of the evidence lower-bound (ELBO) (e.g. Barber, 2017, Chapter 11.2), the bound on the log-likelihood for N incomplete data points is

$$\frac{1}{N} \sum_{i=1}^N \log p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{obs}}^i) \geq \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{q(\mathbf{x}_{\text{mis}}^i | \mathbf{x}_{\text{obs}}^i)} \left[\log \frac{p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{obs}}^i, \mathbf{x}_{\text{mis}}^i)}{q(\mathbf{x}_{\text{mis}}^i | \mathbf{x}_{\text{obs}}^i)} \right],$$

where $q(\mathbf{x}_{\text{mis}}^i | \mathbf{x}_{\text{obs}}^i)$ are the variational distributions. Maximising the ELBO with respect to the model parameters $\boldsymbol{\theta}$ and variational distributions $q \in \mathcal{Q}$ in some distributional family \mathcal{Q} yields an approximate MLE solution to (2). If the variational distributions are equal to the model conditionals $p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{mis}}^i | \mathbf{x}_{\text{obs}}^i)$ for all $\mathbf{x}_{\text{obs}}^i$ then the bound is tight and the solution can be made exact.²

Whilst it is straightforward to obtain a tractable lower bound on the log-likelihood, there is a crucial complication in the missing data problem that sets it apart from standard (amortised) VI problems: for d -dimensional data, there is not one but possibly $2^d - 1$ such variational conditional distributions, namely one for each non-trivial pattern of missingness,

2. Note that to be able to satisfy this condition the variational family should include the model conditionals $p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{mis}}^i | \mathbf{x}_{\text{obs}}^i) \in \mathcal{Q}$. Still, such specification may not guarantee the exactness of the MLE solution, since due to local optima in the ELBO the variational distributions may not perfectly match the model conditionals.

as illustrated in Figure 2. This raises fundamental technical issues that have prevented the application of VI to parameter estimation from incomplete data. We discuss these complications in the next subsection.

2.2 The Difficulty of Amortising Missing Variable Inference

Classical variational inference requires the specification of one variational distribution q per observed data-point, but such an approach is computationally inefficient in modern large-scale use-cases due to a lack of parameter sharing. An amortised version of VI (Gershman and Goodman, 2014) deals with this issue by incorporating global parameter sharing. The key idea in amortised VI is to parametrise the conditional variational distribution by a deterministic function, or an inference network, of the observed inputs with globally shared parameters ϕ . However, in the missing data setting we want to represent all $2^d - 1$ conditional distributions caused by the different missingness patterns. The exponential growth in the number of missingness patterns means that the naïve approach of using one inference network per pattern results in a lack of parameter sharing across data-points even for moderate-dimensional data, thus cancelling the computational advantages of amortised VI.

Efficiently amortising a variational distribution $q_\phi(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})$ for any $(\mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}}) \in \mathbf{x}$ entails simultaneously dealing with two problems: (i) handling all possible combinations of variables in the conditioning set, that is, for all $\mathbf{x}_{\text{obs}} \in \mathbf{x} \setminus \mathbf{x}_{\text{mis}}$, and (ii) constructing pdfs/pmfs for arbitrary sets of target variables $\mathbf{x}_{\text{mis}} \in \mathbf{x} \setminus \mathbf{x}_{\text{obs}}$. Existing work has focused on the first problem, approaching it by either simply fixing the input dimensionality of the inference network to d and then padding the missing inputs with zeroes (Nazábal et al., 2020; Mattei and Frellsen, 2019), or using permutation-invariant network architectures (Ma et al., 2019). However, there is no work in the VI literature that addresses the second problem.³ Dealing with this problem requires care—placing restrictions on the variational family may result in a biased estimate of the target statistical model (as shown below). Hence, to match the possibly complex conditional distributions of the target model we would like to use unrestricted probabilistic models for the variational family whilst still being able to take advantage of the increased efficiency of amortised VI.

We illustrate how restricting the variational family can reduce the quality of the fitted target model. For example, taking inspiration from mean-field VI (e.g. Bishop, 2006, Section 10.1.1) one may assume independence of the missing variables given the observed and work with variational distributions of the form

$$q_\phi(\mathbf{x}_{\text{mis}}^i | \mathbf{x}_{\text{obs}}^i) = \prod_{j \in \text{idx}(\mathbf{m}^i)} q_\phi(x_j^i | \mathbf{x}_{\text{obs}}^i),$$

where $\text{idx}(\mathbf{m}^i)$ denotes the set of indices of the missing values in the data-point. However, in Figure 3 we show that such independence assumption can significantly bias the learnt probabilistic model by artificially reducing correlations in the completed data, with the bias increasing with the fraction of missingness.

In this work, we introduce a novel variational method that can fit statistical models using only d variational conditionals, thus facilitating parameter sharing without introducing strong statistical assumptions or restricting their modelling capability.

3. But see Section 3.6 for related work.

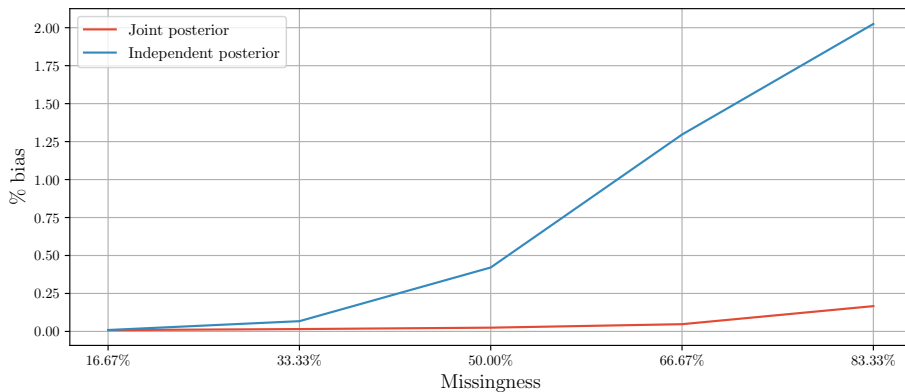


Figure 3: Percent bias measured on the test log-likelihood of a toy factor analysis model with correct (red) and incorrect (blue) modelling assumptions. Log-likelihood was computed on complete test data. We observe that with the incorrect independence assumption $x_j \perp\!\!\!\perp \mathbf{x}_{\text{mis}\setminus j} \mid \mathbf{x}_{\text{obs}}$ for $\forall j \in \text{idx}(\mathbf{m})$ the bias is significantly larger compared to the correctly-specified joint model. Percent bias is computed as $100 \cdot \left| (\mathbb{E} [\ell(p_\theta)] - \ell(p_*)) / \ell(p_*) \right|$ (e.g. van Buuren, 2018, Chapter 2.5), where $\ell(p_\theta)$ and $\ell(p_*)$ are the average log-likelihood on test data using the fitted p_θ and the ground truth p_* models, and the expectation is over different runs of the learning algorithm.

2.3 Imputing Missing Data with Fully-Conditional Models

For multivariate missing data, imputations can be generated by iteratively sampling univariate conditional distributions $f(x_j \mid \mathbf{x}_{\setminus j}; \phi_j)$ for $\forall j \in \text{idx}(\mathbf{m})$ with local parameters ϕ_j . Imputation methods in this family, which learn the conditionals and impute the data iteratively, are known as fully-conditional specification (FCS) or sequential iterative regression methods (Brand, 1999; Rubin, 2003; van Buuren et al., 2006). They have been essential tools in statistical analysis of incomplete data for over two decades.

The most popular of the FCS imputation methods is multivariate imputation by chained equations (MICE, van Buuren and Oudshoorn, 2000). MICE is an iterative framework that starts with a random imputation from the observed data and then sequentially imputes each incomplete variable one-by-one. Each step consists of fitting a conditional distribution $f(x_j \mid \mathbf{x}_{\setminus j}; \phi_j)$ on the data-points where x_j is observed via regression and then imputing the x_j in data-points where it is not observed by sampling the learnt conditional (see Appendix C for more details about MICE). To capture the uncertainty of the missing variables, the MICE procedure is usually independently repeated K number of times to obtain K imputations of the missing data, an approach known as multiple-imputation (MI, Rubin, 1987a). After imputation, statistical models $p_\theta(\mathbf{x})$ can be fit on the imputed data

sets using standard methods for complete data.⁴ We provide more background on such an “impute-then-fit” approach in Appendix D.

Whilst the sampling procedure bears some similarity to Gibbs sampling (Geman and Geman, 1984), see Appendix B for a short review, an important theoretical difference between Gibbs sampling and FCS is that in standard Gibbs sampling one starts with a joint model of the target distribution and then, by decomposing it into full-conditionals, samples a Markov chain that will eventually converge to the joint target distribution. On the other hand, the univariate distributions with disjoint parameters ϕ_j in FCS are not guaranteed to have a joint distribution (Arnold et al., 1999, Section 1.6), which is why this approach has been called pseudo-Gibbs sampling (Heckerman et al., 2000).⁵ Nevertheless, despite the possibly incompatible univariate distributions, it has been empirically found that the procedure can generate good imputations of multivariate missing values in many practical settings (Rubin, 2003; van Buuren et al., 2006), which motivates us to use a similar factorisation to represent the variational distributions of the missing variables.

FCS has several desirable properties: (i) it easily lends itself to the specification of flexible imputation models since even for univariate Gaussian conditionals the joint distribution can be complex and multi-modal (e.g. Arnold et al., 1999, Section 3.4), (ii) heterogeneous data (mixed continuous and discrete) can be handled by using different distribution families for each conditional, (iii) it alleviates the problem of having to learn $2^d - 1$ joint missing variable distributions into one that requires learning only d univariate Gibbs conditionals, and (iv) univariate conditional distributions can be easily constrained to prevent invalid imputation values.

Inspired by the empirical success of FCS methods on multiple-imputation tasks, in the next section we propose a variational approach that characterises the joint variational distributions $q_\phi(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}})$ via d variational full-conditional models. Thus, the joint variational distribution $q_\phi(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}})$ can be made very flexible since we neither require strong statistical independence assumptions nor restrictive distributional assumptions on the variational families (see Section 2.2), which enables us to well optimise the ELBO and hence obtain a good (approximate) MLE estimate of the statistical model $p_\theta(\mathbf{x})$. In this way, we generalise variational inference of univariate missing data to the multivariate case, akin to how FCS generalises univariate regression-based imputation to the multivariate setting (e.g. van Buuren, 2018, Chapter 4.5.1).

3. Variational Gibbs Inference

We present variational Gibbs inference (VGI) to estimate the parameters of statistical models from incomplete data by maximising a variational lower-bound on the log-likelihood. The method requires only d variational conditionals and uses an iterative algorithm that alternates between two steps: (i) the learning step which fits the statistical model and the variational conditionals by optimising the variational objective, and (ii) the imputation-step which updates persistent pseudo-Gibbs chains that provide the imputations for each incom-

4. However, the standard multiple-imputation workflow for parameter estimation is generally not applicable to statistical models when their parameters are not identifiable, as is most often the case in deep generative models, and hence caution must be taken (see Appendix D).

5. Sometimes also called incompatible Gibbs sampling or compound conditional specification (Rubin, 2003).

Term	Reference	Description
$p_{\theta}(\mathbf{x})$	Sec. 2.1	Statistical model of the data
$q_{\phi_j}(x_j \mathbf{x}_{\text{mis}\setminus j}, \mathbf{x}_{\text{obs}})$	Sec. 3.3	Variational conditionals
$\tau_{\phi}(\mathbf{x}_{\text{mis}} \mathbf{x}_{\text{obs}}, \tilde{\mathbf{x}}_{\text{mis}})$	Eq. (4)	Gibbs transition kernel
$\tilde{\tau}_{\phi}(\mathbf{x}_{\text{mis}} \mathbf{x}_{\text{obs}}, \tilde{\mathbf{x}}_{\text{mis}})$	Eq. (15)	Extended(-Gibbs) transition kernel (see Appendix E)
$f_{\phi}^t(\mathbf{x}_{\text{mis}} \mathbf{x}_{\text{obs}})$	Eq. (5)	Marginal distribution of a Markov chain at step t
$\mathcal{J}_{\text{VGI}}^t(\theta, \phi; \mathbf{x}_{\text{obs}})$	Eq. (9)	The VGI objective for a single data-point \mathbf{x}_{obs} at step t

Table 1: A glossary of key terms used in VGI.

plete data-point. The following sections introduce the variational objective (Section 3.1) and the VGI algorithm (Section 3.2), discuss practicalities to consider when modelling the variational conditionals (Section 3.3), explain how VGI can be adapted to specific statistical models, such as latent-variable models (Section 3.4), describe the details of evaluating the method on incomplete held-out data (Section 3.5), and finally discusses the related work (Section 3.6).

3.1 The Variational Objective

We derive a variational ELBO to estimate the statistical model $p_{\theta}(\mathbf{x})$ from incomplete data using the marginal distributions $f_{\phi}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})$ of Markov chains with learnable parameters ϕ as the variational distribution of the missing variables. Maximising the objective allows us to learn the parameters θ of the model $p_{\theta}(\mathbf{x})$ and the parameters ϕ of the Gibbs transition *kernel* $\tau_{\phi}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, \tilde{\mathbf{x}}_{\text{mis}})$,⁶ or equivalently, the parameters of the d univariate Gibbs conditionals (*variational conditionals*) $q_{\phi_j}(x_j | \tilde{\mathbf{x}}_{\text{mis}\setminus j}, \mathbf{x}_{\text{obs}})$ which characterise the kernel. By learning the kernel τ_{ϕ} we can match the marginal distributions of the imputation Markov chains f_{ϕ}^t to the conditional distributions of the target model $p_{\theta}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})$. Consequently, the variational conditionals $q_{\phi_j}(x_j | \mathbf{x}_{\text{mis}\setminus j}, \mathbf{x}_{\text{obs}})$ are learnt to approximate the true conditionals $p_{\theta}(x_j | \mathbf{x}_{\text{mis}\setminus j}, \mathbf{x}_{\text{obs}})$.

Representing the variational distribution implicitly via the samples of a Markov chain allows the VGI method to work with a flexible variational model for adaptive imputations, which in turn enables tightening of the variational lower-bound and thus the maximisation of the likelihood. Moreover, by representing the kernel with d variational conditionals, VGI achieves efficient amortisation of the exponentially-many distributions of missing data, similar to fully-conditional imputation methods discussed in Section 2.3. A summary of the key terms in VGI is provided in Table 1.

We start our derivation with the standard variational ELBO, but with the marginal distribution f_{ϕ}^t of a Markov chain as the variational distribution

$$\log p_{\theta}(\mathbf{x}_{\text{obs}}) \geq \mathcal{L}(\theta, f_{\phi}^t; \mathbf{x}_{\text{obs}}), \quad \mathcal{L}(\theta, f_{\phi}^t; \mathbf{x}_{\text{obs}}) = \mathbb{E}_{f_{\phi}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})} \left[\log \frac{p_{\theta}(\mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}})}{f_{\phi}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})} \right], \quad (3)$$

where t denotes the iteration of our algorithm. To maximise the likelihood of the model p_{θ} we want the bound to be sufficiently tight, and hence f_{ϕ}^t needs be flexible and not biased

6. We use $\tilde{\mathbf{x}}_{\text{mis}}$ to denote the previous imputation value before the transition.

by the choice of the initial imputation distribution f^0 . The flexibility of the marginal distribution of a Markov chain depends on the kernel, and the bias induced by the choice of f^0 decreases with the number of steps t in the Markov chain (e.g. Cover and Thomas, 2006, Chapter 4.4). Hence we want τ_ϕ to be flexible and t to become sufficiently large.

However, we cannot evaluate the above lower-bound since the imputation distribution $f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}})$ is implicitly-defined by the Markov chain. Moreover, maximising the bound with respect to the parameters θ and ϕ with gradient-based methods would be expensive for large t and might suffer from gradient instability. Rather than estimating the gradients of the above lower-bound with respect to ϕ over the full length of the Markov chain, we propose “cutting” the chain just before the current transition and optimising the transition kernel τ_ϕ greedily as we sample. The imputations obtained in iteration t will then be re-used in the next iteration to further improve τ_ϕ and p_θ . This allows us to derive a tractable and efficient variational method.

We choose the kernel, which transforms the imputations from the previous iteration, $\tilde{\mathbf{x}}_{\text{mis}}$, and produces updated imputations \mathbf{x}_{mis} , to be Gibbs (hence the name of the method) and define it as follows

$$\tau_\phi(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, \tilde{\mathbf{x}}_{\text{mis}}) = \sum_{j \in \text{idx}(\mathbf{m})} \pi(j) q_{\phi_j}(x_j \mid \mathbf{x}_{\text{mis} \setminus j}, \mathbf{x}_{\text{obs}}) \delta(\mathbf{x}_{\text{mis} \setminus j} - \tilde{\mathbf{x}}_{\text{mis} \setminus j}), \quad (4)$$

where the dimension index $j \in \text{idx}(\mathbf{m})$ is controlled by a fixed uniform selection distribution $\pi(j)$ of a random-scan Gibbs sampler, x_j and $\mathbf{x}_{\text{mis} \setminus j}$ denote the j -th missing dimension and the remaining missing dimensions of \mathbf{x}_{mis} respectively, and the kernel is specified by the d variational conditionals $q_{\phi_j}(x_j \mid \mathbf{x}_{\text{mis} \setminus j}, \mathbf{x}_{\text{obs}})$. Alternatively, we can also let the updated imputation x_j depend on the previous imputation value \tilde{x}_j , which gives variational conditionals of the form $q_{\phi_j}(x_j^t \mid \mathbf{x}_{\text{obs}}, \tilde{x}_j^{t-1}, \tilde{\mathbf{x}}_{\text{mis} \setminus j}^{t-1})$. We call a kernel $\tilde{\tau}_\phi$ (or subsequently the variational model) that uses this form of conditionals an *extended-Gibbs* kernel (or, for conciseness, *extended* kernel) and provide an analogous derivation of this section in Appendix E.

Marginalising out $\tilde{\mathbf{x}}_{\text{mis}}$ in (4) with respect to the *imputation distribution* f^{t-1} from the previous step in the Markov chain gives the imputation distribution after a single Gibbs update

$$f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}) = \int \tau_\phi(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, \tilde{\mathbf{x}}_{\text{mis}}) f^{t-1}(\tilde{\mathbf{x}}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}) d\tilde{\mathbf{x}}_{\text{mis}} \quad (5)$$

$$\begin{aligned} &= \sum_{j \in \text{idx}(\mathbf{m})} \pi(j) f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, j) \\ &= \mathbb{E}_{\pi(j)} \left[f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, j) \right], \end{aligned} \quad (6)$$

where $f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, j)$ is the imputation distribution f^{t-1} updated in dimension j ,

$$f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, j) = q_{\phi_j}(x_j \mid \mathbf{x}_{\text{mis} \setminus j}, \mathbf{x}_{\text{obs}}) f^{t-1}(\mathbf{x}_{\text{mis} \setminus j} \mid \mathbf{x}_{\text{obs}}). \quad (7)$$

Note that the absence of ϕ in f^{t-1} corresponds to the aforementioned “cutting” of the chain.

Now, we continue with the standard ELBO from (3) and use (6) and (7) to derive the variational Gibbs ELBO at iteration t

$$\begin{aligned}
 \mathcal{L}^t(\boldsymbol{\theta}, \boldsymbol{\phi}; \mathbf{x}_{\text{obs}}) &\stackrel{(6)}{=} \mathbb{E}_{\boldsymbol{\pi}^{(j)} f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, j)} \left[\log \frac{p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})}{f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})} \right] \\
 &= \mathbb{E}_{\boldsymbol{\pi}^{(j)} f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, j)} \left[\log \frac{p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})}{f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, j)} + \log \frac{f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, j)}{f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})} \right] \\
 &= \mathbb{E}_{\boldsymbol{\pi}^{(j)} f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, j)} \left[\log \frac{p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})}{f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, j)} \right] \\
 &\quad + \mathbb{E}_{\boldsymbol{\pi}^{(j)}} D_{\text{KL}}(f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, j) \parallel f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})) \\
 &\geq \mathbb{E}_{\boldsymbol{\pi}^{(j)} f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, j)} \left[\log \frac{p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})}{f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, j)} \right] \\
 &\stackrel{(7)}{=} \mathbb{E}_{\boldsymbol{\pi}^{(j)} f^{t-1}(\mathbf{x}_{\text{mis} \setminus j} | \mathbf{x}_{\text{obs}}) q_{\boldsymbol{\phi}_j}(x_j | \mathbf{x}_{\text{mis} \setminus j}, \mathbf{x}_{\text{obs}})} \left[\log \frac{p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})}{q_{\boldsymbol{\phi}_j}(x_j | \mathbf{x}_{\text{mis} \setminus j}, \mathbf{x}_{\text{obs}})} \right] \\
 &\quad - \mathbb{E}_{\boldsymbol{\pi}^{(j)} f^{t-1}(\mathbf{x}_{\text{mis} \setminus j} | \mathbf{x}_{\text{obs}})} \left[\log f^{t-1}(\mathbf{x}_{\text{mis} \setminus j} | \mathbf{x}_{\text{obs}}) \right], \tag{8}
 \end{aligned}$$

where the inequality follows from the non-negativity of KL divergence. If \mathbf{x}_{mis} is independent of j , which holds for the stationary distribution of the Markov chain characterised by the variational conditionals (see Appendix F), then the KL divergence term is zero and maximising (8) is equivalent to maximising (3).

The ELBO in (8) can be optimised efficiently with respect to $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ since only samples from the imputation distribution $f^{t-1}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})$ are needed and the intractable entropy term in the last line of the equation does not depend on the parameters $\boldsymbol{\theta}$ or $\boldsymbol{\phi}$, and hence does not need to be computed. Removing the entropy term, we obtain the variational Gibbs inference (VGI) objective $\mathcal{J}_{\text{VGI}}^t$ at iteration t for one incomplete data-point \mathbf{x}_{obs} :⁷

$$\mathcal{J}_{\text{VGI}}^t(\boldsymbol{\theta}, \boldsymbol{\phi}; \mathbf{x}_{\text{obs}}) = \mathbb{E}_{\boldsymbol{\pi}^{(j)} f^{t-1}(\mathbf{x}_{\text{mis} \setminus j} | \mathbf{x}_{\text{obs}}) q_{\boldsymbol{\phi}_j}(x_j | \mathbf{x}_{\text{mis} \setminus j}, \mathbf{x}_{\text{obs}})} \left[\log \frac{p_{\boldsymbol{\theta}}(x_j, \mathbf{x}_{\text{mis} \setminus j}, \mathbf{x}_{\text{obs}})}{q_{\boldsymbol{\phi}_j}(x_j | \mathbf{x}_{\text{mis} \setminus j}, \mathbf{x}_{\text{obs}})} \right] \tag{9}$$

Importantly, while cutting the chain removes the entropy of f^{t-1} from the objective function, the entropy of the variational conditionals remains. This prevents the collapse of the variational conditionals to point masses and consequently also prevents the collapse of the joint imputation distribution f^t that is sampled using the fitted conditionals.⁸

For N incomplete data-points we will maximise the averaged objective to obtain the parameter estimates $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\phi}}$ at iteration t

$$\hat{\boldsymbol{\theta}}^t, \hat{\boldsymbol{\phi}}^t = \arg \max_{\boldsymbol{\theta}, \boldsymbol{\phi}} \frac{1}{N} \sum_{i=1}^N \mathcal{J}_{\text{VGI}}^t(\boldsymbol{\theta}, \boldsymbol{\phi}; \mathbf{x}_{\text{obs}}^i). \tag{10}$$

7. The objective using the extended kernel (Appendix E) is almost the same, but the variational conditional additionally includes a dependency on \tilde{x}_j and hence depends on all of the imputed variables from the previous iteration ($\tilde{x}_j, \mathbf{x}_{\text{mis} \setminus j} \sim f^{t-1}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})$).

8. Note that the imputation distribution $f^{t-1}(\mathbf{x}_{\text{mis} \setminus j} | \mathbf{x}_{\text{obs}})$ is kept fixed when we maximise the objective $\mathcal{J}_{\text{VGI}}^t(\boldsymbol{\theta}, \boldsymbol{\phi}; \mathbf{x}_{\text{obs}})$ with respect to $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$, due to the ‘‘cutting’’ of the Markov chain. After updating the parameters, we update $f^{t-1}(\mathbf{x}_{\text{mis} \setminus j} | \mathbf{x}_{\text{obs}})$ via the updated transition kernel according to (5).

In practice, we optimise (10) using stochastic gradient ascent and approximate the expectations in $\mathcal{J}_{\text{VGI}}^t$ with Monte Carlo integration

$$\hat{\mathcal{J}}_{\text{VGI}}^t(\boldsymbol{\theta}, \boldsymbol{\phi}; \mathbf{x}_{\text{obs}}) = \frac{1}{K} \sum_{k=1}^K \frac{1}{M} \sum_{m=1}^M \left[\log \frac{p_{\boldsymbol{\theta}}(\mathbf{x}_{j^m}^k, \mathbf{x}_{\text{mis} \setminus j^m}^k, \mathbf{x}_{\text{obs}})}{q_{\boldsymbol{\phi}_j}(\mathbf{x}_{j^m}^k | \mathbf{x}_{\text{mis} \setminus j^m}^k, \mathbf{x}_{\text{obs}})} \right], \quad (11)$$

where $\mathbf{x}_{j^m}^k \sim q_{\boldsymbol{\phi}_j}(\mathbf{x}_{j^m}^k | \mathbf{x}_{\text{mis} \setminus j^m}^k, \mathbf{x}_{\text{obs}})$, $j^m \sim \boldsymbol{\pi}(j)$, $\mathbf{x}_{\text{mis} \setminus j^m}^k \sim f^{t-1}(\mathbf{x}_{\text{mis} \setminus j^m} | \mathbf{x}_{\text{obs}})$, and K is the number of imputations for each incomplete data-point, M is the number of samples used to approximate the expectation with respect to j (and x_j), and f^{t-1} is represented via samples from a fixed number of Markov chains. We empirically found that using small M and K was sufficient in most of our experiments.⁹

To summarise, we have derived the key VGI objective $\mathcal{J}_{\text{VGI}}^t$, which maximises a lower-bound on the log-likelihood using samples from a marginal distribution of a variational Markov chain at any iteration t . Importantly, the objective uses only d variational conditionals to represent all $2^d - 1$ possible conditional distributions of missing data. We next describe the VGI algorithm which integrates the optimisation of the iteration-dependent variational objective $\mathcal{J}_{\text{VGI}}^t$ and sampling from the Markov chains $f^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})$ into an iterative procedure.

3.2 The VGI Algorithm

The variational Gibbs inference (VGI) algorithm is shown in Algorithm 1 with subroutines summarised in Algorithms 2-5.¹⁰ The core objective of the algorithm is to fit a statistical model $p_{\boldsymbol{\theta}}$ on incomplete data set \mathcal{D} using a variational model $q_{\boldsymbol{\phi}}$ of the Gibbs conditionals that is learnt jointly with $p_{\boldsymbol{\theta}}$. The method uses stochastic gradient optimisation (e.g. Spall, 2003; Ruder, 2017) to learn the parameters $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ efficiently by processing the data set in mini-batches, which are randomly chosen subsets of the full data set \mathcal{D} . Algorithm 1 consists of three stages: initialisation, model warm-up, and the main iterative stage, which we describe below.

Initialisation. In line 1 the algorithm starts with an incomplete data set $\mathcal{D} = \{(\mathbf{x}_{\text{obs}}^i, \mathbf{m}^i)\}_i$ and produces a K -times imputed data set $\mathcal{D}_K = \{(\mathbf{x}_{\text{obs}}^i, \mathbf{m}^i, \mathbf{x}_{\text{mis}}^{(i,1)}, \dots, \mathbf{x}_{\text{mis}}^{(i,K)})\}_i$, where each incomplete data-point is imputed using an initial imputation distribution $\mathbf{x}_{\text{mis}}^{(i,k)} \sim f^0(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}^i)$. The f^0 should be chosen such that it provides a good starting distribution for the Gibbs sampler. Empirically we found that choosing f^0 to be the marginal empirical distribution worked well, which corresponds to the maximum-entropy distribution of the missing data.

Warm-up. The warm-up stage (lines 2-3) has two parts: a variational model $q_{\boldsymbol{\phi}}$ initialisation stage (line 2) and a statistical model $p_{\boldsymbol{\theta}}$ warm-up stage (line 3), which we describe below in the given order. We note that the warm-up stage can be optional, however we empirically found that it allows the models to take reasonable initial values, which can sometimes improve the model fit or stabilise the learning for a small additional initialisation cost.

9. We found $K \in [5, 10]$ and $M \in [1, 10]$ sufficient in our experiments.
 10. Note that in our experiments we use optimised implementation of the algorithms by parallelising the loops where possible. Our implementation is available at <https://github.com/vsimkus/variational-gibbs-inference>.

Algorithm 1 Variational Gibbs inference (VGI) algorithm

Input: $p_{\theta}(\mathbf{x})$, statistical model with parameters θ
 $q_{\phi_j}(x_j | \mathbf{x}_{\setminus j})$ for $j \in \{1 \dots d\}$, variational conditional models with parameters ϕ
 \mathcal{D} , incomplete data set
 K , number of imputations of each incomplete data-point
 $f^0(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})$, initial imputation distribution
 α_{θ} and α_{ϕ} , the parameter learning rates
 max_epochs , number of epochs

Output: θ, ϕ , and K -times imputed data \mathcal{D}_K

- 1: **Create** K -times imputed data \mathcal{D}_K using f^0
 - 2: **Warm up** the parameters ϕ of the variational model q_{ϕ} using Algorithm 2
 - 3: **Warm up** the parameters θ of the statistical model p_{θ} using Algorithm 4
 - 4: **for** t in $[1, max_epochs]$ **do**
 - 5: **for** mini-batch \mathcal{B}_K in \mathcal{D}_K **do**
 - 6: **Update** imputations in \mathcal{B}_K using pseudo-Gibbs sampler (see Algorithm 5)
 - 7: **Store** the updated imputations in \mathcal{B}_K for use in the next epoch
 - 8: **Compute** $\hat{\mathcal{J}}_{\text{VGI}}^t$ in (11) using Algorithm 3
 - 9: $\theta^t = \theta^{t-1} + \alpha_{\theta} \nabla_{\theta} \hat{\mathcal{J}}_{\text{VGI}}^t$ ▷ Update params of p_{θ} with a stochastic gradient step
 - 10: $\phi^t = \phi^{t-1} + \alpha_{\phi} \nabla_{\phi} \hat{\mathcal{J}}_{\text{VGI}}^t$ ▷ Update params of q_{ϕ} with a stochastic gradient step
 - 11: **end for**
 - 12: **end for**
-

The warm-up procedure for the variational model is outlined in pseudo-code in Algorithm 2. In standard variational inference for latent variable models, the variational model is often initialised randomly. However, in the missing-data case, starting-out with randomly-initialised variational conditionals may be sub-optimal since there is usually some observed data that could be used to reasonably initialise the variational model. Therefore, to make use of the available observed data, we suggest pre-training the variational inference networks on the observed values using maximum-likelihood estimation and stochastic gradient ascent (SGA) via

$$\hat{\phi} = \arg \max_{\phi} \mathcal{J}_{\phi}(\mathcal{D}_K), \tag{12}$$

$$\mathcal{J}_{\phi}(\mathcal{D}_K) = \sum_{j=1}^d \frac{1}{|\text{obs}(\mathcal{D}_K, j)|} \sum_{i \in \text{obs}(\mathcal{D}_K, j)} \frac{1}{K} \sum_{k=1}^K \log q_{\phi_j}(x_j^i | \mathbf{x}_{\text{mis}}^{(i,k)}, \mathbf{x}_{\text{obs}\setminus j}^i),$$

where $\text{obs}(\mathcal{D}_K, j)$ represents the set of indices of data-points that are observed in the j -th dimension. This pre-training procedure is probabilistic regression for all dimensions $j \in [1, d]$ using the data where the j -th dimension is observed, with the missing values imputed with a baseline imputation method (for example, samples from the empirical marginals).¹¹ We find that a small number of pre-training iterations is often sufficient to favourably

11. Care must be taken when warming-up the extended-Gibbs kernel to avoid poor initialisation of the conditionals, see technical aside in Appendix E.

initialise the variational model, which can boost the performance of the method and stabilise training in the initial iterations of the main stage of the algorithm.

Next, the algorithm performs SGA on the parameters θ of the statistical model p_θ using the variational Gibbs objective $\hat{\mathcal{J}}_{\text{VGI}}$ in (11) (see Algorithms 3 and 4). This kind of “warm-up” allows the parameters θ to take on reasonable initial values before the main iterative stage starts. We qualitatively find that the warm-up stage generally needs to be performed for a small number of iterations until the change in $\hat{\mathcal{J}}_{\text{VGI}}$ between consecutive iterations falls below a threshold. Note that we keep the variational parameters ϕ fixed at this stage so that the variational model does not deteriorate while the model p_θ is being initialised. Also, in this stage the imputed data in \mathcal{D}_K are not yet updated, they still follow the initial imputation distribution f^0 . This is because we empirically found that early Gibbs sampling can cause divergent imputation chains, training instability, or getting stuck in a local optima.

Main stage. The main stage of the algorithm iterates between updating the imputations in \mathcal{D}_K and fitting the parameters θ and ϕ (see lines 6-10). For each mini-batch \mathcal{B}_K we first update the imputed values $\mathbf{x}_{\text{mis}}^{(i,1)}, \dots, \mathbf{x}_{\text{mis}}^{(i,K)}$ using G Gibbs updates with the variational conditionals (see Algorithm 5) and then update the parameters θ and ϕ using a single stochastic gradient of the variational Gibbs objective $\hat{\mathcal{J}}_{\text{VGI}}^t$ in (11), see Algorithm 3 on computing the objective. These two steps (imputation and parameter update) are then repeated until convergence or until the computational budget is exhausted.

We find that using a small number of Gibbs updates G is generally sufficient and preferable in terms of the trade-off between compute cost and convergence rate.¹² Moreover, a large value of G may cause training instability due to the generalisation gap of the variational conditionals (see Section 3.5 for more details). Updating the imputations via Gibbs sampling using the variational conditionals and storing them for the next iteration corresponds to updating the marginal Markov chain distribution $f^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}})$. Reusing the imputations from the previous iteration, rather than re-sampling from scratch at each iteration, is akin to using persistent chains that have previously been used in a different context in Markov chain Monte Carlo methods (e.g. persistent contrastive divergence, Yunes, 1999; Tieleman, 2008).

3.3 Choosing the Variational Model

The performance of amortised variational methods, and hence also VGI, critically depends on the choice of the functional family and the expressivity of the inference networks used to parametrise them. To maximise the compatibility of the statistical model p_θ and the variational model q_ϕ , we suggest the variational family should be chosen as one that includes or is close to the family of the statistical model. If the target model p_θ is a deep model, the variational inference networks should use architectural blocks in the neural network that are similar to the ones used in the target model.

We must also consider how to specify and parametrise the d variational conditionals. A straightforward way to specify them is to use d inference networks with parameters ϕ_j , one for each variational conditional. However, such an approach would be parameter-inefficient

12. We used $G \in [1, 5]$ in our experiments.

Algorithm 2 Variational model warm-up

Input: $q_{\phi_j}(x_j | \mathbf{x}_{\setminus j})$ for $j \in \{1 \dots d\}$, variational conditionals; \mathcal{D}_K , K -times imputed data; α_ϕ , the parameter learning rate; *var_warmup_epochs*, number of epochs.

Output: ϕ initialised on observed data using (12)

```

1: for  $t_w$  in  $[1, \text{var\_warmup\_epochs}]$  do
2:   for mini-batch  $\mathcal{B}_K$  in  $\mathcal{D}_K$  do
3:      $\mathcal{J}_\phi = 0$ 
4:     for  $j$  in  $[1, d]$  do
5:        $\mathcal{J}_\phi += \frac{1}{|\text{obs}(\mathcal{B}_K, j)|} \mathcal{J}_\phi^j(\mathcal{B}_K, j)$ 
6:     end for
7:      $\phi \leftarrow \phi + \alpha_\phi \nabla_\phi \mathcal{J}_\phi$ 
8:   end for
9: end for
10: def  $\mathcal{J}_\phi^j(\mathcal{B}_K, j)$ :
11:    $\mathcal{J}_\phi^j = 0$ 
12:   for  $\mathbf{x}_{\text{obs}}^i, \mathbf{m}^i, \mathbf{x}_{\text{mis}}^{(i,1)}, \dots, \mathbf{x}_{\text{mis}}^{(i,K)}$  in  $\mathcal{B}_K$  do
13:     if  $x_j^i$  is observed then
14:        $\mathcal{J}_\phi^j += \frac{1}{K} \sum_{k=1}^K \log q_{\phi_j}(x_j^i | \mathbf{x}_{\text{mis}}^{(i,k)}, \mathbf{x}_{\text{obs}\setminus j}^i)$ 
15:     end if
16:   end for
17: return  $\mathcal{J}_\phi^j$ 

```

Algorithm 3 Compute VGI objective in (11)

Input: $p_\theta(\mathbf{x})$, statistical model; $q_{\phi_j}(x_j | \mathbf{x}_{\setminus j})$ for $j \in \{1 \dots d\}$, variational conditionals; \mathcal{B}_K , K -times imputed mini-batch; M , number of missing dimensions to sample.

Output: $\hat{\mathcal{J}}_{\text{VGI}}(\theta, \phi; \mathbf{x}_{\text{obs}})$ averaged over all $\mathbf{x}_{\text{obs}}^i$ in mini-batch \mathcal{B}_K

```

1:  $\hat{\mathcal{J}}_{\text{VGI}} = 0$ 
2: for  $\mathbf{x}_{\text{obs}}^i, \mathbf{m}^i, \mathbf{x}_{\text{mis}}^{(i,1)}, \dots, \mathbf{x}_{\text{mis}}^{(i,K)}$  in  $\mathcal{B}_K$  do
3:   for  $k$  in  $[1, K]$  do
4:      $\hat{\mathcal{J}}_{\text{VGI}} += \frac{1}{|\mathcal{B}_K|K} \hat{\mathcal{J}}_{\text{VGI}}^{(i,k)}(\mathbf{x}_{\text{obs}}^i, \mathbf{x}_{\text{mis}}^{(i,k)})$ 
5:   end for
6: end for
7: def  $\hat{\mathcal{J}}_{\text{VGI}}^{(i,k)}(\mathbf{x}_{\text{obs}}^i, \mathbf{x}_{\text{mis}}^{(i,k)})$ :
8:    $\hat{\mathcal{J}}_{\text{VGI}}^{(i,k)} = 0$ 
9:   for  $m$  in  $[1, M]$  do
10:     $j^m \sim \pi(j) \triangleright$  Sample from  $\text{idx}(\mathbf{m}^i)$ 
11:     $\mathbf{x}_{j^m}^k \sim q_{\phi_j}(x_{j^m}^k | \mathbf{x}_{\text{mis}\setminus j^m}^{(i,k)}, \mathbf{x}_{\text{obs}}^i)$ 
12:     $\hat{\mathcal{J}}_{\text{VGI}}^{(i,k)} += \log \frac{p_\theta(\mathbf{x}_{j^m}^k, \mathbf{x}_{\text{mis}\setminus j^m}^{(i,k)}, \mathbf{x}_{\text{obs}}^i)}{q_{\phi_j}(x_{j^m}^k | \mathbf{x}_{\text{mis}\setminus j^m}^{(i,k)}, \mathbf{x}_{\text{obs}}^i)}$ 
13:   end for
14: return  $\frac{1}{M} \hat{\mathcal{J}}_{\text{VGI}}^{(i,k)}$ 

```

Algorithm 4 Statistical model warm-up

Input: $p_\theta(\mathbf{x})$, statistical model
 $q_{\phi_j}(x_j | \mathbf{x}_{\setminus j})$ for $j \in \{1 \dots d\}$
 \mathcal{D}_K , K -times imputed data
 α_θ , the parameter learning rate
model_warmup_epochs, # of epochs

Output: θ initialised for the main loop

```

1: for  $t_w$  in  $[1, \text{model\_warmup\_epochs}]$  do
2:   for mini-batch  $\mathcal{B}_K$  in  $\mathcal{D}_K$  do
3:     Estimate  $\hat{\mathcal{J}}_{\text{VGI}}$  using Algorithm 3
4:      $\theta \leftarrow \theta + \alpha_\theta \nabla_\theta \hat{\mathcal{J}}_{\text{VGI}}$ 
5:   end for
6: end for

```

Algorithm 5 (Pseudo-)Gibbs sampling

Input: $q_{\phi_j}(x_j | \mathbf{x}_{\setminus j})$ for $j \in \{1 \dots d\}$
 \mathcal{B}_K , K -times imputed mini-batch
 G , number of Gibbs update steps

Output: \mathcal{B}_K with updated imputations

```

1: for  $\mathbf{x}_{\text{obs}}^i, \mathbf{m}^i, \mathbf{x}_{\text{mis}}^{(i,1)}, \dots, \mathbf{x}_{\text{mis}}^{(i,K)}$  in  $\mathcal{B}_K$  do
2:   for  $k$  in  $[1, K]$  do
3:     for  $g$  in  $[1, G]$  do
4:        $j \sim \pi(j) \triangleright$  Sample from  $\text{idx}(\mathbf{m}^i)$ 
5:        $x_j \sim q_{\phi_j}(x_j | \mathbf{x}_{\text{mis}\setminus j}^{(i,k)}, \mathbf{x}_{\text{obs}}^i)$ 
6:        $\mathbf{x}_{\text{mis}}^{(i,k)} \leftarrow (\mathbf{x}_{\text{mis}\setminus j}^{(i,k)}, x_j)$ 
7:     end for
8:   end for
9: end for

```

and scale poorly to higher dimensional data. To address this, we suggest using the extended-Gibbs conditionals $q_{\phi_j}(x_j^t | \mathbf{x}_{\text{obs}}, \tilde{x}_j^{t-1}, \tilde{\mathbf{x}}_{\text{mis}\setminus j}^{t-1})$ (Appendix E), which allow the imputation x_j^t to depend on the previous imputation value \tilde{x}_j^{t-1} . Then, we can use a single partially-shared neural network where the parameters and the computations are shared in the first part of the network for all conditional distributions. These two simple modifications allows us to scale VGI to higher-dimensional data (see Appendix G for a more detailed discussion). We investigate the effects of parameter-sharing and extended conditionals in Section 4.7.

Finally, an important caveat of our approach is that, due to approximation errors, the fitted variational conditionals may not correspond to a joint distribution, which mirrors a similar caveat of the FCS methods outlined in Section 2.3. Hence, pseudo-Gibbs sampling using the fitted conditionals may diverge (we revisit this in Section 3.5 when discussing model selection).¹³ Nevertheless, we show in the experimental sections that the VGI method works well despite the lack of such convergence guarantees, similar to the existing FCS imputations methods (see Section 2.3).

3.4 Variational Block-Gibbs Inference and Latent-Variable Models

Thus far we have considered the general case where the number of possible missingness patterns is 2^d , with d being the dimensionality of the data. In some cases, missingness in a block of dimensions may be coupled such that either all of the dimensions in the block are missing or all are observed, which reduces the number of possible missingness patterns to 2^s , where s is the number of such missing variable blocks, with $s = d$ if the missingness in no dimension is coupled.

We can adapt the VGI lower-bound in (8) and equivalently the VGI objective in (9) to this scenario by letting j denote the index of a block of missing dimensions, rather than the index of a single missing dimension. Hence, the univariate x_j in (8) and (9) become potentially multivariate random variables \mathbf{x}_j . We must then specify a variational model for each block of missing dimensions, where we may choose to specify independent variational models or partially-shared models as discussed in Section 3.3. The Gibbs update step in line 6 of Algorithm 1 then corresponds to the update of a block-Gibbs sampler. Hence, we refer to this method as the variational block-Gibbs inference (VBGI).

A particularly important instance of coupled missingness is the latent-variable model $p_{\theta}(\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}} | \mathbf{x}_z)p_{\theta}(\mathbf{x}_z)$, where all of the latent-variables \mathbf{x}_z are missing together. If the missingness of no other dimension is coupled, then $j \in \text{idx}(\mathbf{m}) \cup \{z\}$ refers to either any one of the missing dimensions or all of the latents. Rewriting $\mathcal{J}_{\text{VGI}}^t$ for this model we get the following objective:¹⁴

$$\mathcal{J}_{\text{VBGI}}^t(\theta, \phi; \mathbf{x}_{\text{obs}}) = \mathbb{E}_{f^{t-1}(\tilde{\mathbf{x}}_{\text{mis}}, \tilde{\mathbf{x}}_z | \mathbf{x}_{\text{obs}})} \left(\sum_{j \in \text{idx}(\mathbf{m})} \pi(j) \mathbb{E}_{q_{\phi_j}(x_j | \tilde{\mathbf{x}}_{\text{mis}\setminus j}, \tilde{\mathbf{x}}_z, \mathbf{x}_{\text{obs}})} \left[\log \frac{p_{\theta}(x_j, \tilde{\mathbf{x}}_{\text{mis}\setminus j}, \mathbf{x}_{\text{obs}} | \tilde{\mathbf{x}}_z) p_{\theta}(\tilde{\mathbf{x}}_z)}{q_{\phi_j}(x_j | \tilde{\mathbf{x}}_{\text{mis}\setminus j}, \tilde{\mathbf{x}}_z, \mathbf{x}_{\text{obs}})} \right] \right)$$

13. Note that, after learning, we advocate using the estimated p_{θ} for generating imputations if they were required and not the variational conditionals.

14. Similar to the standard VGI case, the VBGI objective can optionally be adapted to use the extended-Gibbs conditionals (see Appendix E).

$$+ \pi(j = z) \mathbb{E}_{q_{\phi_z}(\mathbf{x}_z | \tilde{\mathbf{x}}_{\text{mis}}, \mathbf{x}_{\text{obs}})} \left[\log \frac{p_{\theta}(\tilde{\mathbf{x}}_{\text{mis}}, \mathbf{x}_{\text{obs}} | \mathbf{x}_z) p_{\theta}(\mathbf{x}_z)}{q_{\phi_z}(\mathbf{x}_z | \tilde{\mathbf{x}}_{\text{mis}}, \mathbf{x}_{\text{obs}})} \right], \quad (13)$$

where $q_{\phi_z}(\mathbf{x}_z | \tilde{\mathbf{x}}_{\text{mis}}, \mathbf{x}_{\text{obs}})$ is the variational conditional for the latents \mathbf{x}_z . The grouping of missing dimensions allows us to adapt VGI to latent-variable models. It forms the basis for the variational autoencoder-specific version of VGI that we introduce in Section 5.

3.5 Evaluating VGI on Incomplete Held-Out Data

A common step in a machine learning workflow is the validation of the estimated model on a held-out data set, which is used in early stopping or model selection (e.g. Goodfellow et al., 2016). In the incomplete data case, we can expect that the held-out data are also incomplete. One approach is to estimate the marginal log-likelihood $\log p_{\theta}(\mathbf{x}_{\text{obs}})$ using importance sampling or sequential Monte Carlo methods (e.g. Barber, 2017), however these approaches can be too computationally expensive to obtain reliable estimates of the marginal on the held-out data, especially if we wanted to evaluate many models. Alternatively, a validation loss, analogue to the training loss, is commonly used to select a model that minimises it. The evaluation of the VGI objective in (9) on incomplete held-out data requires the iterative evaluation of the variational conditionals (Gibbs sampling) on unseen data. However, evaluating these conditionals on unseen data has known pitfalls. We will first review these pitfalls, then explain how it affects the evaluation of VGI, and finally suggest a simple procedure that enables the evaluation of the VGI objective on held-out data and produces sample imputations.

Cremer et al. (2018) have shown that amortised variational distributions may be significantly biased compared to the optimal variational distribution in the same variational family, a phenomenon they called the amortisation gap. They have also demonstrated that the amortisation gap is typically larger on held-out data if the variational model is fitted only to the training data; we will refer to this as the *inference generalisation gap* (see also the concurrent work by Zhang et al., 2021). The reason for the generalisation gap is overfitting of the variational inference network to its inputs, that is the training data, which is generally a consequence of using finite-sized training data sets. Importantly, the inference generalisation gap can appear even when the target model p_{θ} has not been overfitted.

The same considerations apply to VGI since it is an amortised variational method. In particular, the generalisation gap in VGI can cause the imputation Markov chains to diverge due to the compounding effect of the generalisation errors. This prevents a meaningful evaluation of VGI on the held-out data. We have empirically observed that the chances of divergent behaviour increases with the number of variational conditional models. Mattei and Frellsen (2018) and Cremer et al. (2018) suggested to deal with the inference generalisation gap by fine-tuning the variational models on the held-out data. To avoid information leakage from the held-out data into the model p_{θ} , the estimated model and the corresponding parameters θ are kept fixed during fine-tuning. Moreover, the fine-tuned variational conditionals should not be used as part of the VGI training. A simple way to achieve this is to make a copy of the variational model before fine-tuning and discarding it afterwards.

In Appendix H, Algorithm 6 we show the VGI fine-tuning procedure. The fine-tuning starts with incomplete held-out data and then fills-in the missing values with random draws from the initial imputation distribution f^0 as in Algorithm 1. Then, during the first iter-

ation it performs several rounds of Gibbs sampling over all missing values using the learnt variational conditionals to improve the imputations. To mitigate possible divergent behaviour due to the inference generalisation gap in the first iteration, the procedure rejects any values outside of the observed-data hypercube, defined by the minimum and maximum values in the observed data. After the imputation warm-up in the first iteration, the algorithm continues akin to the learning of ϕ in standard VGI, fine-tuning the variational kernel τ_ϕ and updating the imputations with a small number of Gibbs updates in each iteration, where all proposed imputations are accepted—the acceptance region is no longer necessary, since the kernel is being adapted to the (imputed) held-out data. At the end of the procedure, we obtain the VGI loss on the held-out data as well as imputations of the missing values.

The iterative validation procedure described in this section would be expensive if it were used to continually monitor the validation loss during training. However, the validation loss does not need to be computed at every iteration, computing it only every few iterations is often sufficient. Moreover, we have empirically found that the cost of fine-tuning was often only a small fraction of the training cost.

3.6 Related Work

We here discuss work that is closely related or shares similarities to VGI.

Monte Carlo expectation-maximisation. Monte Carlo EM (MCEM, Wei and Tanner, 1990) has been proposed as an extension of the classical expectation-maximisation (EM, Dempster et al., 1977) algorithm to the setting where the required expectation with respect to the conditional distributions of missing data does not have a closed form expression but is approximated by a Monte Carlo sample average. The method requires sampling from the conditional distribution of missing data at each iteration. Then, like VGI, MCEM iteratively maximises the ELBO using sample imputations of the missing data. In contrast to VGI, MCEM does not use a variational approximation but instead attempts to sample the true conditional distribution, for example, with methods such as rejection sampling (e.g. Barber, 2017, Chapter 27.1.2).

However, sampling from the conditional distributions in higher dimensions usually requires Markov chain Monte Carlo methods (MCMC, e.g. Barber, 2017, Chapter 27.4), which only asymptotically sample from the exact target distribution. In practice, MCMC is used to sample chains of fixed length for computational reasons and hence it may not sample the true conditional distribution. As a consequence, using samples from an unconverged Markov chain in MCEM may adversely affect the learnt model p_θ .

Similar to MCEM with MCMC, VGI also samples a Markov chain of imputations using the variational kernel. However, in contrast to MCEM, where the MCMC sampler needs to be restarted at each iteration and run for a sufficient amount of time, in VGI we use a single persistent Markov chain that we update throughout training. This allows our method to eventually sample better imputations of the missing data given that the algorithm is run for long enough. In Section 6 we estimate a normalising flow model (Rezende and Mohamed, 2015; Papamakarios et al., 2021) from incomplete data with VGI and MCEM using flow-specific MCMC and find that VGI outperforms MCEM both in terms of the computational performance and accuracy. In an attempt to improve the performance of MCEM, we have

empirically investigated the use of persistent chains in a Metropolis-Hastings version of MCMC on normalising flows, but found that this drastically reduced the acceptance rate of the proposed transitions and therefore also reduced the accuracy of the estimated statistical model.

Markov chain variational inference. Salimans et al. (2015) proposed a powerful framework for combining Markov chain Monte Carlo methods and variational inference for latent-variable models called Markov chain variational inference (MCVI). Similar to VGI they propose a lower-bound on the log-likelihood using a Markov chain characterised via a variational transition kernel τ . However, the two methods differ fundamentally in their goals: VGI attempts to *learn* a statistical model p_{θ} from incomplete data while MCVI attempts to efficiently and accurately approximate a conditional distribution under a *fixed* latent-variable model. To achieve their goal, MCVI aims to optimise the variational kernel over the full length of the Markov chain. In order to avoid computing the intractable integral required to obtain the marginal distribution of a Markov chain f_{ϕ}^t they propose a further lower-bound on the ELBO in (3) using a “reverse” transition kernel r , which predicts the reverse path of the Markov chain sampled using the variational kernel τ . A naïve application of MCVI to learn a statistical model would be expensive, since it requires simulating long Markov chains of imputations and then backpropagating the gradients through the sampling path. To simplify the optimisation problem, the authors of MCVI have also briefly considered a sequential (greedy) approach similar to ours. However, as shown in Section 3.1, VGI does not need to learn an auxiliary model of the “reverse” kernel r , and in fact, replacing r in the MCVI lower-bound with the true reverse transition (e.g. Murray, 2007, Section 1.4) recovers the tighter lower-bound in (3) (see Appendix I), which is approximately marginalised in VGI.

A further difference between VGI and MCVI is that MCVI has been developed for the latent-variable and not the missing data setting. Hence, unlike VGI, it does not deal with the problem of how to handle the $2^d - 1$ possible patterns of missingness and the consequent exponential growth in the number of required variational distributions.

Coordinate ascent variational inference. The VGI objective in (9) is related to coordinate ascent variational inference (CAVI, e.g. Bishop, 2006, Chapter 10.1.1; Blei et al., 2017, Section 2.4), which also considers the update of only a single dimension using a univariate variational distribution. However, the previous works only considered fully-factorised variational distributions (mean-field assumption), which can significantly bias the estimate of the statistical model as we have demonstrated in Section 2.2. In contrast, our method works without introducing the mean-field factorisation using a highly-flexible implicit variational distribution given by the marginal of a pseudo-Gibbs sampler, which is adapted to the model p_{θ} by learning the transition kernel (univariate variational conditionals) of the sampler.

Arbitrarily-conditional models. A separate line of research focuses on constructing models such that any conditional distribution $p_{\theta}(\mathbf{x}_u | \mathbf{x}_{\text{obs}})$, where $\mathbf{x}_u \subseteq \mathbf{x} \setminus \mathbf{x}_{\text{obs}}$ may be arbitrarily chosen, is directly available (Li et al., 2020). One could use such models to construct an amortised variational distribution for arbitrary missingness patterns $p_{\theta}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})$ and use it in the variational ELBO to fit a target statistical model. However, unlike the variational conditionals in VGI that we can choose freely, such arbitrarily-conditional

models typically have restricted modelling capabilities compared to their non-conditional counterparts.

Substantive model compatible FCS. As discussed in Section 2.3 our solution to the $2^d - 1$ conditional missing data distributions bears similarity to existing imputation methods in the fully-conditionally specified (FCS) family. Standard FCS methods are independent of the target analysis model. However, Bartlett et al. (2015) have proposed a modified version of the method, called substantive model compatible FCS (SMC-FCS), to generate imputations that are congenial with target (nonlinear) regression model. There is thus some similarity to VGI in the sense that both methods take the target model into account. However, SMC-FCS is about imputing data compatible with regression models (and hence is for supervised learning), while our method estimates joint statistical models (and hence is focused on unsupervised learning). From our understanding SMC-FCS does not apply or readily extend to our setting.

Bayesian data augmentation. Alternatively to maximum-likelihood estimation (MLE) methods, such as the EM or VI, one can also use Bayesian inference to learn models from incomplete data. The advantage of Bayesian methods is that they provide a natural way to evaluate the epistemic uncertainty of the model. One classical method is the data augmentation (DA) algorithm (Tanner and Wong, 1987). In DA, one must first specify a prior over the parameters of the model $p(\boldsymbol{\theta})$ and assume initial imputations $X_{\text{mis}}^{(0)}$.¹⁵ Then, the algorithm iteratively samples the posterior over the parameters $\boldsymbol{\theta}^{(t)} \sim p(\boldsymbol{\theta} \mid X_{\text{mis}}^{(t-1)}, X_{\text{obs}})$ and updates imputations (or augmentations) $X_{\text{mis}}^{(t)} \sim p(X_{\text{mis}}^{(t)} \mid X_{\text{obs}}, \boldsymbol{\theta}^{(t)})$. After an initial warm-up (or burn-in) period, the algorithm produces samples of the missing values and model parameters from the posterior distribution $p(X_{\text{mis}}, \boldsymbol{\theta} \mid X_{\text{obs}})$. Similar to VGI, the iterative procedure in DA is a Gibbs sampler. The main difference between the two is that DA treats the parameters of the model just like the missing variables and samples them accordingly. However, for complex models, just like in the MCEM case, the required distributions are usually only known up to a proportionality and hence computationally expensive MCMC methods would be required to sample them. Moreover, Bayesian inference for modern statistical models, such as VAEs and flows, suffers from scalability issues and is still an active area of research (Gal, 2016; Maddox et al., 2019; Izmailov et al., 2021; Abdar et al., 2021).

4. Experiments on Toy Models

In this section we demonstrate VGI on low- and high-dimensional factor analysis models. We analyse the accuracy of the learnt statistical models and the variational conditionals as well as the effect of the extended-Gibbs variational conditionals. The code for this and the following experimental sections is available at <https://github.com/vsimkus/variational-gibbs-inference>.

15. We use the capital X to denote all data-points \boldsymbol{x}^i in the data set.

4.1 Factor Analysis Model

Factor analysis (FA, e.g. Barber, 2017, Chapter 21.1) is a linear latent-variable model that is often used to discover unobserved factors \mathbf{z} from observed data \mathbf{x} . The prior distribution of \mathbf{z} is assumed to be a multivariate standard Gaussian $p(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \mathbf{0}, \mathbf{I})$, and the distribution of \mathbf{x} given a value of \mathbf{z} is

$$p_{\theta}(\mathbf{x} \mid \mathbf{z}) = \mathcal{N}(\mathbf{x}; \mathbf{F}\mathbf{z} + \boldsymbol{\mu}, \boldsymbol{\Psi}),$$

and the marginal distribution is

$$p_{\theta}(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \mathbf{F}\mathbf{F}^{\top} + \boldsymbol{\Psi}),$$

where the parameters θ are the mean vector $\boldsymbol{\mu}$, the factor matrix \mathbf{F} , and the diagonal matrix $\boldsymbol{\Psi}$ defining the variances of the observation noise.

FA is a linear version of the variational autoencoder (VAE, Kingma and Welling, 2013; Rezende et al., 2014) that is often used as a toy model to analyse new methods (e.g. Williams et al., 2018). Hence, we start our analysis on a FA model before proceeding to evaluate VGI on more complex models in the next sections. Given its relative simplicity, it is possible to fit a FA model on incomplete data with the expectation-maximisation (EM) algorithm (see Appendix J for details). The EM algorithm provides a best-case solution that we can use to gauge the accuracy of the model estimated by VGI (see below). Moreover, reference conditional distributions $p_{\theta}(x_j \mid \mathbf{x}_{\setminus j})$ can be computed analytically (Petersen and Pedersen, 2012), such that we can evaluate the accuracy of the variational conditionals, which is instrumental to producing good imputations and subsequently—to achieving an accurate fit of the model.

4.2 Data

We evaluate VGI on two data sets:

Toy data. A synthetic data set generated with a 6-dimensional FA model with a 2-dimensional latent space (see Appendix K for the ground truth parameters). The training data set has 6400 data-points and the test data set has 5000 data-points.

FA-Frey. A synthetic 560-dimensional data set based on the Frey data,¹⁶ where a FA model with 43 latent dimensions was first fitted on the original data and then used to synthesise a new data set. The training data set has 2400 data-points and the test data set has 3000 data-points.

We consider five fractions of missingness in the training data, ranging from 16.6% to 83.3%, and simulate incomplete training data by generating a binary missingness mask uniformly at random (MCAR). The data-points that are rendered fully missing are removed from the training set, since doing so does not affect the maximum-likelihood estimate under MAR or MCAR missingness.

Where the experiments are repeated multiple times to obtain confidence intervals, the underlying training data are kept constant and only the missingness mask is re-sampled.

16. The original data set is available at https://cs.nyu.edu/~roweis/data/frey_rawface.mat.

Thus, the results demonstrate the accuracy of the estimation methods in the presence of missing data, rather than the variability of the model estimate on different realisations from the ground truth data generating distribution.

4.3 Experimental Settings

In our evaluation, we assume that the statistical model p_{θ} is well-specified. That is, we fit a FA model to data following a FA model with the same latent dimensionality, so that the evaluation can focus on the effect of missing data, rather than on robustness to model specification. We parametrise Ψ as $\Psi = \exp(\gamma)$, initialise the parameter \mathbf{F} with samples from a standard normal distribution, and set $\boldsymbol{\mu}$ and $\boldsymbol{\gamma}$ to $\mathbf{0}$ and $\mathbf{1}$ respectively.

We specify the variational conditionals in VGI to be univariate Gaussians whose parameters $\log \sigma_j^2$ and μ_j are given by the outputs of a fully-connected neural network. For the toy data we use the standard variational conditionals with an independent network of two hidden layers for each conditional. To make the computations more efficient on the FA-Frey data, we use the extended variational conditionals with a partially-shared network: the first two hidden layers share parameters and computations for all conditional distributions and the last hidden layer has independent parameters for each distribution. The networks use leaky ReLU activation functions with negative slope of 0.01 and the weights are initialised using Kaiming initialisation (He et al., 2015).

We compute the univariate Gaussian entropy terms in (9) analytically using (e.g. Norwich, 1993)

$$-\mathbb{E}_{\mathcal{N}(x_j; \mu_j, \sigma_j^2)} \left[\log \mathcal{N}(x_j; \mu_j, \sigma_j^2) \right] = \frac{1}{2} \log(2\pi\sigma_j^2) + \frac{1}{2},$$

where μ_j and $\log \sigma_j^2$ are given by the inference networks with input $(\tilde{\mathbf{x}}_{\text{mis} \setminus j}, \mathbf{x}_{\text{obs}})$, or $(\tilde{\mathbf{x}}_{\text{mis}}, \mathbf{x}_{\text{obs}})$ in the case of the extended variational conditionals.

We fit the model $p_{\theta}(\mathbf{x})$ and the variational model using Algorithm 1. We use $K = 5$ imputation chains for each incomplete data-point, and $G = 3$ (toy-data) and $G = 5$ (FA-Frey) Gibbs updates. In the Monte Carlo averaging in $\hat{\mathcal{J}}_{\text{VGI}}$ we select $M = 1$ (toy-data) and $M = 10$ (FA-Frey) missing dimensions. To compute the gradients with respect to the variational parameters we use the reparametrisation trick (Kingma and Welling, 2013). The model parameters $\boldsymbol{\theta} = (\mathbf{F}, \boldsymbol{\mu}, \boldsymbol{\gamma})$ are optimised using the Adam optimiser (Kingma and Ba, 2014), whereas the variational parameters $\boldsymbol{\phi}$ are fitted using AMSGrad (Reddi et al., 2018), since we found that using Adam on the variational parameters caused training instability.

4.4 Comparison Methods

We compare VGI against the following methods:

EM (Complete). The ideal case where no data is missing, fitted using EM for FA.

Empirical sample imputation. A weak baseline where the incomplete data is $K = 5$ times imputed with random draws from the empirical distribution of the observed values, then the FA model is fitted as described in Appendix D.

EM (Dempster et al., 1977). An optimal method where the model p_{θ} is fitted using EM for FA with missing data (see Appendix J). This method presents the best-case

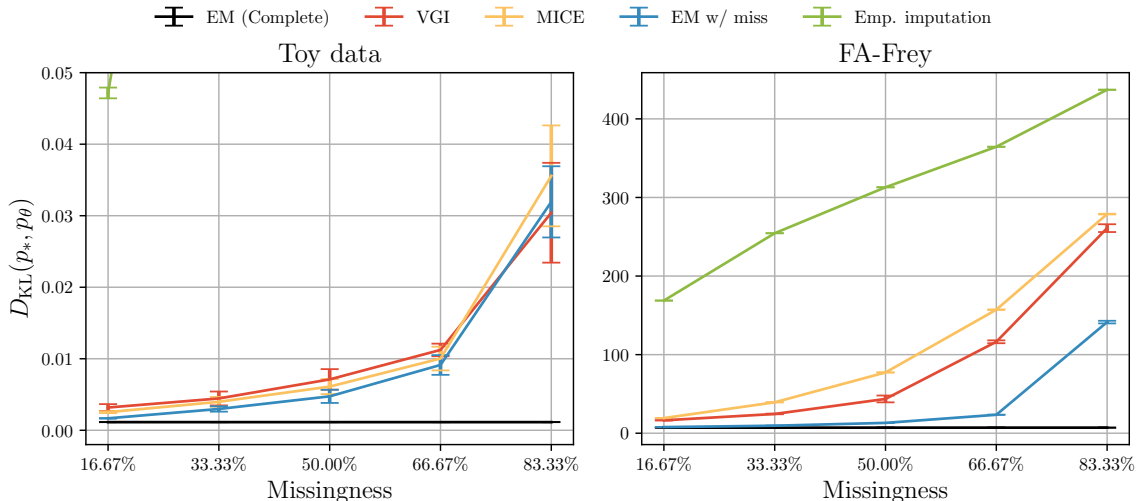


Figure 4: The accuracy of the fitted models p_θ as measured by $D_{\text{KL}}(p_* || p_\theta)$ on toy (left) and FA-Frey (right) data. Smaller values are better. The error bars show standard error on the mean of five experiments with different missingness masks and model initialisation. In the left figure, empirical sample imputation performed significantly worse than the posterior approximating methods and hence is not shown for missingness larger than 16.67%.

performance that could be achieved with a variational method, which is equivalent to setting the variational distribution in VI equal to the true conditional distribution (e.g. Barber, 2017, Chapter 11.2.2). Also, in contrast to the other methods, where SGA is used, we here compute the updated distribution parameters analytically.

MICE (van Buuren and Oudshoorn, 2000). A pseudo-Gibbs sampler as described in Section 2.3 that has been widely used in statistical analysis of incomplete data. MICE code has originally been made available in R (van Buuren and Oudshoorn, 2000). We have implemented MICE in Python using the IterativeImputer and Bayesian linear ridge regression as the conditional imputation models, as implemented in the scikit-learn package (Pedregosa et al., 2011). MICE is used to produce $K = 5$ imputations and then the FA model is fitted as described in Appendix D. MICE is another strong baseline because it should be able to produce imputations that are congenial to the target model, since both MICE conditionals and the target FA model are linear Gaussian models.

4.5 Accuracy of the Fitted FA Model

In this toy setting no over-fitting was observed, hence the model parameters $\hat{\theta}$ from the final training iteration were used in the following evaluation for all methods.

We evaluate the accuracy of the fitted FA model using the Kullback–Leibler divergence $D_{\text{KL}}(p_*(\mathbf{x}) \parallel p_{\theta}(\mathbf{x}))$ between the ground truth model p_* and the fitted statistical model p_{θ} on the two synthetic data sets, shown in Figure 4. It can be immediately seen that the posterior-approximating methods (VGI, MICE, and EM) perform significantly better than the simple empirical imputation baseline showing the clear advantage of these methods over ad-hoc approaches commonly used in practice as a quick fix for missing data.

As expected, MICE performs well on both data sets since the linear imputation model is congenial with the data distribution and the target model. VGI performs comparably to MICE on the toy data (note the overlapping error bars) and shows significant improvement on the FA-Frey data. The better performance on FA-Frey can be attributed to the fact that contrary to MICE, where missing value imputations are generated prior and independently of the model p_{θ} , in VGI the imputations are generated with respect to the model p_{θ} and are not static throughout training, thus better representing the uncertainty of the imputed values.¹⁷ On the toy data, both VGI and MICE achieved a performance that is comparable to the optimal EM solution, but on the FA-Frey data, the gap between them and EM increases with missingness. We show in Figure 15 of Appendix M that a similar gap appears between EM and Monte Carlo EM (MCEM, Wei and Tanner, 1990) with SGA. Hence, we attribute the performance gap to the stochasticity in the optimisation—Monte Carlo averaging and stochastic gradient ascent—used within MCEM, MICE, and VGI, and hence the gap could be reduced, given sufficient compute resources, via standard means in stochastic optimisation.

4.6 Assessing Estimation Consistency

We further evaluate the statistical consistency of VGI on the toy FA model. The learning rate was decayed according to a cosine schedule in order to satisfy the convergence conditions of stochastic gradient ascent (SGA) (e.g. Spall, 2003, Chapter 4.3.2). Note that in these experiments the variational family of the variational conditionals includes the true conditional distributions, and hence the estimator is expected to be unbiased.

In Figure 5 we plot the log-log curves of the number of incomplete training data-points versus the root mean squared error (RMSE) of the fitted model parameters (left subfigure) and the KL divergence (right subfigure).¹⁸ The plots show a linear behavior in the logarithmic domain, which indicates consistency and conforms with the asymptotic normality of the MLE theory (e.g. Wasserman, 2005, Chapter 9.7). We note that a few points in the figure fall above the linear curve. We attribute this to normal sample variation and to false convergence of SGD due to a potentially sub-optimal learning rate decay schedule (Spall, 2003, Chapter 4.3.2) that could be addressed via hyper-parameter search.

4.7 Accuracy of the Variational Model

The imputation accuracy depends on the fit of the variational conditionals, which subsequently influences the accuracy of the fitted model. Hence, we evaluate, and show in

17. We note that the performance of MICE could be improved by generating more imputations, however with additional computations.

18. The factor loading matrix of the fitted model has been rotated using the (orthogonal) Procrustes rotation (Ten Berge, 1977) before computing the RMSE to resolve the partial non-identifiability of the parameter.




Figure 5: The accuracy of the fitted toy FA model in terms of RMSE with the ground truth parameters (left) and KL divergence to the ground truth model (right) as a function of the number of training samples (on a log-log scale). The error bars show standard error on the mean of five experiments with different missingness masks and model initialisation.

Figure 6, the quality of the fitted variational conditional distributions using the KL divergence between the learnt conditionals and the ground truth distribution $D_{\text{KL}}(q_{\phi}(x_j | \mathbf{x}_{\setminus j}) || p_*(x_j | \mathbf{x}_{\setminus j}))$ in blue and the conditionals of the learnt statistical model $D_{\text{KL}}(q_{\phi}(x_j | \mathbf{x}_{\setminus j}) || p_{\theta}(x_j | \mathbf{x}_{\setminus j}))$ in red, where the conditional distributions are computed on test data.¹⁹

The KL divergence to the ground truth indicates how good the imputations sampled from the variational conditionals are. It can be seen that the approximation of the ground truth conditionals gets worse as the missingness increases, which is also in line with the accuracy of the target model in Figure 4. This is expected since the variational conditionals approximate the conditional distribution of the target model p_{θ} .

In contrast, the KL divergence to the learnt model p_{θ} shows how well the variational conditionals approximate the true conditional distribution under the fitted model, this is the objective that is minimised by \mathcal{J}_{VGI} so we expect it to be low. We observe that the variational conditionals are good at the majority of test data-points (see red curves), however we see a long tail where the variational approximations are poor (see the tail of the red violin plot). This under-performance of the variational conditionals without fine-tuning corresponds to the approximation gap discussed in Section 3.5. Hence, when performing Gibbs sampling with the variational conditionals for model selection, some light fine-tuning of the variational model might be necessary in order to prevent divergent Gibbs imputation chains, as discussed in Section 3.5.

¹⁹. For the toy data, we also show results dissected by dimension in Appendix M, Figure 16.

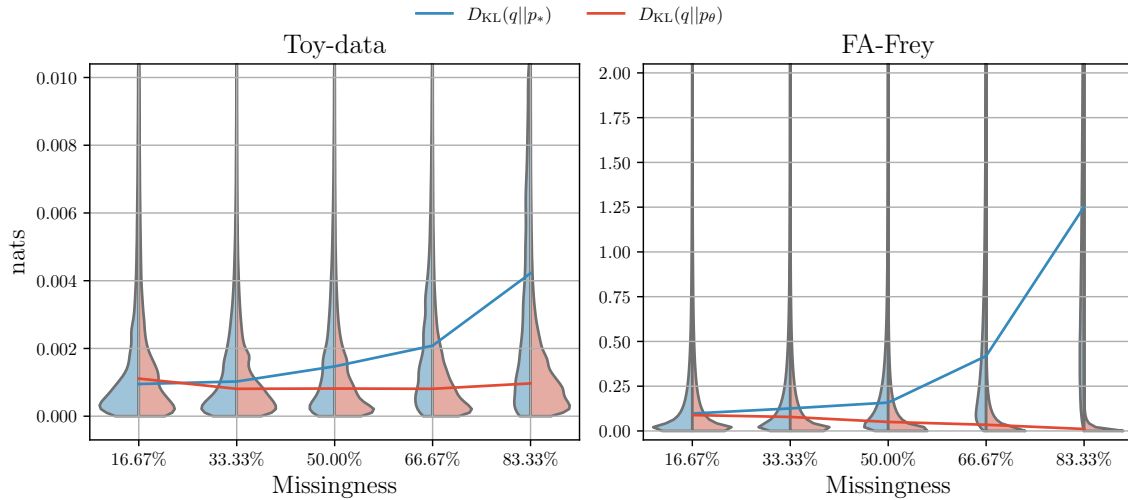


Figure 6: KL divergence, conditioned on the test set (left: toy data, right: FA-Frey data), between the univariate variational conditional distributions and the ground truth distribution $D_{\text{KL}}(q_{\phi_j}(x_j | \mathbf{x}_{\setminus j}) || p_*(x_j | \mathbf{x}_{\setminus j}))$ (blue), and the posterior under the learnt model $D_{\text{KL}}(q_{\phi_j}(x_j | \mathbf{x}_{\setminus j}) || p_{\theta}(x_j | \mathbf{x}_{\setminus j}))$ (red). The lines show the median KL divergence conditioned on the test set. In Appendix M, Figure 17 we show a comparison using Wasserstein distance that displays a similar behaviour.




Figure 7: Median KL divergence of the univariate variational conditional distributions, conditioned on the toy test set. Comparing independent against shared-parameter variational models.

Furthermore, we see that the median KL divergence to the conditionals of the fitted model remains constant for all fractions of missingness on the toy data, but it goes down with increasing missingness on the FA-Frey data (solid red lines). We attribute this to the partial weight sharing used in the variational model on FA-Frey data. We verify this by additionally comparing an independent model (red) against a shared, but not extended, model (blue) on the toy data in Figure 7.²⁰ We can see that with the partially-shared model the KL divergence on toy data also improves with missingness (dashed blue line), whilst it remains almost constant for the independent model (dashed red line). The improved approximation of the conditionals with increasing missingness and partially-shared variational models can be explained by additional amortisation—it is akin to having more data to learn the same number of parameters, which results in a better fit of the variational model.

We further investigated the effect of the extended variational model on the conditional distributions using the toy data (Section 3.3): In Figure 7 we compare a shared variational model with the extended conditionals (green) to models with standard conditionals (red and blue). We observe that the extended conditionals are close approximations of the true Gibbs conditionals, and only slightly worse than the standard variational conditionals. Moreover, the estimate of the model p_{θ} is only slightly affected for missingness larger than 66% (see Figure 18 in Appendix M). Hence, we conclude that conditioning on an additional variable \tilde{x}_j in the extended-Gibbs variational model (Appendix E) results in inconsequential additional variability (“noise”) of the variational distributions. Importantly, in higher dimensional problems, the noise due to conditioning on a single extra variable decreases. There is no significant difference between VGI estimation accuracy with the standard variational conditionals and the extended conditionals when the data dimensionality is larger (see Figure 19 in Appendix M). Hence, when the dimensionality is small and the computational cost is low it is best to use the standard variational conditionals, however, when the dimensionality is large, it is more favourable to use the extended variational model for computational reasons.

5. Experiments on VAE Models

In this section we estimate VAE models from incomplete data. We compare the general-purpose VGI method, VAE-specific methods based on VBGI, and existing VAE-specific methods in the literature in terms of estimation accuracy and computational efficiency.

5.1 Variational Autoencoder Model

The variational autoencoder (VAE) is a non-linear descendant of the factor analysis model. It is a latent-variable model with observables \mathbf{x} and latent variables \mathbf{x}_z , defined via

$$p_{\theta}(\mathbf{x} \mid \mathbf{x}_z) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_{\theta}(\mathbf{x}_z), \Psi_{\theta}(\mathbf{x}_z)), \quad \text{and} \quad p(\mathbf{x}_z) = \mathcal{N}(\mathbf{x}_z; \mathbf{0}, \mathbf{I}),$$

where the $\boldsymbol{\mu}_{\theta}$ and Ψ_{θ} are deep neural networks with parameters θ , and $\Psi_{\theta}(\mathbf{x}_z)$ is diagonal. We use the Gaussian family for the generator, which is the most common case in the VAE literature, although another family of distributions could also be used. The model is

20. We also compare the accuracy of the fitted FA models with both variational models on toy data in Appendix M Figure 18, where we find both approaches comparable in the model p_{θ} accuracy.

typically optimised using variational inference (Kingma and Welling, 2013; Rezende et al., 2014), where a variational distribution $q_\phi(\mathbf{x}_z | \mathbf{x})$ is used to approximate the intractable posterior $p_\theta(\mathbf{x}_z | \mathbf{x})$. The parameters θ and ϕ are then optimised with stochastic gradient ascent by maximising the following lower-bound

$$\log p_\theta(\mathbf{x}) \geq \mathbb{E}_{q_\phi(\mathbf{x}_z | \mathbf{x})} \left[\log \frac{p_\theta(\mathbf{x} | \mathbf{x}_z)p(\mathbf{x}_z)}{q_\phi(\mathbf{x}_z | \mathbf{x})} \right]. \quad (14)$$

5.2 Data

We evaluate VGI on a synthetic 560-dimensional VAE-Frey data set, where a VAE model with one hidden layer in the encoder and decoder and a 10-dimensional latent space was first fitted on the Frey data set, and then it was used to synthesise a new data set that we call VAE-Frey. We used the same VAE architecture as Kingma and Welling (2013). The training data set has 2400 data-points and the test data set has 3000 data-points. The rest of the data setup is identical to Section 4.2.

5.3 Experimental Settings

As before, we specify a target VAE model to fit data following a ground truth VAE model with the same architecture, so that the evaluation can focus on the effect of model estimation from incomplete data, rather than on the robustness of the specified model.

We consider three VGI-based methods for fitting the VAE model:

VGI. This method does not use any assumptions of the statistical model p_θ and hence uses the univariate variational conditional formulation of VGI, as before. We specify the variational encoder architecture to be equivalent to the encoder of the ground truth and replace the $\log p_\theta(\mathbf{x})$ in \mathcal{J}_{VGI} in (9) with the VAE ELBO from (14).

We use the same univariate variational model with extended conditionals and partially-shared parameters as in the FA-Frey data experiments in Section 4.3 and optimise the variational parameters using AMSGrad (Reddi et al., 2018). The other VGI hyper-parameters are also equivalent to those used in the FA-Frey data experiments: $K = 5$ imputation chains, $G = 5$ Gibbs update steps, and $M = 10$ sampled missing dimensions in $\hat{\mathcal{J}}_{\text{VGI}}$.

VBGI-VAE. Here we use the structure and assumptions of a VAE to adapt the VGI method to the model p_θ .

First, VAE is a latent-variable model so we can treat the latent variables as missing and group them together as we discussed in Section 3.4. We then specify a joint variational distribution $q_{\phi_z}(\mathbf{x}_z | \tilde{\mathbf{x}}_{\text{mis}}, \mathbf{x}_{\text{obs}})$, which corresponds to the encoder of the standard VAE.

Second, the missing variables \mathbf{x}_{mis} are assumed to be conditionally independent Gaussians given the latents \mathbf{x}_z , hence we can easily specify the joint distribution of the conditionals $q_{\phi_j}(x_j | \mathbf{x}_{\text{mis} \setminus j}, \tilde{\mathbf{x}}_z, \mathbf{x}_{\text{obs}})$, $j \in \text{idx}(\mathbf{m})$ using a shared variational model $q_{\phi_{\text{mis}}}(\mathbf{x}_{\text{mis}} | \tilde{\mathbf{x}}_z, \mathbf{x}_{\text{obs}})$ for all patterns of missingness, similar to the target generative model. To parametrise $q_{\phi_{\text{mis}}}$ using a neural network, we pad the \mathbf{x}_{obs} with zeros for the missing dimensions to get a fixed-size vector for all patterns of missingness.

Whilst the VAE model assumes that \mathbf{x}_{mis} and \mathbf{x}_{obs} are independent given \mathbf{x}_z , in the variational conditionals $q_{\phi_{\text{mis}}}$ we do not simplify the conditioning set but condition on \mathbf{x}_{obs} as

in the general VGI formulation. Conditioning on \mathbf{x}_{obs} compensates for possible information loss about the true $\mathbf{x}_z^* \sim p_{\theta}(\mathbf{x}_z | \mathbf{x}_{\text{obs}})$ due to the use of an approximation q_{ϕ_z} , and hence eases the learning of $q_{\phi_{\text{mis}}}$. We empirically validated this theoretical argument against a simplification of the variational conditional and observed that conditioning $q_{\phi_{\text{mis}}}$ on \mathbf{x}_{obs} , as dictated by the general VGI methodology, is indeed crucial to the performance of the method.

Then, the index $j \in \{\text{mis}, z\}$ in the VGI objective refers to either all the missing variables \mathbf{x}_{mis} or the latents \mathbf{x}_z . Incorporating the VAE assumptions into the VBGI objective for latent-variable models from (13) we get the VAE-specific objective

$$\mathcal{J}_{\text{VBGI-VAE}}^t(\boldsymbol{\theta}, \boldsymbol{\phi}; \mathbf{x}_{\text{obs}}) = \mathbb{E}_{f^{t-1}(\tilde{\mathbf{x}}_{\text{mis}}, \tilde{\mathbf{x}}_z | \mathbf{x}_{\text{obs}})} \left(\begin{aligned} & \pi(j = \text{mis}) \mathbb{E}_{q_{\phi_{\text{mis}}}(\mathbf{x}_{\text{mis}} | \tilde{\mathbf{x}}_z, \mathbf{x}_{\text{obs}})} \left[\log \frac{p_{\theta}(\mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}} | \tilde{\mathbf{x}}_z) p(\tilde{\mathbf{x}}_z)}{q_{\phi_{\text{mis}}}(\mathbf{x}_{\text{mis}} | \tilde{\mathbf{x}}_z, \mathbf{x}_{\text{obs}})} \right] \\ & + \pi(j = z) \mathbb{E}_{q_{\phi_z}(\mathbf{x}_z | \tilde{\mathbf{x}}_{\text{mis}}, \mathbf{x}_{\text{obs}})} \left[\log \frac{p_{\theta}(\tilde{\mathbf{x}}_{\text{mis}}, \mathbf{x}_{\text{obs}} | \mathbf{x}_z) p(\mathbf{x}_z)}{q_{\phi_z}(\mathbf{x}_z | \tilde{\mathbf{x}}_{\text{mis}}, \mathbf{x}_{\text{obs}})} \right] \end{aligned} \right),$$

where $\pi(j) = \frac{1}{2}$ for all j , meaning that we choose a uniform Gibbs selection probability. We use $K = 5$ imputation chains and $G = 2$ Gibbs updates corresponding to one full update of all unobserved variables \mathbf{x}_{mis} and \mathbf{x}_z , and optimise the above objective using the Adam optimiser.

VBGI-VAE-M. Alternatively, we can use the conditional independence of the observable variables given the latents for the VAE model to marginalise the missing observable dimensions \mathbf{x}_{mis} from the likelihood, which is commonly done in the VAE literature for incomplete data. Then, $j = \text{mis}$ and the objective above simplifies to

$$\mathcal{J}_{\text{VBGI-VAE-M}}^t(\boldsymbol{\theta}, \boldsymbol{\phi}; \mathbf{x}_{\text{obs}}) = \mathbb{E}_{f^{t-1}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}) q_{\phi_z}(\mathbf{x}_z | \mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}})} \left[\log \frac{p_{\theta}(\mathbf{x}_{\text{obs}} | \mathbf{x}_z) p(\mathbf{x}_z)}{q_{\phi_z}(\mathbf{x}_z | \mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}})} \right].$$

To update the Markov chain distribution f^t we perform block-Gibbs sampling using $q_{\phi_z}(\mathbf{x}_z | \mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}})$ for the latents and $p_{\theta}(\mathbf{x}_{\text{mis}} | \mathbf{x}_z)$ for the missing observables. This imputation procedure was initially proposed by Rezende et al. (2014) for missing data imputation on test data, however to the best of our knowledge it has not been considered for learning from incomplete data. We use $K = 5$ imputation chains and $G = 2$ Gibbs updates corresponding to one full update of all unobserved variables \mathbf{x}_{mis} and \mathbf{x}_z , and optimise the above objective using the Adam optimiser.

In all of the above methods we use the reparametrisation trick (Kingma and Welling, 2013) to compute the gradients of the variational models.

5.4 Comparison Methods

We compare the VGI-based methods against:

VAE (Complete). The ideal case where no data is missing, fitted using standard variational inference for VAEs (Kingma and Welling, 2013; Rezende et al., 2014).

MICE (van Buuren and Oudshoorn, 2000). The same as in Section 4.4.

missForest (Stekhoven and Bühlmann, 2012). An iterative FCS method that uses random forest to model the conditionals. Unlike MICE, the imputations are not sampled but instead are replaced with the deterministic values predicted by the random forest. Hence, while the method can be used on non-linear data, the imputations may be biased by lack of uncertainty representation. We implemented the method in Python using IterativeImputer and RandomForestRegressor for the conditional models from the scikit-learn package (Pedregosa et al., 2011). The default settings for this method proved to be too computationally expensive on this data set, hence to make this method computationally feasible we traded-off the expressivity of the random forest for better computational performance by limiting the number of decision trees to 20 and the maximum depth of the trees to 20.²¹ Random forest methods are typically piecewise-constant regressors, which means that by limiting the depth and the number of trees, for reasons above, we may get a too coarse an approximation of the target variable, on the other hand, if the computational budget allows, the trees could be “fully grown” such that each leave corresponds to a unique training data point and hence, in the idealistic scenario of a finely sampled training set, the random forest regressor could be made arbitrarily accurate. Like MICE, missForest was used to produce $K = 5$ imputations and then the VAE model was fitted as described in Appendix D.

MICEForest. One classical way to calibrate the uncertainty of imputations for methods like missForest is predictive mean matching (Little, 1988; van Buuren, 2018, Chapter 3.4). We hence use the MICEForest Python package that provides an implementation of missForest enhanced with predictive mean matching.²² The uncertainty is controlled via a hyperparameter that specifies the number of nearest neighbours (based on the predictive mean) considered for a random imputation draw. As with missForest above, to make this method feasible on the high-dimensional data in this section, the number of estimators was set to 20, the maximum depth of the decision trees was set to 20, and the number of nearest neighbours was set to the default 5.²³

MVAE (Nazábal et al., 2020; Mattei and Frellsen, 2019). The missing dimensions are marginalised in the generator of the VAE and the encoder masks the missing dimensions with zeros.

HIVAE (Nazábal et al., 2020). The missing dimensions are marginalised in the generator of the VAE. The model uses a hierarchical prior and zero masking in the encoder for missing data.

PartialVAE+ (Ma et al., 2019). The missing dimensions are marginalised in the generator of the VAE. Uses a permutation-invariant encoder network with learnable location embeddings, which can handle partially-observed inputs.

21. One imputation of the data set with missForest took about 27 hours on CPU with the scikit-learn implementation.

22. MICEForest implementation can be found here <https://github.com/AnotherSamWilson/miceforest>.

23. One imputation of the data set with MICEForest took about 17 hours on CPU.

MIWAE (Mattei and Frellsen, 2019). The missing dimensions are marginalised in the generator of the VAE. Uses a tighter importance weighted lower-bound and zero masking in the encoder.

The above methods are all fitted using the Adam optimiser.

5.5 Accuracy of the Fitted VAE Models

As is common in machine learning the accuracy of the VAE models increased with training iterations and then decreased due to over-fitting, hence we used model selection on an incomplete held-out validation data set to mitigate this effect (for more details see Appendix L and Figure 20). We then use the selected model parameters $\hat{\theta}$ to assess the accuracy of the fitted model.

In order to evaluate the accuracy of the selected VAE models, we estimated the log-likelihood on complete test data. In general, computing the marginal log-likelihood of a VAE is intractable, therefore we estimate it using the IWAE bound (Burda et al., 2015) that uses samples from a variational proposal distribution q and self-importance weighting,

$$\mathcal{L}_{\text{IWAE}}(\mathbf{x}) = \log \left(\frac{1}{L} \sum_{l=1}^L \frac{p_{\hat{\theta}}(\mathbf{x}, \mathbf{x}_z^l)}{q(\mathbf{x}_z^l | \mathbf{x})} \right), \quad \text{where } \mathbf{x}_z^1, \dots, \mathbf{x}_z^L \sim q(\mathbf{x}_z | \mathbf{x}).$$

The IWAE bound approaches $\log p_{\hat{\theta}}(\mathbf{x})$ monotonically from below as the number of importance samples L increases (Burda et al., 2015, Theorem 1) irrespective of the proposal distribution (subject to minor conditions). However, we note that the non-asymptotic bias of the estimator is in the order of $\mathcal{O}(1/L)$ (Owen, 2013; Paananen et al., 2021) and depends on the accuracy of the proposal distribution $q(\mathbf{x}_z | \mathbf{x})$, and is zero if $q(\mathbf{x}_z | \mathbf{x}) = p_{\theta}(\mathbf{x}_z | \mathbf{x})$. Hence, to attain a good proposal and mitigate the bias we fine-tune the variational encoder q_{ϕ} from the training stage on the complete test data (Mattei and Frellsen, 2018).²⁴ We then estimate the log-likelihood using $\mathcal{L}_{\text{IWAE}}$ and the fine-tuned encoder with a large number $L = 50000$ of importance samples.

In Figure 8 we show the estimated marginal log-likelihood averaged over the test data. We see that the VGI-based methods perform consistently better than the existing VAE-specific methods, particularly in settings with missingness fraction greater than 50%. It is interesting that VGI and VBGI-VAE, which learn a variational imputation model of the missing observable variables, both outperform VBGI-VAE-M, which marginalises the missing variables from the likelihood. This suggests that forcing the VAE model to reconstruct a completed data-point is overall better for the accuracy of the generator than marginalising the missing variable dimensions in the generator as it is often done in the VAE-specific methods and VGI-VAE-M. We conjecture that marginalising the missing variable dimensions allows the latent representation in VAEs to systematically forget (collapse) some of the information presented in the inputs. In contrast, forcing the VAE model to reconstruct a completed data-point provides an adaptive regularisation that prevents such loss of information.

24. We have also considered fitting a randomly-initialised encoder on the complete test data but due to poor initialisation the fitted variational approximations were worse (also observed in previous works e.g. Altsaar et al., 2017).

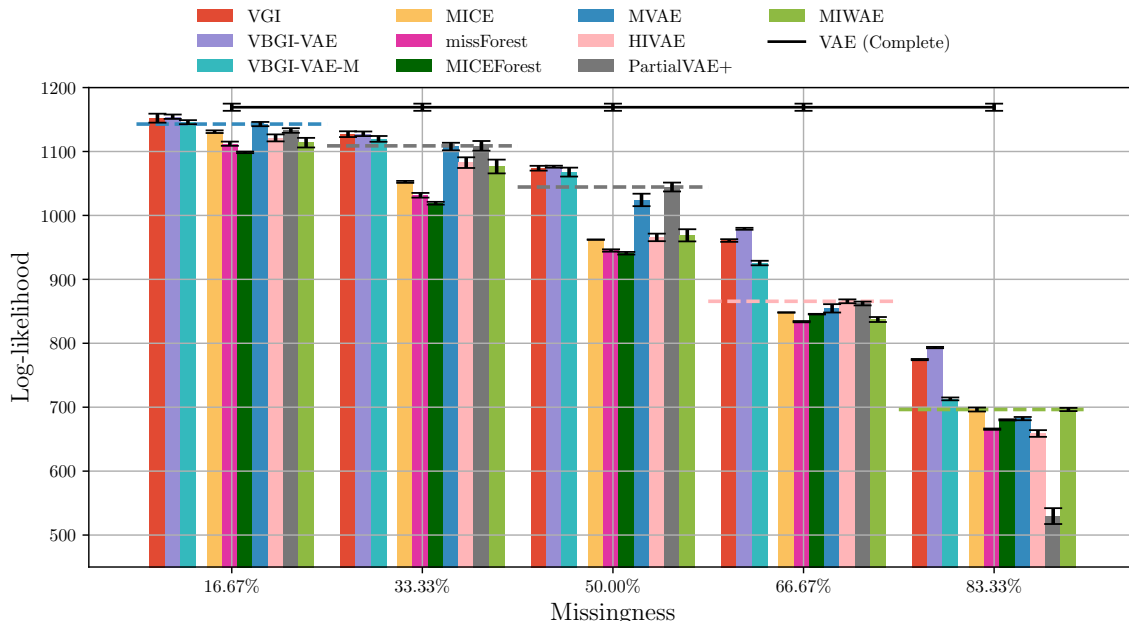


Figure 8: Importance weighted estimates of the marginal log-likelihood on complete test data. The error bars show standard error on the mean of five experiment repetitions with different missingness masks and model initialisation. The dashed horizontal lines show the best performing competitor model, which we note is always outperformed by the VGI-based methods.

The learning curves for the different methods are shown in Figure 21 in Appendix M. We note that for missingness fractions greater than 50%, the methods that marginalise the missing dimensions from the likelihood (VAE-specific and VBGI-VAE-M) start to over-fit on the training data (compare solid and dash-dotted curves) early, whereas the VGI and VBGI-VAE monotonically improve throughout the training, which allows them to surpass the other methods in terms of the model accuracy, especially at greater than 50% missingness, which we also observe in Figure 8.

Moreover, we show in Figure 22 in Appendix M that working with the extended variational model in VGI had no adverse effect on the estimation accuracy while greatly reducing the computational cost. This adds additional evidence to the results in Section 4.7.

Regarding the impute-then-fit methods, we notice that MICE performs well despite being a linear method. This can be explained as follows: the generator is only slightly non-linear since it is parametrised by a one-hidden layer neural network and the data is based on the Frey data set, which can be fitted rather well by linear methods, such as factor analysis. Figure 8 shows that the non-linear imputation methods, missForest or MICEForest, did not perform well in this experiment. We attribute the poor performance of these methods to two factors: modelling limitations due to computational cost restrictions, see Section 5.4, and, possibly, an incorrect representation of imputation uncertainty. If the computational budget

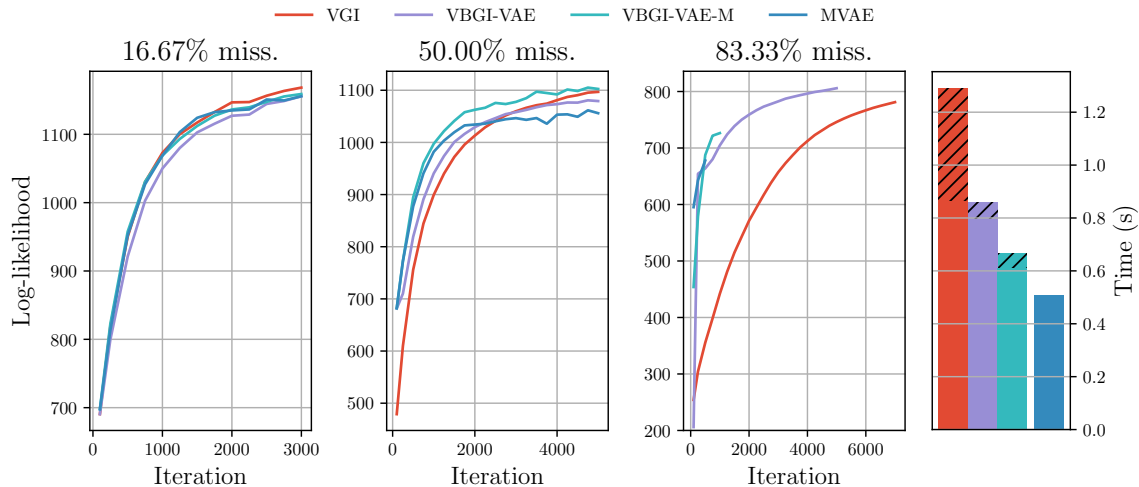


Figure 9: Left: Estimated test log-likelihood on complete data against training iteration. Right: Average duration of one training iteration in seconds. The hatched part of the bars indicates the time spent on updating the imputations via pseudo-Gibbs sampling. Plots including other fractions of missingness and the other VAE-specific methods can be found in the Appendix M Figures 23 and 24.

allows, the former issue can be addressed by using more or larger decision trees although preventing overfitting will then become important. The second issue could be addressed in MICEForest by tuning the number of nearest neighbours in predictive mean matching, although current recommendations are more based on heuristics than firm principles (e.g. van Buuren, 2018, Chapter 3.4), which can make finding an optimal setting difficult.

Finally, we observe that MVAE performs almost always better than MIWAE, even though MIWAE optimises a tighter importance-weighted lower-bound. This suggests that when the statistical model p_{θ} matches the true generative model but, due to missingness, we do not have enough data to actually fit the model, MIWAE might be more prone to overfitting since it can tighten the bound more effectively. To confirm this intuition, we fitted an under-specified model to the incomplete data (by reducing the latent dimension to 5), and in line with our explanation above, MIWAE was then performing better than MVAE. Moreover, PartialVAE+ performed considerably well on missingness levels up to 83.33%, however the performance went down significantly at the largest missingness which points to a difficulty of approximating the variational distribution with a permutation-invariant encoder at very large missingness.

5.6 Computational Aspects

We here discuss the computational aspects of VGI-based methods for VAE estimation, and contrast it with the VAE-specific methods.

In the left-hand side of Figure 9 we plot the estimated test log-likelihood against the training iteration and compare the VGI-based methods to MVAE. We see that out of all VGI-based methods VBGI-VAE-M has the highest rate of likelihood improvement, which is comparable to the rate of MVAE, while achieving consistently better accuracy of the fit. However, the methods that marginalise the missing variables from the likelihood (MVAE and VBGI-VAE-M) quickly over-fit to the training data in the higher missingness settings and hence perform worse than VGI and VBGI-VAE. VBGI-VAE closely follows VBGI-VAE-M at a slightly slower rate of improvement, due to having to learn an additional variational model of the missing observables, but it offers a generally better accuracy of the fit, as seen in the previous section, and does not over-fit to the training data like MVAE and VBGI-VAE-M. The general-purpose VGI method, which does not use the assumptions available in a VAE, has the slowest rate of improvement, which is most clearly seen in high missingness settings. However, as VBGI-VAE, VGI generally outperforms the other methods that marginalise the missing variables from the likelihood in terms of the accuracy of the fit.

The average duration of a training iteration is shown in the right-hand side of Figure 9. We used the same VGI hyper-parameters for all levels of missingness in our experiments, which means that the average per-iteration cost is about the same. While the cost of an iteration for VAE-specific methods does not depend on the missingness, for VGI-based methods, the iteration cost could be adjusted to the level of missingness by, for example, adapting the number of Gibbs update steps (G in Algorithm 5), the number of imputation chains K , or the number of sampled missing dimensions (M in (11)). The figure thus shows results that were not optimised for computational efficiency. The average duration of an iteration of the VGI-based methods are about 2.4 (VGI), 1.6 (VBGI-VAE), and 1.25 (VBGI-VAE-M) times longer than that of the other methods. Moreover, about 0.6 (VGI), 0.2 (VBGI-VAE), and 0.45 (VBGI-VAE-M) of the excess iteration cost over the other methods is due to the Gibbs sampling of imputations, denoted by the hatched bars in the figure. We attribute the rest of the excess cost to the fitting of the variational kernel and evaluating the generative model on multiple imputed samples. We also note that for the experiments in this section we used a shared variational model with extended variational conditionals as discussed in Section 3.3 which resulted in about 2.8 times lower per-iteration cost than VGI with the standard-Gibbs conditionals (see Appendix M Figure 22).

In summary, VGI-based methods offer better estimation accuracy over the other methods at (i) a slightly higher per-iteration computational cost and (ii) a potentially slower rate of likelihood improvement in the high-missingness settings.

6. Experiments on Normalising Flows

In this section we estimate normalising flow models from incomplete UCI data using the general-purpose VGI and a flow-specific Monte Carlo EM method. We compare them in terms of model accuracy and computational efficiency.

6.1 Normalising Flow Model

A normalising flow is a probabilistic model that models complex distributions $p_{\theta}(\mathbf{x})$ via a simple base distribution $p(\mathbf{u})$ and an invertible deterministic transformation $T_{\theta} : \mathbf{u} \mapsto \mathbf{x} =$

$T_{\theta}(\mathbf{u})$ (Rezende and Mohamed, 2015). The density of \mathbf{x} can be obtained from the change of variables formula (e.g. Murphy, 2021) as

$$p_{\theta}(\mathbf{x}) = p(\mathbf{u})|\det J_{T_{\theta}}(\mathbf{u})|^{-1}, \quad \text{where } \mathbf{u} = T_{\theta}^{-1}(\mathbf{x}),$$

and $J_{T_{\theta}}$ is the Jacobian matrix of the transformation. In practice, the transformation T_{θ} is often composed of multiple chained transformations $T_{\theta} = T_L \circ \dots \circ T_1$, where we suppress the dependency of the decomposed transformations T_i on θ for notational simplicity. The parameters of each transformation are given by a neural network, thus enabling complex overall transformations T_{θ} . Moreover, both transformation T_{θ} and T_{θ}^{-1} must be differentiable, which allows us to fit the parameters θ using stochastic gradient ascent.

The expressivity of the flow model greatly depends on the choice of the transformations T_i , hence to show that our method can fit expressive models, we specify the statistical model p_{θ} as a rational-quadratic neural spline flow (RQ-NSF, Durkan et al., 2019), which has been shown to be able to model complex distributions efficiently. In RQ-NSF, each transformation T_i is composed of an invertible linear transformation and a monotonic rational-quadratic spline transformation. The spline transformation maps a fixed interval $[-B, B]$ to $[-B, B]$ using a spline function and an identity mapping outside of this range. The rational-quadratic splines are monotonic piece-wise functions with P sections (*bins*) that are characterised by the boundary coordinates (*knots*) and the (positive) derivatives at the knots, where both the knot coordinates and the derivative values are parametrised by a residual neural network.

6.2 Data

We evaluate VGI on a selection of tabular data sets from the UCI machine-learning repository (Dua and Graff, 2017), which are commonly used to evaluate normalising flow models, and follow the pre-processing in (Papamakarios et al., 2017).

POWER. Measurements of electric power consumption in one household with a one-minute sampling rate collected over a period of 47 months. Contains 6 dimensions and ~2M samples.

GAS. A collection of gas sensor measurements in several gas mixtures. Contains 8 dimensions and ~1M samples.

HEPMASS. A collection of measurements from high-energy physics experiments to detect a new particle of unknown mass. Contains 21 dimensions and ~525K samples.

MINIBOONE. Data taken from a MiniBooNE experiment used to distinguish electron neutrinos from background noise. Contains 43 dimensions and ~36K samples.

The missingness is rendered uniformly at random (MCAR) at three levels of missingness from 16.6% to 83.3%. The rest of the data setup is the same as in the previous sections.

6.3 Experimental Settings

We implement the rational-quadratic neural spline flow (RQ-NSF) with coupling following Durkan et al. (2019, Appendix B.1).

In order to match the expressiveness of the target model p_{θ} we parametrise the univariate conditional distributions using flow-like element-wise distributions. We use the standard normal distribution as the base distribution and R sets of linear and rational-quadratic spline transformations, where the parameters of the spline transformations (knot coordinates and derivatives) are given by a partially-shared residual network. The transformations are defined over the same interval as the target flow model, and we use 4 bins and $R = 3$ element-wise transformations.

As before, we use the extended conditionals with a partially-shared model which consists of a shared network and independent per-conditional networks. The shared network takes a completed data-point and computes a shared representation. This representation is then used to compute the element-wise transformation parameters using independent networks for each of the $R \times D$ transformations. The shared network consists of one residual block with 256 hidden features, and outputs a 128-dimensional shared representation. Each element-wise transformation network consists of 2 residual blocks with 32 hidden features.

Like in the experiments in the previous sections, we optimise the variational parameters ϕ using AMSGrad (Reddi et al., 2018). The other VGI hyper-parameters are as follows: $K = 5/10/20$ imputation chains for corresponding missingness of 16.6%/50%/83.3%, $G = 5$ Gibbs update steps, and $M = 1$ sampled missing dimensions in (11).

6.4 Comparison Methods

We compare VGI against Monte Carlo expectation-maximisation (MCEM, Wei and Tanner, 1990) with the flow-specific MCMC sampler of Cannella et al. (2021). MCEM maximises the observed data log-likelihood by iteratively maximising the ELBO

$$\begin{aligned} \hat{\theta}^t &= \arg \max_{\theta} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{p_{\hat{\theta}^{t-1}}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}^i)} \left[\log p_{\theta}(\mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}}^i) - \log p_{\hat{\theta}^{t-1}}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}^i) \right] \\ &= \arg \max_{\theta} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{p_{\hat{\theta}^{t-1}}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}^i)} \left[\log p_{\theta}(\mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}}^i) \right] \\ &\approx \arg \max_{\theta} \frac{1}{N} \sum_{i=1}^N \frac{1}{K} \sum_{k=1}^K \log p_{\theta}(\mathbf{x}_{\text{mis}}^{(i,k)}, \mathbf{x}_{\text{obs}}^i), \quad \text{where } \mathbf{x}_{\text{mis}}^{(i,k)} \sim p_{\hat{\theta}^{t-1}}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}^i). \end{aligned}$$

We obtain the samples from the conditional distributions of missing data $p_{\hat{\theta}^{t-1}}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}^i)$ using projected latent Markov chain Monte Carlo (PLMCMC, Cannella et al., 2021). Rather than generating proposals in the data space, PLMCMC generates Metropolis-Hastings proposals (e.g. Bishop, 2006, Chapter 11.2.2) in the better-behaved space of the base distribution of a normalising flow using simple proposal distributions, and hence provides a more efficient way to sample from any conditional distribution under a joint normalising flow model.

The use of PLMCMC with MCEM (in the next subsections denoted as PLMCMC for simplicity) for the estimation of normalising flows from incomplete data is not new to this paper and has been demonstrated by Cannella et al. (2021), and in our experiments we follow their experimental configuration. During the first set of iterations the imputations are sampled from a standard Gaussian distribution, and afterwards the conditional imputations

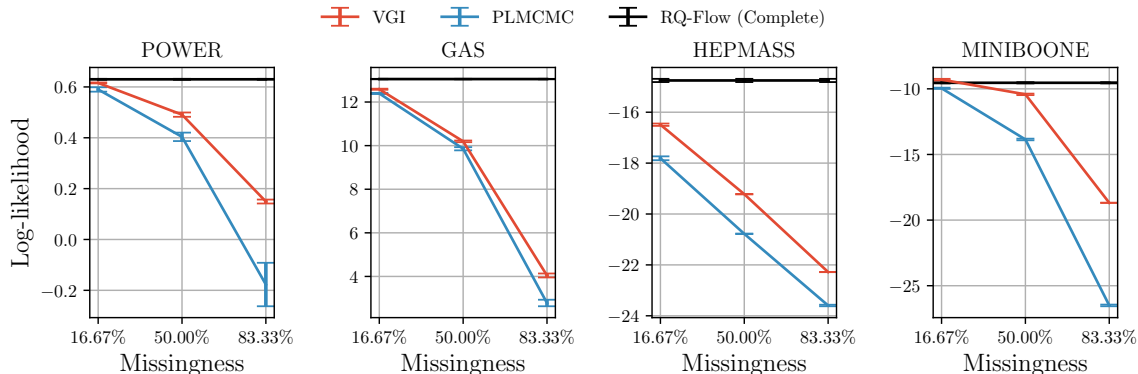


Figure 10: Log-likelihood on complete test UCI data. The error bars show standard error on the mean of five experiment repetitions with different missingness masks and model initialisation. VGI consistently yields a more accurate model p_{θ} than MCEM with the PLMCMC sampler.

are re-sampled using PLMCMC at regular intervals using chains of increasing length (from 200 to 1000 steps), except for the MINIBOONE data for which we could afford running longer chains of up to 3000 steps due to its relatively small size.

After a limited hyper-parameter search we have found that the same proposal distribution as used in the model estimation experiments by Cannella et al. (2021) also performed best in our experiments. The proposal distribution used is a uniform mixture of a conditional normal distribution with mean as the previous state and a standard deviation of 0.01, and an unconditional normal distribution with mean 0 and standard deviation of 1.0.

We observed that the PLMCMC sampler can experience numerical instability due to the inversion of the flow on the GAS and POWER data sets. To stabilise the sampler on the GAS data set we clipped the accepted MCMC proposals to the observed data hypercube, defined by the minimum and maximum observed values for each dimension. A similar fix has been used in some experiments in the original work by Cannella et al. (2021). However, this fix did not help on the POWER data set, and the results in the following sections partly reflect the instability of the sampler.

6.5 Accuracy of the Fitted Normalising Flow Models

We did not observe any over-fitting in the flow model experiments, hence the model parameters $\hat{\theta}$ from the final training iteration were used in the following evaluation for both methods.

We evaluate the accuracy of the fitted models using the log-likelihood on complete test data, shown in Figure 10. Uniformly across the data sets, VGI produces a more accurate fit than MCEM with the PLMCMC sampler. In principle, MCEM should be able to produce model fits that are equally good or better than any variational method due to sampling from the true conditional distributions instead of using an approximation. However, in

practice the performance of MCEM is limited by the performance of the MCMC sampler, which, due to computational constraints, often does not generate exact samples from the true conditional distribution. Even when using PLMCMC as the sampler, which takes advantage of the characteristics of a normalising flow to improve the performance of the sampler, the adverse effect of using finite-length Markov chains is significant. In an attempt to mitigate this pitfall of MCEM, we have empirically investigated the use of persistent imputations chains, similar to VGI, but found that this drastically reduced the acceptance rate of the proposed transitions in the later training epochs and therefore also impaired the accuracy of the estimated statistical model.

In this evaluation we have trained the models for a fixed number of iterations with both methods—we investigate the accuracy-compute trade-off in the following subsection. We thus note that the results could be further improved by running the methods for more iterations (see Appendix M Figure 25 for possible gains when using larger compute budgets), or via other means in stochastic optimisation.

6.6 Computational Aspects

To investigate the computational performance, we plot the test log-likelihood against the training iteration in Figure 11 and further show the average duration of a training iteration for each method. First, we note that the log-likelihood curves of VGI (red) are consistently ahead of PLMCMC (blue). Hence, even without taking the iteration cost into account, VGI improves the statistical model p_{θ} significantly faster than PLMCMC. Moreover, the average per-iteration cost of VGI is also smaller than PLMCMC as shown in the bar plots, making VGI overall more efficient than PLMCMC.

The figure further shows that PLMCMC spends a large fraction of its per-iteration cost on performing MCMC (blue hatched bars). Since the PLMCMC sampler requires the inverse transformation to be efficiently computable, the experiments up to this point used a flow model with coupling layers whose inverse can be computed at a similar cost as the forward transformations. However, many normalising flow models cannot be inverted efficiently. One example of such a model class are the autoregressive flows, where the forward transformation can be computed in constant time with the dimensionality of the model, but the inverse requires d sequential transformations (e.g. Papamakarios et al., 2021, Section 3.1). In Figure 12, we compare the training iteration cost in seconds on a coupling-based (darker colour) and autoregressive (lighter colour) flow models on the MINIBOONE data.²⁵ It can be seen that the cost of PLMCMC on the autoregressive model is significantly larger than for the coupling-based model, which prohibits its application to larger data sets, whilst VGI handles such models much more gracefully.

The results in this section contrast with Section 5.6, where the VGI-based methods were computationally more expensive than the other VAE-specific estimation methods: here we have found that the general-purpose VGI method is more efficient than the existing method for flow model estimation, whilst also producing a more accurate fit.

25. For the autoregressive flow we use the autoregressive rational-quadratic neural spline flow model with the setup of Durkan et al. (2019, Appendix B.1).

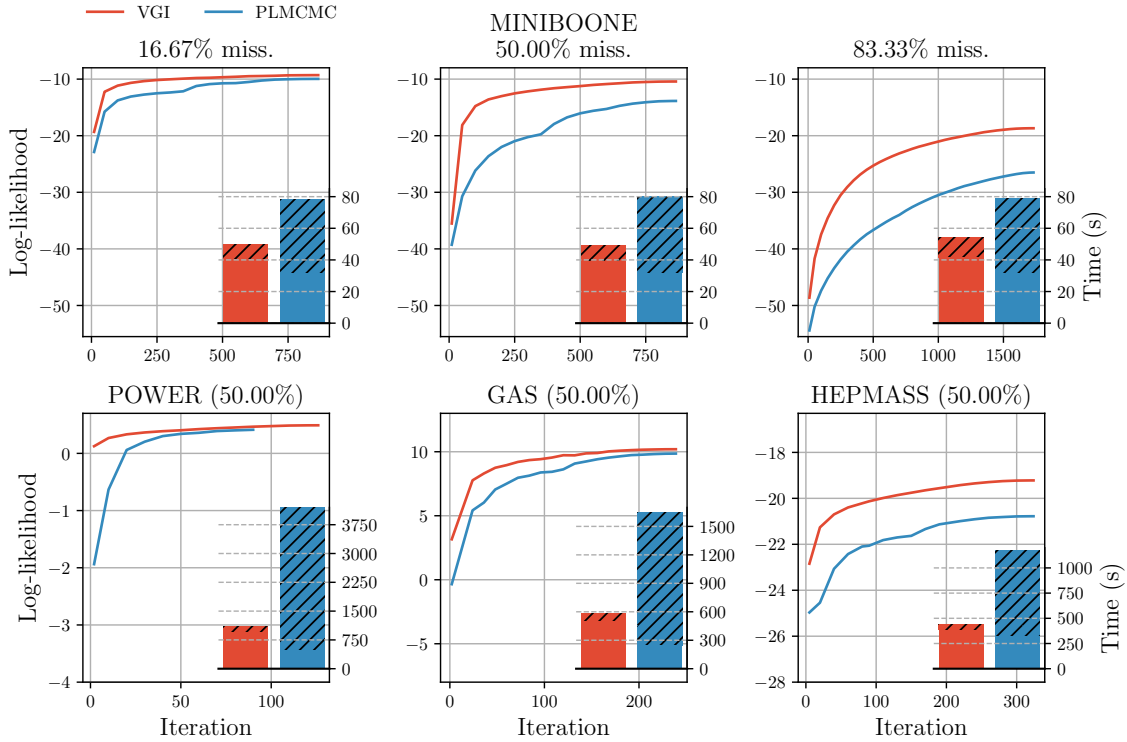


Figure 11: Estimated test log-likelihood on complete data against training iteration and average duration of one training iteration in seconds. The hatched part of the bars indicates the time spent on pseudo-Gibbs sampling in VGI and PLMCMC during MCEM. Top row: results on the MINIBOONE data set for all fractions of missingness. Bottom row: results on the other data sets (POWER, GAS, HEPMASS) for 50% missingness (for the other fractions, see Appendix M Figure 26).

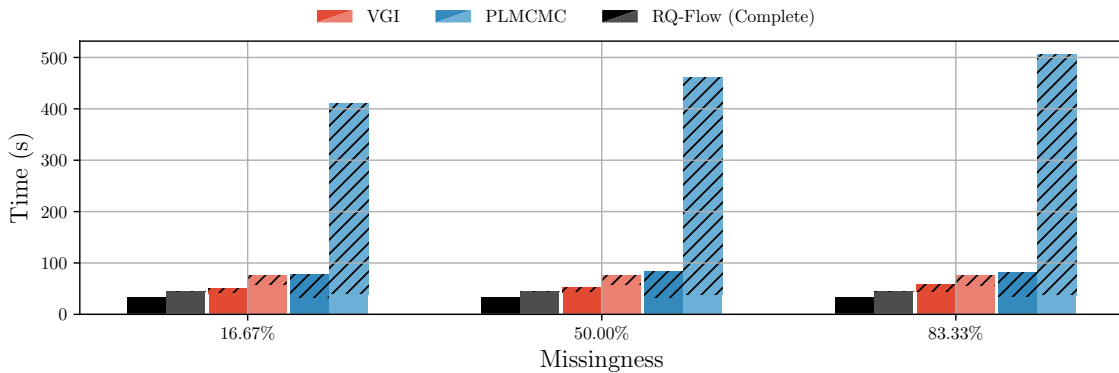


Figure 12: The average duration of a training iteration on the MINIBOONE data set when using coupling-based (darker) and autoregressive (lighter) flow models. VGI easily handles autoregressive models without a significant overhead while PLMCMC scales poorly to such models due to a computationally inefficient inverse.

7. Discussion

In this work we have presented a novel approach to (approximate) maximum-likelihood estimation of statistical models from incomplete data. Our method is applicable to general statistical models and is based on variational inference and Gibbs sampling, both of which have been extensively researched. Hence, the proposed method, termed variational Gibbs inference (VGI), can adopt the diversity of techniques developed for them. We have further introduced a version of VGI for problems where missingness is known to occur in blocks (Section 3.4), for example, in latent-variable models with incomplete data (as demonstrated on a VAE in Section 5).

VGI addresses limitations of (amortised) variational inference (VI) that has strongly limited its applicability to model estimation from incomplete data (see Sections 2.1-2.2): (i) VGI only requires d amortised variational conditionals and thus scales linearly in the dimensionality of the data, gracefully handling the exponential explosion in the number of possible missingness patterns and subsequently the number of conditional variational distributions, and (ii) it enables the specification of a flexible variational family for the conditional distribution of the missing variables using general probabilistic models without restrictions. The parameters of the statistical model p_{θ} and the variational conditionals q_{ϕ_j} are fitted using stochastic gradient ascent, hence it can be used in large-scale data settings that are common in machine learning. The proposed method thus enables efficient learning by drawing on the strengths of amortised variational inference, thanks to property (i), and allows us to well-optimise the ELBO and obtain a good estimate of the target statistical model, thanks to property (ii).

We have validated our method on a toy factor analysis (FA) model against the EM algorithm, which corresponds to the best solution that can be achieved by any variational method, and against a multiple-imputation approach using a popular MICE imputation method in order to evaluate against the two-stage impute-then-fit approaches to statistical model estimation from incomplete data. Our method achieved comparable performance to the classical EM algorithm in a low-dimensional setting and performed similarly to the Monte Carlo EM in the higher-dimensional setting. This suggests that given sufficient compute resources, VGI can achieve good estimates of the model. Moreover, VGI outperformed the impute-then-fit approaches in a higher-dimensional setting even when the (MICE) imputer used the correct family of imputation models relative to the target (FA) model. This suggests that optimising a single variational lower-bound might be a more efficient approach to estimating statistical models from incomplete data than the two-stage impute-then-fit approaches.

We have demonstrated that VGI is particularly well-suited for modern models by successfully fitting important modern statistical models in machine learning, namely VAEs and normalising flows in Sections 5 and 6, respectively.²⁶ In our experiments, VGI achieved better performance than the existing state-of-the-art model-specific estimation methods.

7.1 Future Work

We see a number of interesting future directions following this work.

26. The evaluated models are for continuous random variables, and we focused on them due to their popularity, but the theory presented in Section 3 is general and also holds for discrete distributions.

Due to the generality of the method, it is a candidate to be implemented in probabilistic modelling platforms as a general-purpose inference engine for incomplete data, or as a library in machine learning frameworks. A readily available framework may help improve the current practice of handling missing data, where simple fixes are often preferred over the principled probabilistic methods like VGI.

We also believe that there is area for improvement in the computational efficiency of the method. It is widely known that the mixing of the Markov chain may be inefficient in Gibbs sampling if the variables are highly correlated, which may slow down the rate of improvement in VGI. However, we can use the rich literature on Gibbs sampling to improve VGI: To improve the mixing rate of the Gibbs imputation chains one could use ordered over-relaxation (Neal, 1995) or apply adaptive-scan Gibbs methods to find Gibbs selection probabilities $\pi(j)$ that are better than the uniform ones that we used in this paper (Łatuszyński et al., 2013; Chimisov et al., 2018; Grathwohl et al., 2021), thereby increasing the overall performance of VGI. Alternatively, one could use the (near-)conditional-independencies of the variational conditionals, if such exist, to update some of the dimensions simultaneously, thus breaking the sequential Gibbs update requirement but potentially increasing the mixing rate of the imputation chains (Angelino et al., 2016; Terenin et al., 2020). However, the effect of these changes on the performance of VGI is yet unknown.

Another area of potential improvement is the computational efficiency of evaluating and sampling the variational model. In this work we have implemented the variational conditionals moderately efficiently by using the extended-Gibbs kernel with partially-shared parameters and then computing all conditional distributions simultaneously using optimised matrix operations. However, we usually do not require all conditional distributions for each data-point; only a subset of sampled missing dimensions in the VGI objective in (11) is required. We believe this could be addressed by leveraging the efficient sparse matrix computation libraries built for graph and compressed neural network processing (Bell and Garland, 2008; Han et al., 2016; Huilcen-Baca and Palomino-Valdivia, 2019).

Finally, while our investigation considered statistical model estimation under the ignorable missingness assumption, which is a standard assumption in the missing data field, in practice, the missingness mechanisms are often non-ignorable, and making the ignorability assumption can result in incorrect inferences. Some recent work on VAEs has shown promising results for non-ignorable missingness (Ipsen et al., 2020; Collier et al., 2021; Ma and Zhang, 2021), however no such work has yet been done on normalising flows. Considering the importance of non-ignorable missingness in practice, extending VGI to statistical model estimation from incomplete data that is subject to non-ignorable missingness is an important future direction towards a general-purpose estimation engine for incomplete data.

Acknowledgments

Vaidotas Simkus gratefully acknowledges the support from the University of Edinburgh and the financial support from Huawei through their grant for PhD Studentships in Dialogue Systems and Data Systems. Benjamin Rhodes was supported in part by the EPSRC Centre for Doctoral Training in Data Science, funded by the UK Engineering and Physical Sciences Research Council (grant EP/L016427/1) and the University of Edinburgh. We would also

like to thank the anonymous reviewers for their constructive feedback and valuable suggestions that have helped us improve the paper.

Appendix A. Ignorable Missingness Assumption

In this section we summarise the three classes of missingness as proposed by Rubin (1976) and their effects on statistical model estimation. An important aspect of the incomplete data model in (1) is that the missingness mask \mathbf{m} depends on both observed \mathbf{x}_{obs} and missing \mathbf{x}_{mis} values of the data. This generally means that the estimation of the statistical model $p_{\theta}(\mathbf{x})$ is coupled with the missingness model $p(\mathbf{m} \mid \mathbf{x})$. Rubin (1976) proposed a taxonomy of missing data mechanisms that consists of three levels of simplifying assumptions:

1. Missing not at random (MNAR): no independence assumptions,
2. Missing at random (MAR): $p(\mathbf{m} \mid \mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}}) = p(\mathbf{m} \mid \mathbf{x}_{\text{obs}})$,
3. Missing completely at random (MCAR): $p(\mathbf{m} \mid \mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}}) = p(\mathbf{m})$.

MAR and MCAR assumptions have been widely used in many incomplete data methods since they decouple the model $p_{\theta}(\mathbf{x})$ estimation task from the missingness model $p(\mathbf{m} \mid \mathbf{x})$

$$\begin{aligned} \hat{\theta} &= \arg \max_{\theta} \prod_{i=1}^N \int p_{\theta}(\mathbf{x}_{\text{obs}}^i, \mathbf{x}_{\text{mis}}^i) p(\mathbf{m}^i \mid \mathbf{x}_{\text{obs}}^i) d\mathbf{x}_{\text{mis}}^i \\ &= \arg \max_{\theta} \prod_{i=1}^N p(\mathbf{m}^i \mid \mathbf{x}_{\text{obs}}^i) \prod_{i=1}^N \int p_{\theta}(\mathbf{x}_{\text{obs}}^i, \mathbf{x}_{\text{mis}}^i) d\mathbf{x}_{\text{mis}}^i \\ &= \arg \max_{\theta} \prod_{i=1}^N \int p_{\theta}(\mathbf{x}_{\text{obs}}^i, \mathbf{x}_{\text{mis}}^i) d\mathbf{x}_{\text{mis}}^i, \end{aligned}$$

where the product is due to computing the likelihood for N i.i.d. observed data-points. These assumptions are also known as *ignorable missingness mechanism* because the missingness model can be simply ignored when the only goal is to estimate the model $p_{\theta}(\mathbf{x})$.²⁷

Appendix B. Gibbs Sampling

Gibbs sampling (Geman and Geman, 1984) is a Markov chain Monte Carlo (MCMC, e.g. Barber, 2017, Chapter 27.3) method for (approximately) sampling a joint model $p_{\theta}(\mathbf{x})$ where directly sampling the joint distribution is intractable. The method factorises the distribution using the chain rule $p_{\theta}(\mathbf{x}) = p_{\theta}(x_j \mid \mathbf{x}_{\setminus j}) p_{\theta}(\mathbf{x}_{\setminus j})$, and assumes that the conditional distribution $p_{\theta}(x_j \mid \mathbf{x}_{\setminus j})$ is easy to sample from for any index j . Then, the sampler starts with any random initial sample \mathbf{x}^0 , selects one variable index j to sample, and updates the sample $\mathbf{x}^{t+1} = \{x_j^{t+1}, \mathbf{x}_{\setminus j}^t\}$ with $x_j^{t+1} \sim p_{\theta}(x_j \mid \mathbf{x}_{\setminus j}^t)$. By continuing this procedure, the Gibbs sampler asymptotically in the number of iterations t samples the joint distribution of the model $p_{\theta}(\mathbf{x})$.

27. Assuming the model $p_{\theta}(\mathbf{x})$ is flexible enough, otherwise the estimate can be biased even with MAR missingness (Marlin, 2008, Section 3.4).

The order in which each dimension is updated by the Gibbs sampler is called the “scan”. The two most common scans are: random, where the dimension index j at every iteration t is chosen uniformly at random, and systematic, in which a fixed order of the variables is selected and the sampler repeatedly iterates over the variables in that order. Our VGI method uses the randomised scan as presented in Section 3. However, we feel that it is important to note that neither of these scans are necessarily optimal, and as highlighted in the discussion there are ways to adapt the scan of the Gibbs sampler (Łatuszyński et al., 2013; Chimisov et al., 2018; Grathwohl et al., 2021).

Gibbs sampling methods have been crucial for sampling and learning some important classes of models in physics, such as the Ising model (e.g. Barber, 2017, Section 4.2.5). However, the required conditionals $p_{\theta}(x_j | \mathbf{x}_{\setminus j})$ are typically intractable for most modern statistical models, such as VAEs and normalising flows, and hence, efficient sampling is not tractable either. In VGI, we approximate the Gibbs conditionals with an amortised variational model, which allows us to overcome this issue.

Appendix C. Multiple Imputation by Chained Equations (MICE)

The MICE algorithm starts with random imputations of missing data and then alternates between fitting the univariate conditional distributions $f(x_j | \mathbf{x}_{\setminus j}; \phi_j)$ using the data-points where x_j is observed and imputing the data-points where x_j is missing for $\forall j \in 1 \dots d$. For example, the t -th iteration of the method proceeds as follows

$$\begin{aligned} \hat{\phi}_1^{(t)} &= \arg \max_{\phi_1} \sum_{i \in \text{obs}(\mathcal{D}, 1)} \log f(x_1^{(i,t-1)} | \mathbf{x}_{\setminus 1}^{(i,t-1)}; \hat{\phi}_1^{(t-1)}) \\ x_1^{(i,t)} &\sim f(x_1^{(i,t)} | \mathbf{x}_{\setminus 1}^{(i,t-1)}; \hat{\phi}_1^{(t)}) \text{ for } \forall i \in \text{mis}(\mathcal{D}, 1) \\ &\vdots \\ \hat{\phi}_d^{(t)} &= \arg \max_{\phi_d} \sum_{i \in \text{obs}(\mathcal{D}, d)} \log f(x_d^{(i,t-1)} | \mathbf{x}_{\setminus d}^{(i,t-1)}; \hat{\phi}_d^{(t-1)}) \\ x_d^{(i,t)} &\sim f(x_d^{(i,t)} | \mathbf{x}_{\setminus d}^{(i,t-1)}; \hat{\phi}_d^{(t)}) \text{ for } \forall i \in \text{mis}(\mathcal{D}, d), \end{aligned}$$

where $\text{obs}(\mathcal{D}, j)$ and $\text{mis}(\mathcal{D}, j)$ are respectively the sets of indices of data-points in the data set that are observed and missing in the j -th feature dimension. The procedure is usually repeated several times over all the missing variables to converge to reasonable imputations. Alternatively, if the conditional models $f(x_j | \mathbf{x}_{\setminus j}, \phi_j)$ admit it, the approach can be made Bayesian by specifying a prior over the parameters ϕ_j and then sampling the parameters $\phi_j^{(t)}$ from the posterior distribution $f(\phi_j | \{\mathbf{x}_{<j}^{(t)}, \mathbf{x}_{\geq j}^{(t-1)}\}_{\text{obs}(\mathcal{D}, j)})$ instead of the arg max in the above procedure.

Moreover, as dictated by the multiple-imputation framework, the overall imputation procedure is usually repeated K times, each time starting with random baseline imputations, to obtain K independent imputations of the missing data, which can then be used for many downstream tasks.

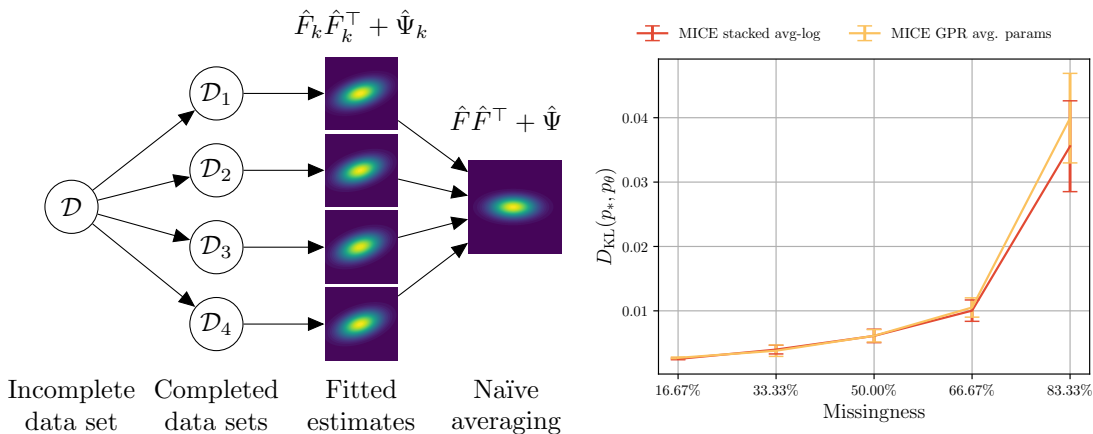


Figure 13: The effect of parameter averaging on a FA model when fitted on imputations from MICE. Left: naïve averaging of parameter estimates from different imputations removes the correlations in the averaged model. Right: comparison of model accuracy after parameter averaging of models trained on independent completed data set (yellow), where generalised Procrustes rotation was applied to the factor loading matrices prior the averaging, and a model trained on several completed data sets stacked together using average log-likelihood (red). Accuracy is measured using KL divergence of the factor analysis model to the ground truth model. Fitting on a stacked completed data set (red) yields equally good parameter estimates as the GPR-based averaging approach.

Appendix D. Combining Multiple Imputations for Statistical Model Estimation

Multiple imputation (MI, Rubin, 1987b) has been the primary tool for the analysis of incomplete data for more than two decades. Many books have covered the standard workflow of principled statistical analysis with multiple imputation and the risks of following invalid procedures (Schafer, 1997; Gelman et al., 2013; van Buuren, 2018).

In this section we first describe the standard analysis workflow with MI and then discuss why it is problematic or cannot be used for statistical models that have non-identifiable parameters. We demonstrate the issue on a classical factor analysis (FA) model, where, unlike for modern deep models (for example, VAEs), there are known methods to resolve the issue. We then explain the procedure we followed in our baseline experiments that use MI-based methods (for example, MICE) and verify on the FA model that it obtains equally good estimates of the statistical model as the standard analysis procedure.

In standard analysis with MI, initially K sets of imputed data are generated using a preferred multiple-imputation method such as the FCS described in Section 2.3. Then, for each imputed data set an independent analysis is performed treating the imputed data as complete and producing K estimates of a desired analysis quantity $\theta^1, \dots, \theta^K$, namely one from each set of imputed data. Following Rubin’s rules (Rubin, 1987b, Result 3.2) the

estimated quantities are then generally combined by averaging

$$\hat{\theta} = \frac{1}{K} \sum_{k=1}^K \theta^k.$$

The estimated quantities $\theta^1, \dots, \theta^K$ can further be used to perform statistical tests and provide confidence intervals of the estimate $\hat{\theta}$.

However, such averaging requires that θ is identifiable. If identifiability is not possible this procedure may result in meaningless estimates. Parametric non-identifiability is a common property of many statistical models starting from classical models, such as mixtures of Gaussians and factor analysis, to modern deep models, such as variational autoencoders and normalising flows, where the non-identifiability stems from the use of neural networks.

We illustrate the issues arising in MI analysis from non-identifiability on a simple example, the factor analysis model, where the factor loading matrix is generally only identifiable up to a multiplication with an orthonormal matrix. Hence, estimating the factor matrix using MI may produce estimates W^1, \dots, W^K that could be written as the product of the true factor matrix F and an orthonormal matrix R^k , $W^k = FR^k$. Naïvely averaging such estimates would give

$$\hat{W} = \frac{1}{K} \sum_{k=1}^K W^k = \frac{1}{K} \sum_{k=1}^K FR^k = F \frac{1}{K} \sum_{k=1}^K R^k,$$

where the average over orthonormal matrices is not a valid orthonormal matrix, and hence the average result is not a valid estimate of the parameters. On the left side of Figure 13 we demonstrate that such naïve averaging approach produces a meaningless model estimate.

For classical models, such as mixtures of Gaussians and factor analysis it is often possible to resolve the indeterminacies of the K estimates before averaging. One solution for the FA model is to transform the factor loading matrices before averaging using the generalised (orthogonal) Procrustes rotation (GPR, Ten Berge, 1977; van Ginkel and Kroonenberg, 2014; Lorenzo-Seva and van Ginkel, 2016). However, no obvious solution exists to resolve the indeterminacy of deep models, which are the main interest of this paper, and hence the standard MI analysis methodology via parameter averaging described above is generally not applicable.

Therefore to avoid averaging, an alternative approach was used in our experiments where MI methods (such as MICE) were used to estimate baseline statistical models. Specifically, we have stacked the imputed data into a single K -times imputed data set and then estimate the model p_{θ} using weighted maximum-likelihood by weighting each imputation by $1/K$

$$\hat{\theta} = \arg \max_{\theta} \frac{1}{N} \sum_{i=1}^N \frac{1}{K} \sum_{k=1}^K \log p_{\theta}(\tilde{\mathbf{x}}_{\text{mis}}^{(i,k)}, \mathbf{x}_{\text{obs}}^i),$$

where $\tilde{\mathbf{x}}_{\text{mis}}^{(i,k)}$ denotes the k -th imputation of the i -th sample.

On the right side of Figure 13 we validate that the stacking approach (red) described above produces equally good parameter estimates as the GPR-based parameter averaging (yellow) for all fractions of missingness. While standard statistical tests may provide invalid

test statistics or confidence intervals if (incorrectly) applied to the stacked data (e.g. van Buuren, 2018, Chapter 5.1) here we are primarily interested in obtaining point estimates of the parameters and hence use the stacked approach since no method (such as the GPR for FA) exists to resolve the parametric indeterminacy of neural networks.

Appendix E. VGI Derivation with Extended-Gibbs Kernel

We here generalise the VGI lower-bound (8) to the case of the extended-Gibbs kernel where the variational conditionals can depend on all variables imputed in the previous iteration. We use this extension to achieve a computationally more efficient method by enabling computation sharing among all variational conditionals (see Section 3.3).

We define the extended-Gibbs kernel by letting the variational conditionals depend on the previous imputation value \tilde{x}_j in dimension j

$$\tilde{\tau}_\phi(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, \tilde{\mathbf{x}}_{\text{mis}}) = \sum_{j \in \text{idx}(\mathbf{m})} \pi(j) q_{\phi_j}(x_j \mid \mathbf{x}_{\text{mis} \setminus j}, \tilde{x}_j, \mathbf{x}_{\text{obs}}) \delta(\mathbf{x}_{\text{mis} \setminus j} - \tilde{\mathbf{x}}_{\text{mis} \setminus j}). \quad (15)$$

Marginalising out $\tilde{\mathbf{x}}_{\text{mis}}$ in the above equation with respect to the previous imputation distribution f^{t-1} yields the imputation distribution after a single Gibbs update

$$\begin{aligned} f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}) &= \int \tilde{\tau}_\phi(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, \tilde{\mathbf{x}}_{\text{mis}}) f^{t-1}(\tilde{\mathbf{x}}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}) d\tilde{\mathbf{x}}_{\text{mis}} \\ &= \sum_{j \in \text{idx}(\mathbf{m})} \pi(j) \int f^{t-1}(\tilde{x}_j \mid \mathbf{x}_{\text{obs}}) \\ &\quad f^{t-1}(\mathbf{x}_{\text{mis} \setminus j} \mid \mathbf{x}_{\text{obs}}, \tilde{x}_j) q_{\phi_j}(x_j \mid \mathbf{x}_{\text{mis} \setminus j}, \tilde{x}_j, \mathbf{x}_{\text{obs}}) d\tilde{x}_j \\ &= \sum_{j \in \text{idx}(\mathbf{m})} \pi(j) \int f^{t-1}(\tilde{x}_j \mid \mathbf{x}_{\text{obs}}) f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, \tilde{x}_j, j) d\tilde{x}_j \\ &= \mathbb{E}_{\pi(j) f^{t-1}(\tilde{x}_j \mid \mathbf{x}_{\text{obs}})} \left[f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, \tilde{x}_j, j) \right] \end{aligned} \quad (16)$$

$$f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, \tilde{x}_j, j) = q_{\phi_j}(x_j \mid \mathbf{x}_{\text{mis} \setminus j}, \tilde{x}_j, \mathbf{x}_{\text{obs}}) f^{t-1}(\mathbf{x}_{\text{mis} \setminus j} \mid \mathbf{x}_{\text{obs}}, \tilde{x}_j). \quad (17)$$

Now we continue from the standard ELBO in (3) and use (16) and (17)

$$\begin{aligned} \mathcal{L}^t(\boldsymbol{\theta}, \boldsymbol{\phi}; \mathbf{x}_{\text{obs}}) &\stackrel{(16)}{=} \mathbb{E}_{\pi(j) f^{t-1}(\tilde{x}_j \mid \mathbf{x}_{\text{obs}}) f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, \tilde{x}_j, j)} \left[\log \frac{p_\theta(\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})}{f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}})} \right] \\ &= \mathbb{E}_{\pi(j) f^{t-1}(\tilde{x}_j \mid \mathbf{x}_{\text{obs}}) f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, \tilde{x}_j, j)} \left[\log \frac{p_\theta(\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})}{f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, \tilde{x}_j, j)} \right. \\ &\quad \left. + \log \frac{f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, \tilde{x}_j, j)}{f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}})} \right] \\ &= \mathbb{E}_{\pi(j) f^{t-1}(\tilde{x}_j \mid \mathbf{x}_{\text{obs}}) f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, \tilde{x}_j, j)} \left[\log \frac{p_\theta(\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})}{f_\phi^t(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}}, \tilde{x}_j, j)} \right] \end{aligned}$$

$$\begin{aligned}
 & + \mathbb{E}_{\boldsymbol{\pi}(j)^{f^{t-1}}(\tilde{x}_j|\mathbf{x}_{\text{obs}})} D_{\text{KL}}(f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, \tilde{x}_j, j) || f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})) \\
 \geq & \mathbb{E}_{\boldsymbol{\pi}(j)^{f^{t-1}}(\tilde{x}_j|\mathbf{x}_{\text{obs}}) f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}}|\mathbf{x}_{\text{obs}}, \tilde{x}_j, j)} \left[\log \frac{p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})}{f_{\boldsymbol{\phi}}^t(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, \tilde{x}_j, j)} \right] \\
 \stackrel{(17)}{=} & \mathbb{E}_{\boldsymbol{\pi}(j)^{f^{t-1}}(\tilde{x}_j|\mathbf{x}_{\text{obs}}) f^{t-1}(\mathbf{x}_{\text{mis}\setminus j}|\mathbf{x}_{\text{obs}}, \tilde{x}_j) q_{\boldsymbol{\phi}_j}(x_j|\mathbf{x}_{\text{mis}\setminus j}, \tilde{x}_j, \mathbf{x}_{\text{obs}})} \left[\right. \\
 & \left. \log \frac{p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})}{q_{\boldsymbol{\phi}_j}(x_j | \mathbf{x}_{\text{mis}\setminus j}, \tilde{x}_j, \mathbf{x}_{\text{obs}})} \right] \\
 & - \mathbb{E}_{\boldsymbol{\pi}(j)^{f^{t-1}}(\tilde{x}_j|\mathbf{x}_{\text{obs}}) f^{t-1}(\mathbf{x}_{\text{mis}\setminus j}|\mathbf{x}_{\text{obs}}, \tilde{x}_j)} \left[\log f^{t-1}(\mathbf{x}_{\text{mis}\setminus j} | \mathbf{x}_{\text{obs}}, \tilde{x}_j) \right] \\
 = & \mathbb{E}_{\boldsymbol{\pi}(j)^{f^{t-1}}(\tilde{x}_j, \mathbf{x}_{\text{mis}\setminus j}|\mathbf{x}_{\text{obs}}) q_{\boldsymbol{\phi}_j}(x_j|\mathbf{x}_{\text{mis}\setminus j}, \tilde{x}_j, \mathbf{x}_{\text{obs}})} \left[\log \frac{p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})}{q_{\boldsymbol{\phi}_j}(x_j | \mathbf{x}_{\text{mis}\setminus j}, \tilde{x}_j, \mathbf{x}_{\text{obs}})} \right] \\
 & - \mathbb{E}_{\boldsymbol{\pi}(j)^{f^{t-1}}(\tilde{x}_j, \mathbf{x}_{\text{mis}\setminus j}|\mathbf{x}_{\text{obs}})} \left[\log f^{t-1}(\mathbf{x}_{\text{mis}\setminus j} | \mathbf{x}_{\text{obs}}, \tilde{x}_j) \right], \tag{18}
 \end{aligned}$$

where the inequality follows from the non-negativity of KL divergence. Dropping the entropy term in the second line of (18) yields the VGI objective with extended conditionals. We note that the VGI ELBO in (18) obtained using the extended-Gibbs kernel is analogous to (8), only the coloured terms are different. To sum up:

- It uses the imputation value \tilde{x}_j from the previous step in the Markov chain, highlighted in red.
- The variational conditionals $q_{\boldsymbol{\phi}_j}$ depend on the previous step in the Markov chain, highlighted in blue.
- The entropy term of the imputation distribution f^{t-1} from (8) became a conditional entropy, highlighted in magenta. Since the term does not include the parameters of interest $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$, it does not need to be computed during optimisation.

The extended-Gibbs conditionals improve the efficiency of the main stage of the VGI algorithm. We note that the extended-Gibbs conditionals need to be separately defined for the variational warm-up stage (line 2 of Algorithm 1 and Algorithm 2), since x_j has no “previous imputation” to condition on. Including the true x_j in the conditioning set for the j -th dimension would cause the variational model to approximate the identity function, which may not be a desirable initialisation. To enable variational model warm-up with the extended conditionals we mask the j -th input to the conditional $q_{\boldsymbol{\phi}_j}$ with a zero.

Furthermore, in the VBGI objective for latent-variable models in (13), we can similarly let the univariate conditionals depend on the previous imputation, which results in the following VBGI objective with extended conditionals

$$\begin{aligned}
 \mathcal{J}_{\text{VBGI}}^t(\boldsymbol{\theta}, \boldsymbol{\phi}; \mathbf{x}_{\text{obs}}) = & \mathbb{E}_{f^{t-1}(\tilde{\mathbf{x}}_{\text{mis}}, \tilde{\mathbf{x}}_z|\mathbf{x}_{\text{obs}})} \left(\right. \\
 & \left. \sum_{j \in \text{idx}(\mathbf{m})} \boldsymbol{\pi}(j) \mathbb{E}_{q_{\boldsymbol{\phi}_j}(x_j|\tilde{\mathbf{x}}_{\text{mis}}, \tilde{\mathbf{x}}_z, \mathbf{x}_{\text{obs}})} \left[\log \frac{p_{\boldsymbol{\theta}}(x_j, \tilde{\mathbf{x}}_{\text{mis}\setminus j}, \mathbf{x}_{\text{obs}} | \tilde{\mathbf{x}}_z) p_{\boldsymbol{\theta}}(\tilde{\mathbf{x}}_z)}{q_{\boldsymbol{\phi}_j}(x_j | \tilde{\mathbf{x}}_{\text{mis}}, \tilde{\mathbf{x}}_z, \mathbf{x}_{\text{obs}})} \right] \right)
 \end{aligned}$$




Figure 14: Partially-shared model of the extended variational conditionals.

$$+ \pi(j = z) \mathbb{E}_{q_{\phi_z}(\mathbf{x}_z | \tilde{\mathbf{x}}_{\text{mis}}, \tilde{\mathbf{x}}_z, \mathbf{x}_{\text{obs}})} \left[\log \frac{p_{\theta}(\tilde{\mathbf{x}}_{\text{mis}}, \mathbf{x}_{\text{obs}} | \mathbf{x}_z) p_{\theta}(\mathbf{x}_z)}{q_{\phi_z}(\mathbf{x}_z | \tilde{\mathbf{x}}_{\text{mis}}, \tilde{\mathbf{x}}_z, \mathbf{x}_{\text{obs}})} \right].$$

It allows us to achieve a similar computational improvement as in the standard VGI case with the extended conditionals.

Appendix F. Properties of the Gibbs Kernel in the Stationary Regime

We show that in the stationary regime the distribution obtained after a Gibbs transition is independent of the updated dimension j . We use a modified kernel from (4), which updates a fixed dimension j

$$\tau_{\phi}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, \tilde{\mathbf{x}}_{\text{mis}}, j) = q_{\phi_j}(x_j | \mathbf{x}_{\text{mis} \setminus j}, \mathbf{x}_{\text{obs}}) \delta(\mathbf{x}_{\text{mis} \setminus j} - \tilde{\mathbf{x}}_{\text{mis} \setminus j}).$$

We now marginalise the kernel with respect to the stationary distribution $q^*(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})$ with conditionals $q_{\phi_j}(x_j | \mathbf{x}_{\text{mis} \setminus j}, \mathbf{x}_{\text{obs}})$

$$\begin{aligned} q^*(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, j) &= \int \tau_{\phi}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}, \tilde{\mathbf{x}}_{\text{mis}}, j) q^*(\tilde{\mathbf{x}}_{\text{mis}} | \mathbf{x}_{\text{obs}}) d\tilde{\mathbf{x}}_{\text{mis}} \\ &= q_{\phi_j}(x_j | \mathbf{x}_{\text{mis} \setminus j}, \mathbf{x}_{\text{obs}}) q^*(\mathbf{x}_{\text{mis} \setminus j} | \mathbf{x}_{\text{obs}}) \\ &= q^*(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}). \end{aligned}$$

Hence, the stationary distribution is invariant to the updated dimension j .

Appendix G. The Variational Inference Network

Using VGI to fit a statistical model p_{θ} first requires us to specify and model d variational distributions. One way to model them is to use d inference networks with parameters ϕ_j , one for each variational distribution. Thus, the neural network for the j -th variational distribution takes an (imputed) data-point $\tilde{\mathbf{x}}_{\setminus j}$ without the target dimension j and outputs the parameters of the distribution $q_{\phi_j}(x_j | \tilde{\mathbf{x}}_{\setminus j})$. Note that all inference networks that parametrise the d Gibbs conditionals have the same input-output dimensionality, hence the computation of all variational distribution parameters can be efficiently parallelised using optimised matrix operations.

Considering that the variational distributions are approximating the conditionals of a joint target distribution we motivate that some parameter sharing in the variational model can be beneficial. One way to incorporate partial parameter sharing is to use a two-part

inference network, where the first part of the network is shared among all conditionals and the second part is independent for each conditional. The shared network can then be used for the different conditionals by fixing the j -th input dimension to zero when computing the intermediate representation for the j -th conditional. In our evaluations we compare an independent model with a partially-shared model and find that they both perform similarly well (see Section 4.7). Importantly, a partially-shared model allows us to scale VGI to higher-dimensional data, by increasing the parameter-efficiency of the variational model.

The VGI objective in (11) requires computing M conditional distributions for each data-point, and the dimensions j^m of those conditionals can differ for each example. One way to compute the parameters of the conditional distributions is to evaluate the inference network for each dimension separately, however, this can require up to d evaluations and hence can be slow in higher dimensions. Alternatively, we could compute all d conditionals in parallel and then select only the required conditionals, however, this would not always be possible, since it would require d -times more memory. To mitigate this, for high-dimensional data we suggest working with the extended-Gibbs variational conditionals where the j -th conditional distribution is allowed to depend on the current imputation in the j -th dimension, thus giving variational distributions of the form $q_{\phi_j}(x_j^t | \mathbf{x}_{\text{obs}}, \tilde{\mathbf{x}}_j^{t-1}, \tilde{\mathbf{x}}_{\text{mis}\setminus j}^{t-1})$. Figure 14 depicts such an extended model. We now need a *single* pass through the shared layers, rather than d passes, which allows for a significant amount of the computation to be shared and hence gives a significant gain in computational efficiency. We empirically investigate the effect of the extended variational conditionals in Section 4.7, where we find no significant detrimental effects of the extended model when estimating higher-dimensional statistical models p_{θ} .

Appendix H. VGI Fine-Tuning Algorithm

In this section we describe the VGI fine-tuning algorithm (Algorithm 6) that is used to evaluate the VGI objective in (9) on held-out data. It is similar to the VGI training algorithm (Algorithm 1) when only the parameters ϕ of the variational model are updated. In our experiments, we use this procedure to select model parameters $\hat{\theta}$ that perform best on the validation data, which is not otherwise possible with amortised variational inference without fine-tuning due to the inference generalisation gap discussed in Section 3.5.

It starts with an incomplete held-out data set and finds the minimum and maximum values in the observed data that defines an imputation acceptance region for the warm-up stage. Then, the incomplete data are imputed using the initial imputation distribution f_0 . In the first iteration of the algorithm, it first “warms up” the incomplete data imputations by performing G_W Gibbs updates, rejecting any imputations outside of the acceptance hypercube. Rejecting imputations that are far from observed data distribution mitigates some of the adverse effects of the inference generalisation gap, but other potentially better rejection strategies, such as the use of a Metropolis-Hastings acceptance step, may also be used. Then, the algorithm continues in the same way as Algorithm 1, fine-tuning the variational parameters ϕ and updating imputations with G Gibbs updates in each iteration, accepting any imputations that are produced by the Gibbs sampler.

Algorithm 6 VGI fine-tuning algorithm

Input: $p_{\hat{\theta}}(\mathbf{x})$, a fitted probabilistic model with parameters $\hat{\theta}$ (kept fixed)
 $q_{\hat{\phi}_j}(x_j | \mathbf{x}_{\setminus j})$ for $j \in \{1 \dots d\}$, var. conditionals fitted on train data with params $\hat{\phi}$
 \mathcal{D}^{val} , incomplete validation data set
 K , number of imputations of each incomplete data-point
 $f^0(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}})$, initial imputation distribution
 G_W , number of Gibbs steps to warm up imputations
 G , number of Gibbs imputation update steps in each epoch
 max_epochs , number of epochs

Output: Validation loss $\hat{\mathcal{J}}_{\text{VGI}}$ and K -times imputed validation data $\mathcal{D}_K^{\text{val}}$

- 1: **Find** the maximum \mathbf{x}_{\max} and minimum \mathbf{x}_{\min} values for each dimension in \mathcal{D}^{val}
- 2: **Create** K -times imputed data $\mathcal{D}_K^{\text{val}}$ using f^0
- 3: $\phi^0 \leftarrow \hat{\phi}$ \triangleright Set the initial parameters to the parameters of the fitted variational model
- 4: **for** t in $[1, max_epochs]$ **do**
- 5: **for** mini-batch \mathcal{B}_K in $\mathcal{D}_K^{\text{val}}$ **do**
- 6: **if** $t = 1$ **then**
- 7: **Update** \mathcal{B}_K with G_W steps of (pseudo-)Gibbs sampler using the variational conditionals, rejecting any Gibbs update values x_j outside $[\mathbf{x}_{\min}(j), \mathbf{x}_{\max}(j)]$
- 8: **else**
- 9: **Update** \mathcal{B}_K with G steps of (pseudo-)Gibbs sampler using Algorithm 5
- 10: **end if**
- 11: **Store** the updated imputations in \mathcal{B}_K for use in the next epoch
- 12: **Compute** $\hat{\mathcal{J}}_{\text{VGI}}^t$ in (11) using Algorithm 3
- 13: $\phi^t = \phi^{t-1} + \alpha_{\phi} \nabla_{\phi} \hat{\mathcal{J}}_{\text{VGI}}^t$ \triangleright Update params of q_{ϕ} with a stochastic gradient step
- 14: **end for**
- 15: **end for**

Appendix I. More on Markov Chain Variational Inference

We review the lower-bound from Markov chain variational inference (MCVI, Salimans et al., 2015), and show that replacing the reverse transition model r with the true reverse operator restores the tighter lower-bound using the marginal of a Markov chain in (3).

The MCVI lower-bound is defined in equation 4 in Salimans et al. (2015) as follows

$$\mathcal{L}_{\text{MCVI}}(\mathbf{x}) = \mathbb{E}_{q(\mathbf{z}_0, \dots, \mathbf{z}^T | \mathbf{x})} \left[\log \frac{p(\mathbf{x}, \mathbf{z}^T)}{q_0(\mathbf{z}^0 | \mathbf{x})} + \sum_{t=1}^T \log \frac{r(\mathbf{z}^{t-1} | \mathbf{x}, \mathbf{z}^t)}{q(\mathbf{z}^t | \mathbf{x}, \mathbf{z}^{t-1})} \right]$$

where T is the total number of transitions in a Markov chain, q and r are transition and reverse transition functions, and \mathbf{z} denotes the latent variable. The joint variational distribution q and the reverse distribution r are assumed to follow a Markov structure $q(\mathbf{z}^0, \dots, \mathbf{z}^T | \mathbf{x}) = q_0(\mathbf{z}^0 | \mathbf{x}) \prod_{t=1}^T q(\mathbf{z}^t | \mathbf{x}, \mathbf{z}^{t-1})$ and $r(\mathbf{z}^0, \dots, \mathbf{z}^{T-1} | \mathbf{x}, \mathbf{z}^T) = \prod_{t=1}^T r(\mathbf{z}^{t-1} | \mathbf{x}, \mathbf{z}^t)$.

We now show that $\mathcal{L}_{\text{MCVI}}$ is a looser lower-bound than the variational ELBO using the marginal of a Markov chain in (3), denoting the marginal distribution $q_T(\mathbf{z}^T | \mathbf{x}) = \int q_0(\mathbf{z}^0 |$

$$\mathbf{x}) \prod_{t=1}^T q(\mathbf{z}^t | \mathbf{x}, \mathbf{z}^{t-1}) d\mathbf{z}^{t-1}.$$

$$\log p_{\theta}(\mathbf{x}) \geq \mathcal{L}_T(\mathbf{x})$$

$$\begin{aligned} \mathcal{L}_T(\mathbf{x}) &= \mathbb{E}_{q_T(\mathbf{z}^T | \mathbf{x})} \left[\log \frac{p(\mathbf{x}, \mathbf{z}^T)}{q_T(\mathbf{z}^T | \mathbf{x})} \right] \\ &= \mathbb{E}_{q_T(\mathbf{z}^T | \mathbf{x})} \left[\log \frac{p(\mathbf{x}, \mathbf{z}^T)}{q_T(\mathbf{z}^T | \mathbf{x})} + \mathbb{E}_{q(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T)} \left[\log \frac{r(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T) q(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T)}{q(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T) r(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T)} \right] \right] \\ &= \mathbb{E}_{q(\mathbf{z}^0, \dots, \mathbf{z}^T | \mathbf{x})} \left[\log \frac{p(\mathbf{x}, \mathbf{z}^T)}{q_T(\mathbf{z}^T | \mathbf{x})} + \log \frac{r(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T)}{q(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T)} \right] \\ &\quad + \mathbb{E}_{q_T(\mathbf{z}^T | \mathbf{x})} \left[D_{\text{KL}}(q(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T) \parallel r(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T)) \right] \\ &= \mathbb{E}_{q(\mathbf{z}^0, \dots, \mathbf{z}^T | \mathbf{x})} \left[\log p(\mathbf{x}, \mathbf{z}^T) + \log \frac{r(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T)}{q(\mathbf{z}^{0\dots T} | \mathbf{x})} \right] \\ &\quad + \mathbb{E}_{q_T(\mathbf{z}^T | \mathbf{x})} \left[D_{\text{KL}}(q(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T) \parallel r(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T)) \right] \\ &= \mathbb{E}_{q(\mathbf{z}^0, \dots, \mathbf{z}^T | \mathbf{x})} \left[\log \frac{p(\mathbf{x}, \mathbf{z}^T)}{q_0(\mathbf{z}^0 | \mathbf{x})} + \log \frac{\prod_{t=1}^T r(\mathbf{z}^{t-1} | \mathbf{x}, \mathbf{z}^t)}{\prod_{t=1}^T q_t(\mathbf{z}^t | \mathbf{x}, \mathbf{z}^{t-1})} \right] \\ &\quad + \mathbb{E}_{q_T(\mathbf{z}^T | \mathbf{x})} \left[D_{\text{KL}}(q(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T) \parallel r(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T)) \right] \\ &= \mathbb{E}_{q(\mathbf{z}^0, \dots, \mathbf{z}^T | \mathbf{x})} \left[\log \frac{p(\mathbf{x}, \mathbf{z}^T)}{q_0(\mathbf{z}^0 | \mathbf{x})} + \sum_{t=1}^T \log \frac{r(\mathbf{z}^{t-1} | \mathbf{x}, \mathbf{z}^t)}{q_t(\mathbf{z}^t | \mathbf{x}, \mathbf{z}^{t-1})} \right] \\ &\quad + \mathbb{E}_{q_T(\mathbf{z}^T | \mathbf{x})} \left[D_{\text{KL}}(q(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T) \parallel r(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T)) \right] \\ &= \mathcal{L}_{\text{MCVI}}(\mathbf{x}) + \mathbb{E}_{q_T(\mathbf{z}^T | \mathbf{x})} \left[D_{\text{KL}}(q(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T) \parallel r(\mathbf{z}^{0\dots T-1} | \mathbf{x}, \mathbf{z}^T)) \right] \end{aligned}$$

$$\log p_{\theta}(\mathbf{x}) \geq \mathcal{L}_T(\mathbf{x}) \geq \mathcal{L}_{\text{MCVI}}(\mathbf{x})$$

Moreover, if we choose the reverse transition function r in MCVI to be the true reverse operator, such that it produces the distribution of the chain at the previous step (e.g. Murray, 2007, Chapter 1.4)

$$r(\mathbf{z}^{t-1} | \mathbf{x}, \mathbf{z}^t) = \frac{q(\mathbf{z}^t | \mathbf{x}, \mathbf{z}^{t-1}) q_{t-1}(\mathbf{z}^{t-1} | \mathbf{x})}{q_t(\mathbf{z}^t | \mathbf{x})},$$

then from MCVI we can recover the tighter lower-bound in (3)

$$\begin{aligned} \mathcal{L}_{\text{MCVI}}(\mathbf{x}) &= \mathbb{E}_{q(\mathbf{z}_0, \dots, \mathbf{z}^T | \mathbf{x})} \left[\log \frac{p(\mathbf{x}, \mathbf{z}^T)}{q(\mathbf{z}^0 | \mathbf{x})} + \sum_{t=1}^T \log \frac{r(\mathbf{z}^{t-1} | \mathbf{x}, \mathbf{z}^t)}{q(\mathbf{z}^t | \mathbf{x}, \mathbf{z}^{t-1})} \right] \\ &= \mathbb{E}_{q(\mathbf{z}_0, \dots, \mathbf{z}^T | \mathbf{x})} \left[\log \frac{p(\mathbf{x}, \mathbf{z}^T)}{q(\mathbf{z}^0 | \mathbf{x})} + \sum_{t=1}^T \log \frac{q(\mathbf{z}^t | \mathbf{x}, \mathbf{z}^{t-1}) q_{t-1}(\mathbf{z}^{t-1} | \mathbf{x})}{q_t(\mathbf{z}^t | \mathbf{x}) q(\mathbf{z}^t | \mathbf{x}, \mathbf{z}^{t-1})} \right] \end{aligned}$$

$$\begin{aligned}
&= \mathbb{E}_{q(\mathbf{z}_0, \dots, \mathbf{z}^T | \mathbf{x})} \left[\log \frac{p(\mathbf{x}, \mathbf{z}^T)}{q(\mathbf{z}^0 | \mathbf{x})} + \sum_{t=1}^T \log \frac{q(\mathbf{z}^t | \mathbf{x}, \mathbf{z}^{t-1}) q_{t-1}(\mathbf{z}^{t-1} | \mathbf{x})}{q_t(\mathbf{z}^t | \mathbf{x}) q(\mathbf{z}^t | \mathbf{x}, \mathbf{z}^{t-1})} \right] \\
&= \mathbb{E}_{q_T(\mathbf{z}^T | \mathbf{x})} \left[\log \frac{p(\mathbf{x}, \mathbf{z}^T)}{q_T(\mathbf{z}^T | \mathbf{x})} \right] \\
&= \mathcal{L}_T(\mathbf{x})
\end{aligned}$$

Hence, having to learn a reverse model r loosens the lower-bound on $\log p_{\theta}(\mathbf{x})$, which makes maximising the likelihood more difficult. The VGI method, on the other hand, does not require learning a model r to reverse the Markov sampling path, and optimises a lower-bound that asymptotically converges to $\mathcal{L}_T(\mathbf{x})$.

Appendix J. Expectation-Maximisation for Factor Analysis with Incomplete Data

We here derive the EM solution for the FA model with incomplete data. From Section 4.1 we know that the FA model is a Gaussian model, with zero-mean standard Gaussian latent variables \mathbf{z} and a linear Gaussian generative model $p(\mathbf{x} | \mathbf{z}) = \mathcal{N}(\mathbf{x}; \mathbf{F}\mathbf{z} + \boldsymbol{\mu}, \boldsymbol{\Psi})$. We assume that \mathbf{x} has observed \mathbf{x}_{obs} and missing \mathbf{x}_{mis} components, and write down the joint distribution

$$p \left(\begin{pmatrix} \mathbf{x}_{\text{obs}} \\ \mathbf{x}_{\text{mis}} \\ \mathbf{z} \end{pmatrix} \right) = \mathcal{N} \left(\begin{pmatrix} \mathbf{x}_{\text{obs}} \\ \mathbf{x}_{\text{mis}} \\ \mathbf{z} \end{pmatrix}; \begin{pmatrix} \boldsymbol{\mu}_{\text{obs}} \\ \boldsymbol{\mu}_{\text{mis}} \\ \mathbf{0} \end{pmatrix}, \begin{bmatrix} \mathbf{F}_{\text{obs}} \mathbf{F}_{\text{obs}}^{\top} + \boldsymbol{\Psi}_{\text{obs}} & \mathbf{F}_{\text{obs}} \mathbf{F}_{\text{mis}}^{\top} & \mathbf{F}_{\text{obs}} \\ \mathbf{F}_{\text{mis}} \mathbf{F}_{\text{obs}}^{\top} & \mathbf{F}_{\text{mis}} \mathbf{F}_{\text{mis}}^{\top} + \boldsymbol{\Psi}_{\text{mis}} & \mathbf{F}_{\text{mis}} \\ \mathbf{F}_{\text{obs}}^{\top} & \mathbf{F}_{\text{mis}}^{\top} & \mathbf{I} \end{bmatrix} \right). \quad (19)$$

Following the standard EM procedure (e.g. Barber, 2017, Chapter 11) we write down the energy, where we remove a constant term due to the prior on \mathbf{z}

$$\begin{aligned}
E(\boldsymbol{\theta}; \boldsymbol{\theta}^{(t-1)}) &= \sum_{i=1}^N \mathbb{E}_{p_{\boldsymbol{\theta}^{(t-1)}}(\mathbf{x}_{\text{mis}}, \mathbf{z} | \mathbf{x}_{\text{obs}}^i)} \left[\log p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{obs}}^i, \mathbf{x}_{\text{mis}} | \mathbf{z}) \right] \\
&= -\frac{1}{2} \sum_{i=1}^N \mathbb{E} \left[d \log(2\pi) + \log \det |\boldsymbol{\Psi}| + (\mathbf{x}^i - \boldsymbol{\mu} - \mathbf{F}\mathbf{z})^{\top} \boldsymbol{\Psi}^{-1} (\mathbf{x}^i - \boldsymbol{\mu} - \mathbf{F}\mathbf{z}) \right] \\
&\propto -\frac{N}{2} \log \det |\boldsymbol{\Psi}| - \frac{1}{2} \sum_{i=1}^N \mathbb{E} \left[(\mathbf{x}^i - \boldsymbol{\mu} - \mathbf{F}\mathbf{z})^{\top} \boldsymbol{\Psi}^{-1} (\mathbf{x}^i - \boldsymbol{\mu} - \mathbf{F}\mathbf{z}) \right], \quad (20)
\end{aligned}$$

where $\boldsymbol{\theta} = (\mathbf{F}, \boldsymbol{\mu}, \boldsymbol{\Psi})$ and $\boldsymbol{\theta}^{(t-1)}$ is the current estimate of $\boldsymbol{\theta}$. The expectation operator $\mathbb{E}[\cdot]$ in (20) means expectation with respect to $p_{\boldsymbol{\theta}^{(t-1)}}(\mathbf{x}_{\text{mis}}, \mathbf{z} | \mathbf{x}_{\text{obs}}^i)$, which is

$$p_{\boldsymbol{\theta}^{(t-1)}} \left(\begin{pmatrix} \mathbf{x}_{\text{mis}} \\ \mathbf{z} \end{pmatrix} \middle| \mathbf{x}_{\text{obs}}^i \right) = \mathcal{N} \left(\begin{pmatrix} \mathbf{x}_{\text{mis}} \\ \mathbf{z} \end{pmatrix}; \begin{pmatrix} \boldsymbol{\mu}_{\text{mis}|\text{obs}} \\ \boldsymbol{\mu}_{\mathbf{z}|\text{obs}} \end{pmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_{\text{mis}|\text{obs}} & \mathbf{C}_{\text{mis}, \mathbf{z}|\text{obs}} \\ \mathbf{C}_{\text{mis}, \mathbf{z}|\text{obs}}^{\top} & \boldsymbol{\Sigma}_{\mathbf{z}|\text{obs}} \end{bmatrix} \right). \quad (21)$$

In the E-step of the algorithm we compute the statistics required in the above from (19) using the rules for conditionals of multivariate Gaussian distributions (Petersen and Pedersen, 2012).

$$\boldsymbol{\Sigma}_{z|\text{obs}} = \mathbf{I} - \mathbf{F}_{\text{obs}}^\top (\mathbf{F}_{\text{obs}} \mathbf{F}_{\text{obs}}^\top + \boldsymbol{\Psi}_{\text{obs}})^{-1} \mathbf{F}_{\text{obs}} = (\mathbf{I} + \mathbf{F}_{\text{obs}}^\top \boldsymbol{\Psi}_{\text{obs}}^{-1} \mathbf{F}_{\text{obs}})^{-1} \quad (22)$$

$$\boldsymbol{\mu}_{z|\text{obs}} = \mathbf{F}_{\text{obs}}^\top (\mathbf{F}_{\text{obs}} \mathbf{F}_{\text{obs}}^\top + \boldsymbol{\Psi}_{\text{obs}})^{-1} (\mathbf{x}_{\text{obs}} - \boldsymbol{\mu}_{\text{obs}}) = \boldsymbol{\Sigma}_{z|\text{obs}} \mathbf{F}_{\text{obs}}^\top \boldsymbol{\Psi}_{\text{obs}}^{-1} (\mathbf{x}_{\text{obs}} - \boldsymbol{\mu}_{\text{obs}}) \quad (23)$$

$$\begin{aligned} \boldsymbol{\Sigma}_{\text{mis}|\text{obs}} &= \mathbf{F}_{\text{mis}} \mathbf{F}_{\text{mis}}^\top + \boldsymbol{\Psi}_{\text{mis}} - \mathbf{F}_{\text{mis}} \mathbf{F}_{\text{obs}}^\top (\mathbf{F}_{\text{obs}} \mathbf{F}_{\text{obs}}^\top + \boldsymbol{\Psi}_{\text{obs}})^{-1} \mathbf{F}_{\text{obs}} \mathbf{F}_{\text{mis}}^\top \\ &= \mathbf{F}_{\text{mis}} \boldsymbol{\Sigma}_{z|\text{obs}} \mathbf{F}_{\text{mis}}^\top + \boldsymbol{\Psi}_{\text{mis}} \end{aligned}$$

$$\boldsymbol{\mu}_{\text{mis}|\text{obs}} = \boldsymbol{\mu}_{\text{mis}} + \mathbf{F}_{\text{mis}} \mathbf{F}_{\text{obs}}^\top (\mathbf{F}_{\text{obs}} \mathbf{F}_{\text{obs}}^\top + \boldsymbol{\Psi}_{\text{obs}})^{-1} (\mathbf{x}_{\text{obs}} - \boldsymbol{\mu}_{\text{obs}}) = \boldsymbol{\mu}_{\text{mis}} + \mathbf{F}_{\text{mis}} \boldsymbol{\mu}_{z|\text{obs}}$$

$$\mathbf{C}_{\text{mis},z|\text{obs}} = \mathbf{F}_{\text{mis}} - \mathbf{F}_{\text{mis}} \mathbf{F}_{\text{obs}}^\top (\mathbf{F}_{\text{obs}} \mathbf{F}_{\text{obs}}^\top + \boldsymbol{\Psi}_{\text{obs}})^{-1} \mathbf{F}_{\text{obs}} = \mathbf{F}_{\text{mis}} \boldsymbol{\Sigma}_{z|\text{obs}},$$

where in (22) we used Woodbury's identity (e.g. Petersen and Pedersen, 2012) and in (23) we used the push-through identity (Henderson and Searle, 1981). Also note that these statistics depend on the observed data $\mathbf{x}_{\text{obs}}^i$ or the observed dimensions \mathbf{m} , we suppressed the index i in the above notation.

In the M-step of the algorithm we maximise the energy with respect to $\boldsymbol{\theta}$. First, to derive the solution for $\boldsymbol{\mu}$ we write down a simplified energy term where \mathbf{z} are marginalised out

$$\begin{aligned} E_{\mathbf{x}}(\boldsymbol{\theta}; \boldsymbol{\theta}^{(t-1)}) &= \sum_{i=1}^N \mathbb{E}_{p_{\boldsymbol{\theta}^{(t-1)}}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}^i)} \left[\log p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{obs}}^i, \mathbf{x}_{\text{mis}}) \right] \\ &= -\frac{1}{2} \sum_{i=1}^N \mathbb{E} \left[d \log(2\pi) + \log \det |\boldsymbol{\Sigma}| + (\mathbf{x}^i - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{x}^i - \boldsymbol{\mu}) \right]. \quad (24) \end{aligned}$$

The expectation operator $\mathbb{E}[\cdot]$ in the above equation depends on $p_{\boldsymbol{\theta}^{(t-1)}}(\mathbf{x}_{\text{mis}} | \mathbf{x}_{\text{obs}}^i)$, which can be directly read from (21) as the marginal of a Gaussian (e.g. Petersen and Pedersen, 2012, Section 8.1.2). Then, taking the partial derivative and setting it to zero we get $\hat{\boldsymbol{\mu}}$

$$\begin{aligned} \frac{\partial E_{\mathbf{x}}(\boldsymbol{\theta}; \boldsymbol{\theta}^{(t-1)})}{\partial \boldsymbol{\mu}} &= -\frac{1}{2} \sum_{i=1}^N \mathbb{E} \left[\frac{\partial}{\partial \boldsymbol{\mu}} \left((\mathbf{x}^i - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{x}^i - \boldsymbol{\mu}) \right) \right] \\ &= -\sum_{i=1}^N \mathbb{E} \left[(\mathbf{x}^i - \boldsymbol{\mu})^\top \right] \boldsymbol{\Sigma}^{-1} \\ &= N \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} - \sum_{i=1}^N \mathbb{E} \left[\mathbf{x}^i \right]^\top \boldsymbol{\Sigma}^{-1} = 0 \\ \implies \hat{\boldsymbol{\mu}} &= \frac{1}{N} \sum_{i=1}^N \hat{\mathbf{x}}^i, \quad (25) \end{aligned}$$

where $\hat{\mathbf{x}}^i = (\mathbf{x}_{\text{obs}}^i, \boldsymbol{\mu}_{\text{mis}|\text{obs}}^i)$.

Now, from (20) we derive the updated $\hat{\mathbf{F}}$

$$\frac{\partial E(\boldsymbol{\theta}; \boldsymbol{\theta}^{(t-1)})}{\partial \mathbf{F}} = -\frac{1}{2} \sum_{i=1}^N \mathbb{E} \left[\frac{\partial}{\partial \mathbf{F}} \left((\mathbf{x}^i - \boldsymbol{\mu} - \mathbf{F} \mathbf{z})^\top \boldsymbol{\Psi}^{-1} (\mathbf{x}^i - \boldsymbol{\mu} - \mathbf{F} \mathbf{z}) \right) \right]$$

$$\begin{aligned}
&= - \sum_{i=1}^N \mathbb{E} \left[(\mathbf{x}^i - \boldsymbol{\mu} - \mathbf{F}\mathbf{z}) \mathbf{z}^\top \boldsymbol{\Psi}^{-1} \right] \\
&= - \sum_{i=1}^N \mathbb{E} \left[(\mathbf{x}^i - \boldsymbol{\mu}) \mathbf{z}^\top \right] \boldsymbol{\Psi}^{-1} + \mathbf{F} \sum_{i=1}^N \mathbb{E} \left[\mathbf{z} \mathbf{z}^\top \right] \boldsymbol{\Psi}^{-1} = 0 \\
\implies \hat{\mathbf{F}} \mathbf{H} &= \mathbf{A} \\
\implies \hat{\mathbf{F}} &= \mathbf{A} \mathbf{H}^{-1}
\end{aligned}$$

where

$$\begin{aligned}
\mathbf{H} &= \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\mathbf{z} \mathbf{z}^\top \right] = \frac{1}{N} \sum_{i=1}^N \left(\boldsymbol{\Sigma}_{z|\text{obs}}^i + \boldsymbol{\mu}_{z|\text{obs}}^i \boldsymbol{\mu}_{z|\text{obs}}^{i\top} \right) \\
\mathbf{A} &= \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[(\mathbf{x}^i - \hat{\boldsymbol{\mu}}) \mathbf{z}^\top \right] = \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\mathbf{x}^i \mathbf{z}^\top \right] - \hat{\boldsymbol{\mu}} \boldsymbol{\mu}_{z|\text{obs}}^{i\top} \\
&= \frac{1}{N} \sum_{i=1}^N \bar{\mathbf{C}}_{\text{mis},z|\text{obs}}^i + \hat{\mathbf{x}}^i \boldsymbol{\mu}_{z|\text{obs}}^{i\top} - \hat{\boldsymbol{\mu}} \boldsymbol{\mu}_{z|\text{obs}}^{i\top} \\
&= \frac{1}{N} \sum_{i=1}^N \bar{\mathbf{C}}_{\text{mis},z|\text{obs}}^i + (\hat{\mathbf{x}}^i - \hat{\boldsymbol{\mu}}) \boldsymbol{\mu}_{z|\text{obs}}^{i\top},
\end{aligned}$$

and $\bar{\mathbf{C}}_{\text{mis},z|\text{obs}}$ is $\mathbf{C}_{\text{mis},z|\text{obs}}$ with 0-rows for the observed dimensions.

Now we derive solution for $\hat{\boldsymbol{\Psi}}$, using $\text{diag}()$ to denote a function that sets off-diagonal elements to zero

$$\begin{aligned}
\frac{\partial E(\boldsymbol{\theta}; \boldsymbol{\theta}^{(t-1)})}{\partial \boldsymbol{\Psi}} &= - \frac{\partial}{\partial \boldsymbol{\Psi}} \frac{N}{2} \log \det |\boldsymbol{\Psi}| \\
&\quad - \text{diag} \left(\sum_{i=1}^N \mathbb{E} \left[\frac{\partial}{\partial \boldsymbol{\Psi}} \left\{ \frac{1}{2} (\mathbf{x}^i - \boldsymbol{\mu} - \mathbf{F}\mathbf{z})^\top \boldsymbol{\Psi}^{-1} (\mathbf{x}^i - \boldsymbol{\mu} - \mathbf{F}\mathbf{z}) \right\} \right] \right) \\
&= - \frac{N}{2} \boldsymbol{\Psi}^{-1} + \text{diag} \left(\sum_{i=1}^N \mathbb{E} \left[\frac{1}{2} \boldsymbol{\Psi}^{-1} (\mathbf{x}^i - \boldsymbol{\mu} - \mathbf{F}\mathbf{z}) (\mathbf{x}^i - \boldsymbol{\mu} - \mathbf{F}\mathbf{z})^\top \boldsymbol{\Psi}^{-1} \right] \right) = 0 \\
\implies \hat{\boldsymbol{\Psi}} &= \text{diag} \left(\frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[(\mathbf{x}^i - \hat{\boldsymbol{\mu}} - \mathbf{F}\mathbf{z}) (\mathbf{x}^i - \hat{\boldsymbol{\mu}} - \mathbf{F}\mathbf{z})^\top \right] \right) \\
&= \text{diag} \left(\frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[(\mathbf{x}^i - \hat{\boldsymbol{\mu}}) (\mathbf{x}^i - \hat{\boldsymbol{\mu}})^\top - 2\mathbf{F}\mathbf{z} (\mathbf{x}^i - \hat{\boldsymbol{\mu}})^\top + \mathbf{F}\mathbf{z} \mathbf{z}^\top \mathbf{F}^\top \right] \right) \\
&= \text{diag} \left(\frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[(\mathbf{x}^i - \hat{\boldsymbol{\mu}}) (\mathbf{x}^i - \hat{\boldsymbol{\mu}})^\top - 2\mathbf{F}\mathbf{A}^\top + \mathbf{F}\mathbf{H}\mathbf{F}^\top \right] \right) \\
&= \text{diag} \left(\frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\mathbf{x}^i \mathbf{x}^{i\top} - 2\mathbf{x}^i \hat{\boldsymbol{\mu}}^\top + \hat{\boldsymbol{\mu}} \hat{\boldsymbol{\mu}}^\top \right] - 2\mathbf{F}\mathbf{A}^\top + \mathbf{F}\mathbf{H}\mathbf{F}^\top \right)
\end{aligned}$$

$$= \text{diag} \left(\mathbf{V} - 2\mathbf{F}\mathbf{A}^\top + \mathbf{F}\mathbf{H}\mathbf{F}^\top \right),$$

with

$$\begin{aligned} \mathbf{V} &= \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\mathbf{x}^i \mathbf{x}^{i\top} \right] - 2\hat{\mathbf{x}}^i \hat{\boldsymbol{\mu}}^\top + \hat{\boldsymbol{\mu}} \hat{\boldsymbol{\mu}}^\top \\ &= \frac{1}{N} \sum_{i=1}^N \bar{\boldsymbol{\Sigma}}_{\text{mis|obs}}^i + \hat{\mathbf{x}}^i \hat{\mathbf{x}}^{i\top} - 2\hat{\mathbf{x}}^i \hat{\boldsymbol{\mu}}^\top + \hat{\boldsymbol{\mu}} \hat{\boldsymbol{\mu}}^\top \\ &= \frac{1}{N} \sum_{i=1}^N \bar{\boldsymbol{\Sigma}}_{\text{mis|obs}}^i + (\hat{\mathbf{x}}^i - \hat{\boldsymbol{\mu}})(\hat{\mathbf{x}}^i - \hat{\boldsymbol{\mu}})^\top, \end{aligned}$$

and where $\bar{\boldsymbol{\Sigma}}_{\text{mis|obs}}^i$ is $\boldsymbol{\Sigma}_{\text{mis|obs}}^i$ with zero rows and columns for the observed dimensions in \mathbf{x}^i . The solutions from the current iteration are then set to $\boldsymbol{\theta}^{(t)} = (\hat{\mathbf{F}}, \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Psi}})$ and used in the next iteration.

Similar EM update rules for FA with incomplete data have been derived by Marlin (2008, Chapter 4.3.2), however our update of the mean vector in (25) uses a tighter lower-bound, by marginalising the latents from the energy in (24), while theirs uses a looser lower-bound that results in the use of the energy in (20). Thus, we can expect that using our update rule should converge faster. Finally, we note that our derivation closely follows the derivation of EM for FA with complete data by Barber (2017, Chapter 21.2.2) with extra terms arising due to incomplete data.

Appendix K. Toy Data Ground Truth

The toy data was generated using a FA model with the following parameters as the ground truth:

$$\mathbf{F} = \begin{bmatrix} -5 & -2 \\ 4 & 0 \\ -3 & -1 \\ -3 & -3 \\ 1 & 5 \\ -1 & 2 \end{bmatrix} \quad \boldsymbol{\mu} = \begin{bmatrix} 3 \\ -1 \\ 0 \\ 2 \\ -1 \\ 0 \end{bmatrix} \quad \text{diag}(\boldsymbol{\Psi}) = \begin{bmatrix} 50.4794 \\ 30.0988 \\ 6.766 \\ 17.3357 \\ 40.9839 \\ 25.1122 \end{bmatrix}.$$

The toy data intentionally has a low signal-to-noise ratio, which made the estimation problem harder.

Appendix L. VAE Model Selection on Incomplete Held-Out Data

The primary goal of our investigation in Section 5 is to evaluate the accuracy of the VAE model $p_{\boldsymbol{\theta}}(\mathbf{x})$ to understand how well it estimates the distribution of the data. Like all other amortised variational methods, the evaluation on held-out data requires fine-tuning of the variational models to the target data due to the inference generalisation gap discussed in Section 3.5. Hence, to perform model selection via validation on held-out data, we

store checkpoints of all parameters (including the generator and the variational models) at predefined intervals during training. Then, we retrospectively fine-tune the variational models to the validation data and select, out of all checkpointed states, the generator parameters that performed best on the fine-tuned validation loss. For VGI we use the validation loss obtained via the fine-tuning algorithm in Appendix H and for the other methods we use their respective loss with the fine-tuned encoder networks.

In Figure 20 of Appendix M we show the corresponding objective curves during fine-tuning of the VGI methods and MVAE,²⁸ where we observe that the fine-tuning of VGI-based methods converge significantly faster than MVAE. Moreover, in Figure 21 of Appendix M we show the validation (dashed) and fine-tuned validation (dash-dotted) learning curves. We can see that the inference generalisation gap affects all evaluated methods (note the significant gap between dashed and dash-dotted curves), hence confirming that fine-tuning of the variational distributions is necessary to perform model selection.

Appendix M. Additional Figures

In this section we show additional figures from the experiments in Sections 4-6.




Figure 15: The accuracy of the fitted statistical models as $D_{\text{KL}}(p_* \parallel p_{\theta})$ on FA-Frey data. MCEM samples imputations from the true conditional $p_{\theta}(\mathbf{x}_{\text{mis}} \mid \mathbf{x}_{\text{obs}})$ using the learnt model p_{θ} and using SGA fits the model on the imputed data. Note that the curves for VGI (red) and MCEM (purple) are close, hence we attribute the performance gap between the EM and VGI to the stochastic optimisation.

²⁸. The fine-tuning loss curves of the other VAE-specific methods were similar to MVAE and hence are not shown.




Figure 16: KL divergence, conditioned on the toy test set, between the univariate variational conditional distributions and the ground truth distribution for each feature dimension $D_{\text{KL}}(q_{\phi_j}(x_j | \mathbf{x}_{\setminus j}) || p_*(x_j | \mathbf{x}_{\setminus j}))$ (blue), and the posterior under the learnt model $D_{\text{KL}}(q_{\phi_j}(x_j | \mathbf{x}_{\setminus j}) || p_\theta(x_j | \mathbf{x}_{\setminus j}))$ (red). The lines show the median conditional KL divergence on the test set.




Figure 17: Wasserstein distance, conditioned on the test set (left: toy data, right: FA-Frey), between the univariate variational conditional distributions and the ground truth distribution $\mathcal{W}(q_{\phi_j}(x_j | \mathbf{x}_{\setminus j}), p_*(x_j | \mathbf{x}_{\setminus j}))$ (blue), and the posterior under the learnt model $\mathcal{W}(q_{\phi_j}(x_j | \mathbf{x}_{\setminus j}), p_\theta(x_j | \mathbf{x}_{\setminus j}))$ (red). The lines show the median Wasserstein distance conditioned on the test set.




Figure 18: The accuracy of the fitted models on toy data, as measured by $D_{\text{KL}}(p_* || p_\theta)$, comparing independent-weight variational models against shared-weight variational models.




Figure 19: Comparison of VGI estimation accuracy on a FA model using a shared standard variational Gibbs model against a shared extended model, as discussed in Section 3.3. Left: estimated KL divergence between the fitted and the ground truth model. Right: average training time per iteration.

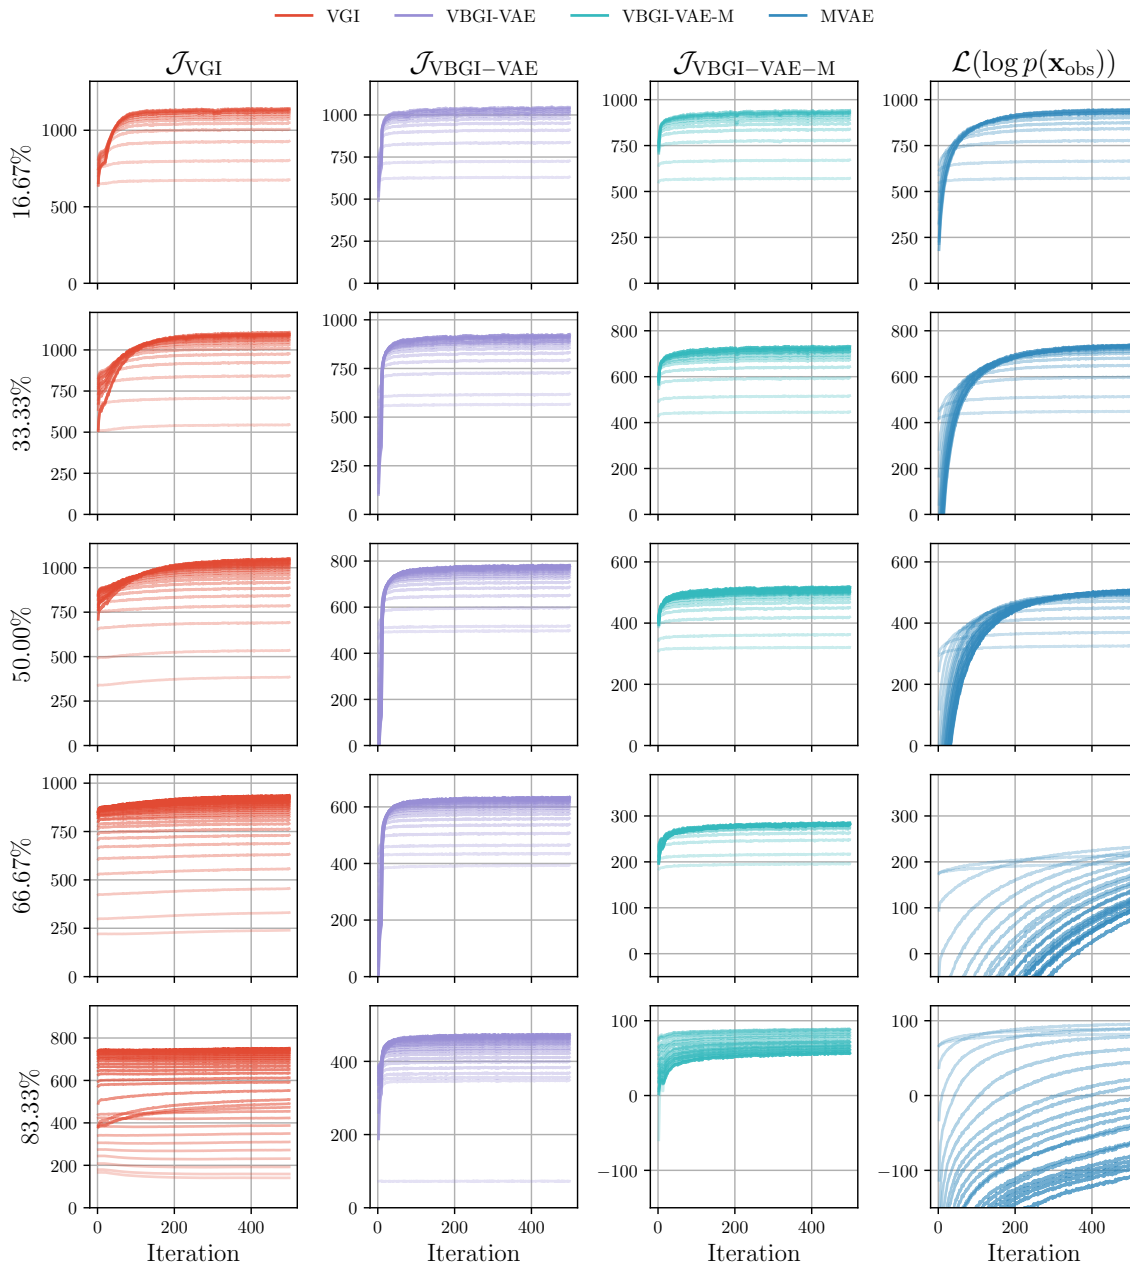


Figure 20: The validation fine-tuning loss curves on VAE-Frey data, where the intensity of the curve represents the training iteration at which the weight snapshot was taken. From an early iteration (less intense) to the last iteration (more intense). The fine-tuning of VGI-based experiments (columns 1-3) is comparatively faster than the VAE-specific methods (right-hand column, the curves of all VAE-specific methods were similar).




Figure 21: The training (solid), validation (dashed), and fine-tuned validation (dash-dotted) loss curves on VAE-Frey data. The gap between the validation and fine-tuned validation curves shows the inference generalisation gap. Note the difference in scale of the objectives, thus they should not be compared directly. Because of the generalisation gap, the validation curves before fine-tuning (dashed red) diverged for VGI when the number of variational models was large. For visualisation purposes, we thus only accepted Gibbs updates in the hypercube defined by the minimum and maximum values in the observed data.




Figure 22: Comparison of VGI estimation on a VAE model using shared standard-Gibbs variational conditionals against shared extended-Gibbs model, as discussed in Section 3.3. Left: estimated test log-likelihood of the fitted model. Right: average training time per iteration.

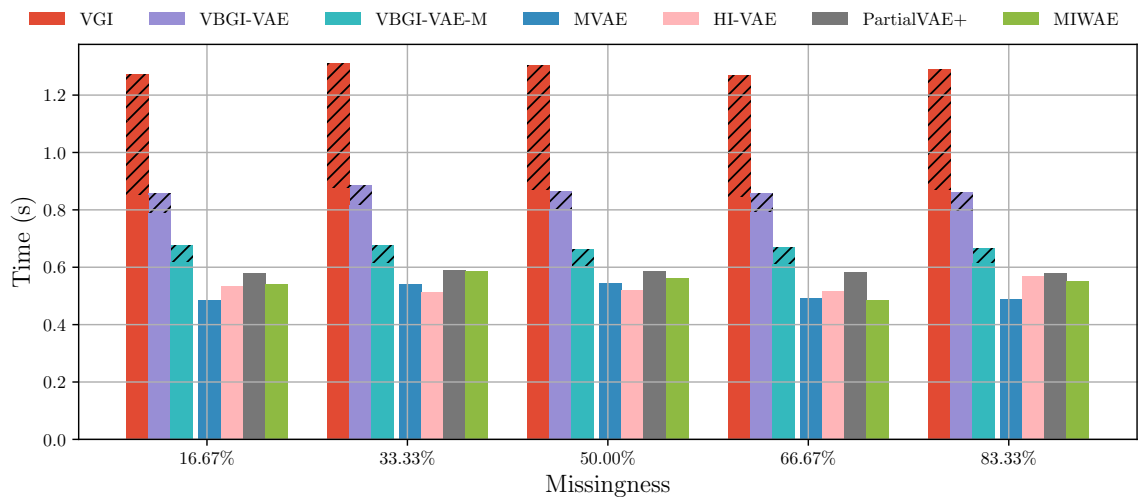


Figure 23: Average time of one training iteration in seconds on VAE-Frey data. The hatched part of the bars indicates the time spent on updating the imputations via pseudo-Gibbs sampling.




Figure 24: Estimated test log-likelihood against the training iteration on complete VAE-Frey data. Estimated using importance sampling as described in Section 5.5, where the variational encoder was refitted to the test data to mitigate the inference generalisation gap.




Figure 25: The percentage gain in test log-likelihood by running the VGI significantly longer. The numbers above the horizontal lines show the test log-likelihood, and the bars show the total training time. Other than the number of maximum iterations, the experimental setup was identical to the main experiments. The results show that the model fits in the main evaluation (Section 6.5) can be further improved by using a larger computational budget.

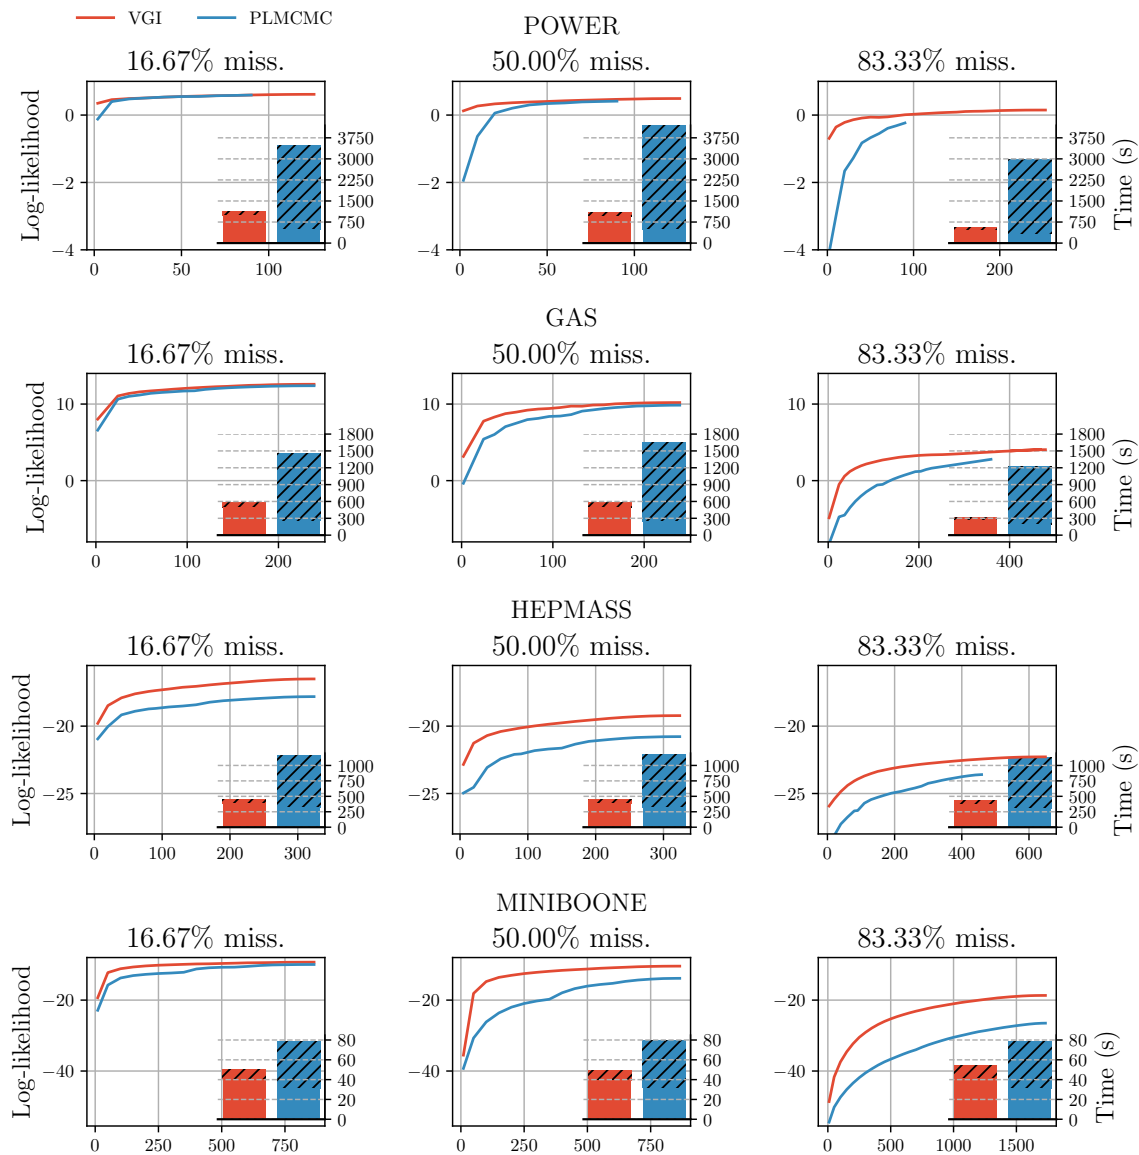


Figure 26: Estimated test log-likelihood on complete data against training iteration and average duration of one training iteration in seconds. The hatched part of the bars indicates the time spent on pseudo-Gibbs sampling in VGI and performing PLMCMC during MCEM. In all of the top row and the third column the PLMCMC log-likelihood curves stop earlier than VGI due to the computational budget, since the average iteration of PLMCMC was significantly more expensive than VGI.

References

- Moloud Abdar, Farhad Pourpanah, Sadiq Hussain, Dana Rezazadegan, Li Liu, Mohammad Ghavamzadeh, Paul Fieguth, Xiaochun Cao, Abbas Khosravi, U. Rajendra Acharya, Vladimir Makarek, and Saeid Nahavandi. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information Fusion*, 76:243–297, 2021. ISSN 1566-2535. doi: 10.1016/j.inffus.2021.05.008.
- Jaan Altosaar, Rajesh Ranganath, and David M. Blei. Proximity Variational Inference. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*. arXiv, May 2017. doi: 10.48550/arXiv.1705.08931.
- Elaine Angelino, Matthew James Johnson, and Ryan P. Adams. Patterns of Scalable Bayesian Inference. *Foundations and Trends in Machine Learning*, 9(2-3):119–247, February 2016. doi: 10.1561/22000000052.
- Barry C. Arnold, Enrique Castillo, and Jose Maria Sarabia. *Conditional Specification of Statistical Models*. Springer-Verlag, New York, 1999. ISBN 978-0-387-77500-5. doi: 10.1007/978-0-387-98135-2.
- David Barber. *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2017. ISBN 978-0-511-80477-9. doi: 10.1017/CBO9780511804779.
- Jonathan W. Bartlett, Shaun R. Seaman, Ian R. White, and James R. Carpenter. Multiple imputation of covariates by fully conditional specification: Accommodating the substantive model. *Statistical Methods in Medical Research*, 24(4):462–487, August 2015. ISSN 0962-2802. doi: 10.1177/0962280214521348.
- Melanie L. Bell, Mallorie Fiero, Nicholas J. Horton, and Chiu-Hsieh Hsu. Handling missing data in RCTs; a review of the top medical journals. *BMC Medical Research Methodology*, 14(1):118, November 2014. ISSN 1471-2288. doi: 10.1186/1471-2288-14-118.
- Nathan Bell and Michael Garland. Efficient Sparse Matrix-Vector Multiplication on CUDA. Technical Report VR-2008-004, NVIDIA Corporation, 2008.
- Christopher M. Bishop. *Machine Learning and Pattern Recognition*. Springer, March 2006. ISBN 0-387-31073-8.
- David M. Blei, Alp Kucukelbir, and Jon D. McAuliffe. Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017. ISSN 1537274X. doi: 10.1080/01621459.2017.1285773.
- Jaap P. L. Brand. *Development, Implementation and Evaluation of Multiple Imputation Strategies for the Statistical Analysis of Incomplete Data Sets*. PhD thesis, Erasmus University Rotterdam, 1999.
- Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. Importance Weighted Autoencoders. In *International Conference on Learning Representations (ICLR)*, San Juan, Puerto Rico, September 2015.

- Chris Cannella, Mohammadreza Soltani, and Vahid Tarokh. Projected Latent Markov Chain Monte Carlo: Conditional Sampling of Normalizing Flows. In *International Conference on Learning Representations (ICLR)*, Austria, June 2021.
- Cyril Chimisov, Krzysztof Latuszynski, and Gareth Roberts. Adapting The Gibbs Sampler. *arXiv preprints*, 2018.
- Mark Collier, Alfredo Nazabal, and Christopher K. I. Williams. VAEs in the Presence of Missing Data. In *Workshop on the Art of Learning with Missing Values (Artemiss) at International Conference on Machine Learning (ICML)*, March 2021.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*, volume 39. Wiley-Interscience, 2006. ISBN 0-471-24195-4. doi: 10.1109/tit.1993.1603955.
- Chris Cremer, Xuechen Li, and David Duvenaud. Inference Suboptimality in Variational Autoencoders. In *International Conference on Machine Learning (ICML)*, May 2018.
- 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. doi: 10.1111/j.2517-6161.1977.tb01600.x.
- Dheeru Dua and Casey Graff. UCI Machine Learning Repository, 2017.
- Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural Spline Flows. In *Advances in Neural Information Processing Systems (NeurIPS)*, December 2019.
- Yarin Gal. *Uncertainty in Deep Learning*. PhD thesis, University of Cambridge, 2016.
- Andrew Gelman, John Carlin, Hal Stern, David Dunson, Aki Vehtari, and Donald B. Rubin. *Bayesian Data Analysis, 3rd Edition*. Chapman & Hall/CRC, 2013.
- Stuart Geman and Donald Geman. Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6):721–741, November 1984. doi: 10.1109/TPAMI.1984.4767596.
- Samuel J. Gershman and Noah D. Goodman. Amortized Inference in Probabilistic Reasoning. In *Annual Meeting of the Cognitive Science Society*, volume 36, 2014.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, Cambridge, MA, USA, 2016.
- Will Grathwohl, Kevin Swersky, Milad Hashemi, David Duvenaud, and Chris J. Maddison. Oops I Took A Gradient: Scalable Sampling for Discrete Distributions. In *International Conference on Machine Learning (ICML)*, volume 139 of *Proceedings of Machine Learning Research*, pages 3831–3841. PMLR, February 2021.
- Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A. Horowitz, and William J. Dally. EIE: Efficient inference engine on compressed deep neural network. *ACM SIGARCH Computer Architecture News*, 44(3):243–254, June 2016. ISSN 0163-5964. doi: 10.1145/3007787.3001163.

- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1026–1034, 2015. ISBN 9781467383912. doi: 10.1109/ICCV.2015.123.
- David Heckerman, David Maxwell Chickering, Christopher Meek, Robert Rounthwaite, and Carl Kadie. Dependency Networks for Inference, Collaborative Filtering, and Data Visualization. *Journal of Machine Learning Research*, 1(Oct):49–75, 2000.
- Harold V. Henderson and Shayle R. Searle. On Deriving the Inverse of a Sum of Matrices. *SIAM Review*, 23(1):53–60, 1981. ISSN 0036-1445. doi: 10.1137/1023004.
- Herwin Alayn Huillcen-Baca and Flor de Luz Palomino-Valdivia. Efficient Sparse Matrix-Vector Multiplication on GPUs using the CSR Format, Pinned Memory and Overlap Data Transfer. In *International Conference on Electronics, Electrical Engineering and Computing (INTERCON)*, pages 1–4, August 2019. doi: 10.1109/INTERCON.2019.8853624.
- Niels Bruun Ipsen, Pierre-Alexandre Mattei, and Jes Frelsen. Not-MIWAE: Deep Generative Modelling with Missing not at Random Data. In *International Conference on Learning Representations (ICLR)*, June 2020.
- Pavel Izmailov, Sharad Vikram, Matthew D. Hoffman, and Andrew Gordon Wilson. What Are Bayesian Neural Network Posteriors Really Like? In *International Conference on Machine Learning (ICML)*, pages 4629–4640. PMLR, July 2021.
- Michael I. Jordan, Zoubin Ghahramani, Tommi S. Jaakkola, and Lawrence K. Saul. An Introduction to Variational Methods for Graphical Models. *Machine Learning*, 37(2): 183–233, 1999. doi: 10.1023/A:1007665907178.
- Diederik P. Kingma and Jimmy Lei Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, December 2014.
- Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations (ICLR)*, December 2013.
- Krzysztof Łatuszyński, Gareth O. Roberts, and Jeffrey S. Rosenthal. Adaptive Gibbs samplers and related MCMC methods. *Annals of Applied Probability*, 23(1):66–98, 2013. ISSN 10505164. doi: 10.1214/11-AAP806.
- Yang Li, Shoaib Akbar, and Junier B. Oliva. Flow Models for Arbitrary Conditional Likelihoods. In *International Conference on Machine Learning (ICML)*, 2020.
- Roderick J. A. Little. Missing-Data Adjustments in Large Surveys. *Journal of Business & Economic Statistics*, 6(3):287–296, 1988. ISSN 0735-0015. doi: 10.2307/1391878.
- Roderick J. A. Little and Donald B. Rubin. *Statistical Analysis with Missing Data: Second Edition*. Wiley-Interscience, 2002. ISBN 0-471-18386-5.

- Urbano Lorenzo-Seva and Joost R. van Ginkel. Multiple Imputation of missing values in exploratory factor analysis of multidimensional scales: Estimating latent trait scores. *Anales de Psicología*, 32(2):596–608, May 2016. ISSN 0212-9728. doi: 10.6018/analesps.32.2.215161.
- Chao Ma and Cheng Zhang. Identifiable Generative Models for Missing Not at Random Data Imputation. In *Neural Information Processing Systems (NeurIPS)*, October 2021.
- 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. In *International Conference on Machine Learning (ICML)*, pages 7483–7504, 2019. ISBN 9781510886988.
- Wesley Maddox, Timur Garipov, Pavel Izmailov, Dmitry Vetrov, and Andrew Gordon Wilson. A Simple Baseline for Bayesian Uncertainty in Deep Learning. In *Neural Information Processing Systems (NeurIPS)*, pages 1–25, Canada, 2019.
- Benjamin M. Marlin. *Missing Data Problems in Machine Learning*. PhD thesis, University of Toronto, 2008.
- Benjamin M. Marlin, Richard S. Zemel, Sam T. Roweis, and Malcolm Slaney. Recommender systems: Missing data and statistical model estimation. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2686–2691, 2011. ISBN 9781577355120. doi: 10.5591/978-1-57735-516-8/IJCAI11-447.
- Pierre-Alexandre Mattei and Jes Frelsen. Refit your Encoder when New Data Comes by. In *Workshop on Bayesian Deep Learning at Neural Information Processing Systems (NeurIPS)*, page 4, Montreal, Canada, 2018.
- Pierre-Alexandre Mattei and Jes Frelsen. MIWAE: Deep Generative Modelling and Imputation of Incomplete Data Sets. In *International Conference on Machine Learning (ICML)*, 2019.
- Kevin P. Murphy. *Probabilistic Machine Learning: An Introduction*. MIT Press, 2021.
- Iain Murray. *Advances in Markov Chain Monte Carlo Methods*. PhD thesis, University College London, 2007.
- Alfredo Nazábal, Pablo M. Olmos, Zoubin Ghahramani, and Isabel Valera. Handling Incomplete Heterogeneous Data using VAEs. *Pattern Recognition*, 107, 2020. ISSN 0031-3203. doi: 10.1016/j.patcog.2020.107501.
- Radford M. Neal. Suppressing Random Walks in Markov Chain Monte Carlo Using Ordered Overrelaxation. Technical Report 9508, Department of Statistics, University of Toronto, Toronto, Ontario, Canada, June 1995.
- Kenneth H. Norwich. *Information, Sensation and Perception*. San Diego: Academic Press, 1993. ISBN 0-12-521890-7.

- Art B. Owen. *Monte Carlo Theory, Methods and Examples*. <https://artowen.su.domains/mc/>, 2013.
- Topi Paananen, Juho Piironen, Paul-Christian Bürkner, and Aki Vehtari. Implicitly adaptive importance sampling. *Statistics and Computing*, 31(2):16, March 2021. ISSN 0960-3174, 1573-1375. doi: 10.1007/s11222-020-09982-2.
- George Papamakarios, Theo Pavlakou, and Iain Murray. Masked Autoregressive Flow for Density Estimation. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017.
- George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing Flows for Probabilistic Modeling and Inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron J. Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. ISSN 2375-0529. doi: 10.1145/2786984.2786995.
- Kaare Brandt Petersen and Michael Syskind Pedersen. The Matrix Cookbook. *Technical University of Denmark*, 16(4):1–16, 2012. ISSN 1611-020X. doi: 10.1017/CBO9780511470943.008.
- Sashank J. Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of Adam and beyond. In *International Conference on Learning Representations (ICLR)*, pages 1–23, 2018.
- Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning (ICML)*, 2015. ISBN 978-1-5108-1058-7.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic Backpropagation and Approximate Inference. In *International Conference on Machine Learning (ICML)*, Beijing, China, 2014.
- Donald B. Rubin. Inference and Missing Data. *Biometrika*, 63(3):581, December 1976. doi: 10.2307/2335739.
- Donald B. Rubin. A Noniterative Sampling/Importance Resampling Alternative to the Data Augmentation Algorithm for Creating a Few Imputations When Fractions of Missing Information are Modest: The SIR Algorithm. *Journal of the American Statistical Association*, 82(398):543–546, 1987a.
- Donald B. Rubin. *Multiple Imputation for Nonresponse in Surveys*. John Wiley & Sons, New York, 1987b. ISBN 0-471-08705-X. doi: 10.2307/3172772.

- Donald B. Rubin. Nested multiple imputation of NMES via partially incompatible MCMC. *Statistica Neerlandica*, 57(1):3–18, 2003. ISSN 00390402. doi: 10.1111/1467-9574.00217.
- Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv:1609.04747 [cs]*, June 2017.
- Tim Salimans, Diederik P. Kingma, and Max Welling. Markov Chain Monte Carlo and Variational Inference: Bridging the Gap. In *International Conference on Machine Learning (ICML)*, pages 1218–1226, 2015.
- Joseph L. Schafer. *Analysis of Incomplete Multivariate Data*. Chapman and Hall, New York, 1997. ISBN 0-412-04061-1.
- James C. Spall. *Introduction to Stochastic Search and Optimization*. Series in Discrete Mathematics and Optimization. Wiley-Interscience, 2003. doi: 10.1007/978-1-4939-1323-7_1.
- Daniel J. Stekhoven and Peter Bühlmann. MissForest - nonparametric missing value imputation for mixed-type data. *Bioinformatics*, 28(1):112–118, January 2012. ISSN 1367-4803, 1460-2059. doi: 10.1093/bioinformatics/btr597.
- Martin Tanner and Wing Hung Wong. The Calculation of Posterior Distributions by Data Augmentation. *Journal of the American Statistical Association*, 82(398):528–540, 1987.
- Jos M. F. Ten Berge. Orthogonal procrustes rotation for two or more matrices. *Psychometrika*, 42(2):267–276, June 1977. ISSN 1860-0980. doi: 10.1007/BF02294053.
- Alexander Terenin, Daniel Simpson, and David Draper. Asynchronous Gibbs Sampling. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- Tijmen Tieleman. Training restricted Boltzmann machines using approximations to the likelihood gradient. In *International Conference on Machine Learning (ICML)*, pages 1064–1071, 2008. ISBN 9781605582054. doi: 10.1145/1390156.1390290.
- Stef van Buuren. *Flexible Imputation of Missing Data*. CRC Press LLC, second edition, 2018. ISBN 978-1-138-58831-8.
- Stef van Buuren and Catharina G. M. Oudshoorn. Multivariate Imputation by Chained Equations. Technical report, TNO Prevention and Health, 2000.
- Stef van Buuren, Jaap P. L. Brand, Catharina G. M. Groothuis-Oudshoorn, and Donald B. Rubin. Fully conditional specification in multivariate imputation. *Journal of Statistical Computation and Simulation*, 76(12):1049–1064, 2006. ISSN 00949655. doi: 10.1080/10629360600810434.
- Joost van Ginkel and Pieter Kroonenberg. Using Generalized Procrustes Analysis for Multiple Imputation in Principal Component Analysis. *Journal of Classification*, 31(2):242–269, 2014.
- Larry Wasserman. *All of Statistics: A Concise Course in Statistical Inference*. Springer Texts in Statistics. Springer, 2005. ISBN 978-1-4419-2322-6.

Greg C. G. Wei and Martin A. Tanner. A Monte Carlo Implementation of the EM Algorithm and the Poor Man's Data Augmentation Algorithms. *Journal of the American Statistical Association*, 85(411):699–704, September 1990. doi: 10.1080/01621459.1990.10474930.

Christopher K. I. Williams, Charlie Nash, and Alfredo Nazábal. Autoencoders and Probabilistic Inference with Missing Data: An Exact Solution for The Factor Analysis Case. *arXiv preprint*, 1801.03851, January 2018.

Laurent Younes. On the convergence of markovian stochastic algorithms with rapidly decreasing ergodicity rates. *Stochastics and Stochastic Reports*, 65(3-4):177–228, February 1999. ISSN 1045-1129. doi: 10.1080/17442509908834179.

Mingtian Zhang, Peter Hayes, and David Barber. Generalization Gap in Amortized Inference. In *Workshop on Bayesian Deep Learning at Neural Information Processing Systems (NeurIPS)*, page 6, 2021.