ESAIM: COCV 32 (2026) 18
https://doi.org/10.1051/cocv/2026002

ESAIM: Control, Optimisation and Calculus of Variations
www.esaim-cocv.org

A HYBRID PHYSICS-INFORMED NEURAL NETWORK BASED MULTISCALE SOLVER AS A PARTIAL DIFFERENTIAL EQUATION CONSTRAINED OPTIMIZATION PROBLEM

Michael Hintermüller[1,2] and Denis Korolev[1,*]

Abstract. In this work, we study physics-informed neural networks (PINNs) constrained by partial differential equations (PDEs) and their application in approximating PDEs with two characteristic scales. From a continuous perspective, our formulation corresponds to a non-standard PDE-constrained optimization problem with a PINN-type objective. From a discrete standpoint, the formulation represents a hybrid numerical solver that utilizes both neural networks and finite elements. For the problem analysis, we introduce a proper function space, and we develop a numerical solution algorithm. The latter combines an adjoint-based technique for the efficient gradient computation with automatic differentiation. This new multiscale method is then applied exemplarily to a heat transfer problem with oscillating coefficients. In this context, the neural network approximates a fine-scale problem, and a coarse-scale problem constrains the associated learning process. We demonstrate that incorporating coarse-scale information into the neural network training process *via* a weak convergence-based regularization term is beneficial. Indeed, while preserving upscaling consistency, this term encourages non-trivial PINN solutions and also acts as a preconditioner for the low-frequency component of the fine-scale PDE, resulting in improved convergence properties of the PINN method. The relevance of our approach to material science is discussed.

Mathematics Subject Classification. 35B27, 68Q32, 68T05, 65K10.

Received November 14, 2024. Accepted January 4, 2026.

1. Introduction

Solving partial differential equations (PDEs) using physics-informed neural networks (PINNs) is currently an active area of research; see, *e.g.* [1] for an overview and references therein. The main principle of physics-informed learning was pioneered by [2] and later reincarnated in its modern computational interpretation by Raissi *et al.* [3]. It consists of integrating physical laws, typically in the form of the residuals of underlying PDEs, into a least-squares objective and finding the approximate solution to the corresponding residual minimization problem. PINN methods are typically meshless, making them potentially useful for numerically solving PDEs on complex domains or in high dimensions [4, 5]. Moreover, (approximate) optimization can be performed rapidly on modern GPU clusters, thanks to the excellent parallelization capabilities of neural networks on

Keywords and phrases: Learning-informed optimal control, PDE constrained optimization, physics-informed neural networks, quasi-minimization, weak convergence, multi-fidelity.

[1] Weierstrass Institute for Applied Analysis and Stochastics (WIAS), Anton-Wilhelm-Amo-Str. 39, 10117 Berlin, Germany.
[2] Institute for Mathematics, Humboldt-Universität zu Berlin, Unter den Linden 6, 10099 Berlin, Germany.

* Corresponding author: korolev@wias-berlin.de

GPUs, advances in automatic differentiation and domain decomposition techniques [6–8]. This progress thereby opens up opportunities to address problems of significant computational complexity, such as those arising in multiscale modelling.

In order to be specific, let $u^\varepsilon \in U$ be a solution in a suitable Sobolev space U to the following PDE with multiscale features:

$$\mathcal{A}^\varepsilon u = f, \quad \text{in} \quad L^2(\Omega), \quad \mathcal{B}u = g, \quad \text{in} \quad L^2(\partial\Omega), \tag{1.1}$$

where ε is a small positive parameter representing the characteristic length scale of fine-scale variations. The operator $\mathcal{A}^\varepsilon : U \to L^2(\Omega)$ is a partial differential operator, and the operator $\mathcal{B} : U \to L^2(\partial\Omega)$ reflects boundary conditions. The functions f and g represent given source and boundary data, respectively. Following [3], we approximate (1.1) with a neural network ansatz $u_{\boldsymbol{\theta},n} \in \mathfrak{N}_{\boldsymbol{\theta},n}$, where $\mathfrak{N}_{\boldsymbol{\theta},n}$ is a (nonlinear) network class parameterized by $\boldsymbol{\theta} \in \mathbb{R}^{N_n}$. The unknown network parameters (the so-called weights and biases) are then determined by solving the following optimization problem:

$$\inf \ \mathcal{J}(u_{\boldsymbol{\theta},n}) = \|\mathcal{A}^\varepsilon u_{\boldsymbol{\theta},n} - f\|^2_{L^2(\Omega)} + \tau_1 \|\mathcal{B}u_{\boldsymbol{\theta},n} - g\|^2_{L^2(\partial\Omega)} \quad \text{over} \quad u_{\boldsymbol{\theta},n} \in \mathfrak{N}_{\boldsymbol{\theta},n}, \tag{1.2}$$

where $\tau_1 > 0$ is a fixed weight. However, when considering such minimization from the perspective of ansatz-generating parameters, its non-convex nature, combined with the complex nonlinear dynamics of the learning process, renders both analytical and numerical treatment of the problem delicate [9–11], with the latter often prone to instabilities and slow convergence.

Motivated by multiscale and multi-fidelity approaches, we stabilize the neural network optimization (1.2) by adding, as a constraint, a numerically well-tractable coarse-scale model of (1.1), learning-informed by the ansatz $u_{\boldsymbol{\theta},n}$, with the respective solution $y(u_{\boldsymbol{\theta},n})$. This gives rise to the following PDE-constrained optimization problem:

$$\begin{cases} \inf \ J(y(u_{\boldsymbol{\theta},n}), u_{\boldsymbol{\theta},n}) := \mathcal{J}(u_{\boldsymbol{\theta},n}) + \tau_2 \mathcal{R}(y(u_{\boldsymbol{\theta},n}), u_{\boldsymbol{\theta},n}) \quad \text{over } (y(u_{\boldsymbol{\theta},n}), u_{\boldsymbol{\theta},n}) \in Y \times \mathfrak{N}_{\boldsymbol{\theta},n}, \\ \text{subject to (s.t.) } \mathcal{L}\big[u_{\boldsymbol{\theta},n}\big] y(u_{\boldsymbol{\theta},n}) = f, \end{cases} \tag{1.3}$$

where $\tau_2 > 0$, and the coupling term $\mathcal{R} : Y \times U \to \mathbb{R}_{\geq 0}$, which stabilizes the training of (1.2), is given by

$$\mathcal{R}(y(u_{\boldsymbol{\theta},n}), u_{\boldsymbol{\theta},n}) = \|\bar{Q}_\delta u_{\boldsymbol{\theta},n} - y(u_{\boldsymbol{\theta},n})\|^2_{L^2(\Omega)}. \tag{1.4}$$

By $\mathcal{L}[u_{\boldsymbol{\theta},n}] : Y \to Z$ we denote a coarse-scale differential operator between Banach spaces Y and Z, which is informed by the neural network ansatz $u_{\boldsymbol{\theta},n}$ *via* a suitable learning-informed upscaling procedure. Together with some given data f, it defines an equality constraint in (1.3). The coupling term $\mathcal{R}$ includes a *compression operator* $\bar{Q}_\delta : U \to L^2(\Omega)$: for stable training of (1.3), one wishes $\bar{Q}_\delta u^\varepsilon = y(u^\varepsilon)$, but some (rather imperfect) choices of $\bar{Q}_\delta$ and coarse-scale approximation $y(u_{\boldsymbol{\theta},n})$ must be made in practice, which may in turn destabilize (1.3). We propose a strategy to estimate these inconsistencies and approximation errors, enabling the construction of a stable approximation scheme and introducing the related *upscaling consistency* concept. In our case, (1.4) incorporates information on the weak convergence of the fine-scale solution of (1.1) to the homogenized solution, the latter serving as a natural reference for the coarse-scale problem as $\varepsilon \to 0$, with a detailed analysis provided.

Extensive work in multiscale modeling (*e.g.*, [12–16]) has led to various upscaling methods, including neural network-based approaches [17–21], for constructing coarse-scale surrogates for (1.1) and similar problems. Learning fine-scale surrogates informed by coarse-scale models and fine-scale data has also been studied, such as in [22]. In our application section, we adopt a numerical homogenization method from material science [23–25], which enables the transition from fine to coarse scales *via* the parametrization $\mathcal{L}\big[u_{\boldsymbol{\theta},n}\big]$. We apply it to the heat conduction problem in heterogeneous media, characterized by an oscillatory heat conductivity coefficient

$\boldsymbol{K}^\varepsilon \in C^{0,1}(\bar{\Omega}, \mathbb{R}_+)$. The following fine- and coarse-scale differential operators are considered:

$$\mathcal{A}^\varepsilon u = -\nabla \cdot (\boldsymbol{K}^\varepsilon \nabla u), \quad \mathcal{B}u = u|_{\partial\Omega}, \quad \langle \mathcal{L}\big[u\big] y(u), v\rangle_{Y^*,Y} = \int_\Omega \widetilde{\boldsymbol{K}}[u] \nabla y(u) \cdot \nabla v \, \mathrm{d}x, \tag{1.5}$$

where the restriction to the boundary $\partial\Omega$ is understood in the sense of traces. At the coarse scale, $\boldsymbol{K}^\varepsilon$ is replaced by an effective medium characterization given by the upscaled coefficient $\widetilde{\boldsymbol{K}} \in L^\infty(\bar{\Omega}, \mathbb{R}^{2\times 2})$. This coefficient is derived by inserting the domain-averaged flux $\langle \boldsymbol{K}^\varepsilon \nabla u^\varepsilon \rangle_\Omega$ and the gradient $\langle \nabla u^\varepsilon \rangle_\Omega$ into Fourier's law. Replacing $\boldsymbol{K}^\varepsilon$ with $\widetilde{\boldsymbol{K}}$ in $\mathcal{A}^\varepsilon$ and integrating by parts yields the coarse-scale operator $\mathcal{L}\big[u\big]$ amenable to finite element discretization. To enable two-way coupling of (1.5) within our learning scheme (1.3), we employ the coupling term (1.4) as a prior to stabilize the fine-scale learning of (1.1). Two-scale settings similar to (1.5) also arise in other contexts, such as porous media flow in complex geometries, where the Stokes equation serves as the fine-scale model and the Stokes–Brinkman equation as the coarse-scale one; see *e.g.* [26] and references therein. In such problems, inferring effective material properties from fine-scale solutions is required in applications, and learning-informed hybrid approaches can provide valuable support in this task. We refer to Section 4 for further details on this application setting.

Discretizing the coarse-scale equation in (1.3), *e.g.*, *via* the finite element method, alongside a PINN-based approach for the fine-scale problem, yields a *hybrid physics-informed multiscale* numerical solver. Note also that in the course of the optimization process for solving the hybrid finite dimensional approximate version of (1.3) possibly requires to frequently solve the discretized coarse-scale PDE. The hope now is that the coarse-scale equation can be solved numerically at a significantly lower cost (compute time) than computing the PINN solution, while still well informing low-frequency components of the fine-scale solution. Our numerical experiments provide evidence that incorporating the coarse-scale solution into the learning process through a coupling term in the objective of (1.3) acts as a preconditioner for the low-frequency component of (1.1), thereby accelerating convergence to such a component and improving the overall performance. We note here that PINNs can be difficult to train for problems exhibiting high-frequency or multiscale behavior [10], particularly due to the so-called "spectral bias" of neural networks. The latter is related to the fact that learning prioritizes low-frequency modes and prevents networks from effectively learning high-frequency functions [27]. To address this, Fourier features [10, 28] and architectures such as FBPINNs [7] have been proposed. FBPINNs, for example, use domain decomposition and local normalization to convert high-frequency content into lower frequencies. While often effective, these methods remain computationally expensive and may suffer from instabilities. Our approach is complementary: although the coarse-scale constraint does not directly address spectral bias, it helps accelerate and stabilize the training of multiscale PINNs when paired with a suitable parameterization $\mathcal{L}\big[u_{\boldsymbol{\theta},n}\big]$.

In the realm of learning-informed optimal control, several works, including [29, 30], have focused on approximating nonlinear constituents or source terms in the state equation using neural networks. PINNs have also been employed in solvers for underlying state and adjoint equations in various PDE-constrained optimization scenarios [31, 32]. Let us point out here that all the aforementioned techniques, while structurally perhaps similar (as $u_{\boldsymbol{\theta},n}$ may be considered a control variable, whereas $y(u_{\boldsymbol{\theta},n})$ acts as a state), differ from this work. Indeed, in the usual approaches the objective is typically not related to a neural network learning problem or PDE residual minimization, but rather to minimizing specific (*e.g.*, tracking-type) cost functionals [33, 34]. This difference has significant implications in analysis and numerical implementation. Besides, to the best of the authors' knowledge, our work is the first to deal in detail with a PINN-based optimization problem constrained by a PDE. Here, problem (1.3) is formulated and analyzed in a function space setting, taking into account the regularity of the fine-scale and coarse-scale PDE solutions as well as the interdependent choices of activation functions and PINN losses. For our new problem (1.3), the concept of quasi-minimization [35] is crucial when studying (approximate) existence of solutions over nonlinear neural network ansatz classes.

2. A hybrid multiscale approach

In this section, we define a function space and a coupling framework for two PDEs, the fine-scale and the coarse-scale problem, respectively. For the treatment of the associated PINNs, we closely follow [35]. For the bounded domain $\Omega \subset \mathbb{R}^d$ with Lipschitz boundary $\partial\Omega$, let $L^p(\Omega)$, $H^1(\Omega)$, $H^1_0(\Omega)$, $H^k(\Omega)$, $W^{k,p}(\Omega)$, *etc.* denote the standard Lebesgue and Sobolev spaces; see, *e.g.*, [36]. We also set $\mathbb{R}_+ := \{x \in \mathbb{R} : \ x > 0\}$ and $\mathbb{R}_{\geq 0} := \{x \in \mathbb{R} : \ x \geq 0\}$.

2.1. Function spaces and PDEs

Let $(U, \|\cdot\|_U)$, $(H, \|\cdot\|_H)$ be Hilbert spaces, $(X, \|\cdot\|_X)$ be a normed vector space, $(V, \|\cdot\|_V)$, $(Z, \|\cdot\|_Z)$ be Banach spaces, and X a dense subspace of U. Suppose also that U is continuously embedded in V and V is continuously embedded in H, *i.e.*, $U \hookrightarrow V \hookrightarrow H$. Moreover, $X \hookrightarrow U$. For given $f^\varepsilon \in H$, we consider the following partial differential equation

$$\mathcal{A}^\varepsilon u = f^\varepsilon, \quad \text{in } H, \quad \mathcal{B}u = 0, \quad \text{in } Z, \tag{2.1}$$

with $\mathcal{A}^\varepsilon : U \to H$ a bounded linear partial differential operator, *i.e.*, $\mathcal{A}^\varepsilon \in L(U, H)$, depending on the fine-scale length $\varepsilon > 0$, and $\mathcal{B} \in L(U, Z)$ defining boundary conditions. For (2.1) we invoke the following solution concept.

Assumption 2.1. For every $\varepsilon > 0$, (2.1) admits a unique solution $u^\varepsilon \in U$ for which there exists an approximating sequence $\{u^\varepsilon_k\}$ in X such that

$$\lim_{k\to\infty} \|u^\varepsilon_k - u^\varepsilon\|_U = 0, \quad \lim_{k\to\infty} \|\mathcal{A}^\varepsilon u^\varepsilon_k - f^\varepsilon\|_H + \|\mathcal{B}u^\varepsilon_k\|_Z = 0. \tag{2.2}$$

Remark 2.2. Let $g \in Z$ and assume that there exists $u^g \in U$ such that $\mathcal{B}u^g = g$ in Z. If $u^\varepsilon \in U$ satisfies (2.2) with $f^\varepsilon = q - \mathcal{A}^\varepsilon u^g$ for some given $q \in H$, then $w^\varepsilon := u^\varepsilon + u^g \in U$ corresponds to (2.1) with $\mathcal{A}^\varepsilon u = q$ in H and $\mathcal{B}u = g$ in Z.

Remark 2.3. In our setting, the operator in (2.1) is linear. While some work on PINNs for nonlinear PDEs exist (*e.g.*, [37, 38]) and can be adapted to our fine-scale problem, extending the two-scale coupling approach to nonlinear cases may require further investigation.

From now on we assume that Assumption 2.1 is satisfied. Note that if $u^\varepsilon \in U$ is the solution to (2.1), then $u^\varepsilon \in U_0 := \{u \in U : \ \mathcal{B}u = 0\} \subseteq U$, and we sometimes use $\|\cdot\|_{U_0}$ instead of $\|\cdot\|_U$. The following stability estimate is standard for least-squares residual minimization, including the least-squares finite element method [39, 40] and physics-informed neural networks [35, 38, 41].

Assumption 2.4 (Fine-scale stability). There exist a stability bound $C^\varepsilon_s \in \mathbb{R}_+$ and an upper bound $C^\varepsilon_b \in \mathbb{R}_+$, both possibly dependent on $\varepsilon > 0$, such that

$$C^\varepsilon_s \|u\|^2_V \leq \|\mathcal{A}^\varepsilon u\|^2_H + \|\mathcal{B}u\|^2_Z \leq C^\varepsilon_b \|u\|^2_U, \quad \forall u \in U. \tag{2.3}$$

The stability bound (2.3) is well-suited for PINN problems which include boundary conditions *via* penalization. In this case, a term penalizing violations of the boundary conditions is added to the PINN objective. Note, however, that boundary conditions may also be imposed exactly, *i.e.*, as (hard) constraints, for PINNs [31, 42, 43]. Then, one may require (2.3) to hold for all $u \in U_0$, while dropping the term $\|\mathcal{B}u\|^2_Z$ and using a stronger norm $\|\cdot\|_U$ in the lower bound. We demonstrate both approaches below in our example section and use (2.3) for our abstract formulation.

Let $(Y, \|\cdot\|_Y)$ be a Hilbert space with $U_0 \subseteq Y$, Y^* the topological dual space of Y and $Y \hookrightarrow H \cong H^* \hookrightarrow Y^*$ a Gelfand triple with compact embedding $Y \hookrightarrow H$. Let $\mathcal{L}[u] \in L(Y, Y^*)$ be parameterized by $u \in U$. The pertinent

bilinear form $b_{\mathcal{L}}[u] : Y \times Y \to \mathbb{R}$ is defined by

$$b_{\mathcal{L}}[u](v, w) := \langle \mathcal{L}[u]v, w\rangle_{Y^*,Y} \tag{2.4}$$

for $v, w \in Y$. The following assumption is needed to secure (local) solvability of the coarse-scale equation.

Assumption 2.5 (Coarse-scale stability). Let $u^\varepsilon \in U$ be the solution of (2.1), $\bar{r}^\varepsilon \in \mathbb{R}_+$ and $B_{\bar{r}^\varepsilon}(u^\varepsilon) := \{v \in U : \ \|u^\varepsilon - v\|_U \le \bar{r}^\varepsilon\}$. We assume that the form (2.4) is uniformly bounded and uniformly coercive for all $u \in B_{\bar{r}^\varepsilon}(u^\varepsilon)$, *i.e.*, there exist $C_b^\varepsilon, C_c^\varepsilon \in \mathbb{R}_+$, independent of u, but possibly dependent on ε, such that

$$b_{\mathcal{L}}[u](v, w) \le C_b^\varepsilon \|v\|_Y \|w\|_Y, \text{ and } \ b_{\mathcal{L}}[u](v, v) \ge C_c^\varepsilon \|v\|_Y^2, \quad \forall v, w \in Y. \tag{2.5}$$

For given $(u^\varepsilon, \tilde{f}) \in U \times H$, we consider the following partial differential equation (in weak form), which we refer to as the coarse-scale problem: find $y(u^\varepsilon) \in Y$ such that

$$b_{\mathcal{L}}[u^\varepsilon](y(u^\varepsilon), v) = \langle \tilde{f}, v\rangle_{Y^*,Y} \quad \forall v \in Y. \tag{2.6}$$

The coarse-scale problem (2.6) is well-posed by the Lax–Milgram lemma. Indeed, there exist a unique solution $y(u^\varepsilon) \in Y$ and a bound $C^\varepsilon \in \mathbb{R}_+$ such that $\|y(u^\varepsilon)\|_Y \le C^\varepsilon \|\tilde{f}\|_{Y^*}$, with C^ε possibly dependent on ε due to (2.5).

2.2. Physics-informed neural networks

For $L \in \mathbb{N}$, an L-layer feed-forward neural network (NN) is a recursively defined function $R_{\boldsymbol{\theta}}(x) : \mathbb{R}^{n_0} \to \mathbb{R}^{n_L}$ with

$$R_{\boldsymbol{\theta}}(x) = z^L(x), \ z^l(x) = W^l \sigma(z^{l-1}(x)) + b^l, \ 2 \le l \le L, \ z^1(x) = W^1 x + b^1,$$

where $W^l \in \mathbb{R}^{n_l \times n_{l-1}}$ is the l-th layer weight matrix, $b_l \in \mathbb{R}^{n_l}$ is the l-th layer bias vector and $\sigma : \mathbb{R} \to \mathbb{R}$ is the activation function, which is applied component-wise in case of input arguments in $\mathbb{R}^{n_l}$. The network architecture is represented by the vector $\vec{\boldsymbol{n}} = (n_0, ..., n_L)$, and the set of all possible network parameters is defined by

$$\Theta(\vec{\boldsymbol{n}}) = \big\{\{(W_j, b_j)\}_{j=1}^L : W_j \in \mathbb{R}^{n_j \times n_{j-1}}, \ b_j \in \mathbb{R}^{n_j}\big\}.$$

For a sequence $\{\vec{\boldsymbol{n}}_n\}_{n\ge1}$ of network architectures such that $\vec{\boldsymbol{n}}_n \le \vec{\boldsymbol{n}}_{n+1}$ for all n, with the vector inequality understood component-wise, we define its corresponding sequence of neural network classes [35], and the related maximal realization map

$$\mathfrak{N}_{\theta,n} := \{R_{\boldsymbol{\theta}} : \boldsymbol{\theta} \in \cup_{\vec{\boldsymbol{v}} \le \vec{\boldsymbol{n}}_n} \Theta(\vec{\boldsymbol{v}})\}, \text{ and } \ \mathfrak{F}_n : \mathbb{R}^{N_n} \to \mathfrak{N}_{\theta,n}, \quad \boldsymbol{\theta} \mapsto \mathfrak{F}_n(\boldsymbol{\theta}) =: v_{\boldsymbol{\theta},n}, \tag{2.7}$$

where N_n denotes the total number of parameters in $\Theta(\vec{\boldsymbol{n}}_n)$. Note that $\mathfrak{N}_{\theta,n} \subset \mathfrak{N}_{\theta,n+1}$ by construction, and that the regularity of $\mathfrak{F}_n$ depends on the one of the underlying activation functions in $\mathfrak{N}_{\theta,n}$. In what follows, $v_{\boldsymbol{\theta},n} \in \mathfrak{N}_{\theta,n}$ refers to a network function, generated by its maximal admissible number of parameters N_n.

The residual minimization of the fine-scale PDE over the class of neural networks leads to the following PINN optimization problem:

$$\inf \mathcal{J}(v_{\boldsymbol{\theta},n}) := \|\mathcal{A}^\varepsilon v_{\boldsymbol{\theta},n} - f^\varepsilon\|_H^2 + \tau_1 \|\mathcal{B} v_{\boldsymbol{\theta},n}\|_Z^2 \quad \text{over } v_{\boldsymbol{\theta},n} \in \mathfrak{N}_{\theta,n}, \tag{2.8}$$

where the penalty parameter $\tau_1 > 0$ is fixed. We observe that $\mathcal{J}(u) \ge 0$ for all $u \in U$, and $\mathcal{J}(u^\varepsilon) = 0$ for the solution u^ε to (2.1). Assuming momentarily that $U = H^2(\Omega)$ and $\mathcal{A}^\varepsilon$ is of second order, we need to ensure that

u^ε can be approximated by a neural network. It is well-known that Sobolev functions can be well approximated *via* popular deep ReLU neural networks; see, *e.g.*, [44]. Note, however, that the ReLU activation function $\sigma(x) = \max(0, x)$ admits only one weak derivative, which renders it infeasible when utilizing the standard least-squares loss with $H = L^2(\Omega)$ and $\mathcal{A}^\varepsilon$ involving higher-order (weak) derivatives. The hyperbolic tangent function $\tanh(x)$ can then be used to approximate functions from Sobolev spaces $H^k(\Omega)$, $k \geq 3$, in the norm of U (see [37], Thm. B.7, but also [45]). We mention that the regularity requirements on activation functions and u^ε can be relaxed by adopting a variational loss function [46, 47]. This involves multiplying the residuals of (2.1) by suitably smooth test functions and integrating by parts. Then, the respective weak residual is minimized over a (discrete) trial class of neural networks. In such a setting, one is confronted with the additional burden of discretizing the space of test functions. Typically, piecewise polynomials of low order, leading to a Petrov–Galerkin discretization, are preferred [48]. Moreover, the Fourier transform can be employed to efficiently evaluate the H^{-1} norm of the PDE residual for simple geometries [49].

In view of the above, in our work we focus on a setting with smooth activation functions. Instead of imposing assumptions on the sufficiently high number of admissible weak derivatives of u^ε, we rather work in a dense subspace X of U, where the approximation with neural networks is less restrictive, such as $X = C^k(\bar{\Omega})$, which motivates the triplet of spaces $\mathfrak{N}_{\theta,n} \subset X \subset U$ and our definition of solution (2.2). Then, we are interested in neural networks that can well approximate elements in X; *cf.* [35] and various neural network approximation results, sometimes referred to as universal approximation theorems [44, 45, 50–52]. Thus, the following assumption on universal approximation is required.

Assumption 2.6 (**Uniform NN approximation of elements in X**). There exists a sequence of neural network classes $\{\mathfrak{N}_{\theta,n}\}$ with $\mathfrak{N}_{\theta,n} \subset X$ and $\mathfrak{N}_{\theta,n} \subset \mathfrak{N}_{\theta,n+1}$ for all $n \in \mathbb{N}$, and $X \subset \overline{\cup_n \mathfrak{N}_{\theta,n}}$ in the topology of $(X, \|\cdot\|_X)$. In addition, it holds that $\mathcal{A}^\varepsilon v_{\boldsymbol{\theta},n} \in L^2(\Omega)$ and $\mathcal{B} v_{\boldsymbol{\theta},n} \in L^2(\partial\Omega)$ for all $v_{\boldsymbol{\theta},n} \in \mathfrak{N}_{\theta,n}$.

2.3. Weak convergence based regularization

Let $H = L^2(\Omega)$ and $H^\delta := L^2(\Omega_\delta)$, where $\Omega_\delta := \{z \in \Omega : \operatorname{dist}(z, \partial\Omega) > \frac{\delta}{2}\} \subset \Omega$ for $\delta > 0$. Weak convergence assumptions are common in the context of homogenization, and here we make use of the following one.

Assumption 2.7 (**Weak convergence**). Assume that for $\gamma > 0$, $u^\varepsilon \rightharpoonup y^0$ in $L^{2+\gamma}(\Omega)$ as $\varepsilon \to 0$, where u^ε is the solution to (2.1) and $y^0 \in Y$ is the solution to the homogenized equation $\mathcal{L}^0 y^0 = f^0$, with $\mathcal{L}^0 \in L(Y, Y^*)$ and $f^0 \in H$ ε-independent, respectively.

Remark 2.8. In our setting, f^ε in (2.1), $\tilde{f}$ in (2.6), and f^0 in Assumption 2.7 are different in general. This is due to the lifting of Dirichlet boundary conditions in all the PDEs, *cf.* Remark 2.2. Besides, we assume that y^0 has at least the same regularity as $y(u^\varepsilon)$, as it follows from Assumption 2.7 and our choice of spaces.

In addition, we need the following; *cf.*, *e.g.* [12], Theorem 4.5 or [25], Section 3.2.

Assumption 2.9 (**Consistent parametrization**). There exist $C_{\mathcal{L}[u]}, C_\varepsilon \in \mathbb{R}_+$ such that $\|y^0 - y(u^\varepsilon)\|_H \leq C_{\mathcal{L}[u]} + C_\varepsilon$, where $C_{\mathcal{L}[u]}$ does not depend on ε and can be made small by an appropriate parametrization of $y(u^\varepsilon)$ by u^ε, and $C_\varepsilon \to 0$ as $\varepsilon \to 0$.

For $v \in H$, $x \in \Omega_\delta$ and $\mathcal{V}_\delta(x) = \{z : \|z - x\|_{\mathbb{R}^d} \leq \frac{\delta}{2}\} \subset \Omega$, define $Q_\delta : H \to H^\delta$ by

$$(Q_\delta v)(x) := \frac{1}{|\mathcal{V}_\delta(x)|} \int_{\mathcal{V}_\delta(x)} v(z)\, dz.$$

Following heterogeneous multiscale methods [12], we fix the compression operator

$$(\bar{Q}_\delta v)(x) := \begin{cases} (Q_\delta v)(x), & x \in \Omega_\delta, \\ v(x), & x \in \Omega \setminus \Omega_\delta. \end{cases} \tag{2.9}$$

Let us next introduce the coupling term $\mathcal{R}_\delta : Y \times U \to \mathbb{R}_{\geq 0}$ as follows:

$$\mathcal{R}_\delta(y(u^\varepsilon), u^\varepsilon) := \|\bar{Q}_\delta u^\varepsilon - y(u^\varepsilon)\|_H^2. \tag{2.10}$$

The purpose of (2.10) is to equip the optimization problem (2.8) with information on weak convergence of the fine-scale solution to the coarse-scale one.

Next we study (2.10). We start with two preparatory result.

Lemma 2.10. *The operator $Q_\delta : H \to H^\delta$ has the following properties:*

1. *$Q_\delta \in L(H, H^\delta)$ and $\bar{Q}_\delta \in L(H) := L(H, H)$.*
2. *Suppose that $u^\varepsilon \rightharpoonup y^0$ in H as $\varepsilon \to 0$. Then, $\lim_{\varepsilon \to 0} \|Q_\delta u^\varepsilon - Q_\delta y^0\|_{H^\delta} = 0$.*

Proof. The linearity of Q_δ is obvious. For $x \in \Omega_\delta$, using the Cauchy–Schwarz inequality, we obtain the estimate $|(Q_\delta u^\varepsilon)(x)| \leq |\mathcal{V}_\delta(x)|^{-1/2} \|u^\varepsilon\|_H$. Since $|\mathcal{V}_\delta(x)|$ is constant for all $x \in \Omega_\delta$, the estimate is uniform. By integrating the estimate over Ω_δ, we get $\|Q_\delta u^\varepsilon\|_{H^\delta} \leq C \|u^\varepsilon\|_H$, where $C \in \mathbb{R}_+$. Therefore, $Q_\delta \in L(H, H^\delta)$ and we readily establish that $\bar{Q}_\delta \in L(H)$.

Let $u^\varepsilon \rightharpoonup y^0$ in H as $\varepsilon \to 0$. Then $\|u^\varepsilon\|_H \leq C$, where $C \in \mathbb{R}_+$ is ε-independent, and for all test functions $v \in H$ it holds that

$$\int_\Omega u^\varepsilon(x) v(x) \, \mathrm{d}x \to \int_\Omega y^0(x) v(x) \, \mathrm{d}x \quad \text{as} \quad \varepsilon \to 0. \tag{2.11}$$

Consider the normalized characteristic function

$$\bar{\chi}_{\mathcal{V}_\delta(x)}(z) := \begin{cases} \frac{1}{|\mathcal{V}_\delta(x)|}, & z \in \mathcal{V}_\delta(x), \\ 0, & z \notin \mathcal{V}_\delta(x), \end{cases}$$

as test function in (2.11). Then one obtains the pointwise convergence of averages $(Q_\delta u^\varepsilon)(x) \to (Q_\delta y^0)(x)$ as $\varepsilon \to 0$ for $x \in \Omega_\delta$. The L^2-convergence follows from the uniform estimate of $|(Q_\delta u^\varepsilon)(x)|$ and Lebesgue's Dominated Convergence Theorem. □

Since $Y \hookrightarrow H$, for a given $y \in Y$, almost every $x \in \Omega_\delta$ is a Lebesgue point, *i.e.*,

$$\lim_{\delta \to 0} \frac{1}{|\mathcal{V}_\delta(x)|} \int_{\mathcal{V}_\delta(x)} y(z) \, dz = y(x). \tag{2.12}$$

For small $\delta > 0$, we consider the following approximation of the above limit:

$$\frac{1}{|\mathcal{V}_\delta(x)|} \int_{\mathcal{V}_\delta(x)} y(z) \, dz \approx \lim_{\delta \to 0} \frac{1}{|\mathcal{V}_\delta(x)|} \int_{\mathcal{V}_\delta(x)} y(z) \, dz = y(x).$$

Additional integrability of ∇y provides a rate for such an approximation.

Lemma 2.11. *Suppose that $y \in W^{1,p}(\Omega)$ with $\Omega \subset \mathbb{R}^d$ and $d < p < \infty$. Then*

$$\|y - Q_\delta y\|_{H^\delta} \leq C_p(y) \delta^{\frac{d(p-2)+1}{p}},$$

where the constant $C_p(y) < \infty$ is independent of δ, but it depends on $\|\nabla y\|_{L^p(\Omega)}$.

Proof. First, we use a well-known trick that controls the deviation of a function from its average on convex sets (see *e.g.* [36], Lem. 4.28). Let $x \in \Omega_\delta$ and define the convex ball $\mathcal{V}_\delta(x) \subset \Omega$. For $y \in C^1(\bar{\Omega})$, $z \in \mathcal{V}_\delta(x)$ and for all $t \in [0,1]$, we get

$$y(z) - y(x) = \int_0^1 \frac{\mathrm{d}}{\mathrm{d}t} y(x + t(z-x))\mathrm{d}t = \int_0^1 \nabla y(x + t(z-x)) \cdot (z-x)\mathrm{d}t.$$

Integrating the above equality over $\mathcal{V}_\delta(x)$ and performing the change of variables $\xi = x + t(z-x)$, we obtain

$$\left| \int_{\mathcal{V}_\delta(x)} y(z)dz - |\mathcal{V}_\delta(x)|\, y(x) \right| \leq \int_0^1 \frac{1}{t^d} \int_{\mathcal{V}_{t\delta}(x)} |\nabla y(\xi)| \left| \frac{\xi - x}{t} \right| \mathrm{d}\xi\, \mathrm{d}t. \tag{2.13}$$

Hölder's inequality with $\frac{1}{p} + \frac{1}{q} = 1$, in conjunction with the fact that $\mathcal{V}_{t\delta}(x) \subseteq \mathcal{V}_\delta(x)$, yields the following estimates:

$$\begin{aligned}
\int_0^1 \frac{1}{t^{d+1}} \int_{\mathcal{V}_{t\delta}(x)} |\nabla y(\xi)||\xi - x| \mathrm{d}\xi\, \mathrm{d}t &\leq \|\nabla y\|_{L^p(\mathcal{V}_\delta(x))} \int_0^1 \frac{1}{t^{d+1}} \left(\int_{\mathcal{V}_{t\delta}(x)} |\xi - x|^q \mathrm{d}\xi \right)^{\frac{1}{q}} \mathrm{d}t \\
&\leq \|\nabla y\|_{L^p(\mathcal{V}_\delta(x))} \int_0^1 \frac{\mathrm{d}t}{t^{\frac{d}{p}}} \left(\int_{\mathcal{V}_\delta(x)} |x - z|^q dz \right)^{\frac{1}{q}} \leq \frac{\|\nabla y\|_{L^p(\Omega)}}{1 - \frac{d}{p}} \left(\int_{\mathcal{V}_\delta(x)} |x - z|^q dz \right)^{\frac{1}{q}},
\end{aligned}$$

where we used the boundedness of the $\mathrm{d}t$-integral for $\frac{d}{p} < 1$. Since $\rho(r) := \rho(|x-z|) = |x-z|^q$ is radially symmetric on $\mathcal{V}_\delta(x)$, integration in polar coordinates yields:

$$\left(\int_{\mathcal{V}_\delta(x)} |x - z|^q dz \right)^{\frac{1}{q}} = |\mathcal{S}_\delta|^{\frac{1}{q}} \left(\int_0^{\delta/2} \rho(r) r^{d-1} dr \right)^{\frac{1}{q}} = C(d) \delta^{\frac{2d-1+q}{q}},$$

where $\mathcal{S}_\delta$ is the $(d-1)$-sphere of radius $\frac{\delta}{2}$ with $|\mathcal{S}_\delta| = \frac{2\pi^{d/2}}{\Gamma(d/2)} (\delta/2)^{d-1}$, Γ is Euler's gamma function and $C(d) \in \mathbb{R}$ is δ-independent. Dividing (2.13) by $|\mathcal{V}_\delta(x)|$ with $|\mathcal{V}_\delta(x)| = \frac{\pi^{d/2}}{\Gamma(d/2+1)} (\delta/2)^d$, combining the estimates and integrating, we get

$$\left(\int_{\Omega_\delta} |y(x) - (Q_\delta y)(x)|^2 \, \mathrm{d}x \right)^{1/2} \leq C_p(y) \delta^{\frac{d(2-q)+q-1}{q}},$$

where $C_p(y) := C(d, p, \Omega) \|\nabla y\|_{L^p(\Omega)}$. The density of $C^1(\bar{\Omega})$ in $W^{1,p}(\Omega)$ extends the above estimate to the desired result. □

Under a mild assumption on the regularity of $y^0 \in Y$, we prove the following result.

Theorem 2.12 (**Upscaling consistency**). *Suppose that Assumptions 2.7, 2.9 hold, $y^0 \in Y \subset H^2(\Omega)$ with $\Omega \subset \mathbb{R}^d$, $d < p < \infty$ and $d \leq 3$. Then, for $\delta > 0$, it holds:*

$$\lim_{\varepsilon \to 0} \mathcal{R}_\delta(y(u^\varepsilon), u^\varepsilon) \leq C_p(y^0)^2 \delta^{\frac{2d(p-2)+2}{p}} + 2C^2_{\mathcal{L}[u]} + |\Omega \setminus \Omega_\delta|^{\frac{2+\gamma}{\gamma}} C_w,$$

where $C_p(y^0) \in \mathbb{R}_+$ and $C_{\mathcal{L}[u]}$ are according to Lemma 2.11 and Assumption 2.9, respectively, and $C_w \in \mathbb{R}_+$ is ε-independent.

Proof. The triangle inequality and Young's inequality yield

$$\mathcal{R}_\delta(y(u^\varepsilon), u^\varepsilon) \leq 2\|\bar{Q}_\delta u^\varepsilon - y^0\|_H^2 + 2\|y^0 - y(u^\varepsilon)\|_H^2. \tag{2.14}$$

Similarly, the estimation of the first term in (2.14) gives

$$\|\bar{Q}_\delta u^\varepsilon - y^0\|_H^2 \leq 2\|Q_\delta u^\varepsilon - Q_\delta y^0\|_{H^\delta}^2 + 2\|y^0 - Q_\delta y^0\|_{H^\delta}^2 + \|u^\varepsilon - y^0\|_{L^2(\Omega\setminus\Omega_\delta)}^2. \tag{2.15}$$

Clearly, Assumptions 2.7 implies that $u^\varepsilon \rightharpoonup y^0$ as $\varepsilon \to 0$ in H. Thus, Lemma 2.10 guarantees that the first term on the right-hand side of (2.15) vanishes as $\varepsilon \to 0$. Since $\nabla y^0 \in H^1(\Omega)$ due to $Y \subset H^2(\Omega)$ and $d \leq 3$, the Sobolev embedding yields $\|\nabla y^0\|_{L^p(\Omega)} < \infty$ for $p \leq 6$. Observe further that y^0 is ε-independent, hence $C_p(y^0)$ is also ε-independent. Therefore, Lemma 2.11 guarantees that $\|y^0 - Q_\delta y^0\|_{H^\delta}^2 \leq C_p(y^0)^2 \delta^{\frac{2d(p-2)+2}{p}}$ and the latter estimate holds in the limit $\varepsilon \to 0$.

We estimate the second term in (2.14) by appealing to Assumption 2.9. Next, let $s = \frac{2+\gamma}{2}$ and $r = \frac{2+\gamma}{\gamma}$. Then $\frac{1}{s} + \frac{1}{r} = 1$ and the Hölder estimate

$$\|u^\varepsilon - y^0\|_{L^2(\Omega\setminus\Omega_\delta)}^2 \leq |\Omega \setminus \Omega_\delta|^{\frac{2+\gamma}{\gamma}} \|u^\varepsilon - y^0\|_{L^{2+\gamma}(\Omega)}^2$$

holds for the third term in (2.15). The uniform boundedness of $\{u^\varepsilon\}$ implies the existence of $C_w \in \mathbb{R}_+$ such that $\|u^\varepsilon - y^0\|_{L^{2+\gamma}(\Omega)}^2 \leq C_w$, hence completing the proof. □

Suppose that $\nabla y^0 \notin L^p(\Omega)$ for some $p > d$. Then Lemma 2.11 is not applicable. In this case, the coupling term (2.10) can be modified as follows:

$$\mathcal{R}_\delta(y(u^\varepsilon), u^\varepsilon) := \|\bar{Q}_\delta u^\varepsilon - \bar{Q}_\delta y(u^\varepsilon)\|_H^2. \tag{2.16}$$

Then, Theorem 2.12 still holds with $C_p(y^0) = 0$.

3. Learning-informed PDE-constrained optimization

We cast our hybrid physics-informed neural network based multiscale approach for a fixed $\varepsilon > 0$ into the following learning-informed PDE-constrained optimization problem:

$$\begin{cases} \inf J(y, v_{\boldsymbol{\theta},n}) := \mathcal{J}(v_{\boldsymbol{\theta},n}) + \tau_2 \mathcal{R}_\delta(y, v_{\boldsymbol{\theta},n}) \text{ over } (y, v_{\boldsymbol{\theta},n}) \in Y \times B_{\bar{r}^\varepsilon,\theta,n}(u^\varepsilon), \\ \text{s.t. } e(y, v_{\boldsymbol{\theta},n}) = 0, \end{cases} \tag{3.1}$$

where $B_{\bar{r}^\varepsilon,\theta,n}(u^\varepsilon) := B_{\bar{r}^\varepsilon}(u^\varepsilon) \cap \mathfrak{N}_{\theta,n}$ with u^ε, $\bar{r}^\varepsilon$ according to Assumption 2.5, $J : Y \times U \to \mathbb{R}_{\geq 0}$ and fixed $\tau_2, \delta \in \mathbb{R}_{\geq 0}$. Further, let $e : Y \times U \to Y^*$ be given by $e : (w, v) \mapsto e(w, v) := b_{\mathcal{L}}[v](w, \cdot) - \langle \tilde{f}, \cdot \rangle_{Y^*,Y}$ and $\mathcal{R}_\delta$ denote the coupling term in (2.16). Note that for u^ε according to Assumption 2.5 there exists $v \in X$ arbitrarily close to u^ε due to the density of the embedding $X \hookrightarrow U$. Then, Assumption 2.6 implies the existence of $\boldsymbol{\theta} \in \mathbb{R}^{N_{n_\varepsilon^*}}$ such that

$$\|u^\varepsilon - v_{\boldsymbol{\theta},n_\varepsilon^*}\|_U \leq \|u^\varepsilon - v\|_U + C\|v - v_{\boldsymbol{\theta},n_\varepsilon^*}\|_X \leq \bar{r}^\varepsilon,$$

for sufficiently large $n_\varepsilon^* \in \mathbb{N}$. Then, for all $n \geq n_\varepsilon^*$, $B_{\bar{r}^\varepsilon,\theta,n}(u^\varepsilon) \neq \emptyset$ and the following *fine-to-coarse scale map* is well-defined:

$$S : B_{\bar{r}^\varepsilon,\theta,n}(u^\varepsilon) \subset U \to Y, \quad u \mapsto y(u) := S(u), \tag{3.2}$$

with $e(y(u),u)=0$. Note that since (2.7) is non-convex in general, finding $\boldsymbol{\theta}$ with $v_{\boldsymbol{\theta},n}\in B_{\bar r^\varepsilon,\theta,n}(u^\varepsilon)$ requires solving an associated non-convex optimization problem, which can be quite delicate in practice; *cf.* [9, 11, 53].

We also need the following.

Assumption 3.1 (**Continuity**). Let $u^\varepsilon\in U$ be the solution of (2.1) and $\{u_k^\varepsilon\}$ its approximating sequence in $B_{\bar r^\varepsilon}(u^\varepsilon)\cap X$. Then $S(u_k^\varepsilon)\rightharpoonup S(u^\varepsilon)$ in Y as $k\to\infty$.

Eliminating y from the set of independent variables in (3.1) results in the reduced optimization problem

$$\inf \widehat{J}(v_{\boldsymbol{\theta},n}) := J(S(v_{\boldsymbol{\theta},n}),v_{\boldsymbol{\theta},n}) \quad \text{over } v_{\boldsymbol{\theta},n}\in B_{\bar r^\varepsilon,\theta,n}(u^\varepsilon). \tag{3.3}$$

We note that guaranteeing the existence of minimizers in $\mathfrak{N}_{\theta,n}$ for (3.3) is not possible, in general, as $\mathfrak{N}_{\theta,n}$ may not be topologically closed in U. In order to cope with this, we resort to the notion of quasi-minimization; *cf.* [35], see also [54]. The latter only requires the existence of an infimum of $\widehat{J}$ over $B_{\bar r^\varepsilon,\theta,n}(u^\varepsilon)\neq\emptyset$.

Clearly, since $\widehat{J}(\cdot)\geq 0$, $\inf\widehat{J}$ exists over $B_{\bar r^\varepsilon,\theta,n}(u^\varepsilon)$ for every $n\geq n_\varepsilon^*$. Now let $\{\gamma_n\}_{n\geq n_\varepsilon^*}$ be a real sequence with $\gamma_n>0$ for all $n\geq n_\varepsilon^*$ and $\gamma_n\downarrow 0$ as $n\to\infty$. Then, for every $n\geq n_\varepsilon^*$, there exists $u_{\boldsymbol{\theta},n}^\varepsilon\in B_{\bar r^\varepsilon,\theta,n}(u^\varepsilon)$ such that

$$\widehat{J}(u_{\boldsymbol{\theta},n}^\varepsilon)\leq \inf_{v_{\boldsymbol{\theta},n}\in B_{\bar r^\varepsilon,\theta,n}(u^\varepsilon)}\widehat{J}(v_{\boldsymbol{\theta},n})+\gamma_n.$$

We refer to $\{u_{\boldsymbol{\theta},n}^\varepsilon\}_{n\geq n_\varepsilon^*}$ as a sequence of quasi-minimizers of (3.3). In the following result, we consider J with the coupling term (2.16), but extending it to (2.10) is straightforward.

Theorem 3.2. *Suppose that Assumptions 2.5, 2.4, 2.6, and 3.1 hold. Let $\{u_{\boldsymbol{\theta},n}^\varepsilon\}_{n\geq n_\varepsilon^*}$, $u_{\boldsymbol{\theta},n}^\varepsilon\in B_{\bar r^\varepsilon,\theta,n}(u^\varepsilon)$ be a quasi-minimizing sequence for $\widehat{J}:U\to\mathbb{R}_{\geq 0}$, where $B_{\bar r^\varepsilon}(u^\varepsilon)$ is chosen according to Assumption 2.5 and $n_\varepsilon^*\in\mathbb{N}$ is chosen according to Assumption 2.6 to guarantee that $B_{\bar r^\varepsilon,\theta,n}(u^\varepsilon)\neq\emptyset$ for all $n\geq n_\varepsilon^*$. Then, $\lim_{n\to\infty}\widehat{J}(u_{\boldsymbol{\theta},n}^\varepsilon)\leq \tau_2\mathcal{R}_\delta(y(u^\varepsilon),u^\varepsilon)$, where u^ε is the solution to (2.1) according to Assumption 2.1. Moreover, for $\tau_1\geq 1$ it holds that*

$$\lim_{n\to\infty}\|u_{\boldsymbol{\theta},n}^\varepsilon-u^\varepsilon\|_V\leq\frac{1}{\sqrt{C_s^\varepsilon}}\sqrt{\tau_2\mathcal{R}_\delta(y(u^\varepsilon),u^\varepsilon)}. \tag{3.4}$$

Proof. Let u^ε be the solution to (2.1) and $\{u_k^\varepsilon\}$, $u_k^\varepsilon\in X$ for all k, be its approximating sequence (2.2). Let $\{r_k\}$ be a positive real sequence with $r_k\downarrow 0$ monotonically as $k\to\infty$ and choose $K_\varepsilon\in\mathbb{N}$ sufficiently large such that $B_{r_k}(u_k^\varepsilon):=\{v\in U:\ \|u_k^\varepsilon-v\|_U\leq r_k\}\subset B_{\bar r^\varepsilon}(u^\varepsilon)$ for all $k\geq K_\varepsilon$. Assumption 2.6 and the dense embedding $X\hookrightarrow U$ imply that $\|u_k^\varepsilon-v_{\boldsymbol{\theta},n_k}^\varepsilon\|_U\leq C\|u_k^\varepsilon-v_{\boldsymbol{\theta},n_k}^\varepsilon\|_X\leq C\epsilon_{n_k}\leq r_k$ for $k\geq K_\varepsilon$ and some $v_{\boldsymbol{\theta},n_k}^\varepsilon\in\mathfrak{N}_{\theta,n_k}$, where $n_k:=n(k)\in\mathbb{N}$, $\epsilon_{n_k}\in\mathbb{R}_+$ and $C>0$ is some embedding constant. Observe further that $B_{r_k}(u_k^\varepsilon)\subset B_{\bar r^\varepsilon}(u^\varepsilon)$ for $k\geq K_\varepsilon$ implies $\epsilon_{n_{K_\varepsilon}}\leq\epsilon_{n_\varepsilon^*}$. Therefore, $n_{K_\varepsilon}\geq n_\varepsilon^*$ with $\mathfrak{N}_{\theta,n_\varepsilon^*}\subseteq\mathfrak{N}_{\theta,n_{K_\varepsilon}}$ holds and there exists a sequence $\{v_{\boldsymbol{\theta},n_k}^\varepsilon\}_{k\geq K_\varepsilon}$ with $v_{\boldsymbol{\theta},n_k}^\varepsilon\in B_{r_k,\theta,n_k}(u_k^\varepsilon)\subset B_{\bar r^\varepsilon}(u^\varepsilon)$. We note that Assumption 2.6 implies that $\{\epsilon_{n_k}\}_{k\geq K_\varepsilon}$ converges to 0 as $k\to\infty$, thus $n_k\to\infty$ as $k\to\infty$. Then, once again the embedding $X\hookrightarrow U$ implies that

$$\|u^\varepsilon-v_{\boldsymbol{\theta},n_k}^\varepsilon\|_U\leq\|u^\varepsilon-u_k^\varepsilon\|_U+C\|u_k^\varepsilon-v_{\boldsymbol{\theta},n_k}^\varepsilon\|_X\to 0\quad\text{as}\quad k\to\infty.$$

In addition, the boundedness of $\mathcal{A}^\varepsilon$ and $\mathcal{B}$, respectively, and the upper bound in (2.3) result in the existence of $C^\varepsilon\in\mathbb{R}_+$, possibly dependent on ε, such that

$$\|\mathcal{A}^\varepsilon v_{\boldsymbol{\theta},n_k}^\varepsilon-f\|_H+\|\mathcal{B}v_{\boldsymbol{\theta},n_k}^\varepsilon\|_Z\leq C^\varepsilon\|v_{\boldsymbol{\theta},n_k}^\varepsilon-u_k^\varepsilon\|_X+\|\mathcal{A}^\varepsilon u_k^\varepsilon-f\|_H+\|\mathcal{B}u_k^\varepsilon\|_Z\to 0$$

as $k\to 0$. Therefore, $\{v_{\boldsymbol{\theta},n_k}^\varepsilon\}_{k\geq K_\varepsilon}$ is also an approximating sequence (2.2).

Next, we study the coupling term and show that

$$\lim_{k\to\infty} \mathcal{R}_\delta(S(v^\varepsilon_{\boldsymbol{\theta},n_k}), v^\varepsilon_{\boldsymbol{\theta},n_k}) \le \mathcal{R}_\delta(y(u^\varepsilon), u^\varepsilon). \tag{3.5}$$

Indeed, invoking the triangle inequality and $\bar{Q}_\delta \in L(H)$, we get

$$\begin{aligned}\mathcal{R}_\delta(S(v^\varepsilon_{\boldsymbol{\theta},n_k}), v^\varepsilon_{\boldsymbol{\theta},n_k}) \le \|\bar{Q}_\delta\|^2 \|v^\varepsilon_{\boldsymbol{\theta},n_k} - u^\varepsilon_k\|^2_H + \|\bar{Q}_\delta u^\varepsilon_k - \bar{Q}_\delta y(v^\varepsilon_{\boldsymbol{\theta},n_k})\|^2_H \\ + 2\|\bar{Q}_\delta\| \|v^\varepsilon_{\boldsymbol{\theta},n_k} - u^\varepsilon_k\|_H \|\bar{Q}_\delta u^\varepsilon_k - \bar{Q}_\delta y(v^\varepsilon_{\boldsymbol{\theta},n_k})\|_H.\end{aligned} \tag{3.6}$$

Since $v^\varepsilon_{\boldsymbol{\theta},n_k} \in B_{\bar{r}^\varepsilon,\theta,n^*_\varepsilon}(u^\varepsilon)$ for $k \ge K_\varepsilon$, the fine-to-coarse mapping (3.2) is well-defined in (3.6). The embedding $U \hookrightarrow H$ guarantees that the first and the last term in (3.6) vanish as $k \to 0$. Similarly, one estimates the intermediate term in (3.6):

$$\begin{aligned}\|\bar{Q}_\delta u^\varepsilon_k - \bar{Q}_\delta y(v^\varepsilon_{\boldsymbol{\theta},n_k})\|^2_H \le \mathcal{R}_\delta(y(u^\varepsilon), u^\varepsilon_k) + \|\bar{Q}_\delta\|^2 \|y(u^\varepsilon) - y(v^\varepsilon_{\boldsymbol{\theta},n_k})\|^2_H \\ + 2\|\bar{Q}_\delta\| \|\bar{Q}_\delta u^\varepsilon_k - \bar{Q}_\delta y(u^\varepsilon)\|_H \|y(u^\varepsilon) - y(v^\varepsilon_{\boldsymbol{\theta},n_k})\|_H.\end{aligned} \tag{3.7}$$

We use Assumption 3.1 and the compactness of $Y \hookrightarrow H$ to obtain $\lim_{k\to\infty} \|y(v^\varepsilon_{\boldsymbol{\theta},n_k}) - y(u^\varepsilon)\|^2_H = 0$. Therefore, (3.5) readily follows from (3.6) and (3.7).

Next, we study the limit $\lim_{n\to\infty} \widehat{J}(u^\varepsilon_{\boldsymbol{\theta},n})$. The stability bound (2.4) gives

$$\mathcal{J}(v^\varepsilon_{\boldsymbol{\theta},n_k}) \le (\|f - \mathcal{A}^\varepsilon u^\varepsilon_k\|_H + C^\varepsilon_b r_k)^2 + \tau_1(\|\mathcal{B}u^\varepsilon_k\|_Z + C^\varepsilon_b r_k)^2. \tag{3.8}$$

Since $B_{\bar{r}^\varepsilon,\theta,n_{K_\varepsilon}}(u^\varepsilon) \neq \emptyset$ and $B_{\bar{r}^\varepsilon,\theta,n_{K_\varepsilon}}(u^\varepsilon) \subset B_{\bar{r}^\varepsilon}(u^\varepsilon)$, we can find a quasi-minimizer $u^\varepsilon_{\boldsymbol{\theta},n_k} \in B_{\bar{r}^\varepsilon,\theta,n_k}(u^\varepsilon)$ for each $k \ge K_\varepsilon$ with $\gamma_{n_k} > 0$ and $\gamma_{n_k} \downarrow 0$ such that

$$\widehat{J}(u^\varepsilon_{\boldsymbol{\theta},n_k}) \le \inf_{v \in B_{\bar{r}^\varepsilon,\theta,n_k}(u^\varepsilon)} \widehat{J}(v) + \gamma_{n_k} \le \widehat{J}(v^\varepsilon_{\boldsymbol{\theta},n_k}) + \gamma_{n_k}. \tag{3.9}$$

Then, it follows from (3.5) and (3.8) that $\lim_{k\to\infty} \widehat{J}(v^\varepsilon_{\boldsymbol{\theta},n_k}) \le \tau_2 \mathcal{R}_\delta(y(u^\varepsilon), u^\varepsilon)$. Hence, for any $\epsilon_J > 0$, there exists $K_{\epsilon_J} \ge K_\varepsilon$ such that $\widehat{J}(u^\varepsilon_{\boldsymbol{\theta},n_k}) \le \epsilon_J/2 + \tau_2 \mathcal{R}_\delta(y(u^\varepsilon), u^\varepsilon)$ for all $k \ge K_{\epsilon_J}$. By resorting to the notion of a quasi-minimizing sequence $\{u^\varepsilon_{\boldsymbol{\theta},n}\}_{n \ge n^*_\varepsilon}$, the existence of $N_{\epsilon_J} \ge n^*_\varepsilon$ follows with $\gamma_n \le \epsilon_J/2$ for all $n \ge N_{\epsilon_J}$. For $\hat{N}_{\epsilon_J} = \max\{n_{K_{\epsilon_J}}, N_{\epsilon_J}\}$ and $n \ge \hat{N}_{\epsilon_J}$, we have $\mathfrak{N}_{\theta,\hat{N}_{\epsilon_J}} \subseteq \mathfrak{N}_{\theta,n}$, and

$$\widehat{J}(u^\varepsilon_{\boldsymbol{\theta},n}) \le \inf_{v \in B_{\bar{r}^\varepsilon,\theta,n}(u^\varepsilon)} \widehat{J}(v) + \gamma_n \le \widehat{J}(u^\varepsilon_{\boldsymbol{\theta},\hat{N}_{\epsilon_J}}) + \gamma_n \le \epsilon_J + \tau_2 \mathcal{R}_\delta(y(u^\varepsilon), u^\varepsilon).$$

Since ϵ_J was arbitrarily chosen, the first limit claim is shown.

The stability estimate (2.3) and $\tau_1 \ge 1$ imply that

$$C^\varepsilon_s \|u^\varepsilon_{\boldsymbol{\theta},n} - u^\varepsilon\|^2_V \le \|\mathcal{A}^\varepsilon u^\varepsilon_{\boldsymbol{\theta},n} - \mathcal{A}^\varepsilon u^\varepsilon\|^2_H + \tau_1 \|\mathcal{B}u^\varepsilon_{\boldsymbol{\theta},n} - \mathcal{B}u^\varepsilon\|^2_Z \le \widehat{J}(u^\varepsilon_{\boldsymbol{\theta},n}).$$

Then, (3.4) follows from the previous result and the a-posteriori bound above. □

Note that it may happen that $C^\varepsilon_s \to 0$ as $\varepsilon \to 0$ for the standard $H = L^2(\Omega)$ PINN objective, making (3.4) merely of qualitative nature.

Remark 3.3. In general, $\inf \widehat{J}$ and $\inf \mathcal{J}$ over $B_{\overline{r}^\varepsilon,\theta,n}(u^\varepsilon)$ are different. Therefore, quasi-minimizers of $\widehat{J}$ and $\mathcal{J}$ are also different. For a quasi-minimizing sequence $\{u^\varepsilon_{\boldsymbol{\theta},n}\}_{n\in\mathbb{N}}$ of $\mathcal{J}$, it holds that $\|u^\varepsilon_{\boldsymbol{\theta},n} - u^\varepsilon\|_V \leq \frac{1}{\sqrt{C^\varepsilon_s}}\sqrt{\mathcal{J}(u^\varepsilon_{\boldsymbol{\theta},n})}$ with $\lim_{n\to\infty} \mathcal{J}(u^\varepsilon_{\boldsymbol{\theta},n}) = 0$; *cf.* [35], Proposition 3.1 and Theorem 3.2.

Comparing Theorem 3.2 with Remark 3.3, we conclude the following: while solving (3.1), one needs to identify parameters $\boldsymbol{\theta}$ in the respective quasi-minimizers $u^\varepsilon_{\boldsymbol{\theta},n}$ such that $\mathcal{J}(u^\varepsilon_{\boldsymbol{\theta},n}) \to 0$ as $n \to \infty$. However, since $\bar{Q}_\delta u^\varepsilon \neq y(u^\varepsilon)$ in $L^2(\Omega)$, we also expect that $\bar{Q}_\delta u^\varepsilon_{\boldsymbol{\theta},n} \neq y(u^\varepsilon_{\boldsymbol{\theta},n})$ in $L^2(\Omega)$ as well. Therefore, for some threshold constant $C^\varepsilon_{tr} \in \mathbb{R}_+$ dependent on u^ε, forcing $\mathcal{R}_\delta(y(u^\varepsilon_{\boldsymbol{\theta},n}), u^\varepsilon_{\boldsymbol{\theta},n}) < C^\varepsilon_{tr}$ by selecting (non-upscaling consistent) parameters $\boldsymbol{\theta}$ through the optimization process may violate the underlying physics of (2.1), rendering convergence $\lim_{n\to\infty} \mathcal{J}(u^\varepsilon_{\boldsymbol{\theta},n}) = 0$ not possible. In our setting, $\mathcal{R}_\delta$ serves the purpose of a regularizer in the learning process, but if no physics violation is caused by $\mathcal{R}_\delta$, it also acts as a preconditioner for (2.8), and, based on Theorem 2.12, we might expect its efficiency to increase as $\varepsilon \to 0$.

3.1. Discrete approximation

The (fully) discrete version of (3.1) comes in three steps: (i) First, we consider $\boldsymbol{\theta} \to v_{\boldsymbol{\theta},n}$, reducing every NN-function to its finite set of generating parameters belonging to $\mathbb{R}^{N_n}$; (ii) then, we replace the state y by a finite dimensional approximation y_h, with $h > 0$ indicating the associated discretization parameter such as, *e.g.*, the mesh width in a finite element method [55]; (iii) finally, the discrete version $\hat{\mathfrak{J}}$ of $\hat{J}$ is replaced by its quadrature approximation $\hat{\mathfrak{J}}^{M,h}_D$, where M indicates the number of collocation or quadrature points, and D is related to a discretization of the compression operator (2.9).

Let us start with (i) and define

$$\mathcal{S} := S \circ \mathfrak{F}_n : \mathbb{R}^{N_n} \to Y, \quad \boldsymbol{\theta} \mapsto y(v_{\boldsymbol{\theta},n}) := \mathcal{S}(\boldsymbol{\theta}), \tag{3.10}$$

$$\widehat{\mathfrak{J}} := \widehat{J} \circ \mathfrak{F}_n : \mathbb{R}^{N_n} \to \mathbb{R}_{\geq 0}, \quad \boldsymbol{\theta} \mapsto \widehat{\mathfrak{J}}(\boldsymbol{\theta}) := J(\mathcal{S}(\boldsymbol{\theta}), \mathfrak{F}_n(\boldsymbol{\theta})), \tag{3.11}$$

where $\mathfrak{F}_n$ is according to (2.7). We invoke the following assumption on differentiability.

Assumption 3.4 (Fréchet differentiability). The fine-to-coarse scale map (3.2) is continuously Fréchet differentiable and $\mathfrak{F}_n \in C^\infty(\mathbb{R}^{N_n}, X)$.

Assumption 3.4 and the chain rule imply that $\mathcal{S}$ is continuously Fréchet differentiable. We also note that $\mathfrak{F}_n$ satisfies Assumption 3.4 for smooth activation functions in a neural network, *i.e.*, $\sigma \in C^\infty(\mathbb{R})$. Accordingly, the differentiability requirement can be reduced *via* reducing the one on σ. *Via* the implicit function theorem for $e(y(v_{\boldsymbol{\theta},n}), \mathfrak{F}_n(\boldsymbol{\theta})) = 0$ we obtain

$$e_y(y(v_{\boldsymbol{\theta},n}), \boldsymbol{\theta})\mathcal{S}'(\boldsymbol{\theta}) + e_v(y(v_{\boldsymbol{\theta},n}), \mathfrak{F}_n(\boldsymbol{\theta}))\mathfrak{F}_n'(\boldsymbol{\theta}) = 0, \tag{3.12}$$

where we have $\mathcal{S}'(\boldsymbol{\theta}) = S'(\mathfrak{F}_n(\boldsymbol{\theta}))\mathfrak{F}_n'(\boldsymbol{\theta}) \in L(\mathbb{R}^{N_n}, Y)$ and, upon obvious identification, $\mathcal{S}'(\boldsymbol{\theta})^* \in L(Y^*, \mathbb{R}^{N_n})$.

By applying the chain rule, we find the gradient of (3.11):

$$\nabla\widehat{\mathfrak{J}}(\boldsymbol{\theta}) = \mathcal{S}'(\boldsymbol{\theta})^*\partial_y J(y(v_{\boldsymbol{\theta},n}), \mathfrak{F}_n(\boldsymbol{\theta})) + \mathfrak{F}_n'(\boldsymbol{\theta})^*\partial_v J(y(v_{\boldsymbol{\theta},n}), \mathfrak{F}_n(\boldsymbol{\theta})), \tag{3.13}$$

In practical realizations of PINNs, the second summand in (3.13) is typically produced by automatic differentiation, and the first summand is realized *via* the adjoint method [34]. For the latter, we need the bilinear form $b_{\mathcal{L}^*}[v_{\boldsymbol{\theta},n}](\cdot,\cdot) : Y \times Y \to \mathbb{R}$ with

$$b_{\mathcal{L}^*}[v_{\boldsymbol{\theta},n}](w, v) := \langle e_y(y(v_{\boldsymbol{\theta},n}), \mathfrak{F}_n(\boldsymbol{\theta}))^* w, v\rangle_{Y^*,Y}. \tag{3.14}$$

The following guarantees that (3.14) is well-defined.

Assumption 3.5. Suppose that $B_{\bar{r}^\varepsilon,\theta,n^*_\varepsilon}(u^\varepsilon) \subset U$, where $\bar{r}^\varepsilon \in \mathbb{R}_+$ is according to Assumption 2.5 and n^*_ε is chosen sufficiently large as in Assumption 2.6. Then, there exist $C^\varepsilon_{b^*}, C^\varepsilon_{c^*} \in \mathbb{R}_+$ such that for all $v_{\boldsymbol{\theta},n} \in B_{\bar{r},\theta,n}(u^\varepsilon)$ with $n \geq n^*_\varepsilon$ it holds that

$$b_{\mathcal{L}^*}[v_{\boldsymbol{\theta},n}](v,w) \leq C^\varepsilon_{b^*}\|v\|_Y\|w\|_Y, \text{ and } b_{\mathcal{L}^*}[v_{\boldsymbol{\theta},n}](v,v) \geq C^\varepsilon_{c^*}\|v\|^2_Y, \ \forall v,w \in Y. \tag{3.15}$$

Note that Assumption 3.5 and (3.12) yield the adjoint equation

$$e_y(y(v_{\boldsymbol{\theta},n}),\mathfrak{F}_n(\boldsymbol{\theta}))^* p(v_{\boldsymbol{\theta},n}) = -\partial_y J(y(v_{\boldsymbol{\theta},n}),\mathfrak{F}_n(\boldsymbol{\theta})) = 2\tau_2 \bar{Q}^*_\delta(\bar{Q}_\delta v_{\boldsymbol{\theta},n} - \bar{Q}_\delta y(v_{\boldsymbol{\theta},n})),$$

or its counterpart in weak form:

$$b_{\mathcal{L}^*}[v_{\boldsymbol{\theta},n}](p(v_{\boldsymbol{\theta},n}),v) = 2\tau_2\langle \bar{Q}_\delta v_{\boldsymbol{\theta},n} - \bar{Q}_\delta y(v_{\boldsymbol{\theta},n}), \bar{Q}_\delta v\rangle_H \quad \forall v \in Y, \tag{3.16}$$

where $p(v_{\boldsymbol{\theta},n}) := p(y(v_{\boldsymbol{\theta},n})) \in Y$ denotes the adjoint variable (or adjoint state, sometimes also called co-state).

Now we come to the second step of discretization. Here we use the finite element (FE) method applied to the coarse-scale equation. More specifically, let $Y_h := \text{span}\{\phi_j,\ 1 \leq j \leq N_h\} \subset Y$, $N_h \in \mathbb{N}$, be the standard finite dimensional space of piecewise-linear and globally continuous functions over a domain $\Omega \subset \mathbb{R}^d$. Of course, other choices of Y_h are possible as well [55]. The finite element approximation of (2.6), which involves the neural network based function $v_{\boldsymbol{\theta},n}$ as data, is then obtained by a standard Galerkin projection: Find $y_h(v_{\boldsymbol{\theta},n}) \in Y_h$ such that

$$b_{\mathcal{L}}[v_{\boldsymbol{\theta},n}](y_h(v_{\boldsymbol{\theta},n}),v_h) = \langle \tilde{f}, v_h\rangle_{Y^*,Y} \quad \forall v_h \in Y_h. \tag{3.17}$$

Assumption 2.5 and Assumption 2.6 imply that (3.17) admits a unique solution $y_h(v_{\boldsymbol{\theta},n}) \in Y_h$ for all $v_{\boldsymbol{\theta},n} \in B_{\bar{r}^\varepsilon,\theta,n}(u^\varepsilon)$ with $n \geq n^*_\varepsilon$.

The adjoint equation is discretized similarly: Find $p_h(v_{\boldsymbol{\theta},n}) := p_h(y_h(v_{\boldsymbol{\theta},n})) \in Y_h$ such that

$$b_{\mathcal{L}^*}[v_{\boldsymbol{\theta},n}](p_h(v_{\boldsymbol{\theta},n}),v_h) = 2\tau_2\langle \bar{Q}_\delta v_{\boldsymbol{\theta},n} - \bar{Q}_\delta y_h(v_{\boldsymbol{\theta},n}), \bar{Q}_\delta v_h\rangle_H \quad \forall v_h \in Y_h. \tag{3.18}$$

Assumption 3.5 implies that (3.18) admits a unique solution $p_h(v_{\boldsymbol{\theta},n}) \in Y_h$ for $v_{\boldsymbol{\theta},n} \in B_{\bar{r}^\varepsilon,\theta,n}(u^\varepsilon)$.

Regarding both of the above equations the following error estimates hold true.

Theorem 3.6. *Suppose that Assumptions 2.4, 2.5, 2.6, 3.1 and 3.5 hold. Furthermore, let $Y \subset H^2(\Omega)$ and for all $u_1,u_2 \in B_{\bar{r}^\varepsilon}(u^\varepsilon)$ and $v,w \in Y$ it holds that*

$$|b_{\mathcal{L}}[u_1](v,w) - b_{\mathcal{L}}[u_2](v,w)| \leq C^\varepsilon\|u_1-u_2\|_V\|v\|_Y\|w\|_Y, \tag{3.19}$$

where $C^\varepsilon \in \mathbb{R}_+$. In addition, (3.19) holds also for the adjoint form $b_{\mathcal{L}^*}[u](\cdot,\cdot)$. Then

$$\lim_{h\to 0}\lim_{n\to\infty}\|y(u^\varepsilon) - y_h(u^\varepsilon_{\boldsymbol{\theta},n})\|_Y \leq C^\varepsilon_y\sqrt{\tau_2\mathcal{R}_\delta(y(u^\varepsilon),u^\varepsilon)}, \tag{3.20}$$

$$\lim_{h\to 0}\lim_{n\to\infty}\|p(u^\varepsilon) - p_h(u^\varepsilon_{\boldsymbol{\theta},n})\|_Y \leq C^\varepsilon_{ad}\sqrt{\tau_2\mathcal{R}_\delta(y(u^\varepsilon),u^\varepsilon)}, \tag{3.21}$$

where $p(u^\varepsilon) := p(y(u^\varepsilon)) \in Y$, $p_h(u^\varepsilon_{\boldsymbol{\theta},n}) := p_h(y_h(u^\varepsilon_{\boldsymbol{\theta},n})) \in Y_h$, $C^\varepsilon_y, C^\varepsilon_{ad} \in \mathbb{R}_+$, and the latter constants may depend on ε.

Proof. For $u^{\varepsilon}_{\boldsymbol{\theta},n} \in B_{\bar{r}^{\varepsilon},\theta,n}(u^{\varepsilon})$, we treat $b_{\mathcal{L}}[u^{\varepsilon}_{\boldsymbol{\theta},n}](\cdot,\cdot)$ as an approximation of $b_{\mathcal{L}}[u^{\varepsilon}](\cdot,\cdot)$ and apply Strang's lemma [56], Lemma 2.27 to get the estimate

$$\begin{aligned}\|y(u^{\varepsilon}) - y_h(u^{\varepsilon}_{\boldsymbol{\theta},n})\|_Y \leq & e^{y,rhs}_{n,h} + \inf_{v_h \in Y_h} \left(\left(1 + \frac{C^{\varepsilon}_b}{C^{\varepsilon}_c}\right) \|y(u^{\varepsilon}) - v_h\|_Y \right. \\ & \left. + \frac{1}{C^{\varepsilon}_c} \sup_{w_h \in Y_h} \frac{|b_{\mathcal{L}}[u^{\varepsilon}](v_h, w_h) - b_{\mathcal{L}}[u^{\varepsilon}_{\boldsymbol{\theta},n}](v_h, w_h)|}{\|w_h\|_Y} \right),\end{aligned} \tag{3.22}$$

where $e^{y,rhs}_{n,h}$ represents a discrepancy measure between the right-hand sides of (2.6) and (3.17), but $e^{y,rhs}_{n,h} = 0$ in our case. Let $\mathcal{I}_h y(u^{\varepsilon})$ be the interpolant of $y(u^{\varepsilon})$ in Y_h. Since $Y \subset H^2(\Omega)$, a well-known [55] interpolation bound yields

$$\inf_{v_h \in Y_h} \|y(u^{\varepsilon}) - v_h\|_Y \leq \|y(u^{\varepsilon}) - \mathcal{I}_h y(u^{\varepsilon})\|_Y \leq Ch \|y(u^{\varepsilon})\|_{H^2(\Omega)} \to 0 \quad \text{as} \quad h \to 0.$$

Assumption 2.5 and the boundedness of $\mathcal{I}_h : Y \to Y_h$ imply that $\|\mathcal{I}_h y(u^{\varepsilon})\|_Y \leq C$, where $C \in \mathbb{R}_+$ does not depend on h. Invoking (3.19) and (3.4), we get

$$|b_{\mathcal{L}}[u^{\varepsilon}](\mathcal{I}_h y(u^{\varepsilon}), w_h) - b_{\mathcal{L}}[u^{\varepsilon}_{\boldsymbol{\theta},n}](\mathcal{I}_h y(u^{\varepsilon}), w_h)| \leq \frac{C}{\sqrt{C^{\varepsilon}_s}} \sqrt{\tau_2 \mathcal{R}_{\delta}(y(u^{\varepsilon}), u^{\varepsilon})} \|w_h\|_Y$$

for $n \to \infty$. The result (3.20) then follows from (3.22) and the above estimates.

The analysis of the adjoint equation is similar to that of the state equation. However, for the application of Strang's lemma, the discrepancy

$$e^{p,rhs}_{n,h} := \sup_{v_h \in Y_h} \frac{|2\tau_2 \langle \bar{Q}_{\delta}(u^{\varepsilon} - u^{\varepsilon}_{\boldsymbol{\theta},n} + y_h(u^{\varepsilon}_{\boldsymbol{\theta},n}) - y(u^{\varepsilon})), \bar{Q}_{\delta} v_h \rangle_H|}{\|v_h\|_Y}$$

between the right-hand sides of (3.16) and (3.18) needs to be additionally estimated. The Cauchy–Schwarz and triangle inequalities, the embeddings $Y \hookrightarrow H$, $V \hookrightarrow H$, and $\bar{Q}_{\delta} \in L(H)$ yield $C \in \mathbb{R}_+$ such that

$$e^{p,rhs}_{n,h} \leq C \left(\|u^{\varepsilon} - u^{\varepsilon}_{\boldsymbol{\theta},n}\|_V + \|y(u^{\varepsilon}) - y_h(u^{\varepsilon}_{\boldsymbol{\theta},n})\|_Y \right).$$

Then, the result (3.21) follows from (3.4) and (3.20). □

Remark 3.7. The coarse-scale finite element approximation influences the upscaling consistency estimate in Theorem 2.12, but this effect can be controlled by choosing a suitable approximation scheme. From (2.14), the finite element discretization contributes an additional term:

$$\mathcal{R}_{\delta}(y_h(u^{\varepsilon}), u^{\varepsilon}) \leq 2\|\bar{Q}_{\delta} u^{\varepsilon} - y^0\|^2_H + 2\|y^0 - y(u^{\varepsilon})\|^2_H + 2\|y_h(u^{\varepsilon}) - y(u^{\varepsilon})\|^2_H.$$

For example, for elliptic problems with $H^2(\Omega)$ regularity, *i.e.*, when $Y \subset H^2(\Omega)$ and $H = L^2(\Omega)$, we obtain from the Aubin–Nitsche trick [55] the estimate $\|y_h(u^{\varepsilon}) - y(u^{\varepsilon})\|^2_H \leq (C^{\varepsilon} h^2)^2 \|\widetilde{f}\|^2_{Y^*}$, where C^{ε} is the coarse-scale stability constant. In practice, C^{ε} should be either independent of ε or uniformly bounded to allow for a moderate mesh size $h \gg \varepsilon$, thereby reducing computational cost while maintaining upscaling consistency.

Both FE discretized equations result in the following algebraic system:

$$\mathbb{B}_h[\boldsymbol{\theta}] \boldsymbol{y}_h = \mathbb{F}_h, \quad \mathbb{B}_h[\boldsymbol{\theta}]^{\top} \boldsymbol{p}_h = 2\tau_2 (\mathbb{P}_h[\boldsymbol{\theta}] - \mathbb{P}_h[\boldsymbol{y}_h]), \tag{3.23}$$

where $\boldsymbol{y}_h \in \mathbb{R}^{N_h}$ and $\boldsymbol{p}_h \in \mathbb{R}^{N_h}$ are the coefficients of the FE functions $y_h = \sum_{i=1}^{N_h} (\boldsymbol{y}_h)_i \phi_i$ and $p_h = \sum_{i=1}^{N_h} (\boldsymbol{p}_h)_i \phi_i$. Moreover, $\mathbb{B}_h[\boldsymbol{\theta}] \in \mathbb{R}^{N_h \times N_h}$, $(\mathbb{B}_h[\boldsymbol{\theta}])_{ij} := b_{\mathcal{L}}[v_{\boldsymbol{\theta},n}](\phi_i, \phi_j)$, $\mathbb{F}_h \in \mathbb{R}^{N_h}$, $(\mathbb{F}_h)_j := \langle \tilde{f}, \phi_j \rangle_{Y^*,Y}$ and $\mathbb{P}_h[\boldsymbol{\theta}] \in \mathbb{R}^{N_h}$, $(\mathbb{P}_h[\boldsymbol{\theta}])_j := \langle \bar{Q}_\delta v_{\boldsymbol{\theta},n}, \bar{Q}_\delta \phi_j \rangle_H$.

For item (iii) we assume that $H = L^2(\Omega)$, $Z = L^2(\partial\Omega)$, and apply a Monte-Carlo approach to approximate the integrals in the PINN objective as follows:

$$\mathcal{J}^M(v^M_{\boldsymbol{\theta},n}) := \frac{|\Omega|}{M_\Omega} \sum_{i=1}^{M_\Omega} \big(\mathcal{A}^\varepsilon v^M_{\boldsymbol{\theta},n}(x^r_i) - f^\varepsilon(x^r_i)\big)^2 + \frac{\tau_1 |\partial\Omega|}{M_{\partial\Omega}} \sum_{i=1}^{M_{\partial\Omega}} \big(\mathcal{B} v^M_{\boldsymbol{\theta},n}(x^b_i)\big)^2, \tag{3.24}$$

where $\{x^r_i\}_{i=1}^{M_\Omega}$ and $\{x^b_i\}_{i=1}^{M_{\partial\Omega}}$ are (uniformly random) collocation points in Ω and on $\partial\Omega$, respectively. Assuming that $M := M_\Omega + M_{\partial\Omega}$ is sufficiently large to guarantee that $v^M_{\boldsymbol{\theta},n} \in B_{\bar{r}^\varepsilon,\theta,n}(u^\varepsilon)$ for all $n \geq n^*_\varepsilon$, we discretize (2.16) using a quadrature rule:

$$\mathcal{R}^h_{\delta,D}(y_h(v^M_{\boldsymbol{\theta},n}), v^M_{\boldsymbol{\theta},n}) := \sum_{i=1}^{N_h} w^h_i \big((\bar{Q}^D_\delta v^M_{\boldsymbol{\theta},n})(x^h_i) - (\bar{Q}^D_\delta y_h(v^M_{\boldsymbol{\theta},n}))(x^h_i)\big)^2, \tag{3.25}$$

where $\{x^h_i,\, w^h_i\}_{i=1}^{N_h}$ are our finite element nodes and quadrature weights, and

$$(\bar{Q}^D_\delta v)(x^h_i) = \begin{cases} \frac{1}{|\mathcal{V}^D_\delta(x^h_i)|} \sum\limits_{x^h_j \in \mathcal{V}^D_\delta(x^h_i)} v(x^h_j), & x^h_i \in \Omega_\delta, \\ v(x^h_i), & x^h_i \in \Omega \setminus \Omega_\delta, \end{cases} \tag{3.26}$$

where $v \in C(\bar{\Omega})$ is a continuous function, $\mathcal{V}^D_\delta(x^h_i) := \{x^r_j : \ |x^r_j - x^h_i| \leq \frac{\delta}{2}\}$ corresponds to the set of residual collocation points around the coarse-scale mesh nodes, and M_D is an associated number of such collocation points. We note that using (2.10) allows us not to apply $\bar{Q}^D_\delta$ to the coarse-scale approximation in (3.25), yielding a simple implementation. The discrete components (3.24), (3.25) and (3.26), as well as the mapping (3.10), finally yield the approximation $\widehat{\mathfrak{J}}^{M,h}_D$ of $\widehat{\mathfrak{J}}$.

Algorithmically and assuming that $\nabla\widehat{\mathfrak{J}}^{M,h}_D(\boldsymbol{\theta})$ is a sufficiently accurate approximation of $\nabla\widehat{\mathfrak{J}}(\boldsymbol{\theta})$ (*i.e.* securing descent properties with respect to $\widehat{\mathfrak{J}}$ at $\boldsymbol{\theta}$) we use the discrete gradient, *e.g.*, in the Adam optimizer [57], to minimize $\widehat{\mathfrak{J}}^{M,h}_D(\boldsymbol{\theta})$. Next we summarize our overall computational procedure in Algorithm 1. Upon successful termination it produces neural network parameters and other outputs by setting $\hat{it} := it + 1$. In Algorithm 1,

Algorithm 1 Hybrid physics-informed NN training

Require: Tolerance tol, max. number of iterations $it_{\max}$, initial NN parameters $\boldsymbol{\theta}^{(0)}$, initial state FEM coefficients, $\boldsymbol{y}_h(\boldsymbol{\theta}^{(0)})$, optimizer hyperparameters.

Output: NN parameters $\hat{\boldsymbol{\theta}} := \boldsymbol{\theta}^{(\hat{it})}$, control variable $v^M_{\hat{\boldsymbol{\theta}},n} \approx u^{\varepsilon,M}_{\boldsymbol{\theta},n}$, state variable $\boldsymbol{y}_h(\hat{\boldsymbol{\theta}}) \in \mathbb{R}^{N_h}$.

1: **while** $0 \leq it \leq it_{\max} - 1$ or $\widehat{\mathcal{J}}^M(v^M_{\boldsymbol{\theta}^{(it)},n}) > tol$ **do**
2: Solve the adjoint system $\mathbb{B}^{\mathbf{ad}}_h[\boldsymbol{\theta}^{(it)}]\boldsymbol{p}_h = 2\tau_2(\mathbb{P}_h[\boldsymbol{\theta}^{(it)}] - \mathbb{P}_h[\boldsymbol{y}_h(\boldsymbol{\theta}^{(it)})])$
3: Compute $\nabla_\theta \widehat{J}^{M,h}_D(y_h(\boldsymbol{\theta}^{(it)}), \boldsymbol{\theta}^{(it)}) := \mathrm{grad}_{\boldsymbol{\theta}}(\widehat{J}^{M,h}_D(\boldsymbol{\theta}^{(it)}))$
4: Assemble the total gradient $\nabla\widehat{\mathfrak{J}}^{M,h}_D(\boldsymbol{\theta}^{(it)})$
5: Update weights $\boldsymbol{\theta}^{(it+1)} \leftarrow \mathrm{Optimizer}\big(\nabla\widehat{\mathfrak{J}}^{M,h}_D(\boldsymbol{\theta}^{(it)}),\ \text{optimizer hyperparameters}\big)$
6: Solve the state system $\mathbb{B}_h[\boldsymbol{\theta}^{(it+1)}]\boldsymbol{y}_h(\boldsymbol{\theta}^{(it+1)}) = \mathbb{F}_h$
7: $it \leftarrow it + 1$
8: **end while**

$\text{grad}_{\boldsymbol{\theta}}$ refers to automatic differentiation with respect to the NN parameters. The typical optimizer of choice also depends on hyperparameters. For the Adam optimizer, these include selecting a learning rate $lr \in \mathbb{R}_+$ or a schedule of learning rates $lr : \mathbb{N} \to \mathbb{R}_+$ with $it \mapsto lr(it)$, as well as specifying values for $\beta_1^{Ad} \in \mathbb{R}_+$ and $\beta_2^{Ad} \in \mathbb{R}_+$ for the moving average update, and setting a batch size. We initialize $\boldsymbol{\theta}^{(0)}$ using the Glorot scheme [58], but $\boldsymbol{\theta}^{(0)}$ can also be obtained by solving a neighboring, yet simpler problem [59, 60]. For an overview of standard NN optimization techniques, we refer exemplarily to [61] and the references therein.

4. An application in heat conduction

In this section, we focus on the example of heat conduction (1.5) with respect to both, the coarse and the fine scale, respectively. For this purpose let $u_{\boldsymbol{\theta},n}$ denote the *parameterization* of $\mathcal{L}\big[u_{\boldsymbol{\theta},n}\big]$. Specifically, we study our hybrid approach in view of its embedded upscaling process.

4.1. Upscaling-based parameterization

Our exemplary stationary heat transfer problem in $\bar{\Omega} = [0,1]^2$ is defined as follows:

$$-\nabla \cdot (\boldsymbol{K}^\varepsilon \nabla w^\varepsilon) = q, \quad \text{in } \Omega, \quad \text{and} \quad w^\epsilon = 0, \quad \text{on } \partial\Omega, \tag{4.1}$$

where w^ε is the temperature field, $q \in L^2(\Omega)$ is the source term, $\boldsymbol{K}^\varepsilon \in C^{0,1}(\bar{\Omega}, \mathbb{R}_+)$ is a Lipschitz continuous coefficient, and $\varepsilon \in \mathbb{R}_+$ is a small parameter indicating the fine scale length. In addition, it holds that

$$\exists \alpha,\ \beta \in \mathbb{R}_+ :\ \alpha \leq \boldsymbol{K}^\varepsilon(x) \leq \beta, \quad \forall x \in \bar{\Omega}, \tag{4.2}$$

We assume further that $\|\nabla \boldsymbol{K}^\varepsilon\|_{L^\infty(\Omega)} \leq \frac{c_{\boldsymbol{K}}}{\varepsilon}$, where $c_{\boldsymbol{K}} \in \mathbb{R}_+$ is ε-independent.

Solving (4.1) using finite elements can be computationally demanding due to the requirement $h \ll \varepsilon$ for a mesh size $h \in \mathbb{R}_+$ and $\varepsilon \ll 1$ in order to resolve the fine scale behaviour numerically. However, learning (4.1) with PINNs for $\varepsilon \ll 1$ also poses a challenge. We aim to efficiently characterize material properties of Ω and determine a computationally feasible coarse-scale (homogenized) counterpart of (4.1). In this vein, the concept of G-convergence is employed to formalize the notion of a homogenized equation and the related effective material; see, *e.g.*, [62].

Definition 4.1. A coefficient sequence $\{\boldsymbol{K}^\varepsilon(\cdot)\}$ is said to G-converge to $\boldsymbol{K}^*(\cdot)$ as $\varepsilon \to 0$, if for any $q \in H^{-1}(\Omega)$ the sequence of solutions $\{w^\varepsilon\}$ of (4.1) satisfies

$$w^\varepsilon \rightharpoonup w^0 \quad \text{in} \quad H_0^1(\Omega), \quad \boldsymbol{K}^\varepsilon \nabla w^\varepsilon \rightharpoonup \boldsymbol{K}^* \nabla w^0 \quad \text{in} \quad L^2(\Omega), \tag{4.3}$$

where w^0 is the solution to the homogenized equation

$$-\nabla \cdot (\boldsymbol{K}^* \nabla w^0) = q, \quad \text{in } \Omega, \quad \text{and} \quad w^0 = 0, \quad \text{on } \partial\Omega.$$

Various techniques exist to find the G-limit. Their respective applicability, however, depends on the specific problem and properties of the underlying medium. If such a G-limit exists, then it does not depend on q and on the boundary data on $\partial\Omega$. We also note that the existence of $\boldsymbol{K}^*$ for general heterogeneous media still remains an open problem. Often, the representative volume element (RVE) technique can be applied to find an approximation of $\boldsymbol{K}^*$. This approach is widely utilized in engineering applications; see, *e.g.*, [23, 24, 63, 64]. Here we state its equivalent characterization *via* weak convergence of gradients and fluxes in (4.3): For any measurable set $\mathcal{V} \subseteq \Omega$ of measure $|\mathcal{V}| > 0$ and $\langle \cdot \rangle_{\mathcal{V}} = \frac{1}{|\mathcal{V}|} \int_{\mathcal{V}} \cdot \, \mathrm{d}x$ (understood component-wise), one has

$$\lim_{\varepsilon \to 0} \langle \nabla w^\varepsilon \rangle_{\mathcal{V}} = \langle \nabla w^0 \rangle_{\mathcal{V}}, \quad \lim_{\varepsilon \to 0} \langle \boldsymbol{K}^\varepsilon \nabla w^\varepsilon \rangle_{\mathcal{V}} = \langle \boldsymbol{K}^* \nabla w^0 \rangle_{\mathcal{V}}. \tag{4.4}$$

The existence of the above limits for general heterogeneous media is difficult to verify. Thus, here we only assume that these limits exist; see also Assumption 2.7. Further, we introduce the following approximations:

$$\langle \nabla w^\varepsilon \rangle_{\mathcal{V}} \approx \lim_{\varepsilon \to 0} \langle \nabla w^\varepsilon \rangle_{\mathcal{V}} = \langle \nabla w^0 \rangle_{\mathcal{V}}, \quad \langle \boldsymbol{K}^\varepsilon \nabla w^\varepsilon \rangle_{\mathcal{V}} \approx \lim_{\epsilon \to 0} \langle \boldsymbol{K}^\varepsilon \nabla w^\varepsilon \rangle_{\mathcal{V}} \approx \widetilde{\boldsymbol{K}} \langle \nabla w^0 \rangle_{\mathcal{V}},$$

where $\widetilde{\boldsymbol{K}}$ is defined as follows.

Definition 4.2 (Upscaled coefficient). The upscaled coefficient $\widetilde{\boldsymbol{K}} := \widetilde{\boldsymbol{K}}(\varepsilon)$ satisfies

$$\langle \boldsymbol{K}^\varepsilon \nabla w^\varepsilon \rangle_{\mathcal{V}} = \widetilde{\boldsymbol{K}}(x) \langle \nabla w^\varepsilon \rangle_{\mathcal{V}}, \quad \text{for all } x \in \mathcal{V}, \tag{4.5}$$

where $\widetilde{\boldsymbol{K}}$ is a 2×2 tensor on $\mathcal{V}$ approximating $\boldsymbol{K}^*$ on $\mathcal{V}$ for a given $\varepsilon > 0$.

Note that $\widetilde{\boldsymbol{K}}$ is constant on $\mathcal{V}$ as a consequence of its definition. Besides, it is beneficial to consider $\widetilde{\boldsymbol{K}}$ as a tensor for anisotropic materials. Further we observe that in our concrete setting, two solutions of the fine-scale problem are required to obtain $\widetilde{\boldsymbol{K}}$ from (4.5). In contrast to $\boldsymbol{K}^*$, boundary conditions may then affect $\widetilde{\boldsymbol{K}}$. Of course, it is very desirable to identify a set of boundary conditions such that the dependence of $\widetilde{\boldsymbol{K}}$ on them is weak. In our application context, we choose the linear temperature drop boundary conditions $w_i^\varepsilon = x_i$ on $\partial\Omega$, $i = 1, 2$, but periodic boundary conditions or temperature drop no-flow conditions can be applied as well. Their related analyses are similar; compare [23, 25, 65]. The calculation of the upscaled thermal conductivity coefficient then leads to the following fine-scale problems: For $i \in \{1, 2\}$, find $w_i^\varepsilon : \Omega \to \mathbb{R}$ with

$$-\nabla \cdot (\boldsymbol{K}^\varepsilon \nabla w_i^\varepsilon) = q, \quad \text{in } \Omega, \quad \text{and} \quad w_i^\varepsilon = x_i, \quad \text{on } \partial\Omega. \tag{4.6}$$

We note that $q = 0$ is typically chosen to prevent any influence of the source term on $\widetilde{\boldsymbol{K}}$. In this case, we refer to [25] for error estimates between $\widetilde{\boldsymbol{K}}$ and $\boldsymbol{K}^*$. However, our multiscale solver is designed for computations of fine-scale solutions with $q \neq 0$, and setting $q = 0$ defines a special case of the NN-based upscaling scheme. Therefore, a more general form of (4.6) is considered, but the influence of q on $\widetilde{\boldsymbol{K}}$ is assumed to be rather weak. The upscaling process is described in Algorithm 2; *cf.* [25], Section 4.

Algorithm 2 Upscaling algorithm

Require: $\Omega = \bigcup_{j=1}^N \mathcal{V}_j$ with $\mathcal{V}_i \cap \mathcal{V}_j = \emptyset$, $i \neq j$, the solutions $\boldsymbol{w}^\varepsilon = \{w_1^\varepsilon, w_2^\varepsilon\}$ of (4.6).

Output: The upcaled coefficient $\widetilde{\boldsymbol{K}}$.

1: **for** $j \in \{1, \ldots, N\}$ **do**

2: Compute $\widetilde{\boldsymbol{F}}_{\boldsymbol{i}}^{\boldsymbol{j}} := \langle \boldsymbol{K}^\varepsilon \nabla w_i^\varepsilon \rangle_{\mathcal{V}_j} \in \mathbb{R}^2$, $\widetilde{\boldsymbol{T}}_{\boldsymbol{i}}^{\boldsymbol{j}} := \langle \nabla w_i^\varepsilon \rangle_{\mathcal{V}_j} \in \mathbb{R}^2$ for $i \in \{1, 2\}$.

3: Insert $\widetilde{\boldsymbol{F}}_{\boldsymbol{i}}^{\boldsymbol{j}}$ and $\widetilde{\boldsymbol{T}}_{\boldsymbol{i}}^{\boldsymbol{j}}$ into (4.5) and find $\widetilde{\boldsymbol{K}}^{\boldsymbol{j}} \in \mathbb{R}^{2\times 2}$ from the matrix equation:

$$\begin{pmatrix} \widetilde{K}_{11}^j & \widetilde{K}_{12}^j \\ \widetilde{K}_{21}^j & \widetilde{K}_{22}^j \end{pmatrix} \begin{pmatrix} (\widetilde{\boldsymbol{T}}_{\boldsymbol{1}}^{\boldsymbol{j}})_1 & (\widetilde{\boldsymbol{T}}_{\boldsymbol{2}}^{\boldsymbol{j}})_1 \\ (\widetilde{\boldsymbol{T}}_{\boldsymbol{1}}^{\boldsymbol{j}})_2 & (\widetilde{\boldsymbol{T}}_{\boldsymbol{2}}^{\boldsymbol{j}})_2 \end{pmatrix} = \begin{pmatrix} (\widetilde{\boldsymbol{F}}_{\boldsymbol{1}}^{\boldsymbol{j}})_1 & (\widetilde{\boldsymbol{F}}_{\boldsymbol{2}}^{\boldsymbol{j}})_1 \\ (\widetilde{\boldsymbol{F}}_{\boldsymbol{1}}^{\boldsymbol{j}})_2 & (\widetilde{\boldsymbol{F}}_{\boldsymbol{2}}^{\boldsymbol{j}})_2 \end{pmatrix}. \tag{4.7}$$

4: **end for**

5: Get $\widetilde{\boldsymbol{K}}$ with $\widetilde{\boldsymbol{K}}(x) = \widetilde{\boldsymbol{K}}^{\boldsymbol{j}}$ for $x \in \mathcal{V}_j$

Clearly, we need to ensure that the matrix of averaged gradients in (4.7) is invertible. This requires the partition $\{\mathcal{V}_j\}$ to be sufficiently heterogeneous with each $\mathcal{V}_j$ of "reasonable" size to prevent linear dependence between $\widetilde{\boldsymbol{T}}_{\boldsymbol{1}}^{\boldsymbol{j}}$ and $\widetilde{\boldsymbol{T}}_{\boldsymbol{2}}^{\boldsymbol{j}}$. Algorithm 2 is based on problem-dependent constitutive relations. In our case, we apply Fourier's law of heat conduction (4.7); *cf.* [64, 65]. In view of other applications, Darcy's law is used for porous

media flows [26, 66] as the relation between fluid velocity and pressure, and Hooke's law is applied [67] in elasticity to relate stress and strain fields.

For $\mathcal{V} = \Omega = (0,1)^2$, we derive a simplified formula for the upscaled coefficient; *cf.* [25]. For this purpose, it is convenient to study problem (4.6) as a problem with homogeneous Dirichlet boundary conditions. Let $u^{x_i} := x_i$ be the extension of our boundary conditions to Ω and consider the problem

$$-\nabla \cdot (\boldsymbol{K}^\varepsilon \nabla u_i^\varepsilon) = q + \nabla \cdot (\boldsymbol{K}^\varepsilon \nabla u^{x_i}), \quad \text{in } \Omega, \quad \text{and} \quad u_i^\varepsilon = 0, \quad \text{on } \partial\Omega. \tag{4.8}$$

Now, we apply Algorithm 2 to $\widetilde{\boldsymbol{F}}_{\boldsymbol{i}} := \langle \boldsymbol{K}^\varepsilon \nabla w_i^\varepsilon \rangle_\Omega$ and $\widetilde{\boldsymbol{T}}_{\boldsymbol{i}} := \langle \nabla w_i^\varepsilon \rangle_\Omega$. Since $w_i^\varepsilon = u_i^\varepsilon + u^{x_i}$ and $\nabla u^{x_i} = \boldsymbol{e}_i$, where $\boldsymbol{e}_i \in \mathbb{R}^2$ denotes the i-th unit vector, the divergence theorem yields

$$\langle \nabla w_i^\varepsilon \rangle_\Omega = \int_{\partial\Omega} u_i^\varepsilon \, \boldsymbol{\eta} \, \mathrm{d}s + \int_\Omega \boldsymbol{e}_i \, \mathrm{d}x = \boldsymbol{e}_i, \tag{4.9}$$

where $\boldsymbol{\eta}(x)$ is the outward unit normal at $x \in \partial\Omega$. Inserting (4.9) into (4.7), we get

$$\widetilde{\boldsymbol{K}}_{ij} = \int_\Omega \boldsymbol{K}^\varepsilon \partial_{x_j} w_i^\varepsilon \, \mathrm{d}x = \int_\Omega \boldsymbol{K}^\varepsilon (\partial_{x_j} u_i^\varepsilon + \delta_{ij}) \, \mathrm{d}x, \tag{4.10}$$

where δ_{ij} denotes the Kronecker delta symbol and $\partial_{x_j} := \frac{\partial}{\partial_{x_j}}$.

4.2. Analysis of the model equations

The upscaling process from Section 4.1 requires solving two equations at the fine scale to compute the upscaled coefficient for anisotropic media. These fine-scale problems can be approximated using two neural networks, as initially assumed in our two-scale hybrid approach. However, this setup would be technically more involved, though not conceptually, compared to our abstract framework. To keep the focus on conceptual development, we assume that certain coarse-scale information is available (*e.g.*, from a finite element simulation or a surrogate model), but $\widetilde{\boldsymbol{K}}_{11}$ may require further refinement.

Assumption 4.3. Assume that $u^\varepsilon := u_1^\varepsilon$ is the solution to problem (4.8). Let $\widetilde{\boldsymbol{K}}[u^\varepsilon] := \widetilde{\boldsymbol{K}} \in \mathbb{R}^{2\times 2}$, where $\widetilde{\boldsymbol{K}}_{12} = \widetilde{\boldsymbol{K}}_{21} \geq 0$ and $\widetilde{\boldsymbol{K}}_{22} \geq 0$ are fixed values, and $\widetilde{\boldsymbol{K}}[u^\varepsilon]_{11} = \int_\Omega \boldsymbol{K}^\varepsilon (\partial_{x_1} u^\varepsilon + 1) \, \mathrm{d}x$ is according to (4.10).

Assumption 4.3 implies that one needs to solve the model fine-scale problem (4.8) only for $i = 1$ to get $\widetilde{\boldsymbol{K}}_{11}$. We note that for periodic coefficients $\boldsymbol{K}^\varepsilon$ with period ε, $\widetilde{\boldsymbol{K}}_{12} = \widetilde{\boldsymbol{K}}_{21} = 0$ and $\widetilde{\boldsymbol{K}}[u^\varepsilon] = \widetilde{\boldsymbol{K}}[u^\varepsilon]_{11} I$, where $I \in \mathbb{R}^{2\times 2}$ is the identity.

Next, we show that the assumptions of Section 2 are satisfied by our fine- and coarse-scale model problems. For this purpose, let $X := C^k(\bar{\Omega})$ for some $k \geq 2$, $U := H^2(\Omega)$, $U_0 := H^2(\Omega) \cap H_0^1(\Omega)$ and $V := H^{1/2}(\Omega)$ be equipped with the norms

$$\|v\|_X = \sum_{|\mathfrak{a}|\leq 2} \sup_{x\in\Omega} |\partial_x^{\mathfrak{a}} v(x)|, \quad \|v\|_U = \|v\|_{U_0} = \Big(\sum_{|\mathfrak{a}|\leq 2} \|\partial_x^{\mathfrak{a}} v\|_{L^2(\Omega)}^2 \Big)^{1/2},$$
$$\|v\|_V = \left(\|v\|_{L^2(\Omega)}^2 + \int_\Omega \int_\Omega \frac{|v(x) - v(y)|^2}{|x-y|^{d+1}} \, \mathrm{d}x\mathrm{d}y \right)^{1/2},$$

respectively. Recall that X is dense in U and $\|u\|_U \leq |\Omega|^{1/2} \|u\|_X$ for $v \in X$, i.e, $X \hookrightarrow U$. Let $H := L^2(\Omega)$, $Z := L^2(\partial\Omega)$, $Y := H_0^1(\Omega)$, $Y^* = H^{-1}(\Omega)$ with $Y \hookrightarrow H \hookrightarrow Y^*$, where Y is compactly embedded into H by the Rellich–Kondrachov Theorem. We define the following linear operators

$$\mathcal{A}^\varepsilon : U \to H, \quad v \mapsto \mathcal{A}^\varepsilon v = -\nabla \cdot (\boldsymbol{K}^\varepsilon \nabla v), \text{ and } \mathcal{B} : U \to Z, \quad v \mapsto \mathcal{B}v = v|_{\partial\Omega}.$$

We cast our model fine-scale problem into the operator form

$$\mathcal{A}^\varepsilon u = f^\varepsilon, \quad \text{in } H, \quad \text{and} \quad \mathcal{B}u = 0, \quad \text{in } Z, \tag{4.11}$$

with $f^\varepsilon := q - \mathcal{A}^\varepsilon u^{x_1} \in L^2(\Omega) = H$. By the Lax–Milgram lemma, (4.11) has a unique solution $u^\varepsilon \in Y$, and we have $w_1^\varepsilon = u^\varepsilon + u^{x_1}$. The next proposition shows that $u^\varepsilon \in U$, thereby confirming Assumption 2.1, and that Assumption 2.4 is satisfied.

Proposition 4.4. *There exist $C_s^\varepsilon \in \mathbb{R}_+$ and $C_b^\varepsilon \in \mathbb{R}_+$ such that*

$$C_s^\varepsilon \|u\|_U^2 \le \|\mathcal{A}^\varepsilon u\|_H^2 \le C_b^\varepsilon \|u\|_U^2, \quad \forall u \in U_0, \tag{4.12}$$

where $C_s^\varepsilon \to 0$, $C_b^\varepsilon \to \infty$ as $\varepsilon \to 0$.

Proof. We multiply (4.11) (with $u = u^\varepsilon$) by u^ε and integrate by parts to get

$$\alpha \int_\Omega |\nabla u^\varepsilon|^2 \, \mathrm{d}x \le \int_\Omega \boldsymbol{K}^\varepsilon \nabla u^\varepsilon \cdot \nabla u^\varepsilon \, \mathrm{d}x = \int_\Omega q u^\varepsilon - \boldsymbol{K}^\varepsilon \partial_{x_1} u^\varepsilon \, \mathrm{d}x.$$

Using the Cauchy–Schwarz and Poincaré inequalities, we get

$$\|\nabla u^\varepsilon\|_H \le \frac{\beta + c_p \|q\|_H}{\alpha}, \tag{4.13}$$

where the Poincaré constant $c_p \in \mathbb{R}_+$ depends only on Ω, and β is according to (4.2). Note that problem (4.11) can be written as follows:

$$-\Delta u^\varepsilon = g^\varepsilon := \frac{f^\varepsilon + \nabla \boldsymbol{K}^\varepsilon \cdot \nabla u^\varepsilon}{\boldsymbol{K}^\varepsilon}, \quad \text{in } \Omega, \quad \text{and} \quad u^\varepsilon = 0, \quad \text{on } \partial\Omega.$$

Invoking a standard $H^2(\Omega)$ regularity result for convex domains [68], we get

$$\|u^\varepsilon\|_U \le \widehat{C} \|g^\varepsilon\|_H \le \frac{\widehat{C}}{\alpha} \big(\|q\|_H + (\|\nabla u^\varepsilon\|_H + 1) \|\nabla \boldsymbol{K}^\varepsilon\|_{L^\infty(\Omega)} \big). \tag{4.14}$$

where $\widehat{C} = \widehat{C}(\Omega) \in \mathbb{R}_+$. Using (4.13) and $\|\nabla \boldsymbol{K}^\varepsilon\|_{L^\infty(\Omega)} \le \frac{c_{\boldsymbol{K}}}{\varepsilon}$, we obtain the estimate

$$\|u^\varepsilon\|_U \le \frac{\widehat{C}}{\alpha} \big((1 + \frac{\beta}{\alpha}) \frac{c_{\boldsymbol{K}}}{\varepsilon} + (1 + \frac{c_{\boldsymbol{K}}}{\varepsilon} \frac{c_p}{\alpha}) \|q\|_H \big). \tag{4.15}$$

The bound (4.15) implies that $\|(\mathcal{A}^\varepsilon)^{-1}\| \le (C_1 + \frac{C_2}{\varepsilon}) < \infty$ for $\varepsilon > 0$ and some ε-independent $C_1, C_2 \in \mathbb{R}_+$, from where the lower bound in (4.12) readily follows. Using the Young and the Cauchy–Schwarz inequalities, for $u \in U$, we get

$$\|\mathcal{A}^\varepsilon u\|_H^2 \le 2 \big(\|\nabla \boldsymbol{K}^\varepsilon \cdot \nabla u\|_H^2 + \|\boldsymbol{K}^\varepsilon \Delta u\|_H^2 \big) \le \left(\frac{C}{\varepsilon^2} + \beta^2 \right) \|u\|_U^2, \tag{4.16}$$

where $C \in \mathbb{R}_+$ is ε-independent. This proves the upper bound in (4.12). □

The estimate (4.12) and interpolation techniques are used to verify the lower bound in Assumption 2.4 including $\|\mathcal{B}u\|_Z$, see [41], Lemma 4.1 and references therein. The upper bound in Assumption 2.4 follows from (4.16) and the trace inequality. These estimates are summarized in the following proposition.

Proposition 4.5. *There exist* $\bar{C}_s^\varepsilon \in \mathbb{R}_+$ *and* $C_b^\varepsilon \in \mathbb{R}_+$ *such that*

$$\bar{C}_s^\varepsilon \|u\|_V^2 \leq \|\mathcal{A}^\varepsilon u\|_H^2 + \|\mathcal{B}u\|_Z^2 \leq C_b^\varepsilon \|u\|_U^2, \quad \forall u \in U, \tag{4.17}$$

where $\bar{C}_s^\varepsilon \to 0$, $C_b^\varepsilon \to \infty$ *as* $\varepsilon \to 0$.

The next result [51], Theorem 3.3 shows that NNs with smooth activation functions are universal approximators of $C^k(\bar{\Omega})$, verifying Assumption 2.6 for $X := C^2(\bar{\Omega})$.

Proposition 4.6. *Suppose that* $\sigma \in C^\infty(\mathbb{R})$, $\sigma^{(s)}(0) \neq 0$ *for* $s = 0, 1, ...,$ *and* $\bar{\Omega} := [0,1]^d$. *If* $v \in C^k(\bar{\Omega})$, *then there exists an architecture* $\vec{\boldsymbol{n}}$ *with one hidden layer and* $v_{\boldsymbol{\theta},n} \in \mathfrak{N}_{\theta,n}$ *such that*

$$\sup_{x\in\Omega} |\partial_x^{\mathfrak{a}} v(x) - \partial_x^{\mathfrak{a}} v_{\boldsymbol{\theta},n}(x)| = \mathcal{O}\left(\frac{1}{n^{(k-|\mathfrak{a}|)/d}} \omega\big(\partial_x^{\mathfrak{b}} v, \frac{1}{n^{1/2}}\big)\right)$$

holds for all multi-indices $\mathfrak{a}, \mathfrak{b}$ *with* $|\mathfrak{a}| \leq k, |\mathfrak{b}| = k$, *where, for* $\delta_c > 0$, $\omega(v, \delta_c) = \sup\{|v(x) - v(y)| : |x - y| \leq \delta_c, x, y \in \bar{\Omega}\}$ *is the modulus of continuity of* v.

For example, the swish function $\sigma(x) = x\ \mathrm{sigmoid}(x)$ satisfies the prerequisites of Proposition 4.6, but the widely-used $\tanh(x)$ does not satisfy $\sigma^{(2)}(0) \neq 0$ rendering Proposition 4.6 non-applicable. The following result [45], Theorem 5.1 is then useful.

Proposition 4.7. *Suppose that* $\sigma(x) = \tanh(x)$ *and* $\bar{\Omega} := [0,1]^d$. *If* $v \in C^k(\bar{\Omega})$, *then there exists an architecture* $\vec{\boldsymbol{n}}$ *with two hidden layers and* $v_{\boldsymbol{\theta},n} \in \mathfrak{N}_{\theta,n}$ *such that* $\|v - v_{\boldsymbol{\theta},n}\|_{W^{m,\infty}(\bar{\Omega})} = \mathcal{O}(n^{-(k-m)})$ *holds*[1], *where* $\|w\|_{W^{m,\infty}(\bar{\Omega})} = \max\limits_{|\mathfrak{a}|\leq m} \sup\limits_{x\in\Omega} |\partial_x^{\mathfrak{a}} w(x)|$.

From the prerequisites of Proposition 4.7 we infer $X := C^k(\bar{\Omega})$ with $k > 2$ as $m = 2$ in our example. Clearly, $\mathfrak{N}_{\theta,n} \subset X$ and $\|v\|_X \leq C\|v\|_{W^{2,\infty}(\bar{\Omega})}$ for $v \in X$ and some $C \in \mathbb{R}_+$. The latter implies $X \subset \overline{\cup_n \mathfrak{N}_{\theta,n}}$, thereby verifying Assumption 2.6 for the hyperbolic tangent activation function.

Our model coarse-scale problem in its weak form reads: Find $y(u^\varepsilon) \in Y$ such that

$$b_{\mathcal{L}}[u^\varepsilon](y(u^\varepsilon), v) = \langle \tilde{f}, v\rangle_{Y^*,Y} \quad \forall v \in Y, \tag{4.18}$$

where $\tilde{f} := q + \nabla \cdot (\widetilde{\boldsymbol{K}}[u^\varepsilon] \nabla u^{x_1})$ and the forms are defined as follows:

$$b_{\mathcal{L}}[u^\varepsilon](w, v) := \int_\Omega \widetilde{\boldsymbol{K}}[u^\varepsilon] \nabla w \cdot \nabla v \ \mathrm{d}x, \quad \langle \tilde{f}, v\rangle_{Y^*,Y} = \int_\Omega qv \ \mathrm{d}x - \int_\Omega \widetilde{\boldsymbol{K}}[u^\varepsilon] \boldsymbol{e}_1 \cdot \nabla v \ \mathrm{d}x.$$

Since $\widetilde{\boldsymbol{K}}[u^\varepsilon]$ is a constant matrix, $\int_\Omega \partial_{x_j} v \ \mathrm{d}x = 0$ for $j \in \{1, 2\}$ and $v \in Y$, it holds:

$$\int_\Omega \widetilde{\boldsymbol{K}}[u^\varepsilon] \boldsymbol{e}_1 \cdot \nabla v \ \mathrm{d}x = \sum_{j=1}^{2} (\widetilde{\boldsymbol{K}}[u^\varepsilon])_{j1} \int_\Omega \partial_{x_j} v \ \mathrm{d}x = 0.$$

To study the coarse-scale problem (4.18), we need the following related map:

$$u \in U \mapsto \widetilde{\boldsymbol{K}}[u] \in \mathbb{R}^{2\times 2}, \quad u \mapsto \begin{pmatrix} \widetilde{\boldsymbol{K}}[u]_{11} & \widetilde{\boldsymbol{K}}_{12} \\ \widetilde{\boldsymbol{K}}_{21} & \widetilde{\boldsymbol{K}}_{22} \end{pmatrix}, \tag{4.19}$$

where $\widetilde{\boldsymbol{K}}[u]_{11} = \int_\Omega \boldsymbol{K}^\varepsilon (\partial_{x_1} u + 1) \ \mathrm{d}x$. Firstly, we prove the following.

[1] We state the asymptotic convergence rate, but explicit constants are estimated in [45].

Lemma 4.8. *The mapping $u \in U \mapsto \widetilde{\boldsymbol{K}}[u] \in \mathbb{R}^{2\times 2}$ defined by (4.19) is Lipschitz continuous, i.e., for $u_1, u_2 \in U$ we have*

$$\|\widetilde{\boldsymbol{K}}[u_1] - \widetilde{\boldsymbol{K}}[u_2]\| \leq |\Omega|^{1/2}\beta\|u_1 - u_2\|_U, \tag{4.20}$$

$$\|\widetilde{\boldsymbol{K}}[u_1] - \widetilde{\boldsymbol{K}}[u_2]\| \leq |\Omega|^{1/2}\varepsilon^{-1}c_{\boldsymbol{K}}\|u_1 - u_2\|_V, \tag{4.21}$$

with $\beta > 0$ from (4.2), and the Frobenius norm $\|\widetilde{\boldsymbol{K}}[u]\| := \big(\sum_{i,j=1}^{2} (\widetilde{\boldsymbol{K}}[u])_{ij}^2\big)^{\frac{1}{2}}$ of $\widetilde{\boldsymbol{K}}[u]$.

Proof. The Cauchy–Schwarz inequality and (4.2) yield

$$|(\widetilde{\boldsymbol{K}}[u_1])_{11} - (\widetilde{\boldsymbol{K}}[u_2])_{11}| \leq |\Omega|^{1/2}\beta\|\partial_{x_1}u_1 - \partial_{x_1}u_2\|_H \leq |\Omega|^{1/2}\beta\|u_1 - u_2\|_U. \tag{4.22}$$

In our case, $\widetilde{\boldsymbol{K}}_{12}$, $\widetilde{\boldsymbol{K}}_{21}$, $\widetilde{\boldsymbol{K}}_{22}$ are fixed, hence we get $\|\widetilde{\boldsymbol{K}}[u_1] - \widetilde{\boldsymbol{K}}[u_2]\| = |(\widetilde{\boldsymbol{K}}[u_1])_{11} - (\widetilde{\boldsymbol{K}}[u_2])_{11}|$ and (4.20) follows.

Next, we write $\widetilde{\boldsymbol{K}}[u]_{11} = \int_\Omega \boldsymbol{K}^\varepsilon - (\partial_{x_1}\boldsymbol{K}^\varepsilon)u \,\mathrm{d}x$. The Cauchy–Schwarz inequality, $\|\nabla \boldsymbol{K}^\varepsilon\|_{L^\infty(\Omega)} \leq \frac{c_{\boldsymbol{K}}}{\varepsilon}$ and the continuous embedding $V \hookrightarrow H$ imply that

$$\|\widetilde{\boldsymbol{K}}[u_1] - \widetilde{\boldsymbol{K}}[u_2]\| \leq \int_\Omega |\partial_{x_1}\boldsymbol{K}^\varepsilon(u_2 - u_1)|\mathrm{d}x \leq |\Omega|^{1/2}\varepsilon^{-1}c_{\boldsymbol{K}}\|u_1 - u_2\|_V.$$

□

We note that (4.21) readily implies (3.19). Next, we show that Assumption 2.5 holds, assuming that the influence of q is rather weak.

Proposition 4.9. *Suppose that $\|q\|_H \leq \frac{\alpha(\nu + \|\nabla u^\varepsilon\|_H^2)}{\|u^\varepsilon\|_H}$ for some $0 < \nu < 1$ and $\alpha \in \mathbb{R}_+$ from (4.2). Then, there exists $B_{\bar r}(u^\varepsilon) = \{v \,:\, \|u^\varepsilon - v\|_U \leq \bar r\} \subset U$ and $C_b, C_c \in \mathbb{R}_+$ such that (2.5) holds for all $v, w \in Y$ and $u \in B_{\bar r}(u^\varepsilon)$.*

Proof. Set $\boldsymbol{v}^\varepsilon = (u^\varepsilon, 0)$. Then $\nabla \cdot \boldsymbol{v}^\varepsilon = \partial_{x_1}u^\varepsilon$ and the divergence theorem yield

$$\int_\Omega \partial_{x_1}u^\varepsilon \,\mathrm{d}x = \int_\Omega \nabla \cdot \boldsymbol{v}^\varepsilon \,\mathrm{d}x = \int_{\partial\Omega} \boldsymbol{v}^\varepsilon \cdot \boldsymbol{\eta} \,\mathrm{d}x = 0, \tag{4.23}$$

since $u^\varepsilon|_{\partial\Omega} = 0$. From (4.10), we deduce

$$\widetilde{\boldsymbol{K}}[u^\varepsilon]_{11} = \boldsymbol{e}_1 \cdot \widetilde{\boldsymbol{K}}[u^\varepsilon]\boldsymbol{e}_1 = \langle \nabla w_1^\varepsilon \cdot \boldsymbol{K}^\varepsilon \nabla w_1^\varepsilon\rangle_\Omega - \langle \nabla u^\varepsilon \cdot \boldsymbol{K}^\varepsilon \nabla w_1^\varepsilon\rangle_\Omega. \tag{4.24}$$

Integrating the last term in (4.24) by parts gives

$$\int_\Omega \nabla u^\varepsilon \cdot \boldsymbol{K}^\varepsilon \nabla w_1^\varepsilon \mathrm{d}x = \int_{\partial\Omega} u^\varepsilon(\boldsymbol{K}^\varepsilon \nabla w_1^\varepsilon \cdot \boldsymbol{\eta}) \,\mathrm{d}s - \int_\Omega \nabla \cdot (\boldsymbol{K}^\varepsilon \nabla w_1^\varepsilon)u^\varepsilon \,\mathrm{d}x, \tag{4.25}$$

where $u^\varepsilon|_{\partial\Omega} = 0$ and $-\nabla \cdot (\boldsymbol{K}^\varepsilon \nabla w_1^\varepsilon) = q$ in Ω. Therefore, from (4.23) we obtain

$$\begin{aligned}\widetilde{\boldsymbol{K}}[u^\varepsilon]_{11} &= \langle \nabla w_1^\varepsilon \cdot \boldsymbol{K}^\varepsilon \nabla w_1^\varepsilon\rangle_\Omega - \int_\Omega qu^\varepsilon \,\mathrm{d}x \geq \alpha \int_\Omega |\nabla w_1^\varepsilon|^2 \mathrm{d}x - \int_\Omega qu^\varepsilon \,\mathrm{d}x \\ &= \alpha \int_\Omega (|\nabla u^\varepsilon|^2 + 2\partial_{x_1}u^\varepsilon + 1)\mathrm{d}x - \int_\Omega qu^\varepsilon \,\mathrm{d}x \geq \alpha(1 + \|\nabla u^\varepsilon\|_H^2) - \int_\Omega qu^\varepsilon \,\mathrm{d}x.\end{aligned}$$

The Cauchy–Schwarz inequality and our (amplitude) assumption on q give us

$$\widetilde{\boldsymbol{K}}[u^\varepsilon]_{11} \geq \alpha(1 + \|\nabla u^\varepsilon\|_H^2) - \|q\|_H \|u^\varepsilon\|_H \geq (1-\nu)\alpha =: C_\alpha.$$

Integration by parts, using the equation satisfied by u^ε, the Cauchy–Schwarz inequality, and (4.13) result in the estimate

$$|\widetilde{\boldsymbol{K}}[u^\varepsilon]_{11}| \leq \beta + \beta\|\nabla u^\varepsilon\|_H \leq \beta + \beta\left(\frac{\beta + c_p\|q\|_H}{\alpha}\right) =: C_\beta.$$

Let $0 \leq \bar{s} \leq \nu(1-\nu)\alpha$ and $\bar{r} = \frac{\bar{s}}{|\Omega|^{1/2}\beta}$. Then (4.20) implies that for $u \in B_{\bar{r}}(u^\varepsilon)$ it holds:

$$0 < (1-\nu)C_\alpha \leq \widetilde{\boldsymbol{K}}[u^\varepsilon]_{ii} - \bar{r} \leq \widetilde{\boldsymbol{K}}[u]_{11} \leq \widetilde{\boldsymbol{K}}[u^\varepsilon]_{11} + \bar{r} \leq C_\beta + \nu(1-\nu)\alpha. \tag{4.26}$$

Then, we set $C_c := (1-\nu)C_\alpha$ and $C_b := C_\beta + \nu(1-\nu)\alpha$ and apply the standard coercivity and continuity estimates to complete the proof. □

The Lax–Milgram lemma and Proposition 4.9 imply that for $u \in B_{\bar{r}}(u^\varepsilon)$ there exist a unique solution $y(u) \in Y$ of (4.18) with $\|y(u)\|_Y \leq C\|q\|_H$, where $C \in \mathbb{R}_+$ is independent of u, but also of ε, since $\bar{r}$, C_c and C_b are independent of ε in Proposition 4.9.

Proposition 4.9, together with either Proposition 4.6 or Proposition 4.7, yields an existence of $B_{\bar{r},\theta,n}(u^\varepsilon)$ for all $n > n_\varepsilon^*$, where $n_\varepsilon^* \in \mathbb{N}$ is chosen sufficiently large. Therefore, for all $v_{\boldsymbol{\theta},n} \in B_{\bar{r},\theta,n}(u^\varepsilon)$ with $n > n_\varepsilon^*$, there exists $y(v_{\boldsymbol{\theta},n}) \in Y$ and it holds:

$$\langle e(y(v_{\boldsymbol{\theta},n}), v_{\boldsymbol{\theta},n}), w\rangle_{Y^*,Y} = b_{\mathcal{L}}[v_{\boldsymbol{\theta},n}](y(v_{\boldsymbol{\theta},n}), w) - \langle \tilde{f}, w\rangle_{Y^*,Y} = 0 \quad \forall w \in Y.$$

Hence, we can cast our two-scale PDE-constrained optimization approach into the reduced form (3.3), which we state here for convenience. It reads as follows:

$$\inf \widehat{J}(v_{\boldsymbol{\theta},n}) := J(S(v_{\boldsymbol{\theta},n}), v_{\boldsymbol{\theta},n}) \ \text{ over } v_{\boldsymbol{\theta},n} \in B_{\bar{r},\theta,n}(u^\varepsilon), \tag{4.27}$$

where the hybrid objective $J : Y \times U \to \mathbb{R}_{\geq 0}$ includes both the PINN term and the coupling term:

$$\mathcal{J}(v_{\boldsymbol{\theta},n}) = \|\mathcal{A}^\varepsilon v_{\boldsymbol{\theta},n} - f^\varepsilon\|_H^2 + \tau_1\|\mathcal{B}v_{\boldsymbol{\theta},n}\|_Z^2, \quad \mathcal{R}_\delta(y(v_{\boldsymbol{\theta},n}), v_{\boldsymbol{\theta},n}) = \|\bar{Q}_\delta v_{\boldsymbol{\theta},n} - y(v_{\boldsymbol{\theta},n})\|_H^2. \tag{4.28}$$

The PINN objective in (4.28) includes the PDE residuals of (4.11), and the compression operator $\bar{Q}_\delta \in L(H)$ is defined by (2.9). The constant coefficient in (4.18) implies that $y(v_{\boldsymbol{\theta},n}) \in H^2(\Omega)$ for all $v_{\boldsymbol{\theta},n} \in B_{\bar{r},\theta,n}(u^\varepsilon)$. Consequently, we employ (2.10) rather than (2.16) as the coupling term, in accordance with the prerequisites of Theorem 2.12; see also Remark 2.8 regarding the regularity assumption for the homogenized solution y^0.

To apply Theorem 3.2 and Theorem 3.6 to (4.27), we need to verify Assumption 3.1 about the continuity of our fine-to-coarse scale map. For this purpose, motivated by Proposition 4.9, we introduce the following rather general assumption.

Assumption 4.10. For $q \in H$, $\boldsymbol{K}^\varepsilon \in C^{0,1}(\bar{\Omega})$, one can find $\bar{r}^\varepsilon \in \mathbb{R}_+$ and respective $B_{\bar{r}^\varepsilon}(u^\varepsilon) \subset U$ such that $y(u) \in Y$ exists and $\|y(u)\|_Y \leq C\|q\|_H$ for all $u \in B_{\bar{r}^\varepsilon}(u^\varepsilon)$, where $C \in \mathbb{R}_+$ is independent of u.

Proposition 4.11. *Suppose that Assumption 4.10 holds and $\{u_k^\varepsilon\} \subset B_{\bar{r}^\varepsilon}(u^\varepsilon) \cap X$ is an approximating sequence of u_ε. Then $S(u_k^\varepsilon) \rightharpoonup S(u^\varepsilon)$ in Y as $k \to \infty$, where $S : U \to Y$ is the fine-to-coarse scale map.*

Proof. Note that $\|y(u_k^\varepsilon)\|_Y \leq C\|q\|_H$ for $\{u_k^\varepsilon\}$, $C \in \mathbb{R}_+$ and $k \in \mathbb{N}$ as stated in Assumption 4.10. The reflexivity of Y and the Banach–-Alaoglu theorem imply that $\{y(u_k^\varepsilon)\}$ admits a weakly convergent subsequence with its

elements still denoted by $y(u_k^\varepsilon)$. Let $\hat{y} \in Y$ denote that weak limit. Rearranging terms in (4.18), we obtain

$$b_{\mathcal{L}}[u^\varepsilon](y(u^\varepsilon), v) + \int_\Omega \big(\widetilde{\boldsymbol{K}}[u_k^\varepsilon] - \widetilde{\boldsymbol{K}}[u^\varepsilon]\big)\nabla y(u^\varepsilon)\cdot\nabla v \,\mathrm{d}x = \int_\Omega qv \,\mathrm{d}x, \quad \forall v \in Y.$$

Note that $b_{\mathcal{L}}[u^\varepsilon](y(u_k^\varepsilon), v) \to b_{\mathcal{L}}[u^\varepsilon](\hat{y}, v)$ as $k \to \infty$ for all $v \in Y$, since $\mathcal{L}[u^\varepsilon] \in L(Y, Y^*)$ and hence it is weakly continuous. The continuity (4.20) and the Cauchy–Schwarz inequality yield

$$\int_\Omega \big(\widetilde{\boldsymbol{K}}[u_k^\varepsilon] - \widetilde{\boldsymbol{K}}[u^\varepsilon]\big)\nabla y(u_k^\varepsilon)\cdot\nabla v \,\mathrm{d}x \le C\|u_k^\varepsilon - u^\varepsilon\|_U \|q\|_H \|v\|_Y \to 0 \text{ as } k \to \infty,$$

since $C \in \mathbb{R}_+$ does not depend on u_k^ε and u^ε. This shows that $\hat{y} = y(u^\varepsilon)$. □

The continuity of $S' : B_{\bar{r}^\varepsilon}(u^\varepsilon) \subset U \to L(U, Y)$ can be established through standard techniques, which we briefly outline here by examining the sensitivity equation. The sensitivity $z := S'(u)h \in Y$ of $S(\cdot)$ at $u \in B_{\bar{r}^\varepsilon}(u^\varepsilon)$ in the direction $h \in U$ is given as the solution of the linearized state equation

$$\langle e_y(y(u), u)z, v\rangle_{Y^*,Y} = -\langle e_u(y(u), u)h, v\rangle_{Y^*,Y} \quad \forall v \in Y, \tag{4.29}$$

where the partial derivatives $e_y(y, u) : Y \to Y^*$ and $e_u(y, u) : U \to Y^*$ read

$$\langle e_y(y, u)w, v\rangle_{Y^*,Y} = \int_\Omega \widetilde{\boldsymbol{K}}[u]\nabla w\cdot\nabla v \,\mathrm{d}x, \quad \langle e_u(y, u)h, v\rangle_{Y^*,Y} = \int_\Omega \widetilde{\boldsymbol{K}}_u[h]\nabla y\cdot\nabla v \,\mathrm{d}x,$$

respectively, and $\widetilde{\boldsymbol{K}}_u[h] \in \mathbb{R}^{2\times 2}$ is given by $(\widetilde{\boldsymbol{K}}_u[h])_{11} = \int_\Omega \boldsymbol{K}^\varepsilon \partial_{x_1} h \,\mathrm{d}x$ and $(\widetilde{\boldsymbol{K}}_u[h])_{21} = (\widetilde{\boldsymbol{K}}_u[h])_{12} = (\widetilde{\boldsymbol{K}}_u[h])_{22} = 0$ with $\|\widetilde{\boldsymbol{K}}_u[h]\| \le |\Omega|^{1/2}\beta\|h\|_U$ for $h \in U$. Here, Proposition 4.9 or Assumption 4.10 implies that $e_y(y, u)$ is boundedly invertible for all $u \in B_{\bar{r}^\varepsilon}(u^\varepsilon)$, verifying Assumption 3.5 for the self-adjoint operator $\mathcal{L}[u]$. For $u_1, u_2 \in B_{\bar{r}^\varepsilon}(u^\varepsilon)$, let $z_1 = S'(u_1)h$ and $z_2 = S'(u_2)h$. Invoking the well-posedness of (4.29) and (4.20) while estimating the right-hand side of (4.29) provide us with the estimate

$$\|\big(S'(u_1) - S'(u_2)\big)h\|_Y \le C_{S'}\|u_1 - u_2\|_U\|h\|_U,$$

where $C_{S'} \in \mathbb{R}_+$ generally depends on u. Therefore, $S(u)$ is continuously Fréchet differentiable and Assumption 3.4 holds true.

4.3. Implementation issues

We provide the implementation details of Algorithm 1 for problem (4.27), where $\boldsymbol{K}^\varepsilon(x)$ in (4.11) is chosen to be a periodic coefficient. Firstly, the derivative $e_{\boldsymbol{\theta}}(y, \boldsymbol{\theta}) \in L(\mathbb{R}^{N_n}, Y^*)$ is given by

$$\langle e_{\boldsymbol{\theta}}(y, \boldsymbol{\theta})\boldsymbol{s}, v\rangle_{Y^*,Y} = \int_\Omega \widetilde{\boldsymbol{K}}_{\boldsymbol{\theta}}[\boldsymbol{s}]\nabla y\cdot\nabla v \,\mathrm{d}x, \quad \widetilde{\boldsymbol{K}}_{\boldsymbol{\theta}}[\boldsymbol{s}] = \int_\Omega \boldsymbol{K}^\varepsilon\langle\nabla_{\boldsymbol{\theta}}(\partial_{x_1} v_{\boldsymbol{\theta},n}), \boldsymbol{s}\rangle_{\mathbb{R}^{N_n}} \,\mathrm{d}x$$

for $\boldsymbol{s} \in \mathbb{R}^{N_n}$. Given the k-th unit vector $\boldsymbol{e}_k \in \mathbb{R}^{N_n}$, one obtains

$$\langle y'(\boldsymbol{\theta})^*\partial_y J(y(\boldsymbol{\theta}), \boldsymbol{\theta}), \boldsymbol{e}_k\rangle_{\mathbb{R}^{N_n}} = \langle e_{\boldsymbol{\theta}}(y(\boldsymbol{\theta}), \boldsymbol{\theta})\boldsymbol{e}_k, p\rangle_{Y^*,Y}. \tag{4.30}$$

The formula (4.30) is useful for the assembly of the first summand in (3.13). The algebraic FE systems are described in Section 2, but we note that $\mathbb{B}_h[\boldsymbol{\theta}] = \widetilde{\boldsymbol{K}}[v_{\boldsymbol{\theta},n}]\mathbb{A}_h$ with $(\mathbb{A}_h)_{ij} := \langle\nabla\phi_i, \nabla\phi_j\rangle_H$. The discrete

counterpart of (4.30) is given by

$$\langle e_{\boldsymbol{\theta}}(y_h(\boldsymbol{\theta}),\boldsymbol{\theta})\boldsymbol{e}_k, p_h\rangle_{Y^*,Y} = \boldsymbol{y}_h^T \mathbb{E}_h[\boldsymbol{\theta}_k]\, \boldsymbol{p}_h, \quad 1 \le k \le N_n,$$

where $\mathbb{E}_h[\boldsymbol{\theta}_k] \in \mathbb{R}^{N_h \times N_h}$, $(\mathbb{E}_h[\boldsymbol{\theta}_k])_{ij} := \langle e_{\boldsymbol{\theta}}(\phi_i,\boldsymbol{\theta})\boldsymbol{e}_k, \phi_j\rangle_{Y^*,Y}$. We get $(\mathbb{E}_h[\boldsymbol{\theta}_k])_{ij} = (\widetilde{\boldsymbol{k}}_M[\boldsymbol{\theta}])_k (\mathbb{A}_h)_{ij}$ where $\widetilde{\boldsymbol{k}}_M[\boldsymbol{\theta}] \in \mathbb{R}^{N_n}$ represents approximations of $\widetilde{\boldsymbol{K}}_{\boldsymbol{\theta}}[\boldsymbol{e}_k]$ by the Monte-Carlo approach and using the residual collocation points

$$(\widetilde{\boldsymbol{k}}_M[\boldsymbol{\theta}])_j = \frac{1}{M_\Omega}\sum_{i=1}^{M_\Omega} \boldsymbol{K}^\varepsilon(x^r_{1,i}, x^r_{2,i}) \frac{\partial^2 v^M_{\boldsymbol{\theta},n}(x^r_{1,i}, x^r_{2,i})}{\partial \boldsymbol{\theta}_j\, \partial x_1}, \quad 1 \le j \le n.$$

We also use a multiscale Fourier feature network [10, 28] to reduce the spectral bias effect. This architecture includes Fourier feature mappings $\mathcal{F}^{(k)} : \mathbb{R} \to \mathbb{R}^{2m}$:

$$\mathcal{F}^{(k)}(x) = \left(\cos(2\pi \boldsymbol{B}^{(k)} x),\ \sin(2\pi \boldsymbol{B}^{(k)} x)\right), \quad 1 \le k \le K,$$

where each entry of $\boldsymbol{B}^{(k)} \in \mathbb{R}^{m\times d}$ is sampled from a Gaussian distribution $\mathcal{N}(0, \varrho_k^2)$ with $\varrho_k > 0$ a specified hyperparameter. These features are used as inputs for the hidden layers, which are defined for $1 \le k \le K$ and $2 \le l \le L-1$ as follows:

$$z_1^{(k)} = \sigma(W_1 \mathcal{F}^{(k)}(x) + b_1),\ z_l^{(k)} = \sigma(W_l z_{l-1}^{(k)} + b_l).$$

Next, the above outputs are concatenated within the linear layer as follows:

$$v^M_{\boldsymbol{\theta},n} = W_L\left[z_L^{(1)}, ..., z_L^{(K)}\right] + b_L,$$

where W_L and b_L are the weights and biases of the output layer. To enforce the boundary conditions exactly [31, 42, 43], one may set $v^{l,M}_{\boldsymbol{\theta},n} := l(x) v^M_{\boldsymbol{\theta},n}$, where l is the corresponding signed distance function for $\partial\Omega$.

4.4. Numerical results

For our numerical experiments, we use the following setup unless otherwise specified. The multiscale Fourier feature network (Ms-PINN) is used as the main architecture of choice, with the two Fourier features initialized by $\varrho_1 = 1$ and $\varrho_2 = 1/\varepsilon$, two hidden layers with 128 neurons each, and $\tanh(x)$ as the activation functions. The full batch is used for training with the Adam algorithm [57], and $\beta_1^{Ad} = 0.9$ and $\beta_2^{Ad} = 0.999$ are chosen as the hyperparameters. The learning rate for the exponential learning rate schedule is initialized as $lr(0) = 5e-4$, with a decay-rate of 0.75 every 1000 training iterations. For PINNs and our hybrid approach, we use automatic loss balancing [53], Algorithm 1(c) at every 100 steps. The latter guarantees that the norm of gradients of each weighted loss term is equal to each other. We choose (2.10) for coupling, the uniform quadrature rule for (3.25), and $\delta = \varepsilon$ for (3.26). To reduce computational costs, we update the coarse-scale information only if $|\widetilde{\boldsymbol{K}}[v^M_{\boldsymbol{\theta}^{(it+T_K)},n}] - \widetilde{\boldsymbol{K}}[v^M_{\boldsymbol{\theta}^{(it)},n}]| > 10^{-2}$, where $T_K = 100$. Indeed, in our first-order optimization setting, the information transferred *via* the fine-to-coarse scale map changes little between Adam iterations. Thus, frequent coarse-scale updates are unnecessary and costly. Our implementation utilizes JAX [69] for the neural network component and FEniCS Dolphin [70] for the finite element component.

The following example is used to illustrate the concept of upscaling consistency:

$$-\partial_x(\boldsymbol{K}^\varepsilon \partial_x u^\varepsilon) = f^\varepsilon, \quad \text{in } \Omega := (0,1), \quad \text{and} \quad u^\varepsilon(0) = 0,\ \ u^\varepsilon(1) = 0, \tag{4.31}$$

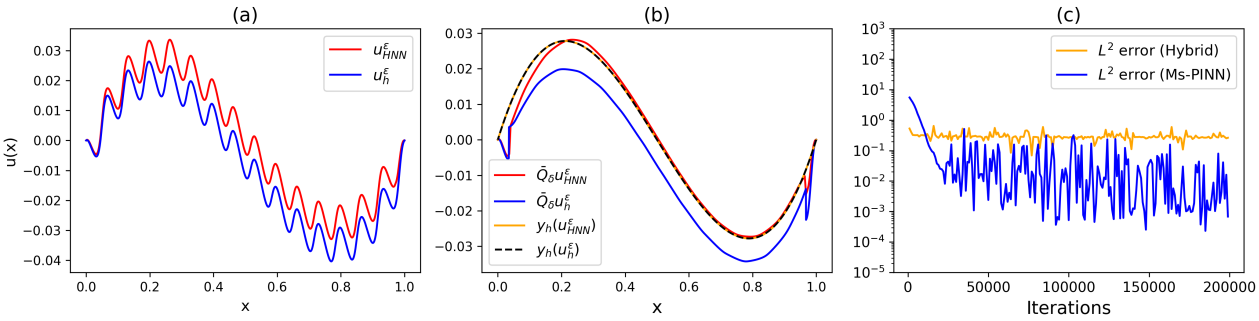

FIGURE 1. Results for $\varepsilon = 1/15$, 1D problem: **(a)** Hybrid fine-scale solution u^ε_{HNN} and FEM solution u^ε_h, **(b)** Predicted state $y_h(u^\varepsilon_{HNN})$, true discrete state $y_h(u^\varepsilon_h)$, compression $\bar{Q}_\delta u^\varepsilon_{HNN}$ and $\bar{Q}_\delta u^\varepsilon_h$, **(c)** Relative L^2 error of the hybrid solver and Ms-PINN vs iterations.

where $\boldsymbol{K}^\varepsilon(x) = 2 + \sin(2\pi\varepsilon^{-1}x)$, $f^\varepsilon := q + \partial_x \boldsymbol{K}^\varepsilon$ and $q = -3(2x-1)$. Then, $\widetilde{\boldsymbol{K}}[u] = \int_0^1 \boldsymbol{K}^\varepsilon(\partial_x u + 1)\,\mathrm{d}x$ and $\boldsymbol{K}^* = (\int_0^1 \frac{\mathrm{d}x}{\boldsymbol{K}^\varepsilon(x)})^{-1} = \sqrt{3} \approx 1.732$, see [62], Section 1.2 for the latter. For $\varepsilon = 1/15$ and $\varepsilon = 0.01$, we get $\widetilde{\boldsymbol{K}}[u^\varepsilon_h] \approx 1.733$ and $\widetilde{\boldsymbol{K}}[u^\varepsilon_h] \approx 1.732$, where u^ε_h are computed using finite elements on equidistant meshes with $h = 5 \cdot 10^{-4}$ and $h = 2 \cdot 10^{-4}$, respectively (these meshes are further used for validation). The analytical solution of the coarse-scale problem can in principle be computed:

$$y(u^\varepsilon_h)(x) = \frac{1}{\widetilde{\boldsymbol{K}}[u^\varepsilon_h]}\big(x^3 - \frac{3}{2}x^2 + \frac{1}{2}x\big). \tag{4.32}$$

For $y_h(u^\varepsilon_h)$, we choose $h = 4 \cdot 10^{-3}$ so that $\frac{\|y(u^\varepsilon_h) - y_h(u^\varepsilon_h)\|_{L^2}}{\|y(u^\varepsilon_h)\|_{L^2}} < 5 \cdot 10^{-4}$. Then, the influence caused by an upscaling method in Assumption 2.9 is negligible in the context of Theorem 2.12, as well as the influence caused by coarse-scale finite element discretization, see also Remark 3.7. Setting $it_{\max} = 2 \cdot 10^5$, we use $M_\Omega = 2000$ and $M_\Omega = 5000$ uniformly random collocation points for $\varepsilon = 1/15$ and $\varepsilon = 0.01$ (both with hard boundary constraints), respectively, and the mesh size $h = 0.004$ for the coarse-scale problem. We use $l(x) = 4x(1-x)$ for hard boundary constraints imposition.

If ε is not small enough, one expects an upscaling consistency gap between $y_h(u)$ and $\bar{Q}_\delta u$ for $u = u^\varepsilon_h$, and this gap is typically overfitted by a neural network $u = v^M_{\boldsymbol{\theta},n}$. This overfitting can introduce discrepancies in the PDE residual of (3.1). The latter can impede the route to convergence on a fine scale, rendering it generally not possible. Indeed, for $\varepsilon = 1/15$ in (4.31), we observe in Fig 1c that the relative L^2 error of the hybrid solver oscillates around 0.3 on our validation set, and does not drop significantly below that value, compared to the standard Ms-PINN approach without coarse-scale constraints, which shows $\sim 10^{-2}$ accuracy. We can see in Fig 1b that $y_h(v^M_{\boldsymbol{\theta},n})$ matches $y_h(u^\varepsilon_h)$ well, but $\bar{Q}_\delta v^M_{\boldsymbol{\theta},n}$ is much closer to $y_h(v^M_{\boldsymbol{\theta},n})$ compared to $\bar{Q}_\delta u^\varepsilon_h$, thus violating the upscaling consistency and causing a significant approximation error in Fig 1a.

Theorem 2.12 shows that the upscaling consistency gap decreases with decreasing ε, thus softening the conflict between the PDE residual term and the coupling term in the optimization and reducing the effect of coarse-scale overfitting. Indeed, we observe a smaller gap between $y_h(u^\varepsilon_h)$ and $\bar{Q}_\delta u^\varepsilon_h$ for $\varepsilon = 0.01$ in Fig 2b, and a lower resulting L^2 error for the hybrid method in Fig 2c, compared to the previous example. In addition, the hybrid solver outperforms the Ms-PINN approach by almost two orders of magnitude in terms of relative error. However, the hybrid solver error still oscillates around 0.1: there is a visible overfitting of the coarse-scale component at $x \approx 0.2$ and $x \approx 0.8$ in Fig 2b, which causes visible approximation errors in Fig 2a at the same spatial points for the fine-scale hybrid approximation.

Next, we study the 2D problem (4.11) with the isotropic coefficient $\boldsymbol{K}^\varepsilon(x) = 3 + \sin(\pi\varepsilon^{-1}x_1) + \sin(2\pi\varepsilon^{-1}x_2)$ with $\varepsilon = 0.1$ and $q(x) = 10\exp(-100\|x - 0.5\|_2^2)$. We note that $\boldsymbol{K}^\varepsilon$ is periodic, hence we need to determine $\widetilde{\boldsymbol{K}}_{11}$ only. In addition to the Ms-PINN, we employ the multilayer perceptron (MLP) PINN with 8 hidden layers

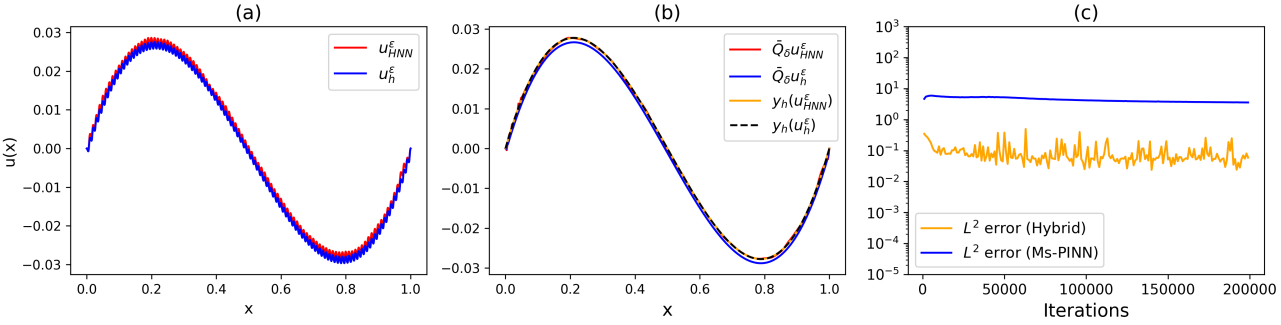

FIGURE 2. Results for $\varepsilon = 0.01$, 1D problem: same description as for Figure 1.

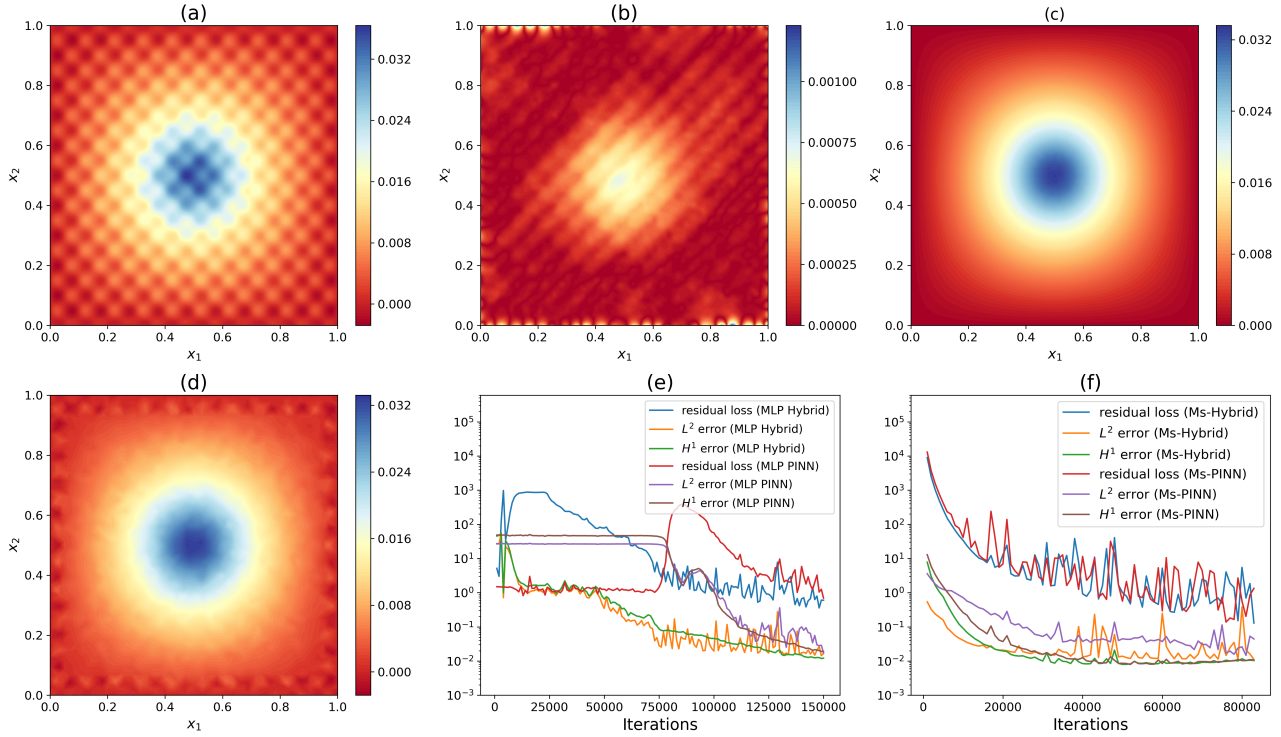

FIGURE 3. Results for $\varepsilon = 0.1$, 2D problem: **(a)** Hybrid MLP-based fine-scale solution u^{ε}_{HNN}, **(b)** Pointwise error $|u^{\varepsilon}_{HNN}(x) - u^{\varepsilon}_{h}(x)|$, **(c)** Predicted state $y_h(u^{\varepsilon}_{HNN})$, **(d)** Compression $\bar{Q}_\delta u^{\varepsilon}_{HNN}$, **(e)** The PDE residual losses and relative L^2 and H^1 errors vs iterations for the MLP-based hybrid solver and the MLP PINN, **(f)** The PDE residual losses and relative L^2 and H^1 errors vs iterations for the Ms-PINN-based hybrid solver and the Ms-PINN.

with 64 neurons each for both the standard and hybrid approaches while setting $lr(0) = 1e-3$ for the initial learning rate. We use $M_\Omega = 25\,000$ and $M_{\partial\Omega}$=2500 uniformly random collocation points and the mesh size $h \approx 0.04$ for the coarse-scale discretization. Figure 3 shows exemplary realizations of the hybrid solver. The MLP network, prone to low-frequency bias, benefits from embedding coarse-scale constraints, which accelerates the approximation of the coarse-scale component and mitigates the bias, as seen in Fig 3e. The MLP approach struggles with being trapped at poor local minima, failing to provide a reasonable approximation for around 75k

TABLE 1. Comparison of the Ms-PINNs and the MLP-based PINNs with the hybrid method for the 2D problem. The quantities are averages over 5 random initializations of NN parameters and collocation points. For the MLP PINN only one realization is presented, as all others failed to deliver reasonable results. All runs are executed on A100 GPUs with 40 GB RAM.

Method	Rel. L^2 error	Rel. H^1 error	iterations	time (s)
Hybrid Ms-PINN	1.28E-02 ± 0.003	2.23E-02 ± 0.008	77 478	554
Hybrid MLP	1.42E-02 ± 0.001	1.44E-02 ± 0.008	139 400	1725
Ms-PINN	1.97E-02 ± 0.005	2.20E-02 ± 0.005	70 009	354
MLP PINN	1.71E-02 ± 1.000	1.89E-02 ± 1.000	150 000	1221

iterations. In contrast, the hybrid approach exhibits steadily decreasing residuals and errors. For the Ms-PINNs, the hybrid method also shows faster decay of the L^2 and H^1 relative errors[2], though the differences are less pronounced compared to the MLP approach, as the Ms-PINN approach is less prone to the spectral bias.

In Table 1, we report the L^2 and H^1 relative errors corresponding to the smallest residual norm value within the learning history, the number of iterations required to achieve it, and the compute time, all averaged over 5 random initializations. According to our selection criteria, the hybrid Ms-PINN requires slightly more iterations than the standard Ms-PINN, but in Fig 3f the convergence to the reported values in Table 1 is factually noticeably faster for the hybrid approach, suggesting more robustness at the cost of compute time. For the standard MLP PINN, one successful realization was obtained for Fig 3e, while many others failed to approximate the solution. The hybrid MLP network consistently yields robust results across all realizations.

Finally, we note that in our 2D setting, computing the coarse-scale solution using finite elements takes approximately 0.01 seconds, whereas computing gradients and performing Adam updates in the PINN approach takes around 0.02 seconds. Therefore, first-order hybrid optimization using Adam is generally more computationally expensive compared to standard PINNs, as it requires many iterations along with frequent coarse-scale updates. However, this may change significantly with second-order optimization, where the cost of the optimization update can substantially exceed that of the coarse-scale finite element solver.

5. CONCLUSION

This paper focuses on the structural properties of a learning-informed PDE-constrained optimization problem with a PINN-based objective, leading to a hybrid two-scale solver. Our approach incorporates a stabilizing constraint in the form of a surrogate PDE model for the PINN objective, enhancing learning in challenging multiscale or multi-fidelity settings. We further apply this framework to multiscale problems by using coarse-scale constraints inspired by multiscale modeling. However, selecting a suitable neural network approximation scheme and developing an efficient optimization algorithm, aimed at enhancing accuracy while minimizing computational time, are both essential for further algorithmic developments. This task requires consideration of recent advancements in PINNs and the field of PDE-constrained optimization. For example, first-order parametric optimization methods are largely agnostic to the infinite-dimensional aspects of the PDE problem and require many iterations, as also evident in the numerical section. These iterations, in turn, necessitate frequent solutions of coarse-scale PDEs in the hybrid setting, thereby increasing the overall cost. To reduce the number of iterations and hence the overall costs, one may employ second-order optimization algorithms that better exploit the PDE structure and the underlying function space geometry. However, ensuring proper scalability in second-order settings, particularly when working with large neural networks and a high density of collocation points, as commonly encountered in multiscale problems, is then crucial for the hybrid solver. One may also consider using non-smooth activation functions, introducing non-smoothness into the related PDE-constrained optimization. Such a setting challenges both the derivation of suitable stationarity conditions for the PDE-constrained multiscale approach and its numerical solution. A comprehensive investigation of such a setting,

[2]The H^1 norm is induced by the L^2 norm of the gradient.

as well as the development of efficient optimization algorithms to address other problems, possibly including nonlinear ones, remain part of our future work in this area.

Acknowledgments

The authors acknowledge the support of the Leibniz Collaborative Excellence Cluster under project ML4Sim (funding reference: K377/2021). MH also acknowledges support by the DFG ExC 2046 MATH+: Berlin Mathematics Research Center under project EF1-17.

Data availability statement

The research data associated with this article are included in the article.

References

[1] S. Cuomo, V.S. Di Cola, F. Giampaolo, G. Rozza, M. Raissi and F. Piccialli, Scientific machine learning through physics–informed neural networks: where we are and what's next. *J. Sci. Comput.* **92** (2022) 88.

[2] I.E. Lagaris, A. Likas and D.I. Fotiadis, Artificial neural networks for solving ordinary and partial differential equations. *IEEE Trans. Neural Netw.* **9** (1998) 987–1000.

[3] M. Raissi, P. Perdikaris and G.E. Karniadakis, Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Computat. Phys.* **378** (2019) 686–707.

[4] Z. Hu, K. Shukla, G.E. Karniadakis and K. Kawaguchi, Tackling the curse of dimensionality with physics-informed neural networks. *Neural Netw.*, **176** (2024) 106369

[5] Y. Zang, G. Bao, X. Ye and H. Zhou, Weak adversarial networks for high-dimensional partial differential equations. *J. Computat. Phys.* **411** (2020) 109409.

[6] A.D. Jagtap and G.E. Karniadakis, Extended physics-informed neural networks (XPINNs): a generalized space-time domain decomposition based deep learning framework for nonlinear partial differential equations, in *AAAI Spring Symposium: MLPS* (2021) 2002–2041.

[7] B. Moseley, A. Markham and T. Nissen-Meyer, Finite basis physics-informed neural networks (fbpinns): a scalable domain decomposition approach for solving differential equations. *Adv. Computat. Math.* **49** (2023) 62.

[8] K. Shukla, A.D. Jagtap and G.E. Karniadakis, Parallel physics-informed neural networks *via* domain decomposition. *J. Computat. Phys.* **447** (2021) 110683.

[9] S. Wang, Y. Teng and P. Perdikaris, Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM J. Sci. Comput.* **43** (2021) A3055–A3081.

[10] S. Wang, H. Wang and P. Perdikaris, On the eigenvector bias of Fourier feature networks: From regression to solving multi-scale PDEs with physics-informed neural networks. *Comput. Methods Appl. Mech. Eng.* **384** (2021) 113938.

[11] S. Wang, X. Yu and P. Perdikaris, When and why PINNs fail to train: a neural tangent kernel perspective. *J. Computat. Phys.* **449** (2022) 110768.

[12] A. Abdulle, E. Weinan, B. Engquist and E. Vanden-Eijnden, The heterogeneous multiscale method. *Acta Numerica* **21** (2012) 1–87.

[13] Y. Efendiev, J. Galvis and T.Y. Hou, Generalized multiscale finite element methods (GMsFEM). *J. Computat. Phys.* **251** (2013) 116–135.

[14] O. Iliev, R. Lazarov and J. Willems, Variational multiscale finite element method for flows in highly porous media. *Multiscale Model. Simul.* **9** (2011) 1350–1372.

[15] A. Målqvist and D. Peterseim, Localization of elliptic multiscale problems. *Math. Computat.* **83** (2014) 2583–2603.

[16] E. Weinan, B. Engquist and Z. Huang, Heterogeneous multiscale method: a general methodology for multiscale modeling. *Phys. Rev. B* **67** (2003) 092101.

[17] H. Arbabi, J.E. Bunder, G. Samaey, A.J. Roberts and I.G. Kevrekidis, Linking machine learning with multiscale numerics: data-driven discovery of homogenized equations. *JOM* **72** (2020) 4444–4457.

[18] J. Han and Y. Lee, A neural network approach for homogenization of multiscale problems. *Multiscale Model. Simul.* **21** (2023) 716–734.

[19] F. Kröpfl, R. Maier and D. Peterseim, Neural network approximation of coarse-scale surrogates in numerical homogenization. *Multiscale Model. Simul.* **21** (2023) 1457–1485.

[20] W.T. Leung, G. Lin and Z. Zhang, NH-PINN: neural homogenization-based physics-informed neural network for multiscale problems. *J. Computat. Phys.* **470** (2022) 111539.

[21] J.S. Richard Park and X. Zhu, Physics-informed neural networks for learning the homogenized coefficients of multiscale elliptic equations. *J. Computat. Phys.* **467** (2022) 111420.

[22] Z. Zhang, C. Moya, W.T. Leung, G. Lin and H. Schaeffer, Bayesian deep operator learning for homogenized to fine-scale maps for multiscale PDE. *Multiscale Model. Simul.* **22** (2024) 956–972.

[23] L.J. Durlofsky, Numerical calculation of equivalent grid block permeability tensors for heterogeneous porous media. *Water Resources Res.* **27** (1991) 699–708.

[24] C. Farmer, Upscaling: a review. *Int. J. Numer. Methods Fluids* **40** (2002) 63–78.

[25] X.-H. Wu, Y. Efendiev and T.Y. Hou, Analysis of upscaling absolute permeability. *Discrete Continuous Dyn. Syst. B* **2** (2002) 185.

[26] M. Griebel and M. Klitz, Homogenization and numerical simulation of flow in geometries with textile microstructures. *Multiscale Model. Simul.* **8** (2010) 1439–1460.

[27] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio and A. Courville, On the spectral bias of neural networks, in *International Conference on Machine Learning*. PMLR (2019) 5301–5310.

[28] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron and R. Ng, Fourier features let networks learn high frequency functions in low dimensional domains. *Adv. Neural Inform. Process. Syst.* **33** (2020) 7537–7547.

[29] G. Dong, M. Hintermüller and K. Papafitsoros, Optimization with learning-informed differential equation constraints and its applications. *ESAIM: Control Optim. Calc. Var.* **28** (2022) 3.

[30] J. Sirignano, J. MacArt and K. Spiliopoulos, PDE-constrained models with neural network terms: optimization and global convergence. *J. Computat. Phys.* **481** (2023) 112016.

[31] L. Lu, R. Pestourie, W. Yao, Z. Wang, F. Verdugo and S.G. Johnson, Physics-informed neural networks with hard constraints for inverse design. *SIAM J. Sci. Comput.* **43** (2021) B1105–B1132.

[32] S. Mowlavi and S. Nabi, Optimal control of PDEs using physics-informed neural networks. *J. Computat. Phys.* **473** (2023) 111731.

[33] E. Casas, Optimal control in coefficients of elliptic equations with state constraints. *Appl. Math. Optim.* **26** (1992) 21–37.

[34] M. Hinze, R. Pinnau, M. Ulbrich and S. Ulbrich, *Optimization with PDE constraints*, vol. 23. Springer Science & Business Media (2008).

[35] Y. Shin, Z. Zhang and G.E. Karniadakis, Error estimates of residual minimization using neural networks for linear PDEs. *J. Mach. Learn. Model. Comput.* **4** (2023) 73–101.

[36] R.A. Adams and J.J.F. Fournier, Sobolev Spaces. Elsevier (2003).

[37] T. De Ryck, A.D. Jagtap and S. Mishra, Error estimates for physics informed neural networks approximating the Navier–Stokes equations. *IMA J. Numer. Anal.* **44** (2024) 83–119.

[38] S. Mishra and R. Molinaro, Estimates on the generalization error of physics-informed neural networks for approximating PDEs. *IMA J. Numer. Anal.* **43** (2022) 1–43.

[39] P. Bochev and M. Gunzburger, Least-squares methods for hyperbolic problems, in *Handbook of Numerical Analysis*, vol. 17. Elsevier (2016) 289–317.

[40] P.B Bochev and M.D. Gunzburger, *Least-squares Finite Element Methods*, vol. 166. Springer Science & Business Media (2009).

[41] M. Zeinhofer, R. Masri and K.-A. Mardal, A unified framework for the error analysis of physics-informed neural networks. *IMA J. Numer. Anal.* **45** (2025) 2988–3025.

[42] P.L. Lagari, L.H. Tsoukalas, S. Safarkhani and I.E. Lagaris, Systematic construction of neural forms for solving partial differential equations inside rectangular domains, subject to initial, boundary and interface conditions. *Int. J. Artif. Intell. Tools* **29** (2020) 2050009.

[43] N. Sukumar and A. Srivastava, Exact imposition of boundary conditions with distance functions in physics-informed deep neural networks. *Comput. Methods Appl. Mech. Eng.* **389** (2022) 114333.

[44] I. Gühring, G. Kutyniok and P. Petersen, Error bounds for approximations with deep ReLU neural networks in W s, p norms. *Anal. Appl.* **18** (2020) 803–859.

[45] T. De Ryck, S. Lanthaler and S. Mishra, On the approximation of functions by tanh neural networks. *Neural Netw.* **143** (2021) 732–750.
[46] E. Kharazmi, Z. Zhang and G.E. Karniadakis, Variational physics-informed neural networks for solving partial differential equations. Preprint arXiv:1912.00873 (2019).
[47] E. Kharazmi, Z. Zhang and G.E. Karniadakis, hp-vpinns: Variational physics-informed neural networks with domain decomposition. *Comput. Methods Appl. Mech. Eng.* **374** (2021) 113547.
[48] S. Berrone, C. Canuto and M. Pintore, Variational physics informed neural networks: the role of quadratures and test functions. *J. Sci. Comput.* **92** (2022) 100.
[49] J.M. Taylor, D. Pardo and I. Muga, A deep Fourier residual method for solving PDEs using neural networks. *Comput. Methods Appl. Mech. Eng.* **405** (2023) 115850.
[50] H.N. Mhaskar and N. Hahm, Neural networks for functional approximation and system identification. *Neural Computat.* **9** (1997) 143–159.
[51] T. Xie and F. Cao, The errors of simultaneous approximation of multivariate functions by neural networks. *Comput. Math. Appl.* **61** (2011) 3146–3152.
[52] D. Yarotsky, Error bounds for approximations with deep ReLU networks. *Neural Netw.* **94** (2017) 103–114.
[53] S. Wang, S. Sankaran, H. Wang and P. Perdikaris, An expert's guide to training physics-informed neural networks. Preprint arXiv:2308.08468 (2023).
[54] I. Brevis, I. Muga and K.G. van der Zee, Neural control of discrete weak formulations: Galerkin, least squares & minimal-residual methods with quasi-optimal weights. *Comput. Methods Appl. Mech. Eng.* **402** (2022) 115716.
[55] S.C. Brenner, The Mathematical Theory of Finite Element Methods. Springer (2008).
[56] E. Alexandre and J.-L. Guermond, Theory and Practice of Finite Elements. Springer (2004).
[57] D.P. Kingma and J. Ba, Adam: a method for stochastic optimization. Preprint arXiv:1412.6980 (2014).
[58] X. Glorot and Y. Bengio, Understanding the difficulty of training deep feedforward neural Networks, in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. JMLR Workshop and Conference Proceedings (2010) 249–256.
[59] S. Goswami, C. Anitescu, S. Chakraborty and T. Rabczuk, Transfer learning enhanced physics informed neural network for phase-field modeling of fracture. *Theoret. Appl. Fract. Mechan.* **106** (2020) 102447.
[60] C. Xu, B.T. Cao, Y. Yuan and G. Meschke, Transfer learning based physics-informed neural networks for solving inverse problems in engineering structures under different loading scenarios. *Comput. Methods Appl. Mech. Eng.* **405** (2023) 115852.
[61] I. Goodfellow, Y. Bengio and A. Courville, Deep Learning. MIT Press (2016).
[62] V. Jikov, S. Kozlov and O. Oleinik, Homogenization of Differential Operators and Integral Functionals. Springer Science & Business Media (2012).
[63] R.M. Christensen, Mechanics of Composite Materials. Courier Corporation (2012).
[64] R. Ewing, O. Iliev, R. Lazarov, I. Rybak and J. Willems, A simplified method for upscaling composite materials with high contrast of the conductivity. *SIAM J. Sci. Comput.* **31** (2009) 2568–2586.
[65] Z.K. Low, N.B., N. Naouar and D. Baillis, Influence of boundary conditions on computation of the effective thermal conductivity of foams. *Int. J. Heat Mass Transfer* **155** (2020) 119781.
[66] O. Iliev, Z. Lakdawala and V. Starikovicius, On a numerical subgrid upscaling algorithm for Stokes–Brinkman equations. *Comput. Math. Appl.* **65** (2013) 435–448.
[67] F. Chalon, M. Mainguy, P. Longuemare and P. Lemonnier, Upscaling of elastic properties for large scale geomechanical simulations. *Int. J. Numer. Anal. Methods Geomech.* **28** (2004) 1105–1119.
[68] P. Grisvard, Elliptic Problems in Nonsmooth Domains. SIAM (2011).
[69] J. Bradbury, R. Frostig, P. Hawkins, M.J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, *et al.*, Jax: composable transformations of python+ numpy programs (v0. 2.5). Software available from https://github.com/google/jax (2018).
[70] A. Logg and G.N. Wells, Dolfin: automated finite element computing. *ACM Trans. Math. Softw.* **37** (2010) 1–28.