CSMA 2026
17ème Colloque National en Calcul des Structures
18-22 Mai 2026, Presqu'île de Giens (Var)

A hybrid numerical framework combining graph neural networks & classical reduced-order models for finite element systems in dynamics

V. Matray[1], D. Néron[1], F. Feyel[1,2], F. Amlani[1]

[1] *Université Paris-Saclay, CentraleSupélec, ENS Paris-Saclay, CNRS, LMPS - Laboratoire de Mécanique Paris-Saclay, 91190, Gif-sur-Yvette {victor.matray, david.neron, faisal.amlani}@ens-paris-saclay.fr*
[2] *Safran Tech, Digital Sciences & Technologies Department, Rue des jeunes bois, Châteaufort, 78114, Magny-les-Hameaux frederic.feyel@safrangroup.com*

Résumé — This contribution presents recent work on building a hybrid Graph Neural Network (GNN)-based reduced-order modeling framework for solving time-dependent partial differential equations on non-parametric geometries. The method exploits graph learning to predict reduced bases in a lightweight architecture that embeds finite element operators, geodesic subspace distance measures, and Gated Recurrent Units (GRUs). A new "Boosted PGD" enrichment step provides fast, on-the-fly error correction. Efficacy is demonstrated on datasets containing wide topological variations and discretization sizes.
Mots clés — graph-based deep learning, reduced-order modeling, non-parametric meshes, digital twins

1 Introduction

Reduced-order modeling (ROM) techniques [1, 2, 3, 4] offer approaches to provide fast and efficient partial differential equation (PDE) approximations for problems that require large numbers of computationally-expensive full-field finite element method (FEM) analyses/simulations (e.g., for iterative design or inverse/optimization problems). By constructing low-dimensional representations of solution fields, ROM methods significantly reduce computational costs in terms of both memory and complexity, while maintaining reasonable accuracy [1, 2]. Common approaches—such as the Proper Orthogonal Decomposition (POD), the Proper Generalized Decomposition (PGD), or the Reduced Basis method (RBM)—utilize parameterized snapshots from select full-field simulations in order to create reduced-order bases that can approximate solutions for new configurations [2, 3, 4]. These methods are well-suited for problems where the geometry and material properties can be parameterized, e.g., where shapes are defined by a finite set of control points or deformation parameters. However, when geometries lack explicit parameterizations (e.g., for large geometric variations or changes in topology), ROM approaches face a significant challenge : classical projection methods often rely on constant spatial discretizations (i.e., fixed degrees of freedom) across training and prediction datasets.

This contribution summarizes recent extensions [5] of a proof-of-concept deep learning-based ROM approach [6] for the numerical solution of time-dependent PDEs on non-parametric finite element meshes. Such a framework is based on a a novel hybridization of both Galerkin projection-based ROM and Graph Neural Networks (GNNs [7]) (where the latter is used to predict the former), enabling treatment of discrete geometries of arbitrary size towards accommodating completely unstructured meshes with variable degrees of freedom (and topologies). The underlying methodology can be applied to predict spatial reduced-order bases (ROBs) for any Galerkin-based method (e.g., POD, PGD), and can be enhanced to [5] : tightly integrate graph-based learning with physical information (by using finite element operators as node and edge level features) ; minimize a Grassmannian (geodesic) distance measure (for penalizing mismatches in subspaces) ; and incorporate a Gated Recurrent Unit (GRU) (for a more efficient and lightweight architecture). A novel, on-the-fly enrichment mechanism, modified from a classical PGD and dubbed "Boosted PGD" [5], can also be used to improve prediction accuracy at low computational cost via additional greedy corrective modes (no matter the original basis model). Solvers constructed by such a hybrid approach demonstrate very competitive accuracy and computational cost in simulating dynamic behavior when compared to conventional full-order FEM, with a capacity to generalize to configurations well outside of the topological scope of the original training and validation geometries/configurations.

2 Context

A number of time-dependent PDEs in both solid and fluid mechanics can be solved using a standard spatial semi-discretization based on a classical finite element method (FEM). Consider a discrete spatial domain (geometry) of interest $\boldsymbol{\Omega} \subset \mathbb{R}^d$ $(d = 1, 2, 3)$ with corresponding boundary $\partial\boldsymbol{\Omega} = \partial\boldsymbol{\Omega}_{\underline{u}} \cup \partial\boldsymbol{\Omega}_{\underline{\sigma}}$ (here, $\partial\boldsymbol{\Omega}_{\underline{u}}$ and $\partial\boldsymbol{\Omega}_{\underline{\sigma}}$ denote boundary regions where Dirichlet or Neumann conditions are imposed, respectively, with $\partial\boldsymbol{\Omega}_{\underline{u}} \cap \partial\boldsymbol{\Omega}_{\underline{\sigma}} = \varnothing$). For a second-order time-dependent PDE (e.g., Eulerian elastodynamics), the classical finite element formulation yields, for discrete unknowns $\mathbf{U}(t)$ (e.g., elastic displacement), the well-known ordinary differential equation (ODE) system in time given by

$$\begin{cases} \mathbb{M}\ddot{\mathbf{U}}(t) + \mathbb{K}\mathbf{U}(t) = \mathbf{f}(t), \\ \mathbf{U}(0) = \mathbf{U}_0, \\ \dot{\mathbf{U}}(0) = \mathbf{V}_0, \end{cases} \tag{1}$$

for temporal coefficient vector $\mathbf{U}(t) \in \mathbb{R}^N$ of a chosen finite element basis representation (evaluated at all discrete values $\boldsymbol{x} \in \boldsymbol{\Omega}$), mass matrix $\mathbb{M} \in \mathbb{R}^{N \times N}$, stiffness matrix $\mathbb{K} \in \mathbb{R}^{N \times N}$, and generalized force vector $\mathbf{f}(t) \in \mathbb{R}^N$ (evaluated at discrete spatial values). Here, N represents the number of degrees-of-freedom (the product of the number of unknown physical quantities and the size of the spatial discretization). The vectors $\mathbf{U}_0 \in \mathbb{R}^N$ and $\mathbf{V}_0 \in \mathbb{R}^N$ correspond to the coefficients of the initial conditions in the finite element basis. For further mathematical details including an explicit presentation of the weak formulation employed in this work, the reader is referred to [5, 6] and the references therein. The final ODE given by Eq. (1) can be solved by any suitable temporal integration scheme, e.g., a conservative and stable Newmark integration scheme [6], where the time domain $\boldsymbol{I} = [0, T]$ is discretized into N_t time steps of size Δt with an implicit formulation employing average acceleration (i.e., $\gamma = 1/2$ and $\beta = 1/4$).

2.1 Galerkin projection-based reduced-order modeling

In order to reduce the computational cost associated with solving the high-dimensional FEM system in Eq. (1) (hereafter referred to as the *full-order model* or *FOM*), which can become cumbersome for large-scale systems or for iterative problems, Galerkin-based reduced-order modeling [1, 2, 3, 4] seeks to approximate the full solution within a low-dimensional subspace spanned by problem-adapted basis functions, i.e., the unknown field is represented as

$$\mathbf{U}(t) \approx \boldsymbol{\Phi}\,\mathbf{a}(t) = \sum_{i=1}^{M} a_i(t)\,\boldsymbol{\phi}_i, \tag{2}$$

where the columns of the matrix $\boldsymbol{\Phi} \in \mathbb{R}^{N \times M}$ contain $M \ll N$ reduced modes obtained, for example, via the Proper Orthogonal Decomposition (POD), the Proper Generalized Decomposition (PGD), or a similar Galerkin-based method. Substituting this ansatz into the full-order system and enforcing a Galerkin projection onto $\text{span}\{\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_M\}$ yields the reduced system given by

$$\mathbf{M}_M \ddot{\mathbf{a}}(t) + \mathbf{K}_M\,\mathbf{a}(t) = \mathbf{f}_M(t), \tag{3}$$

where $\mathbf{M}_M = \boldsymbol{\Phi}^\top \mathbb{M} \boldsymbol{\Phi}$, $\mathbf{K}_M = \boldsymbol{\Phi}^\top \mathbb{K} \boldsymbol{\Phi}$, $\mathbf{f}_M = \boldsymbol{\Phi}^\top \mathbf{f}(t)$. Such projection-based reduction preserves the dominant energetic or dynamical features of the full-order model while drastically decreasing the dimension of the system to be integrated. As a result, the reduced coordinates $\mathbf{a}(t)$ evolve according to a much smaller ODE system.

3 A GNN-based framework for predicting reduced-order bases

Since the reduced basis $\boldsymbol{\Phi}$ in Eq. (2) is constructed from snapshots or modes defined on a particular finite element discretization, the resulting reduced-order model is inherently tied to that fixed mesh. Changes in geometry, topology, or spatial discretization size typically require recomputing (or transferring) the basis, limiting the applicability of classical Galerkin ROMs to such changes (e.g., in non-paremetric or design-driven contexts). However, recent advancements in deep learning, particularly those

based on graph neural networks (GNNs) [5, 6, 7], are well adapted to provide mesh-agnostic predictions by learning the local relationships between graph nodes or verticies (e.g., a mesh), allowing a network trained on one graph to easily generalize and apply to another (including of different size). Hence one can apply a recently-introduced geometry-agnostic strategy [5, 6] in which a GNN predicts the reduced spatial basis $\boldsymbol{\Phi}$ directly on an arbitrary finite element discretization. In particular, the general idea is to learn a mapping

$$\mathcal{M} \;\mapsto\; \boldsymbol{\Phi} = [\phi_1, \ldots, \phi_M],$$

where $\boldsymbol{\Phi}$ is a reduced-order basis (ROB) tailored to the given geometry. The full space-time solution can then be reconstructed by classical Galerkin projection of the predicted basis and the temporal integration described above.

In this framework, a computational mesh $\boldsymbol{\Omega}$ is represented as a graph $G = (V, E)$ whose nodes correspond to finite element vertices and whose edges represent element connectivity. To embed physical information into the learning process, each node is endowed with features such as its spatial coordinates, boundary-condition type, and maximum experienced load, along with the flattened 2×2 diagonal blocks of the full-order mass and stiffness matrices [5]. Similarly, each edge $(i, j) \in E$ includes features formed from the corresponding "cross" mass and stiffness submatrices describing interactions between neighboring nodes. This physics-aware encoding allows the network to learn directly from the structural operators of Eq. (1), enabling an improved capacity to generalize across heterogeneous, unstructured, or topologically distinct meshes. Fig. 1 presents a visual illustration of the analogy between a finite element mesh and a graph representation.

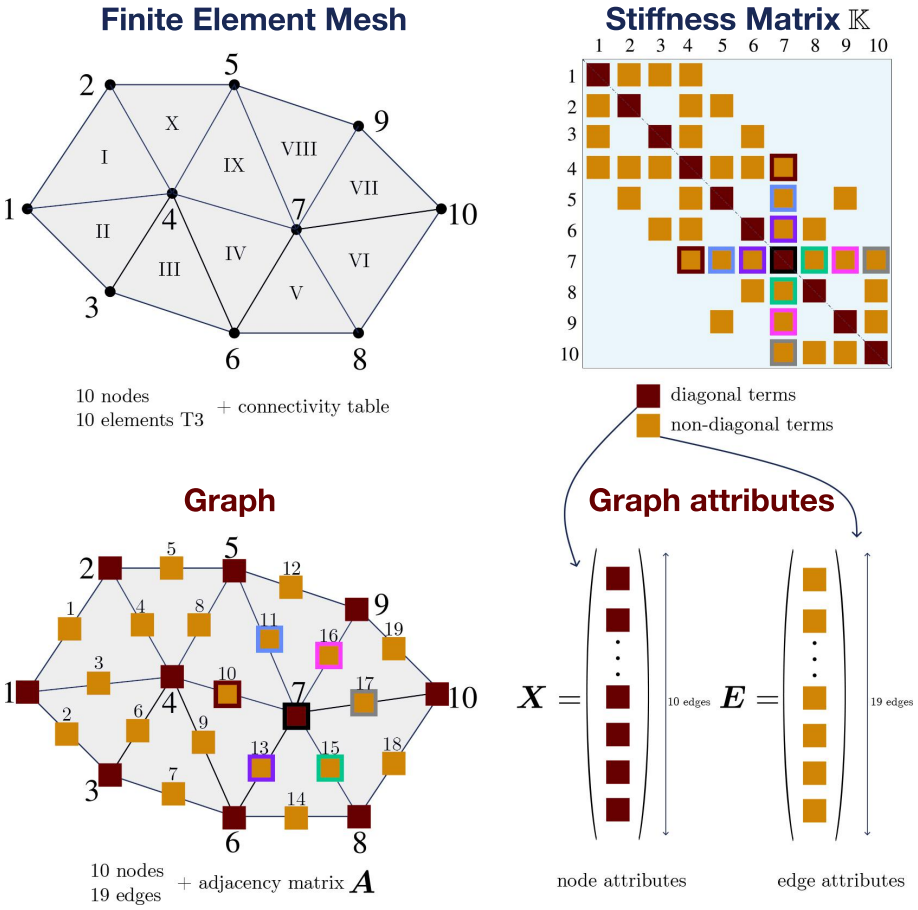

FIGURE 1 – Illustration from [5] of the analogy between a finite element mesh with one of its associated operators (e.g., the stiffness matrix $\mathbb{K}$) and a graph with its node and edge attributes ($\boldsymbol{X}$ and $\boldsymbol{E}$, resp.).

The GNN architecture recently developed in [5] consists of a single message-passing layer coupled with a Gated Recurrent Unit (GRU) that iteratively propagates information across the graph. This recurrent mechanism enables the model to capture both local and long-range geometric relationships while keeping the number of trainable parameters nearly independent of the number of iterations [5]. After

L message-passing steps, the decoder produces $2M$ output values per node (where M is the number of modes to be predicted), which are subsequently reshaped into a predicted reduced basis $\mathbf{\Phi} \in \mathbb{R}^{N\times M}$ consistent with Eq. (2).

Training is performed on a dataset of full-order simulations, where each mesh $\mathbf{\Omega}_{\text{ref}}$ is paired with a reference reduced basis $\mathbf{\Phi}_{\text{ref}}$ obtained by means of a reduced-order basis (ROB) compression of the corresponding FEM solution snapshots (e.g., via POD or PGD). The loss function combines a weighted mean-squared-error (MSE) term with a Grassmannian subspace-alignment penalty and is given by [5]

$$\mathcal{L} = \sum_{m=1}^{M} \eta_m \,\mathrm{MSE}(\phi_m, \phi_{\text{ref},m}) + \kappa d_{\text{G}}(\mathbf{\Phi}, \mathbf{\Phi}_{\text{ref}}), \tag{4}$$

where $\eta_m, \kappa \in \mathbb{R}$ are scalar weights and the Grassmann-inspired distance is approximated by a Frobenius norm as

$$d_{\text{G}}(\mathbf{\Phi}, \mathbf{\Phi}_{\text{ref}}) \;\approx\; \left\| \mathbf{\Phi}\mathbf{\Phi}^{\top} - \mathbf{\Phi}_{\text{ref}}\mathbf{\Phi}_{\text{ref}}^{\top} \right\|_F . \tag{5}$$

This formulation forces some alignment between the vector subspaces spanned by the predicted and reference modes, which is useful for accurate Galerkin projection of Eq. (1).

Once the GNN predicts a basis $\widetilde{\mathbf{\Phi}}$ for a new, previously unseen geometry $\widetilde{\mathbf{\Omega}}$, the reduced dynamics follow from the same projection procedure described in Sec. 2.1, yielding the M-dimensional system

$$\widetilde{\mathbf{\Phi}}^{\top}\mathbb{M}\widetilde{\mathbf{\Phi}}\ddot{\widetilde{\mathbf{a}}}(t) + \widetilde{\mathbf{\Phi}}^{\top}\mathbb{K}\widetilde{\mathbf{\Phi}}\widetilde{\mathbf{a}}(t) = \widetilde{\mathbf{\Phi}}^{\top}\mathbf{f}(t), \tag{6}$$

whose solution provides the temporal coefficients $\mathbf{a}(t)$ and the reconstructed displacement field $\mathbf{U}(t) \approx \widetilde{\mathbf{\Phi}}\widetilde{\mathbf{a}}(t)$.

The overall hybrid GNN–ROM methodology enables reduced-order modeling on geometries exhibiting large variations in discretization, topology, or scale, while retaining (in principle) the physical interpretability and efficiency of classical Galerkin projection. For complete technical details, the reader is referred to [5] and the references therein.

3.1 Optional enrichment

In order to further improve and control accuracy after the initial prediction of ROB modes via GNN, an optional enrichment stage based on a computationally efficient "Boosted PGD" procedure (introduced in [5]) augments the predicted basis with additional modes computed through a quasistatic approximation, enabling rapid correction of residual errors while preserving the overall dynamic fidelity. Details on the fast procedure, including its theoretical/mathematical justification and an analysis on the domain of validity of the quasistatic approximation for a dynamic system, can be found in [5].

4 An example application to structural analysis (elastodynamics)

This section considers an example geometric configuration and its corresponding structural behavior governed by a classical second-order elastodynamics PDE formulation (see [5] for specifics) : an aircraft seat bracket subjected to a crash force as illustrated in Fig. 2. A database of 440 geometries [5], containing *one to two interior circular holes* of different diameters and positions, is considered for training (351 of them), validation (44 of them), and testing (45 of them). The average degrees of freedom is $N = 10\,210$ (STD : $N = 1\,524$). Temporal integration for both the full order model (FOM) reference solutions (via FEM) as well as the space-time reconstruction of the GNN-generated reduced bases is conducted over $N_t = 400$ timesteps (the ROM considered here is POD, whose ground truth modes are constructed via randomized singular value decompositions). All GNN hyperparameters and their values (determined by a systematic design of experiments) are provided in [5] along with further training details.

For a given input test geometry $\widetilde{\mathbf{\Omega}}$, the online phase of the solver [5] consists of : a GNN inference of a POD basis $\widetilde{\mathbf{\Phi}}$ containing $M = 3$ modes (a given hyperparameter); its subsequent full space-time reconstruction $\widetilde{\mathbf{U}}(t) = \widetilde{\mathbf{\Phi}}\,\mathbf{a}(t)$ (obtained via Galerkin projection to determine the corresponding temporal coefficients $\mathbf{a}(t)$); and the optional greedy enrichment via additional Boosted PGD corrective modes.

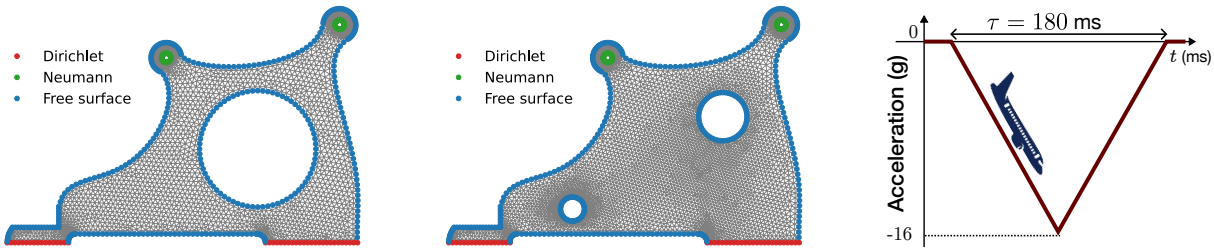

FIGURE 2 – Examples of the training geometries (from [5]) consisting of one to two interior circular holes (left and middle, respectively) of different diameters/positions that are employed to treat an elasto-dynamics problem with a force corresponding to a sudden deceleration/crash (right).

4.1 Results on geometries with topological types unseen during training

In order to evaluate the generalization capability of the proposed approach when applied to new, very different topologies that are not part of the training, validation, or test subsets (which are all comprised of *one to two interior circular holes*), an additional 30 geometries with new topological features and parameter variations can be considered, containing either *three interior circular holes* or *two interior holes of square shape*. The goal is to simulate real-world industrial conditions, where engineers and designers may require greater flexibility in geometry modifications while maintaining computational efficiency. The example mode predictions in Fig. 3 demonstrate that, despite the challenging configuration characterized by highly-heterogeneous meshes and topologies falling completely outside of the training set, the GNN mode inferences are very reasonable, even when changing the number of holes or their shapes.

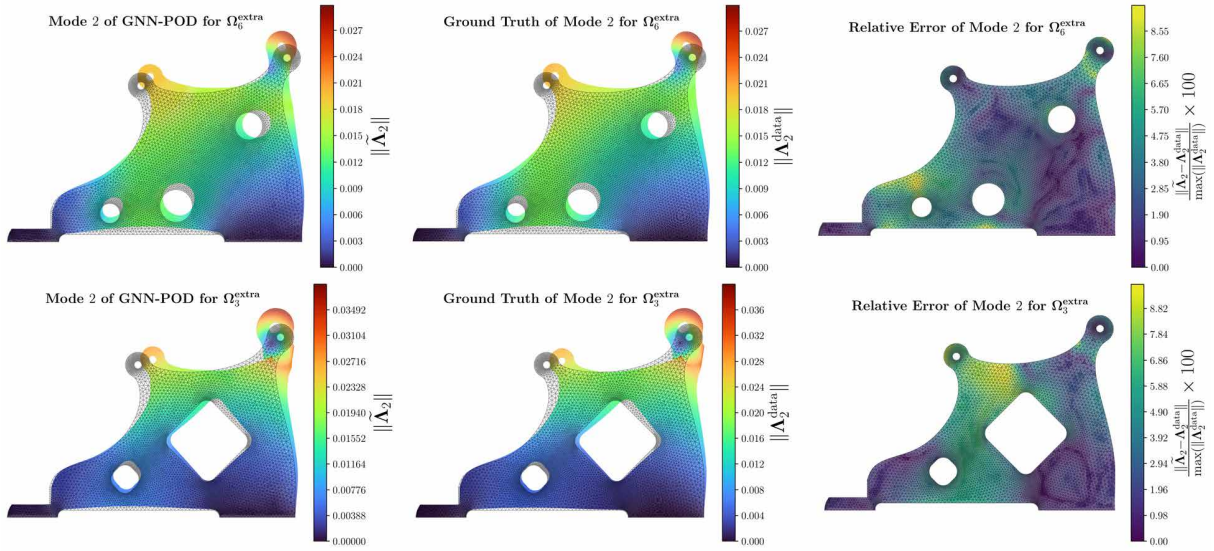

FIGURE 3 – Examples from [5] of modes predicted by GNN for new unseen geometries with fundamentally different topologies (where training is conducted on *one to two circular holes* only).

It is important to note that such changes in input topologies (including their corresponding changes in discretization sizes or degrees of freedom) represent a major challenge for this type of application, where traditional techniques require geometric parameterization or learning-based mesh morphing [9] (which often don't exist for such changes in topology, e.g., from one or two interior holes [used for training] to three interior holes). For a quantitative analysis, Table 1 presents the average $\mathcal{L}^2$ (RMSE) errors of the overall space-time solutions reconstructed from the GNN-predicted basis, as well as total computation times, both with and without enrichment. Indeed, after Galerkin projection and the subsequent resolution of the corresponding ODE system, the overall final errors are reasonably low (remaining below 6% without enrichment). Notably, for the iterative design exploration tasks motivating this work, such a level of error is already potentially highly beneficial, enabling engineers to estimate the dynamic mechanical response (and corresponding performance) of a proposed design with reasonable accuracy and with a significant speedup versus an FOM that is facilitated by FEM. The benefits of the optional enrichment

step are evident as well : with an additional (greedy) Boosted PGD mode added conventionally, errors improve by factors between 30% and 40% (overall space-time errors reduced below 0.2%) with only a relatively small sacrifice in speedup/complexity.

TABLE 1 – Average RMSE space-time errors (relative to the full-order FEM solution) and computational times across the extrapolation/generalization set of 30 geometries (three interior circular holes or two interior square-shaped holes, see Fig. 3), where training is conducted on very different topologies (one to two interior cicular holes only, see Fig. 2).

Solver	Full $\mathcal{L}^2(\boldsymbol{\Omega}, I)$ (space-time RMS) error	Time	Speedup vs FEM
GNN-POD only	5.60%	1.71 s	8.9
Enriched	0.19%	2.28 s	6.6

5 Conclusions

This contribution summarizes recent developments [5] of a new hybrid GNN–based framework [6] for reduced-order modeling of time-dependent PDEs on highly non-parametric geometries. The proposed solver integrates physics-aware finite element operators and a Grassmannian-based loss into a lightweight GRU-driven GNN architecture to predict high-quality reduced bases across heterogeneous meshes, including those with topological characteristics well outside of the training set. The predicted solutions can be enriched by the introduction of an inexpensive greedy enrichment strategy (Boosted PGD [5]) that markedly improves reconstruction accuracies—up to orders of magnitude—while maintaining substantial speedups over full-order/FOM simulations. Overall, the methodology offers a fairly robust and efficient alternative to both classical reduced-order models and conventional finite element analyses, capable of treating a wide class of geometries carrying variable discretizations and topological types while maintaining the physical consistency of traditional techniques. Such a methodology is well adapted to facilitate the construction of digital twins (e.g., for rapid pre-analysis in interative structural design or optimization/inverse problems).

Références

[1] P. Benner, S. Gugercin, K. Willcox, *A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems*, SIAM Review, pp. 483–531, 2015.

[2] F. Chinesta, P. Ladeveze, E. Cueto, *A Short Review on Model Order Reduction Based on Proper Generalized Decomposition*, Archives of Computational Methods in Engineering, vol. 18, no. 4, pp. 395–404, 2011.

[3] L. Sirovich, *Turbulence and the dynamics of coherent structures. I. Coherent structures*, Quarterly of Applied Mathematics, vol. 45, no. 3, pp. 561–571, 1987.

[4] Y. Maday and E. M. Rønquist, *A reduced-basis element method*, Journal of Scientific Computing, vol. 17, pp. 447–459, 2002.

[5] V. Matray, D. Néron, F. Feyel, F. Amlani, *Geometry-agnostic model reduction with GNN-generated reduced POD bases and boosted PGD enrichment for (non)linear structural elastodynamics*, Computer Methods in Applied Mechanics and Engineering, vol. 448, p. 118357, 2026.

[6] V. Matray, F. Amlani, F. Feyel, D. Néron, *A hybrid numerical methodology coupling reduced order modeling and Graph Neural Networks for non-parametric geometries : Applications to structural dynamics problems*, Computer Methods in Applied Mechanics and Engineering, vol. 430, p. 117243, 2024.

[7] S. Deshpande, S. P. A. Bordas, J. Lengiewicz, *Magnet : A graph U-Net architecture for mesh-based simulations*, Engineering Applications of Artificial Intelligence, vol. 133, p. 108055, 2024.

[8] V. Matray and F. Amlani, *SEAT2D-GNN : A Seat-Inspired Elastodynamics Database with Geometric and Topological Variations on 2D Meshes*, Recherche Data Gouv (doi :10.57745/3OADUI), V1, 2025.

[9] D. Ye, V. Krzhizhanovskaya, A. G. Hoekstra, *Data-driven reduced-order modelling for blood flow simulations with geometry-informed snapshots*, Journal of Computational Physics, vol. 497, p. 112639, 2024.