



HAL
open science

A hybrid numerical methodology coupling Reduced Order Modeling and Graph Neural Networks for non-parametric geometries: applications to structural dynamics problems

Victor Matray, Faisal Amlani, Frédéric Feyel, David Néron

► To cite this version:

Victor Matray, Faisal Amlani, Frédéric Feyel, David Néron. A hybrid numerical methodology coupling Reduced Order Modeling and Graph Neural Networks for non-parametric geometries: applications to structural dynamics problems. 2024. hal-04585018v1

HAL Id: hal-04585018

<https://hal.science/hal-04585018v1>

Preprint submitted on 28 May 2024 (v1), last revised 29 Aug 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

A HYBRID NUMERICAL METHODOLOGY COUPLING REDUCED ORDER MODELING AND GRAPH NEURAL NETWORKS FOR NON-PARAMETRIC GEOMETRIES: APPLICATIONS TO STRUCTURAL DYNAMICS PROBLEMS

Victor Matray, Faisal Amlani, David Néron

Université Paris-Saclay, CentraleSupélec, ENS Paris-Saclay, CNRS,
LMPS - Laboratoire de Mécanique Paris-Saclay,
Gif-sur-Yvette, France

Corresponding author: Victor Matray, victor.matray@ens-paris-saclay.fr

Frédéric Feyel

Université Paris-Saclay, CentraleSupélec, ENS Paris-Saclay, CNRS,
LMPS - Laboratoire de Mécanique Paris-Saclay,
Gif-sur-Yvette, France

Safran Tech, Digital Sciences & Technologies Department,
Magny-les-Hameaux, France

ABSTRACT

This work introduces a new approach for accelerating the numerical analysis of time-domain partial differential equations (PDEs) governing complex physical systems. The methodology is based on a combination of a classical reduced-order modeling (ROM) framework and recently-introduced Graph Neural Networks (GNNs), where the latter is trained on highly heterogeneous databases of varying numerical discretization sizes. The proposed techniques are shown to be particularly suitable for non-parametric geometries, ultimately enabling the treatment of a diverse range of geometries and topologies. Performance studies are presented in an application context related to the design of aircraft seats and their corresponding mechanical responses to shocks, where the main motivation is to reduce the computational burden and enable the rapid design iteration for such problems that entail non-parametric geometries. The methods proposed here are straightforwardly applicable to other scientific or engineering problems requiring a large-number of finite element-based numerical simulations, with the potential to significantly enhance efficiency while maintaining reasonable accuracy.

Keywords Reduced-Order Modeling · Non-Parametric Geometries · Graph Neural Networks · Deep Learning · Proper Generalized Decomposition · Finite Element Methods

1 Introduction

Numerical simulations are essential tools in a variety of scientific disciplines, including design engineering, physics, biology, and economics. They often involve resolving complex models governed by physical or phenomenological laws that can be mathematically described by partial differential equations (PDEs). In solid mechanics, the resolution of these models relies mainly on mesh-based numerical methods requiring discretization in space and time, the most widely-used of which is the Finite Element Method (FEM). In recent decades, a great deal of work [1, 2, 3, 4, 5, 6, 7, 8, 9] has been undertaken to improve and extend classical FEM formulations in order to encompass a wider variety of both

linear and nonlinear material behavior to the extent that FEM-based simulations are now ubiquitous in science and engineering [1, 10, 5].

However, despite advances in computing capacity in recent years, the computational cost of FEM-based solvers remains significant and often limiting for highly-sensitive engineering or scientific applications such as those related to aeronautical design [11, 12, 13] (the primary motivating context of this contribution). Iterating over a design space is necessary for engineers to develop optimized configurations [14, 15, 16], and each design may require some sort of computational mechanical analysis. Hence, it is highly advantageous and important to have very fast calculation tools that can enable rapid pre-dimensioning/pre-sizing of mechanical components from the initial design stage, i.e., providing adequately accurate estimations of the corresponding mechanical response and behavior without immediately requiring expensive finite element simulations. Indeed, such tools could help circumvent expensive full-field analyses of modifications incurred during the design phase of a project, potentially resulting in substantial time savings during validation stages. Fast approximate tools can ultimately give development teams the leeway to propose original designs (even those that are radically different from the standard) whose effectiveness can be quickly evaluated without the need for computationally costly *full order modeling* using FEM.

1.1 Acceleration by reduced-order modeling (ROM) approaches

A common technique for reducing the computational burden of large-scale FEM simulations is known as *reduced-order modeling* (ROM) [17, 18, 19]. Reduced-order models exploit the mathematical separability of solution fields in a full physical model, allowing one to solve a smaller, less-costly problem while maintaining control over the error in approximation [20, 21]. A notable class of potent reduction techniques relies on creating a low-dimensional space using a reduced-order basis (ROB) [22]. Construction of such a space involves an initial offline learning phase, where the high-fidelity problem (e.g., full-order FEM) is solved, after which appropriate snapshots are carefully selected. Subsequently, in the online phase, the original high-fidelity problem is often tackled using a projection, typically employing a Galerkin-type approach (although alternative methods are also feasible [23]). For an in-depth review of such projection-based methods, readers are encouraged to consult [17] and the associated references.

The focus of this work is on linear structural dynamics problems, where the most common ROM technique is the use of the Proper Orthogonal Decomposition (POD) [24, 25, 26]. POD has also demonstrated superior capabilities to well-established modal analysis techniques which are also widely used [27]. The principle of this approach is to build a ROB from intelligently-selected snapshots taken from previously-calculated solutions. Techniques such as the Reduced Basis Method (RBM) [20, 21] offer automatic selection procedures to determine the most relevant snapshots. RBMs can also be used to certify the results derived from the reduced model for any parameter contained in the PDE. For a deep overview, the interested reader can refer to [28] and the references therein. Note that ROB methods can also be extended to non-linear PDEs, e.g., via the Discrete Empirical Interpolation Method [29]. These aforementioned ROB methods are generally referred to as *a posteriori* methods and their costs are characterized by the time needed to calculate the database (offline) and the time needed to solve the reduced model (online).

In a mathematically-related but fundamentally-different approach, the Proper Generalised Decomposition (PGD) [30, 18, 31] provides an alternative technique for low-rank resolution, but *a priori*. This is accomplished by generating the ROB *on-the-fly*. Similarly to POD, the PGD method was initially proposed for space-time variables (known as radial approximations [30]). However, PGD has since been extended to other parameters such as those defining material properties [18]. Non-linear equations can also be considered with this approach, e.g., by use of the LATIN-PGD algorithm [32, 33, 34].

Although reduced-order models are becoming more increasingly used in practice [35, 33], they are only truly effective when the problem can be parameterized (e.g., for describing material properties, loading characteristics, or geometries [36]), and even more so in the context of multi-query simulations [31]. Geometric considerations, in particular, often require that the number of degrees-of-freedom remains constant among the offline training set when constructing a ROB (so that the information contained in the previously-calculated snapshots can be re-evaluated). This “geometry parameterization” restriction is the most limiting condition in the context of design innovation and iteration since, in essence, a very original structure born of an engineer’s imagination may not necessarily share the same parameterizations of other (even similar) designs.

When geometry is parametrizable (e.g., by length or width and even more when the topology evolves), ROM techniques are already efficient [36, 37]. However, their applicability is limited to geometries that can be parameterized by a few control points [36]. General techniques include Free-Form Deformation (FFD) [38, 39] as well as interpolation using Radial Basis Functions (RBF) [40, 41]. These methods define parameters as the shifting of specific control points that determine the morphing of the domain. For instance, in [42], airfoil profiles of the NACA family are studied using RBMs in the context of inverse design for aircraft wings. In [43], another approach is employed where a multi-parametric

PGD is used to obtain numerical maps for a rocket launcher component whose geometry is parameterized with two control points that are added to the spatial and temporal parameters of the PGD.

When the geometry is non-parametric (e.g., can be of any arbitrary shape or size), strategies for ROM are difficult to apply. The geometries do not share the same spatial dimension (after FEM discretisation), making classical ROB projection methods infeasible. There are also recently-introduced approaches based on mesh morphing [44, 45], the aim of which is to learn the transformation that transfers any mesh to a reference mesh, which then enables a common spatial dimension to be obtained without a priori knowledge of the geometry parameterization. Searching for this transformation is equivalent to determining an on-the-fly parameterization of the geometries in the database, and assumes that the new geometries tested will have the same parameterization. In [44], this technique is coupled with a ROM to simulate blood flow. Similarly, in [45], a similar strategy is applied to a database generated with five known material parameters and one unknown geometric parameter, which is then coupled to a Gaussian process method. Learning the morphing transformations as well as choosing reference meshes is not very straightforward [45]. Indeed, these transformations cannot necessarily be generalized, particularly when the topology of the mesh changes [46].

The goal of the present work is to propose a more general approach that requires neither parameterizing the geometry with a finite number of control points nor learning the morphing transformation, ultimately enabling consideration of completely unstructured meshes of highly-variable degrees-of-freedom.

1.2 Acceleration by deep learning-based approaches

An alternative class of methods that have been gaining interest for more rapidly and efficiently approximating solutions to PDEs is the burgeoning field of deep learning, which has now touched almost all scientific disciplines [47, 48]. Deep learning has shown remarkable capability in predicting non-linear relationships with physical data [49]. Similarly to classical reduced-order modeling, one needs to distinguish between two computational cost scales: the *offline* learning/training time (which is often quite long) and the *online* model inference time (which is very short—possibly real-time). A great deal of recent work has consisted of trying to enrich deep learning models with physical knowledge [47, 50]. To this end, there are three main types of approaches. The first, proposed in [51] and often referred to as “physics-informed” (or “learning bias”), consists of including physical knowledge in the loss function used to train the neural networks. The idea is to construct a loss function as the sum of a term that accounts for deviations from target data (error) and a term that penalizes non-physical behavior. A second approach, often referred to as “physics-augmented” (or “inductive bias”), consists in introducing physical knowledge *directly* into the architecture of the neural network [52]. For example, in [53], the neural network is constrained to be convex in the same way as the thermodynamic potential that the authors set out to learn. Similarly, in [54], a positive dissipation is constrained to satisfy the second law of thermodynamics. A third approach, classically referred to as “observation bias” [55], aims to intelligently introduce a data structure underlain by physical knowledge. Hence, the deep neural network (DNN) is exposed directly to the observed data and is expected to capture the underlying physical process via training.

Other learning-based approaches that may ignore this physical knowledge a priori—instead focusing on representations adapted to physical simulation data—also show remarkable capabilities. This is the case of approaches based on Graph Neural Networks (GNN), initially proposed in [56], which have become very popular in several communities: in fluid mechanics for Navier-Stokes equations (e.g., [57, 58]), in meteorology for global weather prediction (e.g., [59]), in biomechanics for describing the mechanics of human body organs (e.g., [60]), in Newtonian mechanics (e.g., [61]), and in solid mechanics (e.g., [62, 63]). Several recent works have also improved this approach by explicitly incorporating physical knowledge (see [54] and references therein). Such methods offer interesting extrapolation capabilities and are particularly well-suited for their application to a variety of geometries [64] (which is an important consideration for the structural design context of this work).

GNN and other deep learning approaches are based on training a model from a database which, in the motivating industrial applications of the present work, represents an opportunity to reuse the high-fidelity computations produced by other designs and engineering projects that are kept in databases. It should be noted that in this framework, the amount of available data is often limited, and this constraint is taken into account in the proof-of-concept work presented here. In addition, deep learning models (in general) face uncertainty as to the validity of their outputs which, together with their more-often-than-not lack of interpretability [47, 65], often hinder their certification and widespread acceptance in engineering. For this reason as well, the idea of combining deep learning approaches (e.g., GNN) with well-established solvers used in practice may reveal promising directions.

1.3 Present work

This contribution proposes a new numerical methodology to provide near real-time structural mechanical analysis combining the two distinct approaches discussed above: classical ROM methods and deep learning. The general idea is

to use the latter to predict the former for any arbitrary discretization size and geometric shape (in order to address the limitations of classical approaches as outlined above). More formally, we consider a set of N independent geometries $\{\Omega_p^{\text{data}}\}_{p=\{1,\dots,N\}}$ (which do not share any geometric parameter or discretisation size, as simplified in Figure 1), where each configuration is associated with a reference solution obtained by FEM: $\{\mathbf{U}_p^{\text{data}}\}_{p=\{1,\dots,N\}}$. We then determine the associated ROB designated by $\{\mathbf{P}_p^{\text{data}}\}_{p=1,\dots,N}$ (for example, via PGD). In order to provide predictions for a geometry Ω^{new} that may fall outside of the scope of treatment via the original ROB, this work proposes the use of GNN to provide a surrogate general ROB generator that provides \mathbf{P}_{GNN} , that is trained via supervised learning on $\{\mathbf{P}_p^{\text{data}}\}_{p=\{1,\dots,N\}}$. The solution \mathbf{U}^{new} of the full field is then obtained by a Galerkin projection on \mathbf{P}_{GNN} , similarly to a classical ROM approach. Note that from this \mathbf{U}^{new} solution it is possible to enrich it with new PGD modes, calculated in the usual way if this is necessary to obtain a satisfactory solution. A summary of the overall method, which we call GNN-PGD generator, is presented as a flow chart in Figure 2.

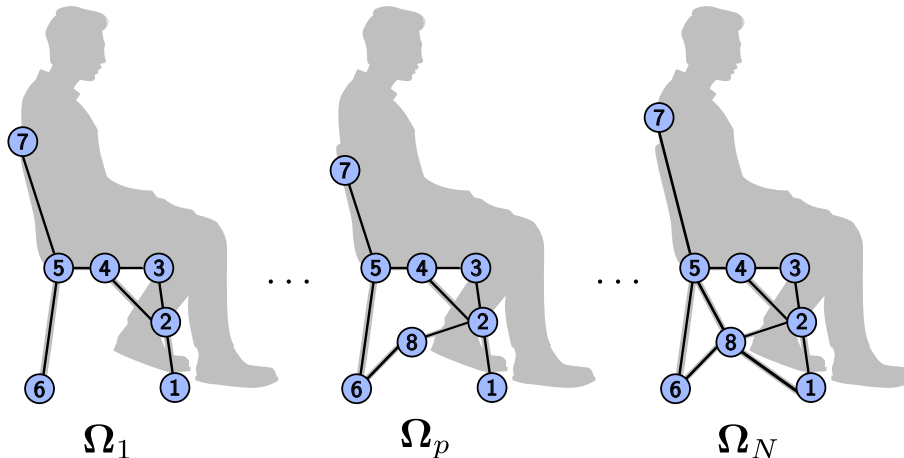


Figure 1: Example illustrations of N non-parameterized Ω_p geometries with possibly different discretization sizes.

Again, the ultimate aim is to enable consideration of much wider and more general ranges of geometries and topologies (i.e., a ROB built on very heterogeneous databases). The algorithmic framework we propose here aims to overcome the constraints mentioned in the previous subsection of traditional ROM approaches by exploiting the flexibility and inference speed of GNNs. The manuscript is organized as follows: section 2 describes the governing PDE of interest as well as the construction of (finite element-based) reference solutions. Thus, section 3 describes the proposed solver combining a GNN architecture with a classical ROM algorithm. Finally, section 4 presents a relevant proof-of-concept case study applying the new method to a problem of aircraft seat design, including performance comparisons with more conventional autoregressive GNN approaches [58].

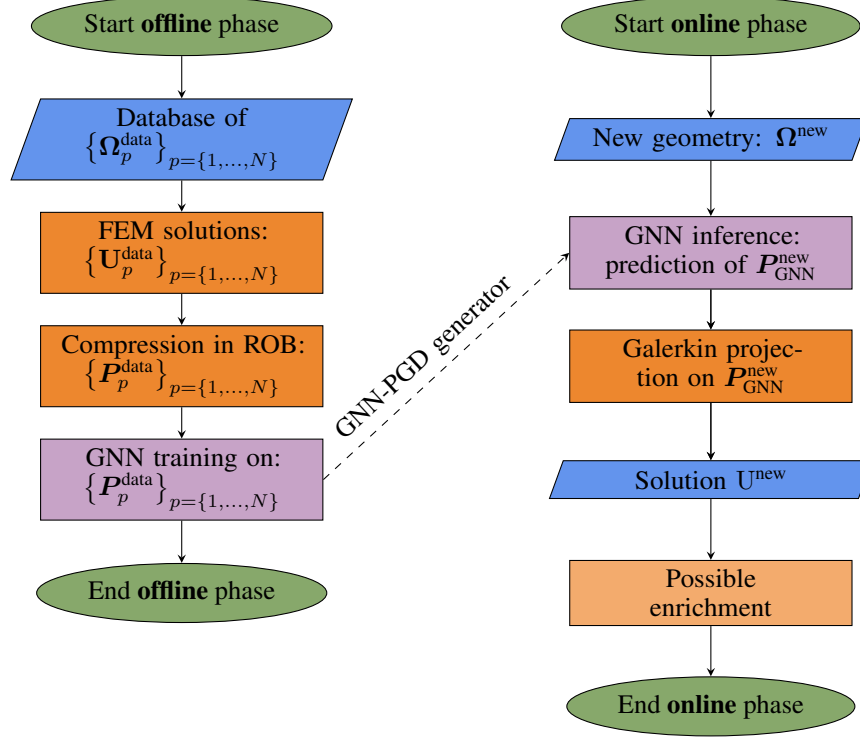


Figure 2: A flowchart diagram introducing the steps of the proposed GNN-PGD methodology (detailed in Section 3).

2 Governing equations

This section details the elastodynamics equations that govern our problem of interest (subsection 2.1) as well as the finite element method employed in this work for the construction of the full-order solutions $\{\mathbf{U}_p^{data}\}_{p=\{1,\dots,N\}}$ associated with non-parameterized geometries $\{\Omega_p^{data}\}_{p=\{1,\dots,N\}}$ (subsection 2.2). The latter is used for the supervised learning of the proposed methodology introduced in section 3.

2.1 Elastodynamics system

Let $\mathbf{x} \in \Omega \subset \{\Omega_p^{data}\}_{p=\{1,\dots,N\}}$, where $\Omega \subset \mathbb{R}^d$ ($d = 1, 2, 3$). We denote the boundary of Ω by $\partial\Omega$ which can be decomposed into $\partial\Omega = \partial\Omega_u \cup \partial\Omega_\sigma$, where $\partial\Omega_u$ (respectively $\partial\Omega_\sigma$) represents the parts of the domain boundary where Dirichlet (respectively Neumann) boundary conditions are imposed, with $\partial\Omega_u \cap \partial\Omega_\sigma = \emptyset$. With a time interval of interest given by $I = [0, T]$ for some final time $T \in \mathbb{R}^+$, we define the following vector spaces:

- $\mathcal{U}(\Omega) = \mathcal{H}^1(\Omega) = \{\mathbf{u} \in \mathbf{L}^2(\Omega), \nabla \mathbf{u} \in \mathbf{L}^2(\Omega)\}$
- $\mathcal{U}(\Omega; \mathbf{u}_d) = \{\mathbf{u} \in \mathcal{H}^1(\Omega), \mathbf{u}|_{\partial\Omega_u} = \mathbf{u}_d\}$
- $\mathcal{U}(\Omega; \mathbf{0}) = \mathcal{H}_0^1(\Omega) = \{\mathbf{u} \in \mathcal{H}^1(\Omega), \mathbf{u}|_{\partial\Omega_u} = \mathbf{0}\}$
- $\mathcal{U}_h(\Omega) = \{\mathbf{u} \in \mathcal{U}(\Omega), \mathbf{u}(\mathbf{x}) = \sum_{i=1}^{N_x} \varphi_i(\mathbf{x}) u_i\}$
- $\mathcal{J} = \mathbf{L}^2(I, \mathbb{R})$ with norm $\|\bullet\|_I = (\int_I \bullet^2 dt)^{1/2}$

The governing elastodynamics PDE is given by

$$\rho \ddot{\mathbf{u}} = \text{div} \boldsymbol{\sigma} + \mathbf{f}_d, \quad (\mathbf{x}, t) \in \Omega \otimes \mathcal{J}, \quad (1)$$

where $\mathbf{f}_d = \mathbf{f}_d(\mathbf{x}, t)$ is an external volumetric force and where, for notational simplicity, we define $\ddot{\mathbf{u}} = \partial^2 \mathbf{u} / \partial t^2$ throughout this paper. Dirichlet and Neumann boundary conditions on $\partial\Omega = \partial\Omega_u \cup \partial\Omega_\sigma$ are given by

$$\begin{cases} \mathbf{u} = \mathbf{u}_d, & (\mathbf{x}, t) \in \partial\Omega_u \otimes \mathcal{J} \\ \boldsymbol{\sigma} \cdot \mathbf{n} = \mathbf{F}_d, & (\mathbf{x}, t) \in \partial\Omega_\sigma \otimes \mathcal{J}, \end{cases} \quad (2a)$$

$$(2b)$$

where $\mathbf{F}_d = \mathbf{F}_d(\mathbf{x}, t)$ is an external surface density and where $\boldsymbol{\sigma}$ is the stress tensor given by $\boldsymbol{\sigma} = \mathbf{K} : \boldsymbol{\varepsilon}$ for $(\mathbf{x}, t) \in \boldsymbol{\Omega} \otimes \mathcal{J}$ (\mathbf{K} is the Hookean elastic operator and $\boldsymbol{\varepsilon}$ is the small deformation strain tensor). For all simulations in this contribution, the material is assumed to be linear elastic, homogeneous, and isotropic.

2.2 Reference solution using the Finite Element Method

This section provides a brief description of the construction of the reference solutions that are used to construct the reduced databases that are used to train the GNN of Section 3.2.

2.2.1 Weak formulation and semi-discretization in space

As in conventional FEM [1, 2, 3, 4, 5, 6, 7, 8, 9], boundary conditions are imposed strongly: the unknown field $\mathbf{u}(\mathbf{x}, t)$ is represented at each instant t in the space of functions defined by $\boldsymbol{\Omega}$ that satisfy a priori the Dirichlet conditions of Equation (2a). Hence, for the weak formulation, we consider test functions $\mathbf{v}^* \in \mathcal{U}(\boldsymbol{\Omega}; \mathbf{0})$ (i.e., the space of kinematically admissible functions at zero), yielding a system for the unknown $\mathbf{u}(\mathbf{x}, t)$ given by

$$\begin{cases} m(\mathbf{v}^*, \mathbf{u}) + k(\mathbf{v}^*, \mathbf{u}) = f(\mathbf{v}^*; t), & \mathbf{u} \in \mathcal{U}(\boldsymbol{\Omega}; \mathbf{u}_d), \forall \mathbf{v}^* \in \mathcal{U}(\boldsymbol{\Omega}; \mathbf{0}), \\ \mathbf{u}(\mathbf{x}, 0) = \mathbf{u}_0, & (\mathbf{x}, t) \in \boldsymbol{\Omega} \otimes \{0\}, \\ \dot{\mathbf{u}}(\mathbf{x}, 0) = \dot{\mathbf{u}}_0, & (\mathbf{x}, t) \in \boldsymbol{\Omega} \otimes \{0\}, \end{cases} \quad (3)$$

where the scalar products $m(\cdot, \cdot)$ and $k(\cdot, \cdot)$ and the linear form $f(\cdot; t)$ are defined respectively by

$$\begin{cases} m(\mathbf{v}^*, \mathbf{u}) = \int_{\boldsymbol{\Omega}} \rho \dot{\mathbf{u}} \mathbf{v}^* d\Omega, \\ k(\mathbf{v}^*, \mathbf{u}) = \int_{\boldsymbol{\Omega}} \boldsymbol{\sigma} : \boldsymbol{\varepsilon}(\mathbf{v}^*) d\Omega, \\ f(\mathbf{v}^*; t) = \int_{\boldsymbol{\Omega}} \mathbf{f}_d \mathbf{v}^* d\Omega + \int_{\partial\boldsymbol{\Omega}} \mathbf{F}_d \mathbf{v}^* dS. \end{cases} \quad (4)$$

The system given by (3) can be resolved by a Galerkin approach that searches for the solution in a subspace $\mathcal{U}_h(\boldsymbol{\Omega}) \subset \mathcal{U}(\boldsymbol{\Omega})$ of finite dimension N_x with a basis given by $\boldsymbol{\Phi} = [\boldsymbol{\varphi}_1, \dots, \boldsymbol{\varphi}_{N_x}]$. Hence an approximate solution at each time t to \mathbf{u} can be represented as $\mathbf{u}_h = \boldsymbol{\Phi} \mathbf{u}(t)$ where $(t) \in \mathbb{R}^{N_x}$ are the corresponding coefficients in the approximation basis $[\boldsymbol{\varphi}_1, \dots, \boldsymbol{\varphi}_{N_x}]$. The semi-discretized problem in space hence consists of solving the following system for the displacement vector $\mathbf{u} : \mathcal{J} \rightarrow \mathbb{R}^{N_x}$:

$$\begin{cases} \mathbb{M} \dot{\mathbf{u}}(t) + \mathbb{K} \mathbf{u}(t) = \mathbf{f}(t), \\ \mathbf{u}(0) = \mathbf{u}_0, \\ \dot{\mathbf{u}}(0) = \mathbf{v}_0, \end{cases} \quad (5)$$

where $\mathbb{M} = m(\boldsymbol{\Phi}, \boldsymbol{\Phi})$ is the mass matrix, $\mathbb{K} = k(\boldsymbol{\Phi}, \boldsymbol{\Phi})$ is the stiffness matrix, and $\mathbf{f}(t) = f(\boldsymbol{\Phi}; t)$ is a generalized force. The vectors $\mathbf{u}_0 \in \mathbb{R}^{N_x}$ and $\mathbf{b}_0 \in \mathbb{R}^{N_x}$ are respective coefficients of \mathbf{u}_0 and $\dot{\mathbf{u}}_0$ in the approximation basis $\boldsymbol{\Phi}$ (where, again, $\dot{\mathbf{u}} = d\mathbf{u}/dt$ and $\ddot{\mathbf{u}} = d^2\mathbf{u}/dt^2$).

2.2.2 Time integration

The final ODEs in Equation (5) can be integrated in time using a classical Newmark scheme [66] (a commonly-employed approach [67, 31]). Specifically, we utilize an implicit scheme with average acceleration, which leads to a conservative and stable algorithm. The time interval of interest $\mathbf{I} = [0, T]$ can be discretized into $t_n = (n-1)\Delta t$, $n = 1, \dots, N_t$ for $\Delta t = T/(N_t - 1)$. Time is discretized uniformly with a timestep Δt . The corresponding discretized space in time (of size N_t), is henceforth denoted $\mathcal{J}_{\Delta t}$, and hence $\mathbf{U} \in \mathcal{U}_h(\boldsymbol{\Omega}) \otimes \mathcal{J}_{\Delta t}$ where \mathbf{U} denotes the discrete space-time solution field.

3 Methodology

This section details the methodology introduced in Figure 2 of the introduction: the generation of the classical ROB sets $\{\mathbf{P}_p^{\text{data}}\}_{p=\{1,\dots,N\}}$ that are each associated with a non-parametric geometry $\{\Omega_p^{\text{data}}\}_{p=\{1,\dots,N\}}$ (Section 3.1); the proposed GNN-PGD model trained on $\{\mathbf{P}_p^{\text{data}}\}$ in order to construct/predict a ROB for an unknown geometry Ω_{test} (Section 3.2); and the Galerkin projection step employed to solve the corresponding time-dependent ODE given by Equation (5) (Section 3.3). A summary of the overall methodology is presented in Section 3.4.

3.1 Database generation

PGD [30, 18, 31] is employed to compress the reference FEM solutions of Section 2.2 in the form of separated variable mode products, i.e., each reduced-order basis contained in $\{\mathbf{P}_p^{\text{data}}\}_{p=\{1,\dots,N\}}$ corresponds to PGD spatial modes. Modes produced by such a decomposition have been shown to provide a better representation of a solution than the eigenmodes of the structure [27], especially for non-linear problems [31].

In order to explain the principles of PGD we temporarily use a continuous formulation of the referential problem of the equation (1). Given a known solution $\mathbf{u}(\mathbf{x}, t) : \Omega \times \mathbf{I} \rightarrow \mathcal{U}(\Omega) \otimes \mathcal{J}$, the method of separating space-time variables via PGD consists of seeking an appropriate approximation $\mathbf{u}_M \in \mathcal{U}(\Omega) \otimes \mathcal{J}$ of $\mathbf{u}(\mathbf{x}, t)$ in the form given by [32]

$$\mathbf{u}(\mathbf{x}, t) \approx \mathbf{u}_M(\mathbf{x}, t) = \sum_{m=1}^M \mathbf{\Lambda}_m(\mathbf{x}) \lambda_m(t), \quad (6)$$

where $\mathbf{\Lambda}_m \in \mathcal{U}(\Omega)$ is a spatial mode and $\lambda_m \in \mathcal{J}$ is a temporal mode (their product, $\mathbf{\Lambda}_m \lambda_m$, is often called a space-time mode [32]). The total number of modes M is called the *rank* of the approximation [32].

The construction of mode m is typically obtained through a greedy algorithm that seeks a rank 1 approximation and minimizes its error with respect to the approximation error of rank $m - 1$ [18, 68, 32], i.e.,

$$\begin{cases} \mathbf{\Lambda}_m, \lambda_m = \arg \min_{\mathcal{U}(\Omega) \times \mathcal{J}} \left(\left\| \underbrace{\mathbf{u}(\mathbf{x}, t) - \mathbf{u}_{m-1}(\mathbf{x}, t)}_{\text{approximation error of rank } m-1} - \underbrace{\mathbf{\Lambda}_m(\mathbf{x}) \lambda_m(t)}_{\text{new approximation of rank 1}} \right\|_{\mathbf{L}^2}^2 \right), \\ \mathbf{u}_{m-1}(\mathbf{x}, t) = \sum_{k=1}^{m-1} \mathbf{\Lambda}_k(\mathbf{x}) \lambda_k(t) \end{cases}, \quad (7)$$

where $\|u\|_{\mathbf{L}^2} = (\int_{\mathcal{J}} \int_{\Omega} \mathbf{u}^2(\mathbf{x}, t) d\Omega dt)^{1/2}$. Defining $\Delta \mathbf{u} = \mathbf{u} - \mathbf{u}_{m-1}$ and exploiting stationarity, Equation (7) can be solved as a coupled problem given by [32]:

$$\begin{cases} \forall \lambda^* \in \mathcal{J}, \int_{\mathbf{I}} \lambda^* (\int_{\Omega} (\Delta \mathbf{u} - \mathbf{\Lambda}_m \lambda_m) \mathbf{\Lambda}_m d\Omega) dt = 0 & \iff \lambda_m = \frac{\int_{\Omega} \Delta \mathbf{u} \mathbf{\Lambda}_m d\Omega}{\int_{\Omega} \mathbf{\Lambda}_m^2 d\Omega} = f_u(\mathbf{\Lambda}_m) \\ \forall \mathbf{\Lambda}^* \in \mathcal{U}(\Omega), \int_{\Omega} \mathbf{\Lambda}^* (\int_{\mathbf{I}} (\Delta \mathbf{u} - \mathbf{\Lambda}_m \lambda_m) \lambda_m dt) d\Omega = 0 & \iff \mathbf{\Lambda}_m = \frac{\int_{\mathbf{I}} \Delta \mathbf{u} \lambda_m dt}{\int_{\mathbf{I}} \lambda_m^2 dt} = g_u(\lambda_m) \end{cases} \quad (8)$$

The solutions $\lambda_m, \mathbf{\Lambda}_m$ exist and are unique up to a multiplicative factor. In order to ensure overall uniqueness of the solutions in (8), the spatial mode is usually normalized [32]. A fixed-point algorithm, using stagnation of the temporal mode as a stopping criterion, can be used to find such solutions with guaranteed convergence since seeking the best rank r approximation with a given $\Delta \mathbf{u}$ can be shown to be equivalent to a generalized eigenvalue problem [68].

In practice, the final rank M is selected based on an error ϵ that serves as the stopping criterion for the greedy algorithm, i.e.,

$$\frac{\|\mathbf{u} - \mathbf{u}_M\|_{\mathbf{L}^2}^2}{\|\mathbf{u}\|_{\mathbf{L}^2}^2} \leq \epsilon. \quad (9)$$

This ultimately enables automatic selection of the optimal number of modes for a given ϵ , which can be computationally advantageous over SVD-based approaches that require calculation of the complete set of modes of the original field \mathbf{u} [67] in order to form an approximation or basis.

In the discrete sense (where the FEM-based reference solution is given by $\mathbf{U} \in \mathcal{U}_h(\Omega) \otimes \mathcal{J}_{\Delta t}$), one seeks the separated basis above in the corresponding finite-dimensional space, i.e., $\forall m, \mathbf{\Lambda}_m \in \mathcal{U}_h(\Omega), \lambda_m \in \mathcal{J}_{\Delta t}$, and hence the reduced-order approximation is given by

$$\mathbf{U} \simeq \mathbf{U}_M = \sum_{m=1}^M \mathbf{\Lambda}_m \otimes \lambda_m. \quad (10)$$

Although $N \times N_t$ real numbers are needed to store the full field \mathbf{U} , the representation for \mathbf{U}_M in Equation (10) requires only $M \times (N + N_t)$ real numbers for $M \ll N$.

Hence for each element of the set of arbitrary geometries $\{\Omega_p^{\text{data}}\}_{p=\{1,\dots,N\}}$, the associated ROB (via a PGD on their corresponding reference solutions as described above) can be given by

$$\mathbf{P}_p^{\text{data}} = \text{Span} \left(\{\Lambda_{p,m}^{\text{data}}\}_{m=\{1,\dots,M\}} \right), \quad (11)$$

where \bullet^{data} denotes the labeled data for a training database. The set of each ROB, $\{\mathbf{P}_p^{\text{data}}\}_{p=\{1,\dots,N\}}$, constitutes the complete database for the supervised learning employed in the GNN component (detailed in the following subsection) of the GNN-PGD generator introduced in this work.

3.2 Deep learning on graphs: Graph Neural Networks (GNNs)

The following presents an overview of GNNs and additionally details the implementation of our proposed architecture. As a reminder, the aim of the GNN-PGD generator is to propose a ROB to describe the solution field of a new, arbitrary, geometry Ω_{test} . In order to accomplish this, the GNN-PGD generator is trained on the database $(\{P_p^{\text{data}}\}_{p=\{1,\dots,N\}})$ presented in the previous subsection. Figure 3 provides an illustrative summary of the main steps of a GNN.

3.2.1 Overview of GNN

The formalism of graph neural networks (GNNs) offers a powerful approach for data processing of numerical problems involving meshes. A mesh can be interpreted as a graph comprising nodes i for $i = 1, \dots, n$ linked by edges. These nodes and edges contain features x_i and e_{ij} respectively. The graph also contains mesh connectivity information in the form of an adjacency matrix A (which is simply a reformulation of the connectivity table deriving from an FEM problem). As an example, the adjacency matrix of the graph in Figure 3 is given by

$$A = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}_{n \times n}$$

In this work, we consider undirected graphs, which result in symmetrical A adjacency matrices. Moreover, we don't add any self-loops to the nodes [69], which enables a diagonal of zeros in A .




Figure 3: Illustration of a Message-Passing layer on a three-stage graph. Neural networks symbolize functions whose parameters can be adjusted by learning (here applied to node three of the example graph of $n = 7$ nodes).

GNNs are recognized as a generalization of Convolutional Neural Networks (CNNs) applied to graph-structured data. Traditionally, CNNs have been used for Cartesian grid-structured data such as images [70]. This generalization enables GNNs to detect and understand complex structures in data by applying convolution-like operations, making them extremely effective for, e.g., regression tasks.

Our proposed architecture uses the Message-Passing GNN (MP-GNN) formalism [69], which is a mathematical formulation local to the node level. In MP-GNN models, the latent feature vector h_i of a given node i is derived by employing a permutation-invariant aggregator function \oplus , such as a sum or a mean, over the feature vectors of its neighbors $\mathcal{N}(i)$. Additionally, each neighbor's feature vector may undergo transformation by a function ψ . Subsequently, \oplus may undergo further transformation by another function ϕ . This sequence of operations constitutes a GNN layer[69]. Formally, the equation governing the feature vector $h_i^{(\ell+1)}$ of a node i in the subsequent GNN layer $\ell + 1$ can be expressed as

$$h_i^{(\ell+1)} = \phi \left(h_i^\ell, \bigoplus_{j \in \mathcal{N}(i)} \psi(h_i^\ell, h_j^\ell) \right), \quad (12)$$

which is illustrated in Figure 3.

Each layer ℓ is thus defined by the three functions (\oplus, ψ, ϕ) . The message transmission function ψ and the update function ϕ are learnable and enable us to exhibit the important graph characteristics required for learning. The aggregation function \oplus is *not* learnable and must be permutation invariant. This latter characteristic ensures that the order and number of neighboring nodes has no impact on the dimension of the aggregation result. This flexibility enables GNN to handle graphs of varying sizes, making them a powerful tool for regression applications on graph-structured data, and therefore potentially on physical problems requiring meshing of any arbitrary size. For further details, the reader is referred to [71].

3.2.2 GNN-PGD implementation and training strategy

The most commonly-found approach for physical problems [60, 57, 59, 61], originally proposed in [58], is known as MESHGRAPHNET. Such a method employs an encode-process-decode architecture of [56], which consists of several multilayer perceptrons (MLPs) [72] shared between all nodes of the graph. This architecture is invoked to design our GNN-PGD algorithm, which is summarized in algorithm 2 (described with a local perspective, at the scale of node i , for clarity, and detailed in this section). In addition to the connectivity matrix \mathbf{A} of the graph, the input to the GNN-PGD procedure contains k characteristics (to be considered) of each node i , i.e., $\mathbf{x}_i \in \mathbb{R}^k$ (forming a row of a matrix we call \mathbf{X}). These characteristics are specific to the Ω geometry and can include, for example, the position of the nodes or the type of boundary conditions to which they are subjected (Dirichlet or Neumann). This is clarified with further details that are presented in section 4.

Algorithm 2: GNN-PGD description

GNN-PGD (\mathbf{X}, \mathbf{A})

inputs : \mathbf{x}_i input feature vector of node i
 \mathbf{A} graph adjacency matrix
output : \mathbf{y}_i output feature vector of node i
Encode
 | $\mathbf{h}_i^0 = \text{MLP}_{\text{node_encoder}}(\mathbf{x}_i)$
Process
 | **for** $0 \leq \ell \leq L - 1$ **do**
 | | $\mathbf{h}_i^{(\ell+1)} = \phi^\ell \left(\mathbf{h}_i^\ell, \oplus_{j \in \mathcal{N}(i)} \psi^\ell \left(\mathbf{h}_i^\ell, \mathbf{h}_j^\ell, \underbrace{\mathbf{x}_i - \mathbf{x}_j}_{\text{skip-connections}} \right) \right)$
 | | Instance Normalisation of $\mathbf{h}_i^{(\ell+1)}$
Decode
 | $\mathbf{p}_i = \text{MLP}_{\text{node_decoder}}(\mathbf{h}_i^L)$
return \mathbf{p}_i ;

For a general output $\mathbf{P} = [\mathbf{p}_1, \dots, \mathbf{p}_n]^\top$, the goal is to predict an image of the ROB associated with the geometry Ω , i.e., at each node i , we predict the output $\mathbf{p}_i \in \mathbb{R}^g$. Since the ROB is associated with M spatial modes, and since GNN provides results at the node level, each node has $g = M \times d$ features (with $d = 1, 2, 3$, the dimension of the problem). In other words, \mathbf{p}_i contains the projections of the ROB \mathbf{P} along the different d dimensions of the Ω geometry. In the training phase, this output \mathbf{P} is compared to the given labeled training data $\{\mathbf{P}_p^{\text{data}}\}_{p=\{1, \dots, N\}}$ in order to adjust the weights of the GNN-PGD generator. In the inference phase (online phase), we designate this output by $\mathbf{P}_{\text{GNN}}^{\text{new}}$. A summary of the notation and nomenclature for the GNN-PGD proposed in this work is given in Table 1.

The three steps in algorithm 2 are as follows. In the first step (**Encode**), an MLP is utilized to transform the input quantities \mathbf{x}_i into a latent space of dimension H . In this work, we employ a two-layer MLP with SiLU activation functions [73] and a normalization layer on the output [74]. The second step (**Process**) employs the GNN layers. Details of the MLP architecture used here are given in table 2. In addition to the information provided by previous layers, the corresponding message functions $\psi^{(\ell)}$ take into account the non-transformed physical quantities \mathbf{x}_i (known as skip-connections, as proposed in DenseNet [75]). After each convolution, an instance normalization layer [76] is applied in order to improve convergence and stability of the learning phase without a dependence on the batch size used during learning [74] (unlike batch normalization layers [77]). Our proposed method employs the same input and output dimensions (of size H) during this process phase, thanks to our choice of MLP architecture and the use of a permutation-invariant aggregation function. The final third step (**Decode**) in algorithm 2 transforms the process output

Structure of graph inputs	
$G = (V, E)$	A graph; V and E are sets of vertices and edges,
n, m	Numbers of vertices and edges in G ; $ V = n, E = m$,
$\mathcal{N}(i)$	1-hop neighborhood of i ,
$\mathbf{A} \in \mathbb{R}^{n \times n}$	The graph adjacency matrices.
Structure of GNN computations	
L ,	The number of GNN layers,
k ,	The number of input node features,
g ,	The number of output node features,
$\mathbf{X} \in \mathbb{R}^{n \times k}$,	Input node features matrix,
$\mathbf{P} \in \mathbb{R}^{n \times g}$,	Outputs node features matrix,
$\mathbf{x}_i, \mathbf{p}_i, \mathbf{h}_i^\ell$,	Input, output, and hidden feature vector of a node i (layer ℓ).

Table 1: Nomenclature of symbols used to describe the proposed GNN-PGD algorithm.

from the latent dimension H to the physical output dimension g . In order to achieve this, we use another two-layer MLP (with SiLU as the activation function for the first layer) and a linear activation function for the output layer. The architectures described above are summarized in Table 2.

Function	Inputs Size	Outputs Size	Details	Learnable?
$\text{MLP}_{\text{node_encoder}}(\bullet)$	k	H	linear ; SiLU ; linear ; SiLU ; LayerNorm	✓
$\psi^\ell(\bullet)$	H	H	linear ; SiLU ; linear ; SiLU	✓
$\oplus(\bullet)$	H	H	mean function	✗
$\phi^\ell(\bullet)$	H	H	linear ; SiLU ; linear ; SiLU	✓
$\text{MLP}_{\text{decoder}}(\bullet)$	H	q	linear ; SiLU ; linear	✓

Table 2: Description of the architectures of the various MLPs employed to build the GNN-PGD algorithm.

For each bath b , the loss function employed in this work for the supervised training of the GNNs is a Mean Squared Error (MSE) function that is weighted by hyperparameters $\{\eta_m\}_{m=\{1, \dots, M\}}$ (each associated with a spatial mode that comprises the ROB \mathbf{P}) and is given by

$$\mathcal{L}_b = \sum_{m=1}^M \eta_m \text{MSE}(\mathbf{\Lambda}_{b,m} - \mathbf{\Lambda}_{b,m}^{\text{data}}), \quad (13)$$

where $\mathbf{\Lambda}_{b,m}$ is the m -th mode associated with the ROB \mathbf{P} of batch b , and where $\mathbf{\Lambda}_{b,m}^{\text{data}}$ is the m -th mode associated with the data $\mathbf{P}_b^{\text{data}}$. Hence the total loss is simply given by

$$\mathcal{L} = \sum_{b=1}^{N_{\text{batch}}} \mathcal{L}_b. \quad (14)$$

Using this loss function, optimization for training proceeds with a stochastic gradient descent algorithm AdamW [78], which incorporates L^2 regularization on the GNN weights to prevent overfitting. This algorithm is controlled by two hyperparameters: the learning rate α and the "weight decay" γ (for regularization). A noise percentage of 5% of the average value of \mathbf{X} is added to the inputs of the GNN-PGD in order to improve the regularisation of the network. [79].

A generalized summary of the GNN-PGD hyperparameters employed in this work is provided in Table 3. These hyperparameters can significantly influence the quality of learning and, consequently, the performance of the final trained GNN-PGD model. In order to optimize these hyperparameters—typically an exploratory and experimental task [45, 80]—we propose the use of a specific method based on a systematic design of experiments (DOE) that is characterized by two types of variables: the maximum number of models to be tested (100 in all subsequent results of this paper), and the intervals of values that the aforementioned hyperparameters are permitted to take in each model (discrete or continuous). The 100 models are generated using the DOE method called maximum projection Latin Hypercube Sampling (LHS), which enables an efficient exploration of all the different possible model combinations [81].

The general idea is to train these 100 models of different values of hyperparameters for a reasonable number of epochs and then select the most relevant model based on a common metric. The models identified as most relevant or accurate are then re-trained for a larger number of epochs.

Hyperparameters	
H	dimension of the latent space,
L	numbers of GNN layers,
η_m	weight of learned modes,
α	the learning rate,
γ	the weight decay,
b_s	the batch size.

Table 3: Hyperparameters that are optimized for the GNN-PGD learning using a Design of Experiments (DOE) procedure known as Latin Hypercube Sampling (LHS) [81].

At the end of the offline phase (see Figure 2) described in this section, a GNN-PGD generator is obtained, which constructs a ROB $\mathbf{P}_{\text{GNN}} = \text{GNN-PGD}(\Omega^{\text{new}})$ for a new (arbitrary) geometry Ω^{new} . As described in the following section, this ROB can then be used to solve the associated ROM for the complete problem.

3.3 Galerkin projection on the trained ROB

To complete the solver, Equation (5) for $n \times d$ degrees-of-freedom is solved in the form of a product of functions with separated variables, similar to common methods of projection onto reduced bases [27, 19, 17], i.e.,

$$\mathbf{u}^{\text{new}}(t) = \mathbf{P}_{\text{GNN}}^{\text{new}} \boldsymbol{\lambda}(t), \quad (15)$$

where $\mathbf{P}_{\text{GNN}}^{\text{new}}$ is the ROB in space obtained by GNN as described in the previous section. In order to obtain the corresponding temporal coefficients of the general solution given by Equation (15), we utilize a standard Galerkin-type projection approach [27] which results in an ODE-system given by

$$\mathbf{P}_{\text{GNN}}^{\text{new} \top} \cdot \mathbb{M} \underbrace{\mathbf{P}_{\text{GNN}}^{\text{new}} \ddot{\boldsymbol{\lambda}}(t)}_{\ddot{\mathbf{u}}^{\text{new}}(t)} + \mathbf{P}_{\text{GNN}}^{\text{new} \top} \cdot \mathbb{K} \underbrace{\mathbf{P}_{\text{GNN}}^{\text{new}} \boldsymbol{\lambda}(t)}_{\mathbf{u}^{\text{new}}(t)} = \mathbf{P}_{\text{GNN}}^{\text{new} \top} \cdot \mathbf{f}(t), \quad (16)$$

which is a system of M ODEs (for each spatial mode $i = 1, \dots, M$) for the corresponding M temporal unknowns $\boldsymbol{\lambda}(t) = (\lambda_1(t), \dots, \lambda_M(t))^\top$. The cost of solving this linear system is very low, since the size of the operators to be inverted is $M \ll n \times d$. Once the unknown $\lambda_i(t)$ are obtained from solving Equation (16), the complete space-time solution can be reconstructed via Equation (15) (or, more explicitly, Equation (6)).

For noise that may result in the $\mathbf{P}_{\text{GNN}}^{\text{new}}$ prediction, what follows is a proposed procedure describing how one can account for such errors. Indeed, one can reasonably expect a prediction of the GNN-PGD generator that is not identical to the reference \mathbf{P} . In the next subsection, we present a noise-reducing strategy and apply it to a simulated example of artificial Gaussian noise to demonstrate its efficacy.

3.3.1 Noise effect and resolution strategy

In order to investigate and ultimately mitigate the impact that noise, which may be present in spatial modes, can have on the quality of the projection and resolution of (16), we consider the ‘‘true’’ (or ‘‘exact’’) modes of the PGD decomposition (i.e., those that are constructed for the database given by $\{\boldsymbol{\Lambda}_m^{\text{data}}\}_{m=\{1, \dots, M\}}$). Gaussian noise of different levels is applied to all degrees-of-freedom of the structure, and the intensity is adjusted by a percentage of the maximum value of the modes amplitude. The corresponding ‘‘noisy’’ modes are denoted $\{\boldsymbol{\Lambda}_m^{\text{noisy}}\}_{m=\{1, \dots, M\}}$ and constitute a new reduced basis given by

$$\mathbf{P}_{\text{noisy}} = \text{Span} \left(\{\boldsymbol{\Lambda}_m^{\text{noisy}}\}_{m=\{1, \dots, M\}} \right). \quad (17)$$

Figure 4 presents the angle formed between the two vector spaces \mathbf{P} and $\mathbf{P}_{\text{noisy}}$ as a function of percentage of artificial noise added for an example ROB solution comprised of three modes. It can be observed that for low levels of noise, these two vector spaces are very close. However, for noise levels close to 100%, these two vector spaces are almost orthogonal.

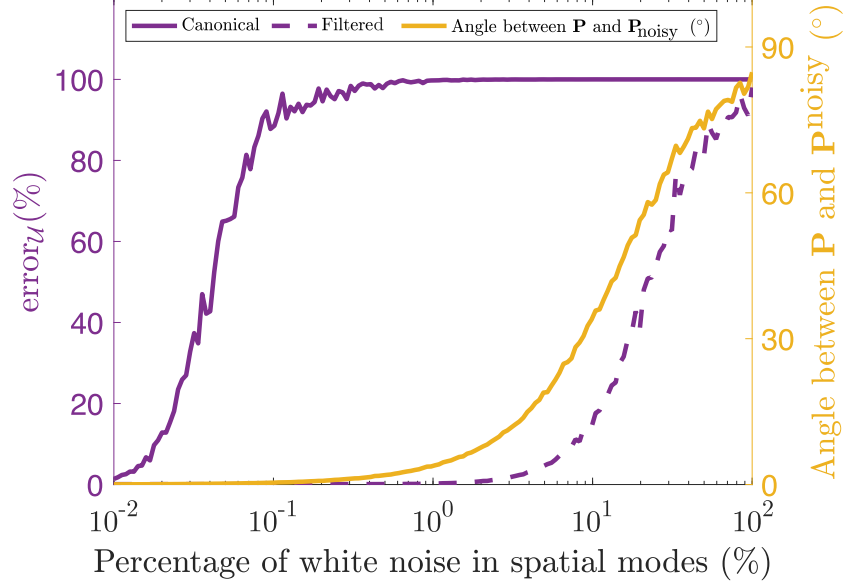


Figure 4: Effect of the presence of noise in the ROB on the resolution of the projected equation. Three modes have been used to generate the ROB.

Figure 4 additionally presents the evolution of the L^2 error in the field \mathbf{U}^{new} (relative to the reference solution \mathbf{U}_{ex}) as a function of the noise added to $\mathbf{P}_{\text{noisy}}$, defined by the expression

$$\text{error}_{\mathcal{U}} = \frac{\|\mathbf{U}_{\text{ex}} - \mathbf{U}^{\text{new}}\|_{\mathcal{U}}}{\|\mathbf{U}_{\text{ex}}\|_{\mathcal{U}}}, \% \quad \|\bullet\|_{\mathcal{U}} = \left(\int_{\Omega \times J} \bullet^2 \, d\Omega dt \right)^{1/2}.$$

Two cases are presented: a *canonical* case and a *filtered* case. The former corresponds to the evolution of the error solving the original form of the ODE system (given by Equation (16)), i.e.,

$$\mathbf{P}_{\text{noisy}}^{\top} \cdot \underbrace{\mathbb{M} \mathbf{P}_{\text{noisy}} \ddot{\boldsymbol{\lambda}}(t)}_{\dot{\mathbf{u}}^{\text{noisy}}(t)} + \mathbf{P}_{\text{noisy}}^{\top} \cdot \underbrace{\mathbb{K} \mathbf{P}_{\text{noisy}} \boldsymbol{\lambda}(t)}_{\mathbf{u}^{\text{noisy}}(t)} = \mathbf{P}_{\text{noisy}}^{\top} \cdot \mathbf{f}(t). \quad (18)$$

The *filtered* curve corresponds to the error when solving a preconditioned formulation of (16) which, in general, is given by

$$\mathbf{P}_{\text{noisy}}^{\top} \cdot \underbrace{\mathbb{S} \mathbb{M} \mathbf{P}_{\text{noisy}} \ddot{\boldsymbol{\lambda}}(t)}_{\dot{\mathbf{u}}^{\text{noisy}}(t)} + \mathbf{P}_{\text{noisy}}^{\top} \cdot \underbrace{\mathbb{S} \mathbb{K} \mathbf{P}_{\text{noisy}} \boldsymbol{\lambda}(t)}_{\mathbf{u}^{\text{noisy}}(t)} = \mathbf{P}_{\text{noisy}}^{\top} \cdot \mathbb{S} \mathbf{f}(t), \quad (19)$$

where $\mathbb{S} \in \mathbb{R}^{(n \times d)^2}$ can be interpreted as a filter or a preconditioner. Here, a preconditioner of $\mathbb{S} = \mathbb{K}^{-1}$ has been utilized, inspired by [82], where a preconditioned norm with operator $\mathbb{S} = \mathbb{K}^{-1}$ is proposed (referred to as the reconditioned equilibrium gap (REG)). In [82], the concept of “spectral sensitivity” is introduced in order to quantify the sensitivity of different norms, constructed with different operators (such as the stiffness operator \mathbb{K}), towards the presence of noise from displacement field measurements (the context of that study is the identification of nonlinear behavior laws from experimental measurements by image correlation). Indeed, the \mathbb{K} operator has a high spectral sensitivity [82], reflecting its high sensitivity to noise, unlike the \mathbb{M} operator. Such sensitivity can be attributed to the differential nature of \mathbb{K} (see Equations (4) and (5)). On the other hand, it’s the integrating nature of \mathbb{S} that can correct the undesirable effects of such noise. Choices other than $\mathbb{S} = \mathbb{K}^{-1}$ can accomplish the same objective, as long as they contain the same regularization properties of \mathbb{K}^{-1} . For the preliminary (linear) solver proposed here, we employ $\mathbb{S} = \mathbb{K}^{-1}$ as a proof-of-concept: despite the prohibitive cost of calculating \mathbb{K}^{-1} in a linear elastodynamics context, such cost remains negligible for nonlinear problems [82] (which is the long-term objective of the methodology presented in this paper).

Figure 4 illustrates the significant improvement that can be attained by using the filtering/preconditioning strategy (Equation (19)) versus solving the original canonical ODE system (Equation (18)). Indeed, for all levels of noise in $\mathbf{P}_{\text{noisy}}$, the former results in significantly lower error; this implies that solving the unfiltered system of Equation (18) is very sensitive to the noise present in $\mathbf{P}_{\text{noisy}}$ and may possibly lead to a very polluted ROB that has poor approximation capabilities of the original FEM system of interest. For all numerical experiments of section 4, the preconditioned system given by Equation (19) is employed to obtain the temporal coefficients $\lambda_i(t)$.

3.4 Summary of our methodology

algorithm 3 presents a summary of the overall methodology (additionally summarized in the introductory Figure 2).

Algorithm 3: Methodology summary

Offline phase $\left(\{\Omega_p^{\text{data}}, \mathbf{U}_p^{\text{data}}\}_{p=\{1, \dots, N\}} \right)$

inputs : $\{\Omega_p^{\text{data}}, \mathbf{U}_p^{\text{data}}\}_{p=\{1, \dots, N\}}$, known geometries and associated high-fidelity solution fields,

output : GNN-PGD generator: trained GNN model,

Database compression

for $0 \leq p \leq N$ **do**

 creation of the ROB $\mathbf{P}_p^{\text{data}}$ from the PGD compression of the high fidelity field $\mathbf{U}_p^{\text{data}}$,

return $\{\mathbf{P}_p^{\text{data}}\}_{p=\{1, \dots, N\}}$

GNN-PGD training

 Supervised learning of GNN-PGD on $\{\mathbf{P}_p^{\text{data}}\}_{p=\{1, \dots, N\}}$

return *GNN-PGD generator*

return *GNN-PGD generator*

Online phase $(\Omega^{\text{new}}, \mathbb{M}, \mathbb{K}, \mathbf{f}, \mathbb{S})$

inputs : $\Omega^{\text{new}}, \mathbb{M}, \mathbb{K}, \mathbf{F}, \mathbb{S}$, new geometry (unknown to GNN-PGD generator and not parameterized) and its associated operators

output : \mathbf{U}^{new} : full displacement field solution

GNN-PGD inference

$\mathbf{P}_{\text{GNN}}^{\text{new}} = \text{GNN-PGD}(\Omega^{\text{new}})$: ROB proposal by GNN inference,

return $\mathbf{P}_{\text{GNN}}^{\text{new}}$

Galerkin projection

$\mathbf{P}_{\text{GNN}}^{\text{new} \top} \cdot \mathbb{M} \mathbf{P}_{\text{GNN}}^{\text{new}} \ddot{\lambda}(t) + \mathbf{P}_{\text{GNN}}^{\text{new} \top} \cdot \mathbb{K} \mathbf{P}_{\text{GNN}}^{\text{new}} \lambda(t) = \mathbf{P}_{\text{GNN}}^{\text{new} \top} \cdot \mathbb{S} \mathbf{f}(t)$

 smart Galerkin projection on $\mathbf{P}_{\text{GNN}}^{\text{new}}$, and resolution of the system ODE,

return $\lambda(t)$

$\mathbf{u}(t) = \mathbf{P}_{\text{GNN}}^{\text{new}} \lambda(t)$: full displacement field

return \mathbf{U}^{new}

4 Performance/case study: simulating aircraft seat response to crash loads

In order to assess the performance of our proposed solver, we consider a case study relevant to structural design, in particular towards the problem of presizing/predimensioning of aircraft parts subjected to sudden loadings. The aircraft industry uses the acceleration (or deceleration) loading profile depicted in Figure 5 in order to assess the strength of aircraft seats according to regulation 49 CFR Part 572, Subpart B [11]. This loading simulates the deceleration experienced by an aircraft in the event of a collision with an obstacle. The sizing criteria focus on both the mechanical integrity of the seat components as well as the physical safety of the passenger. In the proof-of-concept methodology presented here, passenger modeling is not taken into account, and the analysis focuses solely on structural integrity. The acceleration load illustrated in Figure 5 is applied as an inertial force at the seat structure level and also as a loading force at the seating surface (equivalent to the deceleration force experienced by a 90-kg passenger). These two contributions are reflected in the loading term $\mathbf{f}(t)$ in the right-hand-side of Equation (5).




Figure 5: Loading of interest

4.1 Database generation

A database is generated using finite element simulations for the training of the GNN-PGD. In order to accommodate the common restrictions of industrial applications, we use a reasonably small number of geometries: 500, similarly to [45]. Out of these 500 airplane seats, 80% ($N = 400$ seats) are chosen to constitute the training set (i.e., the seats seen by the GNN-PGD for adjusting parameter weights through learning), i.e., $\{\Omega_p^{\text{data}}\}_{p=\{1,\dots,N=400\}}$ different geometries in the training database. Another 10% (50 seats) of the original 500-seat database is used for the validation set, representing seats unknown to the network but used to fine-tune the learning process in order to prevent overfitting by employing early-stopping criteria [83]. The remaining 10% (50 seats) make up the test set on which the GNN performance is evaluated after training, representing new Ω^{new} geometries that are completely unknown to the GNN-PGD model (see algorithm 3). We note that the original high-fidelity space-time fields $\{\mathbf{U}_p^{\text{data}}\}_{p=\{1,\dots,500\}}$, whose solutions are facilitated by FEM, have very similar separability properties. Indeed, by setting a stopping criterion of the PGD mode generation algorithm to $\varepsilon = 10^{-3}$ in Equation 9, we obtain a decomposition rank of $M = 3$ for all set of seats, and hence train with and keep only $M = 3$ modes in the ROB. A summary of the database generation procedure is provided in Figure 6.




Figure 6: Database generation

For the proof-of-concept presented here, we utilize a 2D modeling approach with the assumption of plane deformations (extension of the overall methodology to three-dimensions is straightforward). The material considered is Aluminum 6082-T6, commonly used in the aerospace industry [11], and is further assumed to be homogeneous and linear elastic

isotropic. In order to generate the database from which the seats in Figure 7 are extracted, we employ five geometric parameters: the radii of each of the curved holes and the (x, y) position of the interior hole's center. Although these parameters are crucial for database generation, we recall that the solver proposed here is for non-parametric geometries, and hence are not considered beyond database generation (similarly to other approaches [45, 44]). Indeed, such parameters are specifically unknown to the GNN-PGD model for the motivations of this paper (cf. section 1). The meshes are generated by the open-source software package GMSH [84] with T3 linear elements. In our context, we assume that the input meshes possess good quality in terms of element aspect ratios and node densities, which are adapted to the regularity of the fields of interest. The temporal scheme utilized to solve problem (5) is the Newmark scheme [66], particularly specifically the implicit formulation with average acceleration. A temporal convergence study results in an a choice of an optimal number of time steps corresponding to $N_t = 1000$. The finite element solver for generating the high-fidelity database (described in subsection 2.2.1 and subsection 2.2.2) is self-implemented in an in-house MATLAB code.




Figure 7: Example of three seats from the database to illustrate the embedded nodes and seating surface.

4.1.1 Database characteristics

Figure 8 presents database statistics in the form of histograms. It can be seen that the number of degrees-of-freedom varies considerably from seat to seat, with a total database standard deviation of 58 degrees-of-freedom and an amplitude of 275 degrees-of-freedom between the two extrema. Recall that these differences in degrees-of-freedom render conventional ROM approaches as impossibly to apply on heterogeneous databases such as the one considered here.

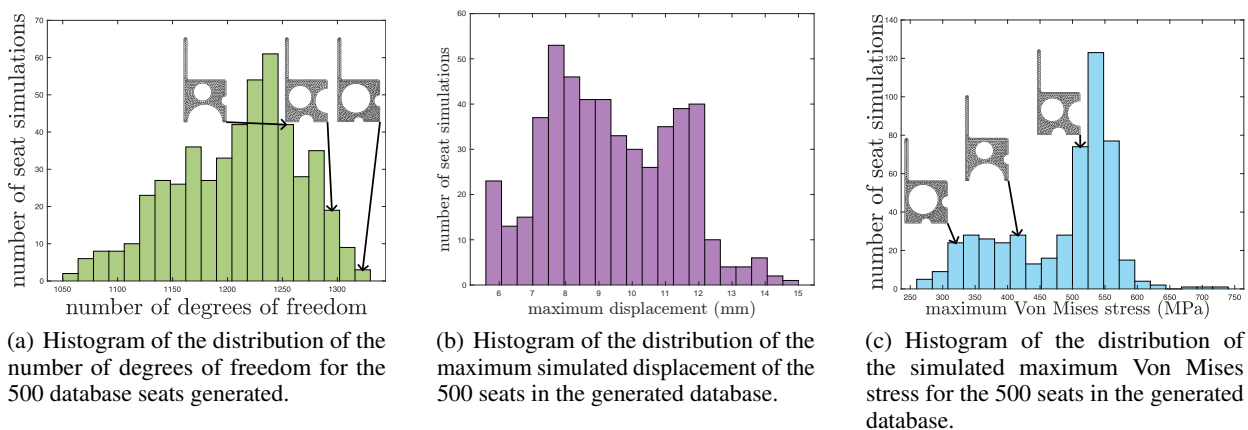


Figure 8: Statistical analysis of the database used for training, validation, and testing.

The middle and right panels of Figure 8 also highlight the diversity in the overall mechanical behavior of the seats based on displacement and maximum Von Mises stress, respectively. Furthermore, the distribution of seats is clearly non-Gaussian. This variability poses a significant challenge for learning and understanding the underlying patterns in seat mechanics.

4.2 GNN-PGD implementation details and parameters

This section details the specific parameters employed in this case study for the GNN-PGD model introduced in subsection 3.2.2. All algorithms were implemented with the pytorch library [85] using pytorch-geometric [86]. The input matrix $\mathbf{X} \in \mathbb{R}^{n \times k}$ (Table 1) is composed of information on the coordinates of the node (two components, x and y), the type of node (free or embedded, as a boolean), the maximum-in-time of the generalized effort $\mathbf{f}(t)$ at node i (two spatial components, f_x and f_y). Hence at each node, there are $k = 5$ elements of input. Recall that the generic output of the GNN-PGD generator is $\mathbf{P} \in \mathbb{R}^{n \times g}$; hence the projection of the $M = 3$ PGD modes here associated with each seat yields $g = 3 \times 2 = 6$.

For improving learning performance of our GNN-PGD method, the input data is also pre-processed as follows. The input quantities (\mathbf{X}) are normalized according to the mean and standard deviation statistics of the training set, and the output quantities (\mathbf{P}) are standardized according to the maximum and minimum of the training set.

Hyperparameters	H	L	η_1	η_2	η_3	α	γ
Values	24	15	10	1	100	3×10^{-3}	5×10^{-3}

Table 4: The final hyperparameters employed for the results of this section.

The hyperparameters determined for these test cases are produced by a DOE (as described in subsection 3.2.2) for 100 different models (i.e., with different hyperparameters on each) over 700 epochs. The table 4 shows the resulting hyperparameters associated with the best model, which is then utilized for training over a larger number of epochs until an early learning stopping criterion is reached. The final GNN-PGD generator has a number of 65 022 learnable parameters.

Figure 9 presents the corresponding evolution of the loss as a function of epochs for the corresponding GNN-PGD model during the training phase (see Equation 14). We note that the GNN-PGD model updates its parameters as many times as there are batches in the training set (400 in our case). The training has been carried out on 4 Nvidia V100 GPUs and lasted 4 hours and 25 minutes (including hyperparameter optimization).

4.3 GNN-PGD learning results

We analyze the results of the GNN-PGD training by evaluating its performance on the 50 seats in the test set $\{\Omega_p^{\text{test}}\}_{p=\{1,\dots,50\}}$, i.e., seats unknown to and unseen by the GNN-PGD. The GNN-PGD's objective is to construct the $\{\mathbf{P}_p^{\text{test}}\}_{p=\{1,\dots,50\}}$ ROBs of the 50 as best as it can. As described earlier, these ROBs are composed of $M = 3$ PGD spatial modes (to be predicted by the GNN-PGD). We quantify the error between the GNN-PGD's inference prediction $\{\mathbf{P}_p^{\text{GNN}}\}_{p=\{1,\dots,50\}}$ and the labeled ROBs $\{\mathbf{P}_p^{\text{test}}\}_{p=\{1,\dots,50\}}$ obtained from the reference field decomposition. The metric used is the normalized root-mean-square error (RMSE) for each of the $M = 3$ modes predicted, and the results are presented in Figure 10.

Results demonstrate the ability of GNN-PGD to predict with good accuracy the ROBs associated with the unknown (test) seat geometries. Indeed, errors on predicted modes are small: less than 5% for Mode 1, less than 10% Mode 2, and less than 15% for Mode 3. We note that, as evidenced in Figure 10, the errors increase on average with the rank of the predicted mode. An analysis of the seats that lie on the higher end of the error plots yields no discernible pattern to the particular geometries that have slightly higher errors. A visual comparison with the corresponding true reference modes ("Ground Truth"), including a plot of the difference, of the first, second and third spatial mode solutions predicted by GNN-PGD are provided in Figures 11, 12, and 13, respectively. These figures demonstrate the ability of the proposed GNN-PGD solver to very accurately capture both coarse and fine behavior found in the very different modes. Figure 14 additionally presents a similar comparison for Mode 3 on a different geometry than Figure 13, further illustrating the large variations that can be found in the various modes between geometries, and the ability of GNN-PGD to accurately capture and predict such disparate behavior between different seats. The results in general indicate the ability of GNN-PGD to predict well a ROB for a new geometry, which can then be used to obtain the complete associated spatio-temporal field via Galerkin projection (subsection 3.3).




Figure 9: Evolution of the loss function of the GNN-PGD generator associated with the hyperparameters of Table 4.




Figure 10: Results on the modes learned by the GNN-PGD.

4.3.1 Error of the fully-reconstructed space-time fields

Using the ROBs predicted above, we analyze the complete space-time field error associated with the test geometries $\{\Omega_p^{\text{test}}\}_{p=\{1,\dots,50\}}$. As described in subsection 3.3, the temporal component can be obtained from a Galerkin projection of the predicted ROBs $\{P_p^{\text{GNN}}\}_{p=\{1,\dots,50\}}$. The resulting space-time solutions, denoted $\{U_p^{\text{GNN}}\}_{p=\{1,\dots,50\}}$, are compared with the FEM reference fields $\{U_p^{\text{test}}\}_{p=\{1,\dots,50\}}$. The resulting normalized space-time $L^2(\Omega, \mathcal{J})$ errors (i.e., over all space and time) are presented in Figure 15. One can observe $< 5\%$ errors for all tested seat geometries $\{\Omega_p^{\text{test}}\}_{p=\{1,\dots,50\}}$. Such performance is promising from an iterative engineering design standpoint, as discussed in the introduction.

4.3.2 Comparison with other approaches

We compare the results of our proposed GNN-PGD solver with commonly-found autoregressive approaches from literature [54, 58, 80]. All other GNN methods, as far as the authors know, construct the temporal solution timestep-by-




Figure 11: Example of the first mode predicted by GNN-PGD for seat 1 from the test database, together with the ground truth and error.




Figure 12: Example of the second mode predicted by GNN-PGD for seat 10 from the test database, together with the ground truth and error.

timestep by feeding back the prediction of the time step t_i as an input to the GNN in order to predict the next time step t_{i+1} . A summary of these standard approaches is illustrated in Figure 16.

We similarly implemented the standard autoregressive approach (GNN-AR) via pytorch, following [80, 58]. The same hyperparameter optimisation strategy (with a DOE, as described in subsection 3.2.2) is employed in order to train the GNN-AR for comparative purposes. Figure 17 presents a comparison of temporal evolution of the spatial L^2 errors between the GNN-PGD introduced in this work and GNN-AR for all 50 test cases. Results demonstrate a well-known drawback of the timestep-by-timestep GNN-AR approach: an error increase as a function of time (*rollout error increase* [80, 58])—for our test cases, up to 15%. Conversely, our proposed GNN-PGD approach does not suffer from such an increase in error over time; in fact, it remains bounded with a maximum error of only 3% over all timesteps. Additionally, we find that the learning time (offline time) associated with the GNN-AR approach (16h40) is significantly higher than the GNN-PGD approach (4h25). This is expected since the autoregressive approach considers all $N_t = 1000$ snapshots of each geometry in the training database, whereas GNN-PGD only considers the $M = 3$ spatial modes associated with these geometries (and reconstructs the temporal component via projection and a computationally-inexpensive resolution of the corresponding ODEs of Equation (16)). Another advantage of the GNN-PGD approach over the




Figure 13: Example of the third mode predicted by GNN-PGD for seat 23 from the test database, together with the ground truth and error.

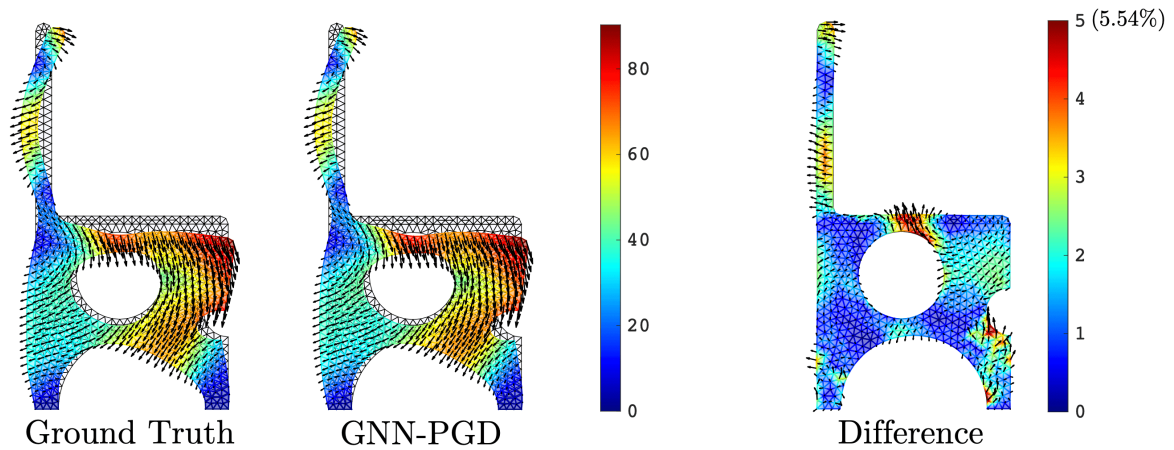


Figure 14: Example of the first mode predicted by GNN-PGD for seat 44 from the test database, together with the ground truth and error.

GNN-AR approach is the number of calls to the GNN network. In the case of an autoregressive approach, the number of calls to the network is linked to the number of time steps required in the simulation (e.g., 1000), whereas for the GNN-PGD approach, only one evaluation is required, which translates into a shorter inference time (online time) for the GNN-PGD approach.

4.3.3 Generalization ability of GNN-PGD to infer other topologies

In order to test GNN-PGD’s generalization capabilities for predicting a ROB on topologies that are highly different from those with which it has been trained (including very different discretizations), we consider the seats illustrated in Figure 18, which contain two or three holes in the interior of the mesh (as opposed to only one used in training, validation, and testing, see Figure 7).

Figures 19 illustrates and compares Mode 1 to the ground truth for two circular holes in a seat. Clearly, GNN-PGD is able to capture the general behavior of the mode, with a maximum error of only 15%. Similarly, 20 illustrates and compares Mode 2 for a geometry containing three holes, which also indicates a good approximation to the ground truth.




Figure 15: Distribution of the normalized space-time $L^2(\Omega, \mathcal{J})$ error for the reconstructed fields of the test geometries $\{\Omega_p^{\text{test}}\}_{p=\{1, \dots, 50\}}$.




Figure 16: Illustration of a standard GNN-AR autoregressive approach employed as a benchmark comparison for the proposed GNN-PGD method.




Figure 17: Comparison of the time evolution of the spatial $L^2(\Omega)$ error between the GNN-AR approach [80, 58] and our proposed GNN-PGD approach for the test set $\{\Omega_p^{\text{test}}\}_{p=\{1, \dots, 50\}}$.




Figure 18: Example of three seats, exterior to the original database, of very different topologies (i.e., multiple holes).

A quantification of the L^2 spatial error for all predicted modes for each of these two cases (two and three holes) is presented in Table 5, along with the complete space-time $L^2(\Omega, \mathcal{J})$ for the final reconstructed solutions.




Figure 19: Mode 1 predicted by GNN-PGD (trained on one hole) on a topology with two holes.



	Mode 1 $L^2(\Omega)$ error	Mode 2 $L^2(\Omega)$ error	Mode 3 $L^2(\Omega)$ error	\mathbf{U}^{GNN} $L^2(\Omega, \mathcal{J})$ error
	16%	15%	30%	16%
	17%	16%	27%	18%

Table 5: Errors of the GNN-PGD approach for topologies containing more holes than the training data set.




Figure 20: Mode 2 predicted by GNN-PGD (trained on one hole) on a topology with three holes.

Although the errors are much higher than those presented for a single hole (on which our GNN-PGD was trained), they are entirely reasonable from an iterative design standpoint. Indeed, for the two new topologies tested, we obtain a final error over the entire space-time field of less than 18%, which demonstrates the remarkable capability of the GNN-PGD solver to generalize to completely different geometries. For pre-dimensioning or pre-sizing of structures, such errors are more than sufficient to approximate various mechanical quantities without having to run full FEM simulations.

5 Conclusion

This contribution introduces a new methodology that combines a Galerkin projection-based model reduction method with a deep learning algorithm based on GNNs in order to generate a ROB for an arbitrary geometry of arbitrary discretization size. Results demonstrate that the proposed GNN-PGD solver can produce reduced order models for non-parametric geometries (where usual model reduction tools are limited [36, 43]). The algorithm proposed and analyzed in this work carries significant advantages over more conventional GNN methods, both in terms of computational cost (training and inference) as well as numerical error. A case study related to the motivating engineering and industrial applications of this work (structural design) has also been presented, using a reasonably-sized database that results in a GNN-PGD solver that provides low errors that are appropriate for pre-dimensioning or pre-sizing of, e.g., aircraft seats. For such configurations, we have additionally demonstrated a strong capability of the GNN-PGD method to generalize to topologies that are completely different from training and validation databases.

The proof-of-concept provided in this contribution demonstrates promise for future extensions and other applications. For example, one can incorporate more physics into the GNN-PGD generator so as to be able to generate more modes with even greater accuracy. Towards this, future work is planned for the inclusion of finite element operators, derived from the physical PDEs, as direct inputs to the graph (inspired by [87]). More importantly, future work entails extending the GNN-PGD method to non-linear problems, e.g., the material non-linearities that are very present in the aeronautical world [11, 12, 13]. For such contexts, a possible enrichment of the predicted space-time solution (post-processing) may also be required.

CRedit authorship contribution statement

Victor Matray: Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Investigation, Formal Analysis, Conceptualization, Visualization. **Faisal Amlani:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Conceptualization. **Frédéric Feyel:** Writing – review & editing, Supervision, Project administration, Conceptualization. **David Néron:** Writing – review & editing, Supervision, Project administration, Conceptualization.

Acknowledgements

This work was granted access to the HPC resources of IDRIS under the allocation 2023-[AD011014260] made by GENCI. This work was performed using HPC resources from the “Mésocentre” computing center of CentraleSupélec, École Normale Supérieure Paris-Saclay and Université Paris-Saclay supported by CNRS and Région Île-de-France (<https://mesocentre.universite-paris-saclay.fr/>). Finally, the authors would like to express their sincere gratitude to Safran Tech for supporting this research and for providing the motivating application contexts.

References

- [1] Lovely Sabat and Chinmay Kumar Kundu. History of Finite Element Method: A Review. In Bibhuti Bhusan Das, Salim Barbhuiya, Rishi Gupta, and Purnachandra Saha, editors, *Recent Developments in Sustainable Infrastructure*, pages 395–404, Singapore, 2021.
- [2] John P Wolf. *The scaled boundary finite element method*. John Wiley & Sons, 2003.
- [3] Erik Burman, Susanne Claus, Peter Hansbo, Mats G Larson, and André Massing. Cutfem: discretizing geometry and partial differential equations. *International Journal for Numerical Methods in Engineering*, 104(7):472–501, 2015.
- [4] John P Wolf and Chongmin Song. The scaled boundary finite-element method—a fundamental solution-less boundary-element method. *Computer Methods in Applied Mechanics and Engineering*, 190(42):5551–5568, 2001.
- [5] Wing Kam Liu, Shaofan Li, and Harold S Park. Eighty years of the finite element method: Birth, evolution, and future. *Archives of Computational Methods in Engineering*, 29(6):4431–4453, 2022.
- [6] Bernardo Cockburn, George E Karniadakis, and Chi-Wang Shu. *Discontinuous Galerkin methods: theory, computation and applications*, volume 11. Springer Science & Business Media, 2012.
- [7] Thomas-Peter Fries and Ted Belytschko. The intrinsic xfem: a method for arbitrary discontinuities without additional unknowns. *International journal for numerical methods in engineering*, 68(13):1358–1385, 2006.
- [8] Usik Lee. *Spectral element method in structural dynamics*. John Wiley & Sons, 2009.
- [9] Ted Belytschko, Robert Gracie, and Giulio Ventura. A review of extended/generalized finite element methods for material modeling. *Modelling and Simulation in Materials Science and Engineering*, 17(4):043001, 2009.
- [10] Bofang Zhu. *The finite element method: fundamentals and applications in civil, hydraulic, mechanical and aeronautical engineering*. John Wiley & Sons, 2018.
- [11] Georgios Tzanakis, Athanasios Kotzakolios, Efthimis Giannaros, and Vassilis Kostopoulos. Structural Analysis of a Composite Passenger Seat for the Case of an Aircraft Emergency Landing. *Applied Mechanics*, 4(1):1–19, March 2023. Number: 1 publisher2: Multidisciplinary Digital Publishing Institute.
- [12] Luca Cavagna, Sergio Ricci, and Lorenzo Travaglini. Neocass: an integrated tool for structural sizing, aeroelastic analysis and mdo at conceptual design level. *Progress in Aerospace Sciences*, 47(8):621–635, 2011.
- [13] B Fredriksson. Advanced numerical methods for analysis and design of aircraft structures. *International journal of vehicle design*, 7(3-4):306–336, 1986.
- [14] Esmā Karagoz, Jimmy Tai, Darshan Sarojini, and Dimitri Mavris. Design space reduction using clustering in aircraft engine design. In *2019 IEEE Aerospace Conference*, pages 1–8. IEEE, 2019.
- [15] L Zhu, N Li, and PRN Childs. Light-weighting in aerospace component and system design. *Propulsion and Power Research*, 7(2):103–119, 2018.
- [16] Sergio Andrés Ardila-Parra, Carmine Maria Pappalardo, Octavio Andrés González Estrada, and Domenico Guida. Finite element based redesign and optimization of aircraft structural components using composite materials. *IAENG International Journal of Applied Mathematics*, 2020.
- [17] Peter Benner, Serkan Gugercin, and Karen Willcox. A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems. *SIAM Review*, 57:483–531, June 2015.
- [18] Francisco Chinesta, Pierre Ladeveze, and Elías Cueto. A Short Review on Model Order Reduction Based on Proper Generalized Decomposition. *Archives of Computational Methods in Engineering*, 18(4):395–404, November 2011.
- [19] Gianluigi Rozza. Reduced basis approximation and error bounds for potential flows in parametrized geometries. *Communications in Computational Physics*, 9(1):1–48, 2011.
- [20] Yvon Maday and Einar M Rønquist. A reduced-basis element method. *Journal of scientific computing*, 17:447–459, 2002.
- [21] Gianluigi Rozza, M Hess, G Stabile, M Tezzele, F Ballarin, C Gräßle, M Hinze, S Volkwein, F Chinesta, P Ladeveze, et al. Volume 2 snapshot-based methods and algorithms, 2020.
- [22] David Amsallem, Matthew J Zahr, and Charbel Farhat. Nonlinear model order reduction based on local reduced-order bases. *International Journal for Numerical Methods in Engineering*, 92(10):891–916, 2012.
- [23] Kevin Carlberg, Matthew Barone, and Harbir Antil. Galerkin v. least-squares petrov–galerkin projection in nonlinear model reduction. *Journal of Computational Physics*, 330:693–734, 2017.

- [24] Lawrence Sirovich. Turbulence and the dynamics of coherent structures. i. coherent structures. *Quarterly of applied mathematics*, 45(3):561–571, 1987.
- [25] Kuan Lu, Yulin Jin, Yushu Chen, Yongfeng Yang, Lei Hou, Zhiyong Zhang, Zhonggang Li, and Chao Fu. Review for order reduction based on proper orthogonal decomposition and outlooks of applications in mechanical systems. *Mechanical Systems and Signal Processing*, 123:264–297, 2019.
- [26] Gaetan Kerschen, Jean-claude Golinval, Alexander F Vakakis, and Lawrence A Bergman. The method of proper orthogonal decomposition for dynamical characterization and order reduction of mechanical systems: an overview. *Nonlinear Dynamics*, 41:147–169, 2005.
- [27] Annika Radermacher and Stefanie Reese. A comparison of projection-based model reduction concepts in the context of nonlinear biomechanics. *Archive of Applied Mechanics*, 83(8):1193–1213, August 2013.
- [28] Jan S Hesthaven, Gianluigi Rozza, Benjamin Stamm, et al. *Certified reduced basis methods for parametrized partial differential equations*, volume 590. Springer, 2016.
- [29] Paolo Tiso and Daniel J. Rixen. Discrete Empirical Interpolation Method for Finite Element Structural Dynamics. In Gaetan Kerschen, Douglas Adams, and Alex Carrella, editors, *Topics in Nonlinear Dynamics, Volume 1*, Conference Proceedings of the Society for Experimental Mechanics Series, pages 203–212, New York, NY, 2013.
- [30] P Ladevèze. Sur une famille d’algorithmes en mécanique des structures. *Sur une famille d’algorithmes en mécanique des structures*, 300(2):41–44, 1985. Place: Paris publisher2: Gauthier-Villars.
- [31] A. Daby-Seesaram, A. Fau, P.-E. Charbonnel, and D. Néron. A hybrid frequency-temporal reduced-order method for nonlinear dynamics. *Nonlinear Dynamics*, 111(15):13669–13689, August 2023.
- [32] Pierre Ladevèze. *Nonlinear Computational Structural Mechanics: New Approaches and Non-Incremental Methods of Calculation*. Springer Science & Business Media, 1999.
- [33] Ronan Scanff, David Néron, Pierre Ladevèze, Philippe Barabinot, Frédéric Cugnon, and Jean-Pierre Delsemme. Weakly-invasive LATIN-PGD for solving time-dependent non-linear parametrized problems in solid mechanics. *Computer Methods in Applied Mechanics and Engineering*, 396:114999, June 2022.
- [34] Christophe Heyberger, Pierre-Alain Boucard, and David Néron. Multiparametric analysis within the proper generalized decomposition framework. *Computational Mechanics*, 49(3):277–289, 2012.
- [35] Fabien Casenave, Asven Gariah, Christian Rey, and Frederic Feyel. A nonintrusive reduced order model for nonlinear transient thermal problems with nonparametrized variability. *Advanced Modeling and Simulation in Engineering Sciences*, 7(1):22, May 2020.
- [36] Gianluigi Rozza, Haris Malik, Nicola Demo, Marco Tezzele, Michele Girfoglio, Giovanni Stabile, and Andrea Mola. Advances in reduced order methods for parametric industrial problems in computational fluid dynamics. *arXiv preprint arXiv:1811.08319*, 2018.
- [37] Andrea Lupini and Bogdan I. Epureanu. On the use of mesh morphing techniques in reduced order models for the structural dynamics of geometrically mistuned blisks. *Mechanical Systems and Signal Processing*, 127:262–275, July 2019.
- [38] Thomas W Sederberg and Scott R Parry. Free-form deformation of solid geometric models. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, pages 151–160, 1986.
- [39] Gianluigi Rozza, Anwar Koshakji, Alfio Quarteroni, et al. Free form deformation techniques applied to 3d shape optimization problems. *Communications in Applied and Industrial Mathematics*, 4:1–26, 2013.
- [40] Martin Dietrich Buhmann. Radial basis functions. *Acta numerica*, 9:1–38, 2000.
- [41] AM Morris, CB Allen, and TCS Rendall. CFD-based optimization of aerofoils using radial basis functions for domain element parameterization and mesh deformation. *International Journal for Numerical Methods in Fluids*, 2008.
- [42] Toni Lassila and Gianluigi Rozza. Parametric free-form shape design with PDE models and reduced basis method. *Computer Methods in Applied Mechanics and Engineering*, 199(23):1583–1592, April 2010.
- [43] Courard Amaury. *PGD-Abaques virtuels pour l’optimisation géométrique des structures*. PhD thesis, Université Paris-Saclay, 2016.
- [44] Dongwei Ye, Valeria Krzhizhanovskaya, and Alfons G. Hoekstra. Data-driven reduced-order modelling for blood flow simulations with geometry-informed snapshots. *Journal of Computational Physics*, 497:112639, January 2024.

- [45] Fabien Casenave, Brian Staber, and Xavier Roynard. Mmgrp: a mesh morphing gaussian process-based machine learning method for regression of physical problems under nonparametrized geometrical variability. *Advances in Neural Information Processing Systems*, 36, 2024.
- [46] Marc Alexa. Recent advances in mesh morphing. In *Computer graphics forum*, volume 21, pages 173–198. Wiley Online Library, 2002.
- [47] George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
- [48] Elias Cueto and Francisco Chinesta. Thermodynamics of learning physical phenomena. *Archives of Computational Methods in Engineering*, 30(8):4653–4666, 2023.
- [49] Lu Lu, Pengzhan Jin, and George Em Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*, 2019.
- [50] Shashank Reddy Vadyala, Sai Nethra Betgeri, John C Matthews, and Elizabeth Matthews. A review of physics-based machine learning in civil engineering. *Results in Engineering*, 13:100316, 2022.
- [51] M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, February 2019.
- [52] Quercus Hernández, Alberto Badías, Francisco Chinesta, and Elías Cueto. Thermodynamics-informed neural networks for physically realistic mixed reality. *Computer Methods in Applied Mechanics and Engineering*, 407:115912, 2023.
- [53] Antoine Benady, Emmanuel Baranger, and Ludovic Chamoin. Nn-mcre: A modified constitutive relation error framework for unsupervised learning of nonlinear state laws with physics-augmented neural networks. *International Journal for Numerical Methods in Engineering*, 125(8):e7439, 2024.
- [54] Quercus Hernández, Alberto Badías, Francisco Chinesta, and Elías Cueto. Thermodynamics-informed graph neural networks. *IEEE Transactions on Artificial Intelligence*, pages 1–1, 2022. arXiv:2203.01874 [cs, math].
- [55] Ali Kashefi, Davis Rempe, and Leonidas J. Guibas. A Point-Cloud Deep Learning Framework for Prediction of Fluid Flow Fields on Irregular Geometries. *Physics of Fluids*, 33(2):027104, February 2021. arXiv:2010.09469 [physics].
- [56] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- [57] Zijie Li and Amir Barati Farimani. Graph neural network-accelerated Lagrangian fluid simulation. *Computers & Graphics*, 103:201–211, April 2022.
- [58] Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter W Battaglia. Learning mesh-based simulation with graph networks. *arXiv preprint arXiv:2010.03409*, 2020.
- [59] Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Alexander Pritzel, Suman Ravuri, Timo Ewalds, Ferran Alet, Zach Eaton-Rosen, Weihua Hu, Alexander Merose, Stephan Hoyer, George Holland, Jacklynn Stott, Oriol Vinyals, Shakir Mohamed, and Peter Battaglia. GraphCast: Learning skillful medium-range global weather forecasting, December 2022. arXiv:2212.12794 [physics].
- [60] David Dalton, Hao Gao, and Dirk Husmeier. Emulation of cardiac mechanics using Graph Neural Networks. *Computer Methods in Applied Mechanics and Engineering*, 401:115645, November 2022.
- [61] Suresh Bishnoi, Ravinder Bhattoo, Sayan Ranu, and NM Krishnan. Enhancing the inductive biases of graph neural ode for modeling dynamical systems. *arXiv preprint arXiv:2209.10740*, 2022.
- [62] Saurabh Deshpande, Stéphane PA Bordas, and Jakub Lengiewicz. Magnet: A graph u-net architecture for mesh-based simulations. *Engineering Applications of Artificial Intelligence*, 133:108055, 2024.
- [63] Rutwik Gulakala, Bernd Markert, and Marcus Stoffel. Graph neural network enhanced finite element modelling. *PAMM*, 22(1):e202200306, 2023.
- [64] Matthieu Nastorg, Michele-Alessandro Bucci, Thibault Faney, Jean-Marc Gratien, Guillaume Charpiat, and Marc Schoenauer. An implicit gnn solver for poisson-like problems. *arXiv preprint arXiv:2302.10891*, 2023.
- [65] Yu Zhang, Peter Tiño, Aleš Leonardis, and Ke Tang. A survey on neural network interpretability. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 5(5):726–742, 2021.

- [66] Nathan M Newmark. A method of computation for structural dynamics. *Journal of the engineering mechanics division*, 85(3):67–94, 1959.
- [67] L. Boucinha, A. Gravouil, and A. Ammar. Space–time proper generalized decompositions for the resolution of transient elastodynamic models. *Computer Methods in Applied Mechanics and Engineering*, 255:67–88, March 2013.
- [68] Anthony Nouy. A priori model reduction through Proper Generalized Decomposition for solving time-dependent partial differential equations. *Computer Methods in Applied Mechanics and Engineering*, 199(23-24):1603–1626, April 2010.
- [69] William L Hamilton. *Graph representation learning*. Morgan & Claypool Publishers, 2020.
- [70] Yann LeCun, Koray Kavukcuoglu, and Clément Farabet. Convolutional networks and applications in vision. In *Proceedings of 2010 IEEE international symposium on circuits and systems*, pages 253–256. IEEE, 2010.
- [71] Maciej Besta and Torsten Hoefer. Parallel and distributed graph neural networks: An in-depth concurrency analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [72] Frank Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386, 1958.
- [73] Stefan Elfving, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural networks*, 107:3–11, 2018.
- [74] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018.
- [75] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [76] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.
- [77] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. pmlr, 2015.
- [78] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [79] Claudio Filipi Gonçalves Dos Santos and João Paulo Papa. Avoiding overfitting: A survey on regularization methods for convolutional neural networks. *ACM Computing Surveys (CSUR)*, 54(10s):1–25, 2022.
- [80] Johannes Brandstetter, Daniel Worrall, and Max Welling. Message passing neural pde solvers. *arXiv preprint arXiv:2202.03376*, 2022.
- [81] V. Roshan Joseph, Evren Gul, and Shan Ba. Maximum projection designs for computer experiments. *Biometrika*, 102(2):371–380, June 2015.
- [82] Stéphane Roux and François Hild. Optimal procedure for the identification of constitutive parameters from experimentally measured displacement fields. *International Journal of Solids and Structures*, 184:14–23, February 2020.
- [83] Leslie Rice, Eric Wong, and Zico Kolter. Overfitting in adversarially robust deep learning. In *International conference on machine learning*, pages 8093–8104. PMLR, 2020.
- [84] Christophe Geuzaine and Jean-François Remacle. Gmsh: A 3-d finite element mesh generator with built-in pre-and post-processing facilities. *International journal for numerical methods in engineering*, 79(11):1309–1331, 2009.
- [85] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- [86] Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. *arXiv preprint arXiv:1903.02428*, 2019.
- [87] Balthazar Donon, Zhengying Liu, Wenzhuo Liu, Isabelle Guyon, Antoine Marot, and Marc Schoenauer. Deep statistical solvers. *Advances in Neural Information Processing Systems*, 33:7910–7921, 2020.