

---

# Missing Data Imputation and Acquisition with Deep Hierarchical Models and Hamiltonian Monte Carlo

---

**Ignacio Peis**  
Universidad Carlos III de Madrid  
Madrid, Spain  
ipeis@tsc.uc3m.es

**Chao Ma**  
Microsoft Research  
University of Cambridge  
Cambridge, UK  
cm905@cam.ac.uk

**José Miguel Hernández-Lobato**  
University of Cambridge  
Cambridge, UK  
jmh233@cam.ac.uk

## Abstract

Variational Autoencoders (VAEs) have recently been highly successful at imputing and acquiring heterogeneous missing data. However, within this specific application domain, existing VAE methods are restricted by using only one layer of latent variables and strictly Gaussian posterior approximations. To address these limitations, we present HH-VAEM, a Hierarchical VAE model for mixed-type incomplete data that uses Hamiltonian Monte Carlo with automatic hyper-parameter tuning for improved approximate inference. Our experiments show that HH-VAEM outperforms existing baselines in the tasks of missing data imputation and supervised learning with missing features. Finally, we also present a sampling-based approach for efficiently computing the information gain when missing features are to be acquired with HH-VAEM. Our experiments show that this sampling-based approach is superior to alternatives based on Gaussian approximations.

## 1 Introduction

Many real-world unsupervised learning tasks require dealing with complicated datasets with mixed types (real, positive-valued, continuous, or discrete) and missing values. For this purpose, variational autoencoders [24, 25, 41] stand out in the recent literature as robust generative models that efficiently handle high-dimensional data. However, in their naive configuration, every data dimension is assumed to have similar statistical properties (i.e., homogeneity), and all dimensions are considered to be completely observed. Both assumptions won't hold in many real-world scenarios. Recent works have adapted VAEs to handle incomplete [9, 14, 34, 36] and mixed-type data [18, 33, 38], and demonstrated improved performance in downstream tasks such as missing data imputation and active information acquisition. Despite these advances, existing approaches are far from optimal as they are based on restrictive design choices: 1), only one layer of latent variables are considered; 2), Gaussian posterior approximations are usually adopted. These will lead to limited flexibility and additional bias, especially under real-world settings with complex mixed-type incomplete data.

In the literature, the issue of model flexibility and inference bias are often addressed separately. For example, approximate inference bias can be reduced by using Monte Carlo sampling [46, 51]. More specifically, Hamiltonian Monte Carlo (HMC) [5, 12] stands out among MCMC methods in machine learning due to its superior efficiency for exploring the target density. In the context of VAE, HMC has also been combined with stochastic variational inference [7] for improving the training of VAEs.

On the other hand, the flexibility of VAEs can be improved by considering hierarchical VAEs with multiple layers of hidden variables [8, 35, 48, 53]. By using a hierarchical structure in the latent space, they enforce the information to flow from high-level representations to more specific observable factors, imitating the way information is often organized in the real world.

However, the issue of modeling flexibility and approximate inference bias are often heavily intertwined, and addressing them simultaneously in a *joint* manner is highly non-trivial. The hierarchical organization of the latent variables creates complicated posterior dependencies that are not straightforward to deal with and require special consideration. To improve Gaussian approximate inference, most works opt by defining shared paths between the recognition and generative networks. While this makes hierarchical VAEs practical, the bias introduced by the Gaussian approximations is still present. To the best of our knowledge, none of the aforementioned hierarchical VAEs has been previously combined with Monte Carlo algorithms for improving over standard Gaussian approximate inference.

To overcome such limitations, we focus on training new hierarchical VAE models for heterogeneous mixed-type data with HMC. Our models can be used for missing data imputation and for supervised learning with missing data. We also present a sampling-based framework that allows our models to perform accurate sequential active information acquisition.

Our main contributions are as follows:

- We present HH-VAEM, a deep hierarchical model for handling mixed-type incomplete data that uses HMC with automatic hyper-parameter tuning for outperforming amortized variational inference by generating low bias samples from the true posterior.
- We propose a sampling-based strategy for missing feature acquisition that benefits from the improved inference of HH-VAEM. By using histograms to estimate the mutual information, this strategy achieves lower bias than other Gaussian-based alternatives.
- We exhaustively evaluate HH-VAEM in the tasks of 1) missing data imputation, 2) supervised learning with missing data and 3) information acquisition with our sampling-based strategy. In all cases we report significant gains with respect to baselines.

## 2 Related work

### 2.1 VAEs for mixed-type incomplete data

Variational Autoencoders [24, 41] are deep generative models that make use of encoder and decoder networks for mapping data into a latent Gaussian distribution, and reconstructing the latent codes into the original observational space, respectively. The parameters of these networks are trained using amortized Variational Inference [10, 56] optimizing a lower bound (ELBO) on the log evidence:  $\mathbb{E}_{q_z(z|\mathbf{x})} \log \frac{p_\theta(\mathbf{x}, z)}{q_\psi(z|\mathbf{x})}$ , where the generative model  $p_\theta(\mathbf{x}, z)$  can be expressed in terms of the likelihood  $p_\theta(\mathbf{x}|z)$  and the prior  $p(z)$ . In mixed-type data, the vector  $\mathbf{x}$  is composed by data from different types: real, positive real, categorical, binary, etc. A naive approach is to consider a factorized decoder using different likelihood contributions  $p_\theta(\mathbf{x}|z) = \prod_d p_\theta(x_d|z)$  [1, 38]. Nonetheless, the problem of handling unbalanced likelihoods leads to the domination of some dimensions during the optimization process. In [33], authors propose a solution using a set of *marginal* VAEs that encode each feature into a Gaussian uni-dimensional space, and a *dependency* VAE that captures the inter-dimensional dependencies more effectively using balanced Gaussian likelihoods.

By marginalizing each dimension of the decoder, incomplete data can be easily handled by dividing the vector  $\mathbf{x}$  into the observed  $\mathbf{x}_O$  and unobserved  $\mathbf{x}_U$  parts. This methodology is completely valid when using the missing-at-random (MAR) assumption [30], i.e. assuming the missing mechanism is independent of the missing values. In this our work, we use the same assumption. As proposed in [38] and [36], the ELBO objective is transformed into a lower bound on the observed data, and the unobserved data is replaced with zeros.

### 2.2 Hierarchical VAEs

Hierarchical models have been successfully employed in deep generative modeling, [2, 44, 45]. In VAEs, defining a hierarchical latent space for VAEs can be straightforward. Nevertheless, potential pitfalls require special attention. Concretely, if the decoder is powerful enough, the model tends

to uniquely use the shallowest layers, ignoring the deepest ones and falling into the well-known problem of *posterior collapse* [35, 40, 54]. In the last few years, several works have investigated possible hierarchical structures for VAEs. In [48], a bottom-up deterministic path is used along with a top-down inference path that shares parameters with the top-down structure of the generative model. In [35], the authors use a bidirectional stochastic inference path. More recently, [53] or [8] have adapted these architectures to complex datasets and high quality images. Possibly motivated by the residual connections in [23], all these works use a shared path between recognition and generative models that helps in tying the divergences between approximations and priors in the ELBO.

### 2.3 Hamiltonian Monte Carlo

HMC [4, 12, 39] is a particularly effective MCMC algorithm for sampling from a target distribution  $p(\mathbf{z}) = \frac{1}{\mathcal{Z}}p^*(\mathbf{z})$  where  $\mathcal{Z}$  is the normalization constant and  $\mathbf{z}$  is a  $d$ -dimensional vector. By augmenting this model to  $p(\mathbf{r}, \mathbf{z}) = \mathcal{N}(\mathbf{r}; \mathbf{0}, \mathbf{M})p(\mathbf{z})$ , and denoting  $\mathbf{r}$  as the *momentum* variable with diagonal covariance matrix  $\mathbf{M}$ , with the same dimensionality as  $\mathbf{z}$ , HMC samples are obtained from the distribution by simulating the time-evolution of a fictitious physical system.

The algorithm starts by firstly sampling  $\mathbf{z}$  and  $\mathbf{r}$  from an initial proposal and the momentum distribution, respectively. Chains with length  $T$  are built by recurrently proposing and accepting new states. To propose a new state, the Hamiltonian dynamics are simulated using a symplectic integrator, Leapfrog being the most common choice. The following updates are repeated for  $l = 1 : LF$  steps:

$$\begin{aligned} \mathbf{r}_{l+\frac{1}{2}} &= \mathbf{r}_l + \frac{1}{2}\phi \odot \nabla_{\mathbf{z}_l} \log p^*(\mathbf{z}_l), \\ \mathbf{z}_{l+1} &= \mathbf{z}_l + \mathbf{r}_{l+\frac{1}{2}} \odot \phi \odot \frac{1}{\mathbf{M}}, \\ \mathbf{r}_{l+1} &= \mathbf{r}_{l+\frac{1}{2}} + \frac{1}{2}\phi \odot \nabla_{\mathbf{z}_{l+1}} \log p^*(\mathbf{z}_{l+1}), \end{aligned} \tag{1}$$

where  $\odot$  refers to the Hadamard product, and  $\phi$  is the *step size* hyperparameter. Although it is typically defined as a scalar for simplicity, a  $d$ -dimensional vector can be considered to apply different step sizes per dimension, or further, as considered in this work, a  $T \times d$  matrix to apply different steps per each proposal of the chain. The new proposal  $(\mathbf{z}', \mathbf{r}')$  is accepted with probability  $\min [1, \exp(-H(\mathbf{z}', \mathbf{r}') + H(\mathbf{z}, \mathbf{r}))]$ , where

$$H(\mathbf{z}, \mathbf{r}) = -\log p^*(\mathbf{z}) + \frac{1}{2}\mathbf{r}^T \mathbf{M}^{-1} \mathbf{r}. \tag{2}$$

For the consecutive  $T$  proposals, a new momentum  $\mathbf{r}$  is resampled and the updates of eq. (1) are repeated for  $LF$  steps to update the state if  $(\mathbf{z}', \mathbf{r}')$  is accepted.

### 2.4 Active Feature Acquisition

Among all the Active Learning techniques, Active Feature Acquisition [20, 37, 43, 50] is of special interest in cost-sensitive applications for modeling a trade-off between the improvement of predictions and the cost of acquiring new data at the feature level. Several works in the recent literature have studied methods for performing the task of sequentially acquiring high-value information by selecting features that maximize an information theoretical reward function and enhance the accuracy of the predictions. This task is denoted by SAIA (Sequential Active Information Acquisition). In [34], an efficient method is proposed for approximating a non-tractable reward by using the encoder of a VAE that handles missing data. In [33] they extend this method for handling mixed-type data. Both works estimate the reward by relying on Gaussian approximations given by the encoder networks.

## 3 Hamiltonian hierarchical VAE for mixed-type incomplete data (HH-VAEM)

The HH-VAEM model (Figure 1 (a)) is a Hierarchical VAE for mixed-type, incomplete data that incorporates HMC with automatic hyper-parameter optimization for sampling from the posterior of the latent variables. In a first stage, the mixed-type data is encoded into marginal Gaussian posterior distributions as given by univariate VAEs fitted to each data dimension. In a second stage, a hierarchical structure captures the dependencies among the standardized, homogeneous dimensions with well-balanced Gaussian likelihoods. The model is trained using samples from the posterior of the

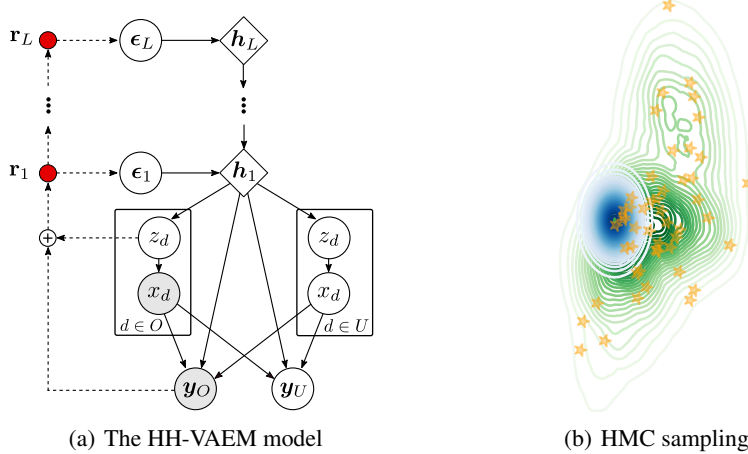


Figure 1: The HH-VAEM model (a). Illustrative example (b): samples  $\epsilon^{(T)}$  obtained with HMC (orange) following the true posterior  $p(\epsilon | \mathbf{x}_O, \mathbf{y}_O)$  (green) using the Gaussian distribution given by the encoder  $q^{(0)}(\epsilon | \mathbf{x}_O, \mathbf{y}_O)$  (blue) as the initial proposal, with latent dimension  $M = 2$ .

hierarchical latent variables by means of HMC, whilst the HMC hyper-parameters are automatically tuned. A more detailed description is provided in the following subsections.

### 3.1 Notation

The model generates both data  $\mathbf{x} \in \mathbb{R}^D$  and output  $\mathbf{y} \in \mathbb{R}^P$ , where each of these variables is divided into observed parts  $\mathbf{x}_O, \mathbf{y}_O$  and unobserved parts  $\mathbf{x}_U, \mathbf{y}_U$ . Each dimension of  $\mathbf{x}$  is denoted by  $x_d$ . A training set is composed of  $N$  observations as tuples  $(\mathbf{x}_O^{(n)}, \mathbf{y}_O^{(n)})$ . For ease of notation, we omit the observation index  $n$ , and the objectives are presented for a single observation. The *dependency* latent space is composed of  $L$  latent variables  $[\epsilon_1, \dots, \epsilon_L]$ , with  $\epsilon_l \in \mathbb{R}^{m_l}$ . The dimension of the joint latent distribution is  $\sum_l m_l = M$ . In the marginal VAEs, the latent variables  $z_d$  are unidimensional.

### 3.2 Handling heterogeneous incomplete data

Following the strategy proposed by [33], we perform a two-staged approach for handling heterogeneous data. The marginal distribution of each feature  $p_{\theta_d}(x_d)$  is modeled by a one-dimensional VAE. First, the  $D$  marginal VAEs are trained independently by maximizing the marginal ELBO over observed points

$$\mathcal{L}_d(x_d; \{\theta_d, \gamma_d\}) = \mathbb{1}(x_d \in \mathbf{x}_O) \mathbb{E}_{q_{\gamma_d}(z_d | x_d)} \log \frac{p_{\theta_d}(x_d, z_d)}{q_{\gamma_d}(z_d | x_d)}, \quad (3)$$

where  $\mathbb{1}(x_d \in \mathbf{x}_O)$  is an indicator function that activates the ELBO when the feature is observed. Under the missing-at-random assumption, which is the one considered in this work, Equation (3) leads to a lower bound on observed data likelihoods. Second, a *dependency* VAE encodes a vector  $\mathbf{z}$  with the concatenated samples from the marginal posteriors  $q_{\gamma_d}(z_d | x_d)$  into the global latent variable  $\mathbf{h}$ , using zero-filling for the unobserved variables. By using this approach,  $\mathbf{z}$  is now homogeneous and can then be easily modeled using a standard Gaussian decoder. By contrast, other works [13, 38] directly operate with different decoding likelihoods per data type. This approach often leads to having very different magnitudes in the ELBO and may reduce learning efficiency. The ELBO for the second stage dependency VAE is

$$\mathcal{L}(\mathbf{x}_O, \mathbf{y}_O; \{\theta, \psi\}) = \mathbb{E}_{q_\psi} \left[ \log \frac{p_\theta(\mathbf{z}_O, \mathbf{y}_O, \epsilon)}{q_\psi(\epsilon | \mathbf{z}_O, \mathbf{x}_O, \mathbf{y}_O)} \right] \quad (4)$$

where  $\epsilon = \{\epsilon_1, \dots, \epsilon_L\}$  is a set of reparameterized hierarchical latent variables. Further details on the design of the hierarchical dependency VAE are provided below.

### 3.3 Predictive enhancement

The combination of generative and discriminative models is an effective well-studied strategy for dealing with predictive models under missing data [15, 52]. In [15], they model  $p(\mathbf{x})$  for imputing missing data using a Gaussian Mixture Model. In a deep learning context, recent supervised VAE models have revisited this combination [21, 47] or used factorisations of type  $p(\mathbf{z})p(\mathbf{x}|\mathbf{z})p(\mathbf{y}|\mathbf{z})$  [22] to learn meaningful representations. In [28] the authors propose a deep generative model with factorisation  $p(\mathbf{z})p(\mathbf{x}|\mathbf{z})p(\mathbf{y}|\mathbf{x}, \mathbf{z})$  for detecting adversarial attacks.

With the aim at reinforcing the prediction of the variable of interest, we turn into a supervised model by including a separate predictor for  $p_{\theta_y}(\mathbf{y}|\hat{\mathbf{x}}, \mathbf{h})$ , apart from the decoder  $p_{\theta_z}(\mathbf{z}|\mathbf{h})$ . The vector  $\hat{\mathbf{x}} = (x_i \in \mathbf{x}_O, \hat{x}_j \in \mathbf{x}_U)$  includes the observed part and imputation of the missing variables  $\hat{x}_j$  by decoding the latent  $\mathbf{h}$  into  $\mathbf{z}$  using  $p(\mathbf{z}|\mathbf{h}_1)$ , and each dimension  $z_j \in \mathbf{z}$  into  $\hat{x}_j$  using  $p(x_j|z_j)$ . The predictor parameters  $\theta_y$  are optimized along with the decoder parameters  $\theta_z$ .

### 3.4 Hierarchical reparameterized latent space

A hierarchical structure over the latent space  $\mathbf{h} = \{\mathbf{h}_1, \dots, \mathbf{h}_L\}$  enriches the prior assumptions and permits a flexible generation of data in a more natural fashion. Nevertheless, as stated in [4, 5], HMC can be pathological when used for sampling from hierarchical densities, where the magnitude of autoregressive variations increase with the depth. For approximating the Hamiltonian dynamics, inside each Leapfrog integrator step, gradients  $\nabla_{\mathbf{h}_{1:L}} \log p^*(\mathbf{h}_{1:L})$  are required. Due to the strong curvature regions, huge norm of high-order derivatives are backpropagated and might eventually explode, ending in overflow issues (Figure 35 in [5]).

If we were to run our HMC method over the hierarchical variables without any reparameterization (Figure 4a), by the time the states reached the aforementioned problematic regions, the integrator would diverge and we would experience the aforementioned overflow problems. By rejecting these problematic states, chains would get stuck close to the proposal and the hierarchical density would not be properly explored (concluding that HMC would not improve the Gaussian proposal). To give an example, in [53], the AR path  $p(\mathbf{h}_l|\mathbf{h}_{<l})$  would provoke instabilities inside the HMC integrator due to huge gradients in  $\nabla_{\mathbf{h}_{1:L}} \log p^*(\mathbf{h}_{1:L})$ .

We successfully solved this issue by introducing a hierarchical reparameterization technique. The representation at each layer is reparameterized from variable  $\epsilon_l$  with standard Gaussian prior  $p(\epsilon_l)$

$$\mathbf{h}_l = f_{\mu_l}(\mathbf{h}_{l+1}) + f_{\sigma_l}(\mathbf{h}_{l+1}) \cdot \epsilon_l, \quad (5)$$

where the functions  $f_{\mu_l}(\mathbf{h}_{l+1})$  and  $f_{\sigma_l}(\mathbf{h}_{l+1})$  are applied by NNs with parameters  $\theta_l = \{\theta_{\mu_l}, \theta_{\sigma_l}\}$ . The result is equivalent as learning the mean and covariance of autoregressive variables (see Figure 4 for illustrative details). However, thanks to this trick, we relax the dependencies among the latent variables, resulting in a smoother joint posterior density  $p(\epsilon|\mathbf{x}_O, \mathbf{y}_O)$ . Performing the inference over  $\epsilon = \{\epsilon_1, \dots, \epsilon_L\}$  leads to a better posed basis for running our HMC optimization, detailed in section 3.5, and avoids the necessity of employing more advanced HMC samplers like [5, 16]. We include further details on the pathological behavior and demonstration of the effectiveness of our solution in Section B.3 of the Supplementary. Provided the promising results we obtain in Section 5, we propose our reparameterization trick as a novel contribution for solving the pathological behavior of HMC combined with hierarchical VAEs.

For the sake of simplicity, we name the generative parameters as  $\theta = \{\theta_z, \theta_y, \theta_1, \dots, \theta_L\}$ . The dependency ELBO under this hierarchical reparameterized model becomes

$$\begin{aligned} \mathcal{L}_{VI}(\mathbf{x}_O, \mathbf{y}_O; \{\theta, \psi\}) &= \mathbb{E}_{q_\psi} \left[ \log \frac{p_\theta(\mathbf{z}_O, \mathbf{y}_O, \epsilon)}{q_\psi(\epsilon|\mathbf{z}_O, \mathbf{x}_O, \mathbf{y}_O)} \right] = \\ &= \mathbb{E}_{q_\psi} [\log p_\theta(\mathbf{z}_O|\mathbf{h}_1) + \log p_\theta(\mathbf{y}_O|\hat{\mathbf{x}}, \mathbf{h}_1)] - \sum_{l=1}^L D_{\text{KL}}(q_\psi(\epsilon_l|\mathbf{x}_O, \mathbf{y}_O) || p(\epsilon_l)). \end{aligned} \quad (6)$$

We name  $\mathbf{r}_l$  the hidden representation at each layer, and defining  $\mathbf{r}_0 = \{\mathbf{x}_O, \mathbf{y}_O\}$ , we employ NNs with parameters  $\psi_{r_l}$  for computing  $\mathbf{r}_l = f_r(\mathbf{r}_{l-1})$ . These vectors are mapped into the parameters of the variational posterior  $q_{\psi_l}(\epsilon_l|\mathbf{x}_O, \mathbf{y}_O)$ , using NNs for computing the mean as  $g_{\mu_l}(\mathbf{r}_l)$  and the covariance as  $g_{\sigma_l}(\mathbf{r}_l)$ , with parameters  $\psi_{\mu_l}$  and  $\psi_{\sigma_l}$ . With compactness purposes, we will denote the encoder parameters as  $\psi = \{\psi_1, \dots, \psi_L\}$ , with  $\psi_l = \{\psi_{r_l}, \psi_{\mu_l}, \psi_{\sigma_l}\}$ .

### 3.5 HMC over the hierarchical density

In recent works, HMC has been combined with deep generative models for improving the inference of the latent variables by obtaining better samples from the posterior [7, 19]. In this work, we propose to transcend these previous approaches and build a generalized method for sampling from complicated, hierarchical latent structures composed by several layers. Inspired by [6] and their method for sampling from the posterior within a vanilla VAE framework while tuning the HMC hyperparameters, we follow a procedure for training the dependency model where i) during a pre-training stage, the encoder and decoder are optimized using standard VI and the ELBO from equation (6), and ii) using the pre-trained encoder for starting from a good proposal [19], HMC samples are obtained to follow the true posterior and jointly optimize the generative model and the HMC hyperparameters. In Figure 1 (b) we include an illustrative example.

We denote by  $q_\phi^{(T)}(\epsilon|z_O, \mathbf{x}_O, \mathbf{y}_O)$  the implicit distribution for the posterior after  $T$  HMC steps. The hyper-parameters of HMC are named  $\phi$ , a  $T \times d$  matrix containing the step sizes for each dimension at each step of the chain. Within this perspective, the hyper-parameters can be optimized using variational inference by maximizing the ELBO

$$\mathbb{E}_{q_\phi^{(T)}(\epsilon)}[\log p(z_O, \mathbf{y}_O, \epsilon)] + H[q_\phi^{(T)}(\epsilon|\mathbf{x}_O, \mathbf{y}_O)], \quad (7)$$

where the first part is the HMC objective, and can be easily estimated via Monte Carlo

$$\mathcal{L}_{HMC}(z_O, \mathbf{y}_O; \{\theta, \psi, \phi\}) = \mathbb{E}_{q_\phi^{(T)}(\epsilon)}[\log p_\theta(z_O|\mathbf{h}_1) + \log p_\psi(\mathbf{y}_O|\hat{\mathbf{x}}, \mathbf{h}_1) + \sum_{l=1}^L p(\epsilon_l^{(T)})]. \quad (8)$$

Nevertheless, the entropy term  $H[q_\phi^{(T)}(\epsilon|\mathbf{x}_O, \mathbf{y}_O)]$  in Equation (7) is not tractable since we are not able to explicitly evaluate the distribution  $q_\phi^{(T)}(\epsilon|\mathbf{x}_O, \mathbf{y}_O)$ . Although optimizing the first term might result in a well-posed algorithm, this would bring consequences that must be considered. Namely, without a proper regularization term, and in case the initial proposal  $q_\phi^{(T)}(\epsilon|\mathbf{x}_O, \mathbf{y}_O)$  is concentrated in high density regions, the chains would barely move from the initial state and only these regions with high density would be explored (see Section B.1 of the Supplementary for illustrative details). To cope with this problem, we define an inflation parameter  $s$  to increase the variance of the proposal  $q_\phi(\epsilon|z_O, \mathbf{x}_O, \mathbf{y}_O)$  given by the Gaussian encoder, ending in the proposal  $\prod_{l=1}^L \mathcal{N}(g_{\mu_l}(\mathbf{r}_l), s_l \cdot g_{\sigma_l}(\mathbf{r}_l))$ . Whilst in [6] the authors define this parameter as a scalar factor applied to all the latent dimensions, in our work, we extend this to apply a different inflation factor at each latent level  $s = \{s_1, \dots, s_L\}$ . In order to tune these inflations we ensure a wider coverage of the density by minimizing the Sliced Kernelized Stein Discrepancy (SKSD) [17]

$$\mathcal{L}_{SKSD}(\mathbf{x}_O, \mathbf{y}_O; \mathbf{s}) = \text{SKSD}\left(q_\phi^{(T)}(\epsilon|z_O, \mathbf{x}_O, \mathbf{y}_O; \mathbf{s}), p(\epsilon|z_O, \mathbf{x}_O, \mathbf{y}_O)\right), \quad (9)$$

which fits perfectly our requirements, since only requires samples from HMC and gradients  $\nabla_\epsilon \log p(z_O, \mathbf{y}_O, \epsilon)$  for measuring a discrepancy between the implicit and the true posterior. Further, the SKSD performs better than other discrepancies like [31] in high dimensional latent spaces.

We provide in Section B.1 of the Supplementary a toy demonstration on the efficacy of the HMC optimization, and in Section B.2 a demonstration of the optimization convergence.

### 3.6 HH-VAEM optimization

The optimization of the HH-VAEM is divided into three stages. In a first stage, we train one independent *marginal* VAE per dimension. In a second stage, the *dependency* VAE is trained, using as inputs the concatenation of the encoded dimensions  $z = \{z_1, \dots, z_D\}$  and the target  $\mathbf{y}$ . Finally, in a third stage, the HMC hyperparameters, the decoder and the predictor are tuned using the HMC objective, the inflation parameter is trained using the SKSD discrepancy, and the encoder is trained using the ELBO. The pseudocode for HH-VAEM training is shown in Algorithm 1.

### 3.7 Computational cost

In a VAE, the data is fed to the encoder with aim at obtaining the variational parameters for sampling from the Gaussian approximated posterior. In our method, the data is similarly encoded to obtain the

initial Gaussian proposal  $q^{(0)}$ , and the samples from this distribution are updated for  $T$  cycles to obtain the implicit  $q^{(T)}$  using HMC. Within each of these iterations,  $L$  leapfrog steps (1) are executed. For each of these steps, the computation of the gradients  $\nabla_{\epsilon_t} \log p(\mathbf{z}, \mathbf{y}, \epsilon_t)$  and  $\nabla_{\epsilon_{t+1}} \log p(\mathbf{z}, \mathbf{y}, \epsilon_{t+1})$  is required. To obtain these gradients, we need to *i*) compute the parameters of the likelihood  $p(\mathbf{z}, \mathbf{y} | \epsilon)$  that are given by the decoder ( $p(\mathbf{z} | \mathbf{h}_1)$ ) and predictor ( $p(\mathbf{y} | \hat{\mathbf{x}}, \mathbf{h}_1)$ ), *ii*) evaluate the likelihood and *iii*) perform the automatic differentiation. Thus, for running our method, an additional cost from both decoding and performing differentiation a total of  $2TL$  times is introduced.

By jointly optimizing the HMC hyperparameters we are able to achieve faster convergence with smaller lengths. To reduce the computational cost, we optimize the hyperparameters in a final training stage, since convergence is rapidly achieved (as demonstrated empirically in Section B.2 of the Supplementary). At test, samples from the Gaussian  $q^{(0)}$  (faster and cheaper), or from HMC  $q^{(T)}$  (slower and better) can be used to fit computational constraints.

## 4 Sampling-based Active Learning

Considering that the input data are tuples of observed and missing features  $\{\mathbf{x}_O, \mathbf{x}_U\}$ , our Active Learning framework follows [33, 34] and determines which feature  $x_i \in \mathbf{x}_U$  will enhance the prediction of the target  $\mathbf{y}$  the most for a particular  $\mathbf{x}_O$ . Concretely, in a Sequential Active Information Acquisition (SAIA) scenario, this decision is taken sequentially to optimally increase knowledge and accurately predict  $\mathbf{y}$ . From an information theoretical perspective, this task can be performed recurrently by maximizing a reward function  $R$  at each step  $d$ . This reward might represent abstract quantities of interest like the cost or benefit of acquiring  $x_i$  (depending on the sign). In Bayesian experimental analysis,  $R$  is the expected gain of information [29]. Following [3], we can define it as

$$R(i, \mathbf{x}_O) = \mathbb{E}_{p(x_i | \mathbf{x}_O)} D_{\text{KL}}(p(\mathbf{y} | x_i, \mathbf{x}_O) p(\mathbf{y} | \mathbf{x}_O)), \quad (10)$$

where  $i$  is the index of each unobserved feature. Intuitively, this quantity can be interpreted as the expected change in the predictive distribution when  $x_i$  is observed. The reward needs to be estimated via Monte Carlo by sampling from  $p(x_i | \mathbf{x}_O)$ . With a robust generative model that handles missing data like HH-VAEM, these samples are easily obtained: first, using HMC, we sample  $\epsilon^{(T)}$  from  $p(\epsilon | \mathbf{x}_O)$ . Second, we decode these samples to obtain  $x_i$  from  $p(\mathbf{z} | \mathbf{h}_1)$  and  $p(x_i | z_d)$ . Nonetheless, the reward defined in (10) is intractable since both  $p(\mathbf{y} | x_i)$  and  $p(\mathbf{y} | x_i, \mathbf{x}_O)$  are intractable: computing them requires to integrate out the latent variables. This motivates the authors of [33, 34] to present a transformation of the reward for being computed in the latent space using the encoder network. Although they prove that this transformation effectively provides a good estimation in several datasets, we demonstrate that for low dimensional targets (commonly one or two dimensions), an approximation using histograms is more effective. Concretely, the reward in (10) can be rewritten as

$$D_{\text{KL}}[p(\mathbf{y}, x_i | \mathbf{x}_O) || p(\mathbf{y} | \mathbf{x}_O) p(x_i | \mathbf{x}_O)] = \mathcal{I}(\mathbf{y}; x_i | \mathbf{x}_O), \quad (11)$$

While a set of advanced non-parametric estimators of the mutual information are available [26, 42], many are not easily adapted for parallelization. We demonstrate that the simplest one,

$$\hat{I}(\mathbf{y}; x_i | \mathbf{x}_O) \approx \sum_{ij} p(i, j) \log \frac{p(i, j)}{p_x(i) p_y(j)}, \quad (12)$$

based on binning the  $x_i$  and  $\mathbf{y}$  domains, is effective and easy to parallelize. In the equation,  $p_x(i) = n_x(i)/N$ ,  $p_y(j) = n_y(j)/N$  and  $p(i, j) = n(i, j)/N$  are the relative frequencies that approximate

---

### Algorithm 1 Training algorithm for HH-VAEM

---

**Input:** data  $(\mathbf{x}_O^{(1:N)}, \mathbf{y}_O^{(1:N)})$ , steps:  $T_d, T_{VI}, T_{HMC}$

**Parameters:**  $\gamma, \theta, \psi, \phi, s$

STAGE 1: MARGINAL VAES

**for**  $d = 1$  **to**  $D$  **do**

    Initialize marginal VAE  $\{\theta_d, \gamma_d\}$

**for**  $t = 1$  **to**  $T_d$  **do**

$\gamma_d^{t+1}, \theta_d^{t+1} \leftarrow \text{Adam}_{\gamma_d^t, \theta_d^t}(\mathcal{L}_d)$

**end for**

**end for**

STAGE 2: DEPENDENCY VAE

**for**  $t = 1$  **to**  $T_{VAE}$  **do**

$\theta^{t+1}, \psi^{t+1} \leftarrow \text{Adam}_{\theta^t, \psi^t}(\mathcal{L}_{VI})$

**end for**

STAGE 3: JOINTLY OPTIMIZING VAE + HMC

**for**  $t = 1$  **to**  $T_{HMC}$  **do**

$\psi^{t+1} \leftarrow \text{Adam}_{\psi^t}(\mathcal{L}_{VI})$

$\theta^{t+1}, \phi^{t+1} \leftarrow \text{Adam}_{\theta^t, \phi^t}(\mathcal{L}_{HMC})$

$s^{t+1} \leftarrow \text{Adam}_{s^t}(\mathcal{L}_{SKSD})$

**end for**

---

	Bank	Insurance	Avocado	Naval	Yatch	Diabetes	Concrete	Wine	Energy	Boston
VAEM	2.84 ± 0.07	1.81 ± 0.03	1.89 ± 0.01	0.55 ± 0.05	3.15 ± 0.28	2.78 ± 0.16	2.45 ± 0.26	3.01 ± 0.61	2.09 ± 0.10	2.01 ± 0.23
MIWAEM	2.74 ± 0.05	1.88 ± 0.04	1.92 ± 0.04	0.57 ± 0.03	2.66 ± 0.11	2.55 ± 0.09	2.34 ± 0.51	2.76 ± 0.48	2.06 ± 0.14	1.94 ± 0.23
H-VAEM	2.82 ± 0.06	1.80 ± 0.04	1.89 ± 0.01	0.48 ± 0.06	3.06 ± 0.31	2.74 ± 0.09	2.42 ± 0.21	2.85 ± 0.56	1.72 ± 0.11	1.89 ± 0.24
HMC-VAEM	2.69 ± 0.05	1.77 ± 0.06	1.89 ± 0.02	0.49 ± 0.07	<b>2.21 ± 0.24</b>	2.72 ± 0.20	2.28 ± 0.29	2.83 ± 0.46	1.73 ± 0.05	1.83 ± 0.16
<b>HH-VAEM</b>	<b>2.63 ± 0.04</b>	<b>1.75 ± 0.03</b>	<b>1.88 ± 0.05</b>	<b>0.40 ± 0.05</b>	2.47 ± 0.27	<b>2.54 ± 0.13</b>	<b>2.28 ± 0.09</b>	<b>1.90 ± 0.17</b>	<b>1.71 ± 0.04</b>	<b>1.83 ± 0.11</b>

Table 1: Test NLL of the unobserved features for our model and baselines.

	Bank	Insurance	Avocado	Naval	Yatch	Diabetes	Concrete	Wine	Energy	Boston
VAEM	0.56 ± 0.06	1.20 ± 0.03	1.18 ± 0.02	2.69 ± 0.01	0.61 ± 0.02	1.59 ± 0.19	1.07 ± 0.09	0.28 ± 0.09	0.61 ± 0.14	0.85 ± 0.21
MIWAEM	0.51 ± 0.03	1.15 ± 0.03	1.15 ± 0.03	2.70 ± 0.01	0.60 ± 0.03	<b>1.36 ± 0.10</b>	0.95 ± 0.22	0.28 ± 0.13	0.54 ± 0.12	0.80 ± 0.21
H-VAEM	0.50 ± 0.03	1.06 ± 0.02	1.18 ± 0.02	2.68 ± 0.01	0.60 ± 0.02	1.71 ± 0.14	1.02 ± 0.09	0.26 ± 0.11	0.46 ± 0.14	0.90 ± 0.22
HMC-VAEM	0.52 ± 0.02	1.00 ± 0.03	1.12 ± 0.03	2.71 ± 0.01	<b>0.52 ± 0.15</b>	1.55 ± 0.29	0.95 ± 0.26	0.28 ± 0.09	0.41 ± 0.07	0.71 ± 0.13
<b>HH-VAEM</b>	<b>0.49 ± 0.03</b>	<b>0.93 ± 0.06</b>	<b>1.10 ± 0.01</b>	<b>2.62 ± 0.01</b>	0.56 ± 0.02	1.38 ± 0.18	<b>0.95 ± 0.08</b>	<b>0.20 ± 0.04</b>	<b>0.32 ± 0.05</b>	<b>0.55 ± 0.04</b>


Table 2: Test NLL of the predicted target for our model and baselines.

$p(x)$ ,  $p(y)$  and  $p(x, y)$  for each bin. Thus,  $n_x(i)$ ,  $n_y(j)$  and  $n(i, j)$  are the number of samples inside each interval. The number of bins defines the width of uniformly distributed intervals over  $x_d$  and  $y$  supports. Since this estimator is sampling-based, under certain conditions (namely, if all densities exist as proper functions), Equation (13) indeed converges to  $\mathcal{I}(y; x_i | x_O)$  if we first let the number of samples  $N \rightarrow \infty$  [26].


## 5 Experiments

The evaluation of the HH-VAEM model is organized into three quantitative experiments and one qualitative experiment. The ablation study includes the validation of our proposed HMC-based, hierarchical model with respect to the Gaussian and one-layered alternatives. Namely, the comparison is performed with the following baseline models:

- *VAEM*: 1 layer, Gaussian-based VAEM [33] (without including the Partial VAE).
- *MIWAEM*: 1 layer, Gaussian-based, importance weighted IWAEM (VAEM + IWAE [36]).
- *H-VAEM*: 2 layers, Gaussian-based VAEM.
- *HMC-VAEM*: 1 layer, HMC-based (with our optimization method) VAEM.



(a) Autoregressive



(b) Reparameterization

	VAE	MIWAE	H-VAE	HMC-VAE	<b>HH-VAE</b>
MNIST	0.124 ± 0.001	0.121 ± 0.001	0.119 ± 0.001	0.101 ± 0.004	<b>0.094 ± 0.003</b>
F-MNIST	0.162 ± 0.002	0.160 ± 0.002	0.156 ± 0.002	0.150 ± 0.002	<b>0.144 ± 0.002</b>

Table 3: Test NLL of the unobserved features for the MNIST datasets.

	VAE	MIWAE	H-VAE	HMC-VAE	<b>HH-VAE</b>
MNIST	0.153 ± 0.009	0.151 ± 0.007	0.146 ± 0.006	0.067 ± 0.007	<b>0.056 ± 0.019</b>
F-MNIST	0.501 ± 0.012	0.496 ± 0.008	0.494 ± 0.007	0.357 ± 0.060	<b>0.337 ± 0.069</b>

Table 4: Test NLL of the predicted target for the MNIST datasets.

	VAE	MIWAE	H-VAE	HMC-VAE	<b>HH-VAE</b>
MNIST	0.953 ± 0.004	0.953 ± 0.003	0.953 ± 0.003	0.978 ± 0.003	<b>0.981 ± 0.005</b>
F-MNIST	0.824 ± 0.005	0.824 ± 0.004	0.824 ± 0.004	0.869 ± 0.015	<b>0.876 ± 0.017</b>

Figure 4: AR (a) vs reparameterized (b). Table 5: Test accuracy of the predicted digits for the MNIST datasets.

For all the models, we manually introduce missing data in the training set by randomly setting per data point a feature as missing with a probability sampled uniformly in the interval  $[0.01, 0.99]$  within each batch. Both the input data  $x$  and the target  $y$  can be missing. For the test set, a fixed probability of 0.5 leads to about half of the input data being observed, whilst the target is completely unobserved.

For the quantitative experiments, a total of 10 UCI datasets [11] that include mixed-type data are employed for the evaluation. We include both MNIST [27] and Fashion-MNIST [55] datasets for evaluating our model in higher dimensional observational and latent spaces and bigger architectures (3 layered convolutional nets for encoder/decoder). For the qualitative results, we evaluate our model in the image inpainting task on MNIST and CelebA [32]. For the three image datasets, the marginal



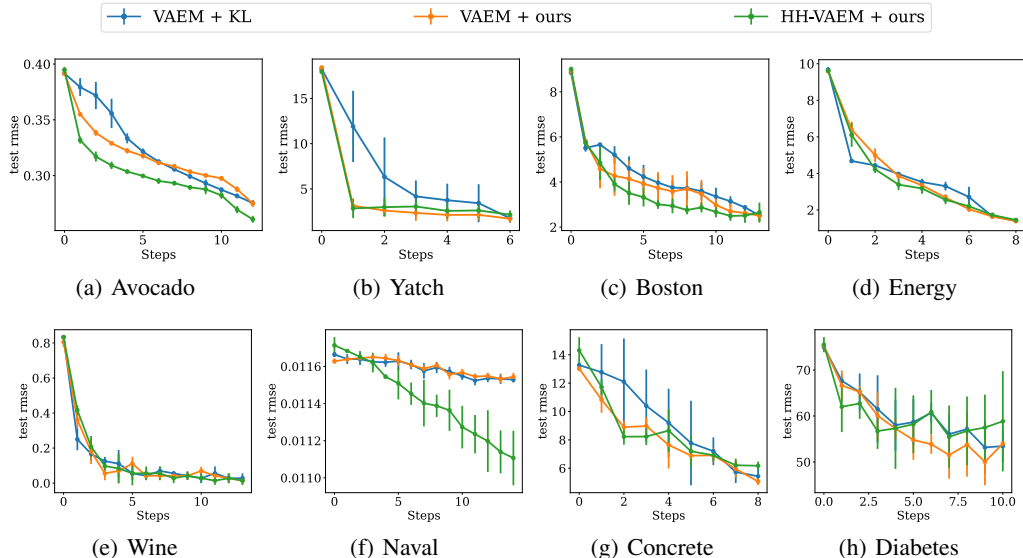


Figure 5: SAIA curves. Horizontal axis shows number of discovered features. Vertical axis is RMSE. VAEs are not included and the dependency VAE is fed directly with the Bernoulli-distributed pixels. We name this model HH-VAE, and similarly, the baselines are renamed as VAE, MIWAE, HMC-VAE and H-VAE. Extended experiments and validations are provided in the Supplementary. The source code for reproducing our work is available at <https://github.com/ipeis/HH-VAEM>.

### 5.1 Mixed type conditional data imputation

In order to evaluate the performance of the model in terms of data imputation, we opt by computing the negative log likelihood of the unobserved features. We make use of the Monte Carlo approximation

$$\log p(\mathbf{x}_U | \mathbf{x}_O) \approx \log \mathbb{E}_{\epsilon \sim q(\epsilon | \mathbf{x}_O)} [p(\mathbf{x}_U | \epsilon)] \approx \log \frac{1}{k} \sum_i^k p(\mathbf{x}_U | \epsilon_i), \quad (13)$$

which is averaged over features in order to compare the imputation performance with the baselines. Additionally, we include in Section C.3 similar results averaging each of the considered likelihoods. Results on the 10 UCI datasets and the MNIST datasets are included in Tables 1 and 3, showing that for most of the datasets, incremental improvement is obtained:  $\text{VAEM} < \text{H-VAEM} < \text{HMC-VAEM} < \text{HH-VAEM}$ . Extended results with the imputation error, included in Section C.1, corroborate this.

### 5.2 Target prediction

For this experiment, we compute the negative log likelihood of the target under the predictive distribution using the same procedure as in the previous section. Results included in Tables 2 and 4 show the same incremental improvement in the prediction task.

### 5.3 Sequential active information acquisition (SAIA)


In this experiment, our HH-VAEM model and our acquisition method are evaluated in a SAIA task. Starting by predicting from completely unobserved inputs, at each step, the missing feature that maximizes the reward is acquired. Figure 5 shows the error curves for the UCI datasets. Blue lines correspond to the Gaussian-based reward proposed by [33]. Orange lines are our sampling-based reward in a VAEM framework. Green lines correspond to HH-VAEM with our reward. In most of the cases, our model and acquisition method obtains lower errors and faster discovery of information.

### 5.4 Conditional image inpainting

We include in this experiment qualitative results when comparing our model with the baselines on the image conditional inpainting task. We include results for MNIST and CelebA in Figure 6 (a)



(a) MNIST



(b) CelebA

Figure 6: Image conditional inpainting on MNIST (a) and CelebA (b). First row: original images from the test set. Second row: input to the model,  $x_O$ , with a black square missing mask  $x_U$  manually introduced. Third, fourth, fifth and sixth rows: imputed  $\hat{x}_U$  for VAE, H-VAE, HMC-VAE and HH-VAE, our model, where  $\hat{x}_U$  is decoded from samples of the approximate posterior (Gaussian for VAE and H-VAE, or HMC-based for HMC-VAE and HH-VAE).

and (b), respectively, that show the superiority of our method. First, the hierarchical Gaussian model (fourth row) considerably improves the one-layered Gaussian alternative. Second, the HMC-based methods (two last rows) vastly improve the Gaussian methods. Third, in some specific cases, an extra improvement is added by HH-VAE with respect to the one-layered HMC-based model (columns 2, 4, 8, 9, 11, 12 and 13 in Figure 6 (a) or columns 1, 2, 5 and 12 in (b)).

## 6 Conclusion

We presented HH-VAEM, to our knowledge, the first hierarchical VAE for mixed-type incomplete data that uses HMC with automatic hyper-parameter tuning for improved inference. We provide both quantitative and qualitative experiments that demonstrate its superiority with respect the baselines in the tasks of missing data imputation and supervised learning, placing HH-VAEM as a robust model for real-world datasets. Further, we have developed a novel sampling-based technique for dynamic feature selection that outperforms the Gaussian-based alternatives and results in an efficient method for active learning in deep generative models.

## Acknowledgments and Disclosure of Funding

Ignacio Peis acknowledges support from Spanish government Ministerio de Ciencia, Innovación y Universidades under grants FPU18/00516, RTI2018-099655-B-I00, EST21/00467 and PID2021-123182OB-I00 and from Comunidad de Madrid under grant Y2018/TCS-4705 PRACTICO-CM. José Miguel Hernández-Lobato acknowledges support from Boltzbit.

## References

- [1] D. Barrejón, P. M. Olmos, and A. Artés-Rodríguez. Medical data wrangling with sequential Variational Autoencoders. *arXiv preprint arXiv:2103.07206*, 2021.
- [2] Y. Bengio. *Learning deep architectures for AI*. Now Publishers Inc, 2009.
- [3] J. M. Bernardo. Expected information as expected utility. *the Annals of Statistics*, pages 686–690, 1979.
- [4] M. Betancourt. A conceptual introduction to Hamiltonian Monte Carlo. *arXiv preprint arXiv:1701.02434*, 2017.
- [5] M. Betancourt and M. Girolami. Hamiltonian Monte Carlo for hierarchical models. *Current trends in Bayesian methodology with applications*, 79(30):2–4, 2015.
- [6] A. Campbell, W. Chen, V. Stimper, J. M. Hernandez-Lobato, and Y. Zhang. A gradient Based Strategy for Hamiltonian Monte Carlo Hyperparameter optimization. In *International Conference on Machine Learning*, pages 1238–1248. PMLR, 2021.
- [7] A. L. Caterini, A. Doucet, and D. Sejdinovic. Hamiltonian Variational Auto-Encoder. *arXiv preprint arXiv:1805.11328*, 2018.
- [8] R. Child. Very deep VAEs generalize autoregressive models and can outperform them on images. *arXiv preprint arXiv:2011.10650*, 2020.
- [9] M. Collier, A. Nazabal, and C. K. Williams. VAEs in the presence of missing data. *arXiv preprint arXiv:2006.05301*, 2020.
- [10] C. Cremer, X. Li, and D. Duvenaud. Inference Suboptimality in Variational Autoencoders. In *International Conference on Machine Learning*, pages 1078–1086. PMLR, 2018.
- [11] D. Dua, C. Graff, et al. UCI machine learning repository. 2017.
- [12] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Physics letters B*, 195(2):216–222, 1987.
- [13] S. Eduardo, A. Nazabal, C. K. Williams, and C. Sutton. Robust Variational Autoencoders for outlier detection and repair of mixed-type data. In *International Conference on Artificial Intelligence and Statistics*, pages 4056–4066. PMLR, 2020.
- [14] M. Garnelo, D. Rosenbaum, C. Maddison, T. Ramalho, D. Saxton, M. Shanahan, Y. W. Teh, D. Rezende, and S. A. Eslami. Conditional Neural Processes. In *International Conference on Machine Learning*, pages 1704–1713. PMLR, 2018.
- [15] Z. Ghahramani and M. I. Jordan. Learning from incomplete data. 1995.
- [16] M. Girolami and B. Calderhead. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.
- [17] W. Gong, Y. Li, and J. M. Hernández-Lobato. Sliced Kernelized Stein Discrepancy. *arXiv preprint arXiv:2006.16531*, 2020.
- [18] Y. Gong, H. Hajimirsadeghi, J. He, T. Durand, and G. Mori. Variational Selective Autoencoder: Learning from partially-observed heterogeneous data. In *International Conference on Artificial Intelligence and Statistics*, pages 2377–2385. PMLR, 2021.
- [19] M. D. Hoffman. Learning deep latent Gaussian models with Markov Chain Monte Carlo. In *International conference on machine learning*, pages 1510–1519. PMLR, 2017.
- [20] S.-J. Huang, M. Xu, M.-K. Xie, M. Sugiyama, G. Niu, and S. Chen. Active feature acquisition with supervised matrix completion. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1571–1579, 2018.

- [21] N. Ipsen, P.-A. Mattei, and J. Frellsen. How to deal with missing data in supervised deep learning? In *Artemiss-ICML Workshop on the Art of Learning with Missing Values*, 2020.
- [22] T. Joy, S. Schmon, P. Torr, S. Narayanaswamy, and T. Rainforth. Capturing label characteristics in vaes. In *Proceedings of the ICLR Conference 2021*. OpenReview, 2021.
- [23] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling. Improved Variational Inference with Inverse Autoregressive Flow. *Advances in neural information processing systems*, 29:4743–4751, 2016.
- [24] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [25] D. P. Kingma and M. Welling. An introduction to Variational Autoencoders. *arXiv preprint arXiv:1906.02691*, 2019.
- [26] A. Kraskov, H. Stögbauer, and P. Grassberger. Estimating mutual information. *Physical review E*, 69(6):066138, 2004.
- [27] Y. LeCun. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [28] Y. Li, J. Bradshaw, and Y. Sharma. Are generative classifiers more robust to adversarial attacks? In *International Conference on Machine Learning*, pages 3804–3814. PMLR, 2019.
- [29] D. V. Lindley. On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, pages 986–1005, 1956.
- [30] R. J. Little and D. B. Rubin. *Statistical analysis with missing data*, volume 793. John Wiley & Sons, 2019.
- [31] Q. Liu, J. Lee, and M. Jordan. A Kernelized Stein Discrepancy for goodness-of-fit tests. In *International conference on machine learning*, pages 276–284. PMLR, 2016.
- [32] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- [33] C. Ma, S. Tschitschek, J. M. Hernández-Lobato, R. Turner, and C. Zhang. VAEM: a Deep Generative Model for Heterogeneous Mixed Type Data. *arXiv preprint arXiv:2006.11941*, 2020.
- [34] C. Ma, S. Tschitschek, K. Palla, J. M. Hernández-Lobato, S. Nowozin, and C. Zhang. EdDI: Efficient Dynamic Discovery of High-Value Information with Partial VAE. *arXiv preprint arXiv:1809.11142*, 2018.
- [35] L. Maaløe, M. Fraccaro, V. Liévin, and O. Winther. BIVA: A very deep hierarchy of latent variables for generative modeling. *arXiv preprint arXiv:1902.02102*, 2019.
- [36] P.-A. Mattei and J. Frellsen. MIWAE: Deep generative modelling and imputation of incomplete data sets. In *International Conference on Machine Learning*, pages 4413–4423. PMLR, 2019.
- [37] P. Melville, M. Saar-Tsechansky, F. Provost, and R. Mooney. Active feature-value acquisition for classifier induction. In *Fourth IEEE International Conference on Data Mining (ICDM'04)*, pages 483–486. IEEE, 2004.
- [38] A. Nazabal, P. M. Olmos, Z. Ghahramani, and I. Valera. Handling Incomplete Heterogeneous Data using VAEs. *Pattern Recognition*, 107:107501, 2020.
- [39] R. M. Neal et al. MCMC using Hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11):2, 2011.
- [40] A. Razavi, A. v. d. Oord, B. Poole, and O. Vinyals. Preventing posterior collapse with delta-VAEs. *arXiv preprint arXiv:1901.03416*, 2019.

- [41] D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pages 1278–1286. PMLR, 2014.
- [42] B. C. Ross. Mutual information between discrete and continuous data sets. *PloS one*, 9(2):e87357, 2014.
- [43] M. Saar-Tsechansky, P. Melville, and F. Provost. Active feature-value acquisition. *Management Science*, 55(4):664–684, 2009.
- [44] R. Salakhutdinov. Learning deep generative models. *Annual Review of Statistics and Its Application*, 2:361–385, 2015.
- [45] R. Salakhutdinov and G. Hinton. Deep Boltzmann machines. In *Artificial intelligence and statistics*, pages 448–455. PMLR, 2009.
- [46] T. Salimans, D. Kingma, and M. Welling. Markov Chain Monte Carlo and Variational Inference: Bridging the gap. In *International Conference on Machine Learning*, pages 1218–1226. PMLR, 2015.
- [47] M. Śmieja, Ł. Struski, J. Tabor, B. Zieliński, and P. Spurek. Processing of missing data by neural networks. *Advances in neural information processing systems*, 31, 2018.
- [48] C. K. Sønderby, T. Raiko, L. Maaløe, S. K. Sønderby, and O. Winther. Ladder Variational Autoencoders. *Advances in neural information processing systems*, 29:3738–3746, 2016.
- [49] D. J. Stekhoven and P. Bühlmann. Missforest—non-parametric missing value imputation for mixed-type data. *Bioinformatics*, 28(1):112–118, 2012.
- [50] M. Thahir, T. Sharma, and M. K. Ganapathiraju. An efficient heuristic method for active feature acquisition and its application to protein-protein interaction prediction. In *BMC proceedings*, volume 6, pages 1–9. BioMed Central, 2012.
- [51] A. Thin, N. Kotelevskii, A. Doucet, A. Durmus, E. Moulines, and M. Panov. Monte carlo variational auto-encoders. In *International Conference on Machine Learning*, pages 10247–10257. PMLR, 2021.
- [52] V. Tresp, S. Ahmad, and R. Neuneier. Training neural networks with deficient data. *Advances in neural information processing systems*, 6, 1993.
- [53] A. Vahdat and J. Kautz. NVAE: A Deep Hierarchical Variational Autoencoder. *arXiv preprint arXiv:2007.03898*, 2020.
- [54] Y. Wang and J. P. Cunningham. Posterior collapse and latent variable non-identifiability. In *Third Symposium on Advances in Approximate Bayesian Inference*, 2020.
- [55] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [56] C. Zhang, J. Bøtche, H. Kjellström, and S. Mandt. Advances in Variational Inference. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):2008–2026, 2018.

## A Experimental details

### A.1 Experimental setup

The networks for the encoder of the model with the MNIST datasets are 2 layered Deep CNNs with  $\{16, 32, 32\}$  output channels, kernel size 5, stride 2 and padding 2. For the experiments with CelebA, we use 5 layered Deep CNNs with  $\{32, 32, 64, 64, 512\}$  output channels, kernel size 4, stride 2 and padding 1, followed by batch norm layers. They are followed by MLPs with 512 hidden units for obtaining the variational parameters for each layer. The decoder that obtains  $p_\theta(\mathbf{x}|\mathbf{h}_1)$  is the symmetric CNN. All the NNs employed in the models trained with UCI datasets are one single layer MLPs with 256 hidden units. The noise variance for Gaussian likelihoods is set up to 0.1.

We employ learning rates of  $1 \times 10^{-3}$  for the models with MLP networks and  $2 \times 10^{-4}$  for the convolutional models. For the inflation parameter  $s$ , we increase to  $1 \times 10^{-2}$  for a faster convergence. A batch size of 100 is used for all the models except for Yatch and Wine dataset, where we use 20 samples per batch. The number of training steps is  $2 \times 10^4$  for Boston, Energy, Wine, Yatch, Concrete, Diabetes and Yatch, and  $5 \times 10^4$  for Naval, Avocado, Bank, and Insurance. For MNIST and Fashion-MNIST, we have  $1 \times 10^5$  training steps. For CelebA, we use  $1, 5 \times 10^5$  training steps. For the marginal VAEs stage, we employ  $1 \times 10^3$  training steps. The dimension of the latent variables is  $[d_1 = 10, d_2 = 5]$  for Boston, Energy, Wine, Naval Avocado, Bank and Insurance  $[d_1 = 4, d_2 = 2]$  for Concrete, Yatch and Diabetes,  $[d_1 = 20, d_2 = 10]$  for MNIST and Fashion-MNIST, and  $[d_1 = 32, d_2 = 16]$  for CelebA.

We use  $LF = 5$  Leapfrog steps in all cases, chains of  $T = 10$  for Boston, Energy, Wine, Naval Avocado, Bank, Insurance, MNIST, Fashion-MNIST and CelebA, and  $T = 5$  for Concrete, Yatch and Diabetes. The SKSD function is estimated using 30 HMC samples.

For the MNIST and CelebA datasets, the use of Nvidia P100 GPU with Pascal architecture sped up the training with the CNN-based architecture. For the UCI datasets, due to the use of small networks, the differences when using CPU or GPU are almost imperceptible.

### A.2 Datasets information

Apart from the MNIST, Fashion-MNIST and CelebA datasets, a total of 10 UCI datasets have been employed in this work including: Bank Marketing, Insurance Company Benchmark, Avocado sales, Naval Propulsion Plants, Yatch Hydrodynamics, Diabetes, Boston Housing, Wines, Energy efficiency and Bank Marketing.

### A.3 Balancing the KLs

Following [53], we define a short initial warming stage (10% of the total training steps) during the optimization where the KLs for the different layers are balanced according to their magnitude and the corresponding latent dimension, preventing the model for falling into posterior collapse by ignoring deepest layers. A factor is applied to each KL, following

$$\gamma_l = \frac{d_l \mathbb{E}_{\mathbf{x} \sim B} [\text{KL}(q(\epsilon_l|\mathbf{x})||p(\epsilon))]}{\sum_{i=1}^L d_i \mathbb{E}_{\mathbf{x} \sim B} [\text{KL}(q(\epsilon_i|\mathbf{x})||p(\epsilon))]} \quad (14)$$

The factors penalises the fact that a layer might be ignored by making the KL smaller when its magnitude is small compared to the rest layers.

## B Hamiltonian Monte Carlo with automatic optimization

### B.1 Efficacy of training HMC

We include in this experiment results that demonstrate the efficacy of training the HMC hyperparameters using our proposed gradient-based strategy on 2D densities. In Figure 7, each row correspond to a different example: first row is a *wave* density, second row is a *dual moon* density. In the first column, the initial set up is showed, including the density contour (dark blue), the initial Gaussian proposal contour (light blue) and samples from HMC (green). In both cases, due to the tightness of the proposal, chains do not properly explore the density and get stuck close the initial state. In the

second column, results after training HMC hyperparameters and the inflation parameter are included. Again, in both cases, and more specially in the first one, the inflation of the horizontal variance of the proposal successfully increases to better cover the surface, and final HMC samples vastly improve the exploration. In the third column, the approximations of the HMC objective, the SKSD discrepancy and the inflation parameters over the optimization steps are included.

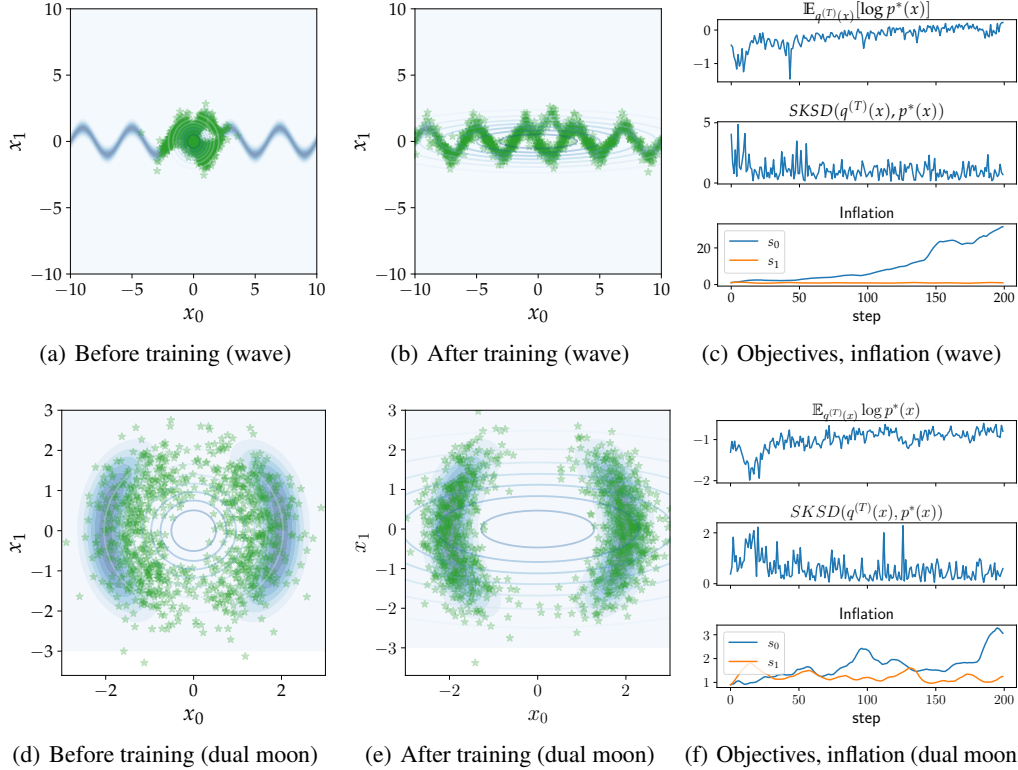


Figure 7: Toy example showing the efficacy of training HMC hyperparameters using our method on two densities: wave (top row) with  $T = 5$  and dual moon (bottom row) with  $T = 10$ . Left column illustrates the non desired behavior when chains hardly explore the density and stuck in a small region close to the mass of the tight Gaussian initial proposal (light blue contour ellipses). Right column shows how optimizing the step sizes and the inflation parameter leads to a vast improvement of the exploration. More specifically, (b) justifies scaling each dimension of the target, since the inflation is bigger on the horizontal axis.

## B.2 HMC optimization

We face the computational cost of running HMC by defining a small percentage of training steps for the last stage in Algorithm 1. A 10% of the total training steps for  $T_{HMC}$  is sufficient for obtaining the convergence. In Figure 8 we include the validation metrics obtained during the optimization of Yatch dataset, where  $T_{HMC} = 2 \times 10^3$ .

In figure 9 (a), the mean acceptance rate of the HMC sampler over the training steps and the step sizes (b) are included. The steps are initialized from  $U(0.05, 0.2)$ . After  $2 \times 10^3$  steps, the mean acceptance rate converges to a value closer to  $\bar{p}_a = 0.65$ , which is defined as the optimal desired acceptance probability [39]. This empirical result provides evidence that, apart from reducing the computational cost, reducing the HMC training step to this value is sufficient for achieving convergence.

## B.3 Reparameterization trick for solving ill-posed HMC with hierarchical densities

We provide in this section strong empirical demonstration on the efficacy of our proposed reparameterization trick. As stated in Section 3.4, naïve implementations of HMC are ill-posed when combined




Figure 8: Validation curves during optimization for Naval dataset.




Figure 9: Evolution of mean acceptance rate  $\bar{p}_a$  (a) and step sizes (b) over the training steps.

with hierarchical densities. The autoregressive correlations lead to non-smooth densities with huge peaks. The large gradients  $\nabla_{\mathbf{h}_{1:L}} \log p^*(\mathbf{h}_{1:L})$  evaluated on these regions inside the Leapfrog steps of Equations 1 make huge modifications of the proposed states. If we denote these diverged states by  $\mathbf{h}_{1:L}^{(ill)}$ , evaluating the objective  $\log p^*(\mathbf{h}_{1:L}^{(ill)})$  lead to overflow issues. This undesired behavior is what we call *divergence of the Leapfrog integrator*, and is also illustrated by [4] in Figure 35.

In order to avoid the aforementioned overflow issues, we can directly reject these problematic states. Nevertheless, by doing this, HMC will not properly explore the density by the time the states fall into the problematic regions, leading to extremely low acceptance rates. To demonstrate this, we include in Figure 10 (a) with blue line the evolution of the acceptance rates when optimizing HMC with the HH-VAEM variant without reparameterization (as illustrated in Figure 4 (a)). The acceptance rate is extremely low as expected, which is a clear indicator of a poorly mixing sampler.

On the contrary, by using our proposed reparameterization trick (Figure 4 (b)), we are able to make HMC work properly and tune the step sizes, getting closer to the ideal acceptance rate  $\bar{p}_a = 0.65$  [39]. Evaluations are applied to the same validation split of the Boston dataset. The step sizes of HMC are initialized from  $U(0.05, 0.2)$  in both cases.

Further, we also include in Figure 10 (b) the evolution of the imputation log likelihood metric of Equation (13) during the whole optimization for both approaches. First  $1, 8 \times 10^4$  steps correspond to the pretraining stage using only the ELBO. When introducing HMC without the reparameterization, due to the low acceptance rate, the states hardly move from the initial proposal, or move away from the density, and the joint optimization fail. We demonstrate again the effectiveness of the reparameterization trick by observing an increase in this metric.

## C Extended experiments

### C.1 Deterministic imputation metrics

We include in Table 6 results on the RMSE obtained with other discriminative validated predictors, using mean imputation under the same missing rates. Additionally, we include here the missForest in the baselines, a wide-spread method for missing data imputation using a Random Forest approach [49]. For classification tasks, the error rate is considered. In almost all cases, HH-VAEM outperforms the baselines.



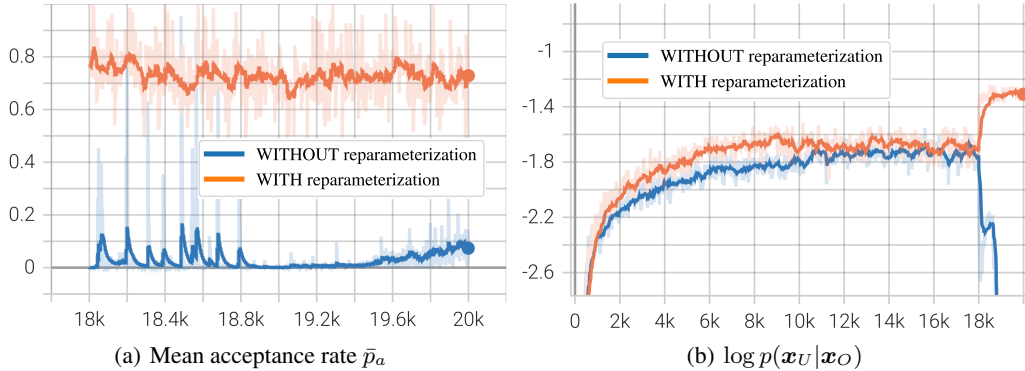


Figure 10: Demonstration of the efficacy of our reparameterization trick. Our model is showed with orange lines, and the same model without the reparameterization with blue lines. In (a), the mean acceptance rate of the HMC proposals over the optimization steps is included. We demonstrate that without the reparameterization, HMC is ill-posed and by rejecting the proposals that make the integrator diverge, the acceptance rate is extremely low, thus not properly exploring the density. In (b), the imputation log likelihood metric is showed for the whole optimization. With the reparameterization trick, we successfully solve the pathological issues, leading to a considerable increase of the metric.

	Bank	Insurance	Avocado	Naval	Yatch	Diabetes	Concrete	Wine	Energy	Boston
missForest	0.64 ± 0.01	0.61 ± 0.06	0.59 ± 0.02	0.30 ± 0.01	<b>0.86 ± 0.12</b>	0.76 ± 0.08	0.76 ± 0.07	0.77 ± 0.11	0.64 ± 0.08	0.59 ± 0.06
VAEM	0.57 ± 0.01	0.40 ± 0.01	0.59 ± 0.00	0.33 ± 0.01	0.93 ± 0.04	0.79 ± 0.02	0.74 ± 0.04	0.64 ± 0.03	0.75 ± 0.02	0.65 ± 0.02
MIWAEM	0.56 ± 0.00	0.39 ± 0.00	0.59 ± 0.01	0.34 ± 0.01	0.94 ± 0.04	0.75 ± 0.01	0.71 ± 0.03	0.63 ± 0.02	0.75 ± 0.02	0.63 ± 0.01
H-VAEM	0.56 ± 0.01	0.39 ± 0.00	0.58 ± 0.01	0.32 ± 0.01	0.92 ± 0.05	0.77 ± 0.01	0.71 ± 0.02	0.60 ± 0.04	0.55 ± 0.03	0.59 ± 0.01
HMC-VAE	0.55 ± 0.01	<b>0.38 ± 0.00</b>	0.58 ± 0.01	0.30 ± 0.02	0.91 ± 0.05	0.76 ± 0.03	0.71 ± 0.02	0.60 ± 0.02	0.55 ± 0.02	0.57 ± 0.02
HH-VAEM	<b>0.54 ± 0.01</b>	<b>0.38 ± 0.00</b>	<b>0.57 ± 0.00</b>	<b>0.29 ± 0.02</b>	0.90 ± 0.00	<b>0.75 ± 0.01</b>	<b>0.70 ± 0.01</b>	<b>0.59 ± 0.04</b>	<b>0.54 ± 0.01</b>	<b>0.56 ± 0.03</b>

Table 6: Test RMSE of the unobserved features for our model and baselines.

## C.2 Likelihood of the observed features

We include in this section results on the negative log likelihood of the observed features

$$\log p(\mathbf{x}_O) \approx \log \mathbb{E}_{\epsilon \sim q^{(T)}(\epsilon | \mathbf{x}_O)} [p(\mathbf{x}_O | \epsilon)] \approx \log \frac{1}{k} \sum_i^k p(\mathbf{x}_O | \epsilon_i). \quad (15)$$

Results are included in Table 7. In almost all the cases, we confirm the incremental superiority when adding each part of our proposed design.

	bank	insurance	avocado	naval	yatch	diabetes	concrete	wine	energy	boston
VAEM	0.51 ± 0.05	0.99 ± 0.05	0.44 ± 0.01	0.21 ± 0.01	0.62 ± 0.13	0.92 ± 0.12	0.63 ± 0.18	0.73 ± 0.18	1.86 ± 0.09	0.56 ± 0.11
MIWAEM	0.63 ± 0.02	1.06 ± 0.03	0.60 ± 0.03	0.33 ± 0.01	0.75 ± 0.07	1.05 ± 0.06	0.76 ± 0.09	0.80 ± 0.06	1.77 ± 0.15	0.67 ± 0.03
H-VAEM	0.40 ± 0.04	0.93 ± 0.04	0.42 ± 0.05	0.19 ± 0.07	0.58 ± 0.09	0.70 ± 0.13	0.53 ± 0.18	0.71 ± 0.15	0.38 ± 0.02	0.49 ± 0.07
HMC-VAE	0.37 ± 0.07	<b>0.92 ± 0.04</b>	0.39 ± 0.06	0.18 ± 0.05	0.54 ± 0.10	<b>0.68 ± 0.07</b>	0.49 ± 0.22	<b>0.55 ± 0.07</b>	0.40 ± 0.06	0.41 ± 0.04
HH-VAEM	<b>0.33 ± 0.03</b>	0.95 ± 0.05	<b>0.36 ± 0.01</b>	<b>0.17 ± 0.04</b>	<b>0.45 ± 0.04</b>	0.68 ± 0.16	<b>0.40 ± 0.16</b>	0.64 ± 0.17	<b>0.37 ± 0.06</b>	<b>0.41 ± 0.04</b>

Table 7: Test NLL of the observed features for our model and baselines.

## C.3 Heterogeneous likelihoods

In experiment 5.1 we reported an average likelihood across heterogeneous variables. Although this quantity is not a valid joint likelihood probability, we employed this average to provide a fair comparison on models that have been trained on the same heterogeneous likelihoods. In this section, we show the comparison on averaging separately the three considered marginal likelihoods (Gaussian, Bernoulli and Categorical) for two of the biggest datasets considered on Tables 8-10. Again, we show the incremental superiority when adding the different design choices of our model.

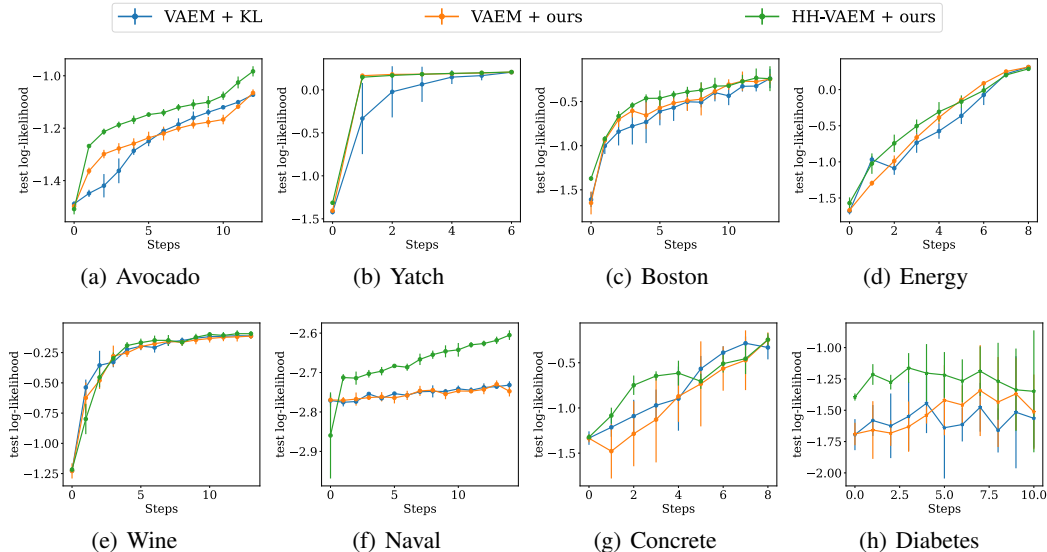


Figure 11: SAIA log-likelihood curves. Horizontal axis shows acquisition steps (number of discovered features). Vertical axis is the log-likelihood of the target  $\log p(\mathbf{y}|\mathbf{x}_O)$ .

	Bank	Avocado		Bank	Avocado		Bank	Avocado
VAEM	$0.36 \pm 0.29$	$0.26 \pm 0.08$	VAEM	$0.13 \pm 0.00$	$0.06 \pm 0.00$	VAEM	$0.24 \pm 0.16$	$0.33 \pm 0.00$
MIWAEM	$0.33 \pm 0.27$	$0.32 \pm 0.06$	MIWAEM	$0.15 \pm 0.00$	$0.09 \pm 0.00$	MIWAEM	$0.26 \pm 0.17$	$0.36 \pm 0.00$
H-VAEM	$0.26 \pm 0.22$	$0.29 \pm 0.07$	H-VAEM	$0.11 \pm 0.00$	$0.07 \pm 0.00$	H-VAEM	$0.23 \pm 0.16$	$0.32 \pm 0.00$
HMC-VAEM	$0.25 \pm 0.21$	$0.25 \pm 0.08$	HMC-VAEM	$0.08 \pm 0.00$	$0.05 \pm 0.00$	HMC-VAEM	$0.22 \pm 0.15$	<b><math>0.30 \pm 0.00</math></b>
<b>HH-VAEM</b>	<b><math>0.20 \pm 0.22</math></b>	<b><math>0.22 \pm 0.07</math></b>	<b>HH-VAEM</b>	<b><math>0.07 \pm 0.00</math></b>	<b><math>0.04 \pm 0.00</math></b>	<b>HH-VAEM</b>	<b><math>0.21 \pm 0.15</math></b>	<b><math>0.30 \pm 0.00</math></b>

Table 8: Average test Gaussian NLL of the observed features.

Table 9: Average test Bernoulli NLL of the observed features.

Table 10: Average test Cat. NLL of the observed features.

#### C.4 SAIA log-likelihoods

In order to extend the results provided in Section 5.3, we include here the log-likelihoods curves when dynamically selecting features using the same procedure (Figure 11).

#### C.5 Training times

Table 11 shows the average training time in minutes for each model in the experiments for Tables 1 and 2. The ratio between training times for our method and the Gaussian baselines is approximately between 5 and 10.

	Bank	Avocado	Yatch	Diabetes	Concrete	Wine	Energy	Boston
VAEM	$29.92 \pm 0.39$	$21.49 \pm 1.64$	$5.89 \pm 0.01$	$8.79 \pm 0.30$	$7.70 \pm 0.58$	$7.92 \pm 0.11$	$8.18 \pm 0.09$	$10.01 \pm 0.34$
MIWAEM	$63.33 \pm 6.20$	$37.17 \pm 2.28$	$8.21 \pm 0.24$	$13.29 \pm 0.33$	$11.51 \pm 0.52$	$11.81 \pm 0.13$	$16.71 \pm 0.16$	$15.01 \pm 0.13$
H-VAEM	$41.44 \pm 0.38$	$33.83 \pm 0.81$	$15.84 \pm 0.11$	$13.38 \pm 0.47$	$11.80 \pm 0.38$	$10.61 \pm 0.54$	$12.71 \pm 1.20$	$13.74 \pm 1.43$
HMC-VAEM	$281.88 \pm 9.14$	$356.22 \pm 5.94$	$23.50 \pm 1.60$	$50.42 \pm 1.18$	$42.70 \pm 3.14$	$63.46 \pm 1.97$	$90.87 \pm 7.05$	$103.72 \pm 7.24$
<b>HH-VAEM</b>	<b><math>316.81 \pm 9.49</math></b>	<b><math>388.28 \pm 6.47</math></b>	<b><math>27.08 \pm 0.12</math></b>	<b><math>68.29 \pm 4.06</math></b>	<b><math>65.78 \pm 0.43</math></b>	<b><math>79.97 \pm 5.53</math></b>	<b><math>140.33 \pm 7.36</math></b>	<b><math>129.30 \pm 4.69</math></b>

Table 11: Training times (in minutes) of our model and baselines.

#### C.6 SAIA times

The times for obtaining the SAIA metric curves in Section 5.3 are included in Figure 12. Although the performance is improved with HH-VAEM, it requires considerably longer time than the baselines to evaluate the reward, due to the HMC algorithm for sampling from the better approximated posterior. Future work might be oriented in proposing ways to measure and reduce this gap.

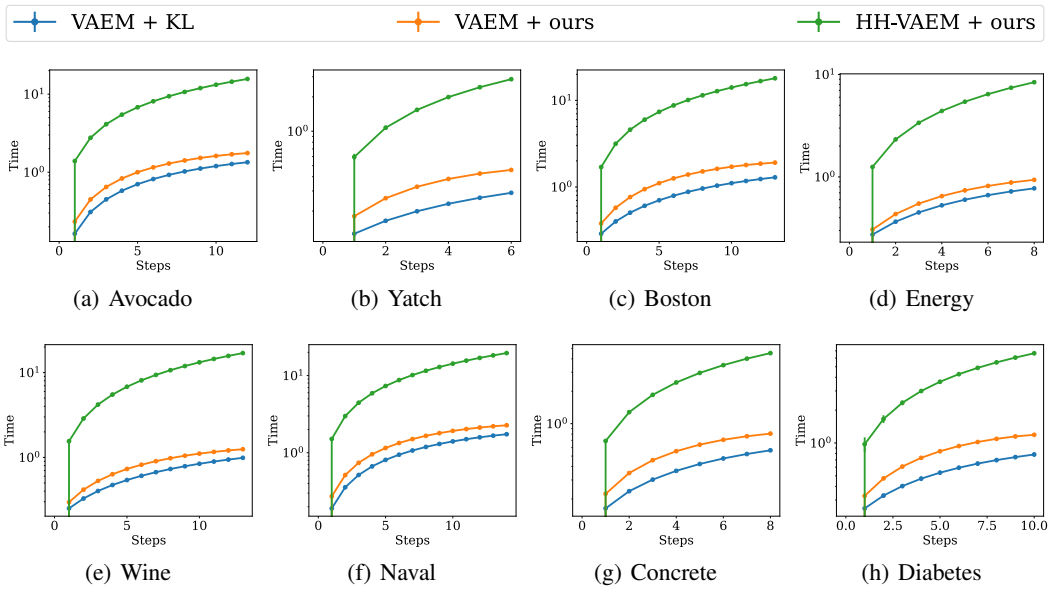


Figure 12: SAIA time curves. Horizontal axis shows acquisition steps (number of discovered features). Vertical axis is the elapsed time.