

Numerical solution of mixed-dimensional PDEs using a neural preconditioner

Nunzio Dimola , Nicola Rares Franco , Paolo Zunino *

MOX, Department of Mathematics, Politecnico di Milano, P.zza Leonardo da Vinci 32, Milan, 20133, Italy

ARTICLE INFO

Keywords:

Mixed-dimensional PDEs
Finite element approximation
Scientific machine learning
Preconditioning

ABSTRACT

Mixed-dimensional partial differential equations (PDEs) are characterized by coupled operators defined on domains of varying dimensions and pose significant computational challenges due to their inherent ill-conditioning. Moreover, the computational workload increases considerably when attempting to accurately capture the behavior of the system under significant variations or uncertainties in the low-dimensional structures such as fractures, fibers, or vascular networks, due to the inevitable necessity of running multiple simulations. In this work, we present a novel preconditioning strategy that leverages neural networks and unsupervised operator learning to design an efficient preconditioner specifically tailored to a class of 3D-1D mixed-dimensional PDEs. The proposed approach is capable of generalizing to varying shapes of the 1D manifold without retraining, making it robust to changes in the 1D graph topology. Moreover, thanks to convolutional neural networks, the neural preconditioner can adapt over a range of increasing mesh resolutions of the discrete problem, enabling us to train it on low resolution problems and deploy it on higher resolutions. Numerical experiments validate the effectiveness of the preconditioner in accelerating convergence in iterative solvers, demonstrating its appeal and limitations over traditional methods. This study lays the groundwork for applying neural network-based preconditioning techniques to a broader range of coupled multi-physics systems.

1. Introduction

Research on mixed-dimensional partial differential equations (PDEs) is advancing to address challenges in complex multiscale applications in fields such as geophysics, biomechanics, and neuroscience. These disciplines often involve thin structures with high aspect ratio, including structural mechanics and geomechanics [1–3], microcirculation and perfusion modeling [4,5], axons in neural networks [6], and environmental flows [7]. Modeling these interactions with traditional 3D meshes would require massive computational resources, making direct simulations impractical. Mixed-dimensional PDEs, which enable lower-dimensional representations to be embedded within higher-dimensional domains, present a promising alternative. However, the computational effort associated with these models rises significantly when trying to capture the behavior of the system for varying configurations or interactions within the low-dimensional structures, the latter representing, e.g., fracture networks in geological formations, fiber distributions in structural materials, or vascular and neural networks in biological systems. Similarly, multiple simulations are needed when the structure of the low-dimensional problem is uncertain or unknown. In all such cases, despite the simplifications provided by mixed-dimensional PDEs, the overall computational demand remains high due to the cumulative cost of repeated individual simulations.

* Corresponding author.

E-mail addresses: nunzio.dimola@polimi.it (N. Dimola), nicolarares.franco@polimi.it (N.R. Franco), paolo.zunino@polimi.it (P. Zunino).

This growing need underscores the importance of developing efficient solvers and preconditioners capable of handling the repeated solutions essential to representing the statistical variability of intricate networks in real-world applications.

Today, the discretization of mixed-dimensional PDEs is established, providing computationally efficient schemes suitable for applications with intricate networks [8,9,19], with rigorous error estimates ensuring accuracy. The well-established literature on block and operator preconditioning techniques, see, e.g., [10], has been instrumental in the development of computational solvers for such problems. Kuchta et al. [11,12] have paved the way for the efficient resolution of discrete mixed-dimensional problems. Along this direction, we also mention the work of [13] on algebraic multigrid methods for metric-perturbed coupled problems, which has significantly advanced the understanding of block preconditioning techniques in mixed-dimensional PDEs. Parallel to this, other researchers have also explored approximate algorithms that can enhance scalability for large systems: see, e.g. [14].

The goal of this study is to develop a novel preconditioning strategy that addresses the twofold challenges inherent in mixed-dimensional PDEs: the ill-conditioned behavior of these problems and the need for adaptability across diverse geometric configurations of the low-dimensional subproblems. We aim to create a preconditioner that not only mitigates the ill-conditioning of the system, but also maintains this improved conditioning across a variety of configurations for 1D or other low-dimensional structures. To accomplish this, the preconditioner is designed as a nonlinear operator rather than a conventional matrix, enabling it to process the entire solution manifold effectively. Our method takes advantage of the latest advances in the approximation capabilities of neural networks and learning approaches related to machine learning [15–17]. These advancements have sparked active research in the creation of novel solvers and preconditioners tailored for different types of problem [18,20–23], this being a non-exhaustive list of examples supporting the potential of neural network-based preconditioners as a viable alternative for solving large, sparse linear systems across various scientific and engineering domains.

Building upon these studies, we develop a nonlinear preconditioner designed to efficiently manage the numerical challenges associated with the complex structure of mixed-dimensional equations. By incorporating a shape descriptor into the learning framework, the preconditioner can generalize across different configurations of the 1D graph or low-dimensional structures without the need for retraining. In order to demonstrate the efficacy of the proposed approach, we focus our attention on a mixed-dimensional 3D-1D model problem, which serves as a representative template for a broader class of coupled systems with similar characteristics. The results show that the learned preconditioner significantly accelerates the convergence of the iterative Generalized Minimum Residual (GMRES) solver, even as the complexity of the problem increases due to changes in the 1D graph's topology or the size of the 3D domain. Unlike traditional preconditioning techniques, which typically require problem-specific tuning or reconfiguration when the underlying system changes, our method maintains consistent performance across diverse problem settings and mesh resolutions. Moreover, the proposed preconditioning approach can be easily generalized, making it a versatile and robust tool for a wide array of mixed-dimensional problems beyond the specific examples addressed in this study.

We finally observe that recent advances in operator learning (e.g., DeepONets [17] and Neural Operators [24]) provide fast surrogates for the parameter-to-solution map of PDEs, making them attractive for approximate evaluations in tasks such as design exploration or uncertainty quantification. However, such surrogates cannot ensure systematic accuracy with mesh refinement, may lose reliability outside the training distribution, and do not deliver the exact discrete PDE solution. By contrast, the approach proposed here is not a surrogate but a learned component embedded in an iterative solver: it preserves the full accuracy of the discretization while accelerating convergence, thereby combining the robustness of classical solvers with the efficiency of machine learning and allowing seamless integration into existing numerical pipelines.

This manuscript is organized as follows. First, we introduce the mathematical formulation of the 3D-1D mixed-dimensional PDE and highlight the challenges associated with solving such systems. Then, we present the proposed preconditioner learning framework, discussing both the theoretical foundations and the practical implementation of the nonlinear preconditioner. Next, we apply the preconditioner to the 3D-1D problem, illustrating its construction and providing a detailed analysis of its efficiency and scalability. Finally, we validate the approach through extensive numerical experiments, comparing its performance against conventional preconditioning methods, and conclude with a discussion of the implications and potential future directions for extending this work.

2. A mixed-dimensional problem

We consider here a mixed-dimensional elliptic problem as a simple template for a wider class of coupled problems with similar characteristics. It will be our reference for comparisons, analyses, and considerations on preconditioner performances. Originally, we define the problem in a three-dimensional (3D) slender domain embedded in an external one, where we set partial differential equations coupled by suitable interface conditions. The mixed-dimensional model is obtained by a *dimensional reduction* strategy that represents the slender interior domain as a one-dimensional (1D) manifold, namely a metric graph. More precisely, we address a coupled problem defined by an exterior domain $\Omega \subset \mathbb{R}^3$ and an interior domain $\Sigma \subset \Omega$. We assume Σ to be a generalized cylinder with a centerline Λ , the latter being a 1D domain with arc length parameter $s \in (0, S)$. We assume that the cylinder has a constant radius $e > 0$, so that if $\lambda : (0, S) \rightarrow \Lambda$ parameterizes the centerline, then $\Sigma = \{\lambda(s) + e : s \in (0, S), \|e\| \leq e, e \perp \tau(s)\}$ with $\tau(s)$ the tangent vector at Λ . We assume that the transverse diameter of Σ , which is equal to $2e$, is small compared to the diameter of Ω . The dimensionality reduction process consists of replacing Σ with Λ while keeping some information about its original 3D structure and its interaction with the external domain Ω . Specifically, let $\Gamma := \partial\Sigma$ be the interface of the coupled problem in the full 3D representation and let \mathcal{T}_Λ be the restriction operator from Ω to Λ . In [25] it is defined as the composition of the trace operator on Γ combined with a projection operator from Γ to Λ based on cross-sectional averages. In particular, the averaging procedure plays a regularizing role: the ill-posed trace of a 3D function onto a 1D manifold (which is not defined in the usual Sobolev spaces) is replaced by the well-defined




Fig. 1. Solution of a 3D-1D coupled problem for a fixed graph geometry Λ . **Left:** plot of 1D solution u_Λ . **Center:** slices of 3D solution u_Ω . **Right:** slices plot of the distance function $d(\Lambda)$ associated to the graph.

operator \mathcal{T}_Λ . Then, the continuous 3D-1D coupled problem can be formulated as follows,

$$\begin{cases} \nabla \cdot (-k_\Omega \nabla u_\Omega) + \sigma_\Omega u_\Omega + 2\pi\epsilon(\mathcal{T}_\Lambda u_\Omega - u_\Lambda)\delta_\Lambda = f_\Omega, & \text{on } \Omega, \\ d_s(-\pi\epsilon^2 k_\Lambda d_s u_\Lambda) + 2\pi\epsilon(u_\Lambda - \mathcal{T}_\Lambda u_\Omega) = 0, & \text{on } \Lambda, \\ -\nabla u_\Omega \cdot \mathbf{n} = 0, & \text{on } \partial\Omega, \\ -d_s u_\Lambda \cdot \mathbf{n} = 0, & \text{on } \partial\Lambda \setminus \partial\Lambda_D, \\ u_\Lambda = 1, & \text{on } \partial\Lambda_D. \end{cases} \tag{1}$$

The functions u_Ω, u_Λ are coupled unknowns in the 3D and 1D domains, respectively, and $k_\Omega, \sigma_\Omega, k_\Lambda$ represent the coefficients of the elliptic operators. Here, $\partial\Lambda_D \subset \partial\Lambda$ is a given subset of the 1D domain holding a Dirichlet boundary condition. We note that the first equation of (1) is only formally defined in its strong form, as the symbol δ_Λ denotes the Dirac distribution supported on the one-dimensional manifold $\Lambda \subset \Omega$. For the weak formulation of (1), given a test function $v : \Omega \rightarrow \mathbb{R}$, $v\delta_\Lambda$ represents the distribution that restricts the integration on Ω to the manifold Λ . The coupling term $2\pi\epsilon(\mathcal{T}_\Lambda u_\Omega - u_\Lambda)\delta_\Lambda$ is obtained by averaging on Γ the Robin-type coupling condition $-k_\Omega \nabla u_\Omega \cdot \mathbf{n} = u_\Omega - u_\Lambda$, similar to the one adopted in [25]. In fact, the coefficients $2\pi\epsilon$ and $\pi\epsilon^2$ represent the cross-sectional measures of Γ and Σ , respectively. An analogous formulation holds when the low-dimensional structure is given by the union of multiple substructures $\Sigma = \cup_i \Sigma_i$ and $\Lambda = \cup_i \Lambda_i$, each being a generalized cylinder. In this case, Λ becomes a 1D graph embedded within a 3D domain. Fig. 1 shows an example of a solution to this mixed-dimensional problem (1D solution on the left, 3D solution in the center).

2.1. Parametrization and discretization of the problem

As we mentioned, we are interested in solving multiple instances of the above problem for different configurations of the low-dimensional structure, including, for instance, changes in the topology, arc-length or density of the 1D graph Λ . To formalize this fact, it is useful to define a suitable *parameter space*, \mathcal{P} , that collects all configurations of interest. To this end, we first note that each 1D graph $\Lambda \subset \Omega$ is uniquely identified by its *3D distance function*, $d(\Lambda) : \Omega \rightarrow \mathbb{R}$, which maps each point $x \in \Omega$ to its distance from the closest point in Λ : see the right panel of Fig. 1 for a visual representation. This allows us to define the parameter space as a subset of the space of continuous functions on Ω , that is, $\mathcal{P} := \{d(\Lambda)\}_\Lambda \subset C(\Omega)$. This representation has several advantages. First, it can be easily transferred to the discrete setting by introducing a suitable discretization of the exterior domain Ω using, for example, Finite Elements (FE). Furthermore, it automatically enriches the parameter space by equipping it with a metrizable topology (the one induced by the supremum norm): this allows us to discuss, for instance, about the continuity of the 3D solution with respect to the 1D graph. Finally, it will be useful when discussing our neural preconditioning approach in Section 4.

We now come to the discretization of (1). To this end, we first address its variational formulation. Let $V := H^1(\Omega) \times H^1_{\partial\Lambda_D}(\Lambda)$ be the mixed-dimensional space where PDE solutions will be sought. Here, $H^1_{\partial\Lambda_D}(\Lambda) = \{w \in H^1(\Lambda) : w|_{\partial\Lambda_D} \equiv 1\}$ accounts for Dirichlet boundary conditions, so that —up to translations— V is isomorphic to a Hilbert space. In the parametric setting, re-writing Eq. (1) in weak form corresponds to saying that, for every $\mu \in \mathcal{P}$ we aim to find $u^\mu \in V$ such that

$$a^\mu(u^\mu, v) = F^\mu(v), \quad \forall v \in V, \tag{2}$$

with $a^\mu : V \times V \mapsto \mathbb{R}$ and $F^\mu : V \mapsto \mathbb{R}$ parameter-dependent operators. Specifically,

$$a^\mu(u, v) = a_\Omega(u_\Omega, v_\Omega) + a_\Lambda^\mu(u_\Lambda, u_\Lambda) + b_\Lambda^\mu(\mathcal{T}_\Lambda u_\Omega - u_\Lambda, \mathcal{T}_\Lambda v_\Omega - v_\Lambda),$$

where having set $(u, v)_\Omega := \int_\Omega u(x)v(x)dx$ we obtain,

$$a_\Omega(u, v) = (k_\Omega \nabla u, \nabla v)_\Omega + (\sigma_\Omega u, v)_\Omega, \tag{3}$$

$$a_\Lambda^\mu(u, v) = (\pi\epsilon^2 k_\Lambda d_s u, d_s v)_\Lambda, \quad b_\Lambda^\mu(u, v) = 2\pi\epsilon(u, v)_\Lambda. \tag{4}$$

While the bilinear form (3) follows a standard derivation, the forms (4) appear after the topological model reduction procedure presented in [25] (taking averages over cross-sections Σ and Γ , thus introducing the factors $\pi\epsilon^2$ and $2\pi\epsilon$, respectively). We note that the adopted notation is slightly redundant because the parametric dependence of a^μ is only due to the fact that a_Λ^μ and b_Λ^μ have support in Λ .

In order to discretize (2), we consider its Galerkin projection onto a (broken) FE space $V_h = V_h^\Omega \times V_h^\Lambda$ of dimension $N_h = N_h^\Omega + N_h^\Lambda$, being $N_h^\Omega = \dim(V_h^\Omega)$ and $N_h^\Lambda = \dim(V_h^\Lambda)$, which is suitably chosen depending on the characteristics of the problem at hand. Assuming for simplicity a fully conformal approximation, given $V_h \subset V$ with $V_h^\Omega \subset H^1(\Omega)$, $V_h^\Lambda \subset H^1_{\partial\Lambda_D}(\Lambda)$, we aim to find $u_h^\mu \in V_h$ such that

$$a^\mu(u_h^\mu, v_h) = F^\mu(v_h) \quad \forall v_h \in V_h. \tag{5}$$

From a discrete point of view, the problem (5) is equivalent to a (large) system of algebraic equations $A_h^\mu \mathbf{u}_h^\mu = \mathbf{F}_h^\mu$. Specifically, we have

$$\underbrace{\begin{bmatrix} K_{h,00} & 0 \\ 0 & K_{h,11} \end{bmatrix}}_{K_h^\mu} + \gamma \underbrace{\begin{bmatrix} M_{h,00}^\mu & M_{h,01}^\mu \\ M_{h,10}^\mu & M_{h,11}^\mu \end{bmatrix}}_{M_h^\mu} \underbrace{\begin{bmatrix} \mathbf{u}_{h,0}^\mu \\ \mathbf{u}_{h,1}^\mu \end{bmatrix}}_{\mathbf{u}_h^\mu} = \underbrace{\begin{bmatrix} \mathbf{f}_{h,0}^\mu \\ \mathbf{f}_{h,1}^\mu \end{bmatrix}}_{\mathbf{F}_h^\mu}, \tag{6}$$

where $\mathbf{u}_h^\mu \in \mathbb{R}^{N_h}$ is the vector of degrees of freedom of the FE approximation; $K_{h,00}$ and $K_{h,11}^\mu$ represent the discretized bilinear forms a_{Ω} and a_Λ^μ where the coefficients k_Ω and $\pi\epsilon^2 k_\Lambda$ have been embedded into the matrices; the block matrix M_h^μ models the coupling enforced by the operator $2\pi\epsilon(\mathcal{T}_\Lambda u_\Omega - u_\Lambda, \mathcal{T}_\Lambda v_\Omega - v_\Lambda)_\Lambda$ where we set from now on $\gamma := 2\pi\epsilon$ because the geometrical meaning of this coefficient is no longer significant as it essentially represents the relative weight between the blocks K_h^μ and M_h^μ ; $\mathbf{f}_{h,0}$ denotes the non-trivial forcing term acting on the bulk domain Ω and $\mathbf{f}_{h,1}^\mu$ takes into account the non-homogeneous Dirichlet condition on the 1D graph.

As we discretize the differential problem, we also take the opportunity to discretize the parameter space, which we do by projecting—or interpolating—all 3D distance functions onto V_h^Ω . This allows us to represent each parametric configuration $\mu \in \mathcal{P}$ through a finite-dimensional vector $\mathbf{d}_h^\mu \in \mathbb{R}^{N_h^\Omega}$, obtained by listing the nodal values of $\mu = d(\Lambda)$ at the degrees of freedom. In this way, the discrete/finite-dimensional parameter space becomes $\mathcal{P}_h \equiv \{\mathbf{d}_h^\mu\} \subset \mathbb{R}^{N_h^\Omega}$.

Remark 1. As apparent from Eqs. (3)–(4), we notice that problem (2) is symmetric positive definite. However, we will not exploit this property when defining the computational solver. This choice is motivated by two main reasons. The first one is that (2) is just a particular case of a general family of operators in which the bilinear form $a_\Omega(\cdot, \cdot)$ may also describe transport phenomena that violate symmetry. The second reason is that, as it will be clearer later on, learned techniques, such as our neural preconditioner, can hardly preserve symmetry, unless this constraint is explicitly built into the network architecture—a strategy that currently lacks a reliable and efficient implementation.

Remark 2. In what follows, we shall make the assumption that the 3D domain has been discretized using a uniform mesh, consisting either of tetrahedrons or bricks. This requirement will be crucial in Section 4 when implementing our neural network-based preconditioner as it will allow us to exploit the relationship between tensor-like structures and convolutional neural networks. Nevertheless, we emphasize that in the context of mixed-dimensional problems, where the primary challenge is representing a complex low-dimensional structure embedded in 3D, this restriction is acceptable. In fact, this requirement does not constrain the low-dimensional domain, which can still exhibit arbitrary geometric complexity.

2.2. A preconditioning strategy for the coupled problem

As we mentioned, solving (6) can be computationally intensive due to poor conditioning of the linear system. Here, we address this fact by using a block-preconditioning approach, specifically tailored for the block-structured matrix arising from the discretization of the coupled 3D-1D problem.

The goal is to employ a right-preconditioner expressed as:

$$A_h^\mu Q_h^\mu \mathbf{z}_h^\mu = \mathbf{F}_h^\mu, \quad Q_h^\mu \mathbf{z}_h^\mu = \mathbf{u}_h^\mu, \tag{7}$$

where Q_h^μ is the inverse of the block upper triangular part of A_h^μ . Employing this right preconditioner to accelerate the convergence of an iterative method, see e.g. [26], involves conditioning the residual from $[\mathbf{z}_{h,0}^\mu, \mathbf{z}_{h,1}^\mu]^\top$ to $[\tilde{\mathbf{z}}_{h,0}^\mu, \tilde{\mathbf{z}}_{h,1}^\mu]^\top$ by solving the following linear system through a two-step back-substitution procedure,

$$\begin{bmatrix} (K_{h,00} + \gamma M_{h,00}^\mu) & \gamma M_{h,01}^\mu \\ 0 & (K_{h,11}^\mu + \gamma M_{h,11}^\mu) \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{z}}_{h,0}^\mu \\ \tilde{\mathbf{z}}_{h,1}^\mu \end{bmatrix} = \begin{bmatrix} \mathbf{z}_{h,0}^\mu \\ \mathbf{z}_{h,1}^\mu \end{bmatrix}. \tag{8}$$

In general, the first step, which is to solve for $\tilde{\mathbf{z}}_{h,1}^\mu$, does not pose significant challenges. This is because the dimension of the discrete 1D problem is typically smaller than the one of the 3D problem, $N_h^\Lambda \ll N_h^\Omega$, making it more manageable to address using traditional

techniques, such as direct solvers. In contrast, the second step, which is to solve for $\bar{\mathbf{z}}_{h,0}^\mu$, can be computationally demanding and can severely hinder the applicability of the approach when considering many-query scenarios, where the linear system is to be solved for multiple instances of the model parameters. This step, in fact, is the one where the main challenges inherent to mixed-dimensional problems become apparent. Indeed, as noted in [13], the latter involves elliptic operators altered by a low-order semidefinite component, referred to as the *metric term*. This term, whose impact is modulated by a scalable intensity parameter ϵ , can disrupt the structure and conditioning of the elliptic component, thereby complicating the numerical approximation of the coupled system. For these reasons, our objective is to develop a nonlinear preconditioner, \mathcal{P} , for the mixed-dimensional system defined on Ω ,

$$(K_{h,00} + \gamma M_{h,00}^\mu) \bar{\mathbf{z}}_{0,h}^\mu = \mathbf{z}_{h,0}^\mu - \gamma M_{01}^\mu \bar{\mathbf{z}}_{1,h}^\mu. \tag{9}$$

This nonlinear preconditioner will take two arguments: one that is conformal to the 3D solution (as any classical preconditioner would do), and one that refers to the problem parameters, allowing the preconditioner to adjust depending on the scenario of interest. We shall write $\mathcal{P} = \mathcal{P}(\mathbf{v}, \mathbf{d})$ in the discrete setting and, occasionally, $\mathcal{P} = \mathcal{P}(v, \mu)$ when considering the more abstract continuous formulation. Using this notation, and having set

$$C_h^\mu := (K_{h,00} + \gamma M_{h,00}^\mu), \quad \mathbf{v}_h^\mu := \bar{\mathbf{z}}_{0,h}^\mu, \quad \mathbf{b}_h^\mu := \mathbf{z}_{h,0}^\mu - \gamma M_{01}^\mu \bar{\mathbf{z}}_{1,h}^\mu, \tag{10}$$

the right-preconditioned family of problems reads $C_h^\mu \mathcal{P}(\mathbf{v}_h^\mu, \mathbf{d}_h^\mu) = \mathbf{b}_h^\mu$. Considering the specific structure of \mathbf{F}_h^μ in (6), we focus on making the preconditioner \mathcal{P} efficient for a limited subset of acceptable right-hand sides within $\mathbb{R}^{N_h^\Omega}$. In particular, we are interested in solving the system for the forcing terms that arise from the solution of the 1D problem, namely all $B = \{\mathbf{b}_h^\mu\}_{\mu \in \mathcal{P}}$ such that

$$B := \left\{ \mathbf{f}_{h,0} - \gamma M_{01,h}^\mu \mathbf{x} : \mu \in \mathcal{P} \text{ and } \mathbf{x} \in \mathbb{R}^{N_h^\Omega} \text{ with } (K_{11}^\mu + \gamma M_{11,h}^\mu) \mathbf{x} = \mathbf{f}_{h,1}^\mu \right\}. \tag{11}$$

As we anticipated, to guarantee generalization across various configurations of the 1D problem, it is essential to develop a suitable nonlinear preconditioner, \mathcal{P} . This involves carefully choosing an appropriate hypothesis space and designing a representative training set, as detailed in the following section.

3. A learning approach for preconditioning parametrized systems

We shall now discuss the problem of learning preconditioners for parametrized systems. Although this task is inherently related to finite-dimensional spaces, we find it convenient to explore and present the idea within an infinite-dimensional context. This shift to a function space setting, as presented in this section, provides a deeper insight into the mathematical structure of the problem. For these reasons, we start this section by fixing some notation. Given two Banach spaces $(V_1, \|\cdot\|_{V_1})$ and $(V_2, \|\cdot\|_{V_2})$ we denote $\mathcal{L}(V_1, V_2)$ the space of bounded linear operators from V_1 to V_2 , equipped with the operator norm

$$\|A\|_{\mathcal{L}(V_1, V_2)} := \sup_{v \in V_1 \setminus \{0\}} \frac{\|Av\|_{V_2}}{\|v\|_{V_1}}.$$

It is also useful to introduce the *General Linear Group*

$$\mathcal{GL}(V_1, V_2) := \{A \in \mathcal{L}(V_1, V_2) : \exists B \in \mathcal{L}(V_2, V_1), BA = \text{Id}_{V_1} \text{ and } AB = \text{Id}_{V_2}\},$$

which consists of all linear operators from $V_1 \rightarrow V_2$ that are both invertible and continuous. Here, Id_{V_i} denotes the identity operator of V_i on itself. Given a Banach space $(V, \|\cdot\|_V)$, we also denote by V' its topological dual, that is, $V' := \mathcal{L}(V, \mathbb{R})$. Our purpose is to address the problem of operator preconditioning for parametrized linear problems in arbitrary (complete) normed spaces, i.e. to find a continuous preconditioner operator $\mathcal{P} = \mathcal{P}(v, \mu)$ for a parametrized problem of the form

$$\text{find } u^\mu \in V : A^\mu u^\mu = b^\mu, \quad \forall \mu \in \mathcal{P}, \tag{12}$$

given a Banach state space $(V, \|\cdot\|_V)$ and a compact metric space \mathcal{P} serving as parameter space. Here, $A^\mu \in \mathcal{GL}(V, V')$ and $b^\mu \in V'$ are parameter-dependent operators and problem data. In general, we think of the latter as a linear system arising from the discretization of a given PDE, such as (5) or (6). Everything, in fact, can be traced back to our discussion in Section 2: we refer to Remark 3 for additional insight on the matter. Furthermore, we anticipate that the upcoming Section 4 will also contain further clarifications on how these concepts translate when put into practice.

We organize the remainder of this Section into distinct parts. Initially, we explore how operator preconditioning is applied to (12), emphasizing the imperative for a nonlinear preconditioner (Section 3.1). Subsequently, we outline the structure of the learning algorithm, presenting a suitable unsupervised training strategy (Section 3.2).

Remark 3. The advantage of working within an abstract environment is that the idea can be easily transferred from the continuous formulation to the discrete one and vice versa. For instance, following our notation in Section 2.2, if we resort to the discrete setting, then $A^\mu \cong C_h^\mu$ is a matrix, while $u^\mu \cong \mathbf{v}_h^\mu$ and $b^\mu \cong \mathbf{b}_h^\mu$ are finite-dimensional vectors; cf. Eq. (9). Conversely, if we lift the idea to the continuous level, then A^μ is the linear operator acting as $u \mapsto a_\Omega(u, \cdot) + b_\Lambda^\mu(\mathcal{T}_\Lambda u, \cdot)$, while u^μ and b^μ are functions defined over Ω : cf. Section 2.1.

3.1. The need for a nonlinear and matrix-free preconditioner

Classical approaches to operator preconditioning focus on finding a linear preconditioner for a given operator of interest $A \in \mathcal{L}(V', V)$. In principle, one could apply this idea to parametric problems by addressing each parametric scenario separately. Indeed, for a generic, but *fixed*, parameter $\bar{\mu} \in \mathcal{P}$ one could consider the corresponding operator $A^{\bar{\mu}} \equiv A$, with $A : V \rightarrow V'$, and then look for a suitable (right) preconditioner that is, any linear operator $P \in \mathcal{L}(V', V)$ for which $\kappa(AP) \ll \kappa(A)$, being κ the operator *condition number* $\kappa(B) := \|B\|_{\mathcal{L}(V_1, V_2)} \|B^{-1}\|_{\mathcal{L}(V_1, V_2)}$ for $B \in \mathcal{GL}(V_1, V_2)$. Then, once a suitable preconditioner is available, the original problem $Au^{\bar{\mu}} = b^{\bar{\mu}}$ can be replaced with a new one, $APz^{\bar{\mu}} = b^{\bar{\mu}}$, which is easier to solve, so that $u^{\bar{\mu}} = Pz^{\bar{\mu}}$. In general, we note that up to isometries, the "optimal" preconditioner is $P = A^{-1}$, regardless of the right hand side $b^{\bar{\mu}} \in V'$. However, this approach becomes computationally unfeasible if the procedure is repeated multiple times for varying $\mu \in \mathcal{P}$. A first alternative could be to look for a surrogate $\tilde{\mathcal{P}} : \mathcal{P} \rightarrow \mathcal{L}(V', V)$ that, for each $\mu \in \mathcal{P}$, returns a suitable preconditioner for $A^\mu u^\mu = b^\mu$. However, we do not find this solution to be very practical as that would entail casting the learning problem within an extremely high-dimensional space. Notice, in fact, that in the discrete setting one has $\dim \mathcal{L}(V', V) = (N_h^\Omega)^2$. Furthermore, even if we restrict our attention to sparse preconditioners, this is still very challenging as the sparsity pattern would likely depend on the structure of the underlying 1D problem, i.e. on $\mu \in \mathcal{P}$. A more practical approach can be to rely on a matrix-free formulation. There, one is not directly interested in the preconditioner by itself, but rather in the way it acts when applied to an input vector. Mathematically speaking, this change of perspective corresponds to looking for a map, potentially *nonlinear*, of the form $\mathcal{P} : V' \times \mathcal{P} \rightarrow V$. This idea is further motivated by the following Lemma.

Lemma 1. *Let $\mathcal{P} \ni \mu \mapsto A^\mu \in \mathcal{L}(V, V')$ be continuous. For any $\mu \in \mathcal{P}$ and any $v \in V'$ let $x^{\mu,v}$ be the solution of $A^\mu x^{\mu,v} = v$. Then, there exists a continuous (nonlinear) operator $\mathcal{P} : V' \times \mathcal{P} \rightarrow V$ such that*

$$\mathcal{P}(v, \mu) = x^{\mu,v} \quad \forall (v, \mu) \in V' \times \mathcal{P}.$$

In particular, $A^\mu \mathcal{P}(v, \mu) = v$ for all $(v, \mu) \in V' \times \mathcal{P}$.

Proof. Let $\mathcal{A} : \mathcal{P} \rightarrow \mathcal{GL}(V, V')$ be the map $\mathcal{A}(\mu) = A^\mu$. Let $\mathcal{I} : \mathcal{GL}(V, V') \rightarrow \mathcal{GL}(V', V)$ be the inversion map, $B \mapsto B^{-1}$. It is well known that \mathcal{I} is continuous — for a rigorous proof, we refer the interested reader to the Appendix, Lemma 2. Let $\mathcal{P} : V' \times \mathcal{P} \rightarrow V$ be defined as $\mathcal{P}(v, \mu) := [(\mathcal{I} \circ \mathcal{A})(\mu)](v)$. Then, \mathcal{P} is continuous and $\mathcal{P}(v, \mu) = (A^\mu)^{-1}v = x^{\mu,v}$, as claimed. \square

As seen in Lemma 1, in order to adapt to different parametric scenarios, the preconditioner \mathcal{P} must be nonlinear in its arguments (see also Remark 4). In light of this, hereon, when speaking of a nonlinear preconditioner for a parametric problem such as (12), we intend any continuous map $\mathcal{P} : V' \times \mathcal{P} \rightarrow V$ for which $A^\mu \mathcal{P}(v, \mu) \approx v$, or equivalently, $(\text{Id} - A^\mu \mathcal{P}(\cdot, \mu))v \approx 0$. Ideally, this should hold for all $\mu \in \mathcal{P}$ and all right-hand sides of interest $v \in \mathcal{K}^\mu \subset V'$, with the latter subset potentially depending on the parameter instance μ . For example, the simplest choice could be $\mathcal{K}^\mu = \{b^\mu\}$, so that $\mathcal{P}(b^\mu, \mu) \approx u^\mu$. Setting $\mathcal{K}^\mu = V'$ is also possible but not strictly necessary: we recall, in fact, that our ultimate goal is to solve (12). We shall return to this in Section 3.2, as we address the details of our learning strategy, and in Section 4.1, where we discuss the actual implementation of the approach and the definition of \mathcal{K}^μ .

Remark 4. In principle, one could restrict oneself to continuous maps $\mathcal{P} : V' \times \mathcal{P} \rightarrow V$ that are linear in their first argument, $\mathcal{P} = \mathcal{P}(v, \mu)$, as it happens for the "ideal" preconditioner in Lemma 1. However, allowing for more general maps gives additional freedom in designing the preconditioner, potentially simplifying the learning process. For example, assume that \mathcal{P} is compact and that the map $\mathcal{F} : \mu \rightarrow b^\mu$ is both continuous and injective. Then, $\mathcal{F}^{-1} : \mathcal{F}(\mathcal{P}) \rightarrow \mathcal{P}$ exists, is continuous and can be extended to $V' \supset \mathcal{F}(\mathcal{P})$ using, e.g., Dugundji's extension of Tietze's Theorem [27]. With this setup, consider the simple scenario $\mathcal{K}^\mu = \{b^\mu\}$ and let \mathcal{P} be the map in Lemma 1. Define $\tilde{\mathcal{P}} : V' \times \mathcal{P} \rightarrow V$ as $\tilde{\mathcal{P}}(v, \mu) := \mathcal{P}(v, \mathcal{F}^{-1}(v))$. Then, $\tilde{\mathcal{P}} \neq \mathcal{P}$ is nonlinear in v , does not depend explicitly on μ (only implicitly), but still $A^\mu \tilde{\mathcal{P}}(v, \mu) = v$ for all $\mu \in \mathcal{P}$ and all $v \in \mathcal{K}^\mu \subset V'$.

3.2. Learning strategy

Following our previous discussion, our purpose here is to learn a (nonlinear) preconditioner for a parametrized problem given as in Eq. (12). In practice, this involves fixing a suitable *hypothesis space* \mathcal{H} and finding the optimal preconditioner $\mathcal{P} \in \mathcal{H}$ by minimizing a suitable functional known as *loss function*. In light of Lemma 1, we let $\mathcal{H} \subset C(V' \times \mathcal{P}, V) = \{\mathcal{P} : V' \times \mathcal{P} \rightarrow V \text{ continuous}\}$ be a subset of the space of continuous functions from $V' \times \mathcal{P}$ to V . For instance, if we resort to deep learning techniques, as we shall do in the remaining of the paper, the latter can be thought of as the space of all possible neural network models associated to a given architecture of choice. In this case, we'll refer to \mathcal{P} as to a *neural preconditioner*. Concerning the loss function, instead, different approaches become available. In principle, one could define the nonlinear preconditioner \mathcal{P} by requiring either $\mathcal{P}(v, \mu) \approx x^{\mu,v}$ or $(\text{Id} - A^\mu \mathcal{P}(\cdot, \mu))v \approx 0$. In fact, we notice that if the original problem admits a unique solution for each $\mu \in \mathcal{P}$, then every \mathcal{P} that fulfills $\mathcal{P}(b^\mu, \mu) = u^\mu$ yields $A^\mu \mathcal{P}(b^\mu, \mu) = b^\mu$ and vice versa. In this sense, the two approaches may seem equivalent. However, our hypothesis space will typically not contain the global minimizer in Lemma 1, reason for which minimizing the *error* $\|\mathcal{P}(v, \mu) - x^{\mu,v}\|$ or the *residual* $\|(\text{Id} - A^\mu \mathcal{P}(\cdot, \mu))v\|$ can provide completely different results. Furthermore, the two approaches differ significantly in their implementation: the first one is *supervised*, as it requires actually solving (12) in order to compute $x^{\mu,v}$ for multiple $\mu \in \mathcal{P}$; in contrast, the second approach is fully *unsupervised*.

This work focuses exclusively on the second approach. This is because preliminary investigations by the authors found the supervised approach to be extremely inefficient, with high costs during the training phase, significant challenges in minimizing the loss function, and unsatisfactory results. In contrast, the unsupervised approach is more practical, as it does not require sampling from

the solution manifold $S = \{u^\mu\}_{\mu \in \mathcal{P}}$. All of this considered, a first definition of the learning problem can be

$$\mathcal{P}_* = \arg \min_{\mathcal{P} \in \mathcal{H}} \int_{\mathcal{P}} \|(\text{Id} - A^\mu \mathcal{P}(\cdot, \mu))b^\mu\|_{V'} \varrho(d\mu), \tag{13}$$

for a given probability measure ϱ defined over the parameter space \mathcal{P} . The downside of this approach is that for each $\mu \in \mathcal{P}$, we focus solely on how well \mathcal{P} operates on b^μ . As anticipated in Section 3.1, a better approach can be to introduce a subset $\mathcal{K}^\mu \subset V'$ of possible right-hand-sides (depending in general on μ), and reformulate (13) as

$$\mathcal{P}_* = \arg \min_{\mathcal{P} \in \mathcal{H}} \int_{\mathcal{P}} \int_{\mathcal{K}^\mu} \|(\text{Id} - A^\mu \mathcal{P}(\cdot, \mu))v\|_{V'} \varrho_\mu(dv) \varrho(d\mu), \tag{14}$$

where ϱ_μ is a suitable probability measure supported over $\mathcal{K}^\mu \subset V'$. Clearly, (13) is a special case of (14), since it can be obtained letting $\mathcal{K}^\mu = \{b^\mu\}$ be a singleton and $\varrho_\mu = \delta_{b^\mu}$ be a Dirac delta measure. The question of defining \mathcal{K}^μ and the corresponding probability measure is generally intriguing, yet quite complex. We will defer this discussion to Section 4.1. We anticipate that \mathcal{K}^μ should contain all the components that are detrimental to the well-posedness of A^μ and, in general, to the convergence rate of the algebraic solver considered.

From a practical point of view, the ideal minimization problem in Eq. (14) can be addressed through *empirical risk minimization* [15, 16]. Specifically, let μ_1, \dots, μ_{N_p} be an independent random sample identically distributed (i.i.d.) with $\mu_j \sim \varrho$. For each $j = 1, \dots, N_p$ let $\mathcal{K}^{\mu_j} := \{v_{1,j}, \dots, v_{N_{\mathcal{K}^{\mu_j}},j}\} \subset \mathcal{K}^{\mu_j}$, be an i.i.d. random sample with $v_{i,j} \sim \varrho_{\mu_j}$. Consider the empirical measures

$$\varrho_{N_p} = \frac{1}{N_p} \sum_{j=1}^{N_p} \delta_{\mu_j}, \quad \varrho_{N_{\mathcal{K}^{\mu_j}}} = \frac{1}{N_{\mathcal{K}^{\mu_j}}} \sum_{v \in \mathcal{K}^{\mu_j}} \delta_v, \quad \varrho_{N_{\mathcal{K}^\mu}} = \sum_{j=1}^{N_p} \mathbf{1}_{\{\mu_j\}}(\mu) \varrho_{N_{\mathcal{K}^{\mu_j}}},$$

where δ_x is the Dirac measure centered at x . Following classical procedures, we define the empirical risk as

$$\begin{aligned} \mathfrak{R}(\mathcal{P}) &:= \int_{\mathcal{P}} \int_{\mathcal{K}^\mu} \|(\text{Id} - A^\mu \mathcal{P}(\cdot, \mu))v\|_{V'} \varrho_{N_{\mathcal{K}^\mu}}(dv) \varrho_{N_p}(d\mu) = \\ &= \frac{1}{N_p} \sum_{j=1}^{N_p} \frac{1}{N_{\mathcal{K}^{\mu_j}}} \sum_{v \in \mathcal{K}^{\mu_j}} \|(\text{Id} - A^{\mu_j} \mathcal{P}(\cdot, \mu_j))v\|_{V'}. \end{aligned} \tag{15}$$

Then, the *training phase* consists in solving the following minimization problem

$$\mathcal{P}_* := \arg \min_{\mathcal{P} \in \mathcal{H}} \mathfrak{R}(\mathcal{P}). \tag{16}$$

Note that while this step might be computationally demanding, it has to be performed only once. After training, the preconditioner can be readily applied to any problem in the parametric class, without further optimizations or assembling stages. In particular, if one designs the hypothesis space to consist of neural network models, this can provide significant speed ups, as the computational cost associated with online evaluations of \mathcal{P}_* becomes nearly negligible (on the order of milliseconds).

4. Realization of the learning approach

In this section, we return to the finite-dimensional setting to make the preconditioner learning strategy specific to the discretization of the parametrized mixed-dimensional problem (9)-(10). As we address here the 3D problem arising from the elimination of the 1D equations from the coupled system, with a slight abuse of notation, we write N_h instead of N_h^Ω . Once a suitable preconditioner $\mathcal{P} : \mathbb{R}^{N_h} \times \mathbb{R}^{N_h} \rightarrow \mathbb{R}^{N_h}$ has been identified (we recall that the finite-dimensional parameter space is $\mathcal{P}_h \subset \mathbb{R}^{N_h}$), it can be used to accelerate the solution of the discrete problem. In general, this can be achieved in multiple ways. Here, we propose the integration of the nonlinear preconditioner within a GMRES solver, motivated by the fact that we cannot guarantee the symmetry of the neural preconditioner. In addition, considering the nonlinear nature of the preconditioner \mathcal{P} , we shall adopt the Flexible GMRES (FGMRES) algorithm [28], as the latter can easily accommodate variations of the preconditioning operator with the input vector.

4.1. The training set \mathcal{K}^μ

The selection of the training set \mathcal{K}^μ is a crucial factor in the success of the operator learning approach for preconditioning. In fact, the training set must match the underlying behavior of the operator to ensure that the learned preconditioner is effective in different instances of the parameter space \mathcal{P} . To start, we notice that inside the FGMRES algorithm, the preconditioner will act only on vectors from the unit sphere $\mathbb{S}^{N_h-1} := \{v \in \mathbb{R}^{N_h} : \|v\| = 1\}$. Thus, we let $\mathcal{K}^\mu \subset \mathbb{S}^{N_h-1}$. Next, based on the findings presented in [13], we propose that the training set should be guided by the kernel structure of the operator. In fact, the coupling between the 3D and 1D domains introduces significant challenges in the construction of efficient solvers, particularly because the kernel of the operator $M_{h,00}^\mu$ contains high-frequency components that must be taken into account by the computational solvers. In the context of the algebraic multigrid (AMG) approach addressed in [13], it is shown that an essential property for the robustness of the AMG solver is the following kernel decomposition condition:

$$\ker(M_{h,00}^\mu) = \ker(M_{h,00}^\mu) \cap V_c + \sum_{j=1}^J \ker(M_{h,00}^\mu) \cap V_j,$$




Fig. 2. Pictorial summary for the construction of the augmented training set \mathcal{K}^μ . The data augmentation set \mathcal{D}^{μ_j} contains vectors in the unit sphere with the desired frequency content. Here, $\mathcal{B}/\|\mathcal{B}\| := \{b/\|b\| : b \in \mathcal{B}\} \subset \mathbb{S}^{N_h-1}$ is the set of normalized right-hand side vectors; $\mu_j \in \mathcal{P} \subset \mathbb{R}^p$ denotes a parameter instance that is part of the training set.

where in the case of AMG algorithms V_c refers to the coarse subspace that captures the low-frequency components of the solution and V_j are the fine subspaces that represent localized corrections in specific regions or aggregates. In the context of a multiresolution method such as AMG, this condition suggests that the numerical solver should operate on all components of the kernel of the coupling operator. The richer the kernel subspace, possibly including low- and high-frequency components, the harder it is to develop good solvers. This issue is closely related to the concept of *spectral bias* in neural network training, which states that lower-frequency components are learned more easily and earlier during training, while higher-frequency components are more challenging to learn [29,30]. For the approximation of mixed-dimensional problems, this spectral bias has direct implications. Specifically, the high-frequency components in the kernel of the coupling operator $M_{h,0}^\mu$ may not be adequately captured if the training process is not designed to address this inherent difficulty. To mitigate this challenge, we include in \mathcal{K}^μ samples that span both the low-frequency and high-frequency components of the kernel of C_h^μ , defined in (10). In the context of operator learning, this insight is reflected in the definition of the measure $\rho_\mu(dv) = \rho_{N_{\mathcal{K}^\mu}}$. This considered, we propose the following training set structure:

$$\mathcal{K}^\mu = \bigcup_{j=1}^{N_p} \mathcal{K}^{\mu_j}, \quad \mathcal{K}^{\mu_j} := \left\{ \frac{b_h^{\mu_j}}{\|b_h^{\mu_j}\|} \right\} \cup \mathcal{D}^{\mu_j}, \quad b_h^{\mu_j} \in \mathcal{B}; \tag{17}$$

where \mathcal{D}^{μ_j} represents a suitable data-augmentation set, given the parameter μ_j (see Fig. 2). In this way, the operator $\mathcal{P}(\cdot, \mu)$ can be trained to act as a preconditioner on a suitable subset of the unit sphere with a specified frequency content. In what follows, we examine two distinct strategies for data augmentation, while the next section will present numerical experiments that support this methodology and demonstrate that incorporating both right-hand side vectors and kernel components into \mathcal{K}^μ improves the convergence rates in FGMRES with the application of the learned preconditioner.

4.1.1. Augmented training set by Krylov subspaces

Our initial data-augmentation strategy involves incorporating high-frequency elements derived from the construction of Krylov subspaces related to the matrix. As a starting point, the Krylov subspace associated with a general operator $O \in \mathcal{L}(X, X)$ and a vector $x \in X$ is defined as:

$$\mathcal{K}_q(O, x) := \text{span}(x, Ox, O^2x, \dots, O^{q-1}x)$$

where q is the order of the Krylov subspace. This subspace contains progressively higher-order components of x under the action of the operator O , capturing information about the behavior of the system across multiple scales. By including vectors from these Krylov subspaces in our training set, we effectively augment the representation of both low- and high-frequency components in the solution space.

Specifically, for a generic parameter μ_j , let $O = C_h^{\mu_j}$ be the discrete operator associated with the mixed-dimensional PDEs described in previous sections. We construct the data-augmentation set \mathcal{D}^{μ_j} as

$$\mathcal{D}^{\mu_j} := \left\{ q_{p'}^{\mu_j} \dots q_{(p+p')}^{\mu_j} \right\},$$

where $0 \leq p'$ is a suitable shift index and $q_i^{\mu_j}$, $i = 0, \dots, p' + p$ is a suitable orthonormal basis of $\mathcal{K}_{(p+p')}(C_h^{\mu_j}, b_h^{\mu_j})$ with $q_0^{\mu_j} \equiv b_h^{\mu_j} / \|b_h^{\mu_j}\|$. These additional basis functions aim to capture a wider spectrum of the operator's behavior, thus enhancing the preconditioner's ability to address different scales and frequencies present in the solution. In this context, it is convenient to keep p and p' small and of the same magnitude (i.e. $p \simeq p'$). This lies in the balance between representing the essential high-frequency characteristics of the operator C_h^μ and maintaining computational efficiency. The choice of small p, p' allows us to efficiently construct the augmented training set without significantly increasing the computational burden during the training phase. As it will be shown later on (see Fig. 6 in Section 5.5), even with a small value of $p + p'$, the Krylov subspace captures important high-order effects that improve the robustness of the preconditioner.

4.1.2. Augmented training set by random vectors

A second approach involves adding to the physics-based training subset \mathcal{B} a new set of completely unrelated random functions. Specifically, we propose augmenting the training set by generating random unity vectors $\mathbf{r}_h \in \mathbb{S}^{N_h-1} \subset \mathbb{R}^{N_h}$, uniformly sampled

from the unit hypersphere. In practice, each \mathbf{r}_h can be generated by drawing and normalizing a random vector $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{N_h})$ with standard multivariate normal distribution, that is, $\mathbf{r}_h = \mathbf{v}/\|\mathbf{v}\|$. Such vectors are designed to introduce high-frequency content that is independent of the physics of the problem, thereby complementing the functions based on physics $\mathbf{b}_h^\mu \in \mathcal{B}$. In fact, given a finite element discretization based on a FEM space $V_h(\Omega)$, where Ω represents the three-dimensional computational domain and h denotes the characteristic element size, the vectors \mathbf{r}_h can be interpreted as the highest frequency modes that can be resolved by the mesh. The frequency associated with these random vectors is proportional to $|\Omega|/h$, which corresponds to the smallest possible wavelength that can be captured in the discretized domain. Mathematically, we can define the data augmentation set as $\mathcal{D}^{\mu_j} := \{\mathbf{r}_{h,1}, \mathbf{r}_{h,2}, \dots\}$. Note that these vectors are independent of μ but a different instance of \mathcal{D} is defined for each μ_j . We expect this strategy to improve robustness with respect to high-frequency vectors that are inside $\ker(M_{h,00}^\mu)$.

4.2. Hypothesis space for \mathcal{P}

At this stage, it remains to define the hypothesis space \mathcal{H} , over which the empirical risk minimization will be performed (16). Following a deep-learning approach, the hypothesis space is specified once a particular *neural architecture* \mathcal{N}_H is selected. This architecture is characterized by a composition of linear and nonlinear functions, which in turn are defined by a set of *hyper-parameters* H . The resulting hypothesis space \mathcal{H} can thus be expressed as $\mathcal{H} = \{\mathcal{N}_H(\cdot; \theta)\}_{\theta \in \mathbb{R}^l}$ where θ is the vector of trainable parameters associated with \mathcal{N}_H , over which the optimization is performed. Drawing inspiration from cutting-edge deep learning techniques in image processing, as proposed in [23], which employ varying resolutions to effectively manage the diverse scales of images through pooling and upscaling operations, the neural architecture chosen to represent the preconditioner will be a *U-net*, see for example [31] for a landmark paper and [32] for recent developments. This architecture is highly capable of addressing multi-scale issues as it effectively captures both global and local features across varying resolutions. This makes it a strong candidate function for representing solutions of mixed-dimensional problems, such as the ones represented in Fig. 1. According to this choice, we define $\mathcal{N}_H(\cdot; \theta) \equiv \mathcal{U}_L(\cdot; \theta)$, where \mathcal{U}_L represents a U-net architecture with L levels. For the sake of simplicity, we will refer only to the hyper-parameter L to indicate the depth of the U-net, although additional parameters such as bottleneck size, number of channels, and convolutional kernel sizes are also involved. Consequently, the desired nonlinear preconditioner operator can be represented as $\mathcal{P} \equiv \mathcal{U}_L(\cdot; \theta^*)$, where θ^* denotes the set of optimized parameters. The U-net architecture is structured as a combination of two components, an encoder function Φ and a decoder function Ψ :

$$\begin{aligned} \Phi &:= \Phi_0 \circ \Phi_1 \circ \dots \circ \Phi_{L-2} \circ \Phi_{L-1} : \mathbb{R}^{c_{in} \times n} \rightarrow \mathbb{R}^{c_b \times m}, \\ \Psi &:= \Psi_{L-1} \circ \Psi_{L-2} \circ \dots \circ \Psi_1 \circ \Psi_0 : \mathbb{R}^{c_b \times m} \rightarrow \mathbb{R}^{c_{out} \times n}. \end{aligned}$$

Then, given an input tensor $\mathbf{X} = [\mathbf{X}_1 | \mathbf{X}_2 | \dots | \mathbf{X}_{c_{in}}]$, where $\mathbf{X}_i \in \mathbb{R}^{n_1 \times n_2 \times \dots \times n_d}$, and a prescribed number of output channels c_{out} , a U-net with j levels can be defined recursively as follows:

$$\mathcal{U}_j(\mathbf{X}) := \Psi_j([\mathcal{U}_{j-1} \circ \Phi_j | \Phi_j](\mathbf{X})), \quad \mathcal{U}_0(\mathbf{X}) := \Psi_0 \circ \Phi_0(\mathbf{X}),$$

where the notation $[a|b]$ represents a *channel stacking* operation that concatenates tensors a and b along the channel dimension. Specifically, if $a \in \mathbb{R}^{c_a \times m}$ and $b \in \mathbb{R}^{c_b \times m}$, then we have $[a|b] \in \mathbb{R}^{(c_a+c_b) \times m}$. Operation $[\mathcal{U}_{j-1} \circ \Phi_j(\mathbf{X}) | \Phi_j(\mathbf{X})]$ is called the j th *skip connection*, which transmits information from the j th encoding layer to the corresponding j th decoding layer. Skip connections are critical in retaining and merging high-resolution features, ensuring that the decoder can effectively recover fine-scale information. In the literature, the term c refers to the *number of channels* at each layer. Specifically, c_b is the number of channels in the bottleneck layer of the U-net, while c_{in} and c_{out} denote the number of input and output channels, respectively. The convolutional layer operates locally by sliding a kernel across the input tensor, extracting features from small spatial neighborhoods. Given a d -dimensional input tensor $\mathbf{X} \in \mathbb{R}^{c_{in} \times n_1 \times n_2 \times \dots \times n_d}$, the convolution operation is defined as:

$$\mathbf{Y}_i = \sum_{j=1}^{c_{in}} \mathbf{k}_{i,j} * \mathbf{X}_j \text{ for } i = 1, 2, \dots, c_{out}, \text{ with } (\mathbf{k}_{i,j} * \mathbf{X}_j)(\mathbf{p}) = \sum_{\mathbf{q} \in \mathbb{Z}^d} \mathbf{k}_{i,j}(\mathbf{q}) \mathbf{X}_j(\mathbf{p} - \mathbf{q}),$$

$\mathbf{k}_{i,j} \in \mathbb{R}^{s_1 \times s_2 \times \dots \times s_d}$ being the kernel for the (i, j) th input-output channel pair, \mathbf{p} indexing the output tensor, and \mathbf{q} spanning the kernel positions. This operation enables convolutional layers to act as local feature extractors, capturing patterns from the input tensor. Input data, often provided as ordered vectors, must be reshaped into tensor format to preserve spatial correlations. Once trained, the kernel can operate on input tensors of arbitrary dimensions, provided the structure is compatible, such as tensor-product meshes. As a result, the U-Net architecture can handle inputs of any size, which may increase its application beyond the dimensions of the training data. However, reliance on structured grids restricts the application of convolutional layers to simple geometries. Future work will address this limitation by integrating mesh-informed neural networks, as proposed in [33], to generalize U-net architectures while retaining their feature extraction capabilities.

5. Numerical solution of mixed-dimensional PDEs using the neural preconditioner

We analyze how the neural preconditioner accelerates the convergence of iterative solvers, in the case of mixed-dimensional PDEs. We discuss the computational setup and analyze the preconditioner’s impact on FGMRES convergence rate, mesh size scalability, and computational time. In doing so, we shall take the opportunity to assess the importance of the several design choices entailed by the implementation of the neural preconditioner, such as, e.g., the data-augmentation strategy and the parametrization of the 1D

Table 1
Summary of the \mathcal{U}_3 architecture. Input data flows from the top row to the bottom row; see also the scheme in Fig. 3. The total number of trainable parameters is $\theta \approx 10^6$.

Layer type	input size	output size	input channels	output channels	# parameters θ
Conv.	21^3	21^3	2	32	13,840
Conv. + Max P.	21^3	10^3	32	64	55,296
Conv. + Max P.	10^3	5^3	64	128	221,184
Conv.	5^3	5^3	128	128	442,368
Trasp. conv. + Skip	5^3	10^3	128	64	286,784
Trasp. conv. + Skip	10^3	21^3	64	32	71,912
Conv.	21^3	21^3	32	1	545




Fig. 3. Schematic representation of the U-Net architecture \mathcal{U}_3 . Tensor data are represented by blocks, where the number of channels corresponds to the depth of the blocks. The tensor shape, $c \times n_1 \times n_2 \times n_3$, is reported on the top of the blocks. Solid line arrows represent layer action, while dashed arrows represent channel-staking skip connection.

geometry. For the sake of this preliminary exploration, in this Section we shall restrict ourselves to numerical solution of problem (10), that is, to the 3D block of the mixed-dimensional system, thus ignoring the 3D-1D coupling. The latter will be addressed right after, in Section 6.

We anticipate that the main message emerging from this numerical exploration is that the neural preconditioner proves competitive when compared to other state-of-the-art algorithms—such as AMG or ILU— thanks to a reduction in the computational cost of each preconditioning step. However, the iteration count and the scalability of the neural preconditioner with respect to the mesh size are not optimal. In fact, as shown later in Table 11, a trade-off between computational cost and algorithmic efficiency appears. However, our findings indicate that, in general, the neural preconditioner offers a competitive approach to solving mixed-dimensional PDEs, which motivates its use in Section 6 when addressing the coupled problem.

5.1. General setup of the numerical tests

The domain Ω is the unit cube, and the problem is discretized using a structured mesh with 9261 nodes (21^3 grid points). Physical constants are set as $k_\Omega = 10^{-3}$ and $\sigma_\Omega = 10^{-3}$, corresponding to a low diffusion-reaction regime in Ω . In addition, we performed tests using $k_\Omega = 10^{-2}$, $\sigma_\Omega = 10^{-2}$, confirming the adaptability of the method. The magnitude of the coupling term is set by fixing the radius $\epsilon = 10^{-3}$ leading to $\gamma = 2\pi\epsilon \approx 6,28 \cdot 10^{-3}$ setting the model in a regime where the coupling term $M_{h,00}^\mu$ is comparable or greater to the block $K_{h,00}$. Concerning the problem in Λ we set $k_\Lambda = 1$, coherently with the general observation that the diffusion in the channels Λ is generally larger than the one in Ω , although the actual magnitude of this coefficient may not exactly correspond to any specific application. Although model (1) is not intended to represent a specific application, choice $k_\Omega, k_\Lambda, \sigma_\Omega, \epsilon$ is consistent with the range of physiologically relevant parameters for oxygen transfer in the microcirculation, described in detail in [5], for example. Here, our aim is to capture the correct balance between diffusion, reaction, and coupling, rather than the precise numerical values of biological constants. The set of right-hand sides considered in the numerical experiments includes non-homogeneous contributions, such as constant and oscillatory functions; however, unless otherwise specified, the tests have been performed with the homogeneous value $\mathbf{f}_{h,0} := \mathbf{0}$. For the training set \mathcal{K}^μ , different 1D graph geometries are considered (see Fig. 4), ranging from $O(1)$ to $O(10^2)$ number of branches.

We considered a three-level U-net architecture, \mathcal{U}_3^μ , with two input channels: one relative to the vector in the training set \mathcal{K}^μ and one carrying the parameter information μ , that is, taking the discrete distance function $\mathbf{d}_h^\mu \in \mathbb{R}^{N_h}$ associated with the graph Λ . More technical details can be found in Table 1 and Fig. 3. The actual training set $\tilde{\mathcal{B}} \approx \mathcal{B}$ is obtained by solving approximately the system in (11), with a relative tolerance of 10^{-4} , thus ensuring computational efficiency in training set generation. The data-augmentations

Table 2
Details on the training dataset and parameters. Due to data augmentation strategies, the total number of training samples is five times the one of the 1D graphs.

Training	1D graphs	Tot. samples	Validation	1D graphs	Tot. samples
Samples	480	2400	Samples	120	600
Batch Size		5	Epochs		250
Learning rate		10^{-3}	Learning strategy		decaying




Fig. 4. Examples of the graphs Λ considered in the numerical tests, with increasing geometrical complexity.

sets D^{μ_j} vary depending on the scenario analyzed, resulting in three different U-Net preconditioners:

$$\begin{aligned}
 D^{\mu_j} = \emptyset &\Rightarrow \mathcal{U}_3^{\emptyset}(\cdot; \theta), \\
 D^{\mu_j} = \{q_5^{\mu_j}, \dots, q_8^{\mu_j}\} &\Rightarrow \mathcal{U}_3^{kry}(\cdot; \theta), \\
 D^{\mu_j} = \{r_{h,1}, \dots, r_{h,4}\} &\Rightarrow \mathcal{U}_3^{hf}(\cdot; \theta).
 \end{aligned}$$

Here, $\{q_5^{\mu_j}, \dots, q_8^{\mu_j}\}$ are orthonormal bases of the Krylov subspace $\mathcal{K}_8(C_h^{\mu_j}, b_h^{\mu_j})$, and $\{r_{h,1}, \dots, r_{h,4}\}$ are high-frequency vectors sampled from the unit sphere \mathbb{S}^{N_h-1} . The networks \mathcal{U}_3 are trained through the minimization of the empirical risk:

$$\mathcal{U}_3^* = \mathcal{U}_3(\theta^*) \text{ with } \theta^* = \arg \min_{\theta} \frac{1}{N_p} \sum_{j=1}^{N_p} \frac{1}{N_{\mathcal{K}_{\mu_j}}} \sum_{v \in \mathcal{K}_{\mu_j}} \left\| v - \alpha^{-1} C_h^{\mu_j} \mathcal{U}_3(v, \mathbf{d}^{\mu_j}; \theta) \right\|^2,$$

where $\alpha := \left(\frac{1}{N_p} \sum_{j=1}^{N_p} \frac{1}{\sqrt{N_h}} \|diag(C_h^{\mu_j})\|_2 \right)$ is a fixed normalizing constant representing the mean matrices diagonal root mean square, which helps stabilize the training process. For comparison purposes, an equal number of graph samples and training epochs is used for $\mathcal{U}_3^{\emptyset}$, \mathcal{U}_3^{kry} , and \mathcal{U}_3^{hf} . More data on the training process are reported in Table 2.

Fig. 5 displays the convergence history of training and validation for the U-Net architecture when the physical constants $k_{\Omega} = 10^{-3}$, $\sigma_{\Omega} = 10^{-3}$ are considered.

The graphs illustrate that the training and validation error for $\mathcal{U}_3^{\emptyset}$, \mathcal{U}_3^{kry} and \mathcal{U}_3^{hf} possesses similar trends but quite different values; nevertheless, as will be clarified later on, the training and validation error alone are not discriminant for the performance of the U-Net preconditioner. What is fundamental is the data set over which the given error is obtained; e.g., even if $\mathcal{U}_3^{\emptyset}$ and \mathcal{U}_3^{kry} possess similar error behavior in the validation set, their performance as a preconditioner is quite different (cf. Table 3).

The performance of the proposed neural preconditioner is evaluated by comparing it with two well-established preconditioning strategies: the Incomplete LU (ILU) factorization and the Algebraic Multigrid (AMG) method [26]. The ILU preconditioner approximates the LU factorization of a matrix by allowing a controlled level of fill-in, which provides a balance between sparsity and

Table 3
Mean FGMRES iterations for different preconditioners, averaged across 100 different test configurations. Δ^{\pm} = difference in iteration counts between the average case (Mean Iter.) and the worst (best) case in the test set, respectively.

Preconditioner	Type	$k_{\Omega} = \sigma_{\Omega} = 10^{-3}$			$k_{\Omega} = \sigma_{\Omega} = 10^{-2}$		
		Mean Iter.	Δ^+	Δ^-	Mean Iter.	Δ^+	Δ^-
None	-	147.72	+83.28	-52.72	187.04	+121.96	-70.04
$\mathcal{U}_3^{\emptyset}$	No Augmentation	92.41	+17.59	-16.41	42.13	+10.87	-6.13
\mathcal{U}_3^{kry}	Krylov Subspace	24.54	+3.89	-4.11	15.54	+1.46	-2.54
\mathcal{U}_3^{hf}	Random Vectors	17.82	+5.18	-2.82	16.63	+2.37	-1.63




Fig. 5. Convergence of train (left) and validation (right) relative errors of the U-Net U_3^* when trained on different datasets, for the physical constants $k_\Omega = 10^{-3}$, $\sigma_\Omega = 10^{-3}$.

approximation accuracy, making it effective for moderately ill-conditioned systems. AMG is a hierarchical method that constructs a sequence of progressively coarser spaces to efficiently capture both low- and high-frequency error components, thereby accelerating convergence for a wide range of linear systems. Both ILU and AMG preconditioners have been implemented using the PETSc library (<https://petsc.org/release/>) [34], which offers efficient and scalable linear algebra routines for large-scale computational problems, together with the PETSc python wrapper `petsc4py` (<https://pypi.org/project/petsc4py>) [35].

All these preconditioners are tested on a flexible GMRES algorithm with right preconditioning, designated as $\text{FGMRES}(k)$, restarted every $k = 20$ steps. Iterations are stopped when the relative residual satisfies $\|r\|/\|r_0\| \leq 10^{-6}$. In all tables, we show the average number of iterations required by the algorithm to converge (denoted as **Mean Iter.**), for an ensemble of linear systems corresponding to 10^2 unseen graphs Λ , ranging from $O(1)$ to $O(10^2)$ branches.

5.2. Effect of the data-augmentation strategy, forcing terms and parameters

We analyze the performance of the proposed preconditioners under three augmentation strategies: (i) no augmentation (U_3^0), (ii) Krylov subspace augmentation (U_3^{kry}), and (iii) random high-frequency vector augmentation (U_3^{hf}). To explore the robustness of the approach, the tests are repeated for two values of the physical parameters. The results are reported in Table 3.

For the most challenging parameter choice, $k_\Omega = \sigma_\Omega = 10^{-3}$, the unaugmented U-Net preconditioner (U_3^0) is less effective, requiring 92.41 iterations on average, which highlights its limited ability to capture the full spectrum of the solution space. Incorporating Krylov subspace augmentation (U_3^{kry}) reduces the mean iteration count significantly (24.54), as this strategy enriches the training set with information aligned with the operator spectrum. Random vector augmentation (U_3^{hf}) achieves the lowest iteration count (17.82), demonstrating superior performance over the other augmentation strategies.

We have further extended this analysis to account for non-homogeneous right-hand sides. In addition to the homogeneous case ($\mathbf{f}_{h,0} \equiv 0$), we considered two representative forcing terms:

$$\mathbf{f}_{h,0,\text{const}} := 10^{-3} \cdot \mathbf{1}, \quad \mathbf{f}_{h,0,\text{osc}} := A(\sin(\pi k_x x) + \sin(\pi k_y y) + \sin(\pi k_z z)),$$

with $A = 10^{-3}$ and $k_x = k_y = k_z = 2$. The numerical results reported in Table 4 confirm that the neural preconditioner remains stable and effective in all these scenarios. In the homogeneous case, iteration counts range from 17 to 25, while in the constant and oscillatory forcing cases, they remain of the same order (approximately 21–28). This shows that the performance of the preconditioner is not significantly affected by the type of right-hand side. Moreover, the comparison between Krylov and random augmentation essentially confirms the conclusion of Table 3.

Finally, we comment on the effect of the coupling term γM_h^μ in the preconditioned system. As the relative weight of the metric contribution increases with respect to the diffusion–reaction terms, the conditioning of the underlying operator is known to deteriorate [36]. Our numerical experiments confirm that the neural preconditioner reflects this behavior: its effectiveness is progressively reduced when the coupling term becomes dominant. This outcome is consistent with the general theory of operator preconditioning for mixed-dimensional PDEs [11,12,36], where the growth of the ratio γ/k_Ω is expected to amplify the condition number of the problem. Thus, while the neural preconditioner remains effective across a broad range of parameter values, it does not circumvent the intrinsic difficulties associated with strongly coupled regimes, highlighting a structural limitation that is shared with classical preconditioning techniques.

5.3. Impact of the 1D graph parameterization

We also investigated to what extent including the distance function \mathbf{d}_h^μ as input enhances the effectiveness of the neural preconditioner $U(\cdot, \mathbf{d}_h^\mu; \theta)$ for different instances of 1D domains. As shown in Table 5, including the distance function \mathbf{d}_h^μ as an input significantly improves its performance, particularly in scenarios where the μ -dependent component $M_{h,00}^\mu$ dominates. In fact, we

Table 4

Mean FGMRES iterations for different preconditioners with zero, constant, and oscillatory right-hand sides. Δ^\pm = difference in iteration counts between the average case (Mean Iter.) and the worst (best) case in the test set, respectively.

Precond.	$\mathbf{f}_{h,0} \equiv 0$			$\mathbf{f}_{h,0} \equiv \mathbf{f}_{h,0, \text{const}}$			$\mathbf{f}_{h,0} \equiv \mathbf{f}_{h,0, \text{osc}}$		
	Mean Iter.	Δ^+	Δ^-	Mean Iter.	Δ^+	Δ^-	Mean Iter.	Δ^+	Δ^-
$\mathcal{U}_3^{\text{krv}}$	24.54	+3.89	-4.11	20.69	+4.31	-3.69	22.4	+4.8	-3.4
$\mathcal{U}_3^{\text{hlf}}$	17.82	+5.18	-2.82	27.92	+7.08	-5.92	22.98	+17.02	-4.98

Table 5

Mean FGMRES iterations with and without influence of the distance function. Δ^\pm = difference in iteration counts between the average case (Mean Iter.) and the worst (best) case in the test set, respectively.

Preconditioner	$k_\Omega = \sigma_\Omega = 10^{-3}$			$k_\Omega = \sigma_\Omega = 10^{-2}$		
	Mean Iter.	Δ^+	Δ^-	Mean Iter.	Δ^+	Δ^-
$\mathcal{U}_3^{\text{d}}(\cdot)$	108.65	+30.35	-29.65	41.50	+8.50	-5.50
$\mathcal{U}_3^{\text{d}}(\cdot, \mathbf{d}_h^\mu)$	92.41	+17.59	-6.41	42.13	+10.87	-6.13
$\mathcal{U}_3^{\text{krv}}(\cdot)$	44.91	+8.09	-7.91	20.69	+2.31	-3.69
$\mathcal{U}_3^{\text{krv}}(\cdot, \mathbf{d}_h^\mu)$	24.54	+3.89	-4.11	15.54	+1.46	-2.54
$\mathcal{U}_3^{\text{hlf}}(\cdot)$	60.31	+26.69	-19.31	24.98	+6.02	-4.98
$\mathcal{U}_3^{\text{hlf}}(\cdot, \mathbf{d}_h^\mu)$	17.82	+5.18	-2.82	16.63	+2.37	-1.63

Table 6

Effect of pre- and post-smoothing on mean FGMRES iterations. Δ^\pm = difference in iteration counts between the average case (Mean Iter.) and the worst (best) case in the test set, respectively. Rel. Gain = relative reduction of the iteration counts when using Jacobi relaxation smoothing.

Precond.	Relax	$k_\Omega = \sigma_\Omega = 10^{-3}$				$k_\Omega = \sigma_\Omega = 10^{-2}$			
		Mean Iter.	Δ^+	Δ^-	Rel. Gain	Mean Iter.	Δ^+	Δ^-	Rel. Gain
\mathcal{U}_3^{d}	None	92.41	+17.59	-16.41	0.64	42.13	+10.87	-6.13	0.47
\mathcal{U}_3^{d}	Jacobi	33.72	+13.28	-7.72		22.47	+3.53	-4.47	
$\mathcal{U}_3^{\text{krv}}$	None	24.54	+3.89	-4.11	0.49	15.54	+1.46	-2.54	0.26
$\mathcal{U}_3^{\text{krv}}$	Jacobi	12.6	+3.40	-2.60		11.46	+0.54	-1.46	
$\mathcal{U}_3^{\text{hlf}}$	None	17.83	+5.17	-2.83	0.26	16.63	+2.37	-1.63	0.22
$\mathcal{U}_3^{\text{hlf}}$	Jacobi	13.12	+2.88	-2.12		12.96	+2.04	-0.96	

recall that the structure of the linear system is defined as $C_h^\mu := K_{h,00} + \gamma M_{h,00}^\mu$, where $K_{h,00}$ is the fixed component, and $M_{h,00}^\mu$ introduces the dependence of parameters. Without providing \mathbf{d}_h^μ , the preconditioner mainly targets $K_{h,00}$, leading to lower performance on the whole system C_h^μ .

5.4. Effect of pre- and post-smoothing

To improve the efficacy of the preconditioning method, similar to other studies [20,23], we can employ techniques for pre- and post-smoothing. The smoothing procedure helps mitigate high-frequency error components, particularly for preconditioners that lack high-frequency training data. For a given residual \mathbf{r} and initial guess \mathbf{x}_0 , we consider a smoothing algorithm with Jacobi relaxation $S_J(\mathbf{A}, \mathbf{r}, \mathbf{x}_0, \text{maxit})$ [37]. The impact of pre-post smoothing is summarized in Table 6. The results show that preconditioning strategies with low-frequency training data (\mathcal{U}_3^{d}) benefit the most from smoothing, reducing iterations by 60%.

This test sheds light on an important property of the data-augmented training strategy. Essentially, in Table 6 we observe that pre- and post-smoothing has an effect comparable to that of data augmentation. This reveals that the data-augmentation strategy is equivalent to adding a pre- and post-smoother to the preconditioner, helping in learning how to process high frequencies and contrast the spectral bias. This property may turn out to be particularly effective in those cases where an efficient smoother is not known for the problem at hand. This conclusion is further supported by the fact that the high-frequency data augmentation strategy $\mathcal{U}_3^{\text{hlf}}$ is less sensitive to smoother application.

5.5. Spectral analysis of training data and operator kernel

As previously observed, the introduction of high-frequency vectors in the training set is beneficial, with the augmentation strategy based on random vectors slightly outperforming that using the Kylov subspace. In this section, we seek to explain the observed behavior by contrasting the two training-set enrichment techniques with the operator’s intrinsic kernel structure. In detail, to quantify the operator spectral content, we analyzed the null spaces of the problem matrices $K_{h,00}$, $\gamma M_{00,h}^\mu$, and C_h^μ using Singular Value

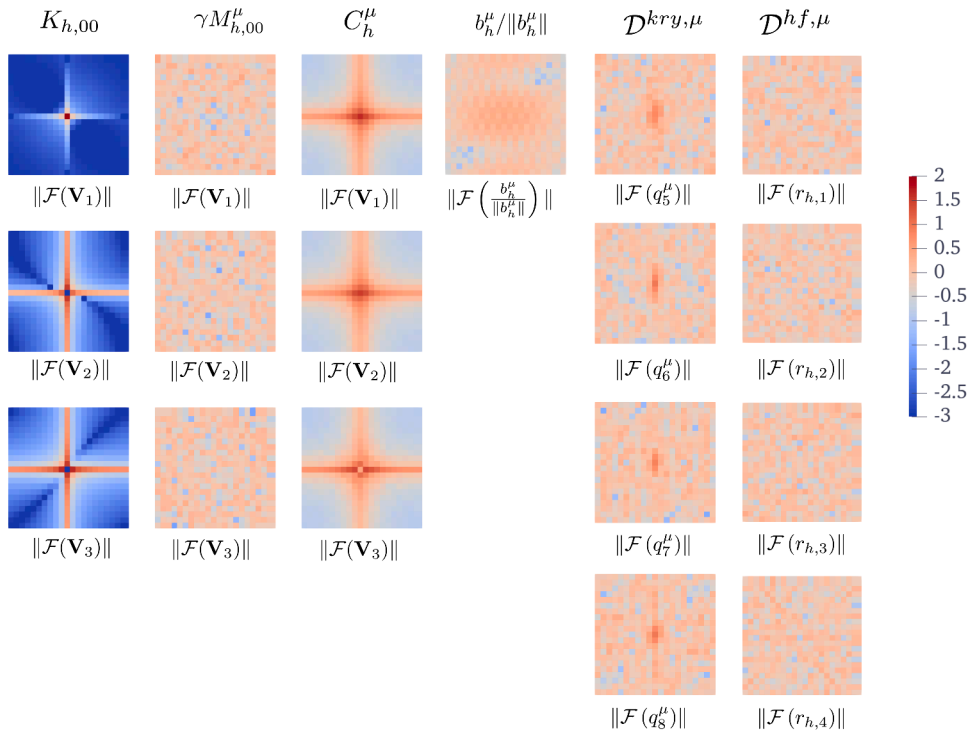


Fig. 6. Logarithmic plot in the frequency domain (k_1, k_2, k_3) of the $k_2 k_3$ -cross section for the spectrum magnitude $\|\mathcal{F}(\mathbf{V}_i)\|$; $k_1 = 0$ is considered and a value for μ is fixed given $k_\Omega = \sigma_\Omega = 10^{-3}$. Frequency components increase with the distance from the center $(0, 0, 0)$. **Column I:** magnitude of first three right-singular vectors associated to $K_{h,00}$; **Column II:** magnitude of first three right-singular vectors associated to $\gamma M_{h,00}^\mu$; **Column III:** magnitude of first three right-singular vectors associated to $C_h^\mu = K_{h,00} + \gamma M_{h,00}^\mu$; **Column IV:** spectrum magnitude of the normalized forcing term b_h^μ ; **Column V:** spectrum magnitude of the orthonormal Krylov bases $\mathcal{D}^{kry;\mu} := \{q_5^\mu, \dots, q_8^\mu\}$ (note that $q_0^\mu \equiv b_h^\mu / \|b_h^\mu\|$); **Column VI:** spectrum magnitude of the random high frequency data augmentation set $\mathcal{D}^{hf;\mu} := \{r_{h,1} \dots r_{h,4}\}$.

Decomposition (SVD). Given a matrix $A \in \mathbb{R}^{m \times n}$ with rank r , the associated SVD is given by $A = U \Sigma V^T$, where $U = [\mathbf{u}_1 \mid \mathbf{u}_2 \mid \dots \mid \mathbf{u}_r]$ and $V = [\mathbf{v}_1 \mid \mathbf{v}_2 \mid \dots \mid \mathbf{v}_r]$ contain the left and right singular vectors, respectively. The right singular vectors \mathbf{v}_i corresponding to zero singular values ($\sigma = 0$) span the null space, capturing kernel components that influence preconditioning performance. Using sparse SVD algorithms (`scipy.sparse.svds`), we extracted some of these vectors and analyzed them in frequency space via the multidimensional discrete Fourier transform. Let us denote by \mathbf{V} the d -dimensional tensor corresponding to a right singular vector \mathbf{v} . Then we denote the DFT as \mathcal{F} , defined as follows,

$$\mathcal{F} : \mathbf{V} \in \mathbb{R}_{n_1} \times \dots \times \mathbb{R}_{n_d} \rightarrow \mathcal{F}(\mathbf{V}) \in \mathbb{C}_{k_1} \times \dots \times \mathbb{C}_{k_d}, \quad \mathcal{F}(\mathbf{V})|_{\mathbf{k}} = \sum_{\mathbf{n}=0}^{\mathbf{N}-1} e^{-i2\pi \mathbf{k} \cdot (\mathbf{n}/\mathbf{N})} \mathbf{V}|_{\mathbf{n}},$$

where $\mathbf{n} = (n_1, n_2, \dots, n_d) \in \mathbb{N}^d$, $\mathbf{k} = (k_1, k_2, \dots, k_d) \in \mathbb{N}^d$ and $\mathbf{N} = (N, N, \dots, N) \in \mathbb{N}^d$ are multi-integers and N is the number of edge subdivision for the considered cubic structured mesh.

In Fig. 6 (columns I-III), two-dimensional slices of the spectrum magnitudes $\|\mathcal{F}(\mathbf{V}_i)\|$ associated with the kernel vectors $\{\mathbf{v}_i\}_{i=1,2,3}$ of matrices $K_{h,00}$, $\gamma M_{h,00}^\mu$ and C_h^μ are shown. For better visualization, the so-called *centered* spectrum is considered, where a suitable phase shift $e^{i\pi \mathbf{k} \cdot (\mathbf{N}-1)}$ in the frequency space is applied to $\mathcal{F}(\mathbf{V})$; in this way, the point $\mathcal{F}(\mathbf{V})|_{(0,\dots,0)}$, is centered on the plot. We note that the frequency content in the spectra of matrices $K_{h,00}$ and $\gamma M_{h,00}^\mu$ is remarkably different. As expected, the kernel of the metric term $\gamma M_{h,00}^\mu$ introduces very noisy high-frequency components that need to be preconditioned. Fig. 6 (columns IV-VI) presents a comparison between the different data enhancement strategies, visualized in the frequency domain. Column IV shows the spectrum magnitude of the normalized forcing term. Column V illustrates the spectrum magnitude of the orthonormal Krylov basis. Finally, the last column presents the spectrum magnitude of the random high-frequency data augmentation set. In Fig. 6 we observe that Krylov subspace augmentation captures higher-order components by repeated application of the operator to a basis vector, producing structured high-frequency content (column V). In contrast, random vector augmentation introduces unstructured high-frequency content by randomly drawing components in an independent manner. This results in a broader and more uniform spectral coverage (Fig. 6, column VI), enabling the preconditioner to handle fine-scale variations in the solution, critical for mixed-dimensional problems like the 3D-1D coupled system. From a comparison between the spectral content of different training sets and the spectrum of the kernel component of the metric term $\gamma M_{h,00}^\mu$ (Fig. 6, column II), one can conclude that the Krylov subspace augmentation introduces structured high-frequency components, while random vector augmentation spans a broader frequency range, closer to the true kernel spectrum of the

Table 7

Mean FGMRES iterations and execution time for different preconditioning strategies of the problem with $k_\Omega = \sigma_\Omega = 10^{-3}$. All times are in milliseconds. $\text{Rate}(N_h) =$ for a given quantity of interest y , it is the estimated exponent q yielding $y \propto N_h^q$. (*) Denotes that, for at least a simulation in the test set, FGMRES did not converge in the maximum number of iterations, here set to 1000.

Precond.	N_h	N	Mean Iter.	Rate (N_h)	Time	Rate (N_h)	Time/Iter	Rate (N_h)
None	9,261	21	147.72	-	365.33	-	-	-
	68,921	41	431.91	0.53	7293.49	1.49	-	-
	531,441	81	1000*	-	-	-	-	-
	9,261	21	17.38	-	45.227	-	2.54	-
$\mathcal{U}_3^{hf}(21^3)$	68,921	41	37.24	0.37	497.86	1.2	13.37	1.21
	531,441	81	100.73	0.49	9212.77	1.43	91.46	1.06
	9,261	21	28.81	-	60.54	-	2.10	-
$\mathcal{U}_3^{hf}(41^3)$	68,921	41	45.54	0.27	519.95	1.07	11.42	1.19
	531,441	81	119.27	0.47	11985.8	1.54	100.49	0.94
	9,261	21	24.73	-	37.7	-	1.52	-
ILU	68,921	41	46.2	0.31	486.04	1.27	10.52	1.04
	531,441	81	102.55	0.39	9422.78	1.45	91.88	0.94
	9,261	21	9.56	-	62.35	-	6.52	-
AMG(3)	68,921	41	17.13	0.29	696.14	1.20	40.64	1.10
	531,441	81	33.66	0.33	9631.04	1.29	286.13	1.05
	9,261	21	3.64	-	55.67	-	15.29	-
AMG(10)	68,921	41	4.01	0.05	329.32	1.02	82.12	1.19
	531,441	81	4.9	0.10	3921.46	1.07	800.30	0.90

operator. This broader frequency coverage can explain the better convergence behavior observed in numerical experiments, where random vector-based augmentation results are more effective than Krylov-based strategies in handling ill-conditioning and diverse parameter variations.

5.6. Scalability with respect to the problem size

Due to the intrinsic properties of the convolutional layers in the U-Net architecture, although trained on a low-resolution input tensor $\mathbf{X} \in \mathbb{R}^{21 \times 21 \times 21 \times 2}$, the proposed neural preconditioner \mathcal{U}_3^{hf} can be applied to inputs of arbitrary dimensions, corresponding to higher resolutions of the problem. This adaptability makes it suitable for evaluating its scalability with respect to increasing problem size, quantified by the number of unknowns N_h of the discrete problem.

We test here the convergence rate of the preconditioned FGMRES for the solution of the 3D problem, for three different levels of mesh resolution N , namely [21, 41, 81] points per side of the cubic domain Ω corresponding respectively to a number of degrees of freedom (d.o.f.), N_h , for a \mathbb{P}^1 FEM discretization equal to $[21^3; 41^3; 81^3]$. We measure the average number of iterations across 100 different test cases, all referring to 1D graphs that were not used during the training phase. Then, we analyze how the average iteration count scales with the number of d.o.f., N_h .

To assess the capacity of the neural preconditioner to generalize on the mesh resolution, we compare the neural preconditioner \mathcal{U}_3^{hf} trained on two different resolutions, $N = 21$ and $N = 41$, respectively, and analyze how these two operators behave in the case of discretization based on 81 points. The comparison also includes incomplete LU factorization (ILU) and Algebraic Multigrid (AMG) with 3 and 10 V-cycles, respectively.

In Tables 7 and 8, we see that the scaling behavior of the neural preconditioner \mathcal{U}_3^{hf} closely matches that of the ILU preconditioner. These methods show a dependence on N_h of order $\approx 0.4 - 0.5$. The AMG(3) preconditioner shows a lower iteration count but features a similar scaling with respect to the size of the problem. None of these methods achieves optimal scaling as the AMG(10) preconditioner, the iteration rates of which are almost independent on N_h .

Finally, we note that the performance of the neural preconditioners $\mathcal{U}_3^{hf}(21^3)$ and $\mathcal{U}_3^{hf}(41^3)$ is almost equivalent, indicating that the size and resolution of the training data do not affect the overall performance of the neural operator. This means that it is possible to train the neural preconditioner on low-resolution data and deploy it on higher resolutions, with considerable savings on the computational cost of the offline training phase.

Previous analyses have focused solely on the number of FGMRES iterations as a performance metric. Although the neural preconditioners demonstrated sub-optimal results in terms of convergence rates, it is also essential to consider execution time, as the forward pass through a neural network is computationally lighter than, e.g., a single evaluation of the AMG algorithm. Although the implementation has not been optimized for performance (e.g., no batch-optimized algorithm or GPU acceleration techniques [38]), we also present in Tables 7-8 the results for execution time. Looking at the scaling with N_h of the mean time required to solve the 3D problem, the advantage of AMG(10) is significantly reduced with respect to the neural preconditioner, ILU, and AMG(3), which provide aligned performances. This can be interpreted by observing that for FGMRES with sparse matrices, the overall computational cost is $O(m \times N_h)$, where m is the number of iterations required for convergence. On the one hand, m is sensitive to the choice of the

Table 8
Mean FGMRES iterations and execution time for different preconditioning strategies of the problem with $k_\Omega = \sigma_\Omega = 10^{-2}$. Table entries read as in Table 7.

Precond.	N_h	N	Mean Iter.	Rate (N_h)	Time	Rate (N_h)	Time/Iter	Rate (N_h)
None	9261	21	187.94	-	216.17	-	-	-
	68,921	41	511.34	0.50	3,329.16	1.36	-	-
	531,441	81	1000*	-	-	-	-	-
	9261	21	16.63	-	41.18	-	2.48	-
$\mathcal{U}_3^{hf}(21^3)$	68,921	41	37.98	0.41	461.18	1.20	12.14	1.26
	531,441	81	113.98	0.54	9,908.77	1.50	86.93	1.04
	9261	21	24.27	-	55.73	-	2.30	-
$\mathcal{U}_3^{hf}(41^3)$	68,921	41	41.39	0.27	458.34	1.05	11.07	1.28
	531,441	81	104.99	0.46	9,586.00	1.49	91.30	0.97
	9261	21	31.46	-	42.09	-	1.34	-
ILU	68,921	41	62.18	0.34	548.30	1.28	8.82	1.06
	531,441	81	147.30	0.42	12,123.32	1.52	82.30	0.91
	9261	21	11.50	-	63.95	-	5.56	-
AMG(3)	68,921	41	20.69	0.29	727.50	1.21	35.16	1.09
	531,441	81	42.14	0.35	10,839.18	1.32	257.22	1.03
	9261	21	3.92	-	53.25	-	13.58	-
AMG(10)	68,921	41	4.00	0.01	404.91	1.01	101.23	0.999
	531,441	81	4.00	0.00	3,450.91	1.05	862.73	0.95

preconditioner; precisely as shown in Tables 7-8 column *Mean Iter.* we have $m = (N_h)^{p_i}$ which being $p_i \simeq 0.5$ for \mathcal{U}_3^{hf} , $p_i \simeq 0.4$ for ILU, $p_i \simeq 0.4$ for AMG(3) and $p_i \simeq 0$ for AMG (10), this being the only case where m is independent of N_h . On the other hand, the time per iteration scales linearly with N_h , that is, $T = c_i N_h$ for all methods, but the proportionality constant may be radically different, namely $c_i = 1.72310^{-4}$, 1.88310^{-4} for neural preconditioners trained on points 21^3 and 41^3 , $c_i = 1.72110^{-4}$ for ILU, $c_i = 5.38310^{-4}$ for AMG(3) and $c_i = 14.9810^{-4}$ for AMG(10), with neural and ILU preconditioners showing advantage over AMG. As a result, the following relation between the total time and the size of the problem holds, $\log T = c_i + (p_i + 1) \log N_h$, showing that methods such as neural preconditioners and ILU provide a potential advantage for systems of small size.

We finally observe that the neural preconditioner struggles to address large systems, for example, the one arising from a discretization of the 3D domain with 101 points per side, corresponding to $N_h \simeq 10^6$, for which $\mathcal{U}_3^{hf}(21^3)$ requires on average 188.65 iterations to solve the 3D problem. The behavior does not improve with increasing the dimension of the training set, as this is true for both $\mathcal{U}_3^{hf}(21^3)$ and $\mathcal{U}_3^{hf}(41^3)$. We hypothesize that this is due to the effect of high frequencies that become dominant components when the discrete solution belongs to a high-dimensional space. For this reason, we tested the preconditioner $\mathcal{U}_3^{hf}(21^3)$ with pre and post-smoothing, as described in Section 5.4, with very encouraging results given by an average number of iterations equal to [13.12; 27.60; 62.32; 91.55] for $N_h = [21^3; 41^3; 81^3; 101^3]$, respectively. This corresponds to the parameters $p_i = 0.31$ for the dependence of the number of iterations on N_h and $c_i = 2.30210^{-4}$ for the proportionality constant between time and iterations. This performance is in line with the ILU preconditioner, actually slightly better.

5.7. Effect of the FEM polynomial order

In order to assess the influence of the polynomial order of the finite element discretization, we have extended our numerical tests to the case of second-order elements. This analysis explores the behavior of both classical and neural preconditioners in higher-order settings. Moving from \mathbb{P}^1 to \mathbb{P}^2 elements causes the system matrices to have a denser sparsity pattern, which increases both setup and application times. As a result, AMG(10) is expected to be significantly slower in terms of wall-clock time, although the number of iterations remains largely unchanged, as demonstrated by numerical experiments (details provided in Table 9). In this context, we perform a comparative analysis between the neural preconditioner and AMG(10). To this end, we regenerated the training data with \mathbb{P}^2 discretizations and trained a dedicated neural preconditioner. We then compared two scenarios: (i) applying a model trained on \mathbb{P}^1 data directly to \mathbb{P}^2 residuals, and (ii) using a model trained specifically on \mathbb{P}^2 data. We denote the two preconditioners by $\mathcal{U}_3^{hf}(\mathbb{P}^1)$ and $\mathcal{U}_3^{hf}(\mathbb{P}^2)$, respectively. The results, summarized in Table 9, confirm that the neural preconditioner is largely insensitive to the polynomial degree in terms of iteration counts. In addition, the performance of the $\mathcal{U}_3^{hf}(\mathbb{P}^1)$ applied to the \mathbb{P}^2 problems is very close to that of the retrained $\mathcal{U}_3^{hf}(\mathbb{P}^2)$, with only minor differences.

Overall, these experiments demonstrate that the neural preconditioner adapts naturally to higher-order discretizations, with stable iteration counts and low per-iteration cost. In contrast, AMG(10), although optimal in iteration numbers, becomes increasingly penalized in terms of computational time as the polynomial order grows. This identifies a practical regime in which the neural preconditioner represents an attractive alternative for higher-order FEM discretizations.

Table 9

Mean FGMRES iterations for different preconditioning strategies with $k_\Omega = \sigma_\Omega = 10^{-3}$. (*) Denotes that, at least once, FGMRES did not converge in the maximum number of iterations (1000); two polynomial orders are considered: \mathbb{P}^1 and \mathbb{P}^2 .

Precond.	N_h	\mathbb{P}^1			\mathbb{P}^2		
		N	Mean Iter.	Time/Iter	N	Mean Iter.	Time/Iter
None	9,261	21	147.72	–	11	173.32	–
	68,921	41	431.91	–	21	440.98	–
	531,441	81	1000*	–	41	1000*	–
ILU	9,261	21	24.73	1.52	11	37.68	1.11
	68,921	41	46.2	10.52	21	77.66	11.50
	531,441	81	102.55	91.88	41	196.01	105.75
AMG(3)	9,261	21	9.56	6.52	11	11.73	6.77
	68,921	41	17.13	40.64	21	21.34	52.26
	531,441	81	33.66	286.13	41	42.11	328.15
AMG(10)	9,261	21	3.64	15.29	11	4.37	19.95
	68,921	41	4.01	82.12	21	4.18	203.67
	531,441	81	4.9	800.30	41	5.84	1405.25
$\mathcal{U}_3^{hf}(\mathbb{P}^1)$	9,261	21	17.38	2.54	11	23.15	1.64
	68,921	41	37.24	13.37	21	37.29	10.87
	531,441	81	100.73	91.46	41	89.17	113.21
$\mathcal{U}_3^{hf}(\mathbb{P}^2)$	9,261	21			11	23.26	1.63
	68,921	41	–		21	35.35	12.02
	531,441	81			41	102.6	101.01

5.8. Coupling the neural preconditioner with ensembled calculations

In the context of parametric simulations of mixed-dimensional PDEs, solving multiple instances of the underlying linear system for varying configurations of the low-dimensional structure can become computationally prohibitive. To mitigate this, the so-called *ensembled* framework has recently been developed to accelerate the solution of families of linear systems that share a common structure [39,40]. Ensembled methods consist of stacking multiple coefficient matrices and right-hand side vectors, each corresponding to a different parameter instance, and solving them simultaneously as a batched linear system. This paradigm is especially well-suited for settings such as uncertainty quantification or parametric sweeps, where many problems must be solved independently but share structural similarities. The main advantage of this approach lies in the optimized memory access.

We note that a key feature of the proposed neural preconditioner lies in its ability to generalize across the parameter space and that this property is well suited to ensemble methods, as the neural operator can process a group of inputs at the same time, each representing various parameter settings.

We have implemented a prototype of the *ensembled* calculation using the algorithmic infrastructure developed for the neural preconditioner, based on the PyTorch library ([41]) together with the CUDA kernel ([42]). In fact, we recall that the input of the neural preconditioner consists of tensors $\mathbf{X}_i \in \mathbb{R}^{c \times n}$ with c representing the number of channels and $\mathbf{n} = n_1 \times n_2 \times \dots \times n_d$ a multi-index encoding the physical dimension of the data (see Section 4.2). Given a set of tensors $\{\mathbf{X}_i\}_{i=1,N}$, it is possible to stack them along an additional *batch dimension* to obtain a new tensor $\mathbf{X}_e = [\mathbf{X}_1 | \dots | \mathbf{X}_N] \in \mathbb{R}^{N \times c \times n}$ that encodes all input tensors to be fed to the neural preconditioner.

We evaluated the speed-up factor achieved by implementing the neural preconditioner \mathcal{U}_3^{hf} on $N \in [1, 100]$ input tensors $\mathbf{X}_i \in \mathbb{R}^{2 \times 21 \times 21 \times 21}$, which correspond to N different 1D graphs in Eq. (9), i.e.

$$s = \frac{\sum_{i=1}^N t(\mathcal{U}_3^{hf}(\mathbf{X}_i))}{t(\mathcal{U}_3^{hf}(\mathbf{X}_e))}$$

with $t(\mathcal{U}_3^{hf}(\cdot))$ representing the execution time of the neural preconditioner pass. The numerical tests are performed on an AMD EPYC 7301 CPU and a NVIDIA Tesla V100 GPU with 64 MB of dedicated RAM.

From the results shown in Table 10 and Fig. 7, we observe a speed-up around 4.5 when up to 100 input tensors are considered. Saturation occurs as the batched input approaches the memory limit of the hardware. The beneficial acceleration is associated with the contiguous memory structure, coupled with CUDA streams and execution via the PyTorch library. We prospect that the synergy between the ensembled approach and the neural preconditioner, along with the use of GPU based architectures, can enable significant gains in solver throughput, especially in multi-query scenarios. A comprehensive study encompassing computational architecture, algorithmic design, memory management, and preconditioner implementation is deferred to future research.




Fig. 7. Left: plot of the speed-up associated to the batched execution of the neural preconditioning pass, using $\mathcal{U}_3^{hf}(\mathbf{X}_e)$ with $\mathbf{X}_e \in \mathbb{R}^{N \times 2 \times 21 \times 21 \times 21}$, and batch size $N \in [1, 100]$. Right: Plot of the execution time, expressed in millisecond, associated respectively to the serial (non-batched) and batched preconditioner pass of \mathcal{U}_3^{hf} .

Table 10

Execution times [ms] and speed-up of the forward pass through the neural preconditioner as a function of the batch size.

Batch size (N)	1	5	10	15	20	30	50	70	90	100
time non-batched [ms]	1.5	6.1	11.8	17.5	23.2	34.7	58.2	81.3	104.7	114.9
time batched [ms]	1.5	2.6	4.1	5.5	6.6	9.7	13.5	18.2	22.9	25.5
Speed-up	1.0	2.3	2.9	3.5	3.7	3.6	4.3	4.5	4.6	4.5

6. Solution of the 3D-1D coupled problem

We now focus on the fully coupled 3D-1D problem, assessing how well the suggested block-preconditioning strategy (7)-(8) performs when applied to the linear system (6). We use a right-preconditioned FGMRES solver with a restart parameter $k = 20$ and a stopping criterion $\|r\|/\|r_0\| \leq 1 \cdot 10^{-15}$. This iterative solver is preconditioned by an approximation \tilde{Q}_h^μ of the block preconditioner defined in (8). More precisely, the block preconditioner \tilde{Q}_h^μ is defined by an approximate solution to the 1D problem, obtained with ILU factorization, returning $\tilde{z}_{h,1}^\mu$ (a direct solver has also been tested, leading to no remarkable differences); then, 3 steps of the preconditioned FGMRES algorithm (precisely P-FGMRES with various preconditioners P) return $\tilde{z}_{h,0}^\mu$, an approximate solution of the 3D system.

The performance of this approach for the 3D-1D coupled problem (6) with a comparison of monolithic solvers and different choices of preconditioners for the 3D problem can be found in Table 11. The latter presents a comparative analysis of the different strategies employed to solve the coupled 3D-1D system using the FGMRES iterative solver, with a focus on the impact of preconditioning.

As seen in the table, the results clearly highlight the importance of preconditioning for the efficient resolution of mixed-dimensional PDEs. Indeed, the first key observation is that the absence of a preconditioner results in a solver failing to converge within the budget of 2000 iterations. This emphasizes the severity of the ill-conditioning that characterizes the 3D-1D coupled problem and establishes the necessity of a robust preconditioning strategy. The block-diagonal preconditioner \tilde{Q}_h^μ emerges as a good option, underlining the effectiveness of a decoupled block-structured approach that aligns with the intrinsic structure of the mixed-dimensional system. Interestingly, the full monolithic preconditioning strategy based on algebraic multigrid with 10 smoothing steps (denoted AMG(10) monolithic) fails to deliver competitive performance. This suggests that generic monolithic approaches are suboptimal for problems with strong 3D-1D coupling, as they are not designed to exploit the specific features of the system. In contrast, the Haznics preconditioner, which involves a combination of AMG in the entire 3D domain and Schwarz smoother on the interface between the 3D and 1D domain [43], shows the best performance in terms of iterations, validating its role as the current state-of-the-art method for mixed-dimensional problems. A degradation of the mean iteration count can be observed for the highest value of N_h , especially due to lack of convergence when very complex graph geometries are considered.

The table also reports the results for strategies that combine the block-diagonal preconditioner \tilde{Q}_h^μ with various solvers applied to its diagonal blocks. In particular, the first block, corresponding to the 3D sub-problem, is preconditioned with ILU, AMG (3), AMG (10) or different instances of the neural preconditioner, namely $\mathcal{U}_3^{hf}(21^3)$ and $\mathcal{U}_3^{hf}(41^3)$. Although the iteration counts vary distinctively between the different strategies, we notice that their overall computational times are remarkably similar. This is a result of a trade-off between the cost per iteration and the number of iterations required. For instance, classical preconditioners like AMG generally result in lower iteration counts but at a higher per-iteration cost. Neural preconditioners, on the other hand, often yield higher iteration counts but with significantly lower cost per iteration, because of their efficient matrix-free representation. This balance ensures that all these strategies achieve similar total run-times, with performance comparable to or superior to monolithic alternatives.

These findings support the potential of the neural preconditioners as a viable and efficient strategy for solving complex mixed-dimensional problems of small-medium scale. The flexibility in trading iteration count for speed suggests that these learned preconditioners can serve as lightweight generalizable components in the broader preconditioning pipeline.

Table 11

Mean FGMRES iterations and execution times with associated rates for the preconditioning strategies considered, in different mesh configurations. Times are reported in seconds. $\text{Rate}(N_h) =$ for a given quantity of interest y , it is the estimated exponent q yielding $y \propto N_h^q$. (*) Denotes that, for at least a simulation in the test set, FGMRES did not converge in the maximum number of iterations, here set to 2000. (n.r.) means that the data was not recorded.

Precond.	N_h	$k_\Omega = \sigma_\Omega = 10^{-3}$				$k_\Omega = \sigma_\Omega = 10^{-2}$			
		Iter	Rate(N_h)	Time	Rate(N_h)	Iter	Rate(N_h)	Time	Rate(N_h)
Non prec. 3D-1D problem	9261	1973.55*				1768.65*			
	68,921	2000*		n.r.		2000*		n.r.	
	531,441	2000*				2000*			
AMG(10) monolithic	9261	1181.48*				1220.83*			
	68,921	1366.02*		n.r.		1423.44*		n.r.	
	531,441	1862.71*				1918.29*			
Haznics-AMG(10)3D-1D problem	9261	38.71	-	1.47	-	35.36	-	1.49	-
	68,921	48.04	0.11	2.14	0.19	41.39	0.08	2.26	0.21
	531,441	725.44*	1.35	84.232	1.80	497.57*	1.22	64.95	1.64
1D: ILU 3D: P-FGMRES P = none	9261	153.82	-	0.50	-	123.25	-	0.49	-
	68,921	417.82	0.50	8.51	1.42	316.88	0.47	4.48	1.10
	531,441	1564.84	0.65	260.41	1.32	916.19	0.52	152.88	1.73
1D: ILU 3D: P-FGMRES P = $\mathcal{U}_3^{hf}(21^3)$	9261	60.74	-	0.47	-	22.28	-	0.27	-
	68,921	123.30	0.35	3.12	0.94	54.33	0.44	1.54	0.87
	531441	438.69	0.62	93.49	1.66	157.69	0.52	32.98	1.50
1D: ILU 3D: P-FGMRES P = $\mathcal{U}_3^{hf}(41^3)$	9261	69.15	-	0.51	-	26.95	-	0.29	-
	68,921	104.80	0.21	2.68	0.83	36.55	0.15	1.06	0.64
	531441	257.91	0.44	54.62	1.48	59.99	0.23	12.78	1.22
1D: ILU 3D: P-FGMRES P = AMG(10)	9,261	50.19	-	0.72	-	16.85	-	0.35	-
	68,921	52.57	0.02	5.40	1.00	17.84	0.002	1.95	0.61
	531441	56.13	0.03	49.17	1.08	18.85	0.02	17.49	0.81
1D: ILU 3D: P-FGMRES P = AMG(3)	9,261	51.53	-	0.63	-	18.23	-	0.33	-
	68,921	56.61	0.05	4.64	0.99	23.00	0.12	2.01	0.90
	531441	68.64	0.08	45.58	1.12	37.1	0.23	24.75	1.23
1D: ILU 3D: P-FGMRES P = ILU	9261	52.14	-	0.4	-	18.72	-	0.26	-
	68,921	58.08	0.05	1.97	0.79	26.52	0.17	1.05	0.70
	531,441	76.33	0.13	20.32	1.14	47.83	0.29	13.25	1.24

7. Conclusions and future perspectives

In this work, we have developed a novel unsupervised neural network-based preconditioning strategy tailored to efficiently solve mixed-dimensional PDEs, specifically targeting the computational challenges arising from the coupling between 3D domains and embedded 1D structures. Using operator learning through convolutional neural networks, our approach successfully generalizes across various topological configurations of the 1D graphs and adapts robustly to different mesh resolutions without requiring re-training.

The neural preconditioner was extensively validated against classical preconditioning methods, including incomplete LU factorization and algebraic multigrid methods. Numerical experiments demonstrated that the proposed neural network preconditioner significantly accelerates convergence of iterative solvers, such as the Flexible GMRES, effectively handling both low- and high-frequency error components. Among the data augmentation strategies investigated, enriching the training set with random high-frequency vectors emerged as the most effective, enabling the neural network to overcome the inherent spectral bias common to most machine learning-based approaches.

Furthermore, our analysis underscored the critical role of including explicit parametric information through the distance function associated with the 1D graph, which improved the effectiveness of the preconditioner in varying geometric configurations. We also explored the complementary effects of pre- and post-smoothing strategies, illustrating their capacity to enhance preconditioning performance, particularly when the training set lacks high-frequency data.

However, there are limitations and areas for future improvement. A significant current limitation is that the proposed approach is applicable only to tensor-product domains, which restricts its use in realistic applications. To address this, we are actively exploring the use of graph- and mesh-informed neural networks [33,44], as suitable alternatives to convolutional layers, with the aim of extending the applicability of this method to more general computational domains. Additionally, developing a solid theoretical framework for selecting data augmentation subsets and, more broadly, for learning-based preconditioning methods is crucial to consolidate this approach and broaden the range of application areas. Although the scalability of the neural preconditioner shows promise, more work is required to enhance its performance in large-scale problems.

The promising results obtained from this study not only establish neural preconditioning as a viable and efficient alternative to traditional preconditioners for mixed-dimensional PDEs, but also lay a robust foundation for further exploration into broader classes of coupled multiphysics problems.

Data availability

Data will be made available on request.

Acknowledgments

The authors thank Prof. Kent-Andre Mardal and Dr. Miroslav Kuchta of the Simula Research Laboratory for useful discussions. PZ acknowledges the support of the MUR PRIN 2022 grant No. 2022WKWZA8 *Immersed methods for multiscale and multiphysics problems* (IMMEDIATE) part of the Next Generation EU program Mission 4, Comp. 2, CUP D53D23006010006. NF has received funding under the project *Reduced Order Modeling and Deep Learning for the real-time approximation of PDEs (DREAM)*, grant no. FIS00003154, funded by the Italian Science Fund (FIS) and by Ministero dell’Università e della Ricerca (MUR). The present research is part of the activities of the *Dipartimento di Eccellenza 2023–2027*, Department of Mathematics, Politecnico di Milano. All authors are members of the Gruppo Nazionale per il Calcolo Scientifico (GNCS), Istituto Nazionale di Alta Matematica (INdAM).

Appendix A. Auxiliary results

Lemma 2. *Let $(V, \|\cdot\|_V)$ and $(W, \|\cdot\|_W)$ be two Banach spaces. Let $\mathcal{GL}(V, W) \subset \mathcal{L}(V, W)$ be the set of invertible bounded linear operators from V to W , equipped with the operator norm*

$$\|A\|_{\mathcal{L}(V,W)} := \sup_{v \in V \setminus \{0\}} \frac{\|Av\|_W}{\|v\|_V}.$$

Then,

1. $\mathcal{GL}(V, W)$ is open in the operator norm topology;
2. for every $A \in \mathcal{GL}(V, W)$ there exists a unique inverse $A^{-1} \in \mathcal{L}(W, V)$. Furthermore, $A^{-1} \in \mathcal{GL}(W, V)$;
3. the inversion map $\mathcal{I} : A \mapsto A^{-1}$ is $\mathcal{GL}(V, W) \mapsto \mathcal{GL}(W, V)$ continuous.

Proof. We mention that a stronger result holds, as \mathcal{I} is in fact holomorphic: see, e.g., the proof in [45, Proposition 4.8]. Nonetheless, we can easily prove the statements in the Lemma as follows.

1. Given $A \in \mathcal{GL}(V, W)$, let $B \in \mathcal{L}(V, W)$ be such that $\|A - B\|_{\mathcal{L}(V,W)} < 1/\|A^{-1}\|_{\mathcal{L}(W,V)}$. We notice that the operator $C = A^{-1}(A - B) = \text{Id}_V - A^{-1}B$ is an endomorphism from V onto itself. Furthermore,

$$\|C\|_{\mathcal{L}(V,V)} \leq \|A^{-1}\|_{\mathcal{L}(W,V)} \cdot \|A - B\|_{\mathcal{L}(V,W)} < 1.$$

It follows that, for every k , the k th power of C , defined via composition, satisfies $\|C^k\|_{\mathcal{L}(V,V)} \leq \|C\|_{\mathcal{L}(V,V)}^k < 1$. As a consequence, the Nuemann series $\sum_{k=0}^{+\infty} C^k$ is shown to be absolutely convergent in $\mathcal{L}(V, V)$. Indeed,

$$\sum_{k=0}^{+\infty} \|C^k\|_{\mathcal{L}(V,V)} \leq \sum_{k=0}^{+\infty} \|C\|_{\mathcal{L}(V,V)}^k = \frac{1}{1 - \|C\|_{\mathcal{L}(V,V)}}.$$

Since $\mathcal{L}(V, V)$ is a Banach space, this ensures that $\sum_{k=0}^{+\infty} C^k$ converges to some element of $\mathcal{L}(V, V)$. In particular, by construction of the Neumann series, $\sum_{k=0}^{+\infty} C^k = (\text{Id}_V - C)^{-1} = (A^{-1}B)^{-1}$. Let

$$B^\dagger := (A^{-1}B)^{-1}A^{-1}. \tag{A.1}$$

Clearly, $B^\dagger \in \mathcal{L}(W, V)$ by composition. Furthermore, we have $A^{-1}BB^\dagger = A^{-1}$. Since A^{-1} is invertible, and thus of full rank, this shows that $BB^\dagger = \text{Id}_W$. Similarly, composing (A.1) with B yields

$$B^\dagger B = (A^{-1}B)(A^{-1}B) = \text{Id}_V.$$

It follows that $B \in \mathcal{GL}(V, W)$ and $B^\dagger = B^{-1}$.

2. This follows immediately from the fact that invertible operators have full rank. To see this, let $A \in \mathcal{GL}(V, W)$ and assume $B_1, B_2 \in \mathcal{L}(W, V)$ are two -possibly different- inverses of A . For all $w \in W$ we have

$$B_1w = B_2A(B_1w) = B_2(AB_1w) = B_2w$$

since $B_1w \in V$ and $AB_1w \in W$. Then $B_1 \equiv B_2 =: A^{-1}$. The fact that $A^{-1} \in \mathcal{GL}(W, V)$ is obvious.

3. Let $A_n \rightarrow A$ in $\mathcal{GL}(V, W)$. Define the error $E_n = A - A_n$. There exists some n_0 such that $\|E_n\|_{\mathcal{L}(V,W)} < 1/\|A\|_{\mathcal{L}(V,W)}$ for all $n \geq n_0$. As for our proof of (1), we notice that this ensures the convergence of the following Neumann series

$$\sum_{k=0}^{+\infty} (A^{-1}E_n)^k = (\text{Id}_V - A^{-1}E_n)^{-1}.$$

whenever $n \geq n_0$. Since

$$\|A^{-1} - A_n^{-1}\|_{\mathcal{L}(W,V)} \leq \|\text{Id}_V - A_n^{-1}A\|_{\mathcal{L}(V,V)} \cdot \|A^{-1}\|_{\mathcal{L}(W,V)},$$

it suffices to prove that $\|\text{Id}_V - A_n^{-1}A\|_{\mathcal{L}(V,V)} \rightarrow 0$ as $A_n \rightarrow A$. To this end, we note that for all $n \geq n_0$

$$\begin{aligned} \|\text{Id}_V - A_n^{-1}A\|_{\mathcal{L}(V,V)} &= \|\text{Id}_V - (A - E_n)^{-1}A\|_{\mathcal{L}(V,V)} \\ &= \|\text{Id}_V - (\text{Id}_V - A^{-1}E_n)^{-1}A^{-1}A\|_{\mathcal{L}(V,V)} \\ &= \|\text{Id}_V - (\text{Id}_V - A^{-1}E_n)^{-1}\|_{\mathcal{L}(V,V)}. \end{aligned}$$

In particular, $\|\text{Id}_V - A_n^{-1}A\|_{\mathcal{L}(V,V)} = \left\| \sum_{k=1}^{+\infty} (A^{-1}E_n)^k \right\|_{\mathcal{L}(V,V)}$, and thus

$$\|\text{Id}_V - A_n^{-1}A\|_{\mathcal{L}(V,V)} \leq \|A^{-1}\|_{\mathcal{L}(W,V)} \sum_{k=1}^{+\infty} \|E_n\|_{\mathcal{L}(V,W)}^k = \frac{\|A^{-1}\|_{\mathcal{L}(W,V)} \|E_n\|_{\mathcal{L}(V,W)}}{1 - \|E_n\|_{\mathcal{L}(V,W)}},$$

which is vanishing for $n \rightarrow +\infty$, as $E_n \rightarrow 0$.

□

References

- [1] C. Meier, M.J. Grill, W.A. Wall, A. Popp, Geometrically exact beam elements and smooth contact schemes for the modeling of fiber-based materials and structures, *Int. J. Solids Struct.* 154 (2018) 124–146.
- [2] F. Lespagnol, C. Grandmont, P. Zunino, M.A. Fernandez, A mixed-dimensional formulation for the simulation of slender structures immersed in an incompressible flow, *Comput. Methods Appl. Mech. Eng.* 432 (2024) 117316.
- [3] A. Fumagalli, A. Scotti, A numerical method for two-phase flow in fractured porous media with non-matching grids, *Adv. Water Resour.* 62 (2013) 454–464.
- [4] E. Vidotto, T. Koch, T. Koeppl, R. Helmig, B. Wohlmuth, Hybrid models for simulating blood flow in microvascular networks, *Multiscale Model. Simul.* 17 (3) (2019) 1076–1102.
- [5] L. Possenti, A. Cicchetti, R. Rosati, D. Cerroni, M.L. Costantino, T. Rancati, P. Zunino, A mesoscale computational model for microvascular oxygen transfer, *Ann. Biomed. Eng.* 29 (12) (2021) 3356–3373.
- [6] A.P. Buccino, M. Kuchta, J. Schreiner, K.-A. Mardal, Improving neural simulations with the EMI model, in: A. Tveito, K.-A. Mardal, M.E. Rognes (Eds.), *Modeling Excitable Tissue: The EMI Framework*, Springer International Publishing, Cham, 2021, pp. 87–98.
- [7] N. Hagemeyer, M. Mayr, I. Steinbrecher, A. Popp, One-way coupled fluid–beam interaction: capturing the effect of embedded slender bodies on global fluid flow and vice versa, *Adv. Model. Simul. Eng. Sci.* 9 (1) (2022).
- [8] W.M. Boon, J.M. Nordbotten, J.E. Vagne, Functional analysis and exterior calculus on mixed-dimensional geometries, *Ann. Mat. Pura Appl.* 200 (2) (2021) 757–789.
- [9] M. Kuchta, F. Laurino, K.-A. Mardal, P. Zunino, Analysis and approximation of mixed-dimensional PDEs on 3D-1D domains coupled with lagrange multipliers, *SIAM J. Numer. Anal.* 59 (1) (2021) 558–582.
- [10] K.-A. Mardal, R. Winther, Preconditioning discretizations of systems of partial differential equations, *Numer. Linear Algebra Appl.* 18 (1) (2011) 1–40.
- [11] M. Kuchta, M. Nordaas, J.C.G. Verschaeve, M. Mortensen, K.A. Mardal, Preconditioners for saddle point systems with trace constraints coupling 2d and 1d domains, *SIAM J. Scient. Comput.* 38 (6) (2016) B962–B987.
- [12] M. Kuchta, K.-A. Mardal, M. Mortensen, Preconditioning trace coupled 3D-1D systems using fractional Laplacian, *Numer. Methods Partial Differ. Equ.* 35 (1) (2019) 375–393. <https://doi.org/https://doi.org/10.1002/num.22304>
- [13] A. Budiša, X. Hu, M. Kuchta, K.-A. Mardal, L. Zikatanov, Algebraic multigrid methods for metric-perturbed coupled problems, *SIAM J. Scient. Comput.* 46 (3) (2024) A1461–A1486.
- [14] M. Firmbach, I. Steinbrecher, A. Popp, M. Mayr, An approximate block factorization preconditioner for mixed-dimensional beam-solid interaction, *Comput. Methods Appl. Mech. Eng.* 431 (2024) 117256.
- [15] N.B. Kovachki, S. Lanthaler, A.M. Stuart, Chapter 9 - Operator learning: algorithms and analysis, in: S. Mishra, A. Townsend (Eds.), *Numerical Analysis Meets Machine Learning*, 25 of *Handbook of Numerical Analysis*, Elsevier, 2024, pp. 419–467. <https://doi.org/https://doi.org/10.1016/bs.hna.2024.05.009>
- [16] N. Boullá@, A. Townsend, Chapter 3 - A mathematical guide to operator learning, in: S. Mishra, A. Townsend (Eds.), *Numerical Analysis Meets Machine Learning*, 25 of *Handbook of Numerical Analysis*, Elsevier, 2024, pp. 83–125. <https://doi.org/https://doi.org/10.1016/bs.hna.2024.05.003>
- [17] L. Lu, P. Jin, G. Pang, Z. Zhang, G.E. Karniadakis, Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators, *Nature Mach. Intell.* 3 (3) (2021) 218–229.
- [18] M. Geist, P. Petersen, M. Raslan, R. Schneider, G. Kutyniok, Numerical solution of the parametric diffusion equation by deep neural networks, *J. Sci. Comput.* 88 (1) (2021) 22.
- [19] I.G. Gjerde, K. Kumar, J.M. Nordbotten, B. Wohlmuth Splitting method for elliptic equations with line sources, *ESAIM: M2AN* 53 (5) (2019) 1715–1739.
- [20] A. Kopaničáková, G.E. Karniadakis, DeepONet based preconditioning strategies for solving parametric linear systems of equations, *SIAM J. Scient. Comput.* 47 (1) (2025) C151–C181. <https://doi.org/10.1137/24M162861X>
- [21] V. Grimm, A. Heinlein, A. Klawonn, Learning the solution operator of two-dimensional incompressible Navier-Stokes equations using physics-aware convolutional neural networks, *arXiv:2308.02137 math.NA* (2023).
- [22] N. Margenberg, D. Hartmann, C. Lessig, T. Richter, A neural network multigrid solver for the Navier-Stokes equations, *J. Comput. Phys.* 460 (2022) 110983.
- [23] Y. Azulay, E. Treister, Multigrid-augmented deep learning preconditioners for the Helmholtz equation, *SIAM J. Scient. Comput.* 45 (3) (2023) S127–S151.
- [24] N. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. Stuart, A. Anandkumar, Neural operator: learning maps between function spaces with applications to PDEs, *J. Mach. Learn. Res.* 24 (2023).
- [25] F. Laurino, P. Zunino, Derivation and analysis of coupled PDEs on manifolds with high dimensionality gap arising from topological model reduction, *ESAIM Math. Modell. Numer. Anal.* 53 (6) (2019) 2047–2080.
- [26] Y. Saad, *Iterative Methods for Sparse Linear Systems*, Applied Mathematics, Society for Industrial and Applied Mathematics, 2003.
- [27] J. Dugundji, An extension of Tietze's theorem, *Pac. J. Math.* 1 (3) (1951) 353–367.
- [28] Y. Saad, A flexible inner-outer preconditioned GMRES algorithm, *SIAM J. Scient. Comput.* 14 (2) (1993) 461–469.

- [29] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F.A. Hamprecht, Y. Bengio, A. Courville, On the Spectral Bias of Neural Networks, arXiv: 1806.08734 stat.ML (2019).
- [30] Z.-Q.J. Xu, Y. Zhang, T. Luo, Overview frequency principle/spectral bias in deep learning, *Commun. Appl. Math. Comput.* (2024).
- [31] O. Ronneberger, P. Fischer, T. Brox, U-Net: convolutional networks for biomedical image segmentation, in: *Medical Image Computing and Computer-assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, Springer, 2015, pp. 234–241.
- [32] C. Williams, F. Falck, G. Deligiannidis, C. Holmes, A. Doucet, S. Syed, A unified framework for U-Net design and analysis, arXiv: 2305.19638 stat.ML (2024).
- [33] N.R. Franco, A. Manzoni, P. Zunino, Mesh-informed neural networks for operator learning in finite element spaces, *J. Sci. Comput.* 97 (35) (2023).
- [34] S. Balay, S. Abhyankar, M.F. Adams, S. Benson, J. Brown, P. Brune, K. Buschelman, E.M. Constantinescu, L. Dalcin, A. Dener, et al., *PETSc/TAO Users Manual (Rev. 3.20)*, Technical Report, Argonne National Laboratory (ANL), Argonne, IL (United States), 2023.
- [35] L.D. Dalcin, R.R. Paz, P.A. Kler, A. Cosimo, Parallel distributed computing using python, *Adv. Water Resour.* 34 (9) (2011) 1124–1139.
- [36] N. Dimola, M. Kuchta, K.-A. Mardal, P. Zunino, Robust preconditioning of mixed-dimensional PDEs on 3d-1d domains coupled with lagrange multipliers, *SEMA SIMAI Springer Series* 36 (2024) 137 – 171.
- [37] U. Trottenberg, C.W. Oosterlee, A. Schuller, *Multigrid*, Elsevier, 2000.
- [38] S. Mittal, A survey on optimized implementation of deep learning models on the nvidia jetson platform, *J. Syst. Archit.* 97 (2019) 428–442.
- [39] E.T. Phipps, M. D’Elia, H.C. Edwards, M. Hoemmen, J. Hu, S. Rajamanickam, Embedded ensemble propagation for improving performance, portability, and scalability of uncertainty quantification on emerging computational architectures, *SIAM J. Scient. Comput.* 39 (2) (2017) C162–C193. <https://doi.org/10.1137/15M1044679>
- [40] K. Liegeois, R. Boman, E.T. Phipps, T.A. Wiesner, M. Arnst, GMRES with embedded ensemble propagation for the efficient solution of parametric linear systems in uncertainty quantification of computational models, *Comput. Methods Appl. Mech. Eng.* 369 (2020) 113188. <https://doi.org/https://doi.org/10.1016/j.cma.2020.113188>
- [41] A. Paszke, Pytorch: An imperative style, high-performance deep learning library, arXiv preprint arXiv:1912.01703 (2019).
- [42] D. Luebke, CUDA: Scalable parallel programming for high-performance scientific computing, in: *2008 5th IEEE International Symposium on Biomedical Imaging: From Nano to Macro, IEEE, 2008*, pp. 836–838.
- [43] A. Budiša, X. Hu, M. Kuchta, K.-A. Mardal, L.T. Zikatanov, HAZniCS – Software Components for multiphysics problems, *ACM Trans. Math. Softw.* 49 (4) (2023). <https://doi.org/10.1145/3625561>
- [44] P. Battaglia, J.B.C. Hamrick, V. Bapst, A. Sanchez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, C. Gulcehre, F. Song, A. Ballard, J. Gilmer, G.E. Dahl, A. Vaswani, K. Allen, C. Nash, V.J. Langston, C. Dyer, N. Heess, D. Wierstra, P. Kohli, M. Botvinick, O. Vinyals, Y. Li, R. Pascanu, Relational inductive biases, deep learning, and graph networks, arXiv (2018).
- [45] B. Adcock, S. Brugiapaglia, C.G. Webster, *Sparse Polynomial Approximation of High-Dimensional Functions*, SIAM, Philadelphia, PA, Philadelphia, PA, 2022.