

LEVERAGING VARIATIONAL AUTOENCODERS FOR MULTIPLE DATA IMPUTATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Missing data persists as a major barrier to data analysis across numerous applications. Recently, deep generative models have been used for imputation of missing data, motivated by their ability to capture highly non-linear and complex relationships in the data. In this work, we investigate the ability of deep models, namely variational autoencoders (VAEs), to account for uncertainty in missing data through multiple imputation strategies. We find that VAEs provide poor empirical coverage of missing data, with underestimation and overconfident imputations, particularly for more extreme missing data values. To overcome this, we employ β -VAEs, which viewed from a generalized Bayes framework, provide robustness to model misspecification. Assigning a good value of β is critical for uncertainty calibration and we demonstrate how this can be achieved using cross-validation. In downstream tasks, we show how multiple imputation with β -VAEs can avoid false discoveries that arise as artefacts of imputation.

1 INTRODUCTION

Missing data persists as a major barrier in large-scale analyses of multivariate data, due to issues like incomplete collection, data availability and low coverage. Early approaches for dealing with missing data tend to reduce the generalizability of results or skew the trends present in the data (Sinharay et al., 2001). These include listwise deletion, where only complete observations are considered, or imputation methods, where the complete data is used to predict plausible values for missing data points. Some of these imputation strategies include substitution by the mean of observed values, stochastic regression techniques and hot deck imputation. Single imputation approaches implicitly assume that the imputation is perfect and thereby fail to account for the uncertainty introduced by the prediction. An attractive solution for this is multiple imputation, which models the uncertainty in the missing values by producing several plausible values for each imputed data point. (Murray, 2018). The complete datasets are then analyzed in downstream tasks, and the results are combined to give estimates and standard errors that acknowledge uncertainty in the missing data.

Recently, deep generative models have become an increasingly popular tool for imputing data, due to their ability to capture highly non-linear relationships and complex dependencies (e.g. Qiu et al., 2020; Camino et al., 2019; Gondara & Wang, 2018; Lewis et al., 2021; García-Laencina et al., 2010; Nelwamondo et al., 2007; Collier et al., 2020; Mattei & Frellsen, 2019; Ipsen et al., 2020; Nazabal et al., 2020; Ma et al., 2018a;b; Ma & Zhang, 2021; Ma et al., 2020). For example, Qiu et al. (2020) use variational autoencoders (VAEs) for imputation of high-dimensional genomic data and find that it performs better than competing methods, such as singular value decomposition (SVD) and K-nearest neighbours (KNN), but focus solely on single imputation. In this work, we aim to investigate the ability of deep models, namely VAEs, to not only reconstruct the missing values but also account for uncertainty through multiple imputation strategies. While expressive and powerful, deep models have been shown to be overconfident even when predictions are incorrect (Szegedy et al., 2013) and also underestimate the variability of out-of-sample test data (Nguyen et al., 2015). In line with these results, we find that VAEs provide poor empirical coverage of the missing data, with underestimation and very overconfident imputations for more extreme missing data values that are far from the mean.

To overcome this, we employ β -VAEs (Higgins et al., 2017), which provide a framework for approximate Bayesian inference of deep generative models under the power likelihood. In statistics,

inference based on the power likelihood has been shown to provide robustness against model misspecification (Bissiri et al., 2016), and thus, in our setting, it is crucial to avoid overfitting and achieve good coverage and well-calibrated uncertainty of the missing data. As is well known in the literature, assigning a good value of β is critical (Holmes & Walker, 2017), and we employ cross-validation to tune β for accurate multiple imputation.

Lastly, we study the implications of multiple imputation in downstream tasks. Failure to account for the uncertainty of the missing data in single imputation can lead to false confidence in downstream analyses, often yielding results that overestimate the significance of relationships between variables. In this paper, we demonstrate how using multiple imputation with β -VAEs yields fewer false positives and more acceptable false discovery rates.

2 BACKGROUND

2.1 VARIATIONAL AUTOENCODERS

Variational autoencoders (Kingma & Welling, 2019) combine graphical models and deep learning. They are made up of two parts, the **encoder** and **decoder**. Firstly, the encoder (also referred to as the inference model) takes an observed data point, $\mathbf{x} \in \mathbb{R}^D$, and computes the posterior distribution, $p_{\theta}(\mathbf{z}|\mathbf{x})$, of the latent variables, $\mathbf{z} \in \mathbb{R}^K$. As the true posterior is intractable in most cases, an approximate model, $q_{\phi}(\mathbf{z}|\mathbf{x})$, is used to approximate the intractable true posterior, $p_{\theta}(\mathbf{z}|\mathbf{x})$, and encode the observed data into the latent variables. The second part is the decoder (also referred to as the generative model) where the latent variables, \mathbf{z} , are used to reconstruct data point, $\hat{\mathbf{x}}$, via the generative model, $p_{\theta}(\mathbf{x}|\mathbf{z})$. The standard choice of distribution for both the inference and generative model is a simple, factorized Gaussian, where the Gaussian mean and variance are parametrized by neural networks, with ϕ and θ containing the weights and biases of the neural networks for the encoder and decoder, respectively. Based on a training data set $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ containing N data points, the neural network parameters ϕ and θ are optimized during training of the VAE by minimizing the reconstruction loss (i.e. minimizing the mean square error between reconstructed values, $\hat{\mathbf{X}} = (\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_N)$, and the observed values, \mathbf{X}) and the latent loss (i.e. minimizing the Kullback-Leibler (KL) divergence between the variational posterior, $q_{\phi}(\mathbf{Z}|\mathbf{X})$, and the standard Gaussian prior, $p(\mathbf{Z})$, with $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_N)$).

From a Bayesian perspective, this is equivalent to approximate variational inference of deep latent variable models, under the generative model $\mathbf{x}_n | \mathbf{z}_n \sim p_{\theta}(\mathbf{x}_n | \mathbf{z}_n)$ with a Gaussian prior on latent variables $\mathbf{z}_n \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. To overcome intractability of the posterior, amortized variational inference (Gershman & Goodman, 2014) is employed, assuming the variational posterior $q_{\phi}(\mathbf{z}_n|\mathbf{x}_n)$ is parametrized by a neural network with ϕ containing the weights and biases. The variational parameters ϕ and generative model parameters θ are optimized by minimizing the KL divergence between the variational posterior $q_{\phi}(\mathbf{Z}|\mathbf{X})$ and the true posterior $p_{\theta}(\mathbf{Z}|\mathbf{X})$, or equivalently maximizing the evidence lower bound (ELBO):

$$\text{ELBO} = \sum_{n=1}^N \mathbb{E}_{\mathbf{z}_n \sim q_{\phi}(\mathbf{z}_n|\mathbf{x}_n)} [\log p_{\theta}(\mathbf{x}_n|\mathbf{z}_n)] - D_{\text{KL}}(q_{\phi}(\mathbf{z}_n|\mathbf{x}_n), p(\mathbf{z}_n)).$$

During training, the ELBO is maximized by backpropagation through the hidden layers of the neural network, randomly sub-sampling the data at each training step and minimizing the loss through stochastic gradient descent. In order to compute the required gradients, we must employ the reparameterization trick (Kingma & Welling; Rezende et al., 2014), which uses a change of variables to obtain independence between the latent noise and ϕ .

2.1.1 β -VAES

An extension of the classic VAE is the β -VAE, which includes a hyperparameter β that enforces a regularization on the latent loss (Higgins et al., 2017). The optimization function is updated as follows:

$$\text{ELBO} = \sum_{n=1}^N \mathbb{E}_{\mathbf{z}_n \sim q_{\phi}(\mathbf{z}_n|\mathbf{x}_n)} [\log p_{\theta}(\mathbf{x}_n|\mathbf{z}_n)] - \beta D_{\text{KL}}(q_{\phi}(\mathbf{z}_n|\mathbf{x}_n), p(\mathbf{z}_n)). \quad (1)$$

While in the machine learning community, β -VAEs are motivated by their improvement in disentangling the latent variables (Chen et al., 2018), we provide an alternative motivation from a statistical perspective. In particular, maximizing the β -VAE bound in equation 1, is equivalent to maximizing:

$$\sum_{n=1}^N \mathbb{E}_{\mathbf{z}_n \sim q_\phi(\mathbf{z}_n | \mathbf{x}_n)} [\log p_\theta(\mathbf{x}_n | \mathbf{z}_n)^{1/\beta}] - D_{\text{KL}}(q_\phi(\mathbf{z}_n | \mathbf{x}_n), p(\mathbf{z}_n)),$$

or minimizing the KL divergence between the variational posterior $q_\phi(\mathbf{Z} | \mathbf{X})$ and the posterior under the power likelihood (see Appendix A.3):

$$p_{\theta, \beta}(\mathbf{Z} | \mathbf{X}) \propto \prod_{n=1}^N p_\theta(\mathbf{x}_n | \mathbf{z}_n)^{1/\beta} p(\mathbf{z}_n).$$

The use of the power likelihood in Bayesian statistics provides frequentist guarantees of posterior consistency in nonparametric models (Walker & Hjort, 2001), while the Bayesian model under the standard updating with $\beta = 1$ may be inconsistent (Barron et al., 1999). Moreover, the power likelihood provides robustness to model misspecification (Bissiri et al., 2016). Given the complex, high-dimensional nature of the deep generative model $p_\theta(\mathbf{x}_n | \mathbf{z}_n)$, this acknowledges and allows for a mismatch between the generative model and the true data generating distribution.

2.2 SINGLE IMPUTATION WITH VAES

Deep generative models have become an increasingly popular tool for imputation of missing data, due to their ability to accommodate highly non-linear relationships and complex dependencies in multivariate data. In this work, we focus on VAEs and, in particular, the approach of Qiu et al. (2020). They first train the VAE using only the subset of complete data to optimize the parameters ϕ and θ . For each data point $n = 1, \dots, N$, \mathbf{x}_n can be split into two parts: $\mathbf{x}_{\text{obs}, n}$ containing the observed features and $\mathbf{x}_{\text{mis}, n}$ containing the missing features, where $\mathbf{X}_{\text{obs}} = (\mathbf{x}_{\text{obs}, 1}, \dots, \mathbf{x}_{\text{obs}, N})$ and $\mathbf{X}_{\text{mis}} = (\mathbf{x}_{\text{mis}, 1}, \dots, \mathbf{x}_{\text{mis}, N})$ list the observed and missing data, respectively. For each data point with missing features, i.e. $\mathbf{x}_n \neq \mathbf{x}_{\text{obs}, n}$, the optimal choice, under the squared error loss, is to impute with the mean under the generative model:

$$\begin{aligned} \hat{\mathbf{x}}_{\text{mis}, n} &= \mathbb{E}[\mathbf{x}_{\text{mis}, n} | \mathbf{x}_{\text{obs}, n}] = \int \mathbf{x}_{\text{mis}, n} p_\theta(\mathbf{x}_{\text{mis}, n} | \mathbf{x}_{\text{obs}, n}) d\mathbf{x}_{\text{mis}, n} \\ &= \int \int \mathbf{x}_{\text{mis}, n} p_\theta(\mathbf{x}_{\text{mis}, n}, \mathbf{z}_n | \mathbf{x}_{\text{obs}, n}) d\mathbf{z}_n d\mathbf{x}_{\text{mis}, n}. \end{aligned}$$

This integral is intractable; thus in Qiu et al. (2020), it is approximated by iteratively computing: 1) the expectation of \mathbf{z}_n (mean of the encoder) given $\hat{\mathbf{x}}_{\text{mis}, n}$ and $\mathbf{x}_{\text{obs}, n}$:

$$\hat{\mathbf{z}}_n = \int \mathbf{z}_n q_\phi(\mathbf{z}_n | \hat{\mathbf{x}}_{\text{mis}, n}, \mathbf{x}_{\text{obs}, n}) d\mathbf{z}_n,$$

and 2) the expectation of $\hat{\mathbf{x}}_{\text{mis}, n}$ (mean of the decoder) given $\hat{\mathbf{z}}_n$:

$$\hat{\mathbf{x}}_{\text{mis}, n} = \int \mathbf{x}_{\text{mis}, n} p_\theta(\mathbf{x}_{\text{mis}, n} | \mathbf{x}_{\text{obs}, n}, \hat{\mathbf{z}}_n) d\mathbf{x}_{\text{mis}, n}.$$

The imputed values are initialized with zero imputation (i.e. mean imputation, after the initial standardization of the data), and these steps are repeated until convergence. Note that when the likelihood factorizes across features (e.g. factorized Gaussian), $p_\theta(\mathbf{x}_{\text{mis}, n} | \mathbf{x}_{\text{obs}, n}, \mathbf{z}_n) = p_\theta(\mathbf{x}_{\text{mis}, n} | \mathbf{z}_n)$. In their paper, Qiu et al. (2020) optimized the model and hyper-parameters through a grid search, claiming that the standard VAE ($\beta = 1$) and training for 250 epochs resulted in the lowest mean absolute error of the imputed values when compared to true values.

2.3 MULTIPLE IMPUTATION

Multiple imputation (see e.g. Little & Rubin, 2019; Murray, 2018; Sinharay et al., 2001) improves upon single imputation by retaining the mean and variance of the overall dataset and accounting for the uncertainty associated with the missing values. It does so by creating M complete datasets with different plausible values for the missing data and then combining inference across all plausible

datasets, e.g. by computing an overall mean and variance estimate for a certain statistic. The general standard is to create five to ten imputed datasets for optimal prediction, but more or less may be required depending on the fraction of data that is missing.

Multiple imputation leverages the missing at random (MAR) case, where variables in the observed data describe the missingness that is present in other variables. We aim to obtain and simulate from the predictive distribution for the missing data given the observed data, i.e. $p(\mathbf{X}_{\text{mis}}|\mathbf{X}_{\text{obs}})$. In particular, we assume that \mathbf{X} follows a distribution, $p(\mathbf{X}|\boldsymbol{\psi})$, where $\boldsymbol{\psi}$ is a collection of all parameters of the model. Then we can write our predictive distribution as:

$$p(\mathbf{X}_{\text{mis}}|\mathbf{X}_{\text{obs}}) = \int p(\mathbf{X}_{\text{mis}}, \boldsymbol{\psi}|\mathbf{X}_{\text{obs}})d\boldsymbol{\psi} = \int p(\mathbf{X}_{\text{mis}}|\mathbf{X}_{\text{obs}}, \boldsymbol{\psi})p(\boldsymbol{\psi}|\mathbf{X}_{\text{obs}})d\boldsymbol{\psi}.$$

To impute the missing data, and thereby simulate one of M plausible datasets, data augmentation (DA) algorithms can be employed. Specifically, DA is a Markov chain method which iteratively samples 1) the parameters $\boldsymbol{\psi}$ from the posterior $p(\boldsymbol{\psi}|\mathbf{X}_{\text{obs}}, \mathbf{X}_{\text{mis}})$ and 2) the missing data \mathbf{X}_{mis} given $\boldsymbol{\psi}$ from $p(\mathbf{X}_{\text{mis}}|\mathbf{X}_{\text{obs}}, \boldsymbol{\psi})$. This ultimately results in sampling from the predictive distribution $p(\mathbf{X}_{\text{mis}}|\mathbf{X}_{\text{obs}})$, producing one of the plausible datasets, denoted as $\mathbf{X}_{\text{mis}}^m = (\mathbf{x}_{\text{mis},1}^m, \dots, \mathbf{x}_{\text{mis},N}^m)$. This procedure is repeated M times to achieve M plausible datasets. Inferences based on these M imputed datasets can be combined via **Rubin’s rules** to compute accurate inference about the entire dataset \mathbf{X} . Note that statistical procedures must be done M times, as there are M datasets.

3 METHODOLOGY

In this work, we generalize single imputation with VAEs in two ways. First, we employ and compare three multiple imputation strategies to account for uncertainty in the missing data. Second, we extend using β -VAEs for improved robustness and uncertainty quantification.

3.1 MULTIPLE IMPUTATION WITH β -VAES

In the case of multiple imputation, the latent variables, \mathbf{Z} , of the β -VAE represent the parameters of our model, previously referred to as $\boldsymbol{\psi}$ in Section 2.3. Therefore, to produce a sample from our target predictive distribution, $p_{\boldsymbol{\theta},\beta}(\mathbf{X}_{\text{mis}}|\mathbf{X}_{\text{obs}})$, we can iteratively sample from the joint distribution $p_{\boldsymbol{\theta},\beta}(\mathbf{X}_{\text{mis}}, \mathbf{Z}|\mathbf{X}_{\text{obs}})$ via a Markov chain Monte Carlo scheme. For β -VAEs, the predictive distribution is constructed from the power likelihood, that is the likelihood of our generative model is raised to the power $1/\beta$ (see Section 2.1.1):

$$\begin{aligned} p_{\boldsymbol{\theta},\beta}(\mathbf{X}_{\text{mis}}|\mathbf{X}_{\text{obs}}) &\propto \int p_{\boldsymbol{\theta}}(\mathbf{X}_{\text{mis}}|\mathbf{X}_{\text{obs}}, \mathbf{Z})^{1/\beta} p_{\boldsymbol{\theta},\beta}(\mathbf{Z}|\mathbf{X}_{\text{obs}})d\mathbf{Z} \\ &= \prod_{n=1}^N \int p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{mis},n}|\mathbf{x}_{\text{obs},n}, \mathbf{z}_n)^{1/\beta} p_{\boldsymbol{\theta},\beta}(\mathbf{z}_n|\mathbf{x}_{\text{obs},n})d\mathbf{z}_n, \end{aligned} \tag{2}$$

where standard VAEs correspond to $\beta = 1$. We note that in the case of the factored Gaussian generative model, the power likelihood $p_{\boldsymbol{\theta}}(\mathbf{x}_{\text{mis},n}|\mathbf{x}_{\text{obs},n}, \mathbf{z}_n)^{1/\beta}$ is simply proportional to a Gaussian with variance rescaled by a factor of β . On the other hand, $p_{\boldsymbol{\theta},\beta}(\mathbf{Z}|\mathbf{X}_{\text{obs}})$ represents the intractable true posterior of the latent variables under the power likelihood given the observed data only.

In the following, we implement and compare three different approaches to sample from our target predictive distribution in equation 2: 1) pseudo-Gibbs (Section 3.1.1), 2) Metropolis-within-Gibbs (Section 3.1.2), and 3) sampling importance resampling (Section 3.1.3). These strategies are proposed in Rezende et al. (2014); Mattei & Frellsen (2018; 2019), respectively, for missing data imputation with deep generative models, and we describe a simple extension based on β -VAEs and the power likelihood. Prior to imputation, we first train the β -VAE, using zero imputation for the missing values, to obtain estimates of generative model parameters $\boldsymbol{\theta}$ and variational parameters ϕ , and thus also an approximation of the true posterior of the latent variables.

3.1.1 PSEUDO-GIBBS

Pseudo-Gibbs sampling was the first strategy developed to generate approximate samples from the predictive distribution in deep generative models (Rezende et al., 2014). In particular, approximate

samples from the joint $p_{\theta,\beta}(\mathbf{X}_{\text{mis}}, \mathbf{Z} \mid \mathbf{X}_{\text{obs}})$ are obtained by iteratively sampling from the encoder and decoder. More specifically, for $s = 1, \dots, S$ iterations and every data point $n \in \{1, \dots, N\}$ with missing features, the pseudo-Gibbs algorithm replaces the expectation steps in the single imputation of Section 2.2 with sampling, as follows: First, Sample \mathbf{z}_n (sample of encoder) given $\mathbf{x}_{\text{mis},n}^{(s-1)}$:

$$\mathbf{z}_n^{(s)} \sim q_{\phi}(\mathbf{z}_n \mid \mathbf{x}_{\text{mis},n}^{(s-1)}, \mathbf{x}_{\text{obs},n}). \quad (3)$$

Next, we sample $\mathbf{x}_{\text{mis},n}^{(s)}$ (sample of decoder) given $\mathbf{z}_n^{(s)}$ based on the power likelihood:

$$\mathbf{x}_{\text{mis},n}^{(s)} \sim p_{\theta,\beta}(\mathbf{x}_{\text{mis},n} \mid \mathbf{x}_{\text{obs},n}, \mathbf{z}_n^{(s)}) \propto p_{\theta}(\mathbf{x}_{\text{mis},n} \mid \mathbf{x}_{\text{obs},n}, \mathbf{z}_n^{(s)})^{1/\beta}.$$

Ideally, in the first step, we would aim to sample from the intractable true posterior of the latent variables. However, if the variational posterior provides a good approximation, the pseudo-Gibbs scheme will produce samples from a distribution close to our target.

3.1.2 METROPOLIS-WITHIN-GIBBS

The pseudo-Gibbs algorithm was improved and extended by Mattei & Frellesen (2018), who derived a Metropolis-within-Gibbs (MWG) sampler that is asymptotically guaranteed to produce samples from the target predictive distribution. This is a simple modification of pseudo-Gibbs that corrects the first step by using the variational posterior as a proposal within a Metropolis-Hastings algorithm. Specifically, in the first step, the sampled value from the encoder in equation 3 represents the proposed value for the latent variables, denoted by \mathbf{z}_n^* , which is then accepted according to the acceptance probability:

$$a(\mathbf{z}_n^{(s-1)} \rightarrow \mathbf{z}_n^*) = \min \left(1, \frac{p_{\theta}(\mathbf{x}_{\text{mis},n}^{(s-1)}, \mathbf{x}_{\text{obs},n} \mid \mathbf{z}_n^*)^{1/\beta} p(\mathbf{z}_n^*)}{p_{\theta}(\mathbf{x}_{\text{mis},n}^{(s-1)}, \mathbf{x}_{\text{obs},n} \mid \mathbf{z}_n^{(s-1)})^{1/\beta} p(\mathbf{z}_n^{(s-1)})} \frac{q_{\phi}(\mathbf{z}_n^{(s-1)} \mid \mathbf{x}_{\text{mis},n}^{(s-1)}, \mathbf{x}_{\text{obs},n})}{q_{\phi}(\mathbf{z}_n^* \mid \mathbf{x}_{\text{mis},n}^{(s-1)}, \mathbf{x}_{\text{obs},n})} \right).$$

Thus, we set:

$$\mathbf{z}_n^{(s)} = \begin{cases} \mathbf{z}_n^* & \text{with prob. } a(\mathbf{z}_n^{(s-1)} \rightarrow \mathbf{z}_n^*) \\ \mathbf{z}_n^{(s-1)} & \text{with prob. } 1 - a(\mathbf{z}_n^{(s-1)} \rightarrow \mathbf{z}_n^*) \end{cases}.$$

If the variational posterior is a perfect approximation of the true posterior, the acceptance probability will be one, and the algorithm reduces to pseudo-Gibbs. In general, MWG acknowledges and corrects for the approximation of the posterior; however, if the variational posterior is far from the true posterior, MWG will suffer from low acceptance rates and slow convergence.

3.1.3 SAMPLING IMPORTANCE RESAMPLING

An alternative to Gibbs is sampling importance resampling (SIR), proposed by Mattei & Frellesen (2019). First, we perform importance sampling using the variational posterior as the importance distribution. In this case, for every data point $n \in \{1, \dots, N\}$ with missing features, we take $s = 1, \dots, S$ samples of the latent variables from our importance distribution:

$$\mathbf{z}_n^{(s)} \sim q_{\phi}(\mathbf{z}_n \mid \mathbf{x}_{\text{mis},n}^{(0)}, \mathbf{x}_{\text{obs},n}),$$

where $\mathbf{x}_{\text{mis},n}^{(0)}$ denotes an initial zero imputation for the missing data. These importance samples $(\mathbf{z}_n^{(s)})$, for $s = 1, \dots, S$, have weights $w_n^{(s)}$ proportional to:

$$\omega_n^{(s)} = \frac{p_{\theta}(\mathbf{x}_{\text{obs},n} \mid \mathbf{z}_n^{(s)})^{1/\beta} p(\mathbf{z}_n^{(s)})}{q_{\phi}(\mathbf{z}_n^{(s)} \mid \mathbf{x}_{\text{mis},n}^{(0)}, \mathbf{x}_{\text{obs},n})},$$

where $w_n^{(s)} = \omega_n^{(s)} / \sum_{s=1}^S \omega_n^{(s)}$ (for further details, see Appendix A.4). Then, for multiple imputation, we obtain M imputations by first sampling (\mathbf{z}_n^m) , for $m = 1, \dots, M$, with replacement from the importance samples $(\mathbf{z}_n^{(s)})$ with probability $w_n^{(s)}$. Next, for each \mathbf{z}_n^m , we impute the missing data by sampling from

$$\mathbf{x}_{\text{mis},n}^m \sim p_{\theta,\beta}(\mathbf{x}_{\text{mis},n} \mid \mathbf{x}_{\text{obs},n}, \mathbf{z}_n^m).$$

In contrast to Gibbs sampling, an advantage of SIR is parallelizability. However, the discrepancy between the variational posterior and true posterior determines the efficiency of the algorithm, and a large discrepancy may result in degeneracy of the weights and require a large number of importance samples (which is required to be exponential in KL divergence between the importance distribution and the target (Chatterjee & Diaconis, 2018)).

3.2 CROSS-VALIDATION TRAINING REGIME

When the generative model and θ match the true data generating distribution exactly, learning is achieved optimally with $\beta = 1$. However, in practice, we have a mismatch and assigning a good value of β becomes critical to achieve robustness and accurate uncertainty quantification. Indeed, if β is set too low, the posterior uncertainty can be underestimated, while if β is set too high, the posterior uncertainty is overestimated. Some directions for assigning a value of β from an information theoretic perspective are provided in Holmes & Walker (2017). Instead, we employ cross-validation to tune β for accurate multiple imputation and coverage of the missing data.

Specifically, the cross-validation approach to tuning β and the number of epochs consists of creating k copies of the data and adding a small proportion of additional missingness in each copy. We then carry out a grid search over the number of epochs and values of β , training k models for each value of β . The final selection is the combination that has acceptable coverage while minimizing the mean absolute error (MAE) over the introduced missing values (averaged across the k models). Once the optimal hyper-parameters for β and epochs are selected, the model is retrained using all of the data. We observed that following this approach results in coverage and MAE on the test set being close to the values estimated through cross-validation. The introduction of additional missing values during cross-validation creates a slight bias towards selecting higher values of β and fewer epochs but this bias is mitigated by increasing the number of copies, k , and thereby reducing the amount of additional missing values in each copy.

3.3 EVALUATING IMPUTATION PERFORMANCE

To evaluate the imputation performance, we consider two quantities: 1) the mean absolute error (MAE) to assess reconstruction accuracy and 2) the empirical coverage (EC) to quantify uncertainty. The MAE compares our imputed values to the ground truth that was originally masked in the complete dataset. Recall that $\hat{\mathbf{X}}_{\text{mis}} = (\hat{\mathbf{x}}_{\text{mis},1}, \dots, \hat{\mathbf{x}}_{\text{mis},N})$ represents the imputed values, while \mathbf{X}_{mis} represents the true (masked) values. The MAE is defined as:

$$\text{MAE} = \frac{1}{N} \sum_{n=1}^N |\hat{\mathbf{x}}_{\text{mis},n} - \mathbf{x}_{\text{mis},n}|, \quad (4)$$

where $|\hat{\mathbf{x}}_{\text{mis},n} - \mathbf{x}_{\text{mis},n}|$ represents the average absolute difference across all missing features for the n th data point. For multiple imputation, the imputed values in equation 4 are averaged across the M imputed datasets, $\hat{\mathbf{x}}_{\text{mis},n} = \frac{1}{M} \sum_{m=1}^M \mathbf{x}_{\text{mis},n}^m$. To evaluate uncertainty in multiple imputation, we first compute $100(1 - \alpha)\%$ confidence intervals (CIs) for each missing value based on the M imputed values. The empirical coverage is then computed as the fraction of times where the true value falls within the predicted interval.

4 RESULTS: GENOMIC DATA IMPUTATION

4.1 PREDICTIONS ARE OVERCONFIDENT AND MISSING DATA VALUES ARE UNDERESTIMATED WITH SINGLE IMPUTATION AND THE STANDARD VAE

We first sought to investigate the impact of single imputation with standard VAEs (Qiu et al., 2020) compared to multiple imputation by our approach through Metropolis-within-Gibbs (MWG), pseudo-Gibbs (PG), and sampling importance resampling (SIR) by evaluating the accuracy of imputed values through MAE and assessing uncertainty through empirical coverage at 95% CIs. We use the same publicly available RNA-sequencing dataset from the Cancer Genome Atlas (TCGA) used in Qiu et al. (2020) in order to benchmark our results against previous work done on genomic data imputation through VAEs. This dataset contains $D = 17,175$ complete features for $N = 667$ glioma patients, comprised of two cancer subtypes, glioblastoma (GBM) and low-grade glioma (LGG). We first simulate missingness in this dataset by masking values with 10% missing completely at random (MCAR) in 20% of samples, and subsequently train the VAE on all remaining observed data, \mathbf{X}_{obs} , with zeros imputed at missing value indices (see Section 3). In order to benchmark against their method, we use the same model and hyper-parameters that were found to be optimal in Qiu et al. (2020), specifically, the standard VAE ($\beta=1$) with 250 training epochs and a

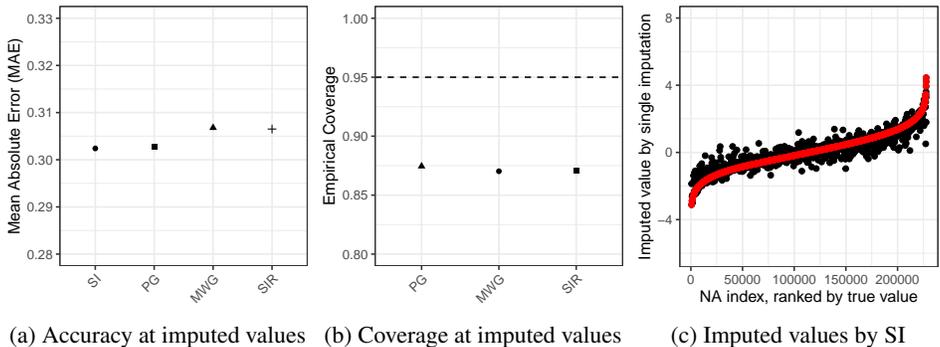


Figure 1: **Standard VAE ($\beta=1$) underestimates the uncertainty in multiple imputation.** Here we report (a) the accuracy of imputed missing data compared to the ground truth by MAE and (b) the fraction of true values that fall within the 95% CIs for the three multiple imputation approaches pseudo-Gibbs (PG), Metropolis-within-Gibbs (MWG) and sampling importance resampling (SIR). Dotted line represents the desired coverage at 0.95. Finally, (c) depicts the imputed values for the missing data by single imputation (SI), ranked by their true values (highlighted in red).

learning rate of 10^{-5} . Once our model is trained, we generate $M = 100$ plausible datasets for each multiple imputation approach and perform single imputation (as described in Section 2.2).

To evaluate imputation of the original masked values, we consider imputation accuracy by MAE and find that the multiple imputation approaches have similar accuracy to single imputation (SI), with pseudo-Gibbs performing slightly better than the other multiple imputation approaches (Figure 1a). Next, we consider the empirical coverage of the masked values based on the 95% CIs computed from the $M = 100$ imputed datasets for all three multiple imputation approaches, and find that the uncertainty is underestimated (Figure 1b). Additionally, the values imputed at masked data points with single imputation are underestimated at more extreme true values (Figure 1c). This is common with neural networks, where model predictions can be poor but still reported with high confidence (Szegedy et al., 2013; Nguyen et al., 2015). Here we see that at more extreme true values, single imputation provides imputed values that are shifted towards the mean. This is likely due to overfitting of the trained model, resulting in overconfident predicted values for masked data points. As Qiu et al. (2020) only used reconstruction accuracy with single imputation to optimize hyperparameters, they were unable to assess uncertainty and overconfidence in the imputations. To overcome this, we explore regularization of the latent space through β -VAEs, optimizing the hyperparameters by considering both reconstruction accuracy and coverage.

4.2 MULTIPLE IMPUTATION WITH β -VAE PROVIDES ACCURATE UNCERTAINTY QUANTIFICATION WHILE STILL RETAINING IMPUTATION ACCURACY

For improved robustness, we employ β -VAEs and the cross-validation scheme described in Section 3.2 to tune β and the number of training epochs, resulting in a value of $\beta = 2$ and 250 training epochs (Supplementary Figure A.1). We then train the β -VAE with these optimal parameters and impute values by single imputation and all three multiple imputation approaches PG, MWG and SIR. This results in good coverage at 95%, with a much lower deviation from the desired coverage than the standard VAE with $\beta = 1$ (Figure 2a-2e, Supplementary Figure A.3, Supplementary Figure A.4). Even with regularization of the latent space, single imputation still results in underestimation at extreme values (Figure 2a). Our multiple imputation by β -VAEs yields good coverage across all missing data, even extreme values, while still retaining comparable accuracy to single imputation (Figure 2f).

4.3 MULTIPLE IMPUTATION REDUCES FALSE POSITIVE RATE IN DOWNSTREAM TASKS.

Lastly, we investigate the impact of all imputation approaches on downstream tasks, namely in identifying discriminating gene sets through logistic regression with the LASSO penalty. In par-

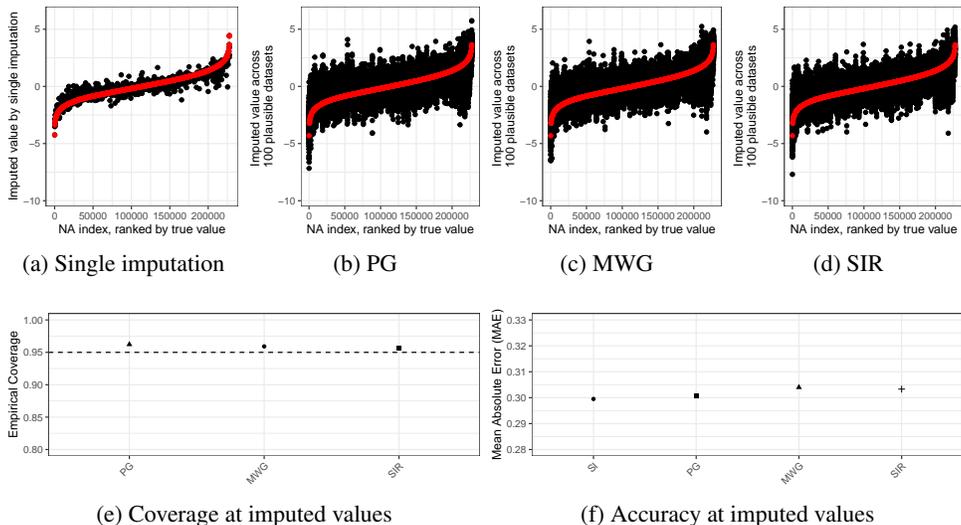


Figure 2: **Multiple imputation with β -VAEs provides proper coverage.** Here we report the (multiple) values imputed for the missing data, ranked by their true values (highlighted in red) for (a) single imputation (SI) and multiple imputation by (b) PG, (c) MWG and (d) SIR. Imputation performance is summarized by (e) the empirical coverage at 95% CIs (dotted line represents the desired coverage at 0.95) and (f) the accuracy at imputed values.

ticular, we run LASSO regression on all imputed datasets to identify the genes which discriminate between the two cancer subtypes, GBM and LGG. This results in one gene set from our ground truth dataset with no missingness (GT), one from single imputation, and 100 discriminating gene sets for each multiple imputation approach, PG, MWG and SIR.

We find that the union across all discriminating gene sets for each multiple imputation approach is much larger than the ground truth set, with the total number of possible non-zero coefficients ranging from 143 to 155, and only 31 discriminating genes in the true dataset (Table 1). When comparing the estimated coefficients from the ground truth data to the (averaged) estimated coefficients based on the imputed data, this results in a slightly higher MAE for multiple imputation approaches (0.066, 0.064 and 0.069 for PG, MWG and SIR, respectively) compared to single imputation (0.053), which also has a set of 31 discriminating genes, although these are not identical to the ground truth set. However, when we inspect the coverage across the multiple imputations, we find that our coverage is close to the desired 95% across PG, MWG and SIR (Table 1, Supplementary Figure A.5).

To identify discriminating gene sets across multiple imputations, we consider two approaches: selecting genes that 1) do not include a coefficient of zero in the 95% CI computed from the 100 imputed datasets, and 2) have an inclusion probability, denoted P_{incl} and defined as the fraction of imputed datasets that the gene has a non-zero LASSO coefficient, greater than a specified threshold. The first approach results in the same set of 12 genes across all three multiple imputation approaches that are all in the true set of non-zero LASSO coefficients (Table 1, Figure 3), giving a false discovery rate (FDR) of 0%. These 12 genes are also contained within the set for single imputation; however, single imputation results in 7 false positives (Figure 3), yielding an FDR of 22.6% (7/31). In the second approach, if we threshold at $P_{\text{incl}} > 0.5$, this results in a final set of 25 discriminating genes for each multiple imputation approach (Table 1, Supplementary Figure A.6). In this case, our gene set contains 2 false positives, yielding an FDR of 8.0% (Figure 3). In summary, we find that multiple imputation with β -VAEs not only provides well-calibrated uncertainty but also results in much more acceptable FDRs in downstream tasks.

Table 1: Performance of different imputation techniques, single imputation (SI) and multiple imputation by PG, MWG and SIR for imputation at missing value indices (first two rows) and downstream impact on LASSO regression (subsequent rows). The final row reports the false discovery rate, based genes with an inclusion probability > 0.5 for multiple imputation.

Metric	SI	PG	MWG	SIR
MAE	0.302	0.301	0.304	0.303
95% CI coverage	N/A	96.2%	95.9%	95.6%
LASSO: MAE	0.053	0.066	0.064	0.069
LASSO: 95% CI coverage	N/A	97.4%	97.2%	96.6%
LASSO: total number of non-zero coefficients	31	155	143	149
LASSO: number of genes without zero in 95% CI	N/A	12	12	12
LASSO: number of genes with $P_{\text{incl}} > 0.5$	N/A	25	25	25
LASSO: False discovery rate	22.6%	8%	8%	8%

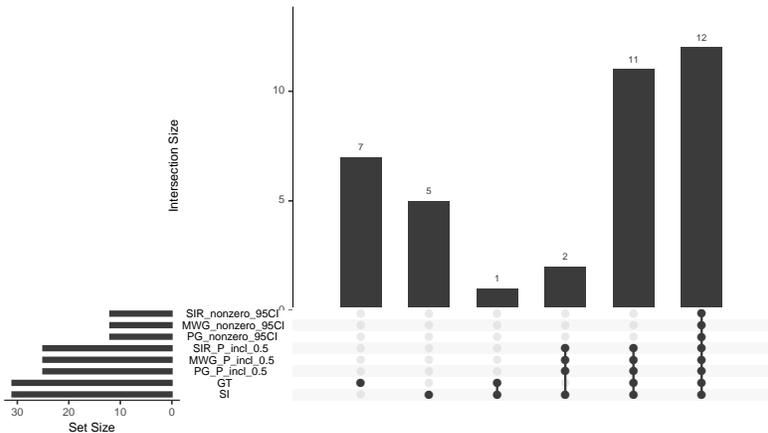


Figure 3: **Upset plot showing overlapping discriminating gene sets by different imputation approaches.** We report the discriminating gene sets by single imputation (SI), ground truth (GT), and all three multiple imputation approaches PG, MWG and SIR with two different inclusion criteria, zero not contained in 95% CI (nonzero_95CI) and inclusion probability, $P_{\text{incl}} > 0.5$ (P_incl.0.5).

5 DISCUSSION

We describe a deep learning framework for multiple imputation using β -VAEs. We propose and compare three multiple imputation methods and develop a new training regime, which uses all observed data to tune hyperparameters by assessing accuracy as well as empirical coverage. Our approach captures the complex, non-linear relationships present in high-dimensional genomic data, imputing values with high accuracy while retaining good coverage. Previous work (Qiu et al., 2020) employed standard VAEs for genomic data imputation by single imputation, resulting in inaccurate and overconfident imputations at extreme missing values. Finally, we investigate the impact of these different imputation approaches on downstream tasks, namely discriminating gene sets identified by logistic regression with the LASSO penalty. We find that multiple imputation through β -VAEs identifies genes that discriminate between the two cancer subtypes with lower false discovery rates than previous methods. All three multiple imputation approaches perform similarly in terms of accuracy and coverage, however SIR is preferable as it is much more computationally efficient.

Future work will investigate missing not at random settings (Ipsen et al., 2020; Collier et al., 2020) and mixed data (Ma et al., 2020). In addition, extensions using ensembles of deep generative models may improve robustness and calibration. Such ensembles can be built from simple approaches, such as training with multiple initializations (Lakshminarayanan et al., 2017), composing models across different epochs (Huang et al., 2017), or Monte Carlo dropout (Gal & Ghahramani, 2016), to more advanced approaches, such as Bayesian methods (Daxberger & Hernández-Lobato, 2019).

REFERENCES

- Andrew Barron, Mark J Schervish, and Larry Wasserman. The consistency of posterior distributions in nonparametric problems. *The Annals of Statistics*, 27(2):536–561, 1999.
- Pier Giovanni Bissiri, Chris C Holmes, and Stephen G Walker. A general framework for updating belief distributions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(5):1103–1130, 2016.
- Ramiro D Camino, Christian A Hammerschmidt, and Radu State. Improving missing data imputation with deep generative models. *arXiv preprint arXiv:1902.10666*, 2019.
- Sourav Chatterjee and Persi Diaconis. The sample size required in importance sampling. *The Annals of Applied Probability*, 28(2):1099–1135, 2018.
- Ricky TQ Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. *Advances in Neural Information Processing Systems*, 31, 2018.
- Mark Collier, Alfredo Nazabal, and Christopher KI Williams. Vaes in the presence of missing data. *arXiv preprint arXiv:2006.05301*, 2020.
- Erik Daxberger and José Miguel Hernández-Lobato. Bayesian variational autoencoders for unsupervised out-of-distribution detection. *arXiv preprint arXiv:1912.05651*, 2019.
- Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, pp. 1050–1059. PMLR, 2016.
- Pedro J García-Laencina, José-Luis Sancho-Gómez, and Aníbal R Figueiras-Vidal. Pattern classification with missing data: a review. *Neural Computing and Applications*, 19(2):263–282, 2010.
- Samuel Gershman and Noah Goodman. Amortized inference in probabilistic reasoning. In *Proceedings of the annual meeting of the cognitive science society*, volume 36, 2014.
- Lovedeep Gondara and Ke Wang. Mida: Multiple imputation using denoising autoencoders. In *Pacific-Asia conference on knowledge discovery and data mining*, pp. 260–272. Springer, 2018.
- Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-VAE: Learning basic visual concepts with a constrained variational framework. *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Sy2fzU9gl>.
- Chris C Holmes and Stephen G Walker. Assigning a value to a power likelihood in a general bayesian model. *Biometrika*, 104(2):497–503, 2017.
- Gao Huang, Yixuan Li, Geoff Pleiss, Zhuang Liu, John E Hopcroft, and Kilian Q Weinberger. Snapshot ensembles: Train 1, get m for free. *arXiv preprint arXiv:1704.00109*, 2017.
- Niels Bruun Ipsen, Pierre-Alexandre Mattei, and Jes Frellsen. not-miwae: Deep generative modelling with missing not at random data. *arXiv preprint arXiv:2006.12871*, 2020.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations*.
- Diederik P Kingma and Max Welling. An introduction to variational autoencoders. *arXiv preprint arXiv:1906.02691*, 2019.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in Neural Information Processing Systems*, 30, 2017.
- Sarah Lewis, Tatiana Matejovicova, Yingzhen Li, Angus Lamb, Yordan Zaykov, Miltiadis Allamanis, and Cheng Zhang. Accurate imputation and efficient data acquisition with transformer-based vaes. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021.

- Roderick JA Little and Donald B Rubin. *Statistical analysis with missing data*, volume 793. John Wiley & Sons, 2019.
- Chao Ma and Cheng Zhang. Identifiable generative models for missing not at random data imputation. *Advances in Neural Information Processing Systems*, 34, 2021.
- Chao Ma, Wenbo Gong, José Miguel Hernández-Lobato, Noam Koenigstein, Sebastian Nowozin, and Cheng Zhang. Partial vae for hybrid recommender system. In *NIPS Workshop on Bayesian Deep Learning*, volume 2018, 2018a.
- Chao Ma, Sebastian Tschitschek, Konstantina Palla, José Miguel Hernández-Lobato, Sebastian Nowozin, and Cheng Zhang. Eddi: Efficient dynamic discovery of high-value information with partial vae. *arXiv preprint arXiv:1809.11142*, 2018b.
- Chao Ma, Sebastian Tschitschek, José Miguel Hernández-Lobato, Richard Turner, and Cheng Zhang. Vaem: a deep generative model for heterogeneous mixed type data. *arXiv preprint arXiv:2006.11941*, 2020.
- Pierre-Alexandre Mattei and Jes Frelsen. Leveraging the exact likelihood of deep latent variable models. *Advances in Neural Information Processing Systems*, 31, 2018.
- Pierre-Alexandre Mattei and Jes Frelsen. Miwae: Deep generative modelling and imputation of incomplete data sets. In *International Conference on Machine Learning*, pp. 4413–4423. PMLR, 2019.
- Jared S Murray. Multiple imputation: a review of practical and theoretical findings. *Statistical Science*, 33(2):142–159, 2018.
- Alfredo Nazabal, Pablo M Olmos, Zoubin Ghahramani, and Isabel Valera. Handling incomplete heterogeneous data using vaes. *Pattern Recognition*, 107:107501, 2020.
- Fulufhelo V Nelwamondo, Shakir Mohamed, and Tshilidzi Marwala. Missing data: A comparison of neural network and expectation maximization techniques. *Current Science*, pp. 1514–1521, 2007.
- Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 427–436, 2015.
- Yeping Lina Qiu, Hong Zheng, and Olivier Gevaert. Genomic data imputation with variational auto-encoders. *GigaScience*, 9(8):giaa082, 2020.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, pp. 1278–1286. PMLR, 2014.
- Sandip Sinharay, Hal S Stern, and Daniel Russell. The use of multiple imputation for the analysis of missing data. *Psychological methods*, 6(4):317, 2001.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- Stephen Walker and Nils Lid Hjort. On Bayesian consistency. *Journal of the Royal Statistical Society: Series B*, 63(4):811–821, 2001.

A APPENDIX

A.1 DATA AND CODE AVAILABILITY

All data used in this manuscript are publicly available. Gene expression data is version 2 of the adjusted pan-cancer gene expression data obtained from Synapse and can be found at <https://www.synapse.org/!Synapse:syn4976369.2>.

All code used to implement the analyses in this manuscript is hosted on GitHub and will be made publicly available upon acceptance.

A.2 SOFTWARE REQUIREMENTS

The analyses carried out in this manuscript require the following software: python v3.10, TensorFlow v2.7.0; R: penalized v0.9, MASS v7.3, caret v6.0.

A.3 β -VAES AND THE POWER LIKELIHOOD

In the following, we show that the variational parameters ϕ which maximize the β -VAE bound equivalently minimize the KL divergence between the variational posterior $q_\phi(\mathbf{Z}|\mathbf{X})$ and the true posterior under the power likelihood. Specifically, the KL divergence between the variational posterior and the posterior under the power likelihood is given by:

$$\begin{aligned} D_{\text{KL}}(q_\phi(\mathbf{Z}|\mathbf{X}), p_{\theta,\beta}(\mathbf{Z}|\mathbf{X})) &= E_{\mathbf{Z} \sim q_\phi(\mathbf{Z}|\mathbf{X})} \left[\log \left(\frac{p_{\theta,\beta}(\mathbf{X}) q_\phi(\mathbf{Z}|\mathbf{X})}{p_\theta(\mathbf{X} | \mathbf{Z})^{1/\beta} p(\mathbf{Z})} \right) \right] \\ &= - \sum_{n=1}^N E_{\mathbf{z}_n \sim q_\phi(\mathbf{z}_n|\mathbf{x}_n)} \left[\log \left(p_\theta(\mathbf{x}_n | \mathbf{z}_n)^{1/\beta} \right) \right] + \sum_{n=1}^N E_{\mathbf{z}_n \sim q_\phi(\mathbf{z}_n|\mathbf{x}_n)} \left[\log \left(\frac{q_\phi(\mathbf{z}_n|\mathbf{x}_n)}{p(\mathbf{z}_n)} \right) \right] \\ &\quad + \log(p_{\theta,\beta}(\mathbf{X})) \\ &= \text{const.} - \sum_{n=1}^N E_{\mathbf{z}_n \sim q_\phi(\mathbf{z}_n|\mathbf{x}_n)} \left[\log \left(p_\theta(\mathbf{x}_n | \mathbf{z}_n)^{1/\beta} \right) \right] + D_{\text{KL}}(q_\phi(\mathbf{z}_n|\mathbf{x}_n), p(\mathbf{z}_n)). \end{aligned}$$

Thus, we can equivalently find ϕ , which maximize the ELBO:

$$\text{ELBO} = \sum_{n=1}^N \mathbb{E}_{\mathbf{z}_n \sim q_\phi(\mathbf{z}_n|\mathbf{x}_n)} [\log p_\theta(\mathbf{x}_n|\mathbf{z}_n)] - \beta D_{\text{KL}}(q_\phi(\mathbf{z}_n|\mathbf{x}_n), p(\mathbf{z}_n)).$$

Example: factorized Gaussian. Assume the generative model is a factorized Gaussian (as is used for the genomic data in Section 4):

$$p_\theta(\mathbf{x}_n|\mathbf{z}_n) = \prod_{d=1}^D \text{N}(x_{n,d} | \mu_d(\mathbf{z}_n), \sigma_d^2(\mathbf{z}_n)),$$

where $(\mu_d(\mathbf{z}_n), \sigma_d^2(\mathbf{z}_n))$ for $d = 1, \dots, D$ represent the output of the final layer of the neural network with weights and biases contained in θ . In this case, the full conditional of the missing data under the power likelihood is

$$\begin{aligned} p_{\theta,\beta}(\mathbf{x}_{\text{mis},n} | \mathbf{x}_{\text{obs},n}, \mathbf{z}_n) &\propto p_\theta(\mathbf{x}_{\text{mis},n} | \mathbf{x}_{\text{obs},n}, \mathbf{z}_n)^{1/\beta} \\ &= \left(\prod_{d \in \mathcal{D}_{\text{mis},n}} \text{N}(x_{n,d} | \mu_d(\mathbf{z}_n), \sigma_d^2(\mathbf{z}_n)) \right)^{1/\beta} \\ &= \left(\prod_{d \in \mathcal{D}_{\text{mis},n}} \frac{1}{\sqrt{2\pi\sigma_d^2(\mathbf{z}_n)}} \exp \left(\frac{1}{2\sigma_d^2(\mathbf{z}_n)} (x_{n,d} - \mu_d(\mathbf{z}_n))^2 \right) \right)^{1/\beta} \\ &\propto \prod_{d \in \mathcal{D}_{\text{mis},n}} \exp \left(\frac{1}{2\beta\sigma_d^2(\mathbf{z}_n)} (x_{n,d} - \mu_d(\mathbf{z}_n))^2 \right) \\ &\propto \prod_{d \in \mathcal{D}_{\text{mis},n}} \text{N}(x_{n,d} | \mu_d(\mathbf{z}_n), \beta\sigma_d^2(\mathbf{z}_n)), \end{aligned}$$

where $\mathcal{D}_{\text{mis},n} \subseteq \{1, \dots, D\}$ contains the indices of the missing features for the n th data point. Thus, in this case, sampling from the full conditional of the missing data under the power likelihood corresponds to sampling from the Gaussian with variance rescaled by a factor of β . Note that for $\beta > 1$ this corresponds to increasing the spread and uncertainty of the missing data, which is critical to improve coverage of the deep generative model.

A.4 SAMPLE IMPORTANCE RESAMPLING

We first note that our SIR scheme differs slightly from the scheme proposed by Mattei & Frellsen (2019), who propose joint importance samples $(\mathbf{z}_n^{(s)}, \mathbf{x}_{\text{mis},n}^{(s)})$ from

$$\mathbf{z}_n^{(s)} \sim q_\phi(\mathbf{z}_n | \mathbf{x}_{\text{mis},n}^{(0)}, \mathbf{x}_{\text{obs},n}), \quad \mathbf{x}_{\text{mis},n}^{(s)} \sim p_\theta(\mathbf{x}_{\text{mis},n} | \mathbf{x}_{\text{obs},n}, \mathbf{z}_n^{(s)}).$$

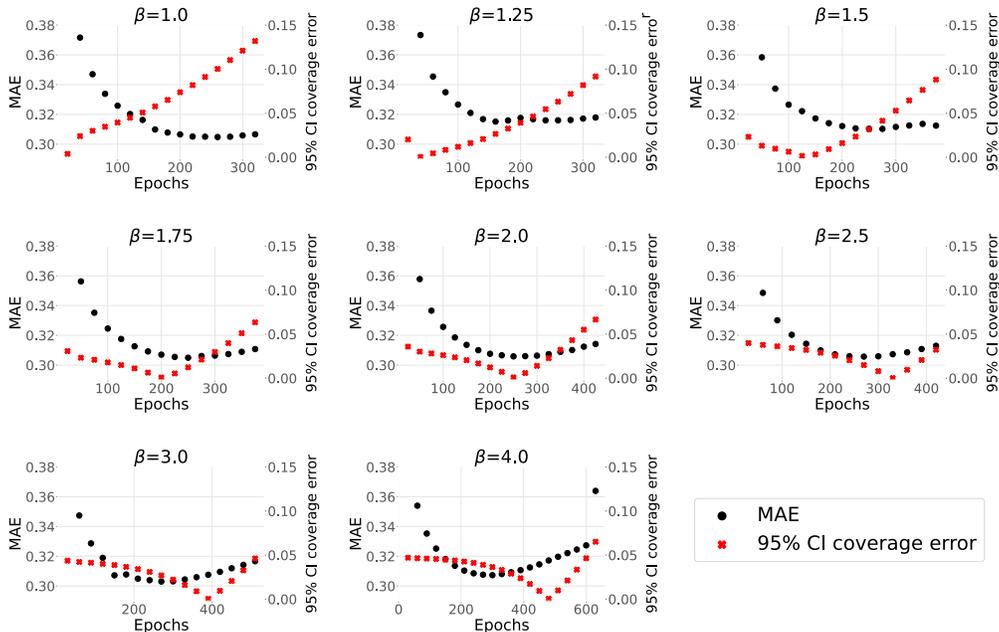
Instead, we only consider importance sampling for $(\mathbf{z}_n^{(s)})$ and subsequently sample the missing data for each of resampled latent variables. Importantly, if the effective sample size is low, resulting in potential duplicates in the M samples of latent variables, we obtain improved variability across multiple imputations of the missing data, compared to the approach of Mattei & Frellsen (2019).

Example: factorized Gaussian. Assume the generative model is a factorized Gaussian, then the importance weights are proportional to :

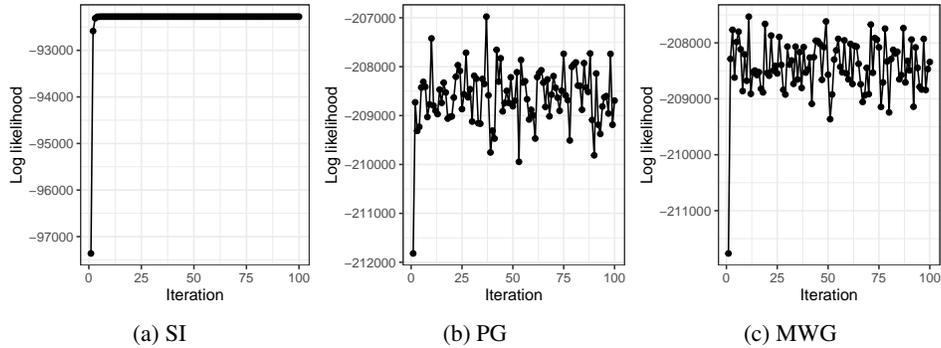
$$\begin{aligned} \omega_n^{(s)} &= \frac{p_\theta(\mathbf{x}_{\text{obs},n} | \mathbf{z}_n^{(s)})^{1/\beta} p(\mathbf{z}_n^{(s)})}{q_\phi(\mathbf{z}_n^{(s)} | \mathbf{x}_{\text{mis},n}^{(0)}, \mathbf{x}_{\text{obs},n})} \\ &= \frac{\left(\prod_{d \in \mathcal{D}_{\text{obs},n}} \mathcal{N}(x_{n,d} | \mu_d(\mathbf{z}_n), \sigma_d^2(\mathbf{z}_n)) \right)^{1/\beta} \mathcal{N}(\mathbf{z}_n^{(s)} | \mathbf{0}, \mathbf{I})}{q_\phi(\mathbf{z}_n^{(s)} | \mathbf{x}_{\text{mis},n}^{(0)}, \mathbf{x}_{\text{obs},n})}, \end{aligned}$$

where $\mathcal{D}_{\text{obs},n} \subseteq \{1, \dots, D\}$ contains the indices of the observed features for the n th data point.

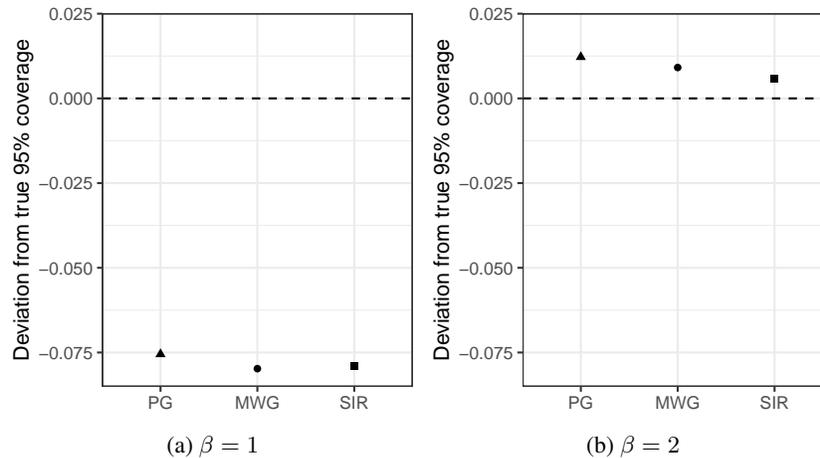
A.5 SUPPLEMENTAL FIGURES



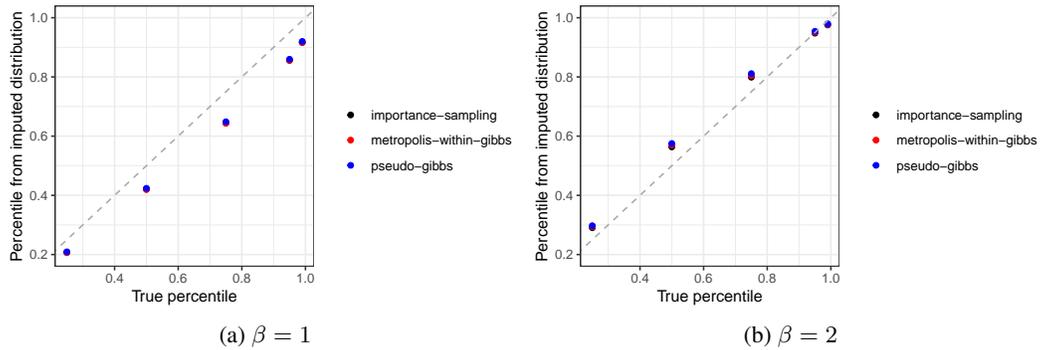
Supplementary Figure A.1: **Results from 5-fold cross-validation to determine optimal model and hyper-parameters.** MAE (black) and EC (red) of 95% CI (computed based on quantiles) at different training epochs. We aim to minimize both of these metrics for optimal training parameters.



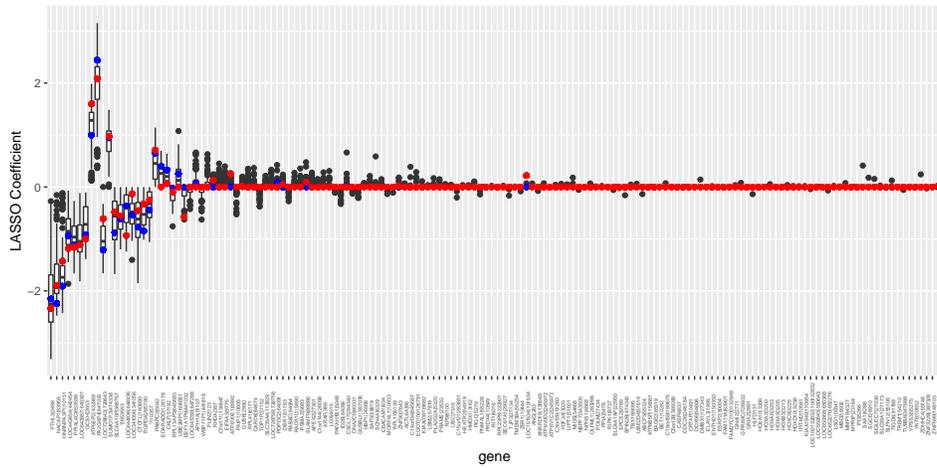
Supplementary Figure A.2: **Trace plots monitoring convergence of Markov Chain Monte Carlo schemes for $\beta = 2$ case.** Here we report the log likelihood of the data under the generative model at each iteration for **(a)** single imputation (SI), **(b)** pseudo-Gibbs (PG) and **(c)** Metropolis-within-Gibbs (MWG). For visualization purposes, we show iterations 1 to 100, but ran to 1000 iterations in implementation. We also show a single data set from the $M = 100$ datasets as an example for PG and MWG.



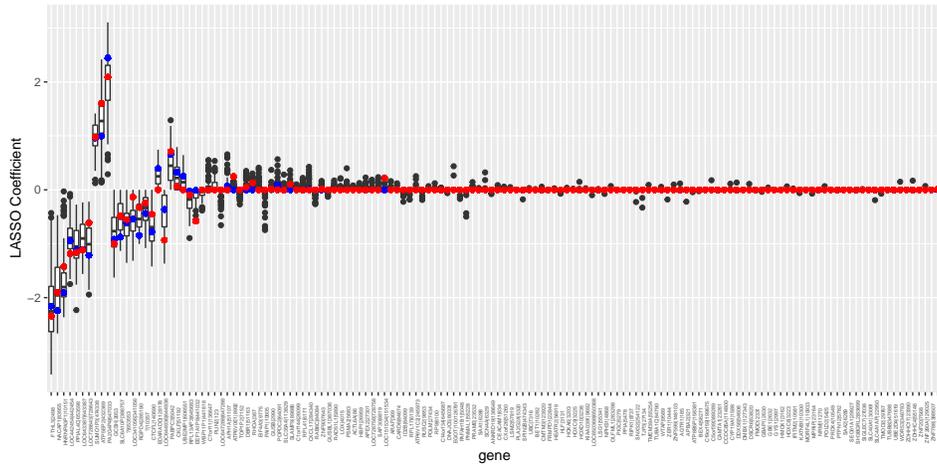
Supplementary Figure A.3: **Deviation from true coverage for all three multiple imputation approaches.** Here we report the deviation from the desired coverage of 95% (shown here by the dotted line) for all three multiple imputation approaches pseudo-Gibbs (PG), Metropolis-within-Gibbs (MWG) and sampling importance resampling (SIR) for **(a)** $\beta = 1$, and **(b)** $\beta = 2$.



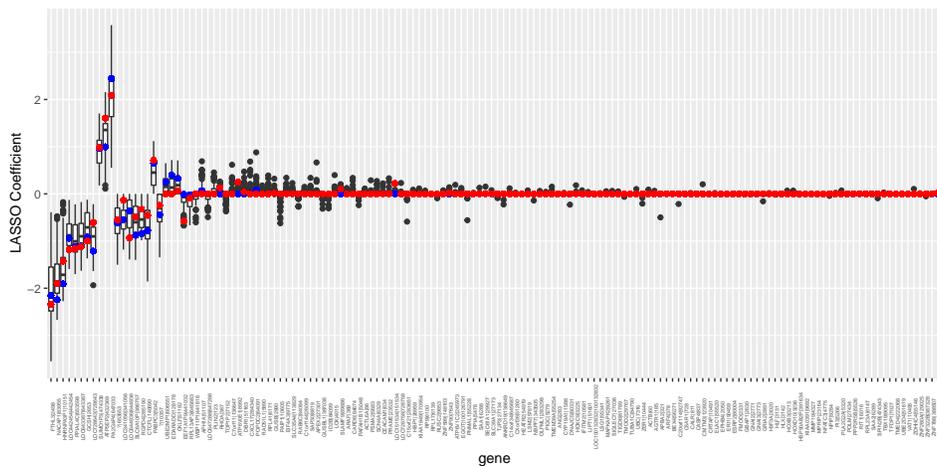
Supplementary Figure A.4: **Coverage by percentiles of all three multiple imputation approaches.** Here we report the coverage evaluated by percentiles across all imputed datasets compared to the true percentiles of 0.25, 0.5, 0.75, 0.95 and 0.99 for all three multiple imputation approaches pseudo-Gibbs (PG), Metropolis-within-Gibbs (MWG) and sampling importance resampling (SIR) for (a) $\beta = 1$, and (b) $\beta = 2$.



(a) PG

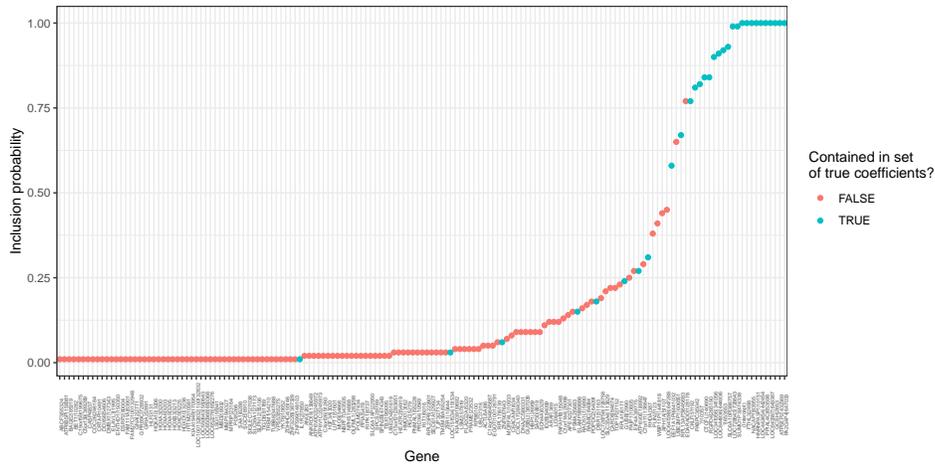


(b) MWG

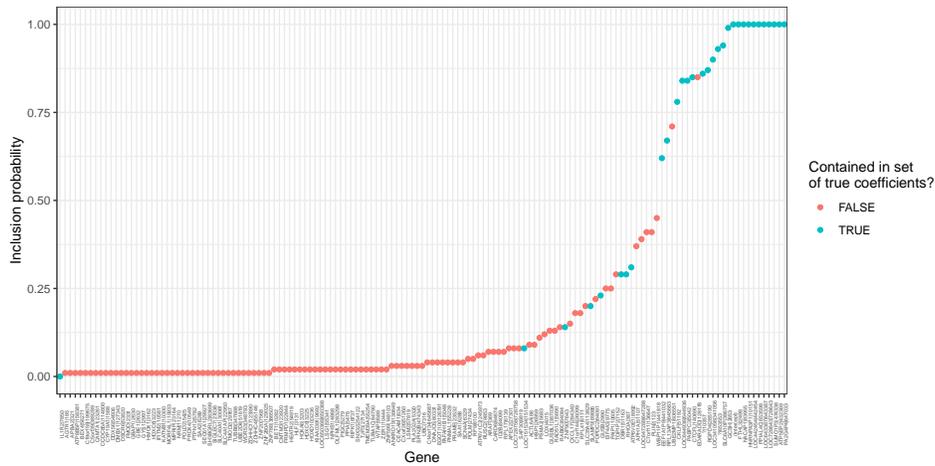


(c) SIR

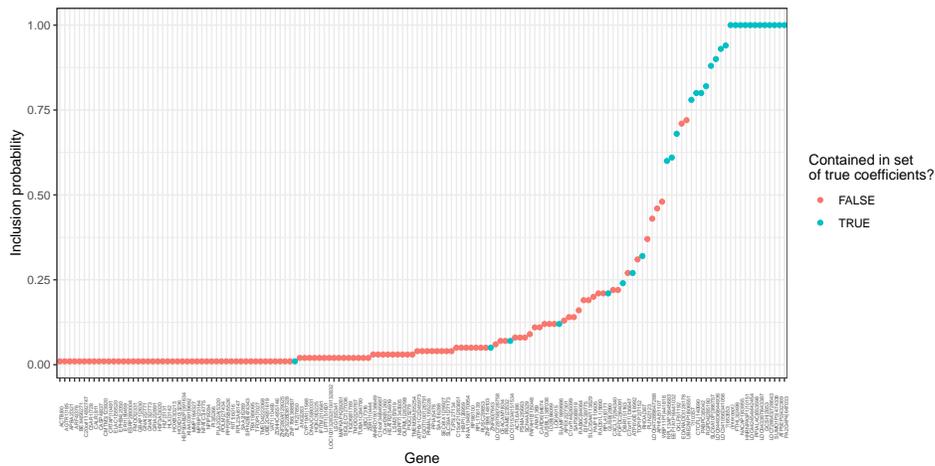
Supplementary Figure A.5: **LASSO regression coefficients across multiple imputation approaches.** Here we report the LASSO regression coefficients across all 100 plausible imputed datasets for (a) pseudo-Gibbs (PG), (b) Metropolis-within-Gibbs (MWG) and (c) sampling importance resampling (SIR). LASSO regression coefficient value for the true dataset is highlighted in red, and for single imputation is highlighted in blue. For the purpose of visualization, the intercept was removed from this plot.



(a) PG



(b) MWG



(c) SIR

Supplementary Figure A.6: **Inclusion probability across all non-zero LASSO coefficients across multiple imputation approaches.** Here we report the inclusion probability P_{incl} for all non-zero LASSO regression coefficients for (a) pseudo-Gibbs (PG), (b) Metropolis-within-Gibbs (MWG) and (c) sampling importance resampling (SIR). Genes in this set that are included in the true LASSO coefficients are highlighted in blue, and those not included in the true discriminating gene set are highlighted in pink.