

MOX-Report No. 21/2026

**Neural Markov chain Monte Carlo: Bayesian inversion via
normalizing flows and variational autoencoders**

Bottacini, G.; Torzoni, M.; Manzoni, A.

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

mox-dmat@polimi.it

<https://mox.polimi.it>

Neural Markov chain Monte Carlo: Bayesian inversion via normalizing flows and variational autoencoders

Giacomo Bottacini^a, Matteo Torzoni^b, Andrea Manzoni^a

^aMOX – Department of Mathematics, Politecnico di Milano, Piazza Leonardo da Vinci 32, I-20133, Milano, Italy

^bDepartment of Civil and Environmental Engineering, Politecnico di Milano,
Piazza Leonardo da Vinci 32, I-20133, Milano, Italy

Abstract

This paper introduces a Bayesian framework that combines Markov chain Monte Carlo (MCMC) sampling, dimensionality reduction, and neural density estimation to efficiently handle inverse problems that (i) must be solved multiple times, and (ii) are characterized by intractable or unavailable likelihood functions. The posterior probability distribution over quantities of interest is estimated via differential evolution Metropolis sampling, empowered by learnable mappings. First, a variational autoencoder performs probabilistic feature extraction from observational data. The resulting latent structure inherently quantifies uncertainty, capturing deviations between the actual data-generating process and the training data distribution. At each step of the MCMC random walk, the algorithm jointly samples from the data-informed latent distribution and the space of parameters to be inferred. These samples are fed into a neural likelihood estimator based on normalizing flows, specifically real-valued non-volume preserving transformations. The scaling and translation functions of the affine coupling layers are modeled by neural networks conditioned on the unknown parameters, allowing the representation of arbitrary observation likelihoods. The proposed methodology is validated on two case studies: (i) structural health monitoring of a railway bridge for damage detection, localization, and quantification, and (ii) estimation of the conductivity field in a steady-state Darcy’s groundwater flow problem. The results demonstrate the efficiency of the inference strategy, while ensuring that model-reality mismatches do not yield overconfident, yet inaccurate, estimates.

Keywords: Bayesian inference, Simulation-based inference, Uncertainty quantification, Normalizing flows, Variational autoencoders

1. Introduction

When analyzing physical systems, key parameters of interest are often difficult or impossible to measure directly. Instead, a forward model, derived from physical laws or empirical relationships, is typically available to describe how hidden parameters and internal states generate observables. This setup allows the estimation of unknown parameters from measured data to be cast as an inverse problem, which is often ill-posed and computationally challenging due to non-uniqueness and sensitivity to noise.

A principled approach to address this class of problems is provided by Bayesian inference [1, 2], which treats unknown parameters as random variables and frames the inverse problem in a probabilistic setting. This formulation inherently enables uncertainty quantification (UQ) and provides a natural mechanism to regularize ill-posedness. In essence, Bayesian inference updates prior beliefs by conditioning them on observational data, yielding a posterior distribution over the parameters of interest [3, 4]. A critical component of this process is the likelihood function, which encapsulates the assumed statistical relationship between model predictions and observations.

However, scientific and engineering applications are increasingly characterized by high-dimensional and highly complex observational data, as seen in fields such as genomics [5], imaging [6, 7], video analysis [8], and natural language processing [9]. High-dimensional data pose several challenges, including overfitting and reduced generalization, which can compromise the performance and robustness of downstream inference models. Additionally, when dealing with high-dimensional observables [10] or a large number of parameters to be inferred (e.g., spatially distributed material properties) [11], the parameter-to-observation map may

exhibit low sensitivity, meaning that significant variations in parameters produce only subtle changes in the outputs. This results in ill-posedness, poorly identifiable solutions, and elevated epistemic uncertainty. Moreover, accurately characterizing the posterior distribution often requires a large number of forward evaluations, which can be computationally prohibitive when each evaluation involves expensive simulations (e.g., finite element analyses). Finally, in many real-world applications, the likelihood function may be intractable, unavailable in closed form, or too costly to evaluate directly.

These challenges have motivated the integration of recent advances in machine learning, particularly deep learning (DL), into Bayesian inference pipelines to enhance computational scalability [12–14]. A prominent direction in this context involves dimensionality reduction [15], through the projection of data into a lower-dimensional space. This approach not only reduces computational costs but can also help uncover the intrinsic structure of the underlying generative process. A wealth of dimensionality reduction techniques have been proposed in the literature [16], and they can generally be categorized into two classes: (i) feature selection and (ii) feature extraction. The former retains a subset of the original variables, whereas the latter maps the data into a new space through learned or predefined transformations. Recent research has increasingly focused on learnable feature extraction [17], aiming to automatically uncover hierarchical latent patterns without relying on manual variable selection or domain-specific assumptions.

Traditional feature extraction methods rely on projection techniques that map high-dimensional data onto lower-dimensional subspaces while seeking to minimize information loss. Among the most widely used approaches are principal component analysis (PCA) and linear discriminant analysis (LDA) [18]. PCA identifies orthogonal directions, or principal components, that capture the maximum variance in the data, providing an unsupervised linear embedding. In contrast, LDA is a supervised method that seeks to optimize class separability by maximizing inter-class variance while minimizing intra-class variance within the projection subspace. Despite their interpretability and computational efficiency, both PCA and LDA are inherently restricted to linear transformations, which limits their ability to model complex, nonlinear relationships.

In this context, autoencoders (AEs) provide a nonlinear DL-based alternative for compressing high-dimensional observations into compact latent representations that capture the most informative features of the data. Extensions of the standard AE structure include sparse AEs, denoising AEs, and variational autoencoders (VAEs) [19, 20]. Variational autoencoders, in particular, have demonstrated superior capabilities in quantifying uncertainty in learned representations – a crucial requirement in applications involving decision-making under uncertainty. Specifically, they impose a probabilistic encoding scheme in which the latent representation not only supports input reconstruction but is also constrained to follow a predefined prior distribution. This probabilistic regularization promotes a smooth, continuous, and semantically meaningful latent space, thereby enabling the generation of plausible samples via latent space interpolation or stochastic sampling.

From a complementary perspective, the emergence of simulation-based inference [21–23] offers a framework for performing Bayesian inference when the likelihood function is intractable or unavailable. These methods bypass explicit likelihood evaluation by leveraging generative or surrogate models, such as neural density estimators [24], to learn the statistical relationship between parameters and data directly from observations.

Several surrogate-based strategies have been proposed to accelerate Bayesian inverse uncertainty quantification. Early works, see [25], have employed polynomial chaos expansions to approximate forward models or likelihood functions, enabling efficient Markov chain Monte Carlo (MCMC) sampling in PDE-constrained inverse problems. Other approaches reduce or bypass sampling by constructing global surrogates of the inverse map or the posterior distribution [26]. More recently, surrogate models have been integrated with advanced MCMC schemes to achieve substantial computational savings while preserving Bayesian consistency [27, 28]. Despite these advances, handling high-dimensional observations while preserving sampling-based uncertainty quantification remains a major challenge.

In this work, we propose *neural* MCMC sampling, a likelihood-free approach [29, 30] that integrates two DL components to enable efficient inference of parameters of interest. The first component is a VAE, trained to extract informative features from high-dimensional observations while explicitly quantifying uncertainty. At each MCMC iteration, a new parameter proposal is generated and paired with a latent code sampled from the VAE encoding. This pair is then passed to a second generative model for density estimation. This second component is a conditional normalizing flow architecture [31], trained to approximate the likelihood function using real-valued non-volume preserving (RealNVP) transformations [32]. The output of this surrogate

likelihood model enables direct computation of the MCMC acceptance ratio, thereby efficiently guiding the sampling process. This hybrid framework enables scalable posterior inference in high-dimensional and computationally intensive inverse problems, while also providing uncertainty quantification through the learned variational structure.

The paper is organized as follows. Section 2 describes the proposed computational methodology. Section 3 presents numerical results for two case studies: (i) a structural health monitoring scenario involving the detection, localization, and quantification of damage in a railway bridge; and (ii) a steady-state groundwater flow problem governed by Darcy’s law, with a focus on estimating a spatially varying hydraulic conductivity field. Finally, Section 4 draws some conclusions.

2. Data assimilation via normalizing flows and variational autoencoders

This section details the proposed simulation-based inference framework. The approach is built upon two components tightly integrated within an MCMC algorithm: (i) a modified VAE for probabilistic feature extraction from observational data, and (ii) a neural likelihood estimator based on a conditional normalizing flow with coupling layers parameterized by neural networks conditioned on the parameters to be inferred. The VAE enables low-dimensional and noise-robust representation of high-dimensional observations, while the conditional normalizing flow provides a flexible likelihood approximation in the latent space, allowing for accurate UQ within a Bayesian setting.

The resulting MCMC scheme jointly samples from the data-informed latent distribution and the parameter space to infer the posterior of target parameters. In the following, section 2.1 outlines the overall framework and explains how the two neural components are embedded within the inference procedure. Sections 2.2 and 2.3 then describe the architectures and training procedures of the RealNVP and the VAE, respectively.

2.1. Simulation-based inference scheme

We formulate the inverse problem in a probabilistic setting. Adopting a Bayesian perspective, we treat N_{par} target unknown parameters $\boldsymbol{\lambda} \in \mathbb{R}^{N_{\text{par}}}$ as random variables to be inferred from data. The prior distribution $p(\boldsymbol{\lambda})$, representing the initial belief about the parameters before assimilating any data, is updated through the likelihood function $p(\mathbf{h} \mid \boldsymbol{\lambda})$ upon observing data $\mathbf{h} \in \mathbb{R}^{N_h}$. By Bayes’ theorem, the posterior distribution is given by:

$$p(\boldsymbol{\lambda} \mid \mathbf{h}) = \frac{p(\mathbf{h} \mid \boldsymbol{\lambda})p(\boldsymbol{\lambda})}{\int p(\mathbf{h} \mid \boldsymbol{\lambda})p(\boldsymbol{\lambda}) d\boldsymbol{\lambda}}. \quad (1)$$

In many applications, evaluating the posterior $p(\boldsymbol{\lambda} \mid \mathbf{h})$ in closed form is infeasible due to both the denominator $p(\mathbf{h}) = \int p(\mathbf{h} \mid \boldsymbol{\lambda})p(\boldsymbol{\lambda}) d\boldsymbol{\lambda}$, which involves an integral over the entire parameter space that may be high-dimensional, and the need for a complex forward model to compute $p(\mathbf{h} \mid \boldsymbol{\lambda})$, making the likelihood function available only numerically. As a result, numerical sampling methods are typically employed.

Monte Carlo methods provide a practical approach for approximating the posterior distribution by drawing samples from $p(\boldsymbol{\lambda} \mid \mathbf{h})$. Several algorithms exist for sampling the posterior, with MCMC methods being among the most widely used [33–35]. A notable example is the Metropolis-Hastings algorithm [36], which constructs a Markov chain by iteratively proposing candidate states $\boldsymbol{\lambda}'$ from a proposal distribution. To obtain the n -th state of the chain, with $n = 1, \dots, N_s$ and N_s denoting the chain length, proposals are drawn from $q_{\text{MH}}(\boldsymbol{\lambda}' \mid \boldsymbol{\lambda}^{(n-1)})$. This proposal is typically chosen as a Gaussian density centered at the current state $\boldsymbol{\lambda}^{(n-1)}$ with covariance matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{N_{\text{par}} \times N_{\text{par}}}$, i.e. $\boldsymbol{\lambda}' \sim \mathcal{N}(\boldsymbol{\lambda}^{(n-1)}, \boldsymbol{\Sigma})$, thereby generating local, symmetric perturbations around the current state.

While Gaussian proposals are widely applicable, they often struggle in high-dimensional or strongly correlated parameter spaces. For this reason, we adopt a more robust and adaptive strategy based on differential evolution MCMC (DE-MCMC) [37]. This family of sampling algorithms employs proposal mechanisms that combine information from past states of the chain to capture parameter correlations and promote global exploration of the parameter space. The standard DE-MCMC proposal takes the following form:

$$\boldsymbol{\lambda}' = \boldsymbol{\lambda}^{(n-1)} + \gamma(\boldsymbol{\Lambda}_{R1} - \boldsymbol{\Lambda}_{R2}) + \boldsymbol{\epsilon}, \quad (2)$$

where $\mathbf{\Lambda}_{R1}, \mathbf{\Lambda}_{R2} \in \mathbb{R}^{N_{\text{par}}}$ are two distinct samples drawn from an archive (or population) of past states $\mathbf{\Lambda}$. This archive collects previously accepted states at the time of generating $\boldsymbol{\lambda}^{(n)}$ and is systematically updated with new $\boldsymbol{\lambda}$ states at intervals of $k_T \in \mathbb{N}$ iterations. The parameter $\gamma \in \mathbb{R}^+$ is a non-negative scaling factor that controls the jump size, and $\boldsymbol{\epsilon} \in \mathbb{R}^{N_{\text{par}}}$ is a small random perturbation introduced to ensure ergodicity. In practice, γ is chosen to scale inversely with the square root of the parameter dimension, with $\gamma = \frac{2.38}{\sqrt{2N_{\text{par}}}}$ being a recommended value that yields near-optimal acceptance rates for multivariate distributions [37].

In this work, we employ the DEMetropolisZ sampler from the PyMC library [38], an adaptive DE-MCMC algorithm that leverages an evolving archive of past states while preserving the standard Metropolis–Hastings acceptance mechanism. The probability of accepting $\boldsymbol{\lambda}'$ as the next state $\boldsymbol{\lambda}^{(n)}$ of the Markov chain targeting $p(\boldsymbol{\lambda} \mid \mathbf{h}) \propto p(\mathbf{h} \mid \boldsymbol{\lambda})p(\boldsymbol{\lambda})$ is therefore defined as:

$$\alpha(\boldsymbol{\lambda}^{(n-1)}, \boldsymbol{\lambda}') = \min \left\{ 1, \frac{p(\mathbf{h} \mid \boldsymbol{\lambda}')p(\boldsymbol{\lambda}')}{p(\mathbf{h} \mid \boldsymbol{\lambda}^{(n-1)})p(\boldsymbol{\lambda}^{(n-1)})} \right\}, \quad (3)$$

which corresponds to the Metropolis–Hastings acceptance probability in the case of a symmetric proposal distribution. Since the stochastic proposal rule [2] is symmetric by construction, this formulation guarantees that the resulting Markov chain satisfies the detailed balance condition, ensuring that its stationary distribution coincides with the target posterior [34].

Evaluating $\alpha(\boldsymbol{\lambda}^{(n-1)}, \boldsymbol{\lambda}')$ only requires computing the likelihood $p(\mathbf{h} \mid \boldsymbol{\lambda}')$ and the prior $p(\boldsymbol{\lambda}')$. However, for many modern problems, directly evaluating the likelihood $p(\mathbf{h} \mid \boldsymbol{\lambda}')$ is computationally prohibitive or even infeasible. Likelihood-free Bayesian methods [39] provide a viable strategy for addressing this limitation. Approximate Bayesian computation (ABC) is a prominent example of a likelihood-free method that bypasses explicit likelihood evaluations by accepting candidate parameter values when the simulated data are sufficiently *close* to the observed data [40]. While conceptually straightforward, ABC can be highly inefficient, often requiring an extremely large number of simulations to obtain a reasonable approximation of the posterior distribution.

To overcome the limitations of common likelihood-free methods such as ABC, we propose a simulation-based inference framework [21] that integrates DE-MCMC sampling with neural density estimation [22]. Specifically, we employ a conditional normalizing flow (CNF) [31] as a surrogate likelihood model to approximate $p(\mathbf{h} \mid \boldsymbol{\lambda})$. Once trained, the CNF provides tractable and differentiable likelihood evaluations that can be directly incorporated into the acceptance ratio [3]. This scheme is further enhanced by a VAE that first compresses the raw observations into a low-dimensional, informative latent representation. This dimensionality reduction step enables the CNF to operate on a substantially lower-dimensional input, improving the overall stability and computational efficiency. By jointly modeling (and sampling from) both the latent VAE representation and the parameter space, the proposed inference scheme allows the CNF to instantaneously estimate observation likelihoods while preserving the reliability of the inference process.

We refer to our architecture as *neural MCMC*, composed of three main components (see Fig. 1): (i) an adaptive MCMC scheme that updates its proposal mechanism based on past samples, thereby improving mixing and sampling efficiency; (ii) a CNF-based estimator, trained to provide fast likelihood approximations during sampling, as detailed in Sec. 2.2 and (iii) a VAE for probabilistic feature extraction from the observed data, as described in Sec. 2.3.

Given raw measurements $\mathbf{x} \in \mathbb{R}^{N_x}$ from an N_x -dimensional input space, with $N_h \ll N_x$, and a prior distribution $p(\boldsymbol{\lambda})$, the algorithm approximates the posterior $p(\boldsymbol{\lambda} \mid \mathbf{x})$. Raw data \mathbf{x} are first encoded into an N_h -dimensional latent distribution using the VAE encoder $q_\theta : \mathbb{R}^{N_x} \mapsto \mathbb{R}^{N_h}$, as:

$$q_\theta(\mathbf{h} \mid \mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_\theta(\mathbf{x}), \boldsymbol{\sigma}_\theta^2(\mathbf{x})), \quad (4)$$

During sampling, parameter proposals are generated through the DE-MCMC mechanism [2]. When generating the n -th state of the chain $\boldsymbol{\lambda}^{(n)}$, a latent sample $\mathbf{h}' \sim q_\theta(\mathbf{h} \mid \mathbf{x})$ is drawn and used to evaluate $\alpha(\boldsymbol{\lambda}^{(n-1)}, \boldsymbol{\lambda}')$ for the proposed parameter $\boldsymbol{\lambda}'$. The likelihood density $p(\mathbf{h}' \mid \boldsymbol{\lambda}')$ is approximated by the surrogate density $p_\xi(\mathbf{h}' \mid \boldsymbol{\lambda}')$ modeled via the CNF, as indicated by subscript ξ and detailed in the next section. The acceptance ratio in Eq. [3] is therefore rewritten as:

$$\alpha = \min \left(1, \frac{p_\xi(\mathbf{h}' \mid \boldsymbol{\lambda}') p(\boldsymbol{\lambda}')}{p_\xi(\mathbf{h}' \mid \boldsymbol{\lambda}^{(n-1)}) p(\boldsymbol{\lambda}^{(n-1)})} \right), \quad (5)$$

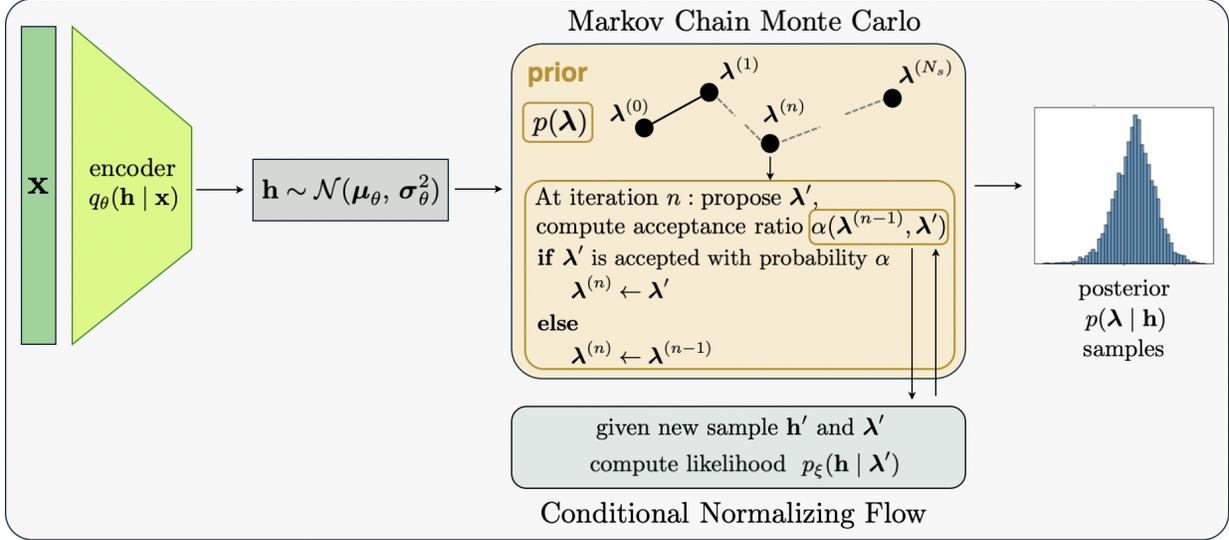


Figure 1: Schematic of the *neural* Markov chain Monte Carlo (MCMC) sampling scheme for simulation-based inference of the parameter posterior $\mathbf{p}(\boldsymbol{\lambda} \mid \mathbf{h})$. The encoder branch of a variational autoencoder (VAE) performs probabilistic feature extraction from the observed data \mathbf{x} , yielding the latent distribution $\mathbf{h} \sim \mathcal{N}(\boldsymbol{\mu}_\theta, \boldsymbol{\sigma}_\theta^2)$. Joint samples $(\mathbf{h}', \boldsymbol{\lambda}')$ are drawn from the latent distribution and the parameter space, respectively. A conditional normalizing flow network then estimates the conditional likelihood density $p(\mathbf{h}' \mid \boldsymbol{\lambda}')$, which is used to efficiently compute the MCMC acceptance ratio $\alpha(\boldsymbol{\lambda}^{(n-1)}, \boldsymbol{\lambda}')$.

which retains the Metropolis-Hastings form under a symmetric proposal. Accepted samples are used to update the Markov chain and, at regular intervals of k_T iterations, to enrich the past-states archive $\boldsymbol{\Lambda}$. After discarding the burn-in samples and thinning the chain, the resulting sequence of samples provides an approximation to the posterior $p(\boldsymbol{\lambda} \mid \mathbf{x})$. A complete algorithmic description of the inference procedure is provided in Algorithm 1.

2.2. Likelihood surrogate modeling using normalizing flows

In our framework, the likelihood approximation $p(\mathbf{h}' \mid \boldsymbol{\lambda}') \approx p_\xi(\mathbf{h}' \mid \boldsymbol{\lambda}')$ exploits a neural surrogate based on a class of invertible generative models known as *normalizing flows*. Normalizing flows are capable of yielding highly expressive yet tractable probability distributions. In essence, they apply a sequence of invertible and differentiable transformations to map a simple base distribution, usually a standard multivariate Gaussian, into a more complex target distribution [41]. This construction supports both efficient sampling and effective density evaluation, and has demonstrated remarkable success across a wide range of domains, from image and audio synthesis [42, 43] to computational statistics and simulation-based Bayesian inference [24].

Formally, let $p_{\mathbf{H}}(\mathbf{h})$ denote a (potentially complex) probability density distribution for \mathbf{h} , as indicated by the subscript \mathbf{H} . In our setting, this density corresponds to the variational approximation $q_\theta(\mathbf{h} \mid \mathbf{x})$ produced by the VAE encoder. Let $f: \mathbb{R}^{N_h} \rightarrow \mathbb{R}^{N_z}$, with $N_h = N_z$, be an invertible transformation that maps an \mathbf{h} sample to a representation $\mathbf{z} = f(\mathbf{h})$, where $\mathbf{z} \in \mathbb{R}^{N_z}$ is drawn from a base distribution $p_{\mathbf{Z}}(\mathbf{z})$, as indicated by the subscript \mathbf{Z} . The inverse mapping is given by $\mathbf{h} = g(\mathbf{z}) = f^{-1}(\mathbf{z})$. By the change-of-variables formula, the density $p_{\mathbf{H}}(\mathbf{h})$ can be expressed as:

$$p_{\mathbf{H}}(\mathbf{h}) = p_{\mathbf{Z}}(f(\mathbf{h})) |\det[\mathbf{D}_f(\mathbf{h})]|, \quad (6)$$

where $\mathbf{D}_f(\mathbf{h}) = \frac{\partial f}{\partial \mathbf{h}} \in \mathbb{R}^{N_z \times N_h}$ is the Jacobian matrix of f , and $\det[\mathbf{D}_f(\mathbf{h})]$ denotes its determinant.

Let the forward transformation f be constructed as a composition of N_f invertible and differentiable bijections (i.e., flows) $f = f_{N_f} \circ f_{N_f-1} \circ \dots \circ f_2 \circ f_1$. We introduce a sequence of intermediate variables $\{\mathbf{h}_k\}_{k=0}^{N_f}$, defined recursively as $\mathbf{h}_k = f_k(\mathbf{h}_{k-1})$, $f_k: \mathbb{R}^{N_h} \rightarrow \mathbb{R}^{N_h}$, for $k = 1, \dots, N_f$, with $\mathbf{h}_0 = \mathbf{h}$ and $\mathbf{h}_{N_f} = \mathbf{z}$. Since each f_k is bijective, the overall transformation f is bijective as well. Denoting $g_k = f_k^{-1}$, the inverse transformation reads $g = g_1 \circ g_2 \circ \dots \circ g_{N_f-1} \circ g_{N_f}$. By the chain rule, the determinant of the

Algorithm 1 Simulation-based inference of parameter posterior via *neural* MCMC.

Input: Trained variational encoder $q_\theta(\mathbf{h} \mid \mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_\theta(\mathbf{x}), \boldsymbol{\sigma}_\theta^2(\mathbf{x}))$
Trained conditional normalizing flow $p_\xi(\mathbf{h} \mid \boldsymbol{\lambda})$
Prior distribution over target parameters $p(\boldsymbol{\lambda})$
Observational data \mathbf{x}
Scaling factor γ
Past-states archive update period k_T
Chain length N_s

▷ Probabilistic feature extraction

1: Compute the latent distribution $q_\theta(\mathbf{h} \mid \mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_\theta(\mathbf{x}), \boldsymbol{\sigma}_\theta^2(\mathbf{x}))$

▷ Sampling scheme

2: Initialize iteration counter $n = 1$

3: Initialize parameters $\boldsymbol{\lambda}^{(1)} \sim p(\boldsymbol{\lambda})$

4: Initialize past-states archive Λ

5: **while** $n < N_s$ **do**

6: $n = n + 1$

7: Sample latent variable $\mathbf{h}' \sim q_\theta(\mathbf{h} \mid \mathbf{x})$

8: Uniformly sample past states $\Lambda_{R1}, \Lambda_{R2}$ from Λ

9: Evaluate candidate parameter $\boldsymbol{\lambda}'$:

$$\boldsymbol{\lambda}' = \boldsymbol{\lambda}^{(n-1)} + \gamma(\Lambda_{R1} - \Lambda_{R2}) + \boldsymbol{\epsilon}$$

10: Estimate conditional likelihood density $p(\mathbf{h}' \mid \boldsymbol{\lambda}') \approx p_\xi(\mathbf{h}' \mid \boldsymbol{\lambda}')$

11: Compute acceptance ratio:

$$\alpha = \min \left(1, \frac{p_\xi(\mathbf{h}' \mid \boldsymbol{\lambda}') p(\boldsymbol{\lambda}')}{p_\xi(\mathbf{h}' \mid \boldsymbol{\lambda}^{(n-1)}) p(\boldsymbol{\lambda}^{(n-1)})} \right)$$

12: **if** $\boldsymbol{\lambda}'$ is accepted with probability α **then**

13: $\boldsymbol{\lambda}^{(n)} \leftarrow \boldsymbol{\lambda}'$

14: **else**

15: $\boldsymbol{\lambda}^{(n)} \leftarrow \boldsymbol{\lambda}^{(n-1)}$

16: **if** $n \equiv 0 \pmod{k_T}$ **then**

17: $\Lambda = [\Lambda, \boldsymbol{\lambda}^{(n)}]$

▷ Post-processing (discard burn-in period and chain thinning)

Output: Markov chain of samples from $p(\boldsymbol{\lambda} \mid \mathbf{h})$

full Jacobian $\mathbf{D}_f(\mathbf{h})$ can be written as the product of the determinants of the Jacobians of the individual sub-transformations:

$$\det[\mathbf{D}_f(\mathbf{h})] = \prod_{k=1}^{N_f} \det \mathbf{D}_{f_k}(\mathbf{h}_{k-1}). \quad (7)$$

Computing Jacobians and their determinants in high-dimensional spaces is generally computationally expensive. This challenge, along with the requirement for each transformation to be strictly bijective, makes the direct application of Eq. (6) impractical for modeling arbitrary distributions. To address this issue, Dinh et al. [44] introduced a coupling-layer-based approach that enables both tractable and expressive normalizing flows. Their key insight is that the determinant of a triangular matrix can be efficiently computed as the product of its diagonal entries. This principle underlies the design of coupling layers, which are simple bijective transformations that update only a subset of the input variables while keeping the remaining components fixed.

To implement each bijection f_k as a coupling layer, we partition the input \mathbf{h}_{k-1} into two subsets $(\mathbf{h}_{k-1}^A, \mathbf{h}_{k-1}^B) \in \mathbb{R}^{N_d} \times \mathbb{R}^{N_h - N_d}$. We then define parametrized bijections $\widehat{f}_k(\cdot; \mathbf{c}_k) : \mathbb{R}^{N_d} \rightarrow \mathbb{R}^{N_d}$, whose parameters $\mathbf{c}_k = C_k(\mathbf{h}_{k-1}^B; \boldsymbol{\xi}_k)$ are given by a learnable conditioning function (or conditioner) $C_k(\cdot; \boldsymbol{\xi}_k)$, with trainable weights $\boldsymbol{\xi}_k$. The resulting coupling transformation f_k is defined as:

$$\mathbf{h}_k = \begin{pmatrix} \mathbf{h}_k^A \\ \mathbf{h}_k^B \end{pmatrix} = f_k(\mathbf{h}_{k-1}) = \begin{pmatrix} \widehat{f}_k(\mathbf{h}_{k-1}^A; C_k(\mathbf{h}_{k-1}^B; \boldsymbol{\xi}_k)) \\ \mathbf{h}_{k-1}^B \end{pmatrix}, \quad (8)$$

where the conditioner C_k uses only the masked input \mathbf{h}_{k-1}^B to parametrize the transformation \widehat{f}_k applied to \mathbf{h}_{k-1}^A . By construction, f_k is invertible provided that \widehat{f}_k is invertible. Moreover, due to the input partitioning, the Jacobian matrix \mathbf{D}_{f_k} is block-triangular:

$$\mathbf{D}_{f_k}(\mathbf{h}_{k-1}) = \begin{bmatrix} \mathbf{D}_{\widehat{f}_k}(\mathbf{h}_{k-1}^A) & \frac{\partial \mathbf{h}_k^A}{\partial \mathbf{h}_{k-1}^B} \\ 0 & \mathbf{I}_{N_h - N_d} \end{bmatrix}, \quad (9)$$

where $\mathbf{I}_{N_h - N_d}$ denotes the identity matrix of size $N_h - N_d$. As a result, the determinant of $\mathbf{D}_{f_k}(\mathbf{h}_{k-1})$ reduces to the determinant of $\mathbf{D}_{\widehat{f}_k}(\mathbf{h}_{k-1}^A)$, which can be computed efficiently.

In this work, we adopt affine coupling layers, as used in RealNVP architectures [32], which are particularly effective for density estimation in high-dimensional settings. The transformation \widehat{f}_k applied to \mathbf{h}_{k-1}^A therefore consists of an element-wise scaling and shifting:

$$\mathbf{h}_k^A = \widehat{f}_k(\mathbf{h}_{k-1}^A; C_k(\mathbf{h}_{k-1}^B; \boldsymbol{\xi}_k)) = \mathbf{h}_{k-1}^A \odot \exp(s_k(\mathbf{h}_{k-1}^B)) + t_k(\mathbf{h}_{k-1}^B), \quad (10)$$

where $s_k : \mathbb{R}^{N_h - N_d} \mapsto \mathbb{R}^{N_d}$ and $t_k : \mathbb{R}^{N_h - N_d} \mapsto \mathbb{R}^{N_d}$ denote the scale and translation functions, respectively, and $\odot : (\mathbb{R}, \mathbb{R}) \mapsto \mathbb{R}$ denotes element-wise multiplication. Both s_k and t_k are produced by the conditioner C_k , which embeds the masked input \mathbf{h}_{k-1}^B into \widehat{f}_k . The corresponding Jacobian $\mathbf{D}_{\widehat{f}_k}(\mathbf{h}_{k-1}^A)$ is diagonal, with entries $\text{diag}[\exp(s_k(\mathbf{h}_{k-1}^B))]$. Consequently, the determinant of the full Jacobian $\mathbf{D}_{f_k}(\mathbf{h}_{k-1})$ simplifies to

$$\det \mathbf{D}_{f_k}(\mathbf{h}_{k-1}) = \exp \left(\sum_{m=1}^{N_d} s_k(\mathbf{h}_{k-1}^B)_m \right), \quad (11)$$

which can be computed efficiently without requiring the Jacobians of s_k or t_k . This feature allows s_k and t_k to be arbitrarily complex functions. Moreover, the inverse transformation can be computed as efficiently as the forward pass. Indeed, inverting the affine coupling layer does not require computing the inverses of s_k or t_k , thereby enabling the use of expressive, potentially non-invertible neural networks. Furthermore, as schematized in Fig. 2, we condition the computation of s_k and t_k on label information $\boldsymbol{\lambda}$, following the approach proposed in [45]. More formally, the coupling transformation \widehat{f}_k is parametrized as $\widehat{f}_k(\mathbf{h}_{k-1}^A; C_k(\mathbf{h}_{k-1}^B, \boldsymbol{\lambda}; \boldsymbol{\xi}_k))$, where the learnable conditioning network C_k takes both the masked features \mathbf{h}_{k-1}^B and label information $\boldsymbol{\lambda}$ as inputs. The scale and translation functions are thus defined as $s_k = s_k(\mathbf{h}_{k-1}^B, \boldsymbol{\lambda})$ and $t_k = t_k(\mathbf{h}_{k-1}^B, \boldsymbol{\lambda})$, respectively. This conditional affine coupling layer architecture enables efficient and accurate modeling of our *neural* MCMC likelihood mechanism.

Let $\mathcal{D}_{\mathbf{h}} = \{(\mathbf{h}_i, \boldsymbol{\lambda}_i)\}_{i=1}^{N_I}$ denote a dataset of N_I feature-label instances, where each feature \mathbf{h}_i (derived from raw observations \mathbf{x}_i) is associated with a target parameter vector $\boldsymbol{\lambda}_i$, $i = 1, \dots, N_I$. The CNF is trained to assign higher likelihoods to correctly labeled samples and lower likelihoods to incorrect ones, implicitly discriminating true or near-true label assignments from incorrect ones. Specifically, the set of learnable weights $\Theta_{\text{CNF}} = \{\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_{N_f}\}$ parametrizing the CFN is optimized by minimizing the negative log-likelihood of the observed data under the model:

$$\Theta_{\text{CNF}}^* = \arg \min_{\Theta_{\text{CNF}}} \mathcal{L}_{\text{CNF}}(\mathcal{D}_{\mathbf{h}}, \Theta_{\text{CNF}}), \quad (12)$$

$$\mathcal{L}_{\text{CNF}} = -\frac{1}{N_I} \sum_{i=1}^{N_I} \log p_{\xi}(\mathbf{h}_i | \boldsymbol{\lambda}_i) = -\frac{1}{N_I} \sum_{i=1}^{N_I} \left[\log p_{\mathbf{z}}(f(\mathbf{h}_i; \boldsymbol{\lambda}_i)) + \log |\det[\mathbf{D}_f(\mathbf{h}_i; \boldsymbol{\lambda}_i)]| \right]. \quad (13)$$

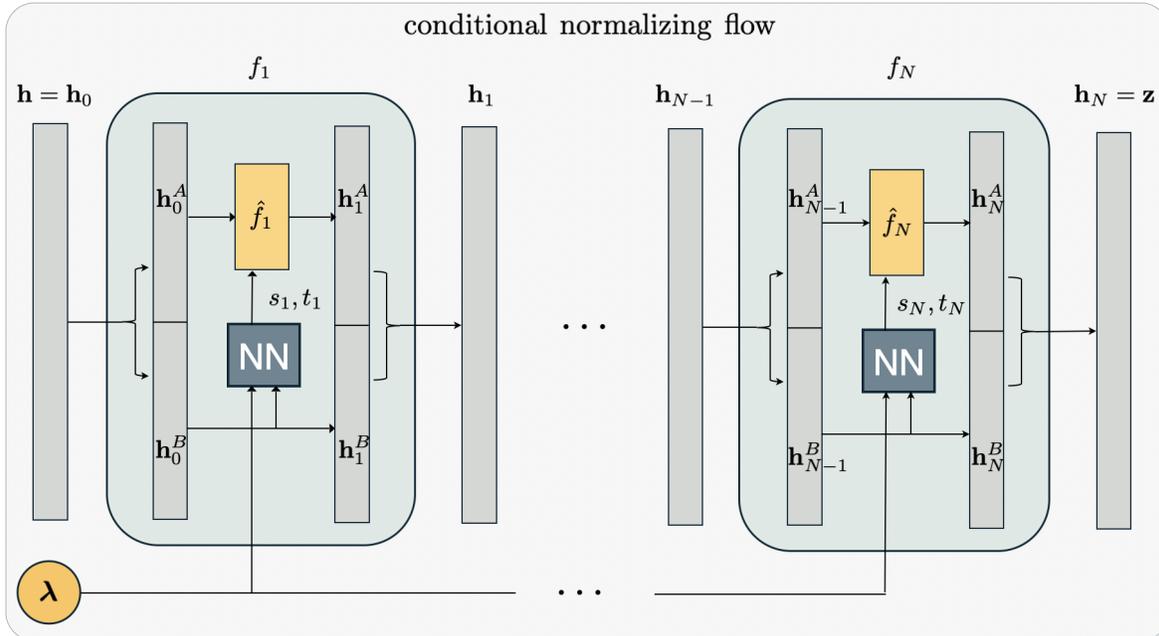


Figure 2: Schematic of the real-valued non-volume preserving normalizing flow with the additional conditioning structure on λ for estimating the likelihood density $p(\mathbf{h} | \lambda)$. Following the normalizing direction $\mathbf{z} = f(\mathbf{h})$, a sample \mathbf{h} from the variational approximation $q_\theta(\mathbf{h} | \mathbf{x})$ is transformed into a realization \mathbf{z} from the base distribution through a sequence of conditional affine coupling layers $f = f_N \circ f_{N-1} \circ \dots \circ f_2 \circ f_1$ conditioned on λ . For notational simplicity, the total number of flow layers is denoted by N in the figure, rather than N_f .

which explicitly depends on the conditioning λ_i . The objective in Eq. (13) corresponds to the empirical negative conditional log-likelihood, obtained via the change-of-variables formula (6), where $f(\cdot; \lambda)$ denotes the invertible transformation induced by the conditional flow and $\mathbf{D}_f(\cdot; \lambda)$ its Jacobian. It is worth noting that the logarithm formulation ensures numerical stability during gradient-based optimization.

This architecture offers a significant advantage over standard feed-forward neural networks as it enables exact likelihood evaluation through the change-of-variables formula. This capability makes conditional normalizing flows particularly well-suited for Bayesian inference, where accurate density evaluation is essential. In contrast, conventional neural networks typically learn deterministic mappings without defining consistent probability density functions over their output space. This difference is also reflected in the training process. Normalizing flows are trained by directly minimizing the negative log-likelihood of the data, as shown in (13), thereby preserving a clear probabilistic interpretation without relying on surrogate objectives. Minimizing this loss encourages the learned transformation to map samples from the latent distribution $q_\theta(\mathbf{h} | \mathbf{x})$ to the base distribution $p_{\mathbf{z}}(\mathbf{z})$, conditioned on λ , while maintaining exact likelihood evaluation. By contrast, standard DL approaches to density estimation often rely on indirect or approximate training objectives. Such approximations can struggle to capture complex dependencies or to represent multimodal, skewed, or highly correlated distributions.

2.3. Variational autoencoder for probabilistic feature extraction

Applying a normalizing flow directly to high-dimensional observational data can be computationally inefficient and prone to overfitting. Moreover, observation noise can obscure the underlying signal structure, degrading the quality of the learned likelihood model. For these reasons, we preliminarily compress data into a lower-dimensional latent space using a variational autoencoder.

Autoencoders are DL architectures designed for self-supervised representation learning. An AE consists of two main components: an encoder and a decoder, typically connected through a central bottleneck layer whose dimensionality is smaller than that of both the input and output layers. The model learns to encode the input data into a compact latent space and then decode it to reconstruct the original input. As a

result, the learned latent representation must preserve the essential information required for reconstruction, enabling feature extraction capabilities beyond traditional linear projection techniques.

In this work, we employ VAE-based feature extraction, which extends standard autoencoders by replacing the deterministic latent representation with a probabilistic one. Instead of mapping an input \mathbf{x} to a single point in the latent space, the encoder outputs the parameters of a distribution over latent variables. Assuming the Gaussian variational posterior $q_\theta(\mathbf{h} \mid \mathbf{x})$, the encoder defines the probabilistic encoding in Eq. (4). To enable stochastic sampling from $q_\theta(\mathbf{h} \mid \mathbf{x})$ while preserving differentiability during training, we adopt the reparameterization trick [46]. Specifically, latent samples are expressed as:

$$\mathbf{h} = \boldsymbol{\mu}_\theta(\mathbf{x}) + \boldsymbol{\sigma}_\theta(\mathbf{x}) \odot \boldsymbol{\epsilon}_{\text{VAE}}, \quad \boldsymbol{\epsilon}_{\text{VAE}} \sim \mathcal{N}(\mathbf{0}_{N_h}, \mathbf{I}_{N_h}), \quad (14)$$

which allows gradients to propagate through the sampling operation. Given a latent realization \mathbf{h} , the decoder $d_\theta : \mathbb{R}^{N_h} \mapsto \mathbb{R}^{N_x}$ is used to reconstruct the original signal in the observation space:

$$\hat{\mathbf{x}} = d_\theta(\mathbf{h}), \quad (15)$$

where $\hat{\mathbf{x}} \in \mathbb{R}^{N_x}$ denotes the reconstructed input.

The encoder and decoder are jointly trained to maximize a variational lower bound on the marginal log-likelihood of the observed data, known as the evidence lower bound (ELBO). This objective provides a tractable surrogate for the generally intractable data likelihood. By balancing reconstruction fidelity with latent-space regularization, ELBO maximization yields a smooth and well-structured latent manifold suitable for generative modeling and downstream inference tasks. For a single observation \mathbf{x} , the ELBO is given by:

$$\log p(\mathbf{x}) \geq \mathbb{E}_{q_\theta(\mathbf{h} \mid \mathbf{x})}[\log p_\theta(\mathbf{x} \mid \mathbf{h})] - \text{D}_{\text{KL}}[q_\theta(\mathbf{h} \mid \mathbf{x}) \parallel \bar{p}(\mathbf{h})], \quad (16)$$

where $\mathbb{E}_{q_\theta(\mathbf{h} \mid \mathbf{x})}[\cdot]$ denotes the expectation with respect to the variational posterior, and $\text{D}_{\text{KL}}[\cdot \parallel \cdot]$ denotes the Kullback–Leibler (KL) divergence between two probability distributions [47]. The conditional Gaussian likelihood $p_\theta(\mathbf{x} \mid \mathbf{h}) = \mathcal{N}(\mathbf{x}; d_\theta(\mathbf{h}), \sigma_x^2 \mathbf{I}_{N_x})$ is induced by the decoder and assumes a mean given by the deterministic decoding function $d_\theta(\mathbf{h})$ and fixed diagonal covariance. The observation noise level $\sigma_x \in \mathbb{R}$ is distinct from the standard deviation $\boldsymbol{\sigma}_\theta$ of the variational posterior and is typically fixed (e.g., $\sigma_x = 1$). Finally, $\bar{p}(\mathbf{h})$ denotes a prescribed prior over the latent variables, commonly chosen as a standard normal distribution.

The first term in Eq. (16) corresponds to a reconstruction objective. Under the assumption of a Gaussian likelihood, it reduces, up to additive constants independent of the model parameters, to a mean squared error (MSE) loss. Let $\mathcal{D}_\mathbf{x} = \{\mathbf{x}_i\}_{i=1}^{N_I}$ denote a dataset of N_I observation instances. Denoting the set of trainable parameters of the VAE by $\boldsymbol{\Theta}_{\text{VAE}}$, the reconstruction loss is defined as

$$\mathcal{L}_{\text{MSE}}(\mathcal{D}_\mathbf{x}, \boldsymbol{\Theta}_{\text{VAE}}) = \frac{1}{N_I} \sum_{i=1}^{N_I} -\frac{1}{2\sigma_x^2} \|\mathbf{x}_i - d_\theta(\mathbf{h}_i)\|_2^2, \quad (17)$$

which simplifies optimization while retaining a probabilistic interpretation.

The second term in Eq. (16) acts as a regularizer that constrains the learned latent distribution to remain close to the prior. For Gaussian variational posteriors, the KL divergence admits the following closed-form expression:

$$\text{D}_{\text{KL}}[q_\theta(\mathbf{h} \mid \mathbf{x}_i) \parallel \bar{p}(\mathbf{h})] = \frac{1}{2} \sum_{j=1}^{N_h} (1 + \log \sigma_j^2(\mathbf{x}_i) - \mu_j^2(\mathbf{x}_i) - \sigma_j^2(\mathbf{x}_i)), \quad (18)$$

where $\mu_j(\mathbf{x}_i)$ and $\sigma_j^2(\mathbf{x}_i)$ denote the mean and variance of the j -th latent dimension, $j = 1, \dots, N_h$. It is worth noting that this expression is equivalent to the standard KL divergence up to a sign convention, consistent with loss minimization. The corresponding loss component is defined as:

$$\mathcal{L}_{\text{KL}}(\mathcal{D}_\mathbf{x}, \boldsymbol{\Theta}_{\text{VAE}}) = \frac{1}{N_I} \sum_{i=1}^{N_I} \text{D}_{\text{KL}}[q_\theta(\mathbf{h} \mid \mathbf{x}_i) \parallel \bar{p}(\mathbf{h})]. \quad (19)$$

Combining reconstruction and regularization terms, the VAE parameters are learned by solving the following optimization problem:

$$\boldsymbol{\Theta}_{\text{VAE}}^* = \arg \min_{\boldsymbol{\Theta}_{\text{VAE}}} [\mathcal{L}_{\text{MSE}}(\mathcal{D}_\mathbf{x}, \boldsymbol{\Theta}_{\text{VAE}}) + \beta \mathcal{L}_{\text{KL}}(\mathcal{D}_\mathbf{x}, \boldsymbol{\Theta}_{\text{VAE}})], \quad (20)$$

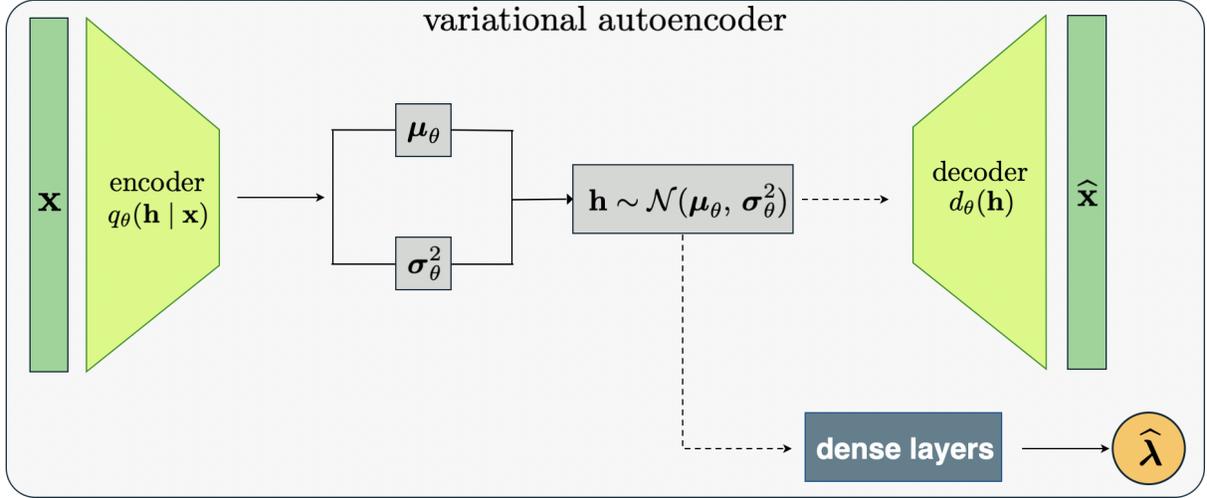


Figure 3: Schematic of the supervised variational autoencoder with the additional prediction block, enabling a label-informed representation for task-specific inference. Hat variables indicate neural network approximations.

where $\beta_{\text{KL}} \in \mathbb{R}^+$ controls the strength of the latent-space regularization.

Forward models are inevitably affected by model-form and epistemic uncertainties. Despite efforts to refine and calibrate them, the inherent complexity of physical and engineering systems often imposes limitations on modeling fidelity. A key concern that arises in this context is the risk of producing overconfident yet inaccurate estimates. Latent variable models such as VAEs can help bridge the gaps left by deterministic feature extraction techniques. In particular, their probabilistic structure can be leveraged to represent discrepancies between the true data-generating process and the distribution learned from training data, thereby enabling more robust and uncertainty-aware feature extraction.

While standard VAEs are self-supervised, our implementation augments this architecture by introducing a dense prediction head branching from the bottleneck layer, as shown in Fig. 3. We leverage this hybrid unsupervised-supervised formulation to learn a regularized latent space that is also enriched with label-relevant information, making the resulting representation more suitable for downstream inference of target variables.

During training, samples drawn from the latent posterior distribution are propagated through two parallel pathways: (i) one feeds the decoder for input reconstruction, and (ii) the other routes the latent samples through a stack of dense layers to produce task-specific predictions. This latter prediction branch introduces an additional supervised loss term into the overall training objective. In our applications, the labels correspond to the vectors of target parameters $\boldsymbol{\lambda}$. Accordingly, the prediction task is formulated as a regression problem, and the associated loss is defined as the MSE between the true and predicted parameter vectors:

$$\mathcal{L}_{\text{pred}} = \frac{1}{N_I} \sum_{i=1}^{N_I} \|\boldsymbol{\lambda}_i - \hat{\boldsymbol{\lambda}}_i\|_2^2, \quad (21)$$

where $\hat{\boldsymbol{\lambda}}_i$ denotes the predicted parameter vector associated with the i -th input sample.

Let $\mathcal{D}_\lambda = \{(\mathbf{x}_i, \boldsymbol{\lambda}_i)\}_{i=1}^{N_I}$ denote the dataset of input-label pairs, where each input \mathbf{x}_i is associated with a corresponding target parameter vector $\boldsymbol{\lambda}_i$. Denoting the set of learnable parameters of the informed-VAE (i-VAE) architecture by $\Theta_{\text{i-VAE}}$, the learning problem is formulated as the following optimization:

$$\Theta_{\text{i-VAE}}^* = \arg \min_{\Theta_{\text{i-VAE}}} \mathcal{L}_{\text{i-VAE}}(\mathcal{D}_\lambda, \Theta_{\text{i-VAE}}), \quad (22)$$

where $\mathcal{L}_{\text{i-VAE}}$ is the total loss function, given by:

$$\mathcal{L}_{\text{i-VAE}} = \mathcal{L}_{\text{MSE}} + \beta_{\text{KL}} \mathcal{L}_{\text{KL}} + \beta_{\text{pred}} \mathcal{L}_{\text{pred}}, \quad (23)$$

where \mathcal{L}_{MSE} enforces input reconstruction, \mathcal{L}_{KL} regularizes the latent space, and $\mathcal{L}_{\text{pred}}$ promotes the learning of latent representations that are informative for parameter prediction. The hyperparameters $\beta_{\text{KL}}, \beta_{\text{pred}} \in \mathbb{R}^+$ modulate the relative contributions of the loss terms, thereby controlling the trade-off between latent space structure, reconstruction fidelity, and inference accuracy.

3. Numerical Experiments

The proposed computational methodology is demonstrated through two case studies: (i) the structural health monitoring of a railway bridge, and (ii) the identification of a spatially varying hydraulic conductivity field in a groundwater flow system.

The full-order and reduced-order models for the first case study have been implemented in the `Matlab` environment using the `redbKIT` library [48]. The full-order problem for the second case study has been solved using the `Python` library `FEniCS` [49]. The DL architectures have been implemented using the `Tensorflow`-based `Keras` API [50], while the MCMC simulations have been carried out with the open-source `Python` library `PyMC` [38]. All computations have been performed on a MacBook Pro (Apple M1, 2020) with a unified memory architecture and 16 GB of RAM.

3.1. Railway bridge case study

The first case study concerns the Hörnefors railway bridge. A schematization of the bridge for structural health monitoring purposes is shown in Fig. 4. The geometrical and mechanical properties, as well as the overall computational setup, are adapted from former research activities [51–54]. The bridge has a span of 15.7 m, a free height of 4.7 m, and a width of 5.9 m (excluding the edge beams). The wing walls extend longitudinally up to 8 m. The structural elements have thicknesses of 0.5 m for the deck, 0.7 m for the frame walls, and 0.8 m for the wing walls. The bridge is founded on two plates connected by two stay beams and supported by pile groups. The concrete is of class C35/45, with Young’s modulus $E = 34$ GPa, Poisson’s ratio $\nu = 0.2$, and density $\rho = 2500$ kg/m³. The superstructure consists of a single track with sleepers spaced 0.65 m apart, resting on a ballast layer 0.6 m deep and 4.3 m wide, with density $\rho_B = 1800$ kg/m³. The bridge supports the transit of 8-axles Gröna Tåget trains traveling at speeds in the range $v \in [160, 215]$ km/h, with each axles carrying a mass of $m_A \in [16, 22]$ ton.

The data for this test case are synthetically generated using the numerical model described below. This approach allows full control over damage scenarios and operational conditions for training and validation of the proposed inference framework. The monitoring system provides displacement recordings

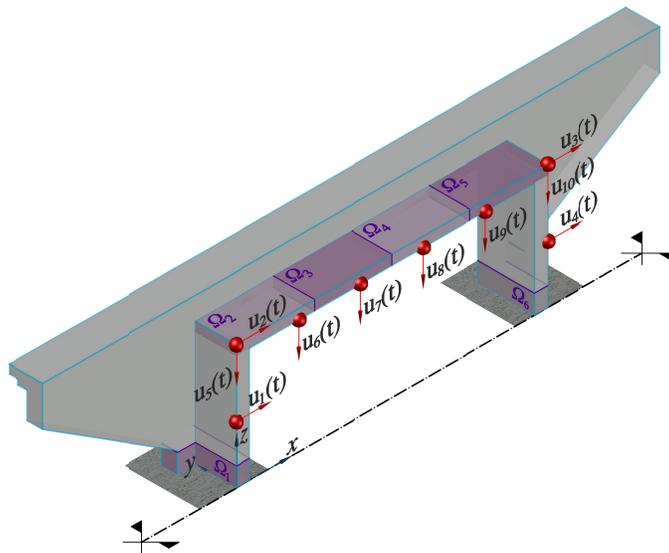


Figure 4: Railway bridge - Schematization for structural health monitoring, including details of synthetic recordings related to displacements $u_1(t), \dots, u_{10}(t)$, and predefined damage regions $\Omega_1, \dots, \Omega_6$.

$\mathbf{U}(\boldsymbol{\mu}) = [\mathbf{u}_1(\boldsymbol{\mu}), \dots, \mathbf{u}_{N_s}(\boldsymbol{\mu})] \in \mathbb{R}^{L \times N_s}$, consisting of $N_s = 10$ time series corresponding to degrees of freedom highlighted in Fig. 4. Each time series contains $L = 600$ displacement measurements uniformly sampled over the time interval $[0, T = 1.5 \text{ s}]$. These recordings coincide with the raw measurement vector \mathbf{x} used as input to the neural MCMC framework described in Sec. 2. The vector $\boldsymbol{\mu} \in \mathbb{R}^{N_\mu}$ collects N_μ parameters characterizing the operational and damage conditions. For the assumed problem setting, each observation window is sufficiently short for these parameters to be regarded as time-invariant, yet long enough to enable a reliable identification of the structural behavior. To account for measurement uncertainty, the displacement recordings are contaminated with additive Gaussian noise, yielding a signal-to-noise ratio of 100.

Changes in the structural dynamic response due to damage are modeled through a localized reduction of the effective stiffness, under the assumption of a clear time-scale separation between damage evolution and health assessment. The local stiffness degradation is parameterized by two variables, $y \in \mathbb{N}$ and $\delta \in \mathbb{R}$, both included in the parameter vector $\boldsymbol{\mu}$. Specifically, $y \in \{0, \dots, 6\}$ identifies the damage region among a set of $N_\Omega = 6$ predefined damageable subdomains $\Omega_1, \dots, \Omega_6$, as illustrated in Fig. 4 with $y = 0$ denoting the undamaged reference configuration. Within a damaged region, the stiffness is reduced of a factor $\delta \in [30\%, 80\%]$, kept fixed while a train travels across the bridge.

The full-order model features 17,292 degrees of freedom, resulting from a FE discretization using linear tetrahedral finite elements with an element size of 0.80 m, locally refined to 0.15 m in the deck region. The ballast layer is modeled by assigning an increased density to the deck and edge beams. The embankments are represented by distributed springs, implemented as Robin boundary conditions (with elastic coefficient 10^8 N/m^3) applied to the surfaces in contact with the ground. Structural dissipation is accounted for through Rayleigh damping, calibrated to yield a 5% damping ratio for the first two structural modes. The solution time interval is uniformly partitioned into 600 time steps, and the time integration is carried out using the implicit, unconditionally stable constant-average-acceleration Newmark scheme [55].

A projection-based reduced-order model is employed to speed up the offline dataset generation phase, following the approach adopted in [52]. Specifically, to reduce the computational cost associated with repeatedly solving the full-order model for different realizations of $\boldsymbol{\mu}$, a reduced-order model is constructed using a proper orthogonal decomposition (POD)-Galerkin reduced basis method [56]. The reduced-order model is derived from a snapshot matrix assembled from 400 full-order solutions corresponding to different parameter realizations $\boldsymbol{\mu} = (v, m_A, y, \delta)^\top$, which are taken as uniformly distributed and sampled via Latin hypercube rule. By enforcing a tolerance 10^{-3} on the fraction of disregarded energy, the resulting reduced basis is made by 133 POD modes. This reduced-order model is then used to populate the training dataset $\mathcal{D}_\lambda^{\text{train}}$ with 10,000 instances. These data are subsequently employed to train the DL models as described in Sec. 2. In contrast, a testing dataset $\mathcal{D}_\lambda^{\text{test}}$ consisting of 4000 instances is generated using the full-order model.

The parameter vector is organized as $\boldsymbol{\lambda} = (\delta(\Omega_1), \dots, \delta(\Omega_6))^\top$, with $N_\Omega = 6$ entries representing the damage magnitude within the damageable subdomains $\Omega_1, \dots, \Omega_6$, where the i -th component is denoted by $\lambda_i = \delta(\Omega_i)$. The low-dimensional feature space is set to dimension $N_h = 4$, which has been found to provide a sufficiently expressive yet compact representation for the CNF to compute accurate log-likelihood values.

A three-dimensional PCA projection of the N_h -dimensional latent mean vectors extracted by the i-VAE for each instance in the test set $\mathcal{D}_\lambda^{\text{test}}$ is shown in Fig. 5 (left). The color coding corresponds to the damage class y , revealing a clear cluster structure that reflects the sensitivity of the measured structural response to the damage location. Beyond damage localization, Fig. 5 (right) presents an analogous visualization for the damage level parameter δ , obtained by displaying only samples belonging to class $y = 1$. These scatter plots qualitatively illustrate how the i-VAE extracts informative latent representations, thereby enabling the CNF to produce accurate log-likelihood evaluations.

To assess the capability of the CNF to assign high log-likelihood values to correctly labeled samples and lower values to incorrect ones, Fig. 6 reports the log-likelihoods computed for the testing data. The CNF is evaluated on each test sample twice: once paired with its true label vector $\boldsymbol{\lambda}^*$, and once paired with a randomly selected label vector $\boldsymbol{\lambda}^{\text{wrong}}$ drawn from the testing dataset $\mathcal{D}_\lambda^{\text{test}}$. Log-likelihood values associated with the correct parameter vectors are shown in blue, while those corresponding to incorrect labels are shown in red. The model consistently assigns higher log-likelihood values to the correct labels, demonstrating its surrogate likelihood model capability.

The DE-MCMC algorithm is initialized with a uniform prior $p(\boldsymbol{\lambda})$, which is iteratively updated to the target posterior distribution $p(\boldsymbol{\lambda} | \mathbf{U})$ by running two chains sequentially, each for $N_s = 20,000$ iterations.

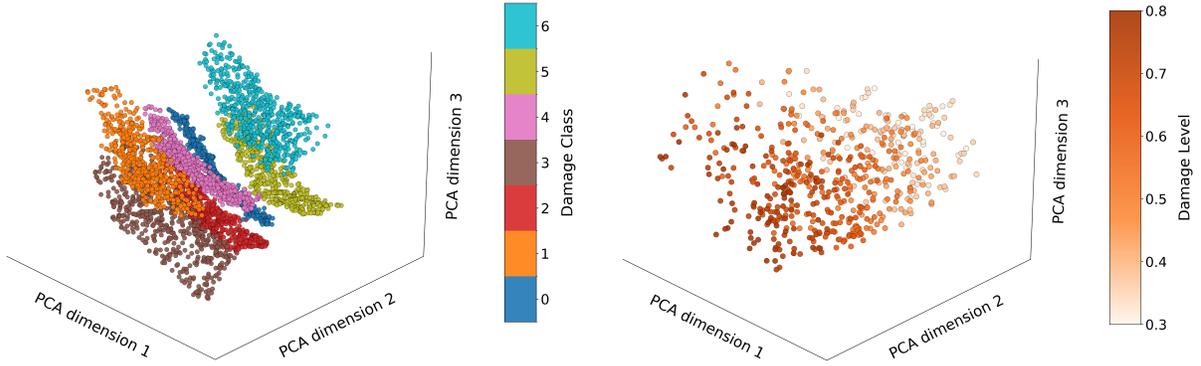


Figure 5: Railway bridge - Three-dimensional principal component projections of the latent-space mean vectors for the testing data. Left: data colored according to the target damage location class. Right: testing data belonging to damage class $y = 1$, colored according to the target damage magnitude.

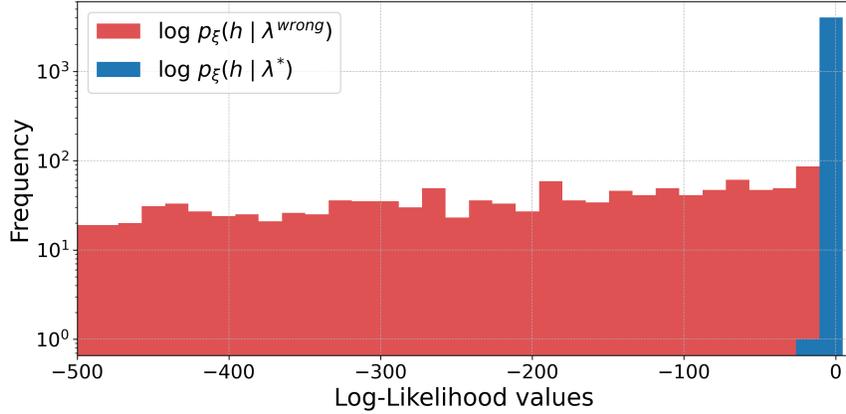


Figure 6: Railway bridge - Frequency distribution of the log-likelihood values for latent representations paired with correct labels (blue) and incorrect labels (red).

This number of samples has been selected to ensure a Gelman-Rubin statistic below 1.01, thereby assessing reliable convergence to a stationary distribution [57]. The runtime for the each chain is approximately 1.5 minutes.

For visualization purposes, Fig. 7 presents an exemplary posterior distribution recovered by DE-MCMC for damage located within Ω_6 . The parameters support ranges from 0 to 0.8, where 0 corresponds to an undamaged state and values from 0.3 to 0.8 indicate stiffness reductions $\delta \in [30\%, 80\%]$. Accordingly, the interval $[0, 0.3]$ represents the no-damage regime.

The results obtained from 100 independent DE-MCMC simulations using the testing dataset $\mathcal{D}_{\lambda}^{\text{test}}$ are summarized in Tab. 1. For each one of them, damage identification is performed using two point estimators from the posterior distribution $p(\lambda | \mathbf{U})$, namely the posterior mean $\hat{\lambda}_{\text{mean}} = \mathbb{E}[p(\lambda | \mathbf{U})]$, and the maximum a posteriori (MAP) estimate $\hat{\lambda}_{\text{MAP}} = \arg \max_{\lambda} p(\lambda | \mathbf{U})$. The ground-truth parameter vector used to generate the observations is denoted by λ^* .

The correct damage class is identified with an accuracy of 93% when using posterior mean and 97% when using MAP. Moreover, the absolute error in the estimated damage magnitude δ is approximately 6% for the posterior mean and 7% for the posterior mode, over an admissible support of 50%. The posterior average standard deviations for each damage class are reported in Tab. 2. Overall, both damage location and magnitude are identified with high accuracy and relatively low uncertainty, with no cases exhibiting

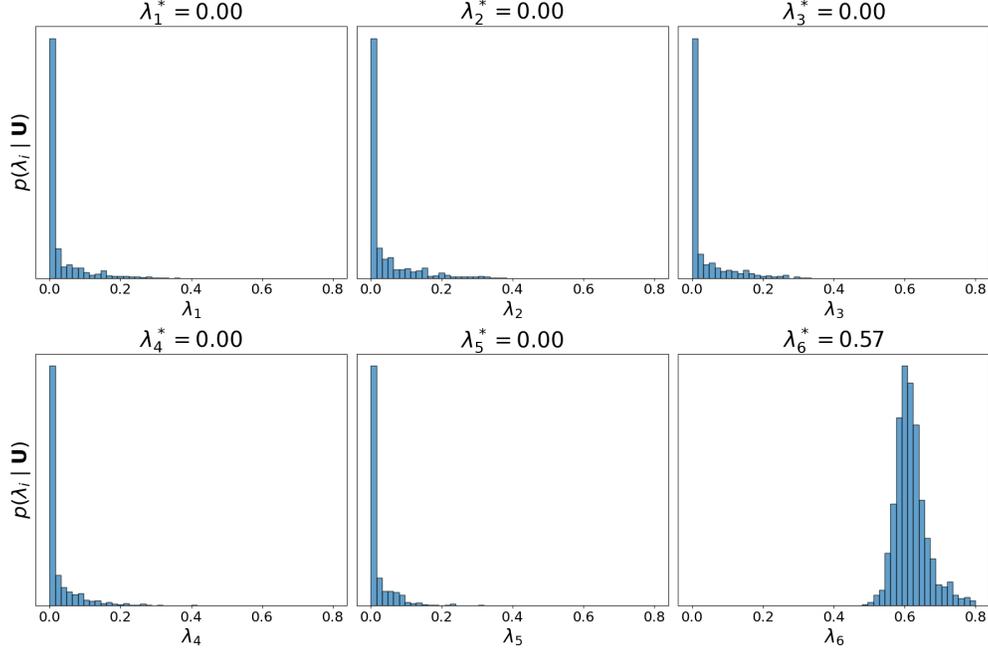


Figure 7: Railway bridge - Exemplary DE-MCMC result showing the posterior marginal distributions from $p(\boldsymbol{\lambda} \mid \mathbf{U})$, each associated with the damage magnitude within a damageable subdomain.

Table 1: Railway bridge - Damage localization and quantification results under different operational and damage conditions, reported in terms of classification accuracy and absolute error for the posterior mean and posterior mode.

Metric	With $\hat{\boldsymbol{\lambda}}_{\text{mean}}$	With $\hat{\boldsymbol{\lambda}}_{\text{MAP}}$
Classification accuracy	93%	97%
Absolute error	0.062	0.072

Table 2: Railway bridge - Posterior standard deviation of the estimated damage magnitude for each damage location.

True damage region	Ω_1	Ω_2	Ω_3	Ω_4	Ω_5	Ω_6
Standard deviation	0.0787	0.0670	0.0907	0.1196	0.1422	0.0899

significant discrepancies between the true damage parameters and the posterior estimates.

3.2. Groundwater flow case study

This second case study concerns a groundwater flow problem, adapted from the work of [58]. It considers a steady-state flow in a confined, inhomogeneous aquifer with domain $\Omega = (0, 1) \times (0, 1)$. Under the assumption of a stationary and incompressible fluid flow, the governing equation that models the pressure distribution $h(\mathbf{x})$ in a porous medium is the following scalar elliptic partial differential equation:

$$\begin{cases} -\nabla \cdot (T(\mathbf{x})\nabla h(\mathbf{x})) = g(\mathbf{x}) & \forall \mathbf{x} \text{ in } \Omega, \\ h(\mathbf{x}) = h_D(\mathbf{x}) & \forall \mathbf{x} \text{ on } \Gamma_D, \\ (T(\mathbf{x})\nabla h(\mathbf{x})) \cdot \mathbf{n} = q_N(\mathbf{x}) & \forall \mathbf{x} \text{ on } \Gamma_N. \end{cases} \quad (24)$$

Here, $T(\mathbf{x})$ denotes the heterogeneous, depth-integrated transmissivity, while in the absence of internal sources or sinks, $g(\mathbf{x}) = 0$. The function $h_D(\mathbf{x})$ prescribes the hydraulic head on the Dirichlet boundary $\Gamma_D = \{0\} \times (0, 1) \cup \{1\} \times (0, 1)$, while $q_N(\mathbf{x})$ denotes the Darcy flux across the Neumann boundary $\Gamma_N = \partial\Omega \setminus \Gamma_D$, with \mathbf{n} being the outward unit normal vector on Γ_N .

To approximate the hydraulic head $\mathcal{h}(\mathbf{x})$, we introduce a finite-dimensional subspace defined over a computational mesh with N_N nodes, spanned by piecewise linear Lagrange basis functions. Substituting the discrete approximation of the hydraulic head into the weak formulation of problem (24) yields a system of sparse linear equations, whose solution vector $\mathbf{h} \in \mathbb{R}^{N_N}$ collects the nodal values of the hydraulic head.

With the aim of identifying a spatially varying hydraulic conductivity field, the aquifer transmissivity is modeled as a log-Gaussian random field. This modeling choice is well established in the literature [59, 60], particularly for problems in which the medium does not exhibit strongly correlated extreme values, geological faults, or sharp discontinuities. Accordingly, the spatial correlation structure of the transmissivity field is characterized by a covariance kernel $C(\mathbf{x}, \mathbf{x}')$. In our numerical experiments, we adopt a squared exponential (Gaussian) kernel:

$$C(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{1}{2} \sum_{j=1}^{N_D} \left(\frac{x_j - x'_j}{\ell}\right)^2\right), \quad (25)$$

where N_D denotes the spatial dimensionality of the problem and $\ell \in \mathbb{R}^+$ is the correlation length controlling the spatial smoothness of the field.

To parametrize the random field, we construct the covariance matrix $\mathbf{C} \in \mathbb{R}^{N_N \times N_N}$, with entries $C_{ij} = C(\mathbf{x}_i, \mathbf{x}_j)$, $i, j = 1, \dots, N_N$, where \mathbf{x}_i and \mathbf{x}_j denote the spatial coordinates of the i -th and j -th nodes of the finite element mesh, respectively. The random field is then parametrized using a Karhunen-Loève (KL) expansion. Let $\langle \pi_i, \psi_i \rangle$, $i = 1, \dots, N_N$, denote the eigenpairs of the covariance matrix \mathbf{C} , sorted in descending order of the eigenvalues. Truncating the expansion after the first N_M dominant modes, the KL representation of the log-transmissivity field reads:

$$\log \mathbf{t} = \boldsymbol{\mu}_t + \sigma_t \boldsymbol{\Psi} \boldsymbol{\Pi}^{1/2} \boldsymbol{\lambda}, \quad (26)$$

where $\mathbf{t} \in \mathbb{R}^{N_N}$ is the discretized transmissivity field, $\boldsymbol{\mu}_t \in \mathbb{R}^{N_N}$ denotes the mean vector of the log-transmissivity field, $\sigma_t \in \mathbb{R}^+$ is its marginal standard deviation, $\boldsymbol{\Psi} = [\boldsymbol{\psi}_1, \dots, \boldsymbol{\psi}_{N_M}] \in \mathbb{R}^{N_N \times N_M}$ collects as columns the retained eigenvectors, $\boldsymbol{\Pi} = \text{diag}[\pi_1, \dots, \pi_{N_M}]$ is a diagonal matrix collecting the corresponding eigenvalues, and $\boldsymbol{\lambda} \sim \mathcal{N}(\mathbf{0}_{N_M}, \mathbf{I}_{N_M})$ is a vector of $N_{\text{par}} = N_M$ independent standard Gaussian random

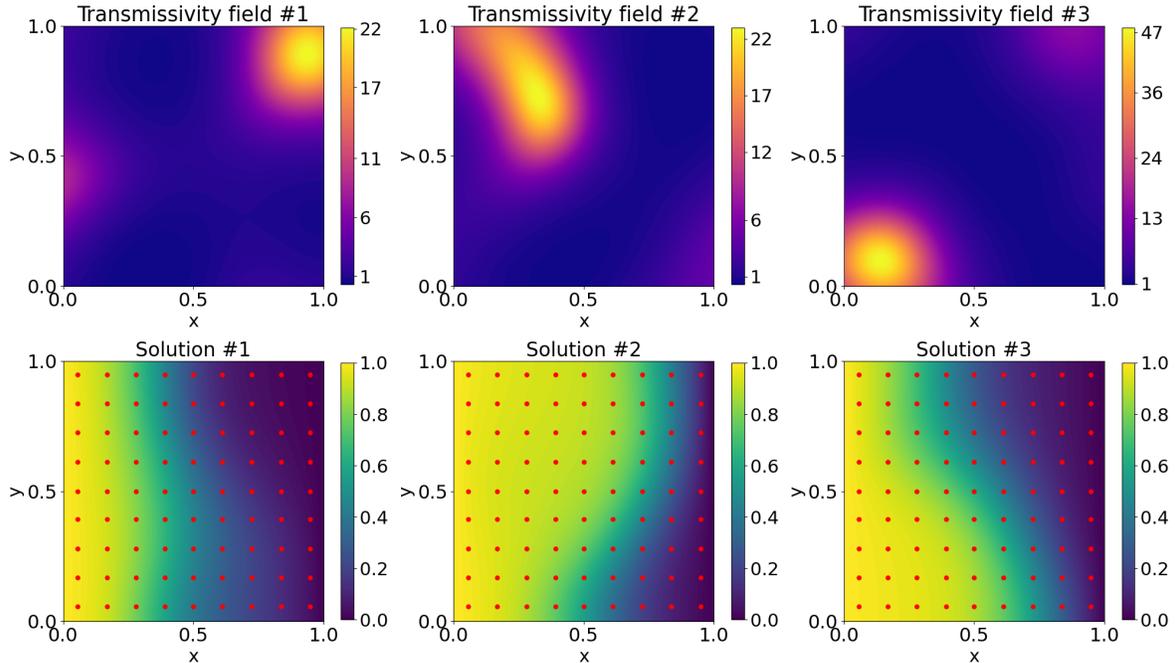


Figure 8: Groundwater flow - Exemplary realizations of the transmissivity field and the corresponding hydraulic head solutions. Sensor locations used for inference are highlighted in red.

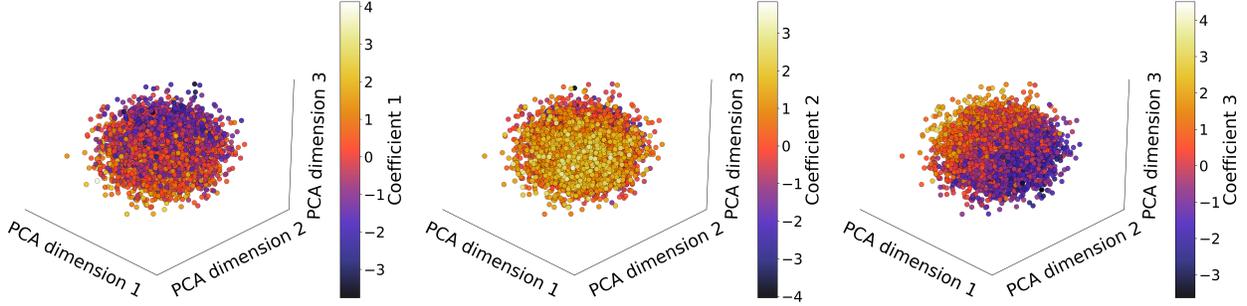


Figure 9: Groundwater flow - Three-dimensional principal component projections of the latent-space mean vectors for the testing data. From left to right, the color coding corresponds to the true values of the first three coefficients of the Karhunen-Loève representation of the log-transmissivity field.

variables. This latter provides a low-dimensional representation of the spatially varying transmissivity field and constitutes the set of target parameters to be inferred.

The computational domain Ω is discretized using linear finite elements with a dimensionless element size of $1/60$. A Dirichlet unit hydraulic head is prescribed on the left boundary, while a homogeneous Dirichlet condition is imposed on the right boundary. Zero-flux Neumann boundary conditions are imposed on the top and bottom boundaries.

The log-transmissivity field is modeled with unit mean μ_t and unit marginal standard deviation σ_t . The spatial correlation length is set to $\ell = 0.25$. The resulting covariance matrix is decomposed via KL expansion and truncated after the first $N_M = 14$ dominant modes, which collectively retain approximately 97% of the total variance. Accordingly, the log-transmissivity field is parametrized by a low-dimensional vector of $N_{\text{par}} = N_M = 14$ independent Gaussian coefficients λ . Figure 8 illustrates exemplary realizations of the transmissivity field together with the corresponding hydraulic head solutions \mathbf{h} .

Inference over the parameter vector λ is performed using hydraulic head measurements collected on a regular 9×9 grid of equispaced sensor locations, highlighted in Fig. 8. These measurements are extracted from the full hydraulic head solution \mathbf{h} through a Boolean selection operator and constitute the raw observation vector, \mathbf{x} , with dimension $N_x = 81$. With this setup, the training dataset $\mathcal{D}_\lambda^{\text{train}}$ is populated with 32,000 realizations to train the DL models, while the testing dataset $\mathcal{D}_\lambda^{\text{test}}$ is populated with 8000 additional instances.

A three-dimensional PCA projection of the $N_h = 20$ -dimensional latent mean vectors extracted by the i-VAE for each instance in $\mathcal{D}_\lambda^{\text{test}}$ is shown in Fig. 9. The color coding corresponds to the values of the first three coefficients of the KL representation of the log-transmissivity field. Compared to the previous case study, the resulting low-dimensional representation appears fuzzier, with only weakly discernible structure and no clear manifold topology. This behavior can be attributed to the increased complexity of the mapping between the parameter vector λ and the observations \mathbf{x} , as well as to the higher dimensionality of λ in this problem.

The ability of the CNF to assign informative log-likelihoods values is summarized in Fig. 10. As in the previous case study, the CNF is evaluated on each test sample twice: once paired with its true coefficient vector λ , shown in blue, and once paired with a randomly selected parameter vector drawn from the testing dataset $\mathcal{D}_\lambda^{\text{test}}$, shown in red. The resulting distribution clearly indicates that the CNF consistently assigns higher log-likelihood values to correctly labeled samples, while penalizing incorrect parameterizations. This behavior confirms the ability of the CNF to reliably discriminate between compatible and incompatible (\mathbf{x}, λ) pairs, even in the presence of a high-dimensional nonlinear inverse mapping.

Starting from a standard multivariate normal prior $p(\lambda)$, the posterior distribution $p(\lambda | \mathbf{x})$ is approximated using the DE-MCMC algorithm. Observed hydraulic head measurements \mathbf{x} are first mapped into a low-dimensional latent space by the i-VAE encoder, producing a probabilistic latent representation $\mathbf{h} \sim \mathcal{N}(\mu_\theta, \sigma_\theta^2)$. Joint samples (\mathbf{h}', λ') are then drawn from the latent and parameter spaces, and the CNF is used to evaluate the conditional likelihood $p(\mathbf{h}' | \lambda')$, which efficiently informs the DE-MCMC acceptance ratio.

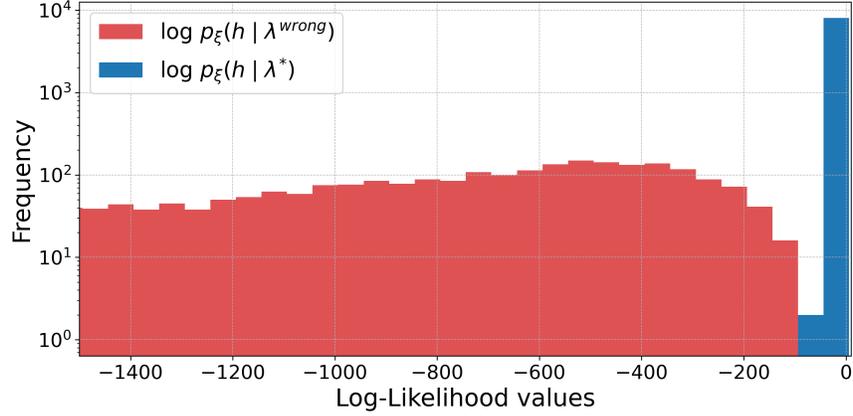


Figure 10: Groundwater flow - Frequency distribution of the log-likelihood values for latent representations paired with correct labels (blue) and incorrect labels (red).

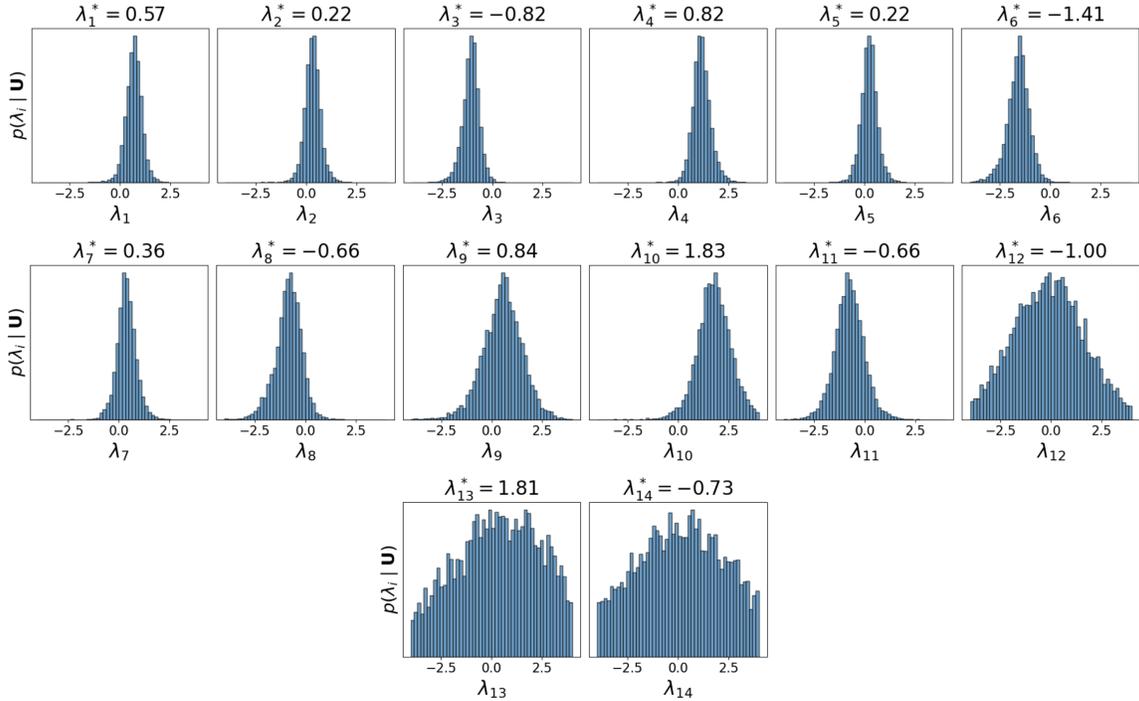


Figure 11: Groundwater flow - Exemplary MCMC result showing the posterior distributions of the 14 coefficients of the Karhunen-Loève representation of the log-transmissivity field.

Two chains are run sequentially, each with $N_s = 30,000$ iterations, ensuring a Gelman-Rubin statistic below 1.01 in all cases. The runtime for each chain is approximately 3 minutes. An exemplary MCMC output is shown in Fig. 11. The corresponding transmissivity field, reconstructed using either the posterior mean or posterior mode of the KL coefficients, is compared with the ground truth in Fig. 12.

The increasing uncertainty observed in the posterior distributions from the first to the last KL coefficients is documented in Fig. 13. This trend is expected, since higher-order KL modes contribute less to the overall structure of the transmissivity field and are therefore less constrained by the observations. However, this increased uncertainty has a negligible impact on the accuracy of the reconstructed transmissivity fields, as the influence of these modes on the solution is limited.

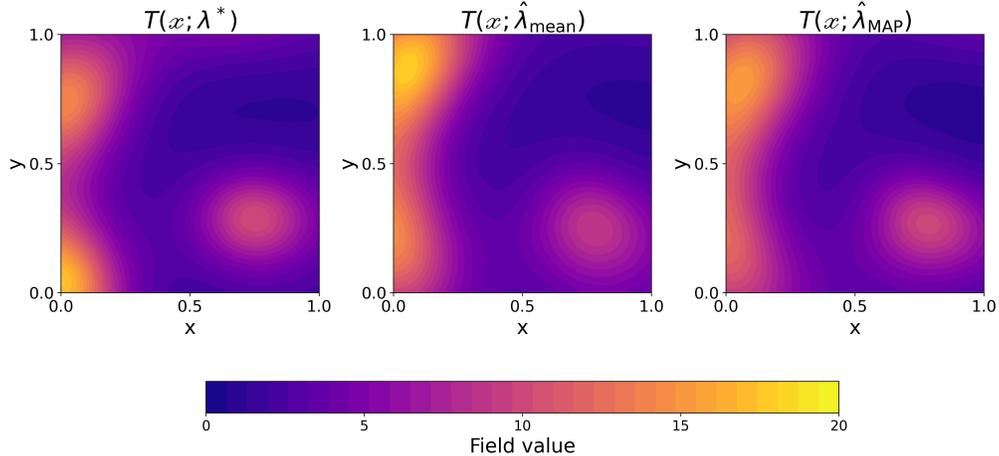


Figure 12: Groundwater flow - Exemplary comparison between (left) ground truth, (middle) estimated transmissivity fields using the posterior mean of the Karhunen-Loève coefficients, and (right) estimated transmissivity fields using the posterior mode of the Karhunen-Loève coefficients.

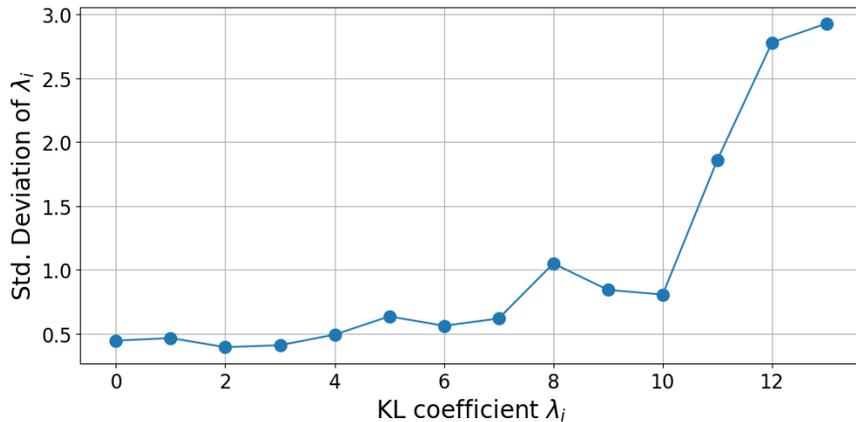


Figure 13: Groundwater flow - Posterior standard deviation of the Karhunen-Loève coefficients, averaged over the testing dataset $\mathcal{D}_{\lambda}^{\text{test}}$.

To quantify the reconstruction performance, 100 DE-MCMC simulations are performed on randomly selected test samples from $\mathcal{D}_{\lambda}^{\text{test}}$. The relative error is computed as

$$\text{relative error} = \frac{\|\mathbf{t}_{\text{true}} - \mathbf{t}_{\text{pred}}\|_2}{\|\mathbf{t}_{\text{true}}\|_2}. \quad (27)$$

The reconstructed transmissivity fields using the posterior mean and posterior mode achieve relative errors of approximately 0.186 and 0.168, respectively; see also Fig. 14 for the corresponding box plots. These results demonstrate that the i-VAE provides a compact and informative latent representation, enabling the CNF to produce accurate log-likelihoods, which in turn supports reliable posterior inference via DE-MCMC.

4. Conclusions

This work introduces a simulation-based inference framework for inverse problems, integrating representation learning, conditional density estimation, and Bayesian sampling. The proposed methodology embeds a supervised variational autoencoder (VAE) and a conditional normalizing flow (CNF) within a Markov chain Monte Carlo (MCMC) scheme, enabling efficient posterior inference in high-dimensional and nonlinear settings.

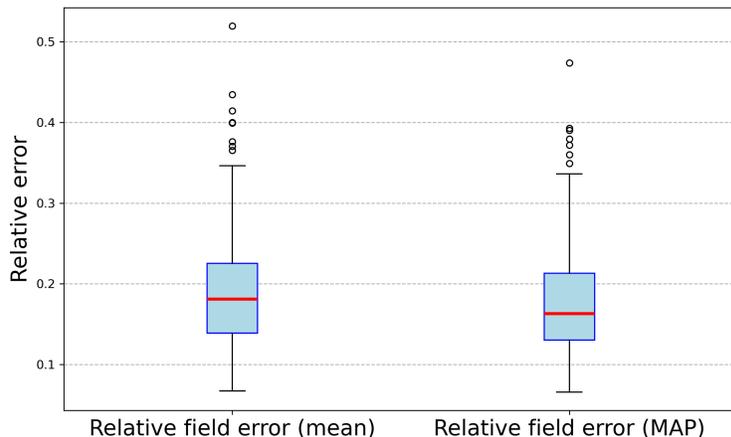


Figure 14: Groundwater flow - Box plots of the relative errors in transmissivity identification for 100 random samples from $\mathcal{D}_{\lambda}^{\text{test}}$, obtained from field reconstructions using (left) the posterior mean and (right) the MAP of the Karhunen-Loève coefficients.

A key contribution is the design of a modified VAE that learns a probabilistic latent representation explicitly informed by the parameters of interest. By augmenting the standard VAE architecture with a prediction branch operating on latent samples, the resulting latent space is both regularized and parameter-informative, providing a compact and uncertainty-aware summary of the observations. A CNF trained in this latent space approximates the conditional likelihood of the latent features given the parameters, replacing repeated forward model evaluations during MCMC sampling while preserving a fully Bayesian treatment of uncertainty.

The framework has been demonstrated on two distinct inverse problems: damage localization and quantification in a railway bridge, and identification of a spatially varying hydraulic transmissivity field in a groundwater flow problem. In both cases, the approach achieved accurate parameter inference and meaningful uncertainty quantification, despite the high dimensionality and ill-posedness of the problems.

The proposed framework is broadly applicable to inverse problems involving expensive forward models and limited observational data. Future work will focus on extensions to time-dependent settings, adaptive learning strategies, and tighter integration within digital twin architectures for real-time data assimilation and decision support [61].

Acknowledgments: The authors thank Dr. Filippo Zacchei (MOX-Department of Mathematics, Politecnico di Milano) for his insightful comments and remarks.

Funding: GB acknowledges the support of European Union - NextGenerationEU within the Italian PNRR program (M4C2, Investment 3.3) for the PhD Scholarship “Physics-informed Artificial Intelligence for surrogate modeling of complex energy systems”.

GB and AM are members of the Gruppo Nazionale Calcolo Scientifico-Istituto Nazionale di Alta Matematica (GNCS-INdAM). AM acknowledges the project “Dipartimento di Eccellenza” 2023–2027, funded by MUR, as well as the financial support from the FIS starting grant DREAM (Grant Agreement FIS00003154), funded by the Italian Science Fund (FIS) - Ministero dell’Università e della Ricerca.

MT and AM acknowledge the financial support from the ERC advanced grant IMMENSE (Grant Agreement 101140720), funded by the European Union. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- [1] A. M. Stuart, Inverse problems: A Bayesian perspective, *Acta Numerica* 19 (2010) 451–559. [doi:10.1017/S0962492910000061](https://doi.org/10.1017/S0962492910000061).
- [2] A. Tarantola, *Inverse Problem Theory and Methods for Model Parameter Estimation*, SIAM: Society for Industrial and Applied Mathematics, Philadelphia, PA, 2005. [doi:10.1137/1.9780898717921](https://doi.org/10.1137/1.9780898717921).
- [3] A. Doucet, A. M. Johansen, A Tutorial on Particle Filtering and Smoothing: Fifteen years later, in: D. Crisan, B. Rozovskii (Eds.), *Handbook of Nonlinear Filtering*, Oxford University Press, Oxford, NY, 2009, pp. 656–704.
- [4] A. Doucet, N. de Freitas, N. Gordon, *Sequential Monte Carlo Methods in Practice*, Springer New York, NY, 2001. [doi:10.1007/978-1-4757-3437-9](https://doi.org/10.1007/978-1-4757-3437-9).
- [5] G. P. Way, C. S. Greene, Extracting a biologically relevant latent space from cancer transcriptomes with variational autoencoders, in: *Proceedings of the 2018 Pacific Symposium on Biocomputing*, Kohala Coast, Hawaii, 2018, pp. 80–91.
- [6] Y. Pu, Z. Gan, R. Henao, X. Yuan, C. Li, A. Stevens, L. Carin, Variational autoencoder for deep learning of images, labels and captions, in: *Proceedings of the 2016 Advances in Neural Information Processing Systems*, Vol. 29, Curran Associates, Inc., Barcelona, Spain, 2016, pp. 2352–2360. [doi:10.5555/3157096.3157360](https://doi.org/10.5555/3157096.3157360).
- [7] J. Ehrhardt, M. Wilms, Autoencoders and variational autoencoders in medical image analysis, in: N. Burgos, D. Svoboda (Eds.), *Biomedical Image Synthesis and Simulation*, Academic Press, 2022, pp. 129–162. [doi:10.1016/B978-0-12-824349-7.00015-3](https://doi.org/10.1016/B978-0-12-824349-7.00015-3).
- [8] Y. Fan, G. Wen, D. Li, S. Qiu, M. D. Levine, F. Xiao, Video Anomaly Detection and Localization via Gaussian Mixture Fully Convolutional Variational Autoencoder, *Computer Vision and Image Understanding* 195 (2020) 102920. [doi:10.1016/j.cviu.2020.102920](https://doi.org/10.1016/j.cviu.2020.102920).
- [9] M. J. Kusner, B. Paige, J. M. Hernández-Lobato, Grammar Variational Autoencoder, in: *Proceedings of the 2017 International Conference on Machine Learning*, Vol. 70, Sydney, Australia, 2017, pp. 1945–1954. [doi:10.5555/3305381.3305582](https://doi.org/10.5555/3305381.3305582).
- [10] S. Arridge, P. Maass, O. Öktem, C.-B. Schönlieb, Solving Inverse Problems Using Data-Driven Models, *Acta Numerica* 28 (2019) 1–174. [doi:10.1017/S0962492919000041](https://doi.org/10.1017/S0962492919000041).
- [11] M. A. Iglesias, K. J. Law, A. M. Stuart, Evaluation of Gaussian Approximations for Data Assimilation in Reservoir Models, *Computational Geosciences* 17 (2013) 851–885. [doi:10.1007/s10596-013-9359-x](https://doi.org/10.1007/s10596-013-9359-x).
- [12] H. Antil, H. C. Elman, A. Onwunta, D. Verma, Novel Deep Neural Networks for Solving Bayesian Statistical Inverse Problems, *arXiv preprint arXiv:2102.03974* (2021). [doi:10.48550/arXiv.2102.03974](https://doi.org/10.48550/arXiv.2102.03974).
- [13] T. Deveney, E. Mueller, T. Shardlow, A Deep Surrogate Approach to Efficient Bayesian Inversion in PDE and Integral Equation Models, *arXiv preprint arXiv:1910.01547* (2019). [doi:10.48550/arXiv.1910.01547](https://doi.org/10.48550/arXiv.1910.01547).
- [14] L. Yan, T. Zhou, An Adaptive Surrogate Modeling Based on Deep Neural Networks for Large-Scale Bayesian Inverse Problems, *arXiv preprint arXiv:1911.08926* (2019). [doi:10.48550/arXiv.1911.08926](https://doi.org/10.48550/arXiv.1911.08926).
- [15] H. Goh, S. Sheriffdeen, J. Wittmer, T. Bui-Thanh, Solving Bayesian Inverse Problems via Variational Autoencoders, *Journal of Machine Learning Research* 20 (2020) 1–49.
- [16] N. Akkarapatty, A. Muralidharan, N. S. Raj, et al., Dimensionality reduction techniques for text mining, in: *Collaborative filtering using data mining and analysis*, IGI Global, 2017, pp. 49–72. [doi:10.4018/978-1-5225-0489-4.ch003](https://doi.org/10.4018/978-1-5225-0489-4.ch003).

- [17] W. K. Mutlag, S. K. Ali, Z. M. Aydam, B. H. Taher, Feature Extraction Methods: A Review, *Journal of Physics: Conference Series* 1591 (2020) 012028. [doi:10.1088/1742-6596/1591/1/012028](https://doi.org/10.1088/1742-6596/1591/1/012028).
- [18] P. Xanthopoulos, P. M. Pardalos, T. B. Trafalis, Linear Discriminant Analysis, in: *Robust Data Mining*, Springer New York, New York, NY, 2013, pp. 27–33. [doi:10.1007/978-1-4419-9878-1_4](https://doi.org/10.1007/978-1-4419-9878-1_4).
- [19] C. Doersch, Tutorial on variational autoencoders, arXiv preprint arXiv:1606.05908 (2016). [doi:10.48550/arXiv.1606.05908](https://doi.org/10.48550/arXiv.1606.05908).
- [20] I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning*, MIT Press, Cambridge, MA, 2016.
- [21] K. Cranmer, J. Brehmer, G. Louppe, The frontier of simulation-based inference, *Proceedings of the National Academy of Sciences* 117 (2020) 30055–30062. [doi:10.1073/pnas.1912789117](https://doi.org/10.1073/pnas.1912789117).
- [22] G. Papamakarios, Neural Density Estimation and Likelihood-free Inference, arXiv preprint arXiv:1910.13233 (2019). [doi:10.48550/arXiv.1910.13233](https://doi.org/10.48550/arXiv.1910.13233).
- [23] G. Papamakarios, D. C. Sterratt, I. Murray, Sequential Neural Likelihood: Fast Likelihood-free Inference with Autoregressive Flows, in: *Proceedings of the 2019 International Conference on Machine Learning*, Long Beach, CA, 2019, pp. 3835–3844.
- [24] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, B. Lakshminarayanan, Normalizing Flows for Probabilistic Modeling and Inference, *Journal of Machine Learning Research* 22 (2021) 1–64.
- [25] Y. M. Marzouk, H. N. Najm, L. A. Rahn, A stochastic Galerkin method for Bayesian inference in inverse problems, *Journal of Computational Physics* 228 (6) (2009) 1862–1902. [doi:10.1016/j.jcp.2008.11.024](https://doi.org/10.1016/j.jcp.2008.11.024).
- [26] P. Kersaudy, B. Sudret, N. Varsier, O. Picon, J. Wiart, Bayesian inference using sparse polynomial chaos expansions, *Journal of Computational Physics* 311 (2016) 267–294. [doi:10.1016/j.jcp.2015.12.047](https://doi.org/10.1016/j.jcp.2015.12.047).
- [27] C. Rossat, D. Azzimonti, D. Ginsbourger, Accelerating Bayesian inverse problems with surrogate-assisted Markov chain Monte Carlo, *Journal of Computational Physics* 459 (2022) 110986. [doi:10.1016/j.jcp.2022.110986](https://doi.org/10.1016/j.jcp.2022.110986).
- [28] Q. Huang, A. Fichtner, Recent advances in surrogate-based Bayesian inference for geophysical inverse problems, *Geophysical Journal International* 238 (1) (2024) 1–26. [doi:10.1093/gji/ggae040](https://doi.org/10.1093/gji/ggae040).
- [29] S. A. Sisson, Y. Fan, Likelihood-Free MCMC, in: S. Brooks, A. Gelman, G. Jones, X.-L. Meng (Eds.), *Handbook of Markov Chain Monte Carlo*, Chapman & Hall/CRC, Boca Raton, FL, 2011, pp. 313–335.
- [30] J. Hermans, V. Begy, G. Louppe, Likelihood-free MCMC with Amortized Approximate Ratio Estimators, in: *Proceedings of the 2020 International Conference on Machine Learning*, Vol. 119, Online, 2020, pp. 4239–4248. [doi:10.5555/3524938.3525333](https://doi.org/10.5555/3524938.3525333).
- [31] C. Winkler, D. E. Worrall, E. Hoogeboom, M. Welling, Learning Likelihoods with Conditional Normalizing Flows, arXiv preprint arXiv:1912.00042 (2019). [doi:10.48550/arXiv.1912.00042](https://doi.org/10.48550/arXiv.1912.00042).
- [32] L. Dinh, J. Sohl-Dickstein, S. Bengio, Density Estimation Using Real NVP, in: *Proceedings of the 2017 International Conference on Learning Representations*, Toulon, France, 2017, p. Poster.
- [33] P. L. Green, K. Worden, Bayesian and Markov chain Monte Carlo methods for identifying nonlinear systems in the presence of uncertainty, *Philosophical Transactions of the Royal Society A* 373 (2015). [doi:10.1098/rsta.2014.0405](https://doi.org/10.1098/rsta.2014.0405).
- [34] W. R. Gilks, S. Richardson, D. Spiegelhalter (Eds.), *Markov Chain Monte Carlo in Practice*, Chapman & Hall/CRC, 1998. [doi:10.1201/b14835](https://doi.org/10.1201/b14835).
- [35] C. P. Robert, G. Casella, *Monte Carlo statistical methods*, Vol. 2, Springer, 1999.

- [36] W. K. Hastings, Monte Carlo Sampling Methods Using Markov Chains and Their Applications, *Biometrika* 57 (1) (1970) 97–109. [doi:10.2307/2334940](https://doi.org/10.2307/2334940).
- [37] C. J. F. ter Braak, A markov chain monte carlo version of the genetic algorithm differential evolution: easy bayesian computing for real parameter spaces, *Statistics and Computing* 16 (3) (2006) 239–249. [doi:10.1007/s11222-006-8769-1](https://doi.org/10.1007/s11222-006-8769-1).
- [38] PyMC Developers, PyMC: Bayesian Modeling in Python (2025).
- [39] D. Wegmann, C. Leuenberger, L. Excoffier, Efficient Approximate Bayesian Computation Coupled With Markov Chain Monte Carlo Without Likelihood, *Genetics* 182 (4) (2009) 1207–1218. [doi:10.1534/genetics.109.102509](https://doi.org/10.1534/genetics.109.102509).
- [40] M. Sunnåker, A. G. Busetto, E. Numminen, J. Corander, M. Foll, C. Dessimoz, Approximate Bayesian Computation, *PLoS Computational Biology* 9 (1) (2013) e1002803. [doi:10.1371/journal.pcbi.1002803](https://doi.org/10.1371/journal.pcbi.1002803).
- [41] I. Kobyzev, S. J. Prince, M. A. Brubaker, Normalizing Flows: An Introduction and Review of Current Methods, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43 (2021) 3964–3979. [doi:10.1109/TPAMI.2020.2992934](https://doi.org/10.1109/TPAMI.2020.2992934).
- [42] D. P. Kingma, P. Dhariwal, Glow: Generative flow with invertible 1×1 convolutions, in: *Proceedings of the 2018 Advances in Neural Information Processing Systems*, Vol. 31, Curran Associates, Inc., Montréal, Canada, 2018, pp. 10236–10245.
- [43] R. Prenger, R. Valle, B. Catanzaro, WaveGlow: A Flow-based Generative Network for Speech Synthesis, in: *Proceedings of the 2019 IEEE International Conference on Acoustics, Speech and Signal Processing*, IEEE, Brighton, UK, 2019, pp. 3617–3621. [doi:10.1109/ICASSP.2019.8683143](https://doi.org/10.1109/ICASSP.2019.8683143).
- [44] L. Dinh, D. Krueger, Y. Bengio, NICE: Non-linear Independent Components Estimation, arXiv preprint arXiv:1410.8516 (2015). [doi:10.48550/arXiv.1410.8516](https://doi.org/10.48550/arXiv.1410.8516).
- [45] Z. Xiao, Q. Yan, Y. Amit, A Method to Model Conditional Distributions with Normalizing Flows, arXiv preprint arXiv:1911.02052 (2019). [doi:10.48550/arXiv.1911.02052](https://doi.org/10.48550/arXiv.1911.02052).
- [46] D. P. Kingma, M. Welling, Auto-Encoding Variational Bayes, arXiv preprint arXiv:1312.6114 (2013). [doi:10.48550/arXiv.1312.6114](https://doi.org/10.48550/arXiv.1312.6114).
- [47] S. Kullback, R. A. Leibler, On Information and Sufficiency, *The Annals of Mathematical Statistics* 22 (1951) 79–86. [doi:10.1214/aoms/1177729694](https://doi.org/10.1214/aoms/1177729694).
- [48] F. Negri, redbKIT, version 2.2, <http://redbkit.github.io/redbKIT> (2016).
- [49] M. Alnæs, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, G. N. Wells, The FEniCS Project Version 1.5, *Archive of Numerical Software* 3 (2015) 9–23. [doi:10.11588/ANS.2015.100.20553](https://doi.org/10.11588/ANS.2015.100.20553).
- [50] F. Chollet, K. Contributors, Keras: Deep learning api for humans (2015). [doi:10.5281/zenodo.14533394](https://doi.org/10.5281/zenodo.14533394).
- [51] M. Torzoni, M. Tezzele, S. Mariani, A. Manzoni, K. E. Willcox, A digital twin framework for civil engineering structures, *Computer Methods in Applied Mechanics and Engineering* 418 (2024) 116584. [doi:10.1016/j.cma.2023.116584](https://doi.org/10.1016/j.cma.2023.116584).
- [52] M. Torzoni, A. Manzoni, S. Mariani, Structural health monitoring of civil structures: A diagnostic framework powered by deep metric learning, *Computers & Structures* 271 (2022) 106858. [doi:10.1016/j.compstruc.2022.106858](https://doi.org/10.1016/j.compstruc.2022.106858).
- [53] M. Torzoni, A. Manzoni, S. Mariani, Enhancing Bayesian model updating in structural health monitoring via learnable mapping, *Engineering with Computers* 41 (2025) 4953–4975. [doi:10.1007/s00366-025-02224-x](https://doi.org/10.1007/s00366-025-02224-x).

- [54] M. Torzoni, A. Manzoni, S. Mariani, A multi-fidelity surrogate model for structural health monitoring exploiting model order reduction and artificial neural networks, *Mechanical Systems and Signal Processing* 197 (2023) 110376. [doi:10.1016/j.ymssp.2023.110376](https://doi.org/10.1016/j.ymssp.2023.110376).
- [55] T. J. Hughes, *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*, Dover Civil and Mechanical Engineering, Dover Publications, 2000.
- [56] A. Quarteroni, A. Manzoni, F. Negri, *Reduced Basis Methods for Partial Differential Equations: An Introduction*, Vol. 92, Springer, Cham, Switzerland, 2015. [doi:10.1007/978-3-319-15431-2](https://doi.org/10.1007/978-3-319-15431-2).
- [57] A. Gelman, D. B. Rubin, Inference from Iterative Simulation Using Multiple Sequences, *Stat Sci* 7 (4) (1992) 457–472. [doi:10.1214/ss/1177011136](https://doi.org/10.1214/ss/1177011136).
- [58] M. B. Lykkegaard, T. J. Dodwell, D. Moxey, Accelerating Uncertainty Quantification of Groundwater Flow Modelling Using a Deep Neural Network Proxy, *Computer Methods in Applied Mechanics and Engineering* 383 (2021) 113895. [doi:10.1016/j.cma.2021.113895](https://doi.org/10.1016/j.cma.2021.113895).
- [59] R. A. Freeze, A Stochastic-Conceptual Analysis of One-Dimensional Groundwater Flow in Nonuniform Homogeneous Media, *Water Resources Research* 11 (1975) 725–741. [doi:10.1029/WR011i005p00725](https://doi.org/10.1029/WR011i005p00725).
- [60] N.-O. Kitterrød, L. Gottschalk, Simulation of Normal Distributed Smooth Fields by Karhunen-Loève Expansion in Combination with Kriging, *Stochastic Hydrology and Hydraulics* 11 (1997) 459–482. [doi:10.1007/BF02428429](https://doi.org/10.1007/BF02428429).
- [61] M. Torzoni, D. Maisto, A. Manzoni, F. Donnarumma, G. Pezzulo, A. Corigliano, Active Digital Twins via Active Inference, arXiv preprint arXiv:2506.14453 (2025). [doi:10.48550/arXiv.2506.14453](https://doi.org/10.48550/arXiv.2506.14453).
- [62] D. Kingma, J. Ba, Adam: A Method for Stochastic Optimization, in: *Int Conf Learn Represent*, Vol. 3, 2015, pp. 1–13.

Appendix A. Implementation details of the deep learning models

In this Appendix, we describe the implementation details of the deep learning models employed in Sec. 3 to validate the proposed inference framework. The considered two case studies deal with the structural health monitoring of a railway bridge and the identification of a spatially varying transmissivity field in a groundwater flow problem. Although these problems involve different physical systems and observation models, they are both addressed using a common learning framework that combines latent-variable representation with conditional density estimation.

The architecture, hyperparameters, and training settings of the informed-VAE — composed of the encoder, decoder, and latent-space predictor — are summarized in Tab. A.3 for the railway bridge case study and in Tab. A.4 for the groundwater flow case study. The encoder consists of a stack of convolutional layers interleaved with max-pooling and dropout operations, followed by fully connected layers that map the input data onto a latent space of dimension N_h . The decoder mirrors the encoding process through fully connected layers and transposed convolutional layers combined with upsampling operations, reconstructing the input from the latent representation. The latent-space predictor consists of three fully-connected layers.

The RealNVP-based conditional normalizing flow architectures, together with the corresponding hyperparameters and training settings, are reported in Tab. A.5 and Tab. A.6 for the railway bridge case and groundwater flow case studies, respectively. The scaling and translation subnetworks, as well as the conditioning network, are implemented as fully connected neural networks.

Training is performed by minimizing either $\mathcal{L}_{\text{i-VAE}}$ or \mathcal{L}_{CNF} for the informed-VAE and the conditional normalizing flow, respectively, using the Adam optimizer [62]. In both cases, the learning rate is reduced over the last 4/5 of the total training steps according to a cosine decay schedule. The optimization uses an 80 : 20 split of the dataset for training and validation, employing an early stopping strategy to prevent overfitting.

Table A.3: Railway bridge - informed-VAE: (a) employed architecture and (b) selected hyperparameters and training settings.

(a)

Layer	Type	Output shape	Activ.	Input layer
Encoder				
0	Input	(B_{VAE}, L, N_s)	None	None
1	Conv1D	$(B_{\text{VAE}}, L, 32)$	Tanh	0
2	MaxPool1D	$(B_{\text{VAE}}, L/2, 32)$	None	1
3	Dropout	$(B_{\text{VAE}}, L/2, 32)$	None	2
4	Conv1D	$(B_{\text{VAE}}, L/2, 64)$	Tanh	3
5	MaxPool1D	$(B_{\text{VAE}}, L/4, 64)$	None	4
6	Dropout	$(B_{\text{VAE}}, L/4, 64)$	None	5
7	Conv1D	$(B_{\text{VAE}}, L/4, 32)$	Tanh	6
8	MaxPool1D	$(B_{\text{VAE}}, L/8, 32)$	None	7
9	Dropout	$(B_{\text{VAE}}, L/8, 32)$	None	8
10	Flatten	$(B_{\text{VAE}}, 32 \cdot L/8)$	None	9
11	Dense	$(B_{\text{VAE}}, 64)$	Tanh	10
12	Dense	$(B_{\text{VAE}}, 16)$	Tanh	11
13	Dense (μ_θ)	(B_{VAE}, N_h)	None	12
14	Dense ($\log \sigma_\theta^2$)	(B_{VAE}, N_h)	None	12
15	Sampling	(B_{VAE}, N_h)	None	13,14
Decoder				
16	Dense	$(B_{\text{VAE}}, 16)$	Tanh	15
17	Dense	$(B_{\text{VAE}}, 64)$	Tanh	16
18	Dense	$(B_{\text{VAE}}, 32 \cdot L/8)$	Tanh	17
19	Reshape	$(B_{\text{VAE}}, L/8, 32)$	None	18
20	Conv1D ^T	$(B_{\text{VAE}}, L/8, 32)$	Tanh	19
21	UpSampling1D	$(B_{\text{VAE}}, L/4, 32)$	None	20
22	Dropout	$(B_{\text{VAE}}, L/4, 32)$	None	21
23	Conv1D ^T	$(B_{\text{VAE}}, L/4, 64)$	Tanh	22
24	UpSampling1D	$(B_{\text{VAE}}, L/2, 64)$	None	23
25	Dropout	$(B_{\text{VAE}}, L/2, 64)$	None	24
26	Conv1D ^T	$(B_{\text{VAE}}, L/2, 32)$	Tanh	25
27	UpSampling1D	$(B_{\text{VAE}}, L, 32)$	None	26
28	Dropout	$(B_{\text{VAE}}, L, 32)$	None	27
29	Conv1D	(B_{VAE}, L, N_s)	Linear	28
Latent-space Predictor				
30	Dense	$(B_{\text{VAE}}, 64)$	Tanh	15
31	Dense	$(B_{\text{VAE}}, 32)$	Tanh	30
32	Dense	$(B_{\text{VAE}}, N_{\text{par}})$	Linear	31

(b)

Latent dimension:	$N_h = 4$
Conv kernel sizes:	25, 13, 7
Dropout rate:	0.05
Weight initializer:	Xavier
L^2 regularization rate:	10^{-3}
Optimizer:	Adam
Batch size:	$B_{\text{VAE}} = 32$
Initial learning rate:	10^{-3}
Max epochs:	250
Learning schedule:	0.8 cosine decay
Weight decay:	0.05
Early stopping patience:	20 epochs
Train-validation split:	80:20
KL weight:	$\beta_{\text{KL}} = 1$
Prediction weight:	$\beta_{\text{pred}} = 15000$

Table A.4: Groundwater flow - informed-VAE: (a) employed architecture and (b) selected hyperparameters and training settings.

(a)

Layer	Type	Output shape	Activ.	Input layer
Encoder				
0	Input	(B_{VAE}, N_x)	None	None
1	Reshape	$(B_{\text{VAE}}, \sqrt{N_x}, \sqrt{N_x}, 1)$	None	0
2	Conv2D	$(B_{\text{VAE}}, \sqrt{N_x}, \sqrt{N_x}, 16)$	Tanh	1
3	Dropout	$(B_{\text{VAE}}, \sqrt{N_x}, \sqrt{N_x}, 16)$	None	2
4	Conv2D	$(B_{\text{VAE}}, \sqrt{N_x}, \sqrt{N_x}, 32)$	Tanh	3
5	Dropout	$(B_{\text{VAE}}, \sqrt{N_x}, \sqrt{N_x}, 32)$	None	4
6	Conv2D	$(B_{\text{VAE}}, \sqrt{N_x}, \sqrt{N_x}, 16)$	Tanh	5
7	Dropout	$(B_{\text{VAE}}, \sqrt{N_x}, \sqrt{N_x}, 16)$	None	6
8	Flatten	$(B_{\text{VAE}}, 16N_x)$	None	7
9	Dense	$(B_{\text{VAE}}, 32)$	Tanh	8
10	Dropout	$(B_{\text{VAE}}, 32)$	None	9
11	Dense (μ_θ)	(B_{VAE}, N_h)	None	10
12	Dense ($\log \sigma_\theta^2$)	(B_{VAE}, N_h)	None	10
13	Sampling	(B_{VAE}, N_h)	None	11,12
Decoder				
14	Input	(B_{VAE}, N_h)	None	13
15	Dense	$(B_{\text{VAE}}, 32)$	Tanh	14
16	Dense	$(B_{\text{VAE}}, 16N_x)$	Tanh	15
17	Reshape	$(B_{\text{VAE}}, \sqrt{N_x}, \sqrt{N_x}, 16)$	None	16
18	Conv2D ^T	$(B_{\text{VAE}}, \sqrt{N_x}, \sqrt{N_x}, 32)$	Tanh	17
19	Dropout	$(B_{\text{VAE}}, \sqrt{N_x}, \sqrt{N_x}, 32)$	None	18
20	Conv2D ^T	$(B_{\text{VAE}}, \sqrt{N_x}, \sqrt{N_x}, 16)$	Tanh	19
21	Dropout	$(B_{\text{VAE}}, \sqrt{N_x}, \sqrt{N_x}, 16)$	None	20
22	Conv2D ^T	$(B_{\text{VAE}}, \sqrt{N_x}, \sqrt{N_x}, 1)$	Tanh	21
23	Dropout	$(B_{\text{VAE}}, \sqrt{N_x}, \sqrt{N_x}, 1)$	None	22
24	Flatten	(B_{VAE}, N_x)	None	23
25	Dense	(B_{VAE}, N_x)	Linear	24
Latent-space Predictor				
26	Input	(B_{VAE}, N_h)	None	13
27	Dense	$(B_{\text{VAE}}, 64)$	Tanh	26
28	Dropout	$(B_{\text{VAE}}, 64)$	None	27
29	Dense	$(B_{\text{VAE}}, 32)$	Tanh	28
30	Dense	$(B_{\text{VAE}}, N_{\text{par}})$	Linear	29

(b)

Latent dimension:	$N_h = 20$
Conv kernel sizes:	3×3
Dropout rate:	0.2
Weight initializer:	Xavier
L^2 regularization rate:	10^{-3}
Optimizer:	Adam
Batch size:	$B_{\text{VAE}} = 32$
Initial learning rate:	10^{-3}
Max epochs:	250
Learning schedule:	0.8 cosine decay
Weight decay:	0.05
Early stopping patience:	25 epochs
Train-validation split:	80:20
KL weight:	$\beta_{\text{KL}} = 2.5^{-4}$
Prediction weight:	$\beta_{\text{pred}} = 10^{-4}$

Table A.5: Railway bridge - conditional normalizing flow: (a) employed architecture and (b) selected hyperparameters and training settings.

(a)

Layer	Type	Output shape	Activ.	Input layer
0	Input (\mathbf{h})	(B_{NF}, N_h)	None	None
1	Input ($\boldsymbol{\lambda}$)	$(B_{\text{NF}}, N_{\text{par}})$	None	None
Parameter conditioning network				
2	Dense	$(B_{\text{NF}}, 32)$	ReLU	1
3	Dense	$(B_{\text{NF}}, 64)$	ReLU	2
4	Dense	$(B_{\text{NF}}, 128)$	ReLU	3
Translation network				
5	Dense	$(B_{\text{NF}}, 64)$	ReLU	0
6	Dropout	$(B_{\text{NF}}, 64)$	None	5
7	Dense	$(B_{\text{NF}}, 128)$	ReLU	4
8	Concatenate	$(B_{\text{NF}}, 192)$	None	6,7
9	Dense	$(B_{\text{NF}}, 64)$	ReLU	8
10	Dense	(B_{NF}, N_h)	Linear	9
Scaling network				
11	Dense	$(B_{\text{NF}}, 64)$	ReLU	0
12	Dropout	$(B_{\text{NF}}, 64)$	None	11
13	Dense	$(B_{\text{NF}}, 128)$	ReLU	4
14	Concatenate	$(B_{\text{NF}}, 192)$	None	12,13
15	Dense	$(B_{\text{NF}}, 64)$	ReLU	14
16	Dense	(B_{NF}, N_h)	Tanh	15

(b)

Latent dimension:	$N_h = 4$
Parameters dimension:	$N_{\text{par}} = 6$
Coupling layers:	$N_f = 16$
Masking strategy:	$N_h/2$ frozen components (alternating)
Dropout rate:	0.1
Weight initializer:	Xavier
L^2 regularization rate:	10^{-3}
Optimizer:	Adam
Batch size:	$B_{\text{NF}} = 32$
Initial learning rate:	10^{-3}
Max epochs:	250
Learning schedule:	0.8 cosine decay
Weight decay:	0.05
Early stopping patience:	15 epochs
Train-validation split:	80:20

Table A.6: Groundwater flow - conditional normalizing flow: (a) employed architecture and (b) selected hyperparameters and training settings.

(a)

Layer	Type	Output shape	Activ.	Input layer
0	Input (\mathbf{x})	(B_{NF}, N_h)	None	None
1	Input ($\boldsymbol{\lambda}$)	$(B_{\text{NF}}, N_{\text{par}})$	None	None
Parameter conditioning network				
2	Dense	$(B_{\text{NF}}, 32)$	ELU	1
3	Dropout	$(B_{\text{NF}}, 32)$	None	2
4	Dense	$(B_{\text{NF}}, 64)$	ELU	3
5	Dropout	$(B_{\text{NF}}, 64)$	None	4
6	Dense	$(B_{\text{NF}}, 128)$	ELU	5
Translation network				
7	Dense	$(B_{\text{NF}}, 64)$	ELU	0
8	Dropout	$(B_{\text{NF}}, 64)$	None	7
9	Dense	$(B_{\text{NF}}, 64)$	ELU	6
10	Concatenate	$(B_{\text{NF}}, 128)$	None	8,9
11	Dense	$(B_{\text{NF}}, 32)$	ELU	10
12	Dense	(B_{NF}, N_h)	Linear	11
Scaling network				
13	Dense	$(B_{\text{NF}}, 64)$	ELU	0
14	Dropout	$(B_{\text{NF}}, 64)$	None	13
15	Dense	$(B_{\text{NF}}, 64)$	ELU	6
16	Concatenate	$(B_{\text{NF}}, 128)$	None	14,15
17	Dense	$(B_{\text{NF}}, 32)$	ELU	16
18	Dense	(B_{NF}, N_h)	Tanh	17

(b)

Latent dimension:	$N_h = 20$
Parameters dimension:	$N_{\text{par}} = 14$
Coupling layers:	$N_f = 30$
Masking strategy:	$N_h/2$ frozen components (alternating)
Dropout rate:	0.2
Weight initializer:	Xavier
L^2 regularization:	$\lambda = 10^{-3}$
Optimizer:	Adam
Batch size:	$B_{\text{NF}} = 32$
Initial learning rate:	10^{-3}
Max epochs:	200
Learning schedule:	0.8 cosine decay
Weight decay:	0.05
Early stopping patience:	20 epochs
Train-validation split:	80:20

MOX Technical Reports, last issues

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

- 17/2026** Caldera, L.; Bottacini, G.; Cavinato, L.
MAGIC-Flow: multiscale adaptive conditional flows for generation and interpretable classification
- 20/2026** Caldera, L.; Cavinato, L.; Cirone, A.; Cama, I.; Garbarino, S.; Lodi, R.; Tagliavini, F.; Nigri, A.; De Francesco, S.; Cappozzo, A.; Piana, M.; Ieva, F.;
DISARM++: Beyond scanner-free harmonization
- 19/2026** Caldera, L.; Cavinato, L.; Ieva, F.
Scanner-agnostic MRI harmonization via SSIM-guided disentanglement
- 15/2026** Zecchi, A. A.; Sanavio, C.; Cappelli, L.; Perotto, S.; Roggero, A.; Succi, S.
Block encoding of sparse matrices with a periodic diagonal structure
- 14/2026** Agasisti, T.; Cannistrà, M.; Paganoni, A.M.
Nudging communication for students at risk: experimental evidence from an Italian university
- 13/2026** Dimola, N.; Coclite, A.; Zunino, P.
Neural Preconditioning via Krylov Subspace Geometry
- 12/2026** Corbetta A; Logan K.M.; Ferro M; Zuccolo L; Perola M.; Ganna A.; Di Angelantonio E.;Ieva F.
Longitudinal patterns of statin adherence and factors associated with decline in over one million individuals in Finland and Italy
- 11/2026** Cicalese, G.; Ciaramella, G.; Mazzieri, I.; Gander, M. J.
Optimized Schwarz Waveform Relaxation for the Damped Wave Equation
- 10/2026** Dimola, N.; Franco, N. R.; Zunino, P.
Numerical Solution of Mixed-Dimensional PDEs Using a Neural Preconditioner
- 09/2026** Manzoni, V.; Ieva, F.;Larranaga, A.C.; Vetrano, D.L.; Gregorio, C.
Hidden multistate models to study multimorbidity trajectories