



MOX-Report No. 56/2025

**Improvements on uncertainty quantification with variational
autoencoders**

Tonini, A.; Bui-Thanh, T.; Regazzoni, F.; Dede', L; Quarteroni, A.

MOX, Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

mox-dmat@polimi.it

<https://mox.polimi.it>

IMPROVEMENTS ON UNCERTAINTY QUANTIFICATION WITH VARIATIONAL AUTOENCODERS

ANDREA TONINI^{1,*}, TAN BUI-THANH^{2,3}, FRANCESCO REGAZZONI¹, LUCA DEDE¹ AND ALFIO
QUARTERONI^{4,5}

ABSTRACT. Inverse problems aim to determine model parameters of a mathematical problem from given observational data. Neural networks can provide an efficient tool to solve these problems. In the context of Bayesian inverse problems, Uncertainty Quantification Variational AutoEncoders (UQ-VAE), a class of neural networks, approximate the posterior distribution mean and covariance of model parameters. This allows for both the estimation of the parameters and their uncertainty in relation to the observational data. In this work, we propose a novel loss function for training UQ-VAEs, which includes, among other modifications, the removal of a sample mean term from an already existing one. This modification improves the accuracy of UQ-VAEs, as the original theoretical result relies on the convergence of the sample mean to the expected value (a condition that, in high dimensional parameter spaces, requires a prohibitively large number of samples due to the curse of dimensionality). Avoiding the computation of the sample mean significantly reduces the training time in high dimensional parameter spaces compared to previous literature results. Under this new formulation, we establish a new theoretical result for the approximation of the posterior mean and covariance for general mathematical problems. We validate the effectiveness of UQ-VAEs through three benchmark numerical tests: a Poisson inverse problem, a non affine inverse problem and a 0D cardiocirculatory model, under the two clinical scenarios of systemic hypertension and ventricular septal defect. For the latter case, we perform forward uncertainty quantification.

Keywords: Bayesian inverse problem, deep learning, partial differential equations, uncertainty quantification, variational autoencoders.

AMS Subject Classification: 68T07, 62C10

1. INTRODUCTION

Forward problems describe mathematical models that map a set of parameters to a corresponding set of observational data. On the other hand, inverse problems aim to determine the model parameters from (noisy) observational data by minimizing a loss function that quantifies the mismatch between the observational data and the model outputs. Deterministic inverse problems seek to identify the single values of the parameters.²⁴ However, even when the forward map from parameters to outputs exists, the inverse map may not, leading to the ill-posedness of deterministic inverse problems. To address this issue,

¹MOX, Dipartimento di Matematica, Politecnico di Milano, Milan, Italy.

²Department of Aerospace Engineering and Engineering Mechanics, the University of Texas at Austin, Austin, Texas, USA..

³The Oden Institute for Computational Engineering and Sciences, the University of Texas at Austin, Austin, Texas, USA..

⁴(Professor Emeritus) MOX, Dipartimento di Matematica, Politecnico di Milano, Piazza Leonardo da Vinci, Milan, 20133, Italy.

⁵(Professor Emeritus) Institute of Mathematics, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland.

*Corresponding author. E-mail: andrea.tonini@polimi.it.

Tikhonov regularization is often applied to the loss function by adding a term that penalizes the norm of the parameters.³⁹

Bayesian inverse problems aim to determine the posterior distribution of the parameters, that is their probability distribution given an observational data.³⁷ The solution to a Bayesian inverse problem accounts for uncertainty in the parameters estimate that depends on the noise in the observational data.

Neural Networks (NNs) have gained increasing attention for solving inverse problems in recent years.^{5,31,32,50} They offer a significant reduction in the computational cost of solving inverse problems that would otherwise require numerous iterations of an optimization method.^{4,6,23} A single iteration of the chosen optimization method requires solving the forward problem, which can make inverse problems prohibitively expensive when the forward one is computationally demanding.

Tikhonov Networks (TNets) have been shown to effectively solve deterministic inverse problems.²⁵ Their loss function consists of two terms: one incorporating the prior knowledge of the parameters and another enforcing a physics-aware constraint. Data randomization has been theoretically and numerically demonstrated to implicitly regularize the weights and biases of the NN. Tikhonov autoencoders (TAEN) further enhance the use of data randomization, enabling the training of the NN generating new data by the randomization of a single sample.⁴⁵ InVAert networks address the ill-posedness nature of deterministic inverse problems by extending the observational data space with latent variables, enabling the identification of different parameters generating the same observational data.^{40,41} NETworks Tikhonov (NETT) have been applied to inverse problems in imaging, leveraging a data-driven regularization approach.²¹ Additionally, convolutional Neural Networks (CNN) have been widely explored for image-based inverse problems.^{8,17,22}

In the context of Bayesian inverse problems, Uncertainty Quantification Variational AutoEncoders (UQ-VAE) have been developed to estimate the mean and covariance of the posterior distribution of parameters. In Ref. 14, the authors established the theoretical foundations of UQ-VAEs, deriving a loss function based on a family of Jensen-Shannon divergences.²⁶ This approach comes from generative modeling that employs Jensen-Shannon divergences to compute the posterior parameter distribution and generate new data samples.^{9,38} In Ref. 42, a novel loss function was proposed to improve the estimation accuracy of UQ-VAEs, while enhancing their generalization capabilities. However, this approach required an increased computational cost for evaluating the training loss function, leading to prolonged training times.

In this work, we build upon the results of Ref. 42 proposing a novel loss function by removing a sample mean term from the original loss function, among other modifications. This approach improves accuracy and reduces the training time compared to previous literature results. The time reduction advantage becomes more pronounced as the dimensionality of the parameter space increases. We obtain a theoretical result analogous to the original one for affine forward problems and extend it to general forward problems under additional assumptions, thereby strengthening the theoretical background of UQ-VAEs. We validate our approach on a Poisson problem and compare it in terms of training time and accuracy with the previous one on a non linear problem. Furthermore, we test our approach on two test cases involving 0D cardiocirculatory models,^{1,10,12,28,43} specifically modeling systemic hypertension²⁷ and ventricular septal defect.^{7,36} For the latter case, we perform forward uncertainty quantification to further assess the performances of the new approach.

The paper is structured as follows: in section 2, we introduce the new loss function and UQ-VAEs; in section 3, we present the numerical tests on the Poisson problem and the 0D cardiocirculatory model; in section 4, we draw our conclusions.

2. BAYESIAN INVERSE PROBLEMS

We present the general framework of Bayesian inverse problems and UQ-VAEs. Additionally, we propose a novel loss function for training the UQ-VAEs addressing previous limitations.⁴² The new loss function reduces the computational cost of training the UQ-VAE while extending the theoretical result of the previous approach.

Bayesian inverse problems are used to estimate the parameters of a mathematical model along with their uncertainties, given noisy observational data. Let U and Y be random variables representing the parameters and the noisy observational data, respectively. We assume the existence of a map \mathcal{F} from the parameters to the observational data in absence of noise. Let E denote the noise random variable. The mathematical model under consideration is:

$$(1) \quad Y = \mathcal{F}(U) + E.$$

We focus on the case where U and Y are finite dimensional real random variables. Specifically, each parameter is a vector $\mathbf{u} \in \mathbb{R}^D$ and each observational data is a vector $\mathbf{y} \in \mathbb{R}^O$, with $D, O \in \mathbb{N}$. Thus, the map is defined as $\mathcal{F} : \mathbb{R}^D \rightarrow \mathbb{R}^O$. The solution of a Bayesian inverse problem is the posterior distribution of the parameters $p_{U|Y}(\mathbf{u}|\mathbf{y})$, which is the distribution of the parameters conditioned to the given noisy observational data.

Using the Bayes' theorem,² it is possible to retrieve an implicit formula for the posterior distribution that, in most of the cases, is not possible to compute explicitly as it requires the evaluation of an integral. Generally, one looks for the mean and covariance, namely the first two moments, of the posterior distribution to characterize the expected parameter and its uncertainty subjected to the noisy observational data. A standard approach to solve Bayesian inverse problems is by computing the Maximum A Posteriori estimate (MAP) parameter $\mathbf{u}_{\text{MAP}} \in \mathbb{R}^D$, that is the most likely parameter to observe given a noisy observational data $\mathbf{y} \in \mathbb{R}^O$, and the Laplace approximation of the covariance matrix $\Gamma_{\text{Lap}} \in \mathbb{R}^{D \times D}$.¹¹ By the Bayes' theorem, the posterior distribution can be expressed as:

$$(2) \quad p_{U|Y}(\mathbf{u}|\mathbf{y}) = \frac{p_{U,Y}(\mathbf{u}, \mathbf{y})}{p_Y(\mathbf{y})} = \frac{p_{Y|U}(\mathbf{y}|\mathbf{u}) p_U(\mathbf{u})}{p_Y(\mathbf{y})} = \frac{p_{Y|U}(\mathbf{y}|\mathbf{u}) p_U(\mathbf{u})}{\int_{\mathbb{R}^D} p_{Y|U}(\mathbf{y}|\mathbf{u}) p_U(\mathbf{u}) d\mathbf{u}},$$

where $p_{U,Y}(\mathbf{u}, \mathbf{y})$, $p_Y(\mathbf{y})$, $p_U(\mathbf{u})$ and $p_{Y|U}(\mathbf{y}|\mathbf{u})$ are the probability density functions (pdf) of the joint probability distribution of U and Y , the marginal pdfs of Y and U , and the conditioned probability of Y with respect to U (likelihood), respectively. We assume that E and U are independent, so the likelihood function is:

$$(3) \quad p_{Y|U}(\mathbf{y}|\mathbf{u}) = p_E(\mathbf{y} - \mathcal{F}(\mathbf{u})),$$

where p_E is the pdf of the noise random variable E . Furthermore, we assume that $U \sim \mathcal{N}(\boldsymbol{\mu}_{\text{pr}}, \Gamma_{\text{pr}})$ and $E \sim \mathcal{N}(\boldsymbol{\mu}_E, \Gamma_E)$, where $\boldsymbol{\mu}_{\text{pr}} \in \mathbb{R}^D$ and $\boldsymbol{\mu}_E \in \mathbb{R}^O$ are the means of the parameters and noise random variables, respectively, and $\Gamma_{\text{pr}} \in \mathbb{R}^{D \times D}$ and $\Gamma_E \in \mathbb{R}^{O \times O}$ are the corresponding covariance matrices. In this work, we restrict our analysis to symmetric positive definite (SPD) covariance matrices, as well as their inverses, to ensure the well definition of the loss function used in training the variational autoencoders. Using

(3) and the normality assumptions on U and E , the joint pdf $p_{U,Y}(\mathbf{u}, \mathbf{y})$ can be rewritten as:

$$p_{U,Y}(\mathbf{u}, \mathbf{y}) = p_E(\mathbf{y} - \mathcal{F}(\mathbf{u})) p_U(\mathbf{u}) = \frac{1}{(2\pi)^{\frac{O+D}{2}} |\Gamma_E|^{\frac{1}{2}} |\Gamma_{pr}|^{\frac{1}{2}}} \exp \left(-\frac{1}{2} \left(\|\mathbf{y} - \mathcal{F}(\mathbf{u}) - \boldsymbol{\mu}_E\|_{\Gamma_E^{-1}}^2 + \|\mathbf{u} - \boldsymbol{\mu}_{pr}\|_{\Gamma_{pr}^{-1}}^2 \right) \right),$$

where $|\cdot|$ denotes the matrix determinant and $\|\cdot\|_{\Gamma^{-1}}$ represents the norm induced by an SPD matrix Γ^{-1} . Maximizing the posterior pdf $p_{U|Y}(\mathbf{u}|\mathbf{y})$ with respect to \mathbf{u} yields the MAP estimate:

$$\begin{aligned} \mathbf{u}_{\text{MAP}} &:= \operatorname{argmax}_{\mathbf{u} \in \mathbb{R}^D} p_{U|Y}(\mathbf{u}|\mathbf{y}) \\ (4) \quad &= \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^D} \left(\|\mathbf{y} - \mathcal{F}(\mathbf{u}) - \boldsymbol{\mu}_E\|_{\Gamma_E^{-1}}^2 + \|\mathbf{u} - \boldsymbol{\mu}_{pr}\|_{\Gamma_{pr}^{-1}}^2 \right). \end{aligned}$$

The first term of the optimization problem derives from the likelihood function and quantifies the mismatch between the observational data and model predictions. The second term encodes prior knowledge about the parameters and serves as a regularization term. Solving this optimization problem typically requires gradient-based methods, which can be computationally expensive, particularly when evaluating the map \mathcal{F} is costly.

To quantify uncertainty in the parameter estimate due to noise, the Laplace approximation Γ_{Lap} of the covariance matrix can be used:¹¹

$$(5) \quad \Gamma_{\text{Lap}} = (J_{\mathcal{F}}(\mathbf{u}_{\text{MAP}})^T \Gamma_E^{-1} J_{\mathcal{F}}(\mathbf{u}_{\text{MAP}}) + \Gamma_{pr}^{-1})^{-1},$$

where $J_{\mathcal{F}}(\mathbf{u}_{\text{MAP}}) \in \mathbb{R}^{O \times D}$ is the Jacobian of \mathcal{F} evaluated at \mathbf{u}_{MAP} . Computing Γ_{Lap} requires evaluating $J_{\mathcal{F}}(\mathbf{u}_{\text{MAP}})$, which is significantly more expensive than a single evaluation of \mathcal{F} .

When evaluating \mathcal{F} is expensive (such as when \mathcal{F} represents the solution operator of a complex partial differential equation), computing \mathbf{u}_{MAP} and Γ_{Lap} becomes infeasible. UQ-VAEs,^{14,42} a class of NNs, offer a way to mitigate this computational burden. Although the most advanced UQ-VAE formulation⁴² presents strong generalization capabilities, even with small datasets, it has two main drawbacks when \mathcal{F} is not affine: an inaccurate estimate of the posterior covariance matrix and a high computational cost for training. Even if the latter is compensated by the efficient solution of Bayesian inverse problems, reducing the training cost is a priority. In this work, we propose a modification to the loss function used for training UQ-VAEs.⁴² The new loss function extends the previous theoretical result while significantly reducing computational costs. This reduction enables also faster testing of the NN architectures, indirectly improving the accuracy of the posterior mean and covariance matrix estimates.

Employing variational inference, we approximate the posterior distribution $p_{U|Y}(\mathbf{u}|\mathbf{y})$ using a Gaussian distribution:

$$(6) \quad q_{\phi}(\mathbf{u}|\mathbf{y}) = \mathcal{N}(\boldsymbol{\mu}_{\text{post}}(\mathbf{y}, \phi), \Gamma_{\text{post}}(\mathbf{y}, \phi)),$$

where $\boldsymbol{\mu}_{\text{post}}(\mathbf{y}, \phi) \in \mathbb{R}^D$, $\Gamma_{\text{post}}(\mathbf{y}, \phi) \in \mathbb{R}^{D \times D}$ and ϕ denotes a set of hyper-parameters (in section 2.2, ϕ is the set of weights and biases of the UQ-VAE). For conciseness, we omit the explicit dependence of $\boldsymbol{\mu}_{\text{post}}$ and Γ_{post} on \mathbf{y} and ϕ .

2.1. Loss function. We describe the loss function and the theoretical result presented in Ref. 42, modifying the former to reduce its computational cost.

Given $\alpha \in (0, 1)$ and an observational data $\mathbf{y} \in \mathbb{R}^O$, the loss function proposed in the previous work is:

$$(7) \quad \begin{aligned} L_\alpha(\boldsymbol{\mu}, C) = & (1 - \alpha) (\|\boldsymbol{\mu} - \boldsymbol{\mu}_{\text{pr}}\|_{\Gamma^{-1}}^2 + \text{tr}(\Gamma^{-1}\Gamma_{\text{pr}})) + \\ & \alpha (\|\boldsymbol{\mu} - \boldsymbol{\mu}_{\text{pr}}\|_{\Gamma_{\text{pr}}^{-1}}^2 + \text{tr}(\Gamma_{\text{pr}}^{-1}\Gamma)) + \\ & \alpha \mathbb{E}_{\mathcal{N}(\boldsymbol{\mu}, \Gamma)} [\|\mathbf{y} - \boldsymbol{\mu}_{\text{E}} - \mathcal{F}(\mathbf{u})\|_{\Gamma_{\text{E}}^{-1}}^2], \end{aligned}$$

where $\mathbb{E}_{\mathcal{N}(\boldsymbol{\mu}, \Gamma)} [\cdot]$ is the expected value with respect to the probability distribution $\mathcal{N}(\boldsymbol{\mu}, \Gamma)$, $\boldsymbol{\mu} \in \mathbb{R}^D$, $\Gamma \in \mathbb{R}^{D \times D}$ is SPD and C is its Cholesky factor (a lower triangular matrix with positive diagonal entries). This loss function is an upper bound for a family of Jensen-Shannon divergences (JSD)²⁶ summed to a Kullback-Leibler divergence, which measure the difference between the posterior distribution $p_{\mathbf{U}|\mathbf{Y}}(\mathbf{u}|\mathbf{y})$ and the distribution $\mathcal{N}(\boldsymbol{\mu}, \Gamma)$. The first two lines come from the Kullback-Leibler divergences between the prior and the distribution $\mathcal{N}(\boldsymbol{\mu}, \Gamma)$ and the third line represents the expected squared error on the observational data.

When \mathcal{F} is affine (i.e., $\mathcal{F}(\mathbf{u}) = F\mathbf{u} + \mathbf{f}$, where $F \in \mathbb{R}^{O \times D}$, $\mathbf{f} \in \mathbb{R}^O$), (4) and (5) can be expressed as:

$$(8) \quad \mathbf{u}_{\text{MAP}} = \Gamma_{\text{Lap}} (F^T \Gamma_{\text{E}}^{-1} (\mathbf{y} - \mathbf{f} - \boldsymbol{\mu}_{\text{E}}) + \Gamma_{\text{pr}}^{-1} \boldsymbol{\mu}_{\text{pr}}),$$

$$(9) \quad \Gamma_{\text{Lap}} = (F^T \Gamma_{\text{E}}^{-1} F + \Gamma_{\text{pr}}^{-1})^{-1}.$$

Moreover, the expected value in (7) is computed exactly:

$$(10) \quad \begin{aligned} & \mathbb{E}_{\mathcal{N}(\boldsymbol{\mu}, \Gamma)} [\|\mathbf{y} - \boldsymbol{\mu}_{\text{E}} - \mathcal{F}(\mathbf{u})\|_{\Gamma_{\text{E}}^{-1}}^2] = \\ & \|\mathbf{y} - \boldsymbol{\mu}_{\text{E}} - \mathbb{E}_{\mathcal{N}(\boldsymbol{\mu}, \Gamma)} [\mathcal{F}(\mathbf{u})]\|_{\Gamma_{\text{E}}^{-1}}^2 + \text{tr}(\Gamma_{\text{E}}^{-1} \text{Cov}_{\mathcal{N}(\boldsymbol{\mu}, \Gamma)} [\mathcal{F}(\mathbf{u})]) = \\ & \|\mathbf{y} - \boldsymbol{\mu}_{\text{E}} - F\boldsymbol{\mu} - \mathbf{f}\|_{\Gamma_{\text{E}}^{-1}}^2 + \text{tr}(\Gamma_{\text{E}}^{-1} F \Gamma F^T), \end{aligned}$$

where $\text{Cov}_{\mathcal{N}(\boldsymbol{\mu}, \Gamma)} [\cdot]$ is the covariance with respect to the probability distribution $\mathcal{N}(\boldsymbol{\mu}, \Gamma)$. Minimizing this loss function yields the following result.⁴²

Theorem 1. Assume that $\mathcal{F}(\mathbf{u}) = F\mathbf{u} + \mathbf{f}$, where $F \in \mathbb{R}^{O \times D}$, $\mathbf{f} \in \mathbb{R}^O$, and $\mathbf{y} \in \mathbb{R}^O$. Suppose that the prior and noise models are Gaussian and independent, $\mathcal{N}(\boldsymbol{\mu}_{\text{pr}}, \Gamma_{\text{pr}})$ and $\mathcal{N}(\boldsymbol{\mu}_{\text{E}}, \Gamma_{\text{E}})$, respectively. Then, the posterior distribution $p_{\mathbf{U}|\mathbf{Y}}(\mathbf{u}|\mathbf{y})$ is $\mathcal{N}(\mathbf{u}_{\text{MAP}}, \Gamma_{\text{Lap}})$, where $(\mathbf{u}_{\text{MAP}}, \Gamma_{\text{Lap}})$ are given by (8) and (9).

Let $\alpha \in (0, 1)$ and $(\hat{\boldsymbol{\mu}}, \hat{C})$ (with $\hat{\Gamma} = \hat{C}\hat{C}^T$) be the stationary points of the loss function

$$(11) \quad \begin{aligned} L_\alpha(\boldsymbol{\mu}, C) = & (1 - \alpha) (\|\boldsymbol{\mu} - \boldsymbol{\mu}_{\text{pr}}\|_{\Gamma^{-1}}^2 + \text{tr}(\Gamma^{-1}\Gamma_{\text{pr}})) + \\ & \alpha (\|\boldsymbol{\mu} - \boldsymbol{\mu}_{\text{pr}}\|_{\Gamma_{\text{pr}}^{-1}}^2 + \text{tr}(\Gamma_{\text{pr}}^{-1}\Gamma)) + \\ & \alpha (\|\mathbf{y} - \boldsymbol{\mu}_{\text{E}} - F\boldsymbol{\mu} - \mathbf{f}\|_{\Gamma_{\text{E}}^{-1}}^2 + \text{tr}(\Gamma_{\text{E}}^{-1} F \Gamma F^T)). \end{aligned}$$

Let

$$(12) \quad \boldsymbol{\mu}_{\text{post}} = \frac{1 - \alpha}{\alpha} \Gamma_{\text{Lap}} \hat{\Gamma}^{-1} (\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_{\text{pr}}) + \hat{\boldsymbol{\mu}},$$

$$(13) \quad \Gamma_{\text{post}} = \hat{\Gamma} A^{-1} \hat{\Gamma},$$

with

$$A = \frac{1 - \alpha}{\alpha} ((\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_{\text{pr}})(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_{\text{pr}})^{\text{T}} + \Gamma_{\text{pr}}).$$

Then $(\boldsymbol{\mu}_{\text{post}}, \Gamma_{\text{post}}) = (\mathbf{u}_{\text{MAP}}, \Gamma_{\text{Lap}})$.

Remark. In practice, we minimize the loss function (11) to find $(\hat{\boldsymbol{\mu}}, \hat{\Gamma})$ and we compute $(\boldsymbol{\mu}_{\text{post}}, \Gamma_{\text{post}})$ using (12) and (13).

Remark. The matrix A is SPD since it is a positive scalar multiple of the sum of two symmetric matrices that are positive semidefinite and positive definite. Consequently, $\hat{\Gamma} A^{-1} \hat{\Gamma}$ is also SPD, as Γ_{post} and Γ_{Lap} are.

When the map \mathcal{F} is not affine, $\mathbb{E}_{\mathcal{N}(\boldsymbol{\mu}, \Gamma)} [\|\mathbf{y} - \boldsymbol{\mu}_{\text{E}} - \mathcal{F}(\mathbf{u})\|_{\Gamma_{\text{E}}^{-1}}^2]$ cannot be computed in a closed form. To approximate this expectation, the central limit theorem is used, leading to the following loss function:

$$(14) \quad \begin{aligned} & (1 - \alpha) (\|\boldsymbol{\mu} - \boldsymbol{\mu}_{\text{pr}}\|_{\Gamma^{-1}}^2 + \text{tr}(\Gamma^{-1} \Gamma_{\text{pr}})) + \\ & \alpha (\|\boldsymbol{\mu} - \boldsymbol{\mu}_{\text{pr}}\|_{\Gamma_{\text{pr}}^{-1}}^2 + \text{tr}(\Gamma_{\text{pr}}^{-1} \Gamma)) + \\ & \alpha \frac{1}{K} \sum_{k=1}^K \|\mathbf{y} - \boldsymbol{\mu}_{\text{E}} - \mathcal{F}(\mathbf{u}_{\text{draw}}^k)\|_{\Gamma_{\text{E}}^{-1}}^2, \end{aligned}$$

where $K \in \mathbb{N}$ and, for $k = 1, \dots, K$, $\mathbf{u}_{\text{draw}}^k = \boldsymbol{\mu} + C \boldsymbol{\varepsilon}^k$, with $\boldsymbol{\varepsilon}^k \in \mathbb{R}^D$ and $\boldsymbol{\varepsilon}^k \sim \mathcal{N}(\mathbf{0}, I)$. Therefore, at each iteration of the optimization algorithm, $\mathbf{u}_{\text{draw}}^k$ are drawn from $\mathcal{N}(\boldsymbol{\mu}, \Gamma)$.

Approximating the expectation in (7) is typically expensive, even when Quasi-Monte Carlo methods are used,¹⁵ as a large number of samples K is required for convergence. Given the Cholesky factorizations of Γ_{pr} and Γ_{E} the computational cost of evaluating (14) is $O(D^3 + KD^2 + Kc(\mathcal{F}) + KO^2)$, where $c(\mathcal{F})$ is the cost of an evaluation of \mathcal{F} . $O(D^3)$ is the cost of solving D linear systems with matrix Γ_{post} or Γ_{pr} . $O(KD^2)$, $O(Kc(\mathcal{F}))$ and $O(KO^2)$ are the costs of computing K samples $\mathbf{u}_{\text{draw}}^k$, of K evaluations of \mathcal{F} and of solving K linear systems with matrix Γ_{E} , respectively. The computational costs of the other terms are lower than the previous ones.

We propose a novel loss function that significantly reduces the previous computational cost:

$$(15) \quad \begin{aligned} L_{\theta}(\boldsymbol{\mu}, C) = & \theta^2 \text{tr}(\Gamma^{-1} \Gamma_{\text{pr}}) + \|\boldsymbol{\mu} - \boldsymbol{\mu}_{\text{pr}}\|_{\Gamma_{\text{pr}}^{-1}}^2 + \theta^2 \text{tr}(\Gamma_{\text{pr}}^{-1} \Gamma) + \\ & \|\mathbf{y} - \boldsymbol{\mu}_{\text{E}} - \mathcal{F}(\boldsymbol{\mu})\|_{\Gamma_{\text{E}}^{-1}}^2 + \\ & \text{tr} \left(\Gamma_{\text{E}}^{-1} \left(\tilde{\mathcal{F}}(\theta C + \boldsymbol{\mu} \mathbf{1}_{\text{D}}^{\text{T}}) - \mathcal{F}(\boldsymbol{\mu}) \mathbf{1}_{\text{D}}^{\text{T}} \right) \left(\tilde{\mathcal{F}}(\theta C + \boldsymbol{\mu} \mathbf{1}_{\text{D}}^{\text{T}}) - \mathcal{F}(\boldsymbol{\mu}) \mathbf{1}_{\text{D}}^{\text{T}} \right)^{\text{T}} \right), \end{aligned}$$

where $\tilde{\mathcal{F}}(A) = [\mathcal{F}(A_{:,1}) \dots \mathcal{F}(A_{:,D})] \in \mathbb{R}^{O \times D}$, for $A \in \mathbb{R}^{D \times D}$, is the map \mathcal{F} applied column wise to A . The vector $\mathbf{1}_{\text{D}} \in \mathbb{R}^D$ consists of ones and the optimal value for $\theta \in \mathbb{R} \setminus \{0\}$ is analyzed in section 3.1.

The new loss function fixes $\alpha = 1/2$ in the previous one (7) and removes the term $\|\boldsymbol{\mu} - \boldsymbol{\mu}_{\text{pr}}\|_{\Gamma^{-1}}^2$, since it vanishes for $\boldsymbol{\mu} = \boldsymbol{\mu}_{\text{pr}}$ as $\|\boldsymbol{\mu} - \boldsymbol{\mu}_{\text{pr}}\|_{\Gamma_{\text{pr}}^{-1}}^2$ does. The coefficient θ is a (small) perturbation needed to obtain the following theoretical result. The coefficient θ^2 balances the derivatives of L_{θ} with respect to C . The expected value of (7) is substituted with the second and third lines of (15) which are more efficient to compute and mimic the expected value in the affine case (10).

With this novel loss function, we are able to establish the following new theorem.

Theorem 2. Assume that \mathcal{F} is a $C^1(\mathbb{R}^D, \mathbb{R}^O)$ function, $\mathbf{y} \in \mathbb{R}^O$ and that the prior and noise models are Gaussian and independent, $\mathcal{N}(\boldsymbol{\mu}_{\text{pr}}, \Gamma_{\text{pr}})$ and $\mathcal{N}(\boldsymbol{\mu}_{\text{E}}, \Gamma_{\text{E}})$, respectively. Let $(\hat{\boldsymbol{\mu}}(\theta), \hat{C}(\theta))$ be the stationary points of $L_\theta(\boldsymbol{\mu}, C)$. Suppose that $\hat{\boldsymbol{\mu}} = \lim_{\theta \rightarrow 0} \hat{\boldsymbol{\mu}}(\theta)$ and $\hat{C} = \lim_{\theta \rightarrow 0} \hat{C}(\theta)$ exist and are bounded in norm, \hat{C} is invertible and the minimization problem (4) has only one stationary point \mathbf{u}_{MAP} . Let $\hat{\Gamma} = \hat{C}\hat{C}^T$ and define:

$$\begin{aligned}\boldsymbol{\mu}_{\text{post}} &= \hat{\boldsymbol{\mu}}, \\ \Gamma_{\text{post}} &= \hat{\Gamma}\Gamma_{\text{pr}}^{-1}\hat{\Gamma}.\end{aligned}$$

Then $(\boldsymbol{\mu}_{\text{post}}, \Gamma_{\text{post}}) = (\mathbf{u}_{\text{MAP}}, \Gamma_{\text{Lap}})$, where $(\mathbf{u}_{\text{MAP}}, \Gamma_{\text{Lap}})$ are given by (4) and (5).

Proof. We compute the gradients of $L_\theta(\boldsymbol{\mu}, C)$:

$$\begin{aligned}\frac{\partial L_\theta}{\partial \boldsymbol{\mu}} &= 2\Gamma_{\text{pr}}^{-1}(\boldsymbol{\mu} - \boldsymbol{\mu}_{\text{pr}}) - 2J_{\mathcal{F}}^T(\boldsymbol{\mu})\Gamma_{\text{E}}^{-1}(\mathbf{y} - \boldsymbol{\mu}_{\text{E}} - \mathcal{F}(\boldsymbol{\mu})) + \\ &\quad 2 \sum_{k=1}^D (J_{\mathcal{F}}(\theta C_{:,k} + \boldsymbol{\mu}) - J_{\mathcal{F}}(\boldsymbol{\mu}))^T \Gamma_{\text{E}}^{-1}(\mathcal{F}(\theta C_{:,k} + \boldsymbol{\mu}) - \mathcal{F}(\boldsymbol{\mu})), \\ \frac{\partial L_\theta}{\partial C} &= -2\theta^2 \Gamma_{\text{pr}}^{-1} \Gamma_{\text{pr}} C^{-T} + 2\theta^2 \Gamma_{\text{pr}}^{-1} C + \\ &\quad 2\theta [J_{\mathcal{F}}^T(\theta C_{:,1} + \boldsymbol{\mu})\Gamma_{\text{E}}^{-1}(\mathcal{F}(\theta C_{:,1} + \boldsymbol{\mu}) - \mathcal{F}(\boldsymbol{\mu})) \quad \dots].\end{aligned}$$

Setting the gradients to vanish and omitting the dependence of $(\hat{\boldsymbol{\mu}}(\theta), \hat{C}(\theta))$ on θ , we get:

$$\begin{aligned}(16) \quad &\Gamma_{\text{pr}}^{-1}(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_{\text{pr}}) - J_{\mathcal{F}}^T(\hat{\boldsymbol{\mu}})\Gamma_{\text{E}}^{-1}(\mathbf{y} - \boldsymbol{\mu}_{\text{E}} - \mathcal{F}(\hat{\boldsymbol{\mu}})) + \\ &\sum_{k=1}^D \left(J_{\mathcal{F}}(\theta \hat{C}_{:,k} + \hat{\boldsymbol{\mu}}) - J_{\mathcal{F}}(\hat{\boldsymbol{\mu}}) \right)^T \Gamma_{\text{E}}^{-1}(\mathcal{F}(\theta \hat{C}_{:,k} + \hat{\boldsymbol{\mu}}) - \mathcal{F}(\hat{\boldsymbol{\mu}})) = \mathbf{0}, \\ &- \hat{\Gamma}^{-1} \Gamma_{\text{pr}} \hat{\Gamma}^{-1} + \Gamma_{\text{pr}}^{-1} + \\ (17) \quad &\frac{1}{\theta} \left[J_{\mathcal{F}}^T(\theta \hat{C}_{:,1} + \hat{\boldsymbol{\mu}})\Gamma_{\text{E}}^{-1}(\mathcal{F}(\theta \hat{C}_{:,1} + \hat{\boldsymbol{\mu}}) - \mathcal{F}(\hat{\boldsymbol{\mu}})) \quad \dots \right] \hat{C}^{-1} = \mathbf{0}.\end{aligned}$$

As $\mathcal{F} \in C^1(\mathbb{R}^D, \mathbb{R}^O)$ and \hat{C} is bounded in norm, we can expand the terms $\mathcal{F}(\theta \hat{C}_{:,k} + \hat{\boldsymbol{\mu}})$ and $J_{\mathcal{F}}(\theta \hat{C}_{:,k} + \hat{\boldsymbol{\mu}})$ for $\theta \rightarrow 0$ as:

$$\begin{aligned}\mathcal{F}(\theta \hat{C}_{:,k} + \hat{\boldsymbol{\mu}}) &= \mathcal{F}(\hat{\boldsymbol{\mu}}) + \theta J_{\mathcal{F}}(\hat{\boldsymbol{\mu}})\hat{C}_{:,k} + o(\theta \|\hat{C}_{:,k}\|_2) = \mathcal{F}(\hat{\boldsymbol{\mu}}) + \theta J_{\mathcal{F}}(\hat{\boldsymbol{\mu}})\hat{C}_{:,k} + o(\theta), \\ J_{\mathcal{F}}(\theta \hat{C}_{:,k} + \hat{\boldsymbol{\mu}}) &= J_{\mathcal{F}}(\hat{\boldsymbol{\mu}}) + o(\|\hat{C}_{:,k}\|_2) = J_{\mathcal{F}}(\hat{\boldsymbol{\mu}}) + o(1),\end{aligned}$$

where $o(\cdot)$ indicates the small-o notation. The boundedness of \hat{C} is needed to remove $\|\hat{C}_{:,k}\|_2$ from $o(\cdot)$. By computing the limit for $\theta \rightarrow 0$ of equation (16), using $\mathcal{F} \in C^1(\mathbb{R}^D, \mathbb{R}^O)$ and that $\hat{\boldsymbol{\mu}}, \hat{C}$ exist and are bounded, we obtain:

$$\begin{aligned}\Gamma_{\text{pr}}^{-1}(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_{\text{pr}}) - J_{\mathcal{F}}^T(\hat{\boldsymbol{\mu}})\Gamma_{\text{E}}^{-1}(\mathbf{y} - \boldsymbol{\mu}_{\text{E}} - \mathcal{F}(\hat{\boldsymbol{\mu}})) + o(\theta) &\rightarrow \\ \Gamma_{\text{pr}}^{-1}(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_{\text{pr}}) - J_{\mathcal{F}}^T(\hat{\boldsymbol{\mu}})\Gamma_{\text{E}}^{-1}(\mathbf{y} - \boldsymbol{\mu}_{\text{E}} - \mathcal{F}(\hat{\boldsymbol{\mu}})) &= \mathbf{0},\end{aligned}$$

By the above equation, we observe that $\boldsymbol{\mu}_{\text{post}} = \hat{\boldsymbol{\mu}}$ satisfies the same equation of the stationary point \mathbf{u}_{MAP} of equation (4), therefore, by the uniqueness of the stationary point of (4), $\boldsymbol{\mu}_{\text{post}}$ is equal to \mathbf{u}_{MAP} .

Since $\hat{\underline{C}}$ is invertible, $\hat{\underline{\Gamma}}^{-1}$ exists. We compute the limit of equation (17) for $\theta \rightarrow 0$:

$$\begin{aligned} & -\hat{\Gamma}^{-1}\Gamma_{\text{pr}}\hat{\Gamma}^{-1} + \Gamma_{\text{pr}}^{-1} + \frac{1}{\theta} (J_{\mathcal{F}}^T(\hat{\underline{\mu}}) + o(1)) \Gamma_{\text{E}}^{-1} (\theta J_{\mathcal{F}}(\hat{\underline{\mu}})\hat{\underline{C}} + o(\theta)) \hat{\underline{C}}^{-1} \rightarrow \\ (18) \quad & -\hat{\underline{\Gamma}}^{-1}\Gamma_{\text{pr}}\hat{\underline{\Gamma}}^{-1} + \Gamma_{\text{pr}}^{-1} + J_{\mathcal{F}}^T(\hat{\underline{\mu}}) \Gamma_{\text{E}}^{-1} J_{\mathcal{F}}(\hat{\underline{\mu}}) = 0. \end{aligned}$$

Since $\hat{\underline{\mu}}$ is equal to \mathbf{u}_{MAP} , we get:

$$\Gamma_{\text{pr}}^{-1} + J_{\mathcal{F}}^T(\hat{\underline{\mu}}) \Gamma_{\text{E}}^{-1} J_{\mathcal{F}}(\hat{\underline{\mu}}) = \Gamma_{\text{Lap}}^{-1}.$$

Therefore, from equation (18), Γ_{post} is equal to Γ_{Lap} . \square

Remark. The above theorem only holds for $\theta \rightarrow 0$. In practice, θ can be chosen arbitrarily small without incurring additional computational cost, as long as machine precision does not hinder the evaluation of the loss.

Remark. The $\lim_{\theta \rightarrow 0} L_{\theta}(\underline{\mu}, C)$ corresponds to the loss function of minimization problem (4), which makes the result $\underline{\mu}_{\text{post}} = \mathbf{u}_{\text{MAP}}$ intuitive. However, it must be shown that the minimization of $L_{\theta}(\underline{\mu}, C)$ with respect to $\underline{\mu}$ and the limit as $\theta \rightarrow 0$ commute, as we have done. Moreover, although taking $\theta \rightarrow 0$ recovers the minimization problem for \mathbf{u}_{MAP} , the θ dependent terms still enable the estimation of the posterior covariance matrix.

Remark. When $\mathcal{F}(\mathbf{u}) = F\mathbf{u} + \mathbf{f}$ is affine, the result of Theorem 2 holds $\forall \theta \in \mathbb{R} \setminus \{0\}$ without any hypotheses on $\hat{\underline{\mu}}, \hat{\underline{C}}$ and the minimization problem (4). It is sufficient to substitute in (16) and (17) the expression of \mathcal{F} and $J_{\mathcal{F}}(\mathbf{u}) = F$.

Remark. The above theorem does not assume a specific posterior distribution therefore the mean and covariance estimates hold in general. We choose to approximate the posterior distribution using a multivariate Gaussian distribution, as shown in (6).

Remark. The assumption of a unique stationary point for (4) is restrictive and generally holds only when the prior dominates the loss function.

The computational cost of evaluating (15) is $O(D^3 + Dc(\mathcal{F}) + DO^2)$, which follows from the same cost analysis as for (14). Thanks to the curse of dimensionality, the number of samples K to accurately estimate an expected value is generally much higher than the sample space dimension D . Consequently, evaluating (15) is significantly less expensive than evaluating (14).

2.2. Uncertainty quantification variational autoencoders. An autoencoder consists in the composition of two maps called encoder and decoder. The encoder maps the inputs of the autoencoder to a latent space, while the decoder maps the outputs of the encoder to the original space. Autoencoders typically approximate the identity function, meaning that their inputs and outputs should match. In our method, the encoder is a NN $\varphi : \mathbb{R}^O \rightarrow \mathbb{R}^{D + \frac{D(D+1)}{2}}$, while the decoder corresponds to the map \mathcal{F} or an approximation of it. The latent space of the UQ-VAE represents the mean $\underline{\mu}$ and the Cholesky factor C of the multivariate normal distribution $\mathcal{N}(\underline{\mu}, \Gamma)$ (Figure 1).

The following Corollary, proven in Ref. 42 and easy to extend to the new loss function (15), demonstrates the convergence of the outputs of a single-layer linear encoder to the posterior mean and covariance in the case of \mathcal{F} affine.

Corollary 3. Consider a single-layer linear encoder. Let $L_{\theta}(\phi) = L_{\theta}(\underline{\mu}(\phi), \Gamma(\phi))$, where ϕ is the set of weights and biases of the NN. Under the assumptions of Theorem 2 and \mathcal{F}




FIGURE 1. Structure of the UQ-VAE when \mathcal{F} is evaluated.

affine, we define:

$$(19) \quad \mu = W_\mu y + \mathbf{b}_\mu,$$

$$(20) \quad C = \text{vec}_L^{-1}(\mathbf{l}) + \text{diag}(\sigma),$$

$$(21) \quad \sigma = \exp(W_\sigma y + \mathbf{b}_\sigma),$$

$$(22) \quad \mathbf{l} = W_1 y + \mathbf{b}_1,$$

where $W_\mu \in \mathbb{R}^{D \times O}$ and $W_\sigma \in \mathbb{R}^{D \times O}$, $W_1 \in \mathbb{R}^{\frac{D(D-1)}{2} \times O}$ are the encoder weight matrices, $\mathbf{b}_\mu \in \mathbb{R}^D$, $\mathbf{b}_\sigma \in \mathbb{R}^D$ and $\mathbf{b}_1 \in \mathbb{R}^{\frac{D(D-1)}{2}}$ are the encoder biases. The functions \exp , vec_L and diag denote the elementwise exponential function, the vectorization of a strictly lower triangular matrix and the construction of a diagonal matrix from its vector diagonal, respectively.

The stationary points $\hat{\phi}$ of $L_\theta(\phi)$ are such that $(\mu_{\text{post}}(\hat{\phi}), \Gamma_{\text{post}}(\hat{\phi})) = (\mathbf{u}_{\text{MAP}}, \Gamma_{\text{Lap}})$.

For deeper encoders with non linear activation functions, (μ, C) is computed as in (19), (20), (21) and (22), using the encoder's outputs in place of $W_\mu y + \mathbf{b}_\mu$, $W_\sigma y + \mathbf{b}_\sigma$ and $W_1 y + \mathbf{b}_1$.

In many applications, evaluating \mathcal{F} is computationally expensive or \mathcal{F} can be unknown. In such cases, we approximate the map \mathcal{F} using a second NN $\psi : \mathbb{R}^D \rightarrow \mathbb{R}^O$ (Figure 2). To account for the approximation error introduced by the NN ψ , we use an error random variable $E_{\text{dec}} \sim \mathcal{N}(\mu_{\text{dec}}, \Gamma_{\text{dec}})$ in the model (1):

$$\mathcal{F}(U) = \psi(U) + E_{\text{dec}} \Rightarrow Y = \psi(U) + E_{\text{dec}} + E.$$

We assume that E_{dec} and E are independent random variables, so that their sum $E_{\text{dec}} + E$ is distributed as $\mathcal{N}(\mu_{\text{dec}} + \mu_E, \Gamma_{\text{dec}} + \Gamma_E)$. Defining $\mu_{\tilde{E}} = \mu_{\text{dec}} + \mu_E$ and $\Gamma_{\tilde{E}} = \Gamma_{\text{dec}} + \Gamma_E$, the corresponding loss function becomes:

$$(23) \quad \theta^2 \text{tr}(\Gamma^{-1} \Gamma_{\text{pr}}) + \|\mu - \mu_{\text{pr}}\|_{\Gamma_{\text{pr}}^{-1}}^2 + \theta^2 \text{tr}(\Gamma_{\text{pr}}^{-1} \Gamma) + \|\mathbf{y} - \mu_{\tilde{E}} - \psi(\mu)\|_{\Gamma_{\tilde{E}}^{-1}}^2 + \text{tr}\left(\Gamma_{\tilde{E}}^{-1} \left(\tilde{\psi}(\theta C + \mu \mathbf{1}_D^T) - \psi(\mu) \mathbf{1}_D^T\right) \left(\tilde{\psi}(\theta C + \mu \mathbf{1}_D^T) - \psi(\mu) \mathbf{1}_D^T\right)^T\right),$$

where $\tilde{\psi}(A) = [\psi(A_{:,1}) \dots \psi(A_{:,D})] \in \mathbb{R}^{O \times D}$, for $A \in \mathbb{R}^{D \times D}$.

When the UQ-VAE consists of two NNs, its training is performed in two stages: at first the decoder NN is trained and after the encoder one.⁴⁵ Simultaneous training of both NNs




FIGURE 2. Structure of the UQ-VAE using a NN to approximate \mathcal{F} .

may lead to unpredictable results, as our analysis of the minima of the loss function (15) assumes a fixed decoder during optimization.

Remark. In practice, we always use fully connected feed forward NNs as encoders and decoders. However, the proposed approach is compatible with other NN architectures.

Remark. We train the decoder NN using the squared error between its outputs and the noiseless observational data.

Remark. After training the decoder NN, we estimate μ_{dec} and Γ_{dec} by computing the sample mean and covariance of $\mathcal{F}(U) - \psi(U)$ on a test set of approximately 29,000 samples.

Remark. When a dataset $\{\mathbf{y}^{(m)}\}_{m=1}^M$, with $M \in \mathbb{N}$, is available, we train the encoder by minimizing the average of the loss functions (15) or (23) over the dataset.

2.2.1. Weights and biases initialization and dataset normalization. The initialization of the weights and biases of the encoder plays a crucial role in its training. A good initial guess is obtained by setting $\mu \sim \mu_{\text{pr}}$ and $C \sim C_{\text{pr}}$ (with C_{pr} the Cholesky factor of Γ_{pr}), as these values minimize the term $\theta^2 \text{tr}(\Gamma^{-1} \Gamma_{\text{pr}}) + \|\mu - \mu_{\text{pr}}\|_{\Gamma_{\text{pr}}^{-1}}^2 + \theta^2 \text{tr}(\Gamma_{\text{pr}}^{-1} \Gamma)$ in the loss function (15). The encoder weights are initialized according to the Xavier uniform initializer.¹³ Then, all the weights in the last layer are scaled by 10^{-4} . This ensures that, up to a small discrepancy, the output of the encoder is given by the biases of the last layer. All encoder biases are set to zero, except for those of the last layer which are fixed to μ_{pr} and \mathbf{c}_{pr} , where \mathbf{c}_{pr} is given by applying the inverse of the right hand side transformation of (20) to $C_{\text{pr}} \in \mathbb{R}^{D \times D}$.

Since the loss function (15) depends on the mean and covariance of both parameters and noise, normalizing the dataset (to enhance the UQ-VAE generalization capability) also affects these quantities.

We consider a dataset $\{(\mathbf{u}^{(m)}, \mathbf{y}^{(m)})\}_{m=1}^M$, where $M \in \mathbb{N}$, $\mathbf{y}^{(m)} = \mathcal{F}(\mathbf{u}^{(m)}) + \varepsilon^{(m)}$, with $\mathbf{u}^{(m)}$ and $\varepsilon^{(m)}$ are sampled from $\mathcal{N}(\mu_{\text{pr}}, \Gamma_{\text{pr}})$ and $\mathcal{N}(\mu_{\text{E}}, \Gamma_{\text{E}})$, respectively. To normalize the data within a given range, e.g., $[0, 1]$, we introduce vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^D$ and $\mathbf{c}, \mathbf{d} \in \mathbb{R}^O$ such that:

$$\begin{aligned} \bar{\mathbf{u}}^{(m)} &= \mathbf{u}^{(m)} \odot \mathbf{a} + \mathbf{b} \in [0, 1]^D, \\ \bar{\mathbf{y}}^{(m)} &= \mathbf{y}^{(m)} \odot \mathbf{c} + \mathbf{d} \in [0, 1]^O. \end{aligned}$$

As shown in Section 2.4.1 of Ref. 42, the normalized random variables \bar{Y} , \bar{U} and \bar{E} satisfy a model analogous to (1) in the normalized scale, with $\bar{U} \sim N(\bar{\mu}_{\text{pr}}, \bar{\Gamma}_{\text{pr}})$ and

$\bar{E} \sim N(\bar{\boldsymbol{\mu}}_E, \bar{\Gamma}_E)$, where:

$$\begin{aligned}\bar{\boldsymbol{\mu}}_{\text{pr}} &= \boldsymbol{\mu}_{\text{pr}} \odot \mathbf{a} + \mathbf{b}, \\ \bar{\Gamma}_{\text{pr}} &= \Gamma_{\text{pr}} \odot (\mathbf{a}\mathbf{a}^T), \\ \bar{\boldsymbol{\mu}}_E &= \boldsymbol{\mu}_E \odot \mathbf{c}, \\ \bar{\Gamma}_E &= \Gamma_E \odot (\mathbf{c}\mathbf{c}^T).\end{aligned}$$

The same normalization is applied to observational data both with and without noise. Once $(\bar{\boldsymbol{\mu}}_{\text{post}}, \bar{\Gamma}_{\text{post}})$ in the normalized scale is computed, we revert to the original scale using the transformation:

$$\begin{aligned}\boldsymbol{\mu}_{\text{post}} &= (\bar{\boldsymbol{\mu}}_{\text{post}} - \mathbf{b}) \oslash \mathbf{a}, \\ \Gamma_{\text{post}} &= \bar{\Gamma}_{\text{post}} \odot ((\mathbf{1}_D \oslash \mathbf{a})(\mathbf{1}_D \oslash \mathbf{a})^T).\end{aligned}$$

3. NUMERICAL RESULTS

We carry out three benchmark numerical tests to show the capabilities of UQ-VAEs: in section 3.1, we solve Bayesian inverse problems for the Poisson equation and study the impact of varying the number of neurons and layers of both the decoder and the encoder, as well as modifying θ and the robustness with respect to the noise level η , the number of observations O , the prior distribution $\mathcal{N}(\boldsymbol{\mu}_{\text{pr}}, \Gamma_{\text{pr}})$; in section 3.2, we face Bayesian inverse problems involving an exponential map \mathcal{F} , investigate the effect of varying the dimensions of the parameter and data space, D and O , and compare the performance, both in terms of computational time and accuracy, of UQ-VAEs trained using the loss function L_α (7) or L_θ (15); in section 3.3, we address Bayesian inverse problems for a 0D cardiocirculatory model in the context of systemic hypertension (section 3.3.1) and ventricular septal defects (section 3.3.2). The activation functions of the NNs are *ReLU* functions and the data are normalized to the range $[0, 1]$. Additionally, we always use the *Adam* optimizer.¹⁸ We tuned the number of neurons and layers of each NN to minimize the validation loss functions. The generated datasets consist of M samples, which are split in $\lfloor 9M/10 \rfloor$ training samples and $M - \lfloor 9M/10 \rfloor$ validation samples, where $\lfloor \cdot \rfloor$ denotes the lower integer part.

In the following sections, the true posterior mean and covariance matrix are approximated using Sobol' sequences,³³ a Quasi-Monte Carlo method that accelerates the convergence of the sample mean to the expected value. Using (2) and (3), we obtain:

$$\begin{aligned}(24) \quad \mathbb{E}_{\text{p}_{\text{U|Y}}(\mathbf{u}|\mathbf{y})}[\mathbf{u}] &= \frac{\mathbb{E}_{\text{p}_{\text{U}}(\mathbf{u})}[\mathbf{u} p_{\text{E}}(\mathbf{y} - \mathcal{F}(\mathbf{u}))]}{\mathbb{E}_{\text{p}_{\text{U}}(\mathbf{u})}[p_{\text{E}}(\mathbf{y} - \mathcal{F}(\mathbf{u}))]}, \\ \text{Cov}_{\text{p}_{\text{U|Y}}(\mathbf{u}|\mathbf{y})}(\mathbf{u}) &= \frac{\mathbb{E}_{\text{p}_{\text{U}}(\mathbf{u})}[\mathbf{u}\mathbf{u}^T p_{\text{E}}(\mathbf{y} - \mathcal{F}(\mathbf{u}))]}{\mathbb{E}_{\text{p}_{\text{U}}(\mathbf{u})}[p_{\text{E}}(\mathbf{y} - \mathcal{F}(\mathbf{u}))]} - \\ (25) \quad &\frac{\mathbb{E}_{\text{p}_{\text{U}}(\mathbf{u})}[\mathbf{u} p_{\text{E}}(\mathbf{y} - \mathcal{F}(\mathbf{u}))] \mathbb{E}_{\text{p}_{\text{U}}(\mathbf{u})}[\mathbf{u} p_{\text{E}}(\mathbf{y} - \mathcal{F}(\mathbf{u}))]^T}{\mathbb{E}_{\text{p}_{\text{U}}(\mathbf{u})}[p_{\text{E}}(\mathbf{y} - \mathcal{F}(\mathbf{u}))]^2}.\end{aligned}$$

This procedure is computationally expensive and can only be performed for simple problems.

3.1. Poisson equation. We consider the Poisson equation on the unit square $\Omega = (0, 1)^2$:

$$(26) \quad \begin{cases} -\nabla(e^{u(\mathbf{x})})\nabla y(\mathbf{x}) = \frac{17}{4}\pi^2 \sin(2\pi x_1) \sin(2\pi x_2) & \mathbf{x} \in \Omega, \\ y(\mathbf{x}) = 0 & \mathbf{x} \in \partial\Omega, \end{cases}$$

where $u(\mathbf{x})$ is the (log) diffusion parameter. We approximate the solution $y(\mathbf{x})$ of the Poisson problem by the linear finite element method on a triangular mesh of $D = 17^2 = 289$ degrees of freedom.

We define a Gaussian random field on the whole domain Ω , with mean and covariance functions $\mu_{\text{pr}}(\mathbf{x})$ and $\mathcal{C}_{\text{pr}}(\mathbf{x}, \mathbf{z})$, for $\mathbf{x}, \mathbf{z} \in (0, 1)$. The parameter U is given by the evaluation of this Gaussian random field at the degrees of freedom of the mesh, obtaining $U \sim \mathcal{N}(\boldsymbol{\mu}_{\text{pr}}, \Gamma_{\text{pr}})$. We define μ_{pr} to be zero in Ω and the covariance function \mathcal{C}_{pr} to ensure the well-posedness of infinite dimensional Bayesian inverse problems.^{3,37} We define the differential operator \mathcal{A} :

$$\mathcal{A}u = -\gamma\Delta u + \delta u \quad \text{in } \Omega,$$

where $\gamma, \delta > 0$. The operator $\mathcal{A}^{-2}(\mathbf{x}, \mathbf{z}) = G_2(\mathbf{x}, \mathbf{z})$, for $(\mathbf{x}, \mathbf{z}) \in \Omega$, is defined as the solution operator of:

$$\begin{cases} \mathcal{A}G_1(\mathbf{x}, \mathbf{z}) = \delta_{\mathbf{x}} & \mathbf{z} \in \Omega, \\ \nabla G_1(\mathbf{x}, \mathbf{z}) \cdot \mathbf{n}(\mathbf{z}) + \beta G_1(\mathbf{x}, \mathbf{z}) = 0 & \mathbf{z} \in \partial\Omega, \\ \mathcal{A}G_2(\mathbf{x}, \mathbf{z}) = G_1(\mathbf{x}, \mathbf{z}) & \mathbf{z} \in \Omega, \\ \nabla G_2(\mathbf{x}, \mathbf{z}) \cdot \mathbf{n}(\mathbf{z}) + \beta G_2(\mathbf{x}, \mathbf{z}) = 0 & \mathbf{z} \in \partial\Omega, \end{cases}$$

where, for $\mathbf{x} \in \Omega$, $\delta_{\mathbf{x}}$ is the Dirac delta centered in \mathbf{x} , \mathbf{n} is the unit exterior normal to $\partial\Omega$ and β is chosen to reduce the boundary artifacts.⁸ We set $\gamma = 0.1$ and $\delta = 0.5$. The covariance function is given by $\mathcal{C}_{\text{pr}}(\mathbf{x}, \mathbf{z}) = \mathcal{A}^{-\frac{3}{2}}(\mathbf{x}, \mathbf{z})$. We compute \mathcal{C}_{pr} using the same mesh and finite element method as for solving the Poisson problem, via the Python library for inverse problems *hIPPYlib*.^{46–48} Finally, we define $(\boldsymbol{\mu}_{\text{pr}})_i = \mu_{\text{pr}}(\mathbf{x}_i)$ and $(\Gamma_{\text{pr}})_{i,j} = \mathcal{C}_{\text{pr}}(\mathbf{x}_i, \mathbf{x}_j)$ for $i, j = 1, \dots, D$, where \mathbf{x}_i are the coordinates of the degrees of freedom of the mesh. The observational data consist of noisy evaluations of the solution $y(\mathbf{x})$ of (26) at $O = 20$ random points in the domain Ω .

We generate a dataset $\{(\mathbf{u}^{(m)}, \mathbf{y}^{(m)})\}_{m=1}^M$, where $M \in \mathbb{N}$ and $\mathbf{u}^{(m)}$ is sampled from $\mathcal{N}(\boldsymbol{\mu}_{\text{pr}}, \Gamma_{\text{pr}})$. Let $\tilde{\mathbf{y}}^{(m)}$ be the noiseless observations of the solution of (26) with parameter $\mathbf{u}^{(m)}$. Then, $\mathbf{y}^{(m)} = \tilde{\mathbf{y}}^{(m)} + \boldsymbol{\varepsilon}^{(m)}$, for $m = 1, \dots, M$, where $\boldsymbol{\varepsilon}^{(m)}$ is sampled from $\mathcal{N}(\boldsymbol{\mu}_{\text{E}}, \Gamma_{\text{E}})$ with $\boldsymbol{\mu}_{\text{E}} = \mathbf{0}$ and Γ_{E} is a diagonal matrix with $(\Gamma_{\text{E}})_{i,i} = (\eta \max_{j,m} |\tilde{y}_j^{(m)}|)^2$, for $i = 1, \dots, O$ and $\eta \in \mathbb{R}$. We fix $\eta = 0.01$.

To normalize the data to $[0, 1]$, we apply a uniform normalization to the parameters and observational data across the entire mesh. We define:

$$\begin{aligned} a &= (\max_{i,m} u_i^{(m)} - \min_{i,m} u_i^{(m)})^{-1}, \\ b &= -a \min_{i,m} u_i^{(m)}, \\ c &= (\max_{i,m} y_i^{(m)} - \min_{i,m} y_i^{(m)})^{-1}, \\ d &= -c \min_{i,m} y_i^{(m)} \end{aligned}$$

and set $\bar{u}(\mathbf{x}) = u(\mathbf{x})a + b$ and $\bar{y}(\mathbf{x}) = y(\mathbf{x})c + d$ in the continuous framework. Referring to section 2.2.1, in the discrete setting, $\mathbf{a} = a\mathbf{1}_D$, $\mathbf{b} = b\mathbf{1}_D$, $\mathbf{c} = c\mathbf{1}_O$ and $\mathbf{d} = d\mathbf{1}_O$. We can derive the Poisson problem in the normalized scale:

$$\begin{cases} -\nabla(e^{\frac{\bar{u}(\mathbf{x})}{a}} \nabla \bar{y}(\mathbf{x})) = ce^{\frac{b}{a}} \frac{17}{4} \pi^2 \sin(2\pi x_1) \sin(2\pi x_2) & \mathbf{x} \in \Omega, \\ \bar{y}(\mathbf{x}) = d & \mathbf{x} \in \partial\Omega. \end{cases}$$

Decoder						
Neurons \ Layers	1	2	3	4	5	6
50	$5.8 \cdot 10^{-4}$	$2.8 \cdot 10^{-4}$	$2.2 \cdot 10^{-4}$	$2.0 \cdot 10^{-4}$	$2.3 \cdot 10^{-4}$	$2.1 \cdot 10^{-4}$
100	$7.4 \cdot 10^{-4}$	$3.1 \cdot 10^{-4}$	$2.1 \cdot 10^{-4}$	$1.8 \cdot 10^{-4}$	$1.5 \cdot 10^{-4}$	$1.2 \cdot 10^{-4}$
150	$6.4 \cdot 10^{-4}$	$2.8 \cdot 10^{-4}$	$2.3 \cdot 10^{-4}$	$1.6 \cdot 10^{-4}$	$1.3 \cdot 10^{-4}$	$1.4 \cdot 10^{-4}$
200	$6.4 \cdot 10^{-4}$	$3.0 \cdot 10^{-4}$	$2.4 \cdot 10^{-4}$	$1.8 \cdot 10^{-4}$	$2.1 \cdot 10^{-4}$	$1.5 \cdot 10^{-4}$
250	$5.9 \cdot 10^{-4}$	$3.4 \cdot 10^{-4}$	$2.3 \cdot 10^{-4}$	$2.5 \cdot 10^{-4}$	$2.1 \cdot 10^{-4}$	$1.8 \cdot 10^{-4}$

Encoder				
Neurons \ Layers	1	2	3	4
500	26.8	24.0	25.2	25.7
750	25.6	24.2	24.6	25.5
1000	24.5	23.5	24.9	25.6
1250	24.5	24.0	25.3	26.5

TABLE 1. Final validation loss functions for the decoder and the encoder for the Poisson problem.

We perform hyperparameter tuning to select the optimal number of layers and neurons per layer for both the decoder and the encoder neural networks. We first determine the best architecture for the decoder network and then use this decoder to train the encoder with parameter $\theta = 10^{-4}$. According to the hyperparameter tuning results (Table 1), we select a decoder with 6 hidden layers and 100 neurons per layer, trained on a dataset of $M = 4096$ samples. The chosen encoder has 2 hidden layers and 1000 neurons per layer and is trained on a subset of 100 samples of the decoder dataset. Although the dimension of the encoder dataset is small with respect to the decoder one, we show in the following that it still generalizes to unseen data.

We generate 2^{19} samples from the distribution $\mathcal{N}(\boldsymbol{\mu}_{\text{pr}}, \Gamma_{\text{pr}})$ to estimate (24) and (25).

We analyze the UQ-VAE performances varying θ in the loss function (23). For the posterior means $\mathbb{E}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$ and variances $\text{Var}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$ (the diagonal of (25)), with $m \in \{1, \dots, M\}$, we compute the relative error with respect to the infinity norms at each degree of freedom:

$$\begin{aligned}
 (27) \quad \mathbf{e}_{\mathbb{E}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]} &= \frac{|\boldsymbol{\mu}_{\text{post}}^{(m)} - \mathbb{E}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]|}{\|\mathbb{E}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]\|_{\infty}}, \\
 \mathbf{e}_{\text{Var}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]} &= \frac{|\text{diag}(\Gamma_{\text{post}}^{(m)}) - \text{Var}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}(\mathbf{u})|}{\|\text{Var}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]\|_{\infty}},
 \end{aligned}$$

where $(\boldsymbol{\mu}_{\text{post}}^{(m)}, \Gamma_{\text{post}}^{(m)})$ are the UQ-VAE mean and covariance estimates for the m -th sample, $|\cdot|$ is applied element-wise and $\|\cdot\|_{\infty}$ is the vector infinite norm. The maximum over the mesh of these errors is the relative error on the m -th sample in the infinity norm.

The average relative errors on the posterior mean for a test set of 100 samples share similar trends and magnitudes for all values of θ (Figure 3). The maximum average relative error on the posterior variance is more than 50% for $\theta = 10^{-5} \vee 10^{-6}$, due to the small gradient of L_{θ} (23) with respect to C which scales with θ^2 . For the other values of θ the errors on the variance are similar. We choose to use in all the tests that follow (also in section 3.2 and section 3.3) $\theta = 10^{-4}$, because Theorem 2 holds for $\theta \rightarrow 0$.

For $\theta = 10^{-4}$, the average relative error on the posterior mean and the posterior variance are less than 11%. Even if there are differences in $\mathbb{E}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$ and $\boldsymbol{\mu}_{\text{post}}^{(m)}$, they estimate the true diffusion coefficients $\mathbf{u}^{(m)}$ with the same accuracy (Figure 4), showing that the




FIGURE 3. Average relative errors (in logarithmic scale) of the UQ-VAE mean $\mu_{\text{post}}^{(m)}$ and variance $\Gamma_{\text{post}}^{(m)}$ on the posterior mean $\mathbb{E}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$ and on the posterior variance $\text{Var}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$ in dependence of θ for a test set of 100 samples for the Poisson problem.

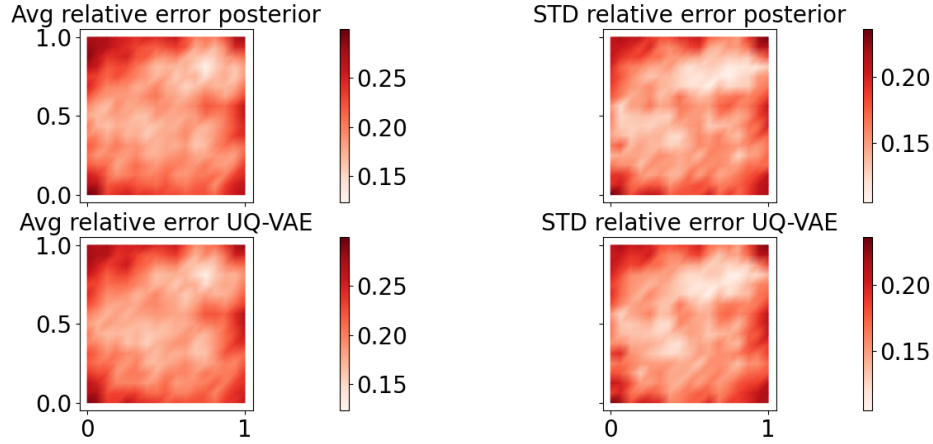


FIGURE 4. Averages and standard deviations of the relative errors of the posterior mean $\mathbb{E}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$ and $\mu_{\text{post}}^{(m)}$ on the true diffusion coefficients $\mathbf{u}^{(m)}$ for a test set of 100 samples for the Poisson problem.

	Offline phase	Online phase (100 samples)
UQ-VAE	5.5h	10.74s
Sobol' sequences	0.6h	8880s

TABLE 2. Comparison between the time required to solve Bayesian inverse problems with UQ-VAE or approximating (24) and (25) with Sobol' sequences.

UQ-VAE approach is effective in estimating parameters in a noisy setting. The relative errors on the true diffusion coefficients are computed by substituting $\mathbb{E}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$ with $\mathbf{u}^{(m)}$ in (27).

The offline computational time for the UQ-VAE (generating the dataset and training the decoder and the encoder) is 5.5h, whereas the one for generating 2^{19} samples to estimate (24) and (25) is 0.6h (Table 2). The UQ-VAE approach is advantageous on the online




FIGURE 5. True solutions $y^{(m)}(\mathbf{x})$ with marked observation points, true diffusion coefficients $\mathbf{u}^{(m)}$, posterior and UQ-VAE mean ($\mathbb{E}_{p_{U|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$ and $\mu_{\text{post}}^{(m)}$) and variance ($\text{Var}_{p_{U|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$ and $\Gamma_{\text{post}}^{(m)}$) estimates for three samples. The last two columns are the posterior or UQ-VAE means (continue blue lines) and true diffusion coefficients (dashed red lines) evaluated at $y = 0.5$. The width of the shaded areas is equal to two times the standard deviation.

phase, when the Bayesian inverse problems are solved, with an average elapsed time of $0.1s$, compared to the $88.8s$ required to estimate (24) and (25).

We analyze, for three test samples, the posterior and UQ-VAE estimates for the mean and variance (Figure 5). The three diffusion coefficients are highly irregular and, when the data are not informative enough, the estimates are poor, like in the third case. For the first two cases, both the posterior and UQ-VAE mean are capable of capturing the trend and the magnitude of the true diffusion coefficient. The two estimates are smoother than the true diffusion coefficient. The UQ-VAE variance is comparable to the posterior one. The estimated mean and variance together are able to catch the irregularity of the true diffusion coefficient, both for the posterior and the UQ-VAE estimates.

We test the UQ-VAE approach on different values for the noise level η and the number of observation points O (Figure 6). For all the values of η and O the posterior and UQ-VAE mean estimates are in agreement. For $\eta = 0.05$ the UQ-VAE variance estimates slightly differ from the posterior ones, whereas for $\eta = 0.01$ the two variances are similar. This also happens for $\eta = 0.005$ with $O = 5 \vee 20$. In all such cases, the variance is smaller near the observational data. For low values of η and large number of observations O , the posterior distribution is less regular and the variance estimates get worse, even with 2^{19} Quasi-Monte Carlo samples, making the comparison with the UQ-VAE estimates unfair (as can be seen for $\eta = 0.005$ and $O = 40$). Indeed, as η decreases or O increases, the determinant of the noise covariance matrix Γ_E decreases. The reduction of the regularity of Γ_E corresponds to a lower regularity for $p_{U|Y}(\mathbf{u}|\mathbf{y})$ (2), due to its dependence on $\|\cdot\|_{\Gamma_E^{-1}}$ and $|\Gamma_E|^{-1/2}$. Therefore, (24) and (25) require an elevated number of samples to be estimated accurately, resulting in a high (and unaffordable) computational cost.

Finally, we test the UQ-VAE approach for different exponents ξ of the covariance function $\mathcal{C}_{\text{pr}}(\mathbf{x}, \mathbf{z}) = \mathcal{A}^{-\xi}$ (Figure 7). We consider three different values of ξ in $(1, 2)$ and use the same neural network architecture for all cases. For $\xi = 3/2$, we obtain the results already analyzed in Figures 3 and 4. For $\xi = 7/4$, the posterior variance is estimated with similar accuracy as in the $\xi = 3/2$, while the errors in the estimates of the true parameters and posterior means are slightly higher. This increase in error is due to hyperparameters tuning being performed exclusively for the case $\xi = 3/2$. Nonetheless, the errors remain




FIGURE 6. Posterior and UQ-VAE estimates for a test sample, with noise level η varying across rows and number of observation points O varying across columns. The first row shows the true diffusion coefficient and the corresponding solution with marked observation points. Estimated means are shown in the second, fourth and sixth rows, while the third, fifth, seventh rows display the corresponding estimated variances.

of the same order of magnitude. In the case $\xi = 5/4$, the average error on the true parameters is comparable to the other two cases, whereas the errors on the posterior means and variances are higher. As for a low noise level η and a high number of observations O , smaller values of ξ reduce the regularity of the prior distribution, which in turn decreases the regularity of the posterior distribution. As a result, (24) and (25) would require a large (and computationally infeasible) number of samples to be estimated accurately.

3.2. Exponential Bayesian inverse problem. We test the UQ-VAE approach on a non linear Bayesian inverse problem with forward map $\mathcal{F}(\mathbf{u}) = e^{\mathbf{u}}$, where the exponentiation is applied element wise. We compare the training time and accuracy of the UQ-VAE trained with loss functions L_α (7) and L_θ (15). To compute L_α , we use a Sobol' sequence with $K = 2^{12}$ samples. In this test, we compare the UQ-VAE estimates with the maximum




FIGURE 7. Average relative errors of UQ-VAE estimates on the true parameters and posterior means and variances on 100 test samples.




FIGURE 8. Performance comparison in terms of computational time and accuracy of UQ-VAEs (tested on 100 samples) trained with the loss function L_α and L_θ on the non linear Bayesian inverse problem with forward map $\mathcal{F}(\mathbf{u}) = e^{\mathbf{u}}$.

a posteriori parameter \mathbf{u}_{MAP} (4) and the Laplace approximation of the covariance matrix Γ_{Lap} (5) because the true posterior mean (24) and covariance (25) are expensive to estimate with an exponential map \mathcal{F} .

We set the dimensions of the parameter and observational data space equal ($D = O$) and vary them from 25 and 150. We set $(\mu_{\text{pr}})_i = -3$, for $i = 1, \dots, D$, and $\Gamma_{\text{pr}} = 4I$, where $I \in \mathbb{R}^{D \times D}$ is the identity matrix.

We generate a dataset $\{(\mathbf{u}^{(m)}, \mathbf{y}^{(m)})\}_{m=1}^M$, where $M \in \mathbb{N}$ and $\mathbf{u}^{(m)}$ is sampled from $\mathcal{N}(\mu_{\text{pr}}, \Gamma_{\text{pr}})$. The noiseless observations are denoted by $\tilde{\mathbf{y}}^{(m)} = \mathcal{F}(\mathbf{u}^{(m)})$. The noisy observations are then defined as $\mathbf{y}^{(m)} = \tilde{\mathbf{y}}^{(m)} + \epsilon^{(m)}$, for $m = 1, \dots, M$, where $\epsilon^{(m)}$ is sampled from $\mathcal{N}(\mu_E, \Gamma_E)$ with $\mu_E = \mathbf{0}$ and Γ_E diagonal has entries $(\Gamma_E)_{i,i} = (\eta \max_m |\tilde{y}_i^{(m)}|)^2$, for $i = 1, \dots, O$ and $\eta \in \mathbb{R}$. We fix $\eta = 0.01$.

We use the true map \mathcal{F} as the decoder. The encoder consists of 3 hidden layers with 1000 neurons each and is trained on a dataset of 100 samples.

While the training time of the UQ-VAE using L_α loss increases with D and O , the training time with L_θ loss is nearly constant (Figure 8). This is because the computational cost of the training with L_θ is mainly driven by the network dimension. In contrast, with




FIGURE 9. 0D cardiocirculatory model. We depict pressures and flow rates in red and blue, respectively, and parameters in black.

L_α , the cost of sampling significantly contributes to the training time, even for small values of D . The average errors (on a dataset of 100 samples) in the estimation of \mathbf{u}_{MAP} , Γ_{Lap} and true parameters follow a similar trend for both loss functions. Nonetheless, training with L_θ yields lower errors across all metrics. The errors in estimating the true parameters with both loss functions are slightly higher than those of \mathbf{u}_{MAP} , but still comparable.

3.3. 0D cardiocirculatory model. 0D (or lumped-parameter) cardiocirculatory models represent the human cardiovascular system as an electrical circuit: the current corresponds to the blood flow through vessels and valves, the electric potential represents the blood pressure, the electric resistance plays the role of the resistance to blood flow, the capacitance represents the vessel compliance and the inductance corresponds to the blood inertia. In this context, Bayesian inverse problems are used to personalize the model for individual patients and provide insights into their clinical condition, while accounting for the uncertainty in the estimated parameters due to measurement errors (and poor observability of the parameters themselves).

A 0D cardiocirculatory model partitions the cardiovascular system into distinct compartments (e.g. right atrium, systemic arteries/veins). The 0D cardiocirculatory model used here is the one proposed in Ref. 28 (Figure 9). The model consists in the four cardiac chambers, the systemic and pulmonary circulation, split into arterial and venous compartments.

The 0D cardiocirculatory model depends on the heart rate (HR), which determines the heartbeat period $T_{\text{HB}} = 60/\text{HR}$, and on the parameters $\tilde{\mathbf{u}}$ listed in Table 3. A subset \mathbf{u} of

	Parameter	Unit	Reference value	Description
GSA	EA_{LA}	mmHg/mL	0.2273	Left atrial active elastance
	EB_{LA}	mmHg/mL	0.209	Left atrial passive elastance
	$V_{U,LA}$	mL	2.0	Left atrial unstressed volume
	EA_{LV}	mmHg/mL	3.0391	Left ventricular active elastance
	EB_{LV}	mmHg/mL	0.10	Left ventricular passive elastance
	$V_{U,LV}$	mL	16.0	Left ventricular unstressed volume
	EA_{RA}	mmHg/mL	0.0429	Right atrial active elastance
	EB_{RA}	mmHg/mL	0.0636	Right atrial passive elastance
	$V_{U,RA}$	mL	2.0	Right atrial unstressed volume
	EA_{RV}	mmHg/mL	0.6683	Right ventricular active elastance
	EB_{RV}	mmHg/mL	0.07	Right ventricular passive elastance
	$V_{U,RV}$	mL	16.0	Right ventricular unstressed volume
	R_{min}	mmHg · s/mL	0.0075	Minimal valve resistance
	R_{max}	mmHg · s/mL	75006.2	Maximal valve resistance
	R_{AR}^{SYS}	mmHg · s/mL	0.588	Systemic arterial resistance
	C_{AR}^{SYS}	mL/mmHg	0.96	Systemic arterial compliance
	L_{AR}^{SYS}	mmHg · s ² /mL	0.005	Systemic arterial inertia
	R_{VEN}^{SYS}	mmHg · s/mL	0.352	Systemic venous resistance
	C_{VEN}^{SYS}	mL/mmHg	60.0	Systemic venous compliance
	L_{VEN}^{SYS}	mmHg · s ² /mL	0.0005	Systemic venous inertia
	R_{PUL}^{AR}	mmHg · s/mL	0.104	Pulmonary arterial resistance
	C_{PUL}^{AR}	mL/mmHg	5.0	Pulmonary arterial compliance
	L_{PUL}^{AR}	mmHg · s ² /mL	0.0005	Pulmonary arterial inertia
	R_{VEN}^{PUL}	mmHg · s/mL	0.0105	Pulmonary venous resistance
	C_{PUL}^{VEN}	mL/mmHg	16.0	Pulmonary venous compliance
	L_{VEN}^{PUL}	mmHg · s ² /mL	0.0005	Pulmonary venous inertia
	HR	s	75	Heart rate
No GSA	tC_{LA}	s	$0.79T_{HB}$	Time of left atrial contraction
	TC_{LA}	s	$0.11T_{HB}$	Duration of left atrial contraction
	tR_{LA}	s	$tC_{LA} + TC_{LA}$	Time of left atrial relaxation
	TR_{LA}	s	$0.8T_{HB}$	Duration of left atrial relaxation
	tC_{LV}	s	0.0	Time of left ventricular contraction
	TC_{LV}	s	$0.35T_{HB}$	Duration of left ventricular contraction
	tR_{LV}	s	$tC_{LV} + TC_{LV}$	Time of left ventricular relaxation
	TR_{LV}	s	$0.4T_{HB}$	Duration of left ventricular relaxation
	tC_{RA}	s	$0.8T_{HB}$	Time of right atrial contraction
	TC_{RA}	s	$0.1T_{HB}$	Duration of right atrial contraction
	tR_{RA}	s	$tC_{RA} + TC_{RA}$	Time of left atrial relaxation
	TR_{RA}	s	$0.7T_{HB}$	Duration of left atrial relaxation
	tC_{RV}	s	0.0	Time of right ventricular contraction
	TC_{RV}	s	$0.3T_{HB}$	Duration of left ventricular contraction
	tR_{RV}	s	$tC_{RV} + TC_{RV}$	Time of right ventricular relaxation
	TR_{RV}	s	$0.4T_{HB}$	Duration of left ventricular relaxation

TABLE 3. List of parameters and their reference values for an ideal healthy individual

these parameters is used to define the model (1), depending on the specific application. The hemodynamics of the cardiovascular system are described by the following dynamical

system:

$$(28) \quad \begin{cases} \dot{V}_{LA}(t) = Q_{VEN}^{PUL}(t) - Q_{MV}(t), \\ \dot{V}_{LV}(t) = Q_{MV}(t) - Q_{AV}(t), \\ C_{AR}^{SYS} \dot{p}_{AR}^{SYS}(t) = Q_{AV}(t) - Q_{AR}^{SYS}(t), \\ L_{AR}^{SYS} \dot{Q}_{AR}^{SYS}(t) = -R_{AR}^{SYS} Q_{AR}^{SYS}(t) + p_{AR}^{SYS}(t) - p_{VEN}^{SYS}(t), \\ C_{VEN}^{SYS} \dot{p}_{VEN}^{SYS}(t) = Q_{AR}^{SYS}(t) - Q_{VEN}^{SYS}(t), \\ L_{VEN}^{SYS} \dot{Q}_{VEN}^{SYS}(t) = -R_{VEN}^{SYS} Q_{VEN}^{SYS}(t) + p_{VEN}^{SYS}(t) - p_{RA}(t), \\ \dot{V}_{RA}(t) = Q_{VEN}^{SYS}(t) - Q_{TV}(t), \\ \dot{V}_{RV}(t) = Q_{TV}(t) - Q_{PV}(t), \\ C_{AR}^{PUL} \dot{p}_{AR}^{PUL}(t) = Q_{PV}(t) - Q_{AR}^{PUL}(t), \\ L_{AR}^{PUL} \dot{Q}_{AR}^{PUL}(t) = -R_{AR}^{PUL} Q_{AR}^{PUL}(t) + p_{AR}^{PUL}(t) - p_{VEN}^{PUL}(t), \\ C_{VEN}^{PUL} \dot{p}_{VEN}^{PUL}(t) = Q_{AR}^{PUL}(t) - Q_{VEN}^{PUL}(t), \\ L_{VEN}^{PUL} \dot{Q}_{VEN}^{PUL}(t) = -R_{VEN}^{PUL} Q_{VEN}^{PUL}(t) + p_{VEN}^{PUL}(t) - p_{LA}(t), \end{cases}$$

coupled with suitable initial conditions \mathbf{x}_0 and with $t \in (0, T]$. The model unknowns are the volumes of the left atrium (V_{LA}) and ventricle (V_{LV}) and of the right atrium (V_{RA}) and ventricle (V_{RV}), the circulatory pressures of the systemic arteries (p_{AR}^{SYS}) and veins (p_{VEN}^{SYS}) and of the pulmonary arteries (p_{AR}^{PUL}) and veins (p_{VEN}^{PUL}), the circulatory fluxes of the systemic arteries (Q_{AR}^{SYS}) and veins (Q_{VEN}^{SYS}) and of the pulmonary arteries (Q_{AR}^{PUL}) and veins (Q_{VEN}^{PUL}).

The fluxes through the valves (tricuspid, pulmonary, mitral and aortic) depend on the pressure jump between the upstream and downstream compartments and the resistance through the valves is given by the non linear function R_{valve} as follows:

$$\begin{aligned} Q_{TV}(t) &= Q_{\text{valve}}(p_{RA}(t) - p_{RV}(t)), & Q_{MV}(t) &= Q_{\text{valve}}(p_{LA}(t) - p_{LV}(t)), \\ Q_{PV}(t) &= Q_{\text{valve}}(p_{RV}(t) - p_{AR}^{PUL}(t)), & Q_{AV}(t) &= Q_{\text{valve}}(p_{LV}(t) - p_{AR}^{SYS}(t)), \\ Q_{\text{valve}}(\Delta p) &= \frac{\Delta p}{R_{\text{valve}}(\Delta p)}, & R_{\text{valve}}(\Delta p) &= \sqrt{R_{\max} R_{\min}} \left(\frac{R_{\max}}{R_{\min}} \right)^{\frac{\arctan(-100\pi\Delta p)}{\pi}}. \end{aligned}$$

Each cardiac chamber is described as a pressure generator and depends on the periodic function $e_c(t)$ that models the periodicity of each heartbeat.

$$\begin{aligned} p_c(t) &= E_c(t)(V_c(t) - V_{U,c}), \\ E_c(t) &= EB_c + EA_c e_c(t), \end{aligned}$$

$$e_c(t) = \begin{cases} \frac{1}{2} \left[1 - \cos \left(\frac{\pi}{TC_c} \bmod(t - tC_c, T_{HB}) \right) \right] & \text{if } \bmod(t - tC_c, T_{HB}) < TC_c, \\ \frac{1}{2} \left[1 + \cos \left(\frac{\pi}{TR_c} \bmod(t - tR_c, T_{HB}) \right) \right] & \text{if } \bmod(t - tR_c, T_{HB}) < TR_c, \\ 0 & \text{otherwise,} \end{cases}$$

for $c \in \{LA, LV, RA, RV\}$.

We solve the discrete dynamical system for 25 heartbeats ($T = 25T_{HB}$) to approach its periodic state and focus on the last heartbeat. We compute the model outputs $\tilde{\mathbf{y}}$ as functions

Model output	Unit	Range	Model value	Description
$LA_{I-V_{\max}}$	mL/m ²	[16,34] ²⁰	33.1	Indexed maximal left atrial volume
$LA_{P_{\max}}$	mmHg	[6,20] ¹⁶	12.0	Maximal left atrial pressure
$LA_{P_{\min}}$	mmHg	[-2,9] ¹⁶	8.1	Minimal left atrial pressure
$LA_{P_{\text{mean}}}$	mmHg	[4,12] ¹⁶	10.9	Mean left atrial pressure
LV_{SV}	mL	[30,80] ²⁰	67.6	Left ventricular stroke volume
CI	L/min/m ²	[2.8,4.2] ¹⁶	2.8	Cardiac index
LV_{I-EDV}	mL/m ²	[50,90] ¹⁶	66.0	Indexed left ventricular end diastolic volume
LV_{ESV}	mL	[18,52] ²⁰	50.6	Left ventricular end systolic volume
LV_{EF}	%	[53,73] ²⁰	57.2	Left ventricular ejection fraction
$LV_{P_{\max}}$	mmHg	[90,140] ¹⁶	113.2	Maximal left ventricular pressure
$LV_{P_{\min}}$	mmHg	[4,12] ¹⁶	5.6	Minimal left ventricular pressure
$RA_{P_{\max}}$	mmHg	[2,14] ¹⁶	8.4	Maximal right atrial pressure
$RA_{P_{\min}}$	mmHg	[-2,6] ¹⁶	5.7	Minimal right atrial pressure
$RA_{P_{\text{mean}}}$	mmHg	[-1,8] ¹⁶	7.2	Mean right atrial pressure
RV_{I-EDV}	mL/m ²	[44,80] ²⁹	65.8	Indexed right ventricular end diastolic volume
RV_{I-ESV}	mL/m ²	[19,46] ²⁹	28.1	Indexed right ventricular end systolic volume
RV_{EF}	%	[44,71] ²⁹	57.3	Right ventricular ejection fraction
$RV_{P_{\max}}$	mmHg	[15,28] ¹⁶	27.4	Maximal right ventricular pressure
$RV_{P_{\min}}$	mmHg	[0,8] ¹⁶	3.9	Minimal right ventricular pressure
SAP_{\max}	mmHg	[-,140] ²⁰	112.2	Systolic systemic arterial pressure
SAP_{\min}	mmHg	[-,80] ²⁰	62.0	Diastolic systemic arterial pressure
PAP_{\max}	mmHg	[15,28] ¹⁶	25.8	Systolic pulmonary arterial pressure
PAP_{\min}	mmHg	[5,16] ¹⁶	16.0	Diastolic pulmonary arterial pressure
PAP_{mean}	mmHg	[10,22] ¹⁶	20.6	Mean pulmonary arterial pressure
PWP_{\min}	mmHg	[1,12] ¹⁶	11.2	Minimal pulmonary wedge pressure
PWP_{mean}	mmHg	[6,15] ¹⁶	11.8	Mean pulmonary wedge pressure
SVR	mmHg · min/L	[11.3,17.5] ¹⁶	15.7	Systemic vascular resistance
PVR	mmHg · min/L	[1.9,3.1] ¹⁶	1.9	Pulmonary vascular resistance

TABLE 4. List of model outputs, their units of measure, the echocardiographic ranges for a healthy individual and the values returned by the 0D cardiocirculatory model with the reference setting of parameters.

(e.g. the maximum or the mean) of the unknowns \mathbf{x} and the parameters $\tilde{\mathbf{u}}$:

$$\tilde{\mathbf{y}} = \mathbf{g}(\{\mathbf{x}(t; \mathbf{x}_0, \tilde{\mathbf{u}}), t \in [24T_{\text{HB}}, 25T_{\text{HB}}]\}; \tilde{\mathbf{u}}).$$

The model outputs are defined in Table 4, where a subscript “I-” refers to the indexed volumes, i.e., normalized by the body surface area (BSA). In general, a subset \mathbf{y} of the model outputs is used as observational data in model (1), depending on the specific application.

We determine the reference setting of parameters $\tilde{\mathbf{u}}^R$ to represent an ideal healthy individual (Table 3). Specifically, we set $\text{HR} = 75 \text{ bpm}$ ⁴⁹ and adjust the values of the other parameters based on modifications of literature values,^{1,10,43} ensuring that the model outputs $\tilde{\mathbf{y}}$ fall within echocardiographic ranges characteristic of a healthy individual (Table 4).

We perform a global sensitivity analysis (GSA) on the $\tilde{\mathbf{u}}$ parameters of the 0D cardiocirculatory model indicated in Table 3 to identify those that most influence a set of observational data \mathbf{y} . This allows to reduce the dimensionality of the parameter space by focusing on the most significant parameters. We compute total Sobol’ indices³⁴ to evaluate the impact of each model parameter on the variance of each observational data y_j , for $j = 1, \dots, O$, accounting for both first-order and higher-order interactions among parameters:

$$S_i^{j,T} = 1 - \frac{\text{Var}_{\tilde{p}(\tilde{\mathbf{u}}_{\sim i})}[\mathbb{E}_{\tilde{p}(\tilde{\mathbf{u}}_i)}[y_j | \tilde{\mathbf{u}}_{\sim i}]]}{\text{Var}_{\tilde{p}(\tilde{\mathbf{u}})}[y_j]},$$

where \tilde{p} is the sampling pdf of the parameters and $\tilde{\mathbf{u}}_{\sim i}$ represents the set of all parameters except for \tilde{u}_i . We estimate the total Sobol’ indices using Saltelli’s method,³⁰ allowing the

Parameter	EA_{LV}	$V_{U,LV}$	R_{AR}^{SYS}	C_{AR}^{SYS}	R_{VEN}^{SYS}	C_{VEN}^{SYS}	R_{AR}^{PUL}	C_{AR}^{PUL}	HR
Modification	+40%	-10%	+50%	-40%	+5%	-5%	+10%	-10%	+10%

TABLE 5. List of further modifications in the parameters ranges for the GSA to model hypertension. The modifications affect the lower bound (-) or the upper bound (+) of the ranges.

parameters $\tilde{\mathbf{u}}$ to vary in a hypercube. The number of samples for the Saltelli's method increases linearly with the number of parameters N_u and a user-defined value $N \in \mathbb{N}$ as $2N(N_u + 1)$. We set $N = 2^{12}$. Since the sensitivity analysis depends on the parameter ranges, we will specify them for the specific application (sections 3.3.1 and 3.3.2). We will solve Bayesian inverse problems for the parameters $\mathbf{u} \in \mathbb{R}^D$ associated with at least one Sobol index greater than 0.1. All the parameters not selected will be kept fixed at their reference value.

Remark. The heart rate (HR) is usually measured in clinical practice. In our framework there are two possibilities to treat HR : ignoring the data and estimating it by means of the UQ-VAE; developing one decoder that can handle different values for HR . Even if the second approach reduces the dimension of the parameter space by one, it would require the training of an encoder for each value of HR because the map \mathcal{F} of the mathematical model (1) would change with HR . Therefore, we restrain ourselves to the first approach.

3.3.1. *Systemic hypertension test case.* We solve the Bayesian inverse problem for the 0D cardiocirculatory model in the context of systemic hypertension.²⁷ The chosen observational data are:

$$\mathbf{y} = [CI, LV_{I-ESV}, LV_{EF}, LV_{Pmax}, SAP_{max}, SAP_{min}].$$

These observational data are directly affected by pressure increases in the systemic circulation.

The hypercube used for the sensitivity analysis is built around the reference parameter setting $\tilde{\mathbf{u}}^R$, with variations of $\pm 25\%$ relative to $\tilde{\mathbf{u}}^R$ to account for intervariability between individuals. Additionally, the parameters listed in Table 5 allow for further adjustments to the lower and upper bounds of the ranges to model systemic hypertension. Based on the total Sobol' indices (Figure 10), the selected parameters are:

$$\mathbf{u} = [EA_{LV}, R_{AR}^{SYS}, C_{AR}^{SYS}, R_{VEN}^{SYS}, HR].$$

Since all entries of the parameter vector \mathbf{u} of the 0D cardiocirculatory model are positive, we work with their logarithms to formulate an unconstrained inverse problem. We assume the parameter vector to be distributed as a lognormal random variable $\log U \sim \mathcal{N}(\boldsymbol{\mu}_{pr}, \Gamma_{pr})$. We set the prior mean as $\boldsymbol{\mu}_{pr} = \log \mathbf{u}^R$, where \mathbf{u}^R is the vector of reference values for \mathbf{u} and \log is applied element-wise. We let Γ_{pr} to be a diagonal matrix. During the sensitivity analysis, for $i = 1, \dots, D$, each entry u_i of the parameter vector varies in the interval $[a_i u_i^R, b_i u_i^R]$ with $a_i, b_i \in \mathbb{R}_+$, $a_i < b_i$, thus $\log u_i \in [\log a_i + \log u_i^R, \log b_i + \log u_i^R]$. We set $(\Gamma_{pr})_{i,i} = (\log b_i - \log a_i)^2 / 12$.

We generate a dataset $\{(\mathbf{u}^{(m)}, \mathbf{y}^{(m)})\}_{m=1}^M$, where $M \in \mathbb{N}$ and $\mathbf{u}^{(m)}$ is sampled from $\mathcal{N}(\boldsymbol{\mu}_{pr}, \Gamma_{pr})$. Let $\tilde{\mathbf{y}}^{(m)}$ be the noiseless observations of the solution of (28) with parameter $\mathbf{u}^{(m)}$. Then, $\mathbf{y}^{(m)} = \tilde{\mathbf{y}}^{(m)} + \boldsymbol{\varepsilon}^{(m)}$, for $m = 1, \dots, M$, where $\boldsymbol{\varepsilon}^{(m)}$ is sampled from $\mathcal{N}(\boldsymbol{\mu}_E, \Gamma_E)$ with $\boldsymbol{\mu}_E = \mathbf{0}$ and Γ_E diagonal with $(\Gamma_E)_{i,i} = (\eta \max_m |\tilde{y}_i^{(m)}|)^2$, for $i = 1, \dots, O$ and $\eta \in \mathbb{R}$. Observe that differently from section 3.1, $(\Gamma_E)_{i,i}$ depends on

	Sobol indices (S_T)					
	CI	LV _{LV}	LV _{RV}	LV _{max}	SA _{max}	SA _{min}
EA _{LA}	0.00	0.00	0.00	0.00	0.00	0.00
EB _{LA}	0.00	0.00	0.00	0.00	0.00	0.00
V0 _{LA}	0.00	0.00	0.00	0.00	0.00	0.00
EA _{LV}	0.01	0.66	0.44	0.01	0.01	0.00
EB _{LV}	0.01	0.00	0.00	0.01	0.01	0.00
V0 _{LV}	0.00	0.07	0.05	0.00	0.00	0.00
EA _{RA}	0.00	0.00	0.00	0.00	0.00	0.00
EB _{RA}	0.00	0.00	0.00	0.00	0.00	0.00
V0 _{RA}	0.00	0.00	0.00	0.00	0.00	0.00
EA _{RV}	0.01	0.00	0.00	0.00	0.00	0.00
EB _{RV}	0.03	0.00	0.00	0.01	0.01	0.00
V0 _{RV}	0.00	0.00	0.00	0.00	0.00	0.00
R _{min}	0.00	0.00	0.00	0.00	0.00	0.00
R _{max}	0.00	0.00	0.00	0.00	0.00	0.00
R _{AR} ^{SYS}	0.01	0.20	0.15	0.55	0.56	0.69
C _{AR} ^{SYS}	0.00	0.03	0.02	0.20	0.20	0.16
L _{AR} ^{SYS}	0.00	0.00	0.00	0.00	0.00	0.00
R _{VEN} ^{SYS}	0.79	0.05	0.05	0.21	0.20	0.04
C _{VEN} ^{SYS}	0.02	0.00	0.00	0.01	0.01	0.01
L _{VEN} ^{SYS}	0.00	0.00	0.00	0.00	0.00	0.00
R _{AR} ^{PUL}	0.00	0.00	0.00	0.00	0.00	0.00
C _{AR} ^{PUL}	0.00	0.00	0.00	0.00	0.00	0.00
L _{AR} ^{PUL}	0.00	0.00	0.00	0.00	0.00	0.00
R _{VEN} ^{PUL}	0.00	0.00	0.00	0.00	0.00	0.00
C _{VEN} ^{PUL}	0.00	0.00	0.00	0.00	0.00	0.00
L _{VEN} ^{PUL}	0.00	0.00	0.00	0.00	0.00	0.00
HR	0.11	0.00	0.29	0.00	0.00	0.12

FIGURE 10. Total Sobol' indices for 0D cardiocirculatory model in the hypertension case.

the specific observational data because they are associated with different quantities. We fix $\eta = 0.05$.

We use a decoder with 5 hidden layers and 250 neurons per layer trained on a dataset of $M = 4096$ samples. The encoder consists of 3 hidden layers and 250 neurons per layer and is trained on a subset of 100 samples of the decoder dataset.

We generate 2^{17} samples from the distribution $\mathcal{N}(\mu_{pr}, \Gamma_{pr})$ to estimate (24) and (25).

The UQ-VAE mean $\mu_{post}^{(m)}$ estimates the posterior mean $\mathbb{E}_{p_{U|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$ with an average relative error less than 2.0% on 100 test samples (Figure 11). The UQ-VAE covariance $\Gamma_{post}^{(m)}$ estimates the posterior covariance $\text{Cov}_{p_{U|Y}(\mathbf{u}|\mathbf{y}^{(m)})}(\mathbf{u})$ with an average relative error less than 0.5%. Both the posterior and the UQ-VAE means achieve similar average relative errors on the true parameters $\mathbf{u}^{(m)}$.

The chosen observational data \mathbf{y} are related to the practical case of systemic hypertension. We now test the UQ-VAE using 25 equally spaced observations of all the 12 unknowns of (28) in $[24T_{HB}, 25T_{HB}]$ ($O = 300$). We perform a new sensitivity analysis to account for the different observational data and we select the parameters:

$$\mathbf{u} = [EA_{LA}, EB_{LA}, EA_{LV}, EB_{LV}, EB_{RA}, EA_{RV}, EB_{RV}, V_{U,RV}, R_{AR}^{SYS}, C_{AR}^{SYS}, R_{VEN}^{SYS}, C_{VEN}^{SYS}, R_{AR}^{PUL}, C_{AR}^{PUL}, L_{AR}^{PUL}, HR].$$

The definition of the prior distribution and of the dataset are analogous to the previous case. The noise distribution has mean $\mu_E = \mathbf{0}$ and diagonal covariance matrix Γ_E . Γ_E is such that all the observations of the same unknown of (28) are related to the same variance obtained by taking the maximum of the absolute values of all the corresponding samples

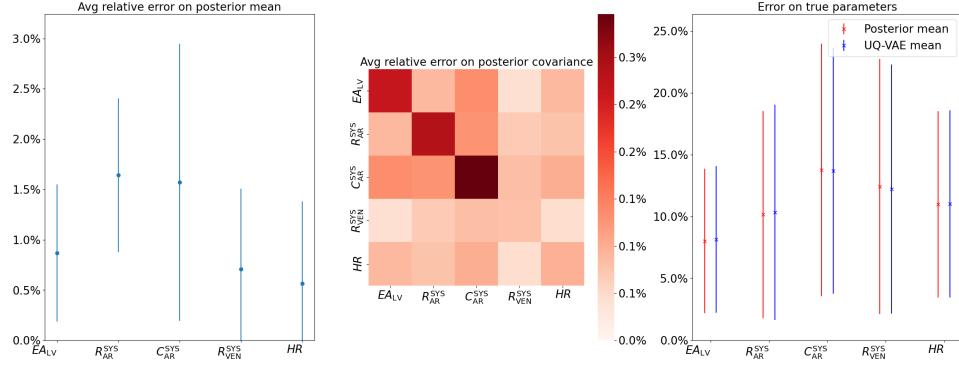


FIGURE 11. From left to right: average and standard deviation of the relative errors of the UQ-VAE mean $\mu_{\text{post}}^{(m)}$ on the posterior one $\mathbb{E}_{p_{U|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$; average relative error of the UQ-VAE covariance $\Gamma_{\text{post}}^{(m)}$ on the posterior one $\text{Var}_{p_{U|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$; average and standard deviations of the relative errors of the posterior and UQ-VAE means on the true parameters $\mathbf{u}^{(m)}$. The test are run for 100 samples for the systemic hypertension case and the relative errors and standard deviations are computed element-wise.

and observations:

$$(\Gamma_E)_{i,i} = \left(\eta \max_m \left| \tilde{y}_k^{(m)} \right| \right)^2_{1+25\text{mod}(i-1,25) \leq k \leq 25(\text{mod}(i-1,25)+1)}.$$

We generate 2^{19} samples from the distribution $\mathcal{N}(\mu_{\text{pr}}, \Gamma_{\text{pr}})$ to estimate (24) and (25).

In this case, the UQ-VAE mean is more accurate than the posterior mean in estimating the true parameters (Figure 12). As in Figure 6, this is because the increase in the number of observational data corresponds to a decrease in the regularity of $p_{U|Y}(\mathbf{u}|\mathbf{y})$ and a poorer estimate of (24) and (25). Consequently, also the estimates of the posterior mean and covariance matrix are worse than Figure 11 with average relative errors less than 15% and 1.6%, respectively. Moreover, the accuracy of the UQ-VAE on the true parameters is unaffected by the higher number of parameters to estimate with respect to the previous case.

3.3.2. Ventricular septal defect test case. We solve the Bayesian inverse problem for the 0D cardiocirculatory model in the context of ventricular septal defect.^{7,36} This congenital defect is characterized by a hole in the septum (wall) between the two ventricles. To account for the blood flow through the ventricular septal defect, the 0D cardiocirculatory model requires a modification in the differential equations for the volumes of the left and right ventricles:

$$(29) \quad \begin{cases} \dot{V}_{LV}(t) = Q_{MV}(t) - Q_{AV}(t) - Q_{VSD}(t), \\ \dot{V}_{RV}(t) = Q_{TV}(t) - Q_{PV}(t) + Q_{VSD}(t), \end{cases}$$

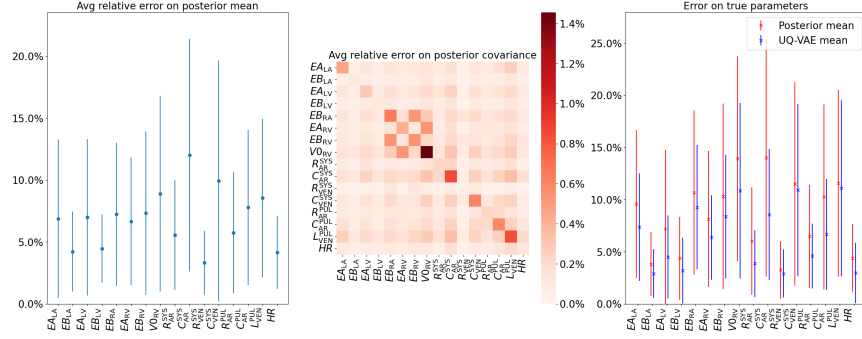


FIGURE 12. From left to right: average and standard deviation of the relative errors of the UQ-VAE mean $\mu_{\text{post}}^{(m)}$ on the posterior one $\mathbb{E}_{\text{p}_{\text{U}|Y}}(\mathbf{u}|\mathbf{y}^{(m)})[\mathbf{u}]$; average relative error of the UQ-VAE covariance $\Gamma_{\text{post}}^{(m)}$ on the posterior one $\text{Var}_{\text{p}_{\text{U}|Y}}(\mathbf{u}|\mathbf{y}^{(m)})[\mathbf{u}]$; average and standard deviations of the relative errors of the posterior and UQ-VAE means on the true parameters $\mathbf{u}^{(m)}$. The test are run for 100 samples for the systemic hypertension case and the relative errors and standard deviations are computed element-wise. The data are 25 equally spaced observations of all the 12 unknowns of (28).

where:

$$Q_{\text{VSD}}(t) = \frac{p_{\text{LV}}(t) - p_{\text{RV}}(t)}{R_{\text{VSD}}},$$

with R_{VSD} representing the resistance to the blood flow through the ventricular septal defect. We use as parameter the radius of the ventricular septal defect r_{VSD} with a reference value of $0.6r_{\text{AV}} = 0.9$ cm, where $r_{\text{AV}} = 1.5$ cm is the aortic annulus radius.⁷ We compute R_{VSD} by assuming that the length of the ventricular septal defect and the thickness of the aortic annulus are equal to L and applying two times the Poiseuille law:

$$R_{\text{VSD}} = \frac{8\mu L}{\pi r_{\text{VSD}}^4} = R_{\text{min}} \left(\frac{r_{\text{AV}}}{r_{\text{VSD}}} \right)^4,$$

where μ denotes the blood viscosity.

The observational data are chosen as follows:

$$\mathbf{y} = [LV_{\text{Pmax}}, LV_{\text{Pmin}}, RA_{\text{Pmean}}, RV_{\text{Pmax}}, RV_{\text{Pmin}}, SAP_{\text{max}}, SAP_{\text{min}}, PAP_{\text{max}}, PAP_{\text{min}}, PAP_{\text{mean}}, PWP_{\text{mean}}, PVR, Q_{\text{P}}, Q_{\text{S}}],$$

where Q_{P} and Q_{S} are the mean flows through the pulmonary and aortic valves. These two observational data are particularly important in the case of ventricular septal defects because their ratio $Q_{\text{P}}/Q_{\text{S}}$ is typically around 1 in healthy individuals, indicating that the same amount of blood flows through the pulmonary and systemic circulation. In the case of ventricular septal defect, this ratio increases beyond one, meaning that some of the oxygenated blood from the lungs flows to the right ventricle, bypassing the systemic circulation and failing to supply oxygen to the body's cells.

The hypercube used for the sensitivity analysis is built around the reference setting of parameters $\tilde{\mathbf{u}}^{\text{R}}$, accounting for intervariability between individuals by allowing a variation

	Sobol indices (S_T)													
	LVmax	LVmin	RAmax	RVmax	RVmin	SAPmax	SAPmin	PAPmax	PAPmin	PAPmean	PWPmean	PVR	Qp	Qs
EA _{LA}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
EB _{LA}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.00
V0 _{LA}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
EA _{LV}	0.01	0.03	0.00	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
EB _{LV}	0.07	0.45	0.02	0.00	0.01	0.07	0.02	0.01	0.03	0.01	0.15	0.00	0.02	0.04
V0 _{LV}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
EA _{RA}	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
EB _{RA}	0.00	0.00	0.03	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
V0 _{RA}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
EA _{RV}	0.02	0.00	0.06	0.01	0.16	0.02	0.00	0.01	0.01	0.01	0.01	0.00	0.01	0.02
EB _{RV}	0.01	0.00	0.25	0.04	0.27	0.01	0.00	0.04	0.02	0.04	0.02	0.00	0.05	0.02
V0 _{RV}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
R _{min}	0.00	0.02	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.00
R _{max}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
R _{AR} ^{SYS}	0.10	0.01	0.00	0.03	0.01	0.10	0.31	0.03	0.03	0.03	0.03	0.00	0.03	0.03
C _{AR} ^{SYS}	0.08	0.01	0.00	0.02	0.00	0.08	0.15	0.02	0.01	0.02	0.02	0.00	0.01	0.01
L _{AR} ^{SYS}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
R _{VEN} ^{SYS}	0.15	0.06	0.17	0.11	0.08	0.15	0.01	0.11	0.11	0.13	0.14	0.00	0.09	0.27
C _{VEN} ^{SYS}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
L _{VEN} ^{SYS}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
R _{AR} ^{PUL}	0.00	0.01	0.00	0.03	0.01	0.00	0.00	0.04	0.20	0.12	0.04	0.99	0.06	0.00
C _{AR} ^{PUL}	0.00	0.00	0.01	0.03	0.01	0.00	0.00	0.05	0.10	0.00	0.00	0.00	0.00	0.00
L _{AR} ^{PUL}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
R _{VEN} ^{PUL}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.01	0.00	0.00
C _{VEN} ^{PUL}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
L _{VEN} ^{PUL}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
rVSD	0.57	0.10	0.25	0.71	0.34	0.57	0.22	0.67	0.46	0.64	0.45	0.00	0.66	0.45
HR	0.01	0.31	0.21	0.02	0.11	0.01	0.30	0.03	0.04	0.01	0.14	0.00	0.07	0.16

FIGURE 13. Total Sobol' indices for 0D cardiocirculatory model in the ventricular septal defect case.

of $\pm 25\%$ with respect to $\check{\mathbf{u}}^R$. Based on the sensitivity analysis results (Figure 13), the selected parameters are:

$$\mathbf{u} = [EB_{LV}, EA_{RV}, EB_{RV}, R_{AR}^{SYS}, C_{AR}^{SYS}, R_{VEN}^{SYS}, R_{AR}^{PUL}, rVSD, HR].$$

Since all entries of the parameter vector \mathbf{u} of the 0D cardiocirculatory model are positive, we work with their logarithms to formulate an unconstrained inverse problem. We assume that the parameter vector follows a lognormal distribution $\log U \sim \mathcal{N}(\boldsymbol{\mu}_{pr}, \Gamma_{pr})$. We set the prior mean as $\boldsymbol{\mu}_{pr} = \log \mathbf{u}^R$, where \mathbf{u}^R is the vector of reference values for \mathbf{u} and \log is applied element-wise. We let Γ_{pr} to be a diagonal matrix. During the sensitivity analysis, for $i = 1, \dots, D$, each entry u_i of the parameter vector varies within the range $[0.75u_i^R, 1.25u_i^R]$, so that $\log u_i \in [\log 0.75 + \log u_i^R, \log 1.25 + \log u_i^R]$. We set $(\Gamma_{pr})_{i,i} = (\log 1.25 - \log 0.75)^2/12$.

We generate a dataset $\{(\mathbf{u}^{(m)}, \mathbf{y}^{(m)})\}_{m=1}^M$, where $M \in \mathbb{N}$ and $\mathbf{u}^{(m)}$ is sampled from $\mathcal{N}(\boldsymbol{\mu}_{pr}, \Gamma_{pr})$. Let $\tilde{\mathbf{y}}^{(m)}$ be the noiseless observations of the solution of (28) modified with (29) with parameter $\mathbf{u}^{(m)}$. Then, $\mathbf{y}^{(m)} = \tilde{\mathbf{y}}^{(m)} + \boldsymbol{\varepsilon}^{(m)}$, for $m = 1, \dots, M$, where $\boldsymbol{\varepsilon}^{(m)}$ is sampled from $\mathcal{N}(\boldsymbol{\mu}_E, \Gamma_E)$ with $\boldsymbol{\mu}_E = \mathbf{0}$ and Γ_E diagonal with $(\Gamma_E)_{i,i} = (\eta \max_m |\tilde{y}_i^{(m)}|)^2$, for $i = 1, \dots, O$ and $\eta \in \mathbb{R}$. We fix $\eta = 0.05$.

We use a decoder with 5 hidden layers and 250 neurons per layer trained on a dataset of $M = 4096$ samples. The encoder consists of 3 hidden layers and 250 neurons per layer and is trained on a subset of 100 samples of the decoder dataset.

We generate 2^{17} samples from the distribution $\mathcal{N}(\boldsymbol{\mu}_{pr}, \Gamma_{pr})$ to estimate (24) and (25).

The UQ-VAE mean $\boldsymbol{\mu}_{post}^{(m)}$ estimates the posterior mean $\mathbb{E}_{p_{U|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$ with an average relative error less than 2.0% on 100 test samples (Figure 14). The UQ-VAE covariance $\Gamma_{post}^{(m)}$ estimates the posterior covariance $\text{Cov}_{p_{U|Y}(\mathbf{u}|\mathbf{y}^{(m)})}(\mathbf{u})$ with an average relative error

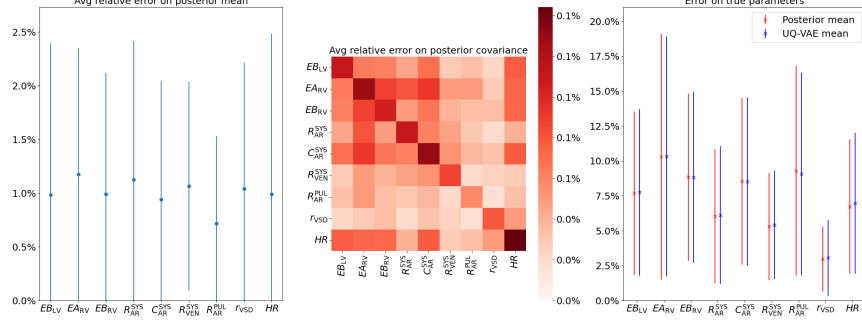


FIGURE 14. From left to right: average and standard deviation of the relative errors of the UQ-VAE mean $\mu_{\text{post}}^{(m)}$ on the posterior one $\mathbb{E}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$; average relative error of the UQ-VAE covariance $\Gamma_{\text{post}}^{(m)}$ on the posterior one $\text{Var}_{\text{p}_{\text{U}|Y}(\mathbf{u}|\mathbf{y}^{(m)})}[\mathbf{u}]$; average and standard deviations of the relative errors of the posterior and UQ-VAE means on the true parameters $\mathbf{u}^{(m)}$. The test are run for 100 samples for the ventricular septal defect case and the relative errors and standard deviations are computed element-wise.

	EB_{LV}	EA_{RV}	EB_{RV}	R_{AR}^{SYS}	C_{AR}^{SYS}	R_{VEN}^{SYS}	R_{AR}^{FUL}	r_{VSD}	HR
\mathbf{u}	0.11	0.70	0.087	0.66	1.1	0.32	0.10	0.79	82
μ_{post}	0.09	0.64	0.074	0.62	1.0	0.32	0.09	0.82	82
Relative error	0.14	0.08	0.15	0.07	0.06	0.01	0.10	0.04	0.01

TABLE 6. True and estimated parameter, along with their relative error, for a test sample of the 0D cardiocirculatory model with VSD.

less than 0.2%. Both the posterior and the UQ-VAE means achieve similar average relative errors on the true parameters $\mathbf{u}^{(m)}$.

We perform forward uncertainty quantification of a single test sample. We use a Sobol' sequence to generate 2^{16} parameter samples from the approximate posterior distribution $q_{\phi}(\mathbf{u}|\mathbf{y}) = \mathcal{N}(\mu_{\text{post}}, \Gamma_{\text{post}})$. For each parameter sample, we run the 0D cardiocirculatory model with the VSD and compute the mean and standard deviation of the resulting time transients pressures and volumes in the four cardiac chambers. We compare the means and standard deviation of the estimated pressure-volume (PV) loops with the true PV loops (Figure 15). PV loops are indicators of cardiac dysfunction and their accurate estimation provides valuable clinical insights into a patient's condition. The estimated PV loops for the right atrium and both ventricles are highly accurate, with smaller uncertainties observed in the left ventricle. The estimation for the left atrium is less accurate compared to the other chambers. This discrepancy may be attributed to the high estimation error in EB_{LV} (Table 6) which affects both the pressure and volume in the left atrium. Another possible explanation is that the true posterior distribution may not be Gaussian, as we assumed in our approximation, potentially introducing a bias in the estimated distribution.




FIGURE 15. Comparison between the true PV loops (blue) and the mean (orange) of those estimated by sampling the parameters from the approximate posterior distribution $q_\phi(\mathbf{u}|\mathbf{y})$. The shaded area represents the standard deviation of the estimated PV loops.

4. CONCLUSIONS

We proposed an improvement (in the form of a novel loss function) of Uncertainty Quantification Variational AutoEncoders (UQ-VAE) for solving Bayesian inverse problems. We introduced a new loss function that significantly reduces the computational cost of a single evaluation, compared to the loss function proposed in Ref. 42, while preserving a relationship between its minimum $(\hat{\boldsymbol{\mu}}, \hat{\Gamma})$ and $(\mathbf{u}_{\text{MAP}}, \Gamma_{\text{Lap}})$ when the forward map \mathcal{F} in (1) is affine (Theorem 1 and Theorem 2) and extending the result to the non affine case.

We showed the capabilities of the UQ-VAEs to approximate both the mean and covariance of the posterior distribution in the cases of the Poisson problem (section 3.1), a non linear problem (section 3.2) and a 0D cardiocirculatory model for systemic hypertension (section 3.3.1) and for ventricular septal defect (section 3.3.2).

With few observational data, the UQ-VAE estimates of the true parameters are as accurate as the posterior mean. With a large number of observational data, the true parameters were estimated better using the UQ-VAE mean rather than the posterior one, due to the latter being poorly approximated by the sample mean in the case of low regularity of the posterior distribution. A similar behavior was observed when the noise level η decreased from high to low. These results show that the UQ-VAE approach is effective in estimating model parameters by solving inverse problems, even for low regular posterior distributions.

The main advantage of a trained UQ-VAEs is in performing inverse uncertainty quantification on the estimated parameters in a short computational time. UQ-VAEs estimated

posterior covariance matrices, with average relative errors less than 1.6% for the 0D cardio-circulatory model and a test set of 100 samples. The average relative error on the posterior variance in the case of the Poisson problem for a test set of 100 samples was less than 11%. For this problem, the UQ-VAE estimated about 42,000 coefficients ($(D + D(D+1))/2$, with $D = 17^2$) for the mean and the covariance matrix, maintaining a good accuracy.

The proposed novel loss function significantly reduced the training time compared to the previous formulation, while also improving accuracy in estimating $(\mathbf{u}_{\text{MAP}}, \Gamma_{\text{Lap}})$. This reduction in computational cost enables the solution of high dimensional Bayesian inverse problems that were previously computationally infeasible.

A limitation of UQ-VAEs arises when the map \mathcal{F} is approximated by a NN ψ . In this case, the decoder is trained on noiseless observational data and the approximation error is also computed using such data, which is only feasible within an in silico framework. In contrast, the encoder is trained using noisy observational data. When the map \mathcal{F} is known, noiseless observational data are no longer required: an advantage in realistic scenarios where observational data are inherently noisy.

Moreover, the decoder is trained using a supervised approach, with paired input and output data, whereas the encoder is trained in a semi-supervised manner, requiring only observational data and prior knowledge about the distribution of the parameters. In real case scenarios, the observational data of a phenomenon are directly measured, while the prior knowledge is represented by a mathematical model describing the phenomenon. The prior mean can be fixed to a reference setting of the model parameters and the prior covariance can be defined by a variation of the prior mean, as we have done in section 3. Therefore, a semi-supervised approach for training the decoder, which requires only observational data, would broaden the applicability of UQ-VAEs, as it would eliminate the need for in silico generated parameters and observational data.

Finally, we approximate the posterior distribution using a Gaussian distribution, which is not correct for general maps \mathcal{F} from parameters to observational data. This limitation can be addressed by using hierarchical VAEs, which enable sampling from more expressive approximate posteriors.^{19,35,44}

ACKNOWLEDGMENTS

AT, FR, LD and AQ are members of the INdAM group GNCS “Gruppo Nazionale per il Calcolo Scientifico” (National Group for Scientific Computing).

AT, FR and LD acknowledge the INdAM GNCS project CUP E53CE53C24001950001.

FR has received support from the project FIS, MUR, Italy 2025-2028, Project code: FIS-2023-02228, CUP: D53C24005440001, “SYNERGIZE: Synergizing Numerical Methods and Machine Learning for a new generation of computational models”.

LD acknowledges the support by the FAIR (“Future Artificial Intelligence Research”) project, funded by the NextGenerationEU program within the PNRR-PE-AI scheme (M4C2, investment 1.3, line on Artificial Intelligence), Italy.

LD acknowledges the project PRIN2022, MUR, Italy, 2023–2025, 202232A8AN “Computational modeling of the heart: from efficient numerical solvers to cardiac digital twins”.

The present research is part of the activities of “Dipartimento di Eccellenza 2023–2027”, MUR, Italy, Dipartimento di Matematica, Politecnico di Milano.

COMPETING INTERESTS

The authors declare no competing interests.

REFERENCES

- [1] A. Albanese, L. Cheng, M. Ursino and N. W. Chbat, An integrated mathematical model of the human cardiopulmonary system: model development, *American Journal of Physiology-Heart and Circulatory Physiology* **310** (2016) H899–H921.
- [2] T. Bayes, An essay towards solving a problem in the doctrine of chances, *Philosophical transactions* **53** (1763) 370–418.
- [3] T. Bui-Thanh, O. Ghattas, J. Martin and G. Stadler, A computational framework for infinite-dimensional bayesian inverse problems part i: The linearized case, with application to global seismic inversion, *SIAM Journal on Scientific Computing* **35** (2013) A2494–A2523.
- [4] R. H. Byrd, P. Lu, J. Nocedal and C. Zhu, A limited memory algorithm for bound constrained optimization, *SIAM Journal on scientific computing* **16** (1995) 1190–1208.
- [5] F. Caforio, F. Regazzoni, S. Pagani, E. Karabelas, C. Augustin, G. Haase, G. Plank and A. Quarteroni, Physics-informed neural network estimation of material properties in soft tissue nonlinear biomechanical models, *Computational Mechanics* **75** (2024) 487–513.
- [6] E. Chong, *An Introduction to Optimization* (Wiley, Hoboken, 2013).
- [7] W. Dakkak, M. H. Alahmadi and T. I. Oliver, Ventricular septal defect, StatPearls Publishing, October 2024.
- [8] Y. Daon and G. Stadler, Mitigating the influence of the boundary on pde-based covariance operators., *Inverse Problems & Imaging* **12**.
- [9] J. Deasy, N. Simidjievski and P. Lió, Constraining variational inference with geometric jensen-shannon divergence, Advances in Neural Information Processing Systems, 2020, https://proceedings.neurips.cc/paper_files/paper/2020/file/78719f11fa2df9917de3110133506521-Paper.pdf.
- [10] L. Dedè, F. Regazzoni, C. Vergara, P. Zunino, M. Guglielmo, R. Scrofani, L. Fusini, C. Cogliati, G. Pontone and A. Quarteroni, Modeling the cardiac response to hemodynamic changes associated with COVID-19: a computational study, *Mathematical Biosciences and Engineering* **18** (2021) 3364–3383.
- [11] M. Evans and T. Swartz, *Approximating integrals via Monte Carlo and deterministic methods* (OUP Oxford, Oxford, 2000).
- [12] A. Formaggio, M. De Luca, S. Borrelli, G. Putame, N. De Vita, F. Minelli, F. Della Corte, R. Vaschetto, A. L. Audenino, C. Olivieri et al., A multidomain 0d model for continuous positive airway pressure ventilation circuit design: Validation and applications, *Biocybernetics and Biomedical Engineering* **45** (2025) 170–180.
- [13] X. Glorot and Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, 2010.
- [14] H. Goh, S. Sherifdeen, J. Wittmer and T. Bui-Thanh, Solving bayesian inverse problems via variational autoencoders, Proceeding of Machine Learning Research, 2nd Annual Conference on Mathematical and Scientific Machine Learning, September 2021.
- [15] Y.-C. Hung, A review of monte carlo and quasi-monte carlo sampling techniques, *Wiley Interdisciplinary Reviews: Computational Statistics* **16** (2024) e1637.
- [16] J. W. Hurst, R. C. Schlant, C. E. Rackley, E. H. Sonnenblick and N. K. Wenger, *The heart, arteries e veins* (McGraw-Hill Inc., New York, 1990).
- [17] K. H. Jin, M. T. McCann, E. Froustey and M. Unser, Deep convolutional neural network for inverse problems in imaging, *IEEE transactions on image processing* **26** (2017) 4509–4522.
- [18] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, January 2014.
- [19] 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, 2016, https://proceedings.neurips.cc/paper_files/paper/2016/file/ddeebdeefdb7e7e7a697e1c3e3d8ef54-Paper.pdf.
- [20] R. M. Lang, L. P. Badano, V. Mor-Avi, J. Afilalo, A. Armstrong, L. Ernande, F. A. Flachskampf, E. Foster, S. A. Goldstein, T. Kuznetsova et al., Recommendations for cardiac chamber quantification by echocardiography in adults: an update from the american society of echocardiography and the european association of cardiovascular imaging, *European Heart Journal-Cardiovascular Imaging* **16** (2015) 233–271.
- [21] H. Li, J. Schwab, S. Antholzer and M. Haltmeier, Nett: Solving inverse problems with deep neural networks, *Inverse Problems* **36**.
- [22] A. Lucas, M. Iliadis, R. Molina and A. K. Katsaggelos, Using deep neural networks for inverse problems in imaging: beyond analytical methods, *IEEE Signal Processing Magazine* **35** (2018) 20–36.
- [23] J. A. Nelder and R. Mead, A simplex method for function minimization, *The computer journal* **7** (1965) 308–313.

- [24] F. D. M. Neto and A. J. da Silva Neto, *An introduction to inverse problems with applications* (Springer Science & Business Media, London, 2012).
- [25] H. V. Nguyen and T. Bui-Thanh, Tnet: A model-constrained tikhonov network approach for inverse problems, *SIAM Journal on Scientific Computing* **46** (2024) C77–C100.
- [26] F. Nielsen, A family of statistical symmetric divergences based on jensen’s inequality, December 2010.
- [27] W. H. Organization, *Global report on hypertension: the race against a silent killer* (World Health Organization, Geneva, 2023).
- [28] F. Regazzoni, M. Salvador, P. C. Africa, M. Fedele, L. Dedè and A. Quarteroni, A cardiac electromechanical model coupled with a lumped-parameter model for closed-loop blood circulation, *Journal of Computational Physics* **457** (2022) 111083.
- [29] L. G. Rudski, W. W. Lai, J. Afilalo, L. Hua, M. D. Handschumacher, K. Chandrasekaran, S. D. Solomon, E. K. Louie and N. B. Schiller, Guidelines for the echocardiographic assessment of the right heart in adults: a report from the american society of echocardiography: endorsed by the european association of echocardiography, a registered branch of the european society of cardiology, and the canadian society of echocardiography, *Journal of the American society of echocardiography* **23** (2010) 685–713.
- [30] A. Saltelli, P. Annoni, I. Azzini, F. Campolongo, M. Ratto and S. Tarantola, Variance based sensitivity analysis of model output. design and estimator for the total sensitivity index, *Computer physics communications* **181** (2010) 259–270.
- [31] M. Salvador, F. Kong, M. Peirlinck, D. W. Parker, H. Chubb, A. M. Dubin and A. L. Marsden, Digital twinning of cardiac electrophysiology for congenital heart disease, *Journal of the Royal Society Interface* **21**.
- [32] M. Salvador, M. Stocchi, F. Regazzoni, C. M. Augustin, L. Dede’, S. A. Niederer and A. Quarteroni, Whole-heart electromechanical simulations using latent neural ordinary differential equations, *NPJ Digital Medicine* **7**.
- [33] I. M. Sobol, On the distribution of points in a cube and the approximate evaluation of integrals, *USSR Computational mathematics and mathematical physics* **7** (1967) 86–112.
- [34] I. M. Sobol, Sensitivity estimates for nonlinear mathematical models, *Math. Model. Comput. Exp.* **1** (1993) 407–414.
- [35] C. K. Sønderby, T. Raiko, L. Maaløe, S. K. Sønderby and O. Winther, Ladder variational autoencoders, *Advances in neural information processing systems*, 2016, https://proceedings.neurips.cc/paper_files/paper/2016/file/6ae07dcb33ec3b7c814df797cbda0f87-Paper.pdf.
- [36] D. E. Spicer, H. H. Hsu, J. Co-Vu, R. H. Anderson and F. J. Fricker, Ventricular septal defect, *Orphanet journal of rare diseases* **9** (2014) 1–16.
- [37] A. M. Stuart, Inverse problems: a bayesian perspective, *Acta numerica* **19** (2010) 451–559.
- [38] T. Sutter, I. Daunhawer and J. Vogt, Multimodal generative learning utilizing jensen-shannon-divergence, *Advances in Neural Information Processing Systems*, 2020, https://proceedings.neurips.cc/paper_files/paper/2020/file/43bb733c1b62a5e374c63cb22fa457b4-Paper.pdf.
- [39] A. N. Tikhonov and V. Y. Arsenin, *Solutions of Ill-Posed Problems* (VH Winston and Sons, Washington DC, 1977).
- [40] G. G. Tong, C. A. S. Long and D. E. Schiavazzi, Invaert networks: A data-driven framework for model synthesis and identifiability analysis, *Computer Methods in Applied Mechanics and Engineering* **423** (2024) 116846.
- [41] G. G. Tong, C. A. Sing-Long and D. E. Schiavazzi, Invaert networks for amortized inference and identifiability analysis of lumped-parameter haemodynamic models, *Philosophical Transactions A* **383** (2025) 20240215.
- [42] A. Tonini and L. Dede’, Enhanced uncertainty quantification variational autoencoders for the solution of bayesian inverse problems, February 2025.
- [43] A. Tonini, F. Regazzoni, M. Salvador, L. Dede’, R. Scrofani, L. Fusini, C. Cogliati, G. Pontone, C. Vergara and A. Quarteroni, Two new calibration techniques of lumped-parameter mathematical models for the cardiovascular system, *International Journal for Numerical Methods in Engineering* **126** (2025) e7648.
- [44] A. Vahdat and J. Kautz, Nvae: A deep hierarchical variational autoencoder, *Advances in neural information processing systems*, 2020, <https://proceedings.neurips.cc/paper/2020/file/e3b21256183cf7c2c7a66be163579d37-Paper.pdf>.
- [45] H. Van Nguyen, T. Bui-Thanh and C. Dawson, Taen: a model-constrained tikhonov autoencoder network for forward and inverse problems, *Computer Methods in Applied Mechanics and Engineering* **446** (2025) 118245.

- [46] U. Villa, N. Petra and O. Ghattas, hippylib: an extensible software framework for large-scale deterministic and bayesian inverse problems, 2016, version 2.1.0.
- [47] U. Villa, N. Petra and O. Ghattas, hippylib: an extensible software framework for large-scale deterministic and bayesian inverse problems, *Journal of Open Source Software* **3** (2018) 940.
- [48] U. Villa, N. Petra and O. Ghattas, Hippylib: An extensible software framework for large-scale inverse problems governed by pdes: Part i: Deterministic inversion and linearized bayesian inference, *ACM Trans. Math. Softw.* **47** (2021) 1–34.
- [49] G. Q. Zhang and W. Zhang, Heart rate, lifespan, and mortality risk, *Ageing research reviews* **8** (2009) 52–60.
- [50] Y. Zhu, J. Zhang, A. Gangrade and C. Scott, Label noise: Ignorance is bliss, *Advances in Neural Information Processing Systems*, 2024, https://proceedings.neurips.cc/paper_files/paper/2024/file/d3696c79d572c995a74eac78037551a8-Paper-Conference.pdf.

MOX Technical Reports, last issues

Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

Tonini, A.; Bui-Thanh, T.; Regazzoni, F.; Dede', L.; Quarteroni, A.

Improvements on uncertainty quantification with variational autoencoders

Gimenez Zapiola, A.; Boselli, A.; Menafoglio, A.; Vantini, S.

55/2025 Gimenez Zapiola, A.; Boselli, A.; Menafoglio, A.; Vantini, S.

Hyper-spectral Unmixing algorithms for remote compositional surface mapping: a review of the state of the art

54/2025 Tomasetto, M.; Williams, J.P.; Braghin, F.; Manzoni, A.; Kutz, J.N.

Reduced order modeling with shallow recurrent decoder networks

53/2025 Zecchi, A. A.; Sanavio, C.; Perotto, S.; Succi, S.

Telescopic quantum simulation of the advection-diffusion-reaction dynamics

51/2025 Tomasetto, M.; Braghin, F.; Manzoni, A.

Latent feedback control of distributed systems in multiple scenarios through deep learning-based reduced order models

50/2025 Bonetti, S.; Botti, M.; Antonietti, P.F.

Conforming and discontinuous discretizations of non-isothermal Darcy–Forchheimer flows

49/2025 Zanin, A.; Pagani, S.; Corti, M.; Crepaldi, V.; Di Fede, G.; Antonietti, P.F.; the ADNI

Predicting Alzheimer's Disease Progression from Sparse Multimodal Data by NeuralODE Models

48/2025 Temellini, E.; Ballarin, F.; Chacon Rebollo, T.; Perotto, S.

On the inf-sup condition for Hierarchical Model reduction of the Stokes problem

47/2025 Gimenez Zapiola, A.; Consolo, A.; Amaldi, E.; Vantini, S.

Penalised Optimal Soft Trees for Functional Data