

Gaussian process regression + deep neural network autoencoder for probabilistic surrogate modeling in nonlinear mechanics of solids

Saurabh Deshpande^a, Hussein Rappel^b, Mark Hobbs^c, Stéphane P.A. Bordas^a, Jakub
Lengiewicz^{a,d,e}

^a*Department of Engineering; Faculty of Science, Technology and Medicine; University of Luxembourg*

^b*Department of Engineering, Faculty of Environment, Science and Economy, University of Exeter, UK*

^c*Future Methods, Rolls-Royce plc, Derby, UK*

^d*Institute of Fundamental Technological Research, Polish Academy of Sciences, Poland*

^e*Luxembourg Institute of Science and Technology, Esch-sur-Alzette, Luxembourg*

Abstract

Many real-world applications demand accurate and fast predictions, as well as reliable uncertainty estimates. However, quantifying uncertainty on high-dimensional predictions is still a severely under-investigated problem, especially when input-output relationships are non-linear. To handle this problem, the present work introduces an innovative approach that combines autoencoder deep neural networks with the probabilistic regression capabilities of Gaussian processes. The autoencoder provides a low-dimensional representation of the solution space, while the Gaussian process is a Bayesian method that provides a probabilistic mapping between the low-dimensional inputs and outputs. We validate the proposed framework for its application to surrogate modeling of non-linear finite element simulations. Our findings highlight that the proposed framework is computationally efficient as well as accurate in predicting non-linear deformations of solid bodies subjected to external forces, all the while providing insightful uncertainty assessments.

Keywords: Surrogate Modeling, Deep Neural Networks, Gaussian Process, Autoencoders, Uncertainty Quantification, Finite Element Method

1. Introduction

In computational mechanics, high-fidelity modeling methods such as the finite element method (FEM) are standard for achieving detailed insights into the mechanical behaviors of materials

*Corresponding author

Email address: jleng@ippt.pan.pl (Jakub Lengiewicz)

and structures. Despite their capability to provide accurate predictions for complex non-linear problems, these methods are often hindered by the substantial computational effort they demand, thereby constraining their applicability in domains requiring real-time solutions. As a remedy, various approximate methods, such as surrogate modeling, have emerged as an indispensable tool to alleviate the computational burden associated with traditional simulations. By providing the balance between computational efficiency and the accuracy, these methods play a crucial role in scenarios requiring fast- or real-time responses.

In conventional methods like FEM, numerical challenges primarily arise from two main factors: the non-linearities present in the problem and the high number of unknowns (degrees of freedom) required to attain the desired level of prediction accuracy. The majority of approximate methods concentrate on addressing the latter aspect by directly reducing the number of unknowns of the problem. For instance, the conventional reduced order modelling (ROM) methods accomplish this by linearly projecting higher-order information to lower-dimensional space using dimensionality reduction techniques such as Principal Component Analysis (PCA) [1, 2], and Proper Orthogonal Decomposition (POD) [3, 4]. However, these approaches yield undesirable results when simulating highly nonlinear phenomena [5–7]. In response to the linearity limitation, numerous nonlinear dimensionality reduction techniques have been proposed [8, 9]. However, obtaining real-time simulations with these methods, in particular when considering high degree of non-linearity, remains a challenge [10].

Rapid advancements in machine learning has naturally fueled developments to approximate modelling. There has been a widespread adoption of machine learning techniques to accelerate computationally intensive numerical methods [11–15]. In particular, deep learning (DL) models have been successfully used as accurate and computationally inexpensive surrogates to solve non-linear problems in mechanics [16–20]. In context of ROM, DL methods are being extensively used to find *non-linear* reduced order representations of the high dimensional data [21, 22]. At the same time, they are computationally inexpensive at the prediction stage, hence, they are highly suitable in a wide area of applications requiring real time simulations [23, 24].

This work focuses on another important aspect, which is the capabilities of surrogate models to quantify uncertainty. Providing reliable uncertainty estimates alongside predictions is needed in various real-life applications, such as surgical simulations [25] or autonomous driving [26], to reduce risks of harmful consequences in these crucial tasks. There exist many machine learning approaches allowing to quantify model- and data uncertainties. In deep learning, the most common approach is to use stochastic deep neural network architectures with relevant suitable training procedures. These are, for instance, [27–29]. They are known to be capable of providing uncertainty information, however, they can be computationally much more intensive as compared to their deterministic versions. Another class of approaches, employed specifically in this work, are probabilistic models that are based on Gaussian process regression [30, 31] (GP regression or GPR). The excellent predictive power and intrinsic capability of GPRs to quantify uncertainties make them a perfect candidate for the application of this contribution. Indeed, GPs have been employed also in mechanics in a wide range of studies both as surrogate models [32–35] and as an approach to model spatially varying parameter fields [36–38]. Ding et al. [39] employed a combination of GPR and PCA to model displacement fields in problems with nonlinear materials (i.e., elastoplastic material), [40] and [41] used GPR to assess uncertainties in the full-field measurement experimental context.

Despite being a powerful technique for constructing probabilistic surrogate models, GPR is known for its poor computational scaling with respect to number data points [42]. The scalability issue is substantially more significant for multioutput GPs (MOGP), where surrogate models gives multiple outputs simultaneously (i.e., $\mathcal{O}(N^3M^3)$, where N and M denote the number of data points and the dimension of the multiple output, respectively [43]. Numerous approximated techniques have been developed to significantly reduce this very poor scaling. The primary effort was to reduce the scaling with respect to N , making the GPR more suitable for big data, see [42] for a more extensive review. For the multi-output GPs, the most common simplification is to apply independent GPs for each dimension of the output (Independent Multi-Output Models), which reduces the scaling to N^3M .

The framework proposed in this contribution focuses on alleviating part of the scaling issue by directly reducing the dimension of outputs, M . This is achieved by applying the compressed latent representation of outputs via, so called, autoencoder deep neural networks—a strategy akin to approaches found in related works [44, 45]. Autoencoder neural networks, as an unsupervised learning technique, have been commonly adopted to efficiently learning compressed representations across various domains [46–48], also in the context of computational fluid dynamics [49, 50] and computational solid mechanics [51, 52]. The primary advantage of autoencoder DNNs is their ability to reduce the dimensionality of data in a non-linear way [53]. As such, they perform much better than POD or PCA when high non-linearities are present in data, see, e.g., the study [54]. Despite the known drawback of DNNs vs. PCA regarding their reduced interpretability of latent representations, in the present work we use DNNs as an integral component of the proposed probabilistic framework. The reason for that is to provide a more general framework, capable of handling a wider class of datasets.

In the present work, we combine GPR with autoencoder DNNs to provide surrogate model for force-displacement predictions in mechanics of solids. In the proposed framework, we reduce the high-dimensional 2D/3D finite element solutions (full-field displacement) into their reduced representation by applying autoencoder neural networks, and then use GPR to provide probabilistic mapping between the input forces and the reduced-order displacements. During the prediction phase, for a given unseen force, GPR predicts the reduced-order (latent space) displacement which is then projected to the full field using the decoder part of autoencoder network. Such general idea to combine GPR with autoencoder DNNs is not new, and has been recently utilized in different applications and contexts. For instance, in the context of high-dimensional outputs, similar strategy has been applied to the spatiotemporal climate modeling [45], in which the full-field climate data is compressed using a DNN autoencoder while GP is used to forecast weather one step ahead. In the context of mechanics, in the works [55, 56], techniques combining GPR and autoencoders have been developed for reliability analysis, and demonstrated for the problem of predicting structural compliance of 2D beams. In the present work, we extend the application to non-linear deformation of 3D solids with unstructured meshes. We emphasize on studying the contribution of errors by individual components (GP and autoencoder), and we provide interpretation of uncertainties in relation to the extrapolated region. We demonstrate the robustness and versatility of our framework by applying it to synthetic datasets generated from the non-linear finite element simulations. Additionally, the source codes and datasets utilized in this work are provided in the following repository: <https://github.com/saurabhdeshpande93/gp-auto-regression>.

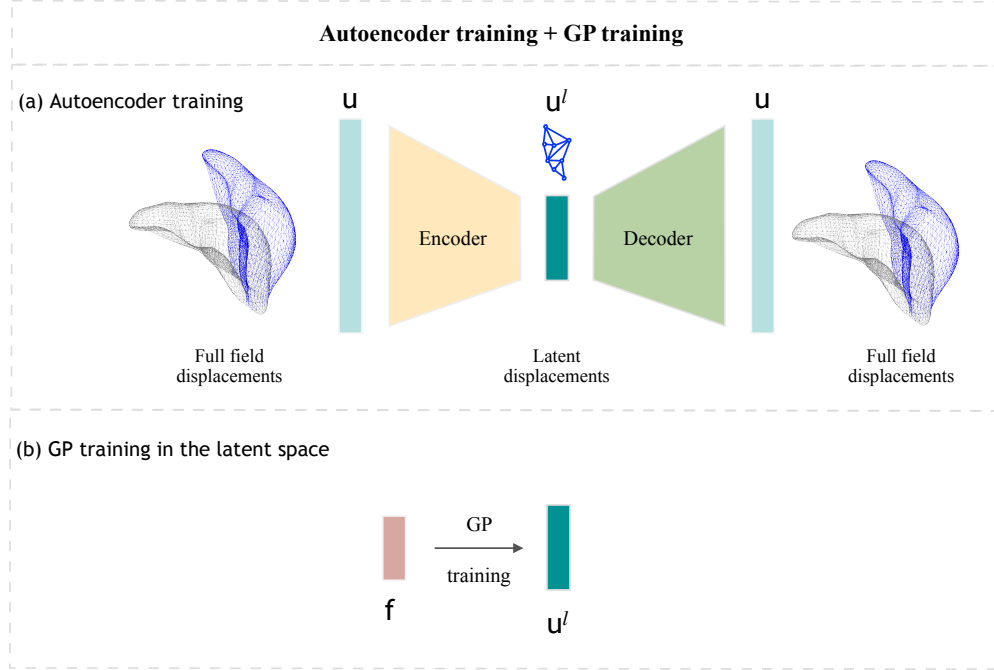


Figure 1: Two-stage training. (a) First, the autoencoder neural network is trained to compress full field displacement data to corresponding latent space representations. (b) Afterwards, the GP model is trained to provide the probabilistic force-displacement mapping in the latent space.

The remainder of this paper is organized as follows. In Section 2 we introduce the surrogate modeling framework, and describe its constituent parts. In Section 3 we present two benchmark examples and study the performance of the framework. The conclusions and future research directions are outlined in Section 4.

2. Methodology

As outlined in the introduction, in this work, we focus on the problem of predicting non-linear deformation of solids under external loads. The goal is to introduce surrogate machine learning framework that is capable of accurate and fast predictions of full field displacements along with the associated uncertainties. The core idea is to combine the autoencoder deep neural network with the Gaussian Process Regression, see Section 2.1 and Section 2.2, respectively. The training of both models is based on the numerically generated force-displacement dataset

$$\mathcal{D}^f = \{(\mathbf{f}_i, \mathbf{u}_i)\}_{i=1,\dots,N}, \quad (1)$$

which is obtained with the classical FEM, see Section 2.4.

The training of the two constituent models is performed sequentially in two stages, see Figure 1. In the first stage, Figure 1a, the autoencoder network is trained only on full-field displacement data to get the corresponding compressed representations, \mathbf{u}^l . In the second stage, Figure 1b, the GP regression model is trained between the forces, \mathbf{f} , and the latent displacements vector \mathbf{u}^l .




Figure 2: For an unseen input force \mathbf{f}^* , the GP is used to predict the latent displacement distribution (i.e., our uncertainty about the displacements is described by a probability distribution). Subsequently, these latent displacements are projected to the full-field displacement distribution by using the decoder component of the autoencoder network.

Following the training, the GPR model is stacked with the Decoder part of the autoencoder network, which provides the probabilistic force-full-displacement model, see also Section 2.3. Specifically, for a given unseen input force, GPR model predicts the latent displacement distribution. These latent displacements are then reconstructed to the full space using the decoder part of the autoencoder network, see Figure 2.

2.1. Obtaining latent representations with autoencoder neural networks

Autoencoders are neural networks that are primarily used for unsupervised learning and dimensionality reduction tasks [53]. They are designed to encode high-dimensional input data into a lower-dimensional latent space representation, and then decode it to the original input space, aiming to reconstruct the input data as accurately as possible.

The architecture of an autoencoder network typically consists of two main components: an encoder and a decoder. The encoder takes the input data and maps it to a compressed representation in the latent space, while the decoder takes this compressed representation and reconstructs the original input data. The encoder and decoder are usually symmetrical, meaning that they have the same number of layers and the layer sizes are mirrored. One of the key aspects of autoencoders is that the latent space, also referred to as the bottleneck layer, has a lower dimensionality than the input space. This forces the network to learn a more efficient and meaningful representation of the data. By compressing the data into a lower-dimensional space and then reconstructing it, autoencoders can capture the most salient features by discarding less relevant information.

The autoencoder network, \mathcal{U} , is constructed by subsequently applying the encoder and decoder networks

$$\mathcal{U}(\mathbf{u}, \boldsymbol{\theta}_{\text{auto}}) = \mathcal{U}_{\text{decoder}}(\mathcal{U}_{\text{encoder}}(\mathbf{u}, \boldsymbol{\theta}_{\text{auto}}), \boldsymbol{\theta}_{\text{auto}}), \quad (2)$$

where \mathbf{u} stands for the full field displacements and $\mathbf{u}^l = \mathcal{U}_{\text{encoder}}(\mathbf{u}, \boldsymbol{\theta}_{\text{auto}})$ is their corresponding latent/compressed representation. For a given full field displacement dataset, $\{\mathbf{u}_i\}_{i=1,\dots,N}$, the autoencoder network is trained by minimizing the following mean squared error loss

$$\mathcal{L}(\mathcal{D}^f, \boldsymbol{\theta}_{\text{auto}}) = \frac{1}{N} \sum_{i=1}^N \|\mathcal{U}(\mathbf{u}_i, \boldsymbol{\theta}_{\text{auto}}) - \mathbf{u}_i\|_2^2, \quad (3)$$

where $\boldsymbol{\theta}_{\text{auto}}$ are trainable parameters of the autoencoder network. The optimal parameters, $\boldsymbol{\theta}_{\text{auto}}^*$, are retrieved by minimizing the loss function:

$$\boldsymbol{\theta}_{\text{auto}}^* = \arg \min_{\boldsymbol{\theta}_{\text{auto}}} \mathcal{L}(\mathcal{D}^f, \boldsymbol{\theta}_{\text{auto}}). \quad (4)$$

Therefore, once the autoencoder network is trained, for a given full field displacement field \mathbf{u} , the corresponding latent representation \mathbf{u}^l is computed by applying the encoder part of the autoencoder network:

$$\mathbf{u}^l = \mathcal{U}_{\text{encoder}}(\mathbf{u}, \boldsymbol{\theta}_{\text{auto}}^*). \quad (5)$$

Hence, we obtain latent state dataset $\mathcal{D}^l = \{(\mathbf{f}_i, \mathbf{u}_i^l)\}_{i=1,\dots,N}$, from the full field dataset $\mathcal{D}^f = \{(\mathbf{f}_i, \mathbf{u}_i)\}_{i=1,\dots,N}$. Note that forces are identical for both datasets and they are always provided as smaller dimensional arrays, which is discussed in detail in Section 3.1. In the next section, we will use the latent dataset, \mathcal{D}^l , to elaborate Gaussian process regression model, which will provide the transformation between the input forces, \mathbf{f} , and the latent solutions, \mathbf{u}^l .

Remark: In our study, we employ two distinct autoencoder architectures, see Section 3.2.1. The first type utilizes fully connected networks, suitable for handling arbitrary unstructured meshes. Meanwhile, the second type utilizes convolutional neural networks (CNNs) specifically designed for structured mesh scenarios [10, 57].

2.2. Gaussian Process Regression in the latent space

As discussed in the introduction, Gaussian Process Regression (GPR) faces scalability challenges as the number of data points, N , and the dimension of outputs, M , increase. In this work we focus on reducing the bottleneck related to M , which is particularly important for problems with a large number of degrees of freedom, such as the examples studied in Section 3. As explained earlier, the main idea behind reducing M is to project the high-dimensional displacements (of dimension M) to a significantly lower-dimensional space, $L \ll M$, and then employ GPR between the input forces and the latent representation of displacements. In addition to that, in this work we follow a common practice to implement independent GPs for each dimension of the output, assuming that all output components are independent of each other. This approach significantly reduces the $\mathcal{O}(N^3 L^3)$ complexity of a multi-output GPR problem to L independent problems of $\mathcal{O}(N^3)$ complexity each (i.e., $\mathcal{O}(N^3 L)$). Because our framework uses independent GPRs for each latent output, in this subsection, we provide a brief introduction to single-output GPs. For more details, interested readers are referred to [30].

GPs are an extension of multivariate normal distributions into an infinite-dimensional Gaussian distribution [58]. Consider two data points $(\mathbf{f}, w(\mathbf{f}))$, $(\mathbf{f}', w(\mathbf{f}'))$, a GP is completely specified by its

mean function and covariance function. We define mean function $m(\mathbf{f})$ and the covariance function $k(\mathbf{f}, \mathbf{f}')$ of a real process $w(\mathbf{f})$ as:

$$\begin{aligned} m(\mathbf{f}) &= \mathbb{E}(w(\mathbf{f})), \\ k(\mathbf{f}, \mathbf{f}') &= \mathbb{E}((w(\mathbf{f}) - m(\mathbf{f}))(w(\mathbf{f}') - m(\mathbf{f}'))). \end{aligned} \quad (6)$$

As mentioned, a GP is used to describe a distribution over functions. Therefore, a realization of a GP reads as:

$$w(\mathbf{f}) \sim \mathcal{GP}(m(\mathbf{f}), k(\mathbf{f}, \mathbf{f}')), \quad (7)$$

where \mathbf{f} and \mathbf{f}' are two GP input vectors of dimension D , which are force vectors in the scope of the present framework. Hence, in general, for a finite collection of inputs $\mathbf{F} = [\mathbf{f}_1, \dots, \mathbf{f}_N]$:

$$\mathbf{w} \sim \mathcal{N}(\mathbf{m}, \mathbf{K}), \quad (8)$$

where $\mathbf{w} = [w_1 \ \dots \ w_N]^T$, $\mathbf{m} = [m(\mathbf{f}_1) \ \dots \ m(\mathbf{f}_N)]^T$ and \mathbf{K} denotes the covariance matrix between two inputs, i.e., $(\mathbf{K})_{ij} = k(\mathbf{f}_i, \mathbf{f}_j)$. Frequent choices for the covariance function are given in [30].

In simple terms, in GPR, we set a GP (Eq.(7)) as a prior for the function that maps inputs to outputs and then update this prior with available observations using Bayes' rule (see [59]). In our framework, the GP predicts displacements corresponding to latent space representations, as shown in Figures 1 and 2. Therefore, we will use the superscript- l notation, $(\cdot)^l$, for all GP outputs.

Since the proposed framework uses independent GPRs for latent outputs, we will focus on a single latent output to explain the GP setup. Let \tilde{u}_i^l be an observation for a particular latent output contaminated by Gaussian noise, i.e., $\tilde{u}_i^l \sim \mathcal{N}(w(\mathbf{f}_i), \sigma^2)$, with $i = 1, \dots, N$. Without loss of generality, we assume that the function mapping inputs to the output is a realization of a GP with zero mean and a covariance function $k(\mathbf{f}, \mathbf{f}')$, defined by its parameter set $\boldsymbol{\theta}_{\text{GP}}$. Furthermore, let $\mathbf{y} = [\tilde{u}_1^l, \dots, \tilde{u}_N^l]$. The predictions \tilde{u}^{l*} of a GPR for a new data point input \mathbf{f}^* will be as follows:

$$\tilde{u}^{l*} | \boldsymbol{\theta}_{\text{GP}}, \sigma^2, \tilde{\mathbf{u}}^l, \mathbf{F}, \mathbf{f}^* \sim \mathcal{N}(\mathbb{E}(\tilde{u}^{l*}), \mathbb{V}(\tilde{u}^{l*})), \quad (9)$$

$$\mathbb{E}(\tilde{u}^{l*}) = \mathbf{k}^{*T}(\mathbf{K} + \sigma^2 \mathbf{I}_N)^{-1} \mathbf{y}, \quad (10)$$

and

$$\mathbb{V}(\tilde{u}^{l*}) = k(\mathbf{f}^*, \mathbf{f}^*) - \mathbf{k}^{*T}(\mathbf{K} + \sigma^2 \mathbf{I}_N)^{-1} \mathbf{k}^*, \quad (11)$$

where $\mathbb{E}(\tilde{u}^{l*})$ denotes mean of the predicted \tilde{u}^{l*} at \mathbf{f}^* , $\mathbb{V}(\tilde{u}^{l*})$ denotes the prediction variance, \mathbf{I}_N is an $N \times N$ identity matrix, \mathbf{K} denotes the $N \times N$ covariance matrix given in Eq. (8), \mathbf{k}^* is an $N \times 1$ vector with $(\mathbf{k}^*)_i = k(\mathbf{f}^*, \mathbf{f}_i)$. Furthermore, to obtain the optimal parameters for GPR, we train the GP by maximizing the following equation:

$$\log p(\mathbf{y} | \mathbf{F}, \boldsymbol{\theta}_{\text{GP}}, \sigma^2) = -\frac{1}{2} \mathbf{y}^T (\mathbf{K} + \sigma^2 \mathbf{I}_n)^{-1} \mathbf{y} - \frac{1}{2} \log[\det(\mathbf{K} + \sigma^2 \mathbf{I}_n)] - \frac{n}{2} \log(2\pi). \quad (12)$$

In Equation (12), $p(\mathbf{y} | \mathbf{F}, \boldsymbol{\theta}_{\text{GP}}, \sigma^2)$ represents the conditional likelihood of the observations. This likelihood depends on the parameter set $\boldsymbol{\theta}_{\text{GP}}$, the input matrix \mathbf{F} , and the noise variance σ^2 . The

optimised parameters $(\boldsymbol{\theta}_{\text{GP}}, \sigma^2)$ are computed by maximising the log likelihood, as discussed in [59]. Note also that the GP and the autoencoder network are trained separately. Autoencoder network needs to be trained first in order to generate the latent state dataset, which is then used to train the GP.

2.3. Projecting latent GP predictions to the full field space using decoder

As depicted in the Figure 2, the latent space predictions obtained using GP are projected to the full space using decoder part of the autoencoder network. At the inference stage, for a new input force \mathbf{f}^* , the GPR first generates latent displacement distribution $p(\tilde{\mathbf{u}}^{*l}|\mathbf{f}^*)$ for each latent output, l , which is a Gaussian distribution with the mean and variance obtained using Eq.(10) and Eq.(11) respectively:

$$p(\tilde{\mathbf{u}}^{*l}|\mathbf{f}^*) = \mathcal{N}(\mathbb{E}_{\text{GP}}, \mathbb{V}_{\text{GP}}). \quad (13)$$

Subsequently, the distributions of full field displacements are obtained by translating samples from the latent predictive distributions through the decoder component. Let the latent vector \mathbf{u}_s^l denote a sample, s , drawn from distributions corresponding to all latent outputs, obtained using independent GPs for an input \mathbf{f}^* . The displacement corresponding to the full field solution for each degree of freedom (DOF) is a Gaussian distribution, whose mean and variance are computed as follows:

$$\begin{aligned} (\tilde{u}_\mu^*)_i &\approx \frac{1}{S} \sum_{s=1}^S (\mathcal{U}_{\text{decoder}}(\mathbf{u}_s^l))_i, \\ (u_\sigma^*)_i^2 &\approx \frac{1}{S-1} \sum_{s=1}^S ((\mathcal{U}_{\text{decoder}}(\mathbf{u}_s^l))_i - (u_\mu^*)_i)^2, \end{aligned} \quad (14)$$

where i represents index of DOF. In order to get the full field displacement distributions, we generate S samples in the latent space, and these samples are then projected to the original space using the decoder. Now the final Gaussian distributions of displacements of the full field solution is represented through the mean $(\tilde{\mathbf{u}}_\mu^*)$ and standard deviation (\mathbf{u}_σ^*) of these S samples. We set $S = 300$ for all the implementations presented in this work.

2.4. Finite element formulation for non-linear deformations of solid bodies

We consider the problem of deformation of a body occupying the domain Ω , that is fixed at a part of its boundary Γ_u (Dirichlet boundary conditions), and loaded over the surface Γ_t (Neumann boundary conditions) and/or by body forces $\bar{\mathbf{b}}$. This boundary value problem is expressed by the virtual work principle

$$\int_{\Omega} \mathbf{P}(\mathbf{F}(\mathbf{u})) \cdot \nabla \delta \mathbf{u} dV - \int_{\Omega} \rho \bar{\mathbf{b}} \cdot \delta \mathbf{u} dV - \int_{\Gamma_t} \bar{\mathbf{t}} \cdot \delta \mathbf{u} dS = 0 \quad \forall \delta \mathbf{u}, \quad (15)$$

where \mathbf{u} and $\delta\mathbf{u}$ belong to appropriate functional spaces, $\mathbf{u} = \bar{\mathbf{u}}$ and $\delta\mathbf{u} = \mathbf{0}$ on Γ_u , and $\mathbf{P}(\mathbf{F})$ is the first Piola-Kirchhoff stress tensor. $\mathbf{P}(\mathbf{F})$ is given by the constitutive relationship through the strain energy potential, $W(\mathbf{F})$, as:

$$\mathbf{P}(\mathbf{F}) = \frac{\partial W(\mathbf{F})}{\partial \mathbf{F}}, \quad (16)$$

where $\mathbf{F} = \mathbf{I} + \nabla\mathbf{u}$ is the deformation gradient tensor. In this work we use following Neo-hookean (hyperelastic) strain energy density:

$$W(\mathbf{F}) = \frac{\mu}{2}(I_c - 3 - 2 \ln J) + \frac{\lambda}{4}(J^2 - 1 - 2 \ln J), \quad (17)$$

where $J = \det(\mathbf{F})$ and $I_c = \text{tr}(\mathbf{F}^T \mathbf{F})$.

Following conventional finite element methodology, the problem represented by Eq. (15) is discretized and posed as a system of nonlinear equations for the unknown nodal displacements, which is then solved iteratively using Newton-Raphson method. As the effect, for an input vector of prescribed external forces, the vector of nodal displacements is computed. The solutions for different input forces provide the necessary force-displacement dataset, in which each element consist of the pair (\mathbf{f}, \mathbf{u}) , where \mathbf{f} encodes the information on external loading and \mathbf{u} is the vector of full-field displacement of FE mesh.

Remark: The vector \mathbf{f} provides a compressed information about the external loading. In our benchmark examples, a body can be subjected to either point loads applied to a single node, or uniform body forces acting throughout the entire volumes. In the 2D example, we apply point loads, and describe the loading \mathbf{f} as its magnitude (f_x, f_y) and the position, d , at which the force is applied (relative to the fixed boundary). In the 3D example, we apply body forces, and the loading \mathbf{f} represents the force density components (b_x, b_y, b_z) . See also more details in Section 3.1.

3. Results

The performance of the framework is demonstrated and studied based on two benchmark examples, as illustrated in Figure 3: the 2D beam example and the 3D liver example. The bodies are subjected to varied loading conditions and undergo large deformations in non-linear regime. The data necessary for training is generated through the finite element simulations, see Section 2.4, and the procedure and data range is described in a more detail in Section 3.1. The finite element simulations are performed with the AceFEM framework [60, 61]. The details of the surrogate framework and its training are provided in Section 3.2. In Sections 3.3-3.6 we study the predictive capabilities of the framework, with a more detailed discussion on sources of errors and on the interpretation of probabilistic predictions. In particular, in Sec. 3.6, we demonstrate the capabilities of our framework to be a convenient indicator of the missing data region.

3.1. Dataset generation information

In the 2D beam example, see Figure 3a, a beam is loaded with point forces of random direction and magnitude, within the ranges described in Table 1. These forces are applied to the nodes located




Figure 3: Schematics of examples considered in this work. (a) 2D beam discretised with quad elements is subjected to point loads on the nodes lying on red line; this example is adapted from [57]. (b) 3D liver is subjected to body forces.

along the top line, which is visually represented by red line in the Figure 3a. For the 2D beam case, displacement data is provided for all degrees of freedom (DOFs) of the mesh, capturing the full-field displacements of the beam. However, the forces used as inputs to the surrogate framework are represented in a compressed format of a 3-component array. Specifically, the applied point load is characterized by two components, namely f_x and f_y , representing the force magnitudes in the x and y directions, respectively. Additionally, the distance, d , from the fixed boundary to the point of application is considered. This information is visually depicted in Figure 3a.

In the 3D liver example, a body representing the liver geometry is subjected to a body force of random direction and magnitude. The liver is also constrained at specific nodes located at the right end, see Figure 3b. Similarly to the beam case, displacement data is provided for all DOFs of the liver mesh, representing the full-field displacements of the organ. The compressed representation of forces consists of the body force density components in the x , y , and z directions, respectively (b_x, b_y, b_z). Further details, such as the range of external forces applied, as well as other information used to generate the datasets are given in the Table 1.

Remark: Note that for each loading case, the data generated by the FE simulation is deterministic. This has been a choice in the scope of the present work, however, it can be extended to the non-deterministic case if needed.

Problem	N.of DOFs (\mathcal{F})	Range (External forces/ body force density)	Young's modulus E [Pa], Poisson's ratio ν , density ρ [kg/m ³]	Dataset size (train+test)
2D beam	128	$f_x, f_y = -2.5$ to 2.5 N	500, 0.4, -	5700 + 300
3D liver	9171	$b_x, b_y, b_z = -0.8$ to 0.8 N/kg	5000, 0.45, 1000	7600 + 400

Table 1: Details of FE datasets.

3.2. Implementation details

3.2.1. Autoencoder networks: Architectures & training details

In this study, we employ two types of autoencoder networks. Firstly, we utilize a CNN-based architecture, illustrated in Figure 4, tailored for the 2D beam example. Secondly, we utilize a fully connected autoencoder network, depicted in Figure 5, implemented for the 3D liver example. Both of these networks consist of an encoder and a decoder component. The primary function of these autoencoder networks is to compress the displacement solutions of the full field space into corresponding latent representations. Note that the forces are initially provided in a simple format, see Sec. 3.1, thus do not need to be compressed by autoencoder networks.

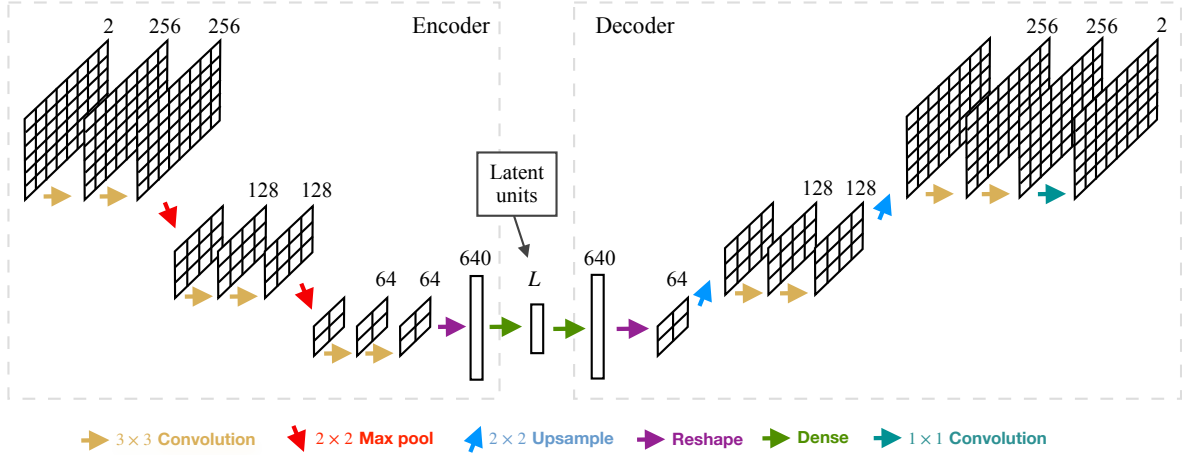


Figure 4: CNN autoencoder architecture for used for the 2D beam case. Input is provided to the network in the structured mesh format, the number at the top of the convolution layer output indicates the number of channels in the respective layer. Input is applied series of convolution and pooling/upsampling layers until the original dimension is retrieved. In the bottleneck level, fully connected layers (dense layer) are applied and the number the top of their output represents the number of units present in it.

Architecture for the 2D beam case

Encoder: The input to the CNN autoencoder is a mesh displacement tensor, which is represented with 2 channels. These channels correspond to the displacements along the x and y DOFs of the mesh. The input tensor undergoes application of two convolutional layers, each with 256 channels and 3×3 filters as shown in Figure 4. Convolutional operations enable feature extraction and dimensionality reduction. Subsequently, a max-pooling operation is applied, which reduces spatial dimensions of the tensors while preserving the number of channels. This process is repeated again with 128 channels, further capturing important patterns in the data. At the latent level of the network, two convolution layers with 64 channels are applied, which is followed by flattening of the tensor. Now the flattened tensor is applied with a dense layer with l units, which stands for the dimension of the compressed/latent state. This compressed representation contains crucial information about the input mesh, capturing the most relevant features.

Decoder: The Decoder takes a flattened tensor of dimension L as its input. Initially, this input tensor undergoes processing via a dense layer. This strategic utilization of the dense layer facilitates a transformation of the arbitrary latent dimensions into a new size that can be conveniently reshaped

into a grid structure, a format that seamlessly aligns with the subsequent convolutional layer operations. For instance, utilizing a dense layer of 640 units as shown in the Figure 4 allows us to reshape the flat tensor to a $64 \times 5 \times 2$ tensor, on which series of upsampling and convolution operations are performed until the original spatial dimensions is achieved. In the end, a 2 channel convolution with 1×1 filter is applied to get back the original mesh shape.

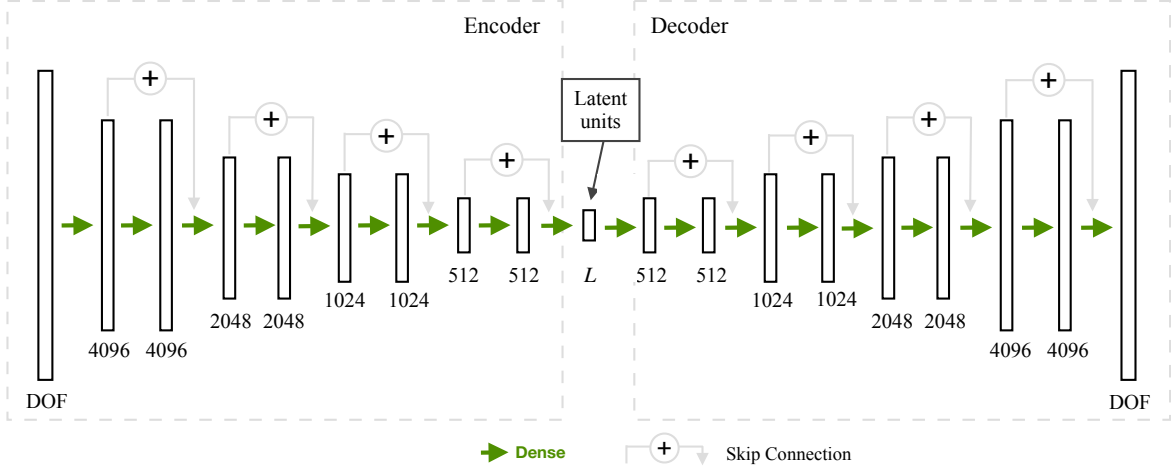


Figure 5: Fully connected autoencoder architecture used for the liver case. Skip connections are indicated with the (+) sign. The number of units for each dense layer output is indicated at the bottom. The latent representation dimension is denoted as L , set to $L = 16$ for the 3D liver case.

Architecture for the 3D liver case

Convolutional neural networks are inherently designed to work with grid-like inputs, such as structured meshes. For that reason, this methodology isn't immediately adaptable to inputs that lack this grid-based nature, such as the unstructured mesh that are present in the liver example. To address this limitation, we introduce an alternative solution in the form of a fully connected autoencoder network as illustrated in Figure 5. This network can handle diverse and unstructured meshes.

Encoder: This component receives input in the form of displacements linked to all degrees of freedom (DOFs) of the mesh. This information is encoded within a flattened tensor, which serves as the input layer. It is further applied with two dense layers with 4096 units each. We also introduce skip connections as illustrated in the Figure 5, which avoid vanishing gradient issues and stabilize the training procedure. This procedure applying blocks of dense layers and skip connections is repeated three more times with 2048, 1024, 512 units for dense layers respectively.

Decoder: Decoder takes the $L = 16$ dimensional 1D tensor as the input (compressed/latent tensor), a similar procedure of applying two dense layers and skip connections is followed until the original size tensor is obtained. The number of units used in each dense layer is detailed in the Figure 5.

Training of autoencoder networks

Autoencoder network is trained on the full field displacement dataset, using the loss function as described by Eq.(3). Minimization is carried out using the Adam optimizer with recommended parameter configurations as presented by Kingma [62]. In both instances, a batch size of 16 is

employed, with the 2D beam case trained across 32000 epochs and the 3D liver case for 4000 epochs. The initial learning rate is set at 1e-4 and linearly decays to 1e-6 during the training process. Autoencoder networks are trained using Tensorflow [63] library on a Tesla V100-SXM2 GPU, utilizing the high-performance computing facilities at the University of Luxembourg [64].

Remark: Besides the size of the latent vector, there are several other hyperparameters that determine performance of both autoencoder architectures. In the case of 2D CNN network, these are: the size of convolution window, the number of channels per layer, the number of convolutions per level, and the type of pooling layers. In the case of 3D fully connected network, these are: the sizes of each layer, and the arrangement of skip connections. The particular choice of hyperparameters that determine both autoencoder architectures is primarily based on our earlier experience and studies, see [19, 57], which is then followed by the grid search optimization in the space of hyperparameters. The presented final architectures have been found to perform the best for the problems analyzed in this work.

3.2.2. Gaussian process details

All the Gaussian Process (GP) implementations presented in this work are performed using the Matérn covariance function, as it is known to be less prone to overfitting. (Note, however, that our framework is not limited to any particular covariance function.) The Matérn covariance function has the following form:

$$C_\nu(r_{ij}) = \lambda^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{r_{ij}}{\psi} \right)^\nu K_\nu \left(\sqrt{2\nu} \frac{r_{ij}}{\psi} \right), \quad (18)$$

where $r_{ij} = \|\mathbf{f}_i - \mathbf{f}_j\|$, \mathbf{f}_i and \mathbf{f}_j denote the i^{th} and j^{th} input vectors, $\|\cdot\|$ denotes the Euclidean norm, ψ is the length scale, $\Gamma(\nu)$ denotes the gamma function, K_ν is a modified Bessel function, λ^2 denotes the variance. The positive parameter ν is critical; the larger it is, the smoother $w(\mathbf{f})$ becomes in Eq. (7). A Gaussian Process (GP) with the Matérn covariance function is $\lceil \nu \rceil - 1$ times mean square differentiable, and for $\nu \rightarrow \infty$, the Matérn covariance function converges to the radial basis function (RBF). If we choose ν in the form of a half-integer, i.e., $\nu = p + 1/2$, it leads to a convenient form of the covariance function that is a product of a polynomial of order p and an exponential function. Often, as in this contribution, ν is fixed due to computational complexities and the extreme changes in the functions that occur as ν changes. The most common choices for ν are 1/2, 3/2, and 5/2. In this work, we decided to use $\nu = 5/2$ by comparing the predictions of the aforementioned covariance functions for the latent space. Accordingly, the 5/2 Matérn covariance function reads:

$$k(r_{ij}) = C_{5/2}(r_{ij}) = \lambda^2 \left(1 + \frac{\sqrt{5}r_{ij}}{\psi} + \frac{5r_{ij}^2}{3\psi^2} \right) \exp \left(- \frac{\sqrt{5}r_{ij}}{\psi} \right). \quad (19)$$

The remaining hyperparameters, ψ and λ^2 , are then inferred or optimized by maximizing the marginal likelihood in Eq. (12).

The kernel hyperparameters are obtained by training GPs using the scikit-learn library [65] on the latent datasets, as described in Section 2.2. Optimization is performed using the L-BFGS-B optimizer [66]. As described in Section 2.2, we implement independent GPs corresponding to each latent component of the autoencoder network.




Figure 6: (a) As described in the Figure 2, for an input force \mathbf{f} , the latent distribution is predicted using GP. \tilde{u}_μ^l and u^l represent the mean of the latent distribution and true latent solution for a particular latent DOF, respectively. (b) Latent distribution is then projected to the full space as described in Section 2.3. \tilde{u}_μ , u_r , u represent the mean of the full field displacement prediction, reconstructed solution and the true FEM solution for a particular full field DOF, respectively. e_f represents the framework error while e_r represents the reconstruction error, for the particular full field DOF.

3.3. Sources of prediction errors

At the prediction phase, for a new unseen case, the GPR is first applied to achieve the latent space prediction, which is subsequently projected to the full space solution using the decoder part of the autoencoder network, see Section 2 and Figure 2. Hence the error in the framework arises from GPR fitting as well as the compression and reconstructions steps done by the autoencoder network. The idea of this split of errors is schematically shown in Figure 6 for a single DOF of the full-field solution.

The mean prediction of the framework, \tilde{u}_μ , is interpreted as the displacement solution. Hence the absolute difference between this mean prediction, \tilde{u}_μ , and the true FEM solution, u , is considered as the error of the framework,

$$e_f = |\tilde{u}_\mu - u|. \quad (20)$$

The error contributed to autoencoder, e_r , is computed by taking absolute difference between the reconstructed solution and the FEM solution. The reconstructed solution is obtained by projecting the true latent solution (as obtained by Eq. (5)) to the full field, thus neglecting any contribution of the GP. The reconstructed solution reads $\mathbf{u}_r = \mathcal{U}_{\text{decoder}}(\mathbf{u}^l)$, and the corresponding reconstruction error is

$$e_r = |u_r - u|. \quad (21)$$

Finally, the error contributed to GP fitting, e_{GP} , is computed by taking the difference between the mean full-field prediction, \tilde{u}_μ , and the reconstruction solution, u_r , as follows:

$$e_{\text{GP}} = |\tilde{u}_\mu - u_r|. \quad (22)$$

Note: All the errors are computed in the decoded space, and are provided per full-field DOF.

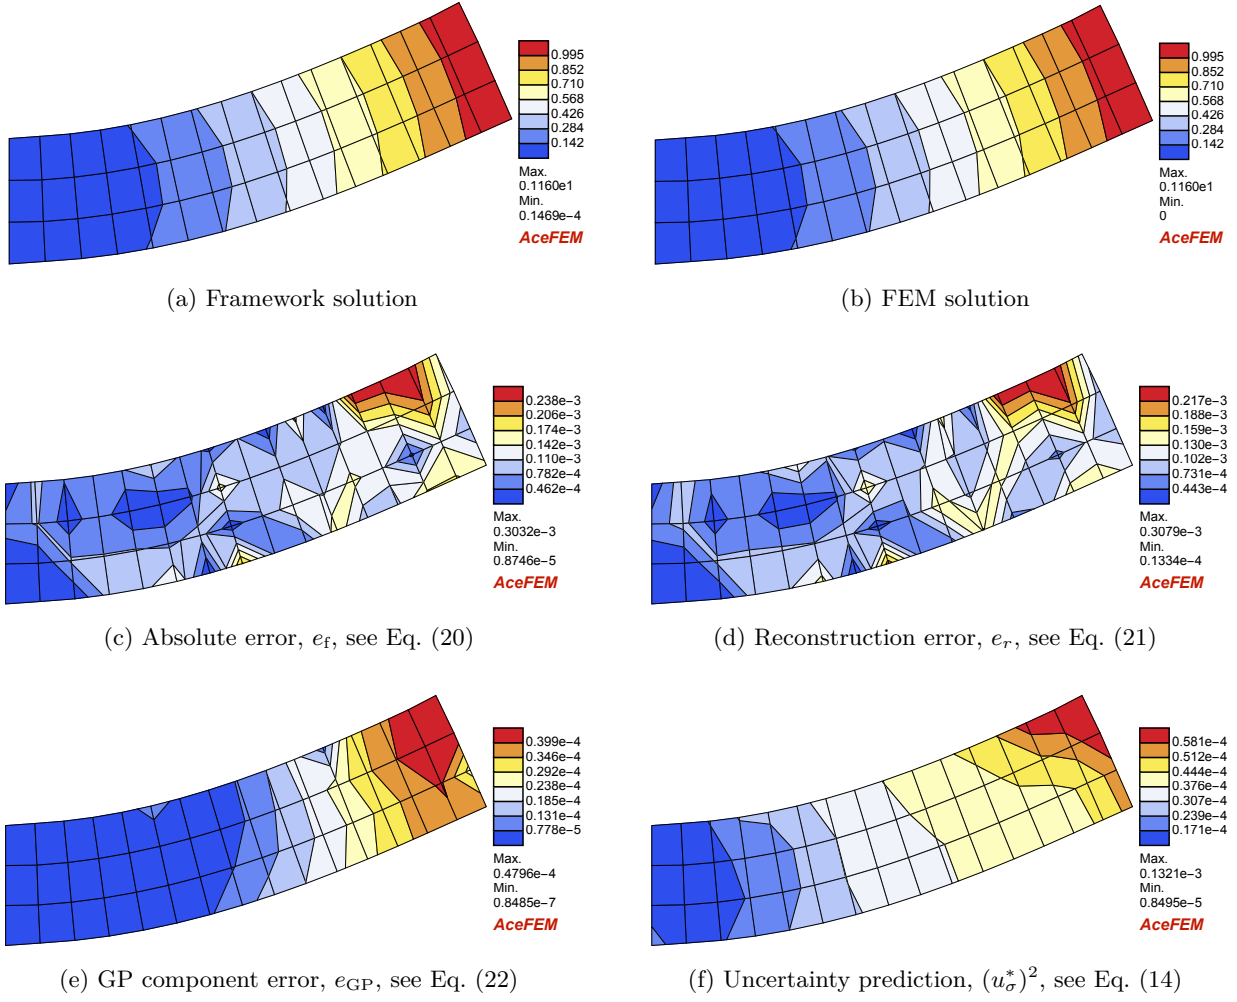


Figure 7: Deformation of the 2D beam subjected to the point load (-2.00, 2.19) N on the top right corner node. Nodal displacements obtained using the (a) proposed framework (b) FEM respectively. (c) Nodal error and (d) reconstruction error of the framework. (e) GP component error. (f) Uncertainty predicted by the framework for the full field solution.

3.4. Performance study of the GP+autoencoder framework

In this section we study the performance of the framework. Firstly, we showcase two particular load cases for the 2D and 3D example representing highest nodal displacement examples of each scenario. We visualize the predictions and corresponding errors, which are plotted on deformed meshes. In particular, we utilize the aforementioned breakdown of the framework error into reconstruction- and GP error, which aims to provide a clear and accessible understanding of how each component contributes to the overall performance of the framework. Figures 7 and 8 illustrate the interpolated node-wise L2 norm of the framework's predictions, including displacements, associated uncertainties, and errors for the 2D beam and 3D liver cases, respectively. Further, we provide overall error metrics for the entire test sets to quantify the framework's performance.

For the 2D beam example, we analyse the test example in with the maximum nodal displacement of 1.16 m, see Figure 7. The prediction of displacement field obtained using the proposed framework which is in a very good agreement with the FEM solution, as presented in the Figures 7a-b. Figure 7c provides nodal error of the entire framework. While Figure 7d and Figure 7e provide reconstruction error and GP error components as described by Eq.(21) and Eq.(22) respectively. Error plots show that all the errors are fairly low when compared to the magnitude of the displacement solution thereby displaying the capability of the framework in making accurate predictions. Figure 7f provides nodal uncertainty estimates predicted by the framework (2 standard deviations), as obtained from Eq.(14). One can observe that the GP component errors are within the uncertainty bounds for respective nodal predictions, i.e., in other words, true solution of the framework is contained within the displacement distribution for all the nodes.

Similarly, we analyse the test example of 3D liver case with the maximum nodal displacement of 0.14 m. Figure 8 shows that the framework is able to accurately predict the deformation responses of the 3D liver geometry subjected to external body forces. Figure 8a, 8b show that the nodal displacements predicted by the framework are in a great agreement with the FEM solution. Figure 8c, 8d present the framework and reconstruction errors which are extremely low compared to the magnitude of displacements. For the maximum displacement node, the prediction error is 0.2% of its magnitude. Directing attention to Figure 8e and Figure 8f, these figures present the error contribution of the GP component and the predicted uncertainty respectively. Again, the nodal GP errors are within the respective nodal uncertainty values thereby indicating that the framework is able to make reliable predictions.

The performance of the framework over the entire test set is computed using the mean absolute error metric. For the m^{th} test case, the mean error is given as follows:

$$e_m = \frac{1}{\mathcal{F}} \sum_{i=1}^{\mathcal{F}} |\tilde{\mathbf{u}}_{\mu}^*(\mathbf{f}_m)^i - \mathbf{u}_m^i|, \quad (23)$$

where \mathcal{F} is the number of DOFs of the mesh representing the full field space, $\tilde{\mathbf{u}}_{\mu}^*(\mathbf{f}_m)$ is the mean prediction of the GP+Decoder framework (full field prediction) as described through Eq.(14) and \mathbf{u}_m is the finite element solution for the full field space. To have a single validation metric over the entire test set, we compute the average mean norm \bar{e} and the corrected sample standard deviation $\sigma(e)$ as follows:

$$\bar{e} = \frac{1}{M} \sum_{m=1}^M e_m, \quad \sigma(e) = \sqrt{\frac{1}{M-1} \sum_{m=1}^M (e_m - \bar{e})^2}. \quad (24)$$

Finally, in addition to that, we also use the maximum error per degree of freedom over the entire test set

$$e_{\max} = \max_{m,i} |\tilde{\mathbf{u}}_{\mu}^*(\mathbf{f}_m)^i - \mathbf{u}_m^i|. \quad (25)$$

The proposed framework predicts probabilistic displacement fields corresponding to the full field mesh. First, we analyse the prediction performance by comparing mean predictions of the framework to the true FEM solutions. Table 2 presents error metrics for both benchmark examples.

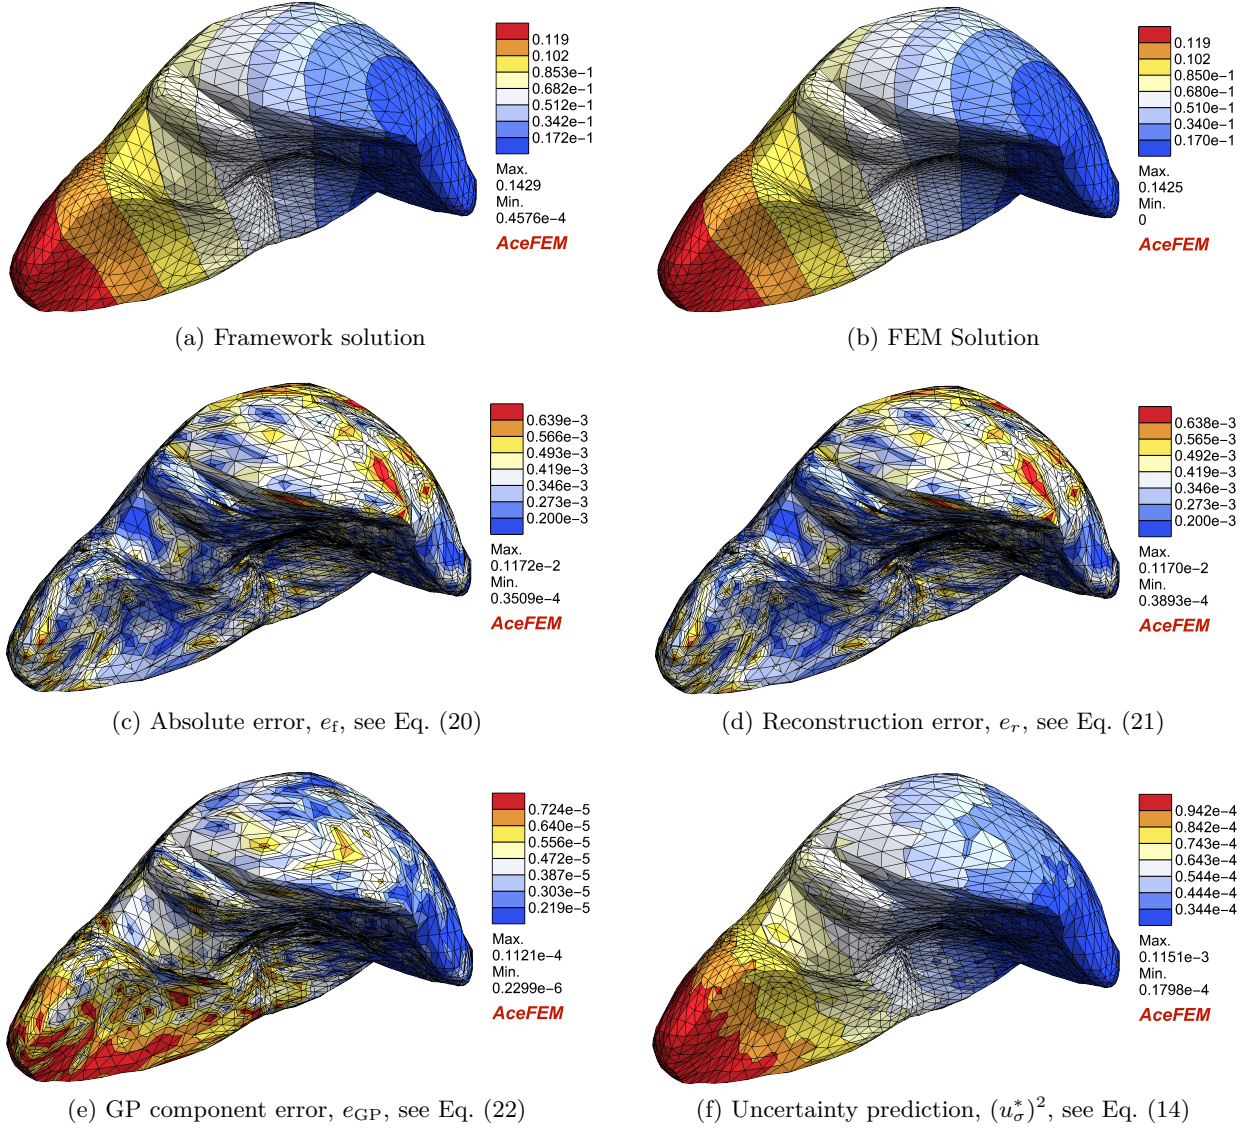


Figure 8: Deformation of the 3D liver for the highest nodal displacement case, subjected to external body force (0.71, 0.69, -0.78) N/kg. Nodal displacements obtained using the (a) proposed framework (b) FEM respectively. (c) Nodal error and (d) reconstruction error of the framework. (e) GP component error. (f) Uncertainty predicted by the framework for the full field solution.

Additionally, the maximum nodal displacements for both examples are also provided, i.e., we compute displacements of all the nodes for every single test case, and provide the maximum nodal displacement in the table. Table 2 shows that both mean and maximum errors are fairly low when compared to the displacement magnitudes, which demonstrates that the proposed framework is capable of predicting mechanical deformation responses with very low errors.

Example	M	\bar{e} [m]	$\sigma(e)$ [m]	e_{\max} [m]	Max nodal disp. [m]
2D Beam	300	0.1 E-3	0.8 E-4	7.5 E-3	1.16
3D liver	400	0.9 E-4	0.3 E-4	2.3 E-3	0.14

Table 2: Error metrics for 2D and 3D test sets for predictions using the proposed GP+autoencoder framework. M stands for the number of test examples, and \bar{e} , $\sigma(e)$ and e_{\max} are error metrics defined in Equations (24)-(25).

3.5. GP performance in the latent space

GPs are utilized to predict outcomes in the latent space. The independent GPs forecast a Gaussian probability distribution for each latent DOF. To ensure the efficiency and accuracy of GP’s predictions, we provide additional accuracy measures. First, we define prediction of a particular DOF as *healthy*, if the corresponding true solution lies within two standard deviations of the mean value.

To give an example, for the liver case, we plot latent probability distributions for a randomly chosen test example and check if the corresponding true latent values lie within the distribution. These true latent values are derived from the autoencoder compression of full field Finite Element Method (FEM) solutions. The displacement variables for all the latent distributions are transformed to an uniform scale. If the true latent solution for that latent distribution lies within $(u_\mu \pm 2u_\sigma)$, it is indicated by the green line, otherwise with the red line as shown in the Figure 9. For a particular test examples, if all the latent predictions are *healthy*, then that prediction is regarded as *correct*, see, e.g., Figure 9a. Even if a single latent prediction is not *healthy*, that prediction is regarded as *incorrect*, see Figure 9b. This is because a single non-healthy latent prediction can affect the full field solution during the reconstruction step.

The percentage of *healthy* and correct predictions together contribute to a comprehensive assessment of the GP’s predictive capabilities in predicting latent distributions. Table 3 provides these accuracy measures for both training and test datasets of both cases. The study on test datasets reveals that even though $\sim 95\%$ of predicted latent components are *healthy*, the percentage of *correct* GPs predictions can be significantly lower. We expect this unwanted effect to be more pronounced with the increasing dimension of the latent space.

Case	% correct predictions		% healthy predictions	
	Train	Test	Train	Test
2D beam	95.1	77.7	99.1	94.1
3D liver	93.1	74.5	99.0	95.6

Table 3: Performance of the GP in correctly predicting latent distributions of both cases.

3.6. Uncertainty for missing data region

As presented in the Table 1, the original 2D beam dataset of 6000 examples was generated by applying random forces of magnitude -2.5 N to 2.5 N in X and Y-direction. In order to further




Figure 9: Latent distribution prediction for two randomly chosen test examples from the liver test dataset. All the 16 true latent values are transformed to an uniform scale, along with the displacement random variables. True latents lying within the two standard deviations ($\pm 2\sigma$ region represented by shaded gray region) are indicated with green lines, while the latents lying outside are indicated with red colors. (a) All true latent displacements lie within $\pm 2\sigma$, which we classify as a correct prediction. (b) Three true latent values don't lie within $\pm 2\sigma$, which we classify as an incorrect prediction.

test the capability of the framework, we remove all the cases whose nodal input force magnitude is less than 1 N (the region represented by white circle in Figure 10). This subset of original data (5019 examples) is now used to train the proposed framework. Now the framework trained on this masked dataset is used to make predictions on two new input datasets.

In the first case, we apply a one dimensional Y-directional force whose magnitude changes from -3 to 3 N, as indicated by red dot line in the Figure 10a. In Figure 11a, we plot the uncertainty of selected latent DOF predicted by GP. For comparison, we also plot the error of latent DOF prediction as well. The figure shows that the uncertainty of prediction for all the latent DOFs is low in the region supported by the data and the prediction errors are within uncertainty bounds. The uncertainty grows quickly in both interpolated and extrapolated missing data region. Especially, the uncertainty grows rapidly in the extrapolated region (for magnitude of the force greater than 2.5 N).

In the second case, we use a new test dataset in which two dimensional input forces are randomly generated as denoted by red dots in Figure 10b. In Figure 11b we plot the uncertainty for the first latent DOF. It is evident that the uncertainty shows an increasing trend in the region not supported by training data. All the 20 latent DOFs show a very similar trend. In the presented examples, these uncertainties correlate with the GP prediction errors, which can provide useful information. However, it's crucial to emphasize that the predictive capabilities in the data-sparse and extrapolated regions can never be assured, with the tendency to deteriorate with the distance from the training data.




Figure 10: Reduced training dataset. The subset of the original dataset is marked with gray color. The distributions of two new test sets are illustrated with red dots in (a) and (b).

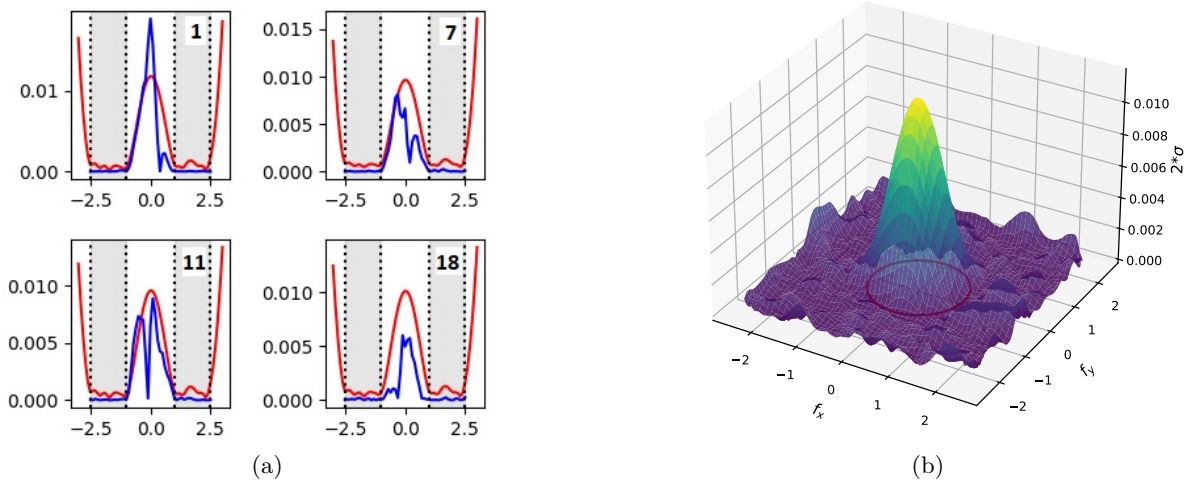


Figure 11: Uncertainty prediction for 2D beam case, for the GP trained using the missing dataset as indicated in Figure 10. In (a), the uncertainties (red line) and true latent prediction errors (blue line) are shown for selected four latent components (component number is indicated in the inset) for the test data depicted in Figure 10a. In (b), the uncertainties are shown for the first latent component for the test data depicted in Figure 10b. In (a) and (b) we observe the trend of increased uncertainty in the region not supported by the training data.

Similarly, higher uncertainty values in regions not supported by data are also observed in the full field space. For instance, one of the test points located in the white region of Figure 10b has its prediction projected into the full space. As shown in Figure 11b, the latent uncertainty for such an input is high. Similarly, Figure 12f shows that the predicted full field uncertainties are also high throughout the 2D domain. As mentioned earlier, uncertainties are only responsible for tracking the GP component error and are not directly correlated with the autoencoder errors. However, higher levels of uncertainty values, which exceed acceptable engineering thresholds, serve as a reliable indicator of whether the full field solution should be trusted or not.

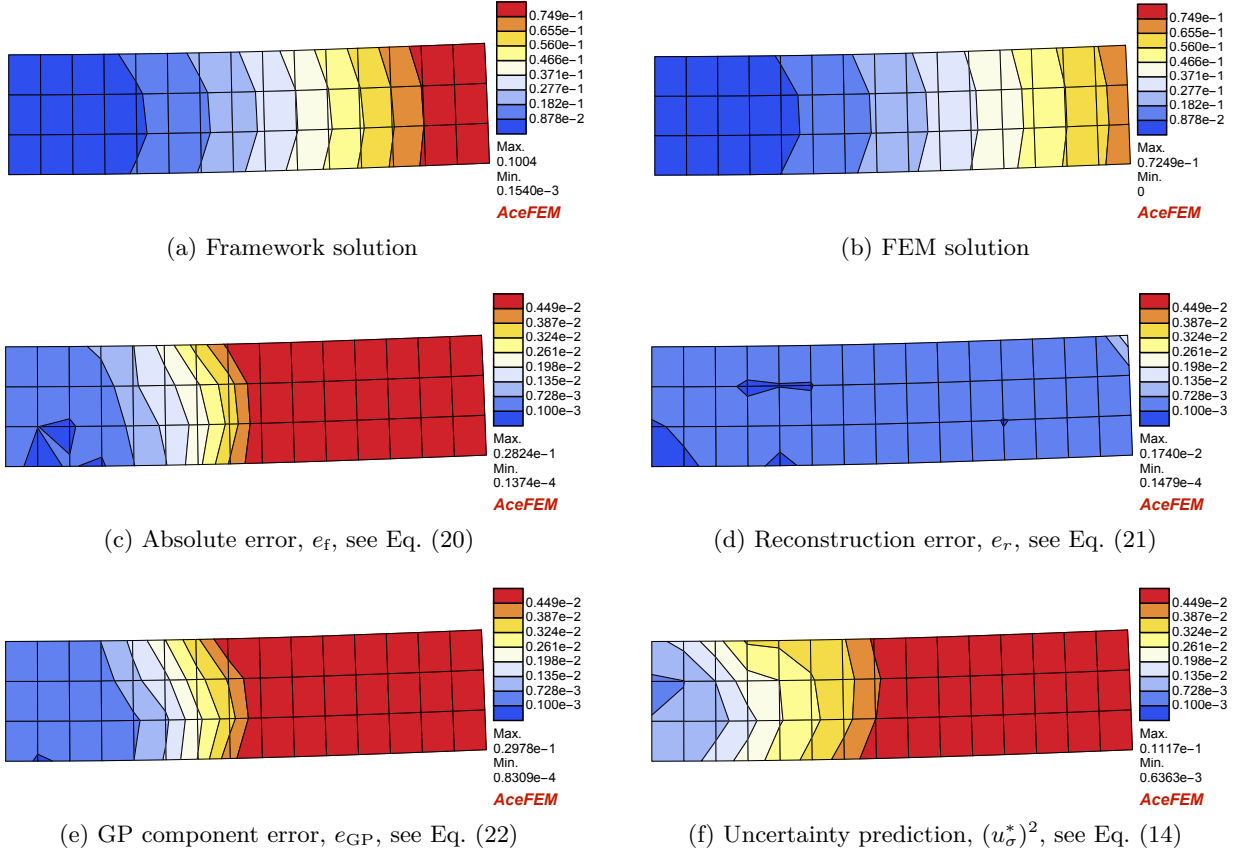


Figure 12: Deformation of the 2D beam subjected to a point load $(-0.14, 0.15)$ N on the corner node, which is present in the missing training data region as described by white circle in Figure 10b. Nodal displacements obtained using the (a) proposed framework (b) FEM respectively. (c) Nodal error and (d) reconstruction error of the framework. (e) GP component error. (f) Uncertainty predicted by the framework for the full field solution.

4. Conclusions

This work presents a novel approach to probabilistic surrogate modeling of high-fidelity simulations. It combines Gaussian processes with reduced-order modeling to efficiently simulate the mechanics of solids while accounting for uncertainties. The reduced representations of high-dimensional displacement data are achieved using autoencoder neural networks, allowing for drastic non-linear compression. For instance, this method reduced approximately 9100 degrees of freedom (DOFs) to 16 latent-space components, as demonstrated with the liver example. This strategy significantly reduces the computational cost of probabilistic predictions using Gaussian processes by lowering the output dimensionality, which is crucial because this cost can scale linearly to cubically with the output dimension size. We demonstrate that the proposed framework can accurately predict non-linear hyper-elastic deformations of solid bodies, along with the associated model uncertainties.

In addition to demonstrating the capabilities of our probabilistic framework, we also studied its possible limitations. We emphasized that while the uncertainties in the missing data region provide a useful indicator that is correlated with the distance from training data, they do not necessarily

need to represent the true solution errors. The development of methods to improve the quantitative nature of those predictions remains an open topic for future research. We also pointed out that even if a single component of the latent solution is predicted incorrectly, it can impact the entire reconstructed solution. This led us to propose more meaningful error metrics for quantifying the performance of the proposed probabilistic framework. This issue can be alleviated by developing a suited training strategy, which poses an interesting future research direction.

This study opens up several other opportunities for future research. One exciting direction is applying this approach to probabilistic simulations of time and path dependent problems. One could use latent state to track how solutions evolve in such cases, which would make the whole process much faster and more efficient [67]. Another direction would be to exploit state of the art graph autoencoder networks to find compressed representations for arbitrary high-dimensional meshes [68]. This advancement would overcome the limitations of fully connected networks, which are recognized for their challenges in handling problems with a high number of dimensions. Finally, we believe that the proposed approach can find direct application in research and development, which is not restricted to problems in mechanics but covers a broader spectrum of engineering and scientific domains. To facilitate this, we have made all the codes and datasets used in this work available as open-access in the GitHub repository at <https://github.com/saurabhdeshpande93/gp-auto-regression>.

Contributions

Saurabh Deshpande: Conceptualization, Data curation, Formal analysis, Methodology, GP+ANN implementation, Validation, Visualization, Writing – original draft, Writing – review & editing.

Hussein Rappel: Formal analysis, Methodology, GP implementation, Writing – original draft, Writing – review & editing.

Mark Hobbs: Reviewing the manuscript.

Stéphane P.A. Bordas: Funding acquisition, Reviewing the manuscript, Supervision.

Jakub Lengiewicz: Conceptualization, Data curation, Investigation, Methodology, Project administration, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing.

Acknowledgements:



This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 764644. Jakub Lengiewicz would like to acknowledge the support from EU Horizon 2020 Marie Skłodowska Curie Individual Fellowship *MOrPhEM* under Grant 800150. Jakub Lengiewicz would like to acknowledge the support of the Luxembourg National Research Fund (FNR) through the project SUMO, Grant INTER/GACR/21/16555380. Stéphane Bordas and Jakub

Lengiewicz are grateful for the support of the Fonds National de la Recherche Luxembourg FNR

grant QuaC C20/MS/14782078. Stéphane Bordas received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 811099 TWINNING Project DRIVEN for the University of Luxembourg. This paper only contains the author’s views and the Research Executive Agency and the Commission are not responsible for any use that may be made of the information it contains.

References

- [1] Lorenzo Grassi, Enrico Schileo, Christelle Boichon, Marco Viceconti, and Fulvia Taddei. Comprehensive evaluation of pca-based finite element modelling of the human femur. *Medical engineering & physics*, 36, 08 2014.
- [2] Gianmarco Aversano, Aurélie Bellemans, Zhiyi Li, Axel Coussement, Olivier Gicquel, and Alessandro Parente. Application of reduced-order models based on pca & kriging for the development of digital twins of reacting flow applications. *Computers & Chemical Engineering*, 121:422–441, 2019.
- [3] O. Gourey and C. Duriez. Fast, generic and reliable control and simulation of soft robots using model order reduction. *IEEE Transactions on Robotics*, 34(6):1565 – 1576, 2018.
- [4] Paris Papavasileiou, Eleni D. Koronaki, Gabriele Pozzetti, Martin Kathrein, Christoph Czettl, Andreas G. Boudouvis, and Stéphane P.A. Bordas. Equation-based and data-driven modeling strategies for industrial coating processes. *Computers in Industry*, 149:103938, 2023.
- [5] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural computation*, 10(5):1299–1319, 1998.
- [6] Gaëtan Kerschen, J.-C Golinval, ALEXANDER VAKAKIS, and LAWRENCE BERGMAN. The method of proper orthogonal decomposition for dynamical characterization and order reduction of mechanical systems: An overview. *Nonlinear Dynamics*, 41:147–169, 08 2005.
- [7] P. Kerfriden, P. Gosselet, S. Adhikari, and S.P.A. Bordas. Bridging proper orthogonal decomposition methods and augmented newton–krylov algorithms: An adaptive model order reduction for highly nonlinear mechanical problems. *Computer Methods in Applied Mechanics and Engineering*, 200(5):850–866, 2011.
- [8] Miguel A Mendez. Linear and nonlinear dimensionality reduction from fluid mechanics to machine learning. *Measurement Science and Technology*, 34(4):042001, jan 2023.
- [9] Pedro Díez, Alba Muixí, Sergio Zlotnik, and Alberto García-González. Nonlinear dimensionality reduction for parametric problems: A kernel proper orthogonal decomposition. *International Journal for Numerical Methods in Engineering*, 122(24):7306–7327, 2021.
- [10] A. Mendizabal, P. Márquez-Neila, and S. Cotin. Simulation of hyperelastic materials in real-time using deep learning. *Medical Image Analysis*, 59:101569, 10 2019.
- [11] Michael Frank, Dimitris Drikakis, and Vassilis Charissis. Machine-learning methods for computational science and engineering. *Computation*, 8(1), 2020.

- [12] Steven L. Brunton, Bernd R. Noack, and Petros Koumoutsakos. Machine learning for fluid mechanics. *Annual Review of Fluid Mechanics*, 52(Volume 52, 2020):477–508, 2020.
- [13] Anina ŠarkiĆ Glumac, Onkar Jadhav, Vladimir Despotović, Bert Blocken, and Stephane P.A. Bordas. A multi-fidelity wind surface pressure assessment via machine learning: A high-rise building case. *Building and Environment*, 234:110135, 2023.
- [14] Saurabh DESHPANDE. *Data Driven Surrogate Frameworks for Computational Mechanics: Bayesian and Geometric Deep Learning Approaches*. PhD thesis, Unilu - University of Luxembourg [Faculty of Science, Technology and Medicine], Luxembourg, 2023.
- [15] Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Ferran Alet, Suman Ravuri, Timo Ewalds, Zach Eaton-Rosen, Weihua Hu, Alexander Merose, Stephan Hoyer, George Holland, Oriol Vinyals, Jacklynn Stott, Alexander Pritzel, Shakir Mohamed, and Peter Battaglia. Learning skillful medium-range global weather forecasting. *Science*, 382(6677):1416–1421, 2023.
- [16] Saurabh Deshpande, Raúl I. Sosa, Stéphane P. A. Bordas, and Jakub Lengiewicz. Convolution, aggregation and attention based deep neural networks for accelerating simulations in mechanics. *Frontiers in Materials*, 10, 2023.
- [17] V. Krokos, V. Bui Xuan, S. P. A. Bordas, P. Young, and P. Kerfriden. A bayesian multiscale cnn framework to predict local stress fields in structures with microscale features. *Computational Mechanics*, 69:733–766, 05 2022.
- [18] Vasilis Krokos, Stéphane P. A. Bordas, and Pierre Kerfriden. A graph-based probabilistic geometric deep learning framework with online physics-based corrections to predict the criticality of defects in porous materials, 2022.
- [19] Saurabh Deshpande, Stéphane P.A. Bordas, and Jakub Lengiewicz. MAGNET: A graph U-Net architecture for mesh-based simulations. *Engineering Applications of Artificial Intelligence*, 133:108055, 2024.
- [20] Karol Frydrych, Maciej Tomczak, and Stefanos Papanikolaou. Crystal plasticity parameter optimization in cyclically deformed electrodeposited copper—a machine learning approach. *Materials*, 17(14), 2024.
- [21] Yasi Wang, Hongxun Yao, and Sicheng Zhao. Auto-encoder based dimensionality reduction. *Neurocomputing*, 184:232–242, 2016.
- [22] Luojie Shi, Kai Zhou, and Zequn Wang. Convolutional dimension-reduction with knowledge reasoning for reliability approximations of structures under high-dimensional spatial uncertainties. *Journal of Mechanical Design*, 146(7):071701, 01 2024.
- [23] G Thippa Reddy, M Praveen Kumar Reddy, Kuruva Lakshmana, Rajesh Kaluri, Dharmendra Singh Rajput, Gautam Srivastava, and Thar Baker. Analysis of dimensionality reduction techniques on big data. *Ieee Access*, 8:54776–54788, 2020.
- [24] Stefania Fresca, Luca Dede’, and Andrea Manzoni. A comprehensive deep learning-based approach to reduced order modeling of nonlinear time-dependent parametrized pdes. *Journal of Scientific Computing*, 87:1–36, 2021.

- [25] Arnaud Mazier, Sophie Ribes, Benjamin Gilles, and Stéphane P.A. Bordas. A rigged model of the breast for preoperative surgical planning. *Journal of Biomechanics*, 128:110645, 2021.
- [26] Antonio Loquercio, Mattia Segù, and Davide Scaramuzza. A general framework for uncertainty estimation in deep learning. *IEEE Robotics and Automation Letters*, 5(2):3153–3160, 2020.
- [27] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra. Weight uncertainty in neural network. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1613–1622, Lille, France, 07–09 Jul 2015. PMLR.
- [28] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1050–1059, New York, New York, USA, 20–22 Jun 2016. PMLR.
- [29] Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2022.
- [30] C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*, volume 1. MIT press Cambridge, 2006.
- [31] S. Rogers and M. Girolami. *A first course in machine learning, second edition*. Chapman & Hall/CRC, 2nd edition, 2016.
- [32] M.C. Kennedy and A. O’Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001.
- [33] M. J. Bayarri, J. O. Berger, R. Paulo, J. Sacks, J. A. Cafeo, J. Cavendish, C. H. Lin, and J. Tu. A framework for validation of computer models. *Technometrics*, 49(2):138–154, 2007.
- [34] P.D. Arendt, D.W. Apley, and W. Chen. Quantification of model uncertainty: Calibration, model discrepancy, and identifiability. *Journal of Mechanical Design*, 134(10):100908, 2012.
- [35] Arnaud Vadeboncoeur, Ieva Kazlauskaitė, Yanni Papandreou, Fehmi Cirak, Mark Girolami, and Omer Deniz Akyildiz. Random grid neural processes for parametric partial differential equations. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 34759–34778. PMLR, 23–29 Jul 2023.
- [36] P. S. Koutsourelakis. A multi-resolution, non-parametric, Bayesian framework for identification of spatially-varying model parameters. *Journal of Computational Physics*, 228(17):6184–6211, 2009.
- [37] H. Rappel, L. Wu, L. Noels, and L. A. A. Beex. A Bayesian Framework to Identify Random Parameter Fields Based on the Copula Theorem and Gaussian Fields: Application to Polycrystalline Materials. *Journal of Applied Mechanics*, 86(12):121009, 10 2019.

- [38] Hussein Rappel, Mark Girolami, and Lars A.A. Beex. Intercorrelated random fields with bounds and the bayesian identification of their parameters: Application to linear elastic struts and fibers. *International Journal for Numerical Methods in Engineering*, 123(15):3418–3463, 2022.
- [39] Chensen Ding, Hussein Rappel, and Tim Dodwell. Full-field order-reduced gaussian process emulators for nonlinear probabilistic mechanics. *Computer Methods in Applied Mechanics and Engineering*, 405:115855, 2023.
- [40] Carl Jidling. Strain field modelling using gaussian processes. *Report: Uppsala University, Division of Systems and Control*, (17004):82, 2017.
- [41] Dario Poloni, Daniele Oboe, Claudio Sbarufatti, and Marco Giglio. Towards a stochastic inverse finite element method: A gaussian process strain extrapolation. *Mechanical Systems and Signal Processing*, 189:110056, 2023.
- [42] Haitao Liu, Yew-Soon Ong, Xiaobo Shen, and Jianfei Cai. When gaussian process meets big data: A review of scalable gps. *IEEE transactions on neural networks and learning systems*, 31(11):4405–4423, 2020.
- [43] Wessel Bruinsma, Eric Perim, William Tebbutt, Scott Hosking, Arno Solin, and Richard Turner. Scalable exact inference in multi-output Gaussian processes. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1190–1201. PMLR, 13–18 Jul 2020.
- [44] Romit Maulik, Themistoklis Botsas, Nesar Ramachandra, Lachlan R. Mason, and Indranil Pan. Latent-space time evolution of non-intrusive reduced-order models using gaussian process emulation. *Physica D: Nonlinear Phenomena*, 416:132797, 2021.
- [45] James Donnelly, Alireza Daneshkhah, and Soroush Abolfathi. Forecasting global climate drivers using gaussian processes and convolutional autoencoders. *Engineering Applications of Artificial Intelligence*, 128:107536, 2024.
- [46] Jonathan Masci, Ueli Meier, Dan Cireşan, and Jürgen Schmidhuber. Stacked convolutional auto-encoders for hierarchical feature extraction. In *Artificial Neural Networks and Machine Learning–ICANN 2011: 21st International Conference on Artificial Neural Networks, Espoo, Finland, June 14–17, 2011, Proceedings, Part I 21*, pages 52–59. Springer, 2011.
- [47] Cheng-Yuan Liou, Wei-Chen Cheng, Jiun-Wei Liou, and Daw-Ran Liou. Autoencoder for words. *Neurocomputing*, 139:84–96, 2014.
- [48] Zilong Wang and Young-Jin Cha. Unsupervised deep learning approach using a deep auto-encoder with a one-class support vector machine to detect damage. *Structural Health Monitoring*, 20(1):406–425, 2021.
- [49] Pranshu Pant, Ruchit Doshi, Pranav Bahl, and Amir Barati Farimani. Deep learning for reduced order modelling and efficient temporal evolution of fluid simulations. *Physics of Fluids*, 33(10), 2021.

- [50] Takaaki Murata, Kai Fukami, and Koji Fukagata. Nonlinear mode decomposition with convolutional neural networks for fluid dynamics. *Journal of Fluid Mechanics*, 882:A13, 2020.
- [51] Stefania Fresca, Giorgio Gobat, Patrick Fedeli, Attilio Frangi, and Andrea Manzoni. Deep learning-based reduced order models for the real-time simulation of the nonlinear dynamics of microstructures. *International Journal for Numerical Methods in Engineering*, 123(20):4749–4777, 2022.
- [52] Krushna Shinde, Vincent Itier, José Mennesson, Dmytro Vasiukov, and Modesar Shakoor. Dimensionality reduction through convolutional autoencoders for fracture patterns prediction. *Applied Mathematical Modelling*, 114:94–113, 2023.
- [53] G.E. Hinton and R.R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science (New York, N.Y.)*, 313:504–7, 08 2006.
- [54] Jasem Almotiri, Khaled Elleithy, and Abdelrahman Elleithy. Comparison of autoencoder and principal component analysis followed by neural network for e-learning using handwritten recognition. In *2017 IEEE Long Island Systems, Applications and Technology Conference (LISAT)*, pages 1–5, 2017.
- [55] Mingyang Li and Zequn Wang. Deep learning for high-dimensional reliability analysis. *Mechanical Systems and Signal Processing*, 139:106399, 2020.
- [56] Mingyang Li and Zequn Wang. Deep reliability learning with latent adaptation for design optimization under uncertainty. *Computer Methods in Applied Mechanics and Engineering*, 397:115130, 2022.
- [57] Saurabh Deshpande, Jakub Lengiewicz, and Stéphane P.A. Bordas. Probabilistic deep learning for real-time large deformation simulations. *Computer Methods in Applied Mechanics and Engineering*, 398:115307, 2022.
- [58] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian data analysis*. Chapman & Hall/CRC Texts in Statistical Science. Chapman & Hall/CRC, 2003.
- [59] H. Rappel, L. A. A. Beex, J. S. Hale, L. Noels, and S. P. A Bordas. A tutorial on bayesian inference to identify material parameters in solid mechanics. *Archives of Computational Methods in Engineering*, 27, 2020.
- [60] J. Korelc. Multi-language and multi-environment generation of nonlinear finite element codes. *Engineering with Computers*, 18:312–327, 01 2002.
- [61] Joze Korelc and Peter Wriggers. *Automation of Finite Element Methods*. Springer, 2016.
- [62] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2017.
- [63] J. V. Dillon, I. Langmore, D. Tran, E. Brevdo, S. Vasudevan, D. Moore, B. Patton, A. Alemi, M. D. Hoffman, and R. A. Saurous. Tensorflow distributions. *CoRR*, abs/1711.10604, 2017.
- [64] S. Varrette, P. Bouvry, H. Cartiaux, and F. Georgatos. Management of an academic hpc cluster: The ul experience. 07 2014.

- [65] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830, 2011.
- [66] Roger Fletcher. *Practical methods of optimization*. John Wiley & Sons, 2000.
- [67] Stefanos Nikolopoulos, Ioannis Kalogeris, and Vissarion Papadopoulos. Non-intrusive surrogate modeling for parametrized time-dependent partial differential equations using convolutional autoencoders. *Engineering Applications of Artificial Intelligence*, 109:104652, March 2022.
- [68] Shivam Barwey, Varun Shankar, Venkatasubramanian Viswanathan, and Romit Maulik. Multiscale graph neural network autoencoders for interpretable scientific machine learning, 2023.