

LEARNING HOMOGENIZATION FOR ELLIPTIC OPERATORS*

KAUSHIK BHATTACHARYA[†], NIKOLA B. KOVACHKI[‡], AAKILA RAJAN[†],
ANDREW M. STUART[†], AND MARGARET TRAUTNER[†]

Abstract. Multiscale partial differential equations (PDEs) arise in various applications, and several schemes have been developed to solve them efficiently. Homogenization theory is a powerful methodology that eliminates the small-scale dependence, resulting in simplified equations that are computationally tractable while accurately predicting the macroscopic response. In the field of continuum mechanics, homogenization is crucial for deriving constitutive laws that incorporate microscale physics in order to formulate balance laws for the macroscopic quantities of interest. However, obtaining homogenized constitutive laws is often challenging as they do not in general have an analytic form and can exhibit phenomena not present on the microscale. In response, data-driven learning of the constitutive law has been proposed as appropriate for this task. However, a major challenge in data-driven learning approaches for this problem has remained unexplored: the impact of discontinuities and corner interfaces in the underlying material. These discontinuities in the coefficients affect the smoothness of the solutions of the underlying equations. Given the prevalence of discontinuous materials in continuum mechanics applications, it is important to address the challenge of learning in this context, in particular, to develop underpinning theory that establishes the reliability of data-driven methods in this scientific domain. The paper addresses this unexplored challenge by investigating the learnability of homogenized constitutive laws for elliptic operators in the presence of such complexities. Approximation theory is presented, and numerical experiments are performed which validate the theory in the context of learning the solution operator defined by the cell problem arising in homogenization for elliptic PDEs.

Key words. homogenization, elliptic PDE, approximation theory, machine learning

MSC codes. 35B27, 35J47, 74H15, 68T07

DOI. 10.1137/23M1585015

1. Introduction. Homogenization theory is a well-established methodology that aims to eliminate fast-scale dependence in partial differential equations (PDEs) to obtain homogenized PDEs which produce a good approximate solution of the problem with fast scales while being more computationally tractable. In continuum mechanics, this methodology is of great practical importance as the constitutive laws derived from physical principles are governed by material behavior at small scales, but the quantities of interest are often relevant on larger scales. These homogenized constitutive laws often do not have a closed analytic form and may have new features not present in the microscale laws. Consequently, there has been a recent surge of interest in employing data-driven methods to learn homogenized constitutive laws.

*Received by the editors July 7, 2023; accepted for publication (in revised form) April 10, 2024; published electronically August 2, 2024.

<https://doi.org/10.1137/23M1585015>

Funding: The work of the first, third, and fourth authors was sponsored, in part, by the Army Research Laboratory under Cooperative Agreement W911NF22-2-0120, and the first author is also sponsored under W911NF-22-1-0269. The work of the fourth author is supported by a Department of Defense Vannevar Bush Faculty Fellowship and by the SciAI Center, funded by the Office of Naval Research (ONR), under grant N00014-23-1-2729. The work of the fifth author is funded by the Department of Energy Computational Science Graduate Fellowship under award DE-SC0021110. The second author is grateful to the NVIDIA Corporation for support through full-time employment. The computations presented here were conducted in the Resnick High Performance Computing Center, a facility supported by Resnick Sustainability Institute at the California Institute of Technology.

[†]California Institute of Technology, Pasadena, CA 91125 USA (bhatta@caltech.edu, arajan@caltech.edu, astuart@caltech.edu, trautner@caltech.edu).

[‡]NVIDIA, Santa Clara, CA 95051 USA (nkovachki@nvidia.com).

The goal of this paper is to study the learnability of homogenized constitutive laws in the context of one of the canonical model problems of homogenization, namely the divergence form elliptic PDE. One significant challenge in applications of homogenization in material science arises from the presence of discontinuities and corner interfaces in the underlying material. This leads to a lack of smoothness in the coefficients and solutions of the associated equations, a phenomenon extensively studied in numerical methods for PDEs. Addressing this challenge in the context of learning remains largely unexplored and is the focus of our work. We develop underlying theory and provide accompanying numerical studies to address learnability in this context.

In subsection 1.1 we establish the mathematical framework and notation for the problem of interest, state the three main contributions of the paper, and overview the contents of each section of the paper. In subsection 1.2 we provide a detailed literature review. Subsection 1.3 states the stability estimates that are key for the approximation theory developed in the paper and discusses the remainder of the paper in the context of these estimates.

1.1. Problem formulation. Consider the following linear multiscale elliptic equation on a bounded domain $\Omega \subset \mathbb{R}^d$:

$$(1.1a) \quad -\nabla_x \cdot (A^\epsilon \nabla_x u^\epsilon) = f, \quad x \in \Omega,$$

$$(1.1b) \quad u^\epsilon = 0, \quad x \in \partial\Omega.$$

Here $A^\epsilon(x) = A\left(\frac{x}{\epsilon}\right)$ for $A(\cdot)$ which is 1-periodic and positive definite: $A : \mathbb{T}^d \rightarrow \mathbb{R}_{\text{sym}, >0}^{d \times d}$, a condition which holds throughout this work. Assume further that $f \in L^2(\Omega; \mathbb{R})$ and has no microscale variation with respect to x/ϵ .

Our focus is on linking this multiscale problem to the homogenized form of (1.1), which is

$$(1.2a) \quad -\nabla_x \cdot (\bar{A} \nabla_x u) = f, \quad x \in \Omega,$$

$$(1.2b) \quad u = 0, \quad x \in \partial\Omega,$$

where \bar{A} is given by

$$(1.3) \quad \bar{A} = \int_{\mathbb{T}^d} (A(y) + A(y) \nabla \chi(y)^T) \, dy,$$

and $\chi : \mathbb{T}^d \rightarrow \mathbb{R}^d$ solves the cell problem

$$(1.4) \quad -\nabla \cdot (\nabla \chi A) = \nabla \cdot A, \quad \chi \text{ is 1-periodic.}$$

All of the preceding PDEs are to be interpreted as holding in the weak sense. For $0 < \epsilon \ll 1$, the solution u^ϵ of (1.1) is approximated by the solution u of (1.2), and the error converges to zero as $\epsilon \rightarrow 0$ in various topologies [5, 10, 54].

We assume that

$$\|A\|_{L^\infty} := \sup_{y \in \mathbb{T}^d} |A(y)|_F < \infty,$$

where $|\cdot|_F$ is the Frobenius norm. Hence $A \in L^\infty(\mathbb{T}^d; \mathbb{R}^{d \times d})$ and $A^\epsilon \in L^\infty(\Omega; \mathbb{R}^{d \times d})$. Similarly, for $A \in L^2(\mathbb{T}^d; \mathbb{R}^{d \times d})$, we define

$$\|A\|_{L^2}^2 := \int_{\mathbb{T}^d} |A(y)|_F^2 \, dy.$$

Also, for given $\beta \geq \alpha > 0$, we define the following subset of 1-periodic, positive-definite, symmetric matrix fields in $L^\infty(\mathbb{T}^d; \mathbb{R}^{d \times d})$ by

$$\text{PD}_{\alpha,\beta} = \{A \in L^\infty(\mathbb{T}^d; \mathbb{R}^{d \times d}) : \forall(y, \xi) \in \mathbb{T}^d \times \mathbb{R}^d, \alpha|\xi|^2 \leq \langle \xi, A(y)\xi \rangle \leq \beta|\xi|^2\}.$$

For open set $\Omega \subset \mathbb{R}^d$, we denote the *variation* of a function $u \in L^1_{\text{loc}}(\Omega)$ by

$$V(u, \Omega) = \sup \left\{ \sum_{i=1}^d \int_{\Omega} \frac{\partial \Phi_i}{\partial x_i} u \, dx : \Phi \in C_0^\infty(\Omega; \mathbb{R}^d), \|\Phi\|_{L^\infty(\Omega; \mathbb{R}^d)} \leq 1 \right\}$$

and the set of functions of bounded variation on \mathbb{T}^d as

$$\text{BV} = \{u \in L^1(\mathbb{T}^d) : V(u, \mathbb{T}^d) < \infty\}.$$

For further information on BV, we refer to [36]. Finally, we often work in the Sobolev space H^1 restricted to spatially mean-zero periodic functions, denoted

$$\dot{H}^1 := \left\{ v \in W^{1,2}(\mathbb{T}^d) \mid v \text{ is 1-periodic, } \int_{\mathbb{T}^d} v \, dy = 0 \right\};$$

the norm on this space is defined by

$$(1.5) \quad \|g\|_{\dot{H}^1} := \|\nabla g\|_{L^2}.$$

Numerically solving (1.1) is far more computationally expensive than solving the homogenized equation (1.2), motivating the wish to find the homogenized coefficient \bar{A} defining (1.2). The difficult part of obtaining (1.2) is solving the cell problem (1.4). Although explicit solutions exist in the one-dimensional setting for piecewise-constant A [9] and in the two-dimensional setting where A is a layered material [54], in general a closed form solution is not available and the cell problem must be solved numerically. Note that in general the action of the divergence $\nabla \cdot$ on terms involving A in the cell problem necessitates the use of weak solutions for $A \notin C^1(\mathbb{T}^d; \mathbb{R}^{d \times d})$; this is a commonly occurring situation in applications such as those arising from porous medium flow, or to vector-valued generalizations of the setting here to elasticity, rendering the numerical solution nontrivial. For this reason, it is potentially valuable to approximate the solution map

$$(1.6) \quad G : A \mapsto \chi,$$

defined by the cell problem, using a map defined by a neural operator (NO). More generally it is foundational to the broader program of learning homogenized constitutive models from data to thoroughly study this issue for the divergence form elliptic equation as the insights gained will be important for understanding the learning of more complex parameterized homogenized models, such as those arising in nonlinear elasticity, viscoelasticity, and plasticity.

The full map from A to the homogenized tensor \bar{A} is expressed by $A \mapsto (\chi, A) \mapsto \bar{A}$, and one could instead learn the map

$$(1.7) \quad F : A \mapsto \bar{A}.$$

Since the map $(\chi, A) \mapsto \bar{A}$ is defined by a quadrature, we focus on the approximation of $A \mapsto \chi$ and state equivalent results for the map $A \mapsto \bar{A}$ that emerge as consequences of the approximation of χ . In this paper we make the following contributions:

1. We state and prove universal approximation theorems for the map G defined by (1.4) and (1.6), and map F defined by (1.3), (1.4), and (1.7).
2. We provide explicit examples of microstructures which satisfy the hypotheses of our theorems; these include microstructures generated by probability measures which generate discontinuous functions in BV .
3. We provide numerical experiments to demonstrate the ability of NOs to approximate the solution map on four different classes of material parameters A , all covered by our theoretical setting.

In subsection 1.2 we provide an overview of the literature, followed in subsection 1.3 by a discussion of stability estimates for (1.4), with respect to variations in A ; these are at the heart of the analysis of universal approximation. The main body of the text then commences with section 2, which characterizes the microstructures of interest to us in the context of continuum mechanics. Section 3 states universal approximation theorems for $G(\cdot)$ and $F(\cdot)$, using the Fourier neural operator (FNO). In section 4 we give numerical experiments illustrating the approximation of the map G defined by (1.6) on microstructures of interest in continuum mechanics. Details of the stability estimates, the proofs of universal approximation theorems, properties of the microstructures, and details of numerical experiments are given in Appendices A, B, C, and D, respectively.

1.2. Literature review. Homogenization aims to derive macroscopic equations that describe the effective properties and behavior of solutions to problems at larger scales given a system that exhibits behavior at (possibly multiple) smaller scales. Although it is developed for the various cases of random, statistically stationary, and periodic small-scale structures, we work here entirely in the periodic setting. The underlying assumption of periodic homogenization theory is that the coefficient is periodic in the small-scale variable and that the scale separation is large compared to the macroscopic scales of interest. Convergence of the solution of the multiscale problem to the homogenized solution is well-studied; see [3, 14]. We refer to the texts [5, 10, 54] for more comprehensive citations to the literature. Homogenization has found extensive application in the setting of continuum mechanics [23] where, for many multiscale materials, the scale-separation assumption is natural. In this work, we are motivated in part by learning constitutive models for solid materials, where crystalline microstructure renders the material parameters discontinuous and may include corner interfaces. This difficulty has been explored extensively in the context of numerical methods for PDEs, particularly with adaptive finite element methods [29, 11, 51, 53].

There is a significant body of work on the approximation theory associated with parametrically dependent solutions of PDEs, including viewing these solutions as a map between the function space of the parameter and the function space of the solution, especially for problems possessing holomorphic regularity [16, 15, 13]. This work could potentially be used to study the cell problem for homogenization that is our focus here. However, there has been recent interest in taking a data-driven approach to solving PDEs via machine learning because of its flexibility and ease of implementation. A particular approach to learning solutions to PDEs is operator learning, a machine learning methodology where the map to be learned is viewed as an operator acting between infinite-dimensional function spaces rather than between finite-dimensional spaces [8, 37, 44, 50, 33]. Determining whether, and then when, operator learning models have advantages over classical numerical methods in solving PDEs remains an active area of research [4]. The paper [47] makes a contribution

to this area, in the context of the divergence form elliptic PDE and the map from coefficient to solution when the coefficient is analytic over its domain; the authors prove that ϵ error is achievable for a DeepONet [44] of size only polylogarithmic in ϵ , leveraging the exponential convergence of spectral collocation methods for boundary value problems with analytic solutions. However, in the setting of learning homogenized constitutive laws in material science, discontinuous coefficients form a natural focus and indeed form the focus of this paper. A few characteristics make operator learning a promising option in this context. First, machine learning has been groundbreaking in application settings with no clear underlying equations, such as computer vision and language models [26, 12]. In constitutive modeling, though the microscale constitutive laws are known, the homogenized equations are generally unknown and can incorporate dependencies that are not present on the microscale, such as history dependence, anisotropy, and slip-stick behavior [55, 7]. Thus, constitutive models lie in a partially equation-free setting where data-driven methods could be useful. Second, machine learned models as surrogates for expensive computation can be valuable when the cost of producing data and training the model can be amortized over many forward uses of the trained model. Since the same materials are often used for fabrication over long time periods, this can be a setting where the upfront cost of data production and model training is justified.

Other work has already begun to explore the use of data-driven methods for constitutive modeling; a general review of the problem and its challenges, in the context of constitutive modeling of composite materials, may be found in [42]. Several works use the popular framework of physics-informed machine learning to approach the problem [21, 58, 45, 24]. In [1], physical constraints are enforced on the network architecture while learning nonlinear elastic constitutive laws. In [38], the model is given access to additional problem-specific physical knowledge. Similarly, the work of [59] predicts the Cholesky factor of the tangent stiffness matrix from which the stress may be calculated; this method enforces certain physical criteria. The paper [31] studies approximation error and uncertainty quantification for this learning problem. In [25], a derivative-free approach is taken to learning homogenized solutions where regularity of the material coefficient is assumed. The work of [39] illustrates the potential of operator learning methodology to model constitutive laws with history dependence, such as those that arise in crystal plasticity. Finally, a number of further works demonstrate empirically the potential of learning constitutive models, including [48, 43, 61, 41].

However, the underlying theory behind operator learning for constitutive models lags behind its empirical application. In [9], approximation theories are developed to justify the use of a recurrent Markovian architecture that performs well in application settings with history dependence. This architecture is further explored in [40] with more complex microstructures. Universal approximation results are a first step in developing theory for learning because they guarantee that there exists an ϵ -approximate operator within the operator approximation class, which is consistent with an assumed true model underlying the data [33, 17, 35, 32]. In addition to universal approximation, further insight may be gained by seeking to quantify the data or model size required to obtain a given level of accuracy; the papers [35, 32, 46] also contain work in this direction, as do the papers [27, 52], which build on the analysis developed in [16, 15, 13] referred to above. In our work we leverage an existing universal approximation theorem for FNOs, a particular practically useful architecture from within the NO class [32]. We take two different approaches to proving approximation theorems based on separate PDE solution stability results in pursuit of a more

robust understanding of the learning problem. Since the field is in its infancy, it is valuable to have different approaches to these analysis problems. Finally, we perform numerical experiments on various microstructures to understand the practical effects of nonsmooth PDE coefficients in learning solutions. We highlight the fact that in this paper we do not tackle issues related to the nonconvex optimization problem at the heart of training neural networks; we simply use state-of-the-art stochastic gradient descent for training, noting that theory explaining its excellent empirical behavior is lacking.

Throughout this paper we focus on (1.1), which describes a conductivity equation in a heterogeneous medium; a natural generalization of interest is to the constitutive law of linear elasticity, in which the solution is vector-valued and the coefficient is a fourth order tensor. Though it is a linear elliptic equation, we echo the sentiment of Blanc and Le Bris [10] with their warning, “do not underestimate the difficulty of equation (1.1).” There are many effects to be understood in this setting, and resolving learning challenges is a key step toward understanding similar questions for the learning of parametric dependence in more complex homogenized constitutive laws where machine learning may prove particularly useful.

1.3. Stability estimates. At the heart of universal approximation theorems is stability of the solution map (1.6), in particular continuity of the map for certain classes of A . In this subsection, we present three key stability results that are used to prove the approximation theorems in section 3. The proofs of the following stability estimates may all be found in Appendix A.

A first strike at the stability of the solution map (1.6) is a modification of the classic L^∞/H^1 Lipschitz continuity result for dependence of the solution of elliptic PDEs on the coefficient; here generalization is necessary because the coefficient also appears on the right-hand side of the equation defining $G(\cdot)$.

PROPOSITION 1.1. *Consider the cell problem defined by (1.4). The following hold:*

1. *If $A \in \text{PD}_{\alpha,\beta}$, then (1.4) has a unique solution $\chi \in \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)$ and*

$$\|\chi\|_{\dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)} \leq \frac{\sqrt{d}\beta}{\alpha}.$$

2. *For $\chi^{(1)}$ and $\chi^{(2)}$ solutions to the cell problem in (1.4) associated with coefficients $A^{(1)}, A^{(2)} \in \text{PD}_{\alpha,\beta}$, respectively, it follows that*

$$(1.8) \quad \|\chi^{(2)} - \chi^{(1)}\|_{\dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)} \leq \frac{\sqrt{d}}{\alpha} \left(1 + \frac{\beta}{\alpha}\right) \|A^{(1)} - A^{(2)}\|_{L^\infty(\mathbb{T}^d; \mathbb{R}^{d \times d})}.$$

However, this perturbation result is insufficient for approximation theory because the space L^∞ is not separable and it is not natural to develop approximation theory in such spaces [19, Chapter 9]. While it is possible to define the problem on a separable subspace of L^∞ (see Lemma A.1), such spaces are not particularly useful in applications to micromechanics. Many natural models for realistic microstructures work with classes of discontinuous functions in which the boundary of material discontinuity can occur anywhere in the domain. Such functions cannot be contained in any separable subspace of L^∞ ; see Lemma A.2. To deal with this issue it is desirable to establish continuity from L^q to \dot{H}^1 for some $q \in [2, \infty)$. To this end, we provide two additional stability results. The first stability result gives continuity, but not Lipschitz continuity, from L^2 to \dot{H}^1 . The second stability result gives Lipschitz continuity from L^q to \dot{H}^1 , some $q \in (2, \infty)$.

PROPOSITION 1.2. Endow $\text{PD}_{\alpha,\beta}$ with the $L^2(\mathbb{T}^d; \mathbb{R}^{d \times d})$ induced topology and let $K \subset \text{PD}_{\alpha,\beta}$ be a closed set. Define the mapping $G : K \rightarrow \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)$ by $A \mapsto \chi$ as given by (1.4). Then there exists a bounded continuous mapping

$$\mathcal{G} \in C(L^2(\mathbb{T}^d; \mathbb{R}^{d \times d}); \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d))$$

such that $\mathcal{G}(A) = G(A)$ for any $A \in K$.

The preceding L^2 continuity proposition is used to prove the approximation results for the FNO in Theorems 3.3 and 3.4. While not necessary for the approximation theory proofs, the following proposition on Lipschitz continuity from L^q to \dot{H}^1 establishes a more concrete bound on the approximation error, which allows for additional analysis such as providing rough bounds on grid error as discussed in subsection 4.3.

PROPOSITION 1.3. There exists $q_0 \in (2, \infty)$ such that, for all q satisfying $q \in (q_0, \infty]$, the following holds. Endow $\text{PD}_{\alpha,\beta}$ with the $L^q(\mathbb{T}^d; \mathbb{R}^{d \times d})$ topology and let $K \subset \text{PD}_{\alpha,\beta}$ be a closed set. Define the mapping $G : K \rightarrow \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)$ by $A \mapsto \chi$ as given by (1.4). Then there exists a bounded Lipschitz-continuous mapping

$$\mathcal{G} : L^q(\mathbb{T}^d; \mathbb{R}^{d \times d}) \rightarrow \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)$$

such that $\mathcal{G}(A) = G(A)$ for any $A \in K$.

Remark 1.4. Explicit upper bounds for q_0 in Proposition 1.3 exist and are discussed in Remark A.14.

2. Microstructures. The main application area of this work is constitutive modeling. In this section we describe various classes of microstructures that our theory covers. In particular, we describe four classes of microstructures in two dimensions:

1. Smooth microstructures generated via truncated, rescaled log-normal random fields.
2. Discontinuous microstructures with smooth interfaces generated by Lipschitz star-shaped inclusions.
3. Discontinuous microstructures with square inclusions.
4. Voronoi crystal microstructures.

Visualizations of examples of these microstructures may be found in Figure 1. We emphasize that all four examples lead to functions in BV , a fact that we exploit in section 4 when showing that our abstract analysis from section 3 applies to them all.

Smooth microstructures. The smooth microstructures are generated by exponentiating a rescaled Gaussian random field. A is symmetric and coercive everywhere in the domain with a bounded eigenvalue ratio. Furthermore, the smooth function A and its derivatives are Lipschitz. Our theory is developed specifically to analyze nonsmooth microstructures, so this example is used mainly as a point of comparison.

Star inclusions. For the star inclusion microstructure, A is taken to be constant inside and outside the star-shaped boundary. The boundary function is smooth and



FIG. 1. Microstructure examples.

Lipschitz in each of its derivatives. A is positive and coercive in both regions with a bounded eigenvalue ratio. This microstructure introduces discontinuities, but the boundary remains smooth.

Square inclusions. For the square inclusion microstructure, A is taken to be constant inside and outside the square boundary. Since we assume periodicity, without loss of generality the square inclusion is centered. The size of the square inclusion within the cell is varied between samples, as are the constant values of A . This microstructure builds on the complexity of the star inclusion microstructure by adding corners to the inclusion boundary.

Voronoi interfaces. The Voronoi crystal microstructures are generated by assuming a random Voronoi tessellation and letting A be piecewise-constant taking a single value on each Voronoi cell. The values of A on the cells and locations of the cell centers may be varied. This is the most complex microstructure among our examples and is a primary motivation for this work as Voronoi tessellations are a common model for crystal structure in materials.

3. Universal approximation results. In this section we state the two approximation theorems for learning solution operators to the cell problem. Theorem 3.3 concerns learning the map $A \rightarrow \chi$ in (1.4), and Theorem 3.4 concerns learning the map $A \rightarrow \bar{A}$ described by the combination of (1.4) and (1.3). Theorems 3.3 and 3.4 are specific to learning an FNO, which is a subclass of the general NO. The proofs of the theorems in this section may be found in Appendix B.

3.1. Definitions of neural operators. First, we define a general NO and the FNO. The definitions are largely taken from [33], and we refer to this work for a more in-depth understanding of these operators. In this work, we restrict the domain to the torus.

DEFINITION 3.1 (general neural operator). *Let \mathcal{A} and \mathcal{U} be two Banach spaces of real vector-valued functions over domain \mathbb{T}^d . Assume input functions $a \in \mathcal{A}$ are \mathbb{R}^{d_a} -valued while the output functions $u \in \mathcal{U}$ are \mathbb{R}^{d_u} -valued. The NO architecture $\mathcal{G}_\theta : \mathcal{A} \rightarrow \mathcal{U}$ is*

$$\begin{aligned} \mathcal{G}_\theta &= \mathcal{Q} \circ \mathcal{L}_{T-1} \circ \dots \circ \mathcal{L}_0 \circ \mathcal{P}, \\ v_{t+1} &= \mathcal{L}_t v_t = \sigma_t(W_t v_t + \mathcal{K}_t v_t + b_t), \quad t = 0, 1, \dots, T-1, \end{aligned}$$

with $v_0 = \mathcal{P}(a)$, $u = \mathcal{Q}(v_T)$ and $\mathcal{G}_\theta(a) = u$. Here, $\mathcal{P} : \mathbb{R}^{d_a} \rightarrow \mathbb{R}^{d_{v_0}}$ is a local lifting map, $\mathcal{Q} : \mathbb{R}^{d_{v_T}} \rightarrow \mathbb{R}^{d_u}$ is a local projection map, and the σ_t are fixed nonlinear activation functions acting locally as maps $\mathbb{R}^{d_{v_{t+1}}} \rightarrow \mathbb{R}^{d_{v_{t+1}}}$ in each layer (with all of \mathcal{P} , \mathcal{Q} and the σ_t viewed as operators acting pointwise, or pointwise almost everywhere, over the domain \mathbb{T}^d), $W_t \in \mathbb{R}^{d_{v_{t+1}} \times d_{v_t}}$ are matrices, $\mathcal{K}_t : \{v_t : \mathbb{T}^d \rightarrow \mathbb{R}^{d_{v_t}}\} \rightarrow \{v_{t+1} : \mathbb{T}^d \rightarrow \mathbb{R}^{d_{v_{t+1}}}\}$ are integral kernel operators, and $b_t : \mathbb{T}^d \rightarrow \mathbb{R}^{d_{v_{t+1}}}$ are bias functions. For any $m \in \mathbb{N}_0$, the activation functions σ_t are restricted to the set of continuous $\mathbb{R} \rightarrow \mathbb{R}$ maps which make real-valued, feed-forward neural networks dense in $C^m(\mathbb{R}^d)$ on compact sets for any fixed network depth.¹ The integral kernel operators \mathcal{K}_t are defined as

$$(\mathcal{K}_t v_t)(x) = \int_{\mathbb{T}^d} \kappa_t(x, y) v_t(y) dy$$

with standard multilayered perceptrons (MLP) $\kappa_t : \mathbb{T}^d \times \mathbb{T}^d \rightarrow \mathbb{R}^{d_{v_{t+1}} \times d_{v_t}}$. We denote by θ the collection of parameters that specify \mathcal{G}_θ , which include the weights W_t , biases

¹We note that all globally Lipschitz, nonpolynomial, $C^m(\mathbb{R})$ functions belong to this class.

b_t , parameters of the kernels κ_t , and parameters describing the lifting and projection maps \mathcal{P} and \mathcal{Q} (usually also MLPs).

The FNO is a subclass of the NO.

DEFINITION 3.2 (Fourier neural operator). *The FNO inherits the structure and definition of the NO in Definition 3.1, together with some specific design choices. We fix $d_{v_t} = d_v$ for all t , where d_v is referred to as the number of channels, or model width, of the FNO. We fix $\sigma_t = \sigma$ to be a globally Lipschitz, nonpolynomial, C^∞ function.² Finally, the kernel operators \mathcal{K}_t are parameterized in the Fourier domain in the following manner. Let*

$$\psi_k(x) = e^{2\pi i \langle k, x \rangle}, \quad x \in \mathbb{T}^d, k \in \mathbb{Z}^d,$$

denote the Fourier basis for $L^2(\mathbb{T}^d; \mathbb{C})$ where $i = \sqrt{-1}$ is the imaginary unit. Then, for each t , the kernel operator \mathcal{K}_t is parameterized by

$$(\mathcal{K}_t v_t)_l(x) = \sum_{\substack{k \in \mathbb{Z}^d \\ |k| \leq k_{\max}}} \left(\sum_{j=1}^{d_v} P_{lj}^k \langle \psi_k, (v_t)_j \rangle_{L^2(\mathbb{T}^d; \mathbb{C})} \right) \psi_k(x).$$

Here, $l = 1, \dots, d_v$ and each $P^k \in \mathbb{C}^{d_v \times d_v}$ constitute the learnable parameters of the integral operator.

From the definition of the FNO, we note that parameterizing the kernels in the Fourier domain allows for efficient computation using the FFT. We refer to [37, 33] for additional details.

Finally we observe that in numerous applications, an example being learning of the map $A \mapsto \bar{A}$ (1.3), (1.4), it is desirable to modify the FNO so that the output space is simply a Euclidean space, and not a function space; this generalization is explored in [30]. An alternative approach, exemplified by Theorem 3.4 in the next subsection, is to allow the FNO output to be a function that may be evaluated at any point in the domain to yield an approximation of the point in Euclidean space.

3.2. Main theorems. These two theorems guarantee the existence of an FNO approximating the maps $A \mapsto \chi$ and $A \mapsto \bar{A}$ and are based on the stability estimate for continuity from $L^2 \rightarrow \dot{H}^1$ obtained in Proposition 1.2. Both theorems are proved in Appendix B.

THEOREM 3.3. *Let $K \subset \text{PD}_{\alpha, \beta}$ and define the mapping $G : K \rightarrow \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)$ by $A \mapsto \chi$ as given by (1.4). Assume in addition that K is compact in $L^2(\mathbb{T}^d; \mathbb{R}^{d \times d})$. Then, for any $\epsilon > 0$, there exists an FNO $\Psi : K \rightarrow \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)$ such that*

$$\sup_{A \in K} \|G(A) - \Psi(A)\|_{\dot{H}^1} < \epsilon.$$

THEOREM 3.4. *Let $K \subset \text{PD}_{\alpha, \beta}$ and define the mapping $F : K \rightarrow \mathbb{R}^{d \times d}$ by $A \mapsto \bar{A}$ as given by (1.3), (1.4). Assume in addition that K is compact in $L^2(\mathbb{T}^d; \mathbb{R}^{d \times d})$. Then, for any $\epsilon > 0$, there exists an FNO $\Phi : K \rightarrow L^\infty(\mathbb{T}^d; \mathbb{R}^{d \times d})$ such that*

$$\sup_{A \in K} \sup_{x \in \mathbb{T}^d} |F(A) - \Phi(A)(x)|_F < \epsilon.$$

²In this work, in all numerical experiments we use the GeLU activation function as in [37].

The above approximation results can also be formulated to hold, on average, over any probability measure with a finite second moment that is supported on $\text{PD}_{\alpha,\beta}$. In particular, if we let μ be such a probability measure, then there exists an FNO or an NO Ψ such that

$$(3.1) \quad \mathbb{E}_{A \sim \mu} \|G(A) - \Psi(A)\|_{\dot{H}^1} < \epsilon.$$

This follows by applying Theorem 18 from [32] in the respective proofs instead of Theorem 5 from the same work. We do not carry out the full details here. While this allows approximation over the noncompact set $\text{PD}_{\alpha,\beta}$, the error can only be controlled on average instead of uniformly. In section 4, inputs are generated via probability measures supported on compact subsets of L^2 ; thus both the approximation Theorem 3.3 and its analogue in the form (3.1) are relevant.

4. Numerical experiments. In this section, we show that it is possible to find good operator approximations of the homogenization map (1.6), defined by (1.4), in practice. We focus on use of the FNO and note that while Theorems 3.3 and 3.4 assert the existence of desirable operator approximations, they are not constructive and do not come equipped with error estimates. We find approximations using standard empirical loss minimization techniques and, by means of numerical experiments, quantify the complexity with respect to volume of data and with respect to size of parametric approximation.

We work with the microstructures from section 2. In this context we note that Theorems 3.3 and 3.4 apply. To demonstrate this it is necessary to establish that the subsets of coefficient functions employed are compact in L^2 . We achieve this by noting that all our sets of coefficient functions are contained in $\text{PD}_{\alpha,\beta} \cap \text{BV}$, as defined in subsection 1.1. Then we use Lemma C.1 to establish compactness of these subsets of coefficient functions in L^2 . The smooth microstructure example serves as a comparison case for examining the impact of discontinuous coefficients on the learning accuracy. The remaining three examples present different approximation theoretic challenges including curved boundaries (star inclusions), corners (square inclusions), and junctions of several domains (Voronoi).

The experiments are all conducted using an FNO with a fixed number $T = 4$ of hidden layers. The two remaining parameters to vary are the channel width d_v and the number of Fourier modes k_{max} . For implementation details, see Appendix D. We make the following observations based on the numerical experiments.

1. The effective \bar{A} tensors computed from the model predicted solutions exhibit relative error under 1% for all examples; the effective \bar{A} is computed from the learned cell problem solution χ using (1.3).
2. The error in the learned χ is significantly higher along discontinuous material boundaries and corner interfaces, as expected. However, the FNO operator approximation is able to approximate the solution with reasonable relative error even for the most complex case; this most complex case concerns the set of input functions with varying Voronoi geometry and varying microstructural properties within the domain.
3. In comparison with the smooth microstructure case, learning the map for the Voronoi microstructure requires substantially more data to avoid training a model which plateaus at a poor level of accuracy.
4. When compared with the smooth microstructure case, the error for the Voronoi microstructure decreases more slowly with respect to increasing model width,

but shows more favorable response with respect to increasing the number of Fourier modes.

5. Models trained at one discretization may be evaluated at different discretizations for both the smooth and Voronoi microstructures as is characteristic of the FNO. The Voronoi microstructure exhibits, empirically, greater robustness to changes in discretization.

We first describe implementation details of each of the microstructures in subsection 4.1. Then we show outcomes of the numerical experiments in subsection 4.2, discussing them in subsection 4.3.

4.1. Microstructure implementation. For each microstructure, two positive eigenvalues and three components of the two eigenvectors are randomly generated, and the final eigenvector component is chosen to enforce symmetry. All eigenvalue ratios are at most e^2 by construction. In this manner, A is symmetric and coercive and has a bounded eigenvalue ratio.

Smooth microstructures. The smooth microstructures are generated by exponentiating a rescaled approximation of a Gaussian random field. The random field used to generate the eigenvalues and three eigenvector components of $A(x)$ is as follows:

$$\begin{aligned}\widehat{\lambda}_i(x) &= \sum_{k_1, k_2=1}^4 \xi_{k_1, k_2}^{(1)} \sin(2\pi k_1 x_1) \cos(2\pi k_2 x_2) + \xi_{k_1, k_2}^{(2)} \cos(2\pi k_1 x_1) \sin(2\pi k_2 x_2), \\ \lambda_i(x) &= \exp\left(\frac{\widehat{\lambda}_i(x)}{\max_{x' \in [0, 1]^2} |\widehat{\lambda}_i(x')|}\right),\end{aligned}$$

where $\xi_{k_1, k_2}^{(j)}$ are independent and identically distributed (i.i.d.) normal Gaussian random variables.

Star-shaped inclusions. The star-shaped inclusions are generated by defining a random Lipschitz polar boundary function as

$$r(\theta) = \mathbf{a} + \mathbf{b} \sum_{k=1}^5 \xi_k \sin(k\theta),$$

where ξ_k are i.i.d. uniform random variables $U[-1, 1]$, and \mathbf{a} and \mathbf{b} are constants that guarantee $0 < \epsilon < r < 0.5 - \epsilon$ for some fixed $\epsilon > 0$. Then $A(x)$ is constant inside and outside the boundary. We randomly sample eigenvalues for A on each domain via $\lambda_i \sim U[e^{-1}, e]$. The three components of the eigenvectors are i.i.d. normal random variables.

Square inclusions. The radius of the square is randomly generated via

$$r = \mathbf{a} + \mathbf{b}\zeta,$$

where ζ is a uniform random variable on $[0, 1]$ and \mathbf{a} and \mathbf{b} are positive constants that guarantee $0 < \epsilon < r < 0.5 - \epsilon$ for some fixed $\epsilon > 0$. The values of A on each of the constant domains are chosen in the same manner as in the star-shaped inclusion case.

Voronoi interfaces. The Voronoi crystal microstructure has constant A on each Voronoi cell and is chosen uniformly at random in the same manner as for the star inclusions. Voronoi tessellations are a common model for crystal structure in materials. In one Voronoi example, we fix the geometry for all data, and in a second Voronoi example we vary the geometry by randomly sampling five cell centers from a uniform distribution on the unit square.

4.2. Results. Each FNO model is trained using the empirical estimate of the mean squared H^1 norm:

$$(4.1) \quad \text{Loss}(\theta) = \frac{1}{N} \sum_{n=1}^N \left(\|\chi^{(n)} - \hat{\chi}^{(n)}\|_{L^2}^2 + \|\nabla\chi^{(n)} - \nabla\hat{\chi}^{(n)}\|_{L^2}^2 \right),$$

where n is the sample index, χ is the true solution, and $\hat{\chi}$ is the FNO approximation of the solution parameterized by θ . In the analysis, we examine several different measures of error, including the following relative H^1 and relative $W^{1,10}$ errors:

$$(4.2a) \quad \text{Relative } H^1 \text{ Error (RHE)} = \frac{1}{N} \sum_{n=1}^N \left(\frac{\|\chi^{(n)} - \hat{\chi}^{(n)}\|_{L^2}^2 + \|\nabla\chi^{(n)} - \nabla\hat{\chi}^{(n)}\|_{L^2}^2}{\|\chi^{(n)}\|_{L^2}^2 + \|\nabla\chi^{(n)}\|_{L^2}^2} \right)^{\frac{1}{2}},$$

$$(4.2b) \quad \text{Relative } W^{1,10} \text{ Error (RWE)} = \frac{1}{N} \sum_{n=1}^N \left(\frac{\|\chi^{(n)} - \hat{\chi}^{(n)}\|_{L^{10}}^{10} + \|\nabla\chi^{(n)} - \nabla\hat{\chi}^{(n)}\|_{L^{10}}^{10}}{\|\chi^{(n)}\|_{L^{10}}^{10} + \|\nabla\chi^{(n)}\|_{L^{10}}^{10}} \right)^{\frac{1}{10}}.$$

The $W^{1,10}$ norm gives a sense of the higher errors that occur at interfaces, corners, and functions. We could have used $W^{1,p}$ for any p large enough.

Finally, we also look at error in \bar{A} , which we scale by the difference between the arithmetic and harmonic mean of A . Any effective \bar{A} should have a norm in this range; these are known in mechanics as Voigt–Reuss bounds and have a physical interpretation as bounds obtained via energy principles by ignoring equilibrium for the upper bound (arithmetic mean) and ignoring compatibility for the lower bound (harmonic mean) [28]. The resulting error measure is given by

$$(4.3) \quad \text{Relative } \bar{A} \text{ Error (RAE)} = \frac{\|\bar{A} - \hat{\bar{A}}\|_F}{a_m - a_h},$$

where the arithmetic mean a_m and harmonic mean a_h are given by

$$a_m = \left\| \int_{\mathbb{T}^2} A(x) \, dx \right\|_F, \quad a_h = \left\| \left(\int_{\mathbb{T}^2} A^{-1}(x) \, dx \right)^{-1} \right\|_F.$$

We note that using $a_m - a_h$ rather than $\|\bar{A}\|_F$ as a scaling factor in (4.3) leads to a larger error value, so achieving low error in this measure of distance is harder.

We train models on five different datasets. Visualizations of the median-error test samples for each example may be viewed in Figure 2, and the numerical errors are shown in Table 1. Each of these models is trained on 9500 data samples generated using an FE solver on a triangular mesh with the solution interpolated to a 128×128 grid. Additional model details may be found in Appendix D.

We perform an experiment to test the discretization-robustness of the FNO model, results of which are shown in Figure 3. The models are trained with data from the resolution 128×128 and evaluated on test data with different resolution. We emphasize that evaluating the FNO on different resolution is trivial in implementation by design.




FIG. 2. Visualization of the trained models evaluated on test samples that gave median relative H^1 error for each microstructure. The microstructure inputs of each row correspond to those of Figure 1. The first shows the true χ_1 , the second shows the FNO predicted χ_1 , and the third shows the absolute value of the error between the true and predicted χ_1 . The fourth column shows the 2-norm of the gradient of the true χ_1 , and the fifth shows the 2-norm of the gradient of the predicted χ_1 . The last column shows the 2-norm of the difference between the two gradients.

TABLE 1

Errors for each numerical experiment; five sample models are trained for each microstructure. The expressions for the RHE (relative H^1 error), RWE (relative $W^{1,10}$ error), and RAE (relative \bar{A} error) may be found in (4.2) and (4.3). The errors are evaluated over a test set of size 500. All examples have varying geometry except the second Voronoi example.

Microstructure	Mean RHE	Mean RWE	Median RAE
Smooth	$0.0062 \pm 1 \cdot 10^{-4}$	$0.0091 \pm 1 \cdot 10^{-4}$	$0.0007 \pm 1 \cdot 10^{-5}$
Star	$0.0313 \pm 1 \cdot 10^{-4}$	$0.1318 \pm 5 \cdot 10^{-4}$	$0.0014 \pm 3 \cdot 10^{-5}$
Square	$0.1012 \pm 5 \cdot 10^{-4}$	$0.2741 \pm 2 \cdot 10^{-3}$	$0.0047 \pm 1 \cdot 10^{-4}$
Voronoi	$0.0565 \pm 4 \cdot 10^{-4}$	$0.2129 \pm 3 \cdot 10^{-3}$	$0.0027 \pm 8 \cdot 10^{-5}$
Voronoi (fixed geometry)	$0.0073 \pm 3 \cdot 10^{-5}$	$0.0140 \pm 3 \cdot 10^{-4}$	$0.0007 \pm 2 \cdot 10^{-5}$




FIG. 3. Five sample models trained on smooth and Voronoi data at 128×128 grid resolution evaluated at different resolutions.

We also investigate the effects of the number of training data and the model size on the error for the smooth and Voronoi microstructures; similar experiments, for different operator learning problems, are presented in [18]. A plot of error versus training data may be found in Figure 4, and plots of error versus the number of Fourier modes for fixed total model size, as measured by (model width) \times (number Fourier modes), may be found in Figure 5. Figure 5 addresses the question of how to optimally distribute computational budget through different parameterizations to achieve minimum error at given cost as measured by number of parameters; it should be compared to similar experiments in [34].




FIG. 4. A comparison of test error for different amounts of training data for models trained on Voronoi and smooth data. Five sample models are used for each data point.




FIG. 5. Relative H^1 error versus model size for the smooth and Voronoi examples with varying geometry. The number of Fourier modes in each direction and the model width were varied. Each line indicates a constant product of modes \times width. Training data size was fixed at 9500 samples, and five samples were used for each data point.

4.3. Discussion. As can be seen from the data in Table 1, the microstructures exhibiting discontinuities lead to higher model error than the smooth microstructure, and the introduction of corner interfaces leads to further increase in error. The visualizations of the median-error test samples in Figure 2 give some intuition; error is an order of magnitude higher along discontinuous boundaries, which is most apparent in the gradient. The true solution gradient often takes its most extreme values along the discontinuities, and the RWE gives an indication of how well the model captures the most extreme values in the solution. Unsurprisingly, this error is much higher than the RHE, but we note that it is confined to a small area of the domain along discontinuous boundaries and corner interfaces.

In the discretization-robustness experiment described in Figure 3, we observe that the Voronoi model exhibits greater robustness to changes in discretization. We hypothesize that, in the direction of decreasing resolution, the smaller error increase for the Voronoi model, in comparison with the smooth model, could be due to the piecewise-constant nature of the Voronoi microstructure on faces; improved resolution here does not help. On the other hand, for larger grid sizes, increased resolution on corners and discontinuities can help, which could explain the decrease in error from grid edge size of 128 to 256 for the Voronoi model while the smooth model increases in error. One could fine-tune the trained models with small amounts of data from different resolutions, but we leave this transfer learning exploration to future work.

We also examine the effect of the number of training data samples and the FNO size on model accuracy for the smooth and Voronoi microstructures. For data size dependence, we observe in Figure 4 that for these two microstructures, the test error scales $\approx N^{-0.65}$ and $\approx N^{-0.25}$, respectively, where N is the number of training data. In theory, we do not expect to beat the Monte Carlo error decay of $\frac{1}{\sqrt{N}}$ [49]. We note that this is comparable to the behavior during training over 400 epochs; the test error for the smooth microstructure continues to decrease over the entire training periodic, but the test error for the Voronoi microstructure plateaus by around 100 epochs. The model size also presents a qualitatively different effect on error for the smooth and Voronoi microstructures. In Figure 5, we see the trade-off between the number of Fourier modes and the model width for approximately constant model size, measured as the product of the width and number of modes. The Voronoi example benefits from additional Fourier modes, whereas the smooth example worsens. On the other hand, the smooth model benefits more from an increase in model width. We refer

to [18, 34] for in-depth numerical studies of errors, choice of hyperparameters, and parameter distributions for FNO; here we highlight only the qualitative differences between the model behavior for different microstructures.

We also note that a significant portion of the model error may be attributed to grid ambiguity; with a 128×128 grid, the FNO does not know where between gridpoints a discontinuity may fall. This may be quantified empirically in the case of the square microstructure. We perform an experiment in which we create data of square microstructure inclusions whose boundary falls exactly on the gridpoints. One dataset treats the boundary as open, and the other treats the boundary as closed; the input grid points that fall on the boundary differ between the two datasets. We quantify grid ambiguity error by the difference in the outputs of a model given both the open square data and the closed square data. We find that the absolute H^1 norm of the difference between these two outputs is 0.041, which is slightly under twice the absolute H^1 norm of the output compared to the true solution, which has a value of 0.025. We hypothesize that the model learns to assume the boundary falls near the middle of the grid square, which explains why the output difference between the two datasets is roughly twice the true error. From a theory standpoint, one could bound the Lipschitz constant of the FNO and compare it to the Lipschitz constant of the true map described by Proposition 1.3. However, we leave the theoretical estimates of error rates to future work.

Finally, we compare the error in the effective \bar{A} defined in (1.3). This error is scaled by a difference between the Frobenius norms of the arithmetic and harmonic means of the true A because the Frobenius norm of the true \bar{A} should fall within that range. For this reason, in the case where the arithmetic and harmonic means are very close, as is frequently the case for the square and star inclusions, it is not valuable to learn the true \bar{A} . On the other hand, the varying-geometry Voronoi microstructure example on average has about 100 times greater difference between the arithmetic and harmonic means, in comparison with the star and square microstructure examples. This characteristic of the Voronoi microstructure further underscores the value of learning in this setting.

5. Conclusions. In this work, we establish approximation theory for learning the solution operator arising from the elliptic homogenization cell problem (1.4), viewed as a mapping from the coefficient to the solution; the theory allows for discontinuous coefficients. We also perform numerical experiments that validate the theory, explore qualitative differences between various microstructures, and quantify error/cost trade-offs in the approximation. We provide two different stability results for the underlying solutions that build understanding of the underlying map. These stability results, when combined with existing universal approximation results for neural operators, result in rigorous approximation theory for learning in this problem setting. On the empirical side we provide, and then study numerically, examples of various microstructures that satisfy the conditions of the approximation theory. We observe that model error is dominated by error along discontinuous and corner interfaces, and that discontinuous microstructures give rise to qualitatively different learning behavior. Finally, we remark that the learned effective properties are highly accurate, especially in the case of the Voronoi microstructure that we regard as the most complex. Since discontinuous microstructures arise naturally in solid mechanics, understanding learning behavior in this context is an important prerequisite for using machine learning for applications. In this area and others, numerous questions remain which address the rigor necessary for use of machine learning in scientific applications.

We have confined our studies to one of the canonical model problems of homogenization theory, the divergence form elliptic setting with periodic microstructure, to obtain deeper understanding of the learning constitutive laws. One interesting potential extension of this work is the setting in which the material coefficient A is not periodic but random with respect to the microstructure. Another is where it is only locally periodic and has dependence on the macroscale variable as well; thus $A^\epsilon = A(x, \frac{x}{\epsilon})$. In this case, the form of the cell problem (1.4) and homogenized coefficient (1.3) remain the same, but A and χ both have parametric dependence on x . The approximation theory and the empirical learning problem would grow in complexity in comparison to what is developed here, but the resulting methodology could be useful and foundational for understanding more complex constitutive models in which the force balance equation couples to other variables. Indeed, the need for efficient learning of constitutive models is particularly pressing in complex settings such as crystal plasticity. We anticipate that the potential use of machine learning to determine parametric dependence of constitutive models defined by homogenization will be for these more complex problems. The work described in this paper provides an underpinning conceptual approach, foundational analysis, and set of numerical experiments that serve to underpin more applied work in this field.

Appendix A. Proofs of stability estimates. In this section, we prove the stability estimates stated in section 1.3. The following lemma is a modification of the standard estimate for parametric dependence of elliptic equations on their coefficient. We include it here for completeness.

PROPOSITION 1.1. *Consider the cell problem defined by (1.4). The following hold:*

1. *If $A \in \text{PD}_{\alpha,\beta}$, then (1.4) has a unique solution $\chi \in \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)$ and*

$$\|\chi\|_{\dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)} \leq \frac{\sqrt{d}\beta}{\alpha}.$$

2. *For $\chi^{(1)}$ and $\chi^{(2)}$ solutions to the cell problem in (1.4) associated with coefficients $A^{(1)}, A^{(2)} \in \text{PD}_{\alpha,\beta}$, respectively, it follows that*

$$(1.8) \quad \|\chi^{(2)} - \chi^{(1)}\|_{\dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)} \leq \frac{\sqrt{d}}{\alpha} \left(1 + \frac{\beta}{\alpha}\right) \|A^{(1)} - A^{(2)}\|_{L^\infty(\mathbb{T}^d; \mathbb{R}^{d \times d})}.$$

Proof. For existence and uniqueness of the solution to the cell problem using Lax–Milgram, we refer to the texts [10, 54]; we simply derive the bounds and stability estimate. First, note that (1.4) decouples, in particular,

$$(A.1) \quad -\nabla \cdot (A \nabla \chi_\ell) = \nabla \cdot A e_\ell, \quad y \in \mathbb{T}^d,$$

for $l = 1, \dots, d$, where e_ℓ is the ℓ th standard basis vector of \mathbb{R}^d and each $\chi_\ell \in \dot{H}^1(\mathbb{T}^d; \mathbb{R})$. Multiplying by χ_ℓ and integrating by parts shows

$$\begin{aligned} \alpha \|\nabla \chi_\ell\|_{L^2}^2 &\leq \int_{\mathbb{T}^d} \langle A \nabla \chi_\ell, \nabla \chi_\ell \rangle \, dy \\ &= - \int_{\mathbb{T}^d} \langle A e_\ell, \nabla \chi_\ell \rangle \, dy \\ &\leq \int_{\mathbb{T}^d} |A e_\ell| |\nabla \chi_\ell| \, dy \\ &\leq \left(\int_{\mathbb{T}^d} |A e_\ell|^2 \, dy \right)^{\frac{1}{2}} \left(\int_{\mathbb{T}^d} |\nabla \chi_\ell|^2 \, dy \right)^{\frac{1}{2}} \\ &\leq \|A\|_{L^\infty} \|\nabla \chi_\ell\|_{L^2}. \end{aligned}$$

Therefore

$$\|\nabla\chi\|_{L^2}^2 = \sum_{l=1}^d \|\nabla\chi_l\|_{L^2}^2 \leq \frac{d\|A\|_{L^\infty}^2}{\alpha^2} \leq \frac{d\beta^2}{\alpha^2},$$

which implies the first result.

To prove the second result, we denote the right-hand side of (A.1) by $f_\ell^{(i)} = \nabla \cdot A^{(i)}e_\ell$ in what follows. For any $v \in \dot{H}^1(\mathbb{T}^d; \mathbb{R})$, we have that

$$\begin{aligned} & - \int_{\mathbb{T}^d} \nabla \cdot (A^{(1)}\nabla\chi_\ell^{(1)})v \, dy = \int_{\mathbb{T}^d} f_\ell^{(1)}v \, dy, \\ & - \int_{\partial\mathbb{T}^d} vA^{(1)}\nabla\chi_\ell^{(1)} \cdot \hat{n} \, dy + \int_{\mathbb{T}^d} \nabla v \cdot A^{(1)}\nabla\chi_\ell^{(1)} \, dy = \int_{\mathbb{T}^d} f_\ell^{(1)}v \, dy. \end{aligned}$$

Since v , $A^{(1)}$, and the solution $\chi_\ell^{(1)}$ are all periodic on \mathbb{T}^d , the first term is 0. Combining with the equation for $\chi_\ell^{(2)}$, we get

$$\begin{aligned} & \int_{\mathbb{T}^d} \nabla v \cdot (A^{(1)} - A^{(2)})\nabla\chi_\ell^{(1)} \, dy \\ & = \int_{\mathbb{T}^d} (f_\ell^{(1)} - f_\ell^{(2)})v + \nabla v \cdot (A^{(2)}(\nabla\chi_\ell^{(2)} - \nabla\chi_\ell^{(1)})) \, dy. \end{aligned}$$

Setting $v = \chi_\ell^{(2)} - \chi_\ell^{(1)}$, we have

$$\begin{aligned} & \int_{\mathbb{T}^d} (\nabla\chi_\ell^{(2)} - \nabla\chi_\ell^{(1)}) \cdot ((A^{(1)} - A^{(2)})\nabla\chi_\ell^{(1)}) \, dy = \int_{\mathbb{T}^d} (f_\ell^{(1)} - f_\ell^{(2)})(\chi_\ell^{(2)} - \chi_\ell^{(1)}) \, dy \\ & + \int_{\mathbb{T}^d} (\nabla\chi_\ell^{(2)} - \nabla\chi_\ell^{(1)}) \cdot (A^{(2)}(\nabla\chi_\ell^{(2)} - \nabla\chi_\ell^{(1)})) \, dy, \\ & \alpha\|\nabla\chi_\ell^{(2)} - \nabla\chi_\ell^{(1)}\|_{L^2}^2 \leq \|A^{(1)} - A^{(2)}\|_{L^\infty}\|\nabla\chi_\ell^{(1)}\|_{L^2}\|\nabla\chi_\ell^{(2)} - \nabla\chi_\ell^{(1)}\|_{L^2} \\ & + \|f_\ell^{(1)} - f_\ell^{(2)}\|_{\dot{H}^{-1}}\|\nabla\chi_\ell^{(2)} - \nabla\chi_\ell^{(1)}\|_{L^2}, \end{aligned}$$

$$(A.2) \quad \|\chi_\ell^{(2)} - \chi_\ell^{(1)}\|_{\dot{H}^1} \leq \frac{1}{\alpha} \left(\|A^{(1)} - A^{(2)}\|_{L^\infty}\|\nabla\chi_\ell^{(1)}\|_{L^2} + \|f_\ell^{(1)} - f_\ell^{(2)}\|_{\dot{H}^{-1}} \right).$$

Evaluating,

$$(A.3) \quad \|f_\ell^{(1)} - f_\ell^{(2)}\|_{\dot{H}^{-1}} = \|\nabla \cdot A^{(1)}e_\ell - \nabla \cdot A^{(2)}e_\ell\|_{\dot{H}^{-1}}$$

$$(A.4) \quad = \sup_{\|\xi\|_{\dot{H}^1}=1} \int_{\mathbb{T}^d} \xi \nabla \cdot (A^{(1)} - A^{(2)})e_\ell \, dy$$

$$(A.5) \quad \leq \sup_{\|\xi\|_{\dot{H}^1}=1} \|(A^{(1)} - A^{(2)})e_\ell\|_{L^2} \|\nabla\xi\|_{L^2}$$

$$(A.6) \quad \leq \|A^{(1)} - A^{(2)}\|_{L^2} \leq \|A^{(1)} - A^{(2)}\|_{L^\infty}$$

since our domain is \mathbb{T}^d . Combining this with (A.2) and the bound of $\|\nabla\chi_\ell\|_{L^2} \leq \frac{\beta}{\alpha}$ obtained in the first part of this proposition, we have

$$(A.7) \quad \|\chi_\ell^{(2)} - \chi_\ell^{(1)}\|_{\dot{H}^1} \leq \frac{1}{\alpha} \left(1 + \frac{\beta}{\alpha} \right) \|A^{(1)} - A^{(2)}\|_{L^\infty}.$$

Returning to d vector components yields the result. \square

The following result shows that the mapping $A \mapsto \bar{A}$ is continuous on separable subspaces of $L^\infty(\mathbb{T}^d; \mathbb{R}^{d \times d})$.

LEMMA A.1. Let $\mathcal{A} \subset L^\infty(\mathbb{T}^d; \mathbb{R}^{d \times d})$ be a separable subspace and $K \subset \mathcal{A} \cap \text{PD}_{\alpha, \beta}$ a closed set in L^∞ . Define the mapping $F : K \rightarrow \mathbb{R}^{d \times d}$ by $A \mapsto \bar{A}$ as given by (1.3). Then there exists a continuous mapping $\mathcal{F} \in C(\mathcal{A}; \mathbb{R}^{d \times d})$ such that $\mathcal{F}(A) = F(A)$ for any $A \in K$.

Proof. Let $A^{(1)}, A^{(2)} \in K$; then, by Proposition 1.1,

$$\begin{aligned} |F(A^{(1)}) - F(A^{(2)})|_F &\leq \int_{\mathbb{T}^d} |A^{(1)} - A^{(2)}|_F (1 + |\nabla \chi^{(1)}|_F) \, dy \\ &\quad + \int_{\mathbb{T}^d} |A^{(2)}|_F |\nabla \chi^{(1)} - \nabla \chi^{(2)}|_F \, dy \\ &\leq \|A^{(1)} - A^{(2)}\|_{L^\infty} (1 + \|\nabla \chi^{(1)}\|_{L^2}) + \|A^{(2)}\|_{L^\infty} \|\nabla \chi^{(1)} - \nabla \chi^{(2)}\|_{L^2} \\ &\leq \left(1 + \frac{\sqrt{d}}{\alpha} \left(\|A^{(1)}\|_{L^\infty} + \|A^{(2)}\|_{L^\infty} \left(\frac{\min(\|A^{(1)}\|_{L^\infty}, \|A^{(2)}\|_{L^\infty})}{\alpha} + 1 \right) \right) \right) \\ &\quad \cdot \|A^{(1)} - A^{(2)}\|_{L^\infty}, \end{aligned}$$

hence $F \in C(K; \mathbb{R}^{d \times d})$. Applying the Tietze extension theorem [20] to F implies the existence of \mathcal{F} . □

The following lemma shows that, unfortunately, separable subspaces of $L^\infty(\mathbb{T}^d; \mathbb{R}^{d \times d})$ are not very useful. Indeed, in the desired area of application of continuum mechanics, we ought to be able to place a boundary of material discontinuity anywhere in the domain. The following result shows that doing so is impossible for a subset of $\text{PD}_{\alpha, \beta}$ which lies only in a separable subspace of $L^\infty(\mathbb{T}^d; \mathbb{R}^{d \times d})$.

LEMMA A.2. For any $t \in [0, 1]$ define $c_t : [0, 1] \rightarrow \mathbb{R}$ by

$$c_t(x) = \begin{cases} 1, & x \leq t, \\ 0, & x > t, \end{cases} \quad \forall x \in [0, 1].$$

Define $E = \{c_t : t \in [0, 1]\} \subset L^\infty([0, 1])$. There exists no separable subspace $\mathcal{A} \subset L^\infty([0, 1])$ such that $E \subseteq \mathcal{A}$.

Proof. Suppose otherwise. Since $(\mathcal{A}, \|\cdot\|_{L^\infty})$ is a separable metric space, $(E, \|\cdot\|_{L^\infty})$ must be separable since $E \subseteq \mathcal{A}$; this is a contradiction since $(E, \|\cdot\|_{L^\infty})$ is not separable. To see this, let $\{c_{t_j}\}_{j=1}^\infty$ be an arbitrary countable subset of E . Then for any $t \notin \{t_j\}_{j=1}^\infty$, we have

$$\inf_{\{t_j\}_{j=1}^\infty} \|c_t - c_{t_j}\|_{L^\infty} = 1.$$

Hence no countable subset can be dense. □

Instead of working on a compact subset of a separable subspace of $L^\infty(\mathbb{T}^d; \mathbb{R}^{d \times d})$, we may instead try to find a suitable probability measure which contains the discontinuous functions of interest. The following remark makes clear why such an approach would still be problematic for the purposes of approximation.

Remark A.3 (Gaussian thresholding). Let μ be a Gaussian measure on $L^2([0, 1])$. Define

$$T(x) = \begin{cases} 1, & x \geq 0, \\ 0, & x < 0, \end{cases} \quad \forall x \in [0, 1],$$

and consider the corresponding Nemytskii operator $N_T : L^2([0, 1]) \rightarrow L^\infty([0, 1])$. Then, working with the definitions in Lemma A.2, it is easy to see that $E \subset \text{supp } N_T^\# \mu$. Therefore there exists no separable subspace of $L^\infty([0, 1])$ which contains $\text{supp } N_T^\# \mu$.

We therefore abandon L^∞ and instead show continuity and Lipschitz continuity for some L^q with $q < \infty$ to \dot{H}^1 . The following lemma is a general result for convergence of sequences in metric spaces which is used in a more specific context in the next lemma.

LEMMA A.4. *Let (M, d) be a metric space and $(a_n) \subset M$ a sequence. If every subsequence $(a_{n_k}) \subset (a_n)$ contains a subsequence $(a_{n_{k_l}}) \subset (a_{n_k})$ such that $(a_{n_{k_l}}) \rightarrow a \in M$, then $(a_n) \rightarrow a$.*

Proof. Suppose otherwise. Then, there exists some $\epsilon > 0$ such that, for every $N \in \mathbb{Z}^+$, there exists some $n = n(N) > N$ such that

$$d(a_n, a) \geq \epsilon.$$

Then we can construct a subsequence $(a_{n_j}) \subset (a_n)$ such that $d(a_{n_j}, a) \geq \epsilon$ for all n_j . Therefore a_{n_j} does not have a subsequence converging to a , which is a contradiction. \square

The following lemma proves existence of a limit in $L^2(D; \mathbb{R}^d)$ of a sequence of outputs of operators in $L^\infty(D; \mathbb{R}^{d \times d})$.

LEMMA A.5. *Let $D \subseteq \mathbb{R}^d$ be an open set and $(A_n) \subset L^\infty(D; \mathbb{R}^{d \times d})$ a sequence satisfying the following.*

1. $A_n \in \text{PD}_{\alpha, \beta}$ for all n .

2. There exists $A \in L^\infty(D; \mathbb{R}^{d \times d})$ such that $(A_n) \rightarrow A$ in $L^2(D; \mathbb{R}^{d \times d})$.

Then, for any $g \in L^2(D; \mathbb{R}^d)$, we have that $(A_n g) \rightarrow Ag$ in $L^2(D; \mathbb{R}^d)$.

Proof. We have

$$\|A_n g\|_{L^2} \leq \beta \|g\|_{L^2},$$

hence $(A_n g) \subset L^2(D; \mathbb{R}^d)$, and, similarly, by finite-dimensional norm equivalence, there is a constant $C_1 > 0$ such that

$$\|Ag\|_{L^2} \leq C_1 \|A\|_{L^\infty} \|g\|_{L^2},$$

hence $Ag \in L^2(D; \mathbb{R}^d)$. Again, by finite-dimensional norm equivalence, we have that there exists a constant $C_2 > 0$ such that, for $j \in \{1, \dots, d\}$ and almost every $y \in D$, we have

$$(A_n g)_j(y)^2 \leq |A_n^{(j)}(y)|^2 |g(y)|^2 \leq C_2 \beta^2 |g(y)|^2,$$

where $A_n^{(j)}(y)$ denotes the j th row of $A_n^{(j)}(y)$. In particular,

$$|(A_n g)_j(y)| \leq \sqrt{C_2} \beta |g(y)|.$$

Let $(A_{n_k}) \subset (A_n)$ be an arbitrary subsequence. Since $(A_n) \rightarrow A$, we have that $(A_{n_k}) \rightarrow A$ in $L^2(D; \mathbb{R}^{d \times d})$. Therefore, there exists a subsequence $(A_{n_{k_l}}) \subset (A_{n_k})$ such that $A_{n_{k_l}}(y) \rightarrow A(y)$ for almost every $y \in D$. Then $A_{n_{k_l}}(y)g(y) \rightarrow A(y)g(y)$ for almost every $y \in D$. Since $|g| \in L^2(\mathbb{R}^d)$, we have, by the dominated convergence theorem, that $(A_{n_{k_l}} g)_j \rightarrow (Ag)_j$ in $L^2(D)$ for every $j \in \{1, \dots, d\}$. Therefore $(A_{n_{k_l}} g) \rightarrow Ag$ in $L^2(D; \mathbb{R}^d)$. Since the subsequence (A_{n_k}) was arbitrary, Lemma A.4 implies the result. \square

Finally, we may prove Proposition 1.2.

PROPOSITION 1.2. *Endow $\text{PD}_{\alpha,\beta}$ with the $L^2(\mathbb{T}^d; \mathbb{R}^{d \times d})$ induced topology and let $K \subset \text{PD}_{\alpha,\beta}$ be a closed set. Define the mapping $G : K \rightarrow \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)$ by $A \mapsto \chi$ as given by (1.4). Then there exists a bounded continuous mapping*

$$\mathcal{G} \in C(L^2(\mathbb{T}^d; \mathbb{R}^{d \times d}); \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d))$$

such that $\mathcal{G}(A) = G(A)$ for any $A \in K$.

Proof. Consider the PDE

$$(A.8) \quad -\nabla \cdot (A \nabla u) = \nabla \cdot A e, \quad y \in \mathbb{T}^d,$$

where e is some standard basis vector of \mathbb{R}^d . Let $(A_n) \subset K$ be a sequence such that $(A_n) \rightarrow A \in K$ in $L^2(\mathbb{T}^d; \mathbb{R}^{d \times d})$. Denote by $u_n \in \dot{H}^1(\mathbb{T}^d)$ the solution to (A.8) corresponding to each A_n and by $u \in \dot{H}^1(\mathbb{T}^d)$ the solution corresponding to the limiting A . A similar calculation as in the proof of Proposition 1.1 shows

$$\begin{aligned} \alpha \|u_n - u\|_{\dot{H}^1}^2 &\leq \int_{\mathbb{T}^d} \langle (A - A_n)(\nabla u + e), \nabla u_n - \nabla u \rangle \, dy \\ &\leq \|u_n - u\|_{\dot{H}^1} \|(A_n - A)(\nabla u + e)\|_{L^2}. \end{aligned}$$

Since $\nabla u + e \in L^2(\mathbb{T}^d; \mathbb{R}^d)$, by Lemma A.5, $(A_n(\nabla u + e)) \rightarrow A(\nabla u + e)$ in $L^2(\mathbb{T}^d; \mathbb{R}^d)$, hence $(u_n) \rightarrow u$ in $\dot{H}^1(\mathbb{T}^d)$. In particular, the mapping $A \mapsto u$ defined by (A.8) is continuous. Since the problem (1.4) decouples as shown by (A.1), we have that each component mapping $G_l : K \rightarrow \dot{H}^1(\mathbb{T}^d)$ defined by $A \mapsto \chi_\ell$ is continuous, thus G is continuous. Applying the Tietze extension theorem [20] to G implies the existence of \mathcal{G} . \square

The following is a straightforward consequence of Proposition 1.2 that establishes continuity of the map $A \mapsto \bar{A}$ defined in (1.3) as well.

LEMMA A.6. *Endow $\text{PD}_{\alpha,\beta}$ with the $L^2(\mathbb{T}^d; \mathbb{R}^{d \times d})$ induced topology and let $K \subset \text{PD}_{\alpha,\beta}$ be a closed set. Define the mapping $F : K \rightarrow \mathbb{R}^{d \times d}$ by $A \mapsto \bar{A}$ as given by (1.3). Then there exists a bounded continuous mapping $\mathcal{F} \in C(L^2(\mathbb{T}^d; \mathbb{R}^{d \times d}); \mathbb{R}^{d \times d})$ such that $\mathcal{F}(A) = F(A)$ for any $A \in K$.*

Proof. Since $\nabla : \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{d \times d})$ is a bounded operator, Lemma 1.2 implies that the mapping $A \mapsto A + A \nabla \chi^T$ is continuous as compositions, sums, and products of continuous functions are continuous. Now let $A \in \text{PD}_{\alpha,\beta}$; then $A \in L^1(\mathbb{T}^d; \mathbb{R}^{d \times d})$ since $A \in L^\infty(\mathbb{T}^d; \mathbb{R}^{d \times d})$. Thus

$$\left| \int_{\mathbb{T}^d} A \, dy \right|_F \leq \int_{\mathbb{T}^d} |A|_F \, dy \leq \|A\|_{L^2}$$

by Hölder’s inequality and the fact that $\int_{\mathbb{T}^d} \, dy = 1$. Hence $F \in C(K; \mathbb{R}^{d \times d})$ as a composition of continuous maps. Again applying the Tietze extension theorem [20] to F implies the existence of \mathcal{F} . \square

To prove Proposition 1.3, we need to establish Lipschitz continuity. We first establish the following result, which is similar to the one proved in [11] in Theorem 2.1. We show it again here both for completeness and because we specialize to the case of the cell problem (1.4) with periodic boundary conditions rather than the system (1.1) with Dirichlet boundary conditions.

LEMMA A.7. Let $A^{(1)}, A^{(2)} \in \text{PD}_{\alpha, \beta}$ and let $\chi^{(1)}, \chi^{(2)}$ be the corresponding solutions to (1.4). Then

$$(A.9) \quad \|\chi^{(1)} - \chi^{(2)}\|_{\dot{H}^1} \leq \frac{\sqrt{d}}{\alpha} \left(\|A^{(2)} - A^{(1)}\|_{L^2} + \|\nabla \chi^{(2)}\|_{L^p} \|A^{(2)} - A^{(1)}\|_{L^q} \right)$$

for $p \geq 2$ and $q = \frac{2p}{p-2}$.

Proof. As in the proof of Proposition 1.1, we denote $f^{(i)} = \nabla \cdot A^{(i)}$ for $i \in \{1, 2\}$ for simplicity of notation and to be easily comparable to the proof of Theorem 2.1 in [11]. Since both sides of the cell problem equation (1.4) depend on $A^{(i)}$, we introduce $\tilde{\chi}$ as the solution of

$$(A.10) \quad -\nabla \cdot (\nabla \tilde{\chi} A^{(2)}) = \nabla \cdot A^{(1)}, \quad \tilde{\chi} \in \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)$$

as an intermediate function. We obtain bounds using $\tilde{\chi}$ and apply the triangle inequality to

$$\|(\chi^{(1)} - \tilde{\chi}) + (\tilde{\chi} - \chi^{(2)})\|_{\dot{H}^1}$$

to obtain a bound on $\|\chi^{(1)} - \chi^{(2)}\|_{\dot{H}^1}$. From the naïve perturbation bound in (A.2) we have

$$\|\tilde{\chi}_\ell - \chi_\ell^{(2)}\|_{\dot{H}^1} \leq \frac{1}{\alpha} \|f_\ell^{(1)} - f_\ell^{(2)}\|_{\dot{H}^{-1}},$$

so we are left to bound $\|\chi_\ell^{(1)} - \tilde{\chi}_\ell\|_{\dot{H}^1}$. We note that

$$\begin{aligned} \nabla \cdot (A^{(2)} \nabla \tilde{\chi}_\ell) &= \nabla \cdot (A^{(1)} \nabla \chi_\ell^{(1)}), \\ \int_{\mathbb{T}^d} A^{(2)} \nabla \tilde{\chi}_\ell \cdot \nabla v \, dy &= \int_{\mathbb{T}^d} A^{(1)} \nabla \chi_\ell^{(1)} \cdot \nabla v \, dy \quad \forall v \in \dot{H}^1(\mathbb{T}^d; \mathbb{R}). \end{aligned}$$

Letting $v = \chi_\ell^{(1)} - \tilde{\chi}_\ell$,

$$\begin{aligned} \int_{\mathbb{T}^d} A^{(2)} \nabla \tilde{\chi}_\ell \cdot (\nabla \chi_\ell^{(1)} - \nabla \tilde{\chi}_\ell) \, dy &= \int_{\mathbb{T}^d} A^{(1)} \nabla \chi_\ell^{(1)} \cdot (\nabla \chi_\ell^{(1)} - \nabla \tilde{\chi}_\ell) \, dy, \\ \int_{\mathbb{T}^d} A^{(2)} (\nabla \tilde{\chi}_\ell - \nabla \chi_\ell^{(1)}) \cdot (\nabla \tilde{\chi}_\ell - \nabla \chi_\ell^{(1)}) \, dy &= \int_{\mathbb{T}^d} (A^{(2)} - A^{(1)}) \nabla \chi_\ell^{(1)} \cdot (\nabla \chi_\ell^{(1)} - \nabla \tilde{\chi}_\ell) \, dy, \\ \alpha \|\tilde{\chi}_\ell - \chi_\ell^{(1)}\|_{\dot{H}^1} &\leq \|(A^{(2)} - A^{(1)}) (\nabla \chi_\ell^{(1)})\|_{L^2}. \end{aligned}$$

Applying Hölder for L^2 , we get

$$(A.11) \quad \|\tilde{\chi}_\ell - \chi_\ell^{(1)}\|_{\dot{H}^1} \leq \frac{1}{\alpha} \|\nabla \chi_\ell^{(1)}\|_{L^p} \|A^{(2)} - A^{(1)}\|_{L^q}$$

for $q = \frac{2p}{p-2}$, where $p \in [2, \infty]$. Putting the two parts together, we have that

$$\begin{aligned} \|\chi_\ell^{(2)} - \chi_\ell^{(1)}\|_{\dot{H}^1} &\leq \frac{1}{\alpha} \|\nabla \cdot A^{(2)} e_\ell - \nabla \cdot A^{(1)} e_\ell\|_{\dot{H}^{-1}} + \frac{1}{\alpha} \|\nabla \chi_\ell^{(1)}\|_{L^p} \|A^{(2)} - A^{(1)}\|_{L^q} \\ &\leq \frac{1}{\alpha} \|A^{(2)} - A^{(1)}\|_{L^2} + \frac{1}{\alpha} \|\nabla \chi_\ell^{(1)}\|_{L^p} \|A^{(2)} - A^{(1)}\|_{L^q}. \end{aligned}$$

Combining bounds for all d dimensions yields the result. \square

Remark A.8. Since $L^q(\Omega) \hookrightarrow L^2(\Omega)$ for bounded $\Omega \subset \mathbb{R}^d$ and $q \geq 2$, we could also write the bound of Lemma A.7 as

$$\|\chi_\ell^{(2)} - \chi_\ell^{(1)}\|_{\dot{H}^1} \leq \frac{1}{\alpha} \left(C + \|\nabla \chi_\ell^{(1)}\|_{L^p} \right) \|A^{(2)} - A^{(1)}\|_{L^q}$$

for some C dependent only on q and Ω .

The result of Lemma A.7 is unhelpful if $\|\nabla \chi\|_{L^p}$ is unbounded. In this setting, it is not possible for Lemma A.7 to result in Lipschitz continuity as a map from L^2 to \dot{H}^1 . Instead, we seek to bound $\|\nabla \chi\|_{L^p}$ for some p satisfying $2 < p < \infty$.

Before continuing, we establish a bound on the gradient of the solution to the Poisson equation on the torus. This follows the strategy of [11] for the Dirichlet problem. In order to avoid extra factors of 2π in all formulae, we work on the rescaled torus denoted $\mathbb{Y}^d = [0, 2\pi]^d$ with opposite faces identified for the following result of Lemma A.9. As we work on the torus, it is useful to first set up notation for the function spaces of interest. Let

$$\mathcal{D}(\mathbb{Y}^d) = C_c^\infty(\mathbb{Y}^d) = C^\infty(\mathbb{Y}^d)$$

be the space of test functions, where the last equality follows from compactness of the torus. Functions can be either \mathbb{R} or \mathbb{C} valued, hence we do not explicitly specify the range. We equip $\mathcal{D}(\mathbb{Y}^d)$ with a locally convex topology generated by an appropriate family of seminorms; see, for example, [56, section 3.2.1]. Any function $g \in \mathcal{D}(\mathbb{Y}^d)$ can be represented by its Fourier series

$$g(x) = \sum_{k \in \mathbb{Z}^d} \hat{g}(k) e^{ix \cdot k},$$

where \hat{g} denotes the Fourier transform of g and convergence of the right-hand-side sum is with respect to the topology of $\mathcal{D}(\mathbb{Y}^d)$, and i denotes the imaginary unit. It holds that $\hat{g} \in \mathcal{S}(\mathbb{Z}^d)$, the Schwartz space of rapidly decreasing functions on the integer lattice, so we have

$$|\hat{g}(k)| \leq c_m (1 + |k|)^{-m}, \quad m = 0, 1, \dots,$$

for some constants c_m . We may then define the topological (continuous) dual space of $\mathcal{D}(\mathbb{Y}^d)$, the space of distributions, denoted $\mathcal{D}'(\mathbb{Y}^d)$, which can be described as follows: the condition that $f \in \mathcal{D}'(\mathbb{Y}^d)$ is characterized by the property

$$|\hat{f}(k)| \leq b_m (1 + |k|)^m, \quad m = 0, 1, \dots,$$

for some constants b_m . We take the weak-* topology on $\mathcal{D}'(\mathbb{Y}^d)$ and generally use the prime notation for any such dual space. For any $-\infty < s < \infty$, we define the fractional Laplacian as

$$(A.12) \quad (-\Delta)^s f = \sum_{k \in \mathbb{Z}^d \setminus \{0\}} |k|^{2s} \hat{f}(k) e^{ik \cdot x},$$

where the right-hand-side sum converges in the topology of $\mathcal{D}'(\mathbb{Y}^d)$. It is easy to see that $(-\Delta)^s : \mathcal{D}'(\mathbb{Y}^d) \rightarrow \mathcal{D}'(\mathbb{Y}^d)$ is continuous. Furthermore, for any $j \in \{1, \dots, d\}$, we define the family of operators $\tilde{R}_j : \mathcal{D}'(\mathbb{Y}^d) \rightarrow \mathcal{D}'(\mathbb{Y}^d)$, defining periodic Riesz transforms, by

$$(A.13) \quad \tilde{R}_j f = \sum_{k \in \mathbb{Z}^d} -\frac{ik_j}{|k|} \hat{f}(k) e^{ik \cdot x},$$

where we identify $\frac{k_j}{|k|}|_{k=0} = \lim_{|k| \rightarrow 0} \frac{k_j}{|k|} = 0$. Again, we stress that convergence of the right-hand-side sum is in the topology of $\mathcal{D}'(\mathbb{Y}^d)$. Last, we denote by $\mathcal{S}(\mathbb{R}^d)$ and $\mathcal{S}'(\mathbb{R}^d)$ the Schwartz space and the space of tempered distributions on \mathbb{R}^d , respectively; see, for example, [57, Chapter 1] for the precise definitions.

The following lemma establishes boundedness of the periodic Riesz transform on $L^p(\mathbb{Y}^d)$. It is essential in proving boundedness of the gradient to the solution of the Poisson equation on the torus. The result is essentially proven in [57]. We include it here, in our specific torus setting, giving the full argument for completeness.

LEMMA A.9. *There exists a constant $c = c(d, p) > 0$ such that, for any $j \in \{1, \dots, d\}$ and any $f \in L^p(\mathbb{Y}^d)$ for some $2 \leq p < \infty$, we have*

$$\|\tilde{R}_j f\|_{L^p(\mathbb{Y}^d)} \leq c \|f\|_{L^p(\mathbb{Y}^d)}.$$

Proof. Let $g \in L^2(\mathbb{R}^d) \cap L^p(\mathbb{R}^d)$ for some $1 < p < \infty$. For any $j \in \{1, \dots, d\}$, define the family of operators R_j by

$$(R_j g)(x) = \lim_{\delta^{-1}, \epsilon \rightarrow 0^+} \int_{\delta \geq |t| \geq \epsilon} g(x-t) K_j(t) dt,$$

where

$$K_j(t) = \frac{\Gamma((d+1)/2) t_j}{\pi^{(d+1)/2} |t|^{d+1}}$$

and Γ denotes the Euler-Gamma function. By [57, Chapter 4, Theorem 4.5], $K_j \in \mathcal{S}'(\mathbb{R}^d)$ and its Fourier transform satisfy

$$\widehat{K}_j(t) = -\frac{it_j}{|t|}.$$

Therefore, for any $\phi \in \mathcal{S}(\mathbb{R}^d)$, we have

$$(K_j * \phi)(t) = -\frac{it_j}{|t|} \widehat{\phi}(t),$$

where $*$ denotes convolution; see, for example, [57, Chapter 1, Theorem 3.18]. Since $g \in L^2(\mathbb{R}^d)$, we therefore find that, by [57, Chapter 6, Theorem 2.6],

$$(A.14) \quad (R_j g)(x) = -\frac{ix_j}{|x|} \widehat{g}(x)$$

for Lebesgue almost every $x \in \mathbb{R}^d$. The result [57, Chapter 6, Theorem 2.6] further shows that there exists a constant $c = c(d, p) > 0$ such that

$$\|R_j g\|_{L^p(\mathbb{R}^d)} \leq c \|g\|_{L^p(\mathbb{R}^d)}.$$

We note from (A.14) and the definition (A.13) that \tilde{R}_j may be viewed as R_j with the restriction of the Fourier multiplier $-\frac{ix_j}{|x|}$ to the lattice \mathbb{Z}^d . We can therefore use the transference theory of [57] to establish boundedness of \tilde{R}_j from the boundedness of R_j . In particular, note that the mapping $x \mapsto -\frac{ix_j}{|x|}$ is continuous at all $x \in \mathbb{R}^d$ except $x = 0$. However, by symmetry, we have that, for all $\epsilon > 0$,

$$\int_{|x| \leq \epsilon} -\frac{ix_j}{|x|} dx = 0.$$

Therefore we can apply [57, Chapter 7, Theorem 3.8, Corollary 3.16] to conclude that, since R_j is bounded from $L^p(\mathbb{R}^d)$ to $L^p(\mathbb{R}^d)$, \tilde{R}_j is bounded from $L^p(\mathbb{Y}^d)$ to $L^p(\mathbb{Y}^d)$ with

$$\|\tilde{R}_j\|_{L^p(\mathbb{Y}^d) \rightarrow L^p(\mathbb{Y}^d)} \leq \|R_j\|_{L^p(\mathbb{R}^d) \rightarrow L^p(\mathbb{R}^d)}.$$

This implies the desired result. □

We define the Bessel potential spaces by

$$L^{s,p}(\mathbb{Y}^d) = \{u \in \mathcal{D}'(\mathbb{Y}^d) \mid \|u\|_{L^{s,p}(\mathbb{Y}^d)} := \|(I - \Delta)^{s/2}u\|_{L^p(\mathbb{Y}^d)} < \infty\}$$

for any $-\infty < s < \infty$ and $1 < p < \infty$. We also define the homogeneous version of these spaces, sometimes called the Riesz potential spaces, by

$$\dot{L}^{s,p}(\mathbb{Y}^d) = \left\{ u \in \mathcal{D}'(\mathbb{Y}^d) \mid \|u\|_{\dot{L}^{s,p}(\mathbb{Y}^d)} := \|(-\Delta)^{s/2}u\|_{L^p(\mathbb{Y}^d)} < \infty, \int_{\mathbb{Y}^d} u \, dy = 0 \right\}.$$

It is clear that $\dot{L}^{s,p}(\mathbb{Y}^d) \subset L^{s,p}(\mathbb{Y}^d)$ is closed subspace. We then have the following result for the Poisson equation.

LEMMA A.10. *For each $f \in L^{s,p}(\mathbb{Y}^d)$, for $-\infty < s < \infty$ and $2 \leq p < \infty$, the solution u of the equation*

$$(A.15) \quad -\Delta u = f, \quad u \text{ 1-periodic}, \quad \int_{\mathbb{Y}^d} u \, dy = 0$$

satisfies

$$(A.16) \quad \|\nabla u\|_{\dot{L}^{s+1,p}(\mathbb{Y}^d)} \leq K \|f\|_{\dot{L}^{s,p}(\mathbb{Y}^d)}$$

for some finite $K > 0$ depending only on p and d .

Proof. From the definitions (A.12) and (A.13), it is easy to see that the Riesz transform can be written as

$$\tilde{R}_j = -\partial_{x_j}(-\Delta)^{-1/2}$$

in the sense of distributions. Consider now (A.15) with $f \in L^{s,p}(\mathbb{Y}^d)$ for $2 \leq p < \infty$. We have that

$$\begin{aligned} \|\partial_{x_j} u\|_{\dot{L}^{s+1,p}(\mathbb{Y}^d)} &= \|\partial_{x_j}(-\Delta)^{-1} f\|_{\dot{L}^{s+1,p}(\mathbb{Y}^d)} \\ &= \|\partial_{x_j}(-\Delta)^{-1/2}(-\Delta)^{s/2} f\|_{L^p(\mathbb{Y}^d)} \\ &= \|\tilde{R}_j(-\Delta)^{s/2} f\|_{L^p(\mathbb{Y}^d)}. \end{aligned}$$

It is clear that

$$\|(-\Delta)^{s/2} f\|_{L^p(\mathbb{Y}^d)} = \|f\|_{\dot{L}^{s,p}(\mathbb{Y}^d)} < \infty,$$

hence $(-\Delta)^{s/2} f \in L^p(\mathbb{Y}^d)$. We can thus apply Lemma A.9 to find a constant $c = c(d, p) > 0$ such that

$$\|\partial_{x_j} u\|_{\dot{L}^{s+1,p}(\mathbb{Y}^d)} \leq c \|(-\Delta)^{s/2} f\|_{L^p(\mathbb{Y}^d)} = c \|f\|_{\dot{L}^{s,p}(\mathbb{Y}^d)}.$$

The result follows by finite-dimensional norm equivalence. □

Next we define the homogeneous Sobolev spaces on the torus as

$$(A.17) \quad \dot{W}^{k,p}(\mathbb{T}^d) = \left\{ u \in W^{k,p}(\mathbb{T}^d) \mid u \text{ is 1-periodic, } \int_{\mathbb{T}^d} u \, dy = 0 \right\}$$

for $k = 0, 1, \dots$, and $1 \leq p \leq \infty$ with the standard norm on $W^{k,p}$; see, for example, [2].

Remark A.11. By [56, section 3.5.4], we have that, for any $k = 0, 1, \dots$ and $1 < p < \infty$,

$$L^{k,p}(\mathbb{T}^d) = W^{k,p}(\mathbb{T}^d), \quad \dot{L}^{k,p}(\mathbb{T}^d) = \dot{W}^{k,p}(\mathbb{T}^d).$$

Furthermore, by [56, section 3.5.6],

$$\begin{aligned} W^{-k,p'}(\mathbb{T}^d) &= (W^{k,p}(\mathbb{T}^d))' = (L^{k,p}(\mathbb{T}^d))' = L^{-k,p'}(\mathbb{T}^d), \\ \dot{W}^{-k,p'}(\mathbb{T}^d) &= (\dot{W}^{k,p}(\mathbb{T}^d))' = (\dot{L}^{k,p}(\mathbb{T}^d))' = \dot{L}^{-k,p'}(\mathbb{T}^d), \end{aligned}$$

where p' is the Hölder conjugate of p , i.e., $1/p + 1/p' = 1$.

In the following, we use the notation

$$(A.18) \quad [K_0, K_1]_{\theta,q}$$

to denote the real interpolation between two Banach spaces continuously embedded in the same Hausdorff topological space, as described in [2]. We also need Lemma A1 from [22], which we have copied below as Lemma A.12 to ease readability. Although this lemma was written only for $q = 2$, the result still holds for our $q > 2$ with a very similar proof.

LEMMA A.12. *Let $E_1 \subset E_0$ be two Banach spaces with E_1 continuously embedded in E_0 . Let $T : E_j \rightarrow E_j$ be a bounded operator with closed range and assume that T is a projection, $j \in \{0, 1\}$. Denote by K_0 and K_1 the ranges of $T|_{E_0}$ and $T|_{E_1}$, respectively. Then the following two spaces coincide with equivalent norms:*

$$[K_0, K_1]_{\theta,q} = [E_0, E_1]_{\theta,q} \cap K_0 \quad \forall \theta \in (0, 1).$$

We now state the result for the bound on $\|\nabla \chi\|_{L^p}$ with a proof largely developed in [11].

LEMMA A.13. *Let χ solve (1.4) for $A \in \text{PD}_{\alpha,\beta}$. Then*

$$(A.19) \quad \|\nabla \chi\|_{L^p} \leq \frac{K^{\eta(p)}}{1 - K^{\eta(p)} \left(1 - \frac{\alpha}{\beta}\right)}$$

for $2 \leq p < p^*(\frac{\alpha}{\beta})$ where

$$(A.20) \quad p^*(t) := \max \left\{ p \mid K^{-\eta(p)} \geq 1 - t, 2 < p < Q \right\}$$

for $\eta(p) = \frac{1/2-1/p}{1/2-1/Q}$ and $K = K(d, Q)$ is the constant in Lemma A.10 for any choice of $Q > p$.

Proof. The operator $T = -\Delta$ is invertible from H^{-1} to \dot{H}^1 , and the inverse T^{-1} is bounded with norm 1 since the Poisson equation with periodic boundary conditions has a unique solution in \dot{H}^1 for $f \in H^{-1}$ with bound $\|u\|_{\dot{H}^1} \leq \|f\|_{H^{-1}}$.

From Lemma A.10 it is also bounded with norm $K = K(d, Q)$ from $W^{-1, Q}$ to $\dot{W}^{1, Q}$ for any $Q > 2$. By the real method of interpolation [2], for $2 < p < Q$ we have that

$$(A.21) \quad W^{1, p} = [H^1, W^{1, Q}]_{\eta(p), p}$$

using the notation of [2] where $\eta(p) = \frac{1/2-1/p}{1/2-1/Q}$. From the duality theorem (Theorem 3.7.1 of [6]), we have that

$$(A.22) \quad [H^{-1}, W^{-1, Q}]_{\eta(p), p} = \left([H^1, W^{1, Q'}]_{\eta(p), p'} \right)'$$

From real interpolation, the right-hand side equals $(W^{1, p'})' = W^{-1, p}$ in our notation. Therefore, we have the necessary dual statement that parallels (A.21):

$$(A.23) \quad W^{-1, p} = [H^{-1}, W^{-1, Q}]_{\eta(p), p}$$

Next we restrict these spaces to functions with periodic boundary conditions. Using the projection onto the space of continuous, periodic functions on \mathbb{T}^d and noticing that $W^{1, Q} \hookrightarrow H^1$, we apply Lemma A.12 with $K_0 = \dot{H}^1$ and have

$$(A.24) \quad \dot{W}^{1, p} = [\dot{H}^1, \dot{W}^{1, Q}]_{\eta(p), p}$$

Using the exact interpolation theorem, Theorem 7.23 of [2], T^{-1} is also a bounded map from $W^{-1, p}$ to $\dot{W}^{1, p}$ with norm $K^{\eta(p)}$:

$$(A.25) \quad \|T^{-1}f\|_{\dot{W}^{1, p}} \leq K^{\eta(p)} \|f\|_{W^{-1, p}}$$

The remainder of the proof is identical to that of the proof of Proposition 1 in [11], but we state it here in our notation for completeness. Define $S: \dot{W}^{1, p} \rightarrow W^{-1, p}$ as the operator $Su = -\nabla \cdot (\frac{1}{\beta} A \nabla u)$. Let V be the perturbation operator $V := T - S$. Since $A \in \text{PD}_{\alpha, \beta}$, S and V are bounded operators from $\dot{W}^{1, p}$ to $W^{-1, p}$, with the operator norms $\|S\| \leq 1$ and $\|V\| \leq 1 - \frac{\alpha}{\beta}$. Therefore,

$$(A.26) \quad \|T^{-1}V\|_{\dot{W}^{1, p} \rightarrow \dot{W}^{1, p}} \leq \|T^{-1}\|_{W^{-1, p} \rightarrow \dot{W}^{1, p}} \|V\|_{\dot{W}^{1, p} \rightarrow W^{-1, p}} \leq K^{\eta(p)} \left(1 - \frac{\alpha}{\beta}\right),$$

where the input and output spaces defining the operator norms are included for clarity. Since T is invertible, $S = T(I - T^{-1}V)$ is invertible provided $K^{\eta(p)}(1 - \frac{\alpha}{\beta}) < 1$. Moreover, for S^{-1} as a mapping from $W^{-1, p}$ to $\dot{W}^{1, p}$,

$$(A.27) \quad \|S^{-1}\| \leq \|(I - T^{-1}V)^{-1}\|_{\dot{W}^{1, p} \rightarrow \dot{W}^{1, p}} \|T^{-1}\|_{W^{-1, p} \rightarrow \dot{W}^{1, p}} \leq \frac{K^{\eta(p)}}{1 - K^{\eta(p)} \left(1 - \frac{\alpha}{\beta}\right)}.$$

Therefore,

$$(A.28) \quad \|\nabla \chi\|_{L^p} = \|\chi\|_{\dot{W}^{1, p}} \leq \frac{1}{\beta} \|S^{-1}\| \|\nabla \cdot A\| \leq \frac{K^{\eta(p)}}{1 - K^{\eta(p)} \left(1 - \frac{\alpha}{\beta}\right)}$$

provided $K^{\eta(p)}(1 - \frac{\alpha}{\beta}) < 1$. The bound and specified range of p follow. □

Finally, we may prove Proposition 1.3

PROPOSITION 1.3. *There exists $q_0 \in (2, \infty)$ such that, for all q satisfying $q \in (q_0, \infty]$, the following holds. Endow $\text{PD}_{\alpha, \beta}$ with the $L^q(\mathbb{T}^d; \mathbb{R}^{d \times d})$ topology and let $K \subset \text{PD}_{\alpha, \beta}$ be a closed set. Define the mapping $G: K \rightarrow \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)$ by $A \mapsto \chi$ as given by (1.4). Then there exists a bounded Lipschitz-continuous mapping*

$$\mathcal{G} : L^q(\mathbb{T}^d; \mathbb{R}^{d \times d}) \rightarrow \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)$$

such that $\mathcal{G}(A) = G(A)$ for any $A \in K$.

Proof. Lemma A.13 guarantees a $p_0 > 2$ such that $\|\nabla \chi^{(2)}\|_{L^p}$ in Lemma A.7 is bounded above by a constant for $2 < p < p_0$. Then Lemma A.7 gives Lipschitz continuity of the solution map from $L^q(\mathbb{T}^d) \mapsto \dot{H}^1(\mathbb{T}^d)$ for q satisfying $q_0 < q < \infty$ for some $q_0 > 2$. \square

Remark A.14. From the results of Lemmas A.13 and A.7, we have that we can take $q_0 = \frac{2p_0}{p_0-2}$ where

$$p_0 = \max\{p \mid K^{-\eta(p)} \geq 1 - t, 2 < p < Q\}.$$

Therefore, bounds on p_0 may be inherited from bounds on K that appear in Lemma A.10.

Appendix B. Proofs of approximation theorems. In this section we prove the approximation theorems stated in section 3.

THEOREM 3.3. *Let $K \subset \text{PD}_{\alpha, \beta}$ and define the mapping $G : K \rightarrow \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)$ by $A \mapsto \chi$ as given by (1.4). Assume in addition that K is compact in $L^2(\mathbb{T}^d; \mathbb{R}^{d \times d})$. Then, for any $\epsilon > 0$, there exists an FNO $\Psi : K \rightarrow \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)$ such that*

$$\sup_{A \in K} \|G(A) - \Psi(A)\|_{\dot{H}^1} < \epsilon.$$

Proof. By Proposition 1.2, there exists a continuous map $\mathcal{G} \in C(L^2(\mathbb{T}^d; \mathbb{R}^{d \times d}); \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d))$ such that $\mathcal{G}(A) = G(A)$ for any $A \in K$. By [32, Theorem 5], there exists an FNO $\Psi : L^2(\mathbb{T}^d; \mathbb{R}^{d \times d}) \rightarrow \dot{H}^1(\mathbb{T}^d; \mathbb{R}^d)$ such that

$$\sup_{A \in K} \|\mathcal{G}(A) - \Psi(A)\|_{\dot{H}^1} < \epsilon.$$

Therefore

$$\sup_{A \in K} \|G(A) - \Psi(A)\|_{\dot{H}^1} = \sup_{A \in K} \|\mathcal{G}(A) - \Psi(A)\|_{\dot{H}^1} < \epsilon$$

as desired. \square

THEOREM 3.4. *Let $K \subset \text{PD}_{\alpha, \beta}$ and define the mapping $F : K \rightarrow \mathbb{R}^{d \times d}$ by $A \mapsto \bar{A}$ as given by (1.3), (1.4). Assume in addition that K is compact in $L^2(\mathbb{T}^d; \mathbb{R}^{d \times d})$. Then, for any $\epsilon > 0$, there exists an FNO $\Phi : K \rightarrow L^\infty(\mathbb{T}^d; \mathbb{R}^{d \times d})$ such that*

$$\sup_{A \in K} \sup_{x \in \mathbb{T}^d} |F(A) - \Phi(A)(x)|_F < \epsilon.$$

Proof. The result follows as in Theorem 3.3 by applying Lemma A.6 instead of Proposition 1.2. \square

Appendix C. Proofs for microstructure examples. The following lemma establishes the compactness of subsets of $\text{PD}_{\alpha, \beta}$ generated by the probability measures from section 4. As we are unaware of a proof in the literature, we have provided one below. The proof uses the L^1 -Lipschitz spaces, which are defined as

$$\text{Lip}_\alpha(L^1) = \{u \in L^1 : \exists M(u) > 0 : \omega(u, t)_1 \leq Mt^\alpha\},$$

where $\omega(u, t)_1$ is the 1-modulus of continuity, defined via

$$\omega(u, t)_1 = \sup_{0 \leq |h| \leq t} \|\tau_h u - u\|_{L^1(\mathbb{T}^d)}.$$

LEMMA C.1. $\text{BV}(\mathbb{T}^d) \cap L^\infty(\mathbb{T}^d)$ is compactly embedded in $L^2(\mathbb{T}^d)$.

Proof. Let $u \in B$, where B is a bounded subset of $\text{BV}(\mathbb{T}^d) \cap L^\infty(\mathbb{T}^d)$ with L^∞ norm and BV seminorm bounded by M , and let $\tau_h f$ denote the translation of f by h , i.e., $\tau_h f(x) = f(x - h)$. Then

$$(C.1) \quad \|\tau_h u - u\|_{L^2} \leq \|\tau_h u - u\|_{L^1}^{1/2} \|\tau_h u - u\|_{L^\infty}^{1/2}.$$

Since $\text{BV}(\mathbb{T}^d) \equiv \text{Lip}_1(L^1(\mathbb{T}^d))$, $\|\tau_h u - u\|_{L^1} \leq \|u\|_{\text{BV}}|h|$. We have then

$$\|\tau_h u - u\|_{L^2} \leq \|u\|_{\text{BV}}^{1/2} |h|^{1/2} (2M)^{1/2}.$$

By the Fréchet–Kolmogorov theorem [60], this equicontinuity result is sufficient for compactness of B in $L^2(\mathbb{T}^d)$. \square

Using the result of Lemma C.1, we see that any set of microstructure coefficients bounded in $L^\infty(\mathbb{T}^d) \cap BV(\mathbb{T}^d)$ satisfies the compactness assumption of the approximation theorems in section 3. It is clear that the method of construction of the examples in subsection 4.1 leads to such sets.

Appendix D. Numerical implementation details. All FNO models are implemented in PyTorch using Python 3.9.7. Unless otherwise specified, the models have 18 modes in each dimension, a width of 64, and 4 hidden layers. The lifting layer is a linear transformation with trainable parameters, and the projecting layer is a pointwise multilayer perceptron with trainable parameters. The batch size is 20, the learning rate is 0.001, and the number of epochs is 400. These hyperparameters are chosen with a small grid search, but we emphasize that the FNO does not drastically change in performance unless these parameters are changed by an order of magnitude. For a model trained on 9500 data using these hyperparameters and accelerated with an Nvidia P100 GPU, the training time is approximately 7 hours. In Figures 3, 4, and 5, the error bars shown correspond to two standard deviations in each direction over the five samples. All code for this work may be found at github.com/mtrautner/LearningHomogenization/.

Acknowledgments. The authors are grateful to Samuel Lanthaler and Endre Süli for essential discussions leading to the proof of Lemma C.1.

REFERENCES

- [1] R. A. ADAMS AND J. J. FOURNIER, *Sobolev Spaces*, Elsevier, Amsterdam, 2003.
- [2] G. ALLAIRE, *Homogenization and two-scale convergence*, SIAM J. Math. Anal., 23 (1992), pp. 1482–1518.
- [3] F. AS'AD, P. AVERY, AND C. FARHAT, *A mechanics-informed artificial neural network approach in data-driven constitutive modeling*, Internat. J. Numer. Methods Engrg., 123 (2022), pp. 2738–2759.
- [4] P. BATLLE, M. DARCY, B. HOSSEINI, AND H. OWHADI, *Kernel methods are competitive for operator learning*, J. Comput. Phys., 496 (2024), 112549.
- [5] A. BENSOUSSAN, J.-L. LIONS, AND G. PAPANICOLAOU, *Asymptotic Analysis for Periodic Structures*, AMS, Providence, RI, 2011.
- [6] J. BERGH AND J. LÖFSTRÖM, *Interpolation Spaces: An Introduction*, Grundlehren Math. Wiss. 223, Springer, Berlin, 2012.
- [7] K. BHATTACHARYA, *Phase boundary propagation in a heterogeneous body*, Proc. A, 455 (1999), pp. 757–766.
- [8] K. BHATTACHARYA, B. HOSSEINI, N. B. KOVACHKI, AND A. M. STUART, *Model reduction and neural networks for parametric pdes*, SMAI J. Comput. Math., 7 (2021), pp. 121–157.
- [9] K. BHATTACHARYA, B. LIU, A. STUART, AND M. TRAUTNER, *Learning Markovian homogenized models in viscoelasticity*, Multiscale Model. Simul., 21 (2023), pp. 641–679.

- [10] X. BLANC AND C. LE BRIS, *Homogenization Theory for Multiscale Problems: An Introduction*, MS&A 21, Springer, New York, 2023.
- [11] A. BONITO, R. A. DEVORE, AND R. H. NOCHETTO, *Adaptive finite element methods for elliptic problems with discontinuous coefficients*, SIAM J. Numer. Anal., 51 (2013), pp. 3106–3134.
- [12] T. BROWN, B. MANN, N. RYDER, M. SUBBIAH, J. D. KAPLAN, P. DHARIWAL, A. NEELAKANTAN, P. SHYAM, G. SASTRY, A. ASKELL, ET AL., *Language models are few-shot learners*, Adv. Neural Inf. Process. Syst., 33 (2020), pp. 1877–1901.
- [13] A. CHKIFA, A. COHEN, R. DEVORE, AND C. SCHWAB, *Sparse adaptive taylor approximation algorithms for parametric and stochastic elliptic pdes*, ESAIM Math. Model. Numer. Anal., 47 (2013), pp. 253–280.
- [14] D. CIORANESCU AND P. DONATO, *An Introduction to Homogenization*, Oxford Lecture Ser. Math. Appl. 17, Oxford University Press, Oxford, UK, 1999.
- [15] A. COHEN, R. DEVORE, AND C. SCHWAB, *Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDE's*, Anal. Appl., 9 (2011), pp. 11–47.
- [16] A. COHEN, R. DEVORE, AND C. SCHWAB, *Convergence rates of best n -term Galerkin approximations for a class of elliptic spdes*, Found. Comput. Math., 10 (2010), pp. 615–646.
- [17] G. CYBENKO, *Approximation by superpositions of a sigmoidal function*, Math. Control Signals Systems, 2 (1989), pp. 303–314.
- [18] M. V. DE HOOP, Z. HUANG, E. QIAN, AND A. M. STUART, *The cost-accuracy trade-off in operator learning with neural networks*, J. Mach. Learn., 1 (2022), pp. 299–341.
- [19] R. DEVORE AND G. LORENTZ, *Constructive Approximation*, Grundlehren Math. Wiss. 303, Springer, Berlin, 1993.
- [20] J. DUGUNDJI, *An extension of Tietze's theorem*, Pac. J. Math., 1 (1951), pp. 353–367.
- [21] J. N. FUHG AND N. BOUKLAS, *On physics-informed data-driven isotropic and anisotropic constitutive models through probabilistic machine learning and space-filling sampling*, Comput. Methods Appl. Mech. Engrg., 394 (2022).
- [22] J.-L. GUERMOND, *The LBB condition in fractional Sobolev spaces and applications*, IMA J. Numer. Anal., 29 (2009), pp. 790–805.
- [23] M. E. GURTIN, *An Introduction to Continuum Mechanics*, Academic Press, New York, 1982.
- [24] E. HAGHIGHAT, S. ABOUALI, AND R. VAZIRI, *Constitutive model characterization and discovery using physics-informed deep learning*, Eng. Appl. Artif. Intell., 120 (2023), 105828.
- [25] J. HAN AND Y. LEE, *A neural network approach for homogenization of multiscale problems*, Multiscale Model. Simul., 21 (2023), pp. 716–734.
- [26] K. HE, X. ZHANG, S. REN, AND J. SUN, *Deep residual learning for image recognition*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 770–778.
- [27] L. HERRMANN, C. SCHWAB, AND J. ZECH, *Neural and GPC Operator Surrogates: Construction and Expression Rate Bounds*, preprint, arXiv:2207.04950, 2022.
- [28] R. HILL, *The elastic behaviour of a crystalline aggregate*, Proc. Phys. Soc. Sect. A, 65 (1952), 349.
- [29] T. Y. HOU AND X.-H. WU, *A multiscale finite element method for elliptic problems in composite materials and porous media*, J. Comput. Phys., 134 (1997), pp. 169–189.
- [30] D. Z. HUANG, N. H. NELSEN, AND M. TRAUTNER, *An Operator Learning Perspective on Parameter-to-Observable Maps*, preprint, arXiv:2402.06031, 2024.
- [31] D. Z. HUANG, K. XU, C. FARHAT, AND E. DARVE, *Learning constitutive relations from indirect observations using deep neural networks*, J. Comput. Phys., 416 (2020), 109491.
- [32] N. KOVACHKI, S. LANTHALER, AND S. MISHRA, *On universal approximation and error bounds for Fourier neural operators*, J. Mach. Learn. Res., 22 (2021), pp. 1–76.
- [33] N. KOVACHKI, Z. LI, B. LIU, K. AZIZZADENESHELI, K. BHATTACHARYA, A. STUART, AND A. ANANDKUMAR, *Neural operator: Learning maps between function spaces with applications to PDEs*, J. Mach. Learn. Res., 24 (2023), pp. 1–97.
- [34] S. LANTHALER, Z. LI, AND A. M. STUART, *The Nonlocal Neural Operator: Universal Approximation*, preprint, arXiv:2304.13221, 2023.
- [35] S. LANTHALER, S. MISHRA, AND G. E. KARNIADAKIS, *Error estimates for DeepONets: A deep learning framework in infinite dimensions*, Trans. Math. Appl., 6 (2022).
- [36] G. LEONI, *A First Course in Sobolev Spaces*, AMS, Providence, RI, 2017.
- [37] Z. LI, N. KOVACHKI, K. AZIZZADENESHELI, B. LIU, K. BHATTACHARYA, A. STUART, AND A. ANANDKUMAR, *Fourier neural operator for parametric partial differential equations*, in Proceedings of the International Conference on Learning Representations, 2021.
- [38] K. LINKA, M. HILLGÄRTNER, K. P. ABDOLAZIZI, R. C. AYDIN, M. ITSKOV, AND C. J. CYRON, *Constitutive artificial neural networks: A fast and general approach to predictive data-driven constitutive modeling by deep learning*, J. Comput. Phys., 429 (2021), 110010.

- [39] B. LIU, N. KOVACHKI, Z. LI, K. AZIZZADENESHELI, A. ANANDKUMAR, A. M. STUART, AND K. BHATTACHARYA, *A learning-based multiscale method and its application to inelastic impact problems*, *J. Mech. Phys. Solids*, 158 (2022).
- [40] B. LIU, E. OCEGUEDA, M. TRAUTNER, A. M. STUART, AND K. BHATTACHARYA, *Learning macroscopic internal variables and history dependence from microscopic models*, *J. Mech. Phys. Solids*, 178 (2023), 105329.
- [41] X. LIU, F. TAO, AND W. YU, *A neural network enhanced system for learning nonlinear constitutive law and failure initiation criterion of composites using indirectly measurable data*, *Compos. Struct.*, 252 (2020), 112658.
- [42] X. LIU, S. TIAN, F. TAO, AND W. YU, *A review of artificial neural networks in the constitutive modeling of composite materials*, *Compos. B Eng.*, 224 (2021), 109152.
- [43] H. J. LOGARZO, G. CAPUANO, AND J. J. RIMOLI, *Smart constitutive laws: Inelastic homogenization through machine learning*, *Comput. Methods Appl. Mech. Engrg.*, 373 (2021), 113482.
- [44] L. LU, P. JIN, G. PANG, Z. ZHANG, AND G. E. KARNIADAKIS, *Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators*, *Nature Mach. Intell.*, 3 (2021), pp. 218–229.
- [45] P. MA, P. Y. CHEN, B. DENG, J. B. TENENBAUM, T. DU, C. GAN, AND W. MATUSIK, *Learning neural constitutive laws from motion observations for generalizable PDE dynamics*, in *International Conference on Machine Learning*, PMLR, 2023, pp. 23279–23300.
- [46] C. MARCATI, J. A. OPSCHOOR, P. C. PETERSEN, AND C. SCHWAB, *Exponential ReLU neural network approximation rates for point and edge singularities*, *Found. Comput. Math.*, 23 (2023), pp. 1043–1127.
- [47] C. MARCATI AND C. SCHWAB, *Exponential convergence of deep operator networks for elliptic partial differential equations*, *SIAM J. Numer. Anal.*, 61 (2023), pp. 1513–1545.
- [48] M. MOZAFFAR, R. BOSTANABAD, W. CHEN, K. EHMANN, J. CAO, AND M. BESSA, *Deep learning predicts path-dependent plasticity*, *Proc. Natl. Acad. Sci. USA*, 116 (2019), pp. 26414–26420.
- [49] K. P. MURPHY, *Machine Learning: A Probabilistic Perspective*, MIT Press, Cambridge, MA, 2012.
- [50] N. H. NELSEN AND A. M. STUART, *The random feature model for input-output maps between Banach spaces*, *SIAM J. Sci. Comput.*, 43 (2021), pp. A3212–A3243.
- [51] R. H. NOCHETTO, K. G. SIEBERT, AND A. VEESER, *Theory of adaptive finite element methods: An introduction*, in *Multiscale, Nonlinear and Adaptive Approximation: Dedicated to Wolfgang Dahmen on the Occasion of his 60th Birthday*, Springer, New York, 2009, pp. 409–542.
- [52] J. A. OPSCHOOR, C. SCHWAB, AND J. ZECH, *Exponential ReLU DNN expression of holomorphic maps in high dimension*, *Constr. Approx.*, 55 (2022), pp. 537–582.
- [53] H. OWHADI AND L. ZHANG, *Numerical homogenization of the acoustic wave equations with a continuum of scales*, *Comput. Methods Appl. Mech. Engrg.*, 198 (2008), pp. 397–406, <https://doi.org/10.1016/j.cma.2008.08.012>.
- [54] G. PAVLIOTIS AND A. STUART, *Multiscale Methods: Averaging and Homogenization*, Springer, New York, 2008.
- [55] R. PHILLIPS, *Crystals, Defects and Microstructures: Modeling Across Scales*, Cambridge University Press, Cambridge, UK, 2001.
- [56] H.-J. SCHMEISSER AND H. TRIEBEL, *Topics in Fourier Analysis and Function Spaces*, Vol. 1, Wiley, New York, 1987.
- [57] E. M. STEIN AND G. WEISS, *Introduction to Fourier Analysis on Euclidean Spaces*, Vol. 1, Princeton University Press, Princeton, NJ, 1971.
- [58] A. M. TARTAKOVSKY, C. O. MARRERO, P. PERDIKARIS, G. D. TARTAKOVSKY, AND D. BARAJAS-SOLANO, *Physics-informed deep neural networks for learning parameters and constitutive relationships in subsurface flow problems*, *Water Resources Res.*, 56 (2020), e2019WR026731.
- [59] K. XU, D. Z. HUANG, AND E. DARVE, *Learning constitutive relations using symmetric positive definite neural networks*, *J. Comput. Phys.*, 428 (2021), 110072.
- [60] K. YOSIDA, *Functional Analysis*, Springer, New York, 2012.
- [61] H. YOU, Y. YU, S. SILLING, AND M. D’ELIA, *Data-driven learning of nonlocal models: From high-fidelity simulations to constitutive laws*, *CEUR-WS*, 2964 (2021).