



Distributional inverse homogenization

Arnaud Vadeboncoeur ^{a,1,*}, Mark Girolami ^{a,b,2}, Kaushik Bhattacharya ^{c,3},
Andrew M. Stuart ^{d,4}

^a Department of Engineering, University of Cambridge, 7a JJ Thomson Ave, Cambridge, CB3 0FA, UK

^b The Alan Turing Institute, 96 Euston Rd., London, NW1 2DB, UK

^c Mechanical and Civil Engineering, California Institute of Technology, 1200 E California Blvd, Pasadena, CA 91125, US

^d Computing and Mathematical Sciences, California Institute of Technology, 1200 E California Blvd, Pasadena, CA 91125, US

ARTICLE INFO

Keywords:

Homogenization
Inverse problems
Distributional inference

ABSTRACT

For many materials, macroscopic mechanical behavior is determined by an intricate microstructure. Understanding the relation between these two scales helps scientists and engineers design better materials. The relation which maps microstructure to bulk material properties can be understood via the well-established theory of homogenization. However inverting the homogenization process, to recover microstructural information from measured macroscopic properties, is fraught with difficulties because of the averaging processes that underlie homogenization. Therefore, scientists and engineers usually need recourse to more invasive, often highly localized, investigations to estimate the microstructure. In this work, we develop a noninvasive methodology by which one can leverage large collections of measured bulk material properties to infer information about the *statistics* of microstructure at a global level. We call this, *distributional inverse homogenization*. We study this problem in one and two dimensions, considering both periodic and stochastic homogenization. We demonstrate the methodology in the context of 2D Voronoi constructions and underpin the observed empirical success with theory in 1D. We also show how the natural spatial variability of microstructure can be exploited to gather data that enables distributional inversion. And we concurrently learn a surrogate model, approximating the homogenization map, that accelerates the resulting computations in this setting. The work formulates a new class of inverse problems, bridging ideas from probability and homogenization to facilitate the inference of microstructural material variability from macroscopic measurements.

1. Introduction

The theory of homogenization can be used to define a map from microscopic material properties to the resulting macroscopic material properties. Inverting this map is difficult, because it involves a complex averaging procedure, meaning that multiple

* Corresponding author.

E-mail addresses: av537@cam.ac.uk (A. Vadeboncoeur), mag92@cam.ac.uk (M. Girolami), bhatta@caltech.edu (K. Bhattacharya), astuart@caltech.edu (A.M. Stuart).

¹ Supported through the EPSRC ROSEHIPS grant EP/W005816/1.

² Supported by a Royal Academy of Engineering Research Chair and UKRI grants EP/X037770/1, EP/Y028805/1, EP/W005816/1, EP/V056522/1, EP/V056441/1.

³ Supported by a US Office of Naval Research DDCR MURI (award N00014-23-1-2654.)

⁴ Supported by a Department of Defense (DoD) Vannevar Bush Faculty Fellowship (award N00014-22-1-2790) and US Office of Naval Research DDCR MURI (award N00014-23-1-2654.)

microstructures map onto the same macroscopic response. However, inferring information about the microstructure from macroscopic properties, obviating the need for intrusive methods such as etching and microscopy, is potentially of great interest. Whilst inverting the homogenization map is severely ill-posed, we show in this paper that inverting for the *statistics* of the microstructure is a tractable problem. We refer to this as *distributional inverse homogenization*.

Statistical characterization of microstructure has many applications. In alloy development, one may wish to understand the impact of certain processes on the distribution of grain size, shape, orientation (Rollett et al., 2007). In metamaterials, regularity of periodic microstructure is paramount for quality control (Catalanotti, 2016). Natural and engineered cellular materials such as wood and foam are highly heterogeneous; characterizing this inherent variability can help gauge suitability of production batches for industrial use (Gibson, 2003; Fratzl and Weinkamer, 2007). In plastics, the length of polymer chains and crystallinity determine many mechanical properties (Gent, 1970). In concrete, void sizes and geometric disposition affect both mechanical properties and electrolyte solution transport behavior, affecting in turn the strength and longevity of structures (Martín-Pérez et al., 2000; Martín-Pérez and Thomas, 2001; Martín-Pérez, 1999). In such applications statistical information is highly relevant to understanding how manufacturing controls material and how operational conditions may affect materials at a microstructural level.

In this work, we frame the task of inferring microstructure from a generative modelling perspective. That is, given expert knowledge, one poses a statistical generative model for a posited microstructural class. Such a model class can be hypothesized on the basis of a combination of experience, first principles and collection of small numbers of microscale images. It may help the reader to consider the following example: posit a model class for a material as that of a polycrystalline microstructure, of M crystal types, generated via a Voronoi diagram with a given number of uniformly distributed nucleation sites. The (random) map taking the properties of the M crystal types to the random Voronoi diagram of the microstructure is a physically interpretable generative model. The hypothesized microstructural generative model, for given sets of physically interpretable parameters, generates a distribution on microstructure. One then relates the sampled microstructures from the hypothesized model to macroscopic material properties using the theory of homogenization. We are then left with the task of calibrating the physically interpretable model parameters for our hypothesized model class to a dataset of macroscopic material properties from the material specimen(s) of interest. This model calibration/data-fitting stage is performed through matching the distribution of the homogenized coefficients sampled from the generative model with the distribution of observed macroscopic responses. This is the procedure that we refer to as *distributional inverse homogenization*.

We develop our methodology in the setting of both periodic and stochastic homogenization. In the former, one assumes a regular repeating microstructure extending over the entire domain to be homogenized; in the latter, one assumes instead that the distribution over the microstructure is given by a random process whose n -point statistics are invariant to translation and that any single random realization of the field is statistically representative of any other random realization. In both cases, homogenization boils down to mapping the material configuration in a representative computational *cell* to a single summary coefficient. This coefficient can be anisotropic, even if the microstructural material constituent properties are isotropic, when the underlying geometry of the material layout is anisotropic. Central to our investigation is the relation between microstructural constituent material properties, the distribution of material volume fractions, their microscale geometric organization, and the resulting distribution on bulk (homogenized) properties. As will be seen, the distribution on macroscopic material properties for randomization of the microstructure is extremely information rich. Distributional inverse homogenization exploits this information.

In this paper we focus on homogenizing operators for scalar solution fields in one and two spatial dimensions. These are representative of physical applications such as steady-state heat flow and electrostatics. We confine our attention to scalar problems in one and two spatial dimensions because of analytical tractability (one dimension) and straightforward but informative computations (Voronoi microstructures in two dimensions); this leads to a straightforward exposition of the ideas. However, we expect the work herein to generalize to distributional inverse homogenization for elasticity in two and three spatial dimensions, where a 4-tensor field characterizes the microstructure. Indeed investigation of this inverse problem in the context of elasticity constitutes an interesting avenue for further work.

1.1. Contributions and outline

We now summarize the main contributions of this work.

- (C1) Inferring a unique microstructure from measurement of a bulk material property is not well posed. We demonstrate however, both with theory and numerical experiment, that inferring the statistics of microstructure, from measurements of a bulk material properties, is well-posed.
- (C2) We demonstrate the effectiveness and practical relevance of inverting for the statistics of microstructure in a suite of numerical experiments focused on Voronoi microstructures in the periodic and stochastic homogenization setting in two dimensions. The periodic i.i.d. case is extended to the setting of elastic homogenization.
- (C3) We develop a novel surrogate learning strategy to overcome computational challenges arising in distributional inversion for stochastic homogenization, making the methodology practical.
- (C4) We develop statistical models for locally periodic and locally stationary-ergodic Voronoi fields to approximate microstructural variations in large material specimens. Using this construction, we show how distributional inversion can exploit the natural spatial variability of materials to statistically characterize microstructure.

Fig. 1 illustrates (C1). The basic approach that we adopt in this paper can be described at a high level as follows. Let datadistribution and gendistribution denote two probability distributions. The first is defined by the collection of measurements at the bulk (homoge-




Fig. 1. The left half of this figure illustrates why inverting the homogenization map for microstructure is hard. The right half of the figure shows why inverting the homogenization map for statistics of the microstructure is tractable. Both left and right show materials comprised of the same three constituent microstructures with $(a_1, a_2, a_3) = (1, 2, 3)$. On the left each constituent occupies the same volume fraction in each realization of the material, and hence all map to a single bulk property (displayed, lower left). On the right the volume fractions are chosen at random from a statistical distribution; this distribution may be estimated from the (displayed, lower right) distribution on bulk properties. The particular distribution of volume fractions in the right half of this figure is described and used in [Appendix A](#).

nized) scale. The second is defined by a generative model that generates microstructures randomly and maps them to bulk properties through forward homogenization (periodic or stochastic). The generative process is parameterized by θ , a parameter which controls the generation of microstructures. Let D denote some form of distance measure between probability distributions. Then the methodology underlying the computations in (C2)–(C4) is to solve an optimization problem for parameter θ^* that minimizes the dissimilarity between these distributions:

$$\theta^* = \arg \min_{\theta} D(\text{datadistribution}, \text{gendistribution}(\theta)). \quad (1)$$

Explaining the details of the two distributions `datadistribution` and `gendistribution` and the choice of D defines the basic methodology defined and deployed in this paper. The remainder of this section is organized as follows: [Section 1.2](#) reviews related work on the topics of generative modelling, distributional inference and homogenization; [Section 1.3](#) highlights important notation. In [Section 2](#) we provide background on homogenization and distributional inversion; [Section 3](#) then presents the proposed methodological contribution of this work, addressing (C1) and explaining (1) in detail; [Section 4](#) is devoted to the 1D component of (C1). [Section 5](#) extends numerical investigation to the 2D periodic case, [Section 6](#) in turn focuses on 2D stochastic homogenization, and [Section 7](#) extends some of results to the case of elasticity; together these sections address the 2D component of (C1), together with (C2) and (C3). Contribution (C3) amounts to learning a judicious approximation of the forward homogenization map that makes evaluation of `gendistribution`(θ) orders of magnitude faster.

1.2. Related works

In this section we review relevant literature. We first discuss recent developments in generative modelling which provides the tools necessary for our work. This is followed by discussion of distributional inversion. Finally we discuss select works in microstructure design, making links between material science and various objective functionals used for distributional inversion.

1.2.1. Generative modelling

In recent years, statistical inference and machine learning has focused considerable attention on tasks related to modelling datasets at the distributional level. Many works in this direction focus on generating new samples from a distribution implied by a dataset (Bond-Taylor et al., 2021; Song et al., 2020; Goodfellow et al., 2014). A subset of these methods use divergences and metrics on the space of probability distributions which are well behaved when comparing empirical distributions: mixtures of Dirac measures. Such metrics/divergences include Wasserstein distances (Panaretos and Zemel, 2019), Energy Distance (Sejdinovic et al., 2013; Székely and Rizzo, 2013), Maximum Mean Discrepancy (MMD) (Gretton et al., 2012) and Sliced-Wasserstein (Bonnel et al., 2015) (SW). The SW distance, which plays an important role in this paper, has many variants (Flamary et al., 2026, 2021; Kolouri et al., 2019; Chen et al., 2020; Deshpande et al., 2019; Nguyen and Ho, 2023; Nguyen et al., 2021). Some works make use of generative modelling in the context of physical simulation (Gao et al., 2024) or for solving inverse problems (Goh et al., 2019; Glyn-Davies et al., 2025). However, these works do not target distributional inversion.

1.2.2. Distributional inversion

The concept of distributional inversion as applied here was introduced and developed, primarily for applications in the physical systems in Akyildiz et al. (2025), Vadeboncoeur et al. (2025b). In these works, problems typically approached via hierarchical Bayes (Gelman et al., 1995) and empirical Bayes (Robbins, 1992) are re-interpreted as distributional inversion tasks. The resulting task requires a very large number of forward PDE solves over a relatively small set of parameter values. The methodology introduced in Akyildiz et al. (2025), Vadeboncoeur et al. (2025b) establishes an approach to concurrently learn surrogate models of the forward solve, needed at each of the optimization algorithm defining distributional inversion. Such concurrently learned surrogates are also explored in Zhang et al. (2026b). Related to the presented work are Li et al. (2025a, 2024, 2025b), Bernton et al. (2019). Relevant application of distributionally inferred priors can be found in Vadeboncoeur et al. (2025a).

1.2.3. Microstructure inverse design and indirect investigation

We now highlight works in the context of homogenization where objectives related to ours are explored. Of particular relevance to our work are Cherkaev (2001), Cherkaev and Golden (1998), in which the authors develop methodology for recovering spectral representations of structural information of microstructure for two-phase materials using electromagnetic measurement. In Quey and Renversade (2018), design and microstructural description questions are explored by optimizing polycrystalline structure to match desired distributional properties, such as for grain size. More broadly the tools in Quey and Kasemer (2022), Quey et al. (2011), Sundararaghavan and Zabaras (2006), Sternfels and Koutsourelakis (2011) may be of great use in answering such questions. We also highlight (Al-Ostaz et al., 2007), which develops related methodology to understand microstructure and the associated local stress fields. In Sigmund (1994) the author presents methodology for obtaining desired bulk material properties by designing microstructure. Questions related to forward UQ and random microstructure are explored in Yi and Bessa (2025). Other noteworthy works focus on microstructural investigation and design via machine learning (Rixner and Koutsourelakis, 2022; Senthilnathan et al., 2024; Liu et al., 2015; Chatzopoulos and Koutsourelakis, 2024; Sundararaghavan and Zabaras, 2005), some also use generative modelling methodology (Zang and Koutsourelakis, 2026).

1.3. Notation

This work is naturally expressed in terms of the notation of probability theory, nevertheless we prioritize wherever possible notation from mechanics. We now provide an introduction to relevant notation used throughout the paper.

We denote by v_i the i^{th} element, $i = 1, \dots, n$, of an ordered n -tuple v . We define functions either by specifying domain and co-domain as $f : X \rightarrow Z$, or, sometimes more conveniently, by specifying $f : x \mapsto z$ where it should be understood from context what is the implied domain and co-domain. On occasion we also overload notation for functions: we include the (co-)domain in the definition of a function so that a function denoted with the same symbol but with different (co-)domain should be understood as a different (likely very related) function.

We define the simplex

$$\Delta_{M-1} = \{\rho \in \mathbb{R}^M : \sum_{m=1}^M \rho_m = 1, \rho_m \geq 0 \text{ for } m = 1, \dots, M\}. \quad (2)$$

One important application of the simplex is that elements in it define a discrete probability measure on M discrete elements: each ρ_i is the probability of element i . The simplex is also useful for describing volume fractions which, like the probabilities, must add to 1 and be non-negative.

For exposition, let $(X, \mathcal{F})$ be a measurable space and $\mathbb{P}, \mathbb{P}'$ be probability measures. We denote superscripts in parenthesis when the values these objects take are random draws: $v^{(i)} \sim \mathbb{P}$ is the i^{th} random draw from $\mathbb{P}$. We use “ $\sim$ ” both to mean “drawn from”, and “distributed according to” depending on there being a superscript in parenthesis on the variable. Given function f and probability measure $\mathbb{P}$, the pushforward $f_{\#}\mathbb{P}$ is defined by the identity $f_{\#}\mathbb{P}(B) = \mathbb{P}(f^{-1}(B))$, holding for any $B \in \mathcal{F}$ and measurable f ; a random $z^{(\cdot)} \sim f_{\#}\mathbb{P}$ is drawn by first sampling $x^{(\cdot)} \sim \mathbb{P}$, then setting $z^{(\cdot)} = f(x^{(\cdot)})$.

To draw a sample from the convolution of probability measures, denoted $z \sim (\mathbb{P} * \mathbb{P}')$, one samples $x^{(\cdot)} \sim \mathbb{P}$ and $x'^{(\cdot)} \sim \mathbb{P}'$ and sets $z^{(\cdot)} = x^{(\cdot)} + x'^{(\cdot)}$. By $\mathbb{P}_n \rightarrow \mathbb{P}$ we mean converges in distribution. We denote empirical expectations with $\hat{\mathbb{E}}$. From context it should be clear if (sub/super)scripts denotes labels or other operations like exponentiation. The symbol e^m is a vector of all 0's with a 1 in position m . We use $\langle \cdot, \cdot \rangle$ to denote the Euclidean inner product.

For a map $f : X \rightarrow Z$ to be *injective* means that $f(x_1) = f(x_2)$ automatically implies that $x_1 = x_2$, i.e. the map is invertible for elements in its image. We say that identifying x from $f(x)$ is identifiable if f is injective.

In this work we make use of two statistical approaches to modelling random microstructure: (i) *independent identically distributed* (i.i.d.) models, whereby a random realization of a microstructure contains no information pertaining to any other realization from the model; and (ii) *statistical copulas*, for which it is possible specify the one- and two-point statistics of random spatial functions, and measurements of material properties at different locations are statistically identical and correlated, but decorrelate on lengthscales defined by the copula construction. In this work, the generative models used for the inference task will always be of type (i). However we apply our methodology to synthetic data of both types (i) and (ii). Data of type (i) enables us to conduct proof-of-concept studies; data of type (ii) enables us to mimic realistic data acquisition scenarios. The construction of these two statistical models are detailed in [Sections 5-6](#).

2. Background

In this section we give a brief overview of homogenization for scalar elliptic operators, both for periodic and stationary coefficient fields. We then introduce distributional inversion, and finalize this section with a note on surrogate-model-enabled acceleration to amortize the compute cost of homogenization in the context of distributional inversion.

2.1. Homogenization of scalar elliptic operators

The exposition herein is largely based on [Bensoussan et al. \(2011\)](#), [Blanc et al. \(2016\)](#), [Pavliotis and Stuart \(2008\)](#). We are interested in problems of the form

$$-\nabla \cdot (A^\varepsilon \nabla u^\varepsilon) = f, \quad \text{for } x \in D \subset \mathbb{R}^d, \quad (3a)$$

$$u^\varepsilon = 0, \quad \text{for } x \in \partial D, \quad (3b)$$

where $A^\varepsilon(x) = A(x/\varepsilon)$, for small ε , and $A : Q \rightarrow M_d$ where M_d is the space of real, symmetric, and uniformly elliptic $d \times d$ matrices; Q is the “material cell” whose definition depends on the type of homogenization procedure used to model the dependence between the material micro-structure and corresponding bulk properties. We assume the forcing f to be bounded. As A^ε varies on a much faster scale than does the solution u^ε , one can employ homogenization to identify a constant coefficient $\bar{A}$, on D , for which the solution to (3) is well approximated by solution u of the PDE

$$-\nabla \cdot (\bar{A} \nabla u) = f, \quad \text{for } x \in D \subset \mathbb{R}^d, \quad (4a)$$

$$u = 0, \quad \text{for } x \in \partial D. \quad (4b)$$

The type of homogenization used will depend what structure of A^ε can be exploited for averaging. In this work we study periodic and stochastic homogenization.

2.1.1. Periodic homogenization

In periodic homogenization, we make the identification $Q := \mathbb{T}_{(0,1)}^d$, so $A : Q \rightarrow M_d$ is 1-periodic. The homogenized coefficient $\bar{A}$ corresponding the cell micro-structure A is

$$\bar{A} = \int_Q (A(y) + A(y) \nabla \chi(y)^\top) dy, \quad (5)$$

where the corrector function $\chi : Q \rightarrow \mathbb{R}^d$ is in H_{per}^1 and weakly satisfies the cell problem

$$-\nabla_y \cdot (\nabla_y \chi A^\top) = \nabla_y \cdot A^\top, \quad 0 = \int_Q \chi(y) dy. \quad (6)$$

Note that the map $A \mapsto \chi$ may be viewed as a complex nonlocal averaging procedure; the homogenization map is then defined by using χ and A in a simple integration over the material cell. It is useful to define the homogenization map

$$F^\dagger : A \mapsto \chi \mapsto \bar{A}. \quad (7)$$

This map may be readily approximated, for example using finite elements to solve for χ , and quadrature to form $\bar{A}$. The latter map induces a complicated nonlinear averaging of A . If A is chosen at random from a probability distribution, parameterized by θ , then (approximate) evaluation of $F^\dagger$ defines $\text{gendistribution}(\theta)$, appearing in the optimization problem (1).

2.1.2. Stochastic homogenization

Let $(\Omega, \mathcal{F}, \mathbb{P}_\omega)$ be a probability triple. In stochastic homogenization the cell microstructure A is now defined on the whole of $\mathbb{R}^d$, and depends on a randomly chosen parameter in Ω . Specifically we take $A : \mathbb{R}^d \times \Omega \rightarrow M_d$. Furthermore we assume that the field is stationary-ergodic: $\forall k \in \mathbb{Z}^d, A(x+k, \omega) = A(x, h_k \omega)$, where h_k is an ergodic group action ([Blanc et al., 2016](#)). Then the homogenized coefficient $\bar{A}$ is well-defined through an almost sure limiting procedure. To define $\bar{A}$, we approximate the unbounded domain $\mathbb{R}^d$ by a family of finite domains, and limits taken in the size of this domain. The resulting methodology is also the basis of computational methods for stochastic homogenization.

We introduce the truncated domain $Q_n = (0, n)^d$, with n large and $A_n = A|_{Q_n}$. For given $\omega \in \Omega$ we may then define

$$\bar{A}_n(\omega) = \int_{Q_n} (A_n(y, \omega) + A_n(y, \omega) \nabla \chi_n(y, \omega)^\top) dy, \quad (8)$$

where $\chi_n : \mathbb{T}_{[0,n]}^d \rightarrow \mathbb{R}^d$ is periodic and viewed as an element in $H_{n\text{-per}}^1$ which weakly satisfies the n -cell problem

$$-\nabla_y \cdot (\nabla_y \chi_n A_n^\top) = \nabla_y \cdot A_n^\top, \quad 0 = \int_{Q_n} \nabla \chi_n(y, \omega) dy. \quad (9)$$

Here, as in the periodic case, one can use, for example, finite elements to obtain the solution of the corrector problem and then quadrature to define $\bar{A}_n(\omega)$. We denote by $F_n^\dagger$ the map from A_n to $\bar{A}_n$, analogously to (7).

The homogenized coefficient $\bar{A}$ is obtained as the almost sure limit $n \rightarrow \infty$ of $\bar{A}_n(\omega)$. We can define the stochastic homogenization map $G^\dagger : \mathbb{P}_\omega \mapsto \bar{A}$ because $\mathbb{P}_\omega$ contains all information of A which is not lost in the almost sure limiting process. In practice, however, we work with finite n , assuming that n is large enough so that Q_n itself is large with respect to typical lengthscales of A_n . Using the fact that A_n is approximately stationary-ergodic in Q_n , we see that $F_n^\dagger$ thus implements an effective averaging procedure over both variation in $\mathbb{R}^d$ and in Ω . To reduce the variance in estimating $\bar{A}$ from finite n it is common to average over many random realizations of $A_n(\omega)$ via

$$\bar{A}_{\text{finite}} = \frac{1}{J_{\text{avg}}} \sum_{j=1}^{J_{\text{avg}}} \bar{A}_n^{(j)}, \quad \bar{A}_n^{(j)} \sim F_n^\dagger(A_n(Q_n, \cdot))_{\#} \mathbb{P}_\omega. \quad (10)$$

The notation $F_n^\dagger(A_n(Q_n, \cdot))$ expresses the fact that $F_n^\dagger$ acts on $A_n(y, \omega)$ as a function of $y \in Q_n$, for fixed ω , to produce a constant (with respect to y) object; and that the resulting function of ω alone can be used to pushforward the underlying distribution on ω . The precise definition of $\bar{A}$ is recovered by taking the large (n, J_{avg}) limit of $\bar{A}_{\text{finite}}$ as defined in the preceding identity. In the remainder of this paper, whenever $G^\dagger$ is invoked, one should assume the computation is performed via (10) to obtain $\bar{A}_{\text{finite}} \approx \bar{A}$. There is a tradeoff in approximation efficiency between the size of n and J_{avg} (Gloria, 2012; Gloria and Otto, 2012, 2011). We have worked with choices of the pair which work well empirically.

2.2. Distributional inversion

In distributional inversion – equipped with some forward model $G^\dagger$ – we assume to have access to data of the form

$$y^{(n)} = G^\dagger(z^{(n)}) + \xi^{(n)}, \quad (11a)$$

$$\xi^{(n)} \sim \mathbb{P}_\xi := \mathcal{N}(0, \sigma^2 \mathbf{I}), \quad (11b)$$

$$z^{(n)} \sim \mathbb{P}_z^\dagger, \quad (11c)$$

with $y^{(n)} \in \mathbb{R}^d$. (Other noise models for $\xi^{(n)}$ could be used, but we choose this specific one for expository simplicity.) Probability measure $\mathbb{P}_z^\dagger$ is the object of interest, and we would like to infer it from the data $y^{(n)}$. We collect data $\{y^{(n)}\}_{n=1}^N$ into an empirical distribution

$$\mathbb{P}_y^N = \frac{1}{N} \sum_{n=1}^N \delta_{y^{(n)}}; \quad \delta_{y^{(n)}}(B) = \begin{cases} 1, & y^{(n)} \in B, \\ 0, & y^{(n)} \notin B. \end{cases} \quad (12)$$

This distribution is simply a collection of points; to sample from $\mathbb{P}_y^N$ one simply selects points at random.

2.2.1. Optimization problem

We assume to have a class of parametrized probability measures $\mathcal{P}_\Theta$ and a map $\mathbb{P}_z : \Theta \rightarrow \mathcal{P}_\Theta$. The inference task is defined by the optimization problem

$$\theta^* \in \arg \min_{\theta \in \Theta} D(\mathbb{P}_y^N, \mathbb{P}_\xi * (G^\dagger)_{\#} \mathbb{P}_z(\theta)). \quad (13)$$

Here $D : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}_{\geq 0}$ is a divergence between probability measures which is well defined under empiricalization. Recall from Section 1.3 that “ $*$ ” denotes the convolution operator between distributions (we add noise to the output of the generative model according to the observational noise distribution, $\mathbb{P}_\xi$). Also recall from Section 1.3 that “ $\#$ ” denotes the pushforward of the distribution on the right, under the map on the left. Computing $D(\mathbb{P}_y^N, \mathbb{P}_\xi * (G^\dagger)_{\#} \mathbb{P}_z(\theta))$ involves three steps: (i) sampling a collection of i.i.d. realization of the form $z^{(\cdot)} \sim \mathbb{P}_z(\theta)$, applying $G^\dagger(z^{(\cdot)})$ to them (the pushforward), and adding noise $\xi^{(\cdot)} \sim \mathbb{P}_\xi$ (the convolution); (ii) selecting a collection of data from $\mathbb{P}_y^N$; (iii) evaluating D for the two collections of particles.

If the true $\mathbb{P}_z^\dagger$ is obtained from the map $\mathbb{P}_z$ by choosing $\theta^\dagger \in \Theta$ and if N is large then it is natural to expect $\theta^* \approx \theta^\dagger$. This is because, if $G^\dagger$ does not compress variations in the input, then $\mathbb{P}_y^N$ will contains information about the choice of θ .

In this work we employ, for D , the Sliced-Wasserstein metric, defined by

$$SW_2^2(\mathbb{P}, \mathbb{P}') = \int_{\mathbb{S}^{d-1}} W_2^2(P_{\rho \#} \mathbb{P}, P_{\rho \#} \mathbb{P}') d\rho, \quad (14)$$




Fig. 2. Distributional inverse homogenization schematic workflow in the locally stationary-ergodic microstructure setting described in Section 6.2. Here (α, a) are the sought parameters, where (α, ℓ, ν) parametrize the copula description of volume fractions, ρ is a spatially varying volume fraction, and a are the material properties.

where $\mathbb{S}^{d-1}$ is the sphere embedded in $\mathbb{R}^d$, and $P_\rho(\cdot) = \langle \rho, \cdot \rangle$. Numerically, the integral is approximated with N_ρ Monte-Carlo samples

$$\widehat{\text{SW}}_2^2(\mathbb{P}, \mathbb{P}') = \frac{1}{N_\rho} \sum_{i=1}^{N_\rho} W_2^2(P_{\rho^{(i)}} \# \mathbb{P}, P_{\rho^{(i)}} \# \mathbb{P}'); \quad (15a)$$

$$\rho^{(i)} = \tilde{\rho}^{(i)} / \|\tilde{\rho}^{(i)}\|, \quad \tilde{\rho}^{(i)} \sim \mathcal{N}(0, \mathbb{I}_{d_y}). \quad (15b)$$

This metric is widely used because the Monte Carlo approximation requires evaluation only of one dimensional Wasserstein distances, which are analytically tractable. In contrast the Wasserstein distance is hard to compute in higher dimensional spaces or when there is large amounts of data. Thus, assuming $G^\dagger$ and $\mathbb{P}_z$ are differentiable, we can approximate (13) via standard gradient descent through stochastic optimization methods such as Adam (Kingma and Ba, 2014).

2.2.2. Surrogate acceleration

One might be interested in replacing the computation of $G^\dagger$ using a surrogate model. This serve two purposes: i) it amortizes the cost of the forward map; and ii) it enables the differentiation of possibly “black-box”/discontinuous forward maps $G^\dagger$. Ideally, we would like our surrogate model, G^ϕ with parameters ϕ , to be accurate for $z \sim \mathbb{P}_z^\dagger$. As $\mathbb{P}_z^\dagger$ is not known, we instead choose ϕ so that $G^\phi(z) \approx G^\dagger(z)$ for $z \sim \mathbb{P}_z(\theta')$ with θ' an estimate of θ^* . We formalize this as a bilevel optimization problem seeking

$$\theta^* \in \arg \min_{\theta \in \Theta} D(\mathbb{P}_y^N, (G^{\phi^*(\theta)}) \# \mathbb{P}_z(\theta)), \quad (16a)$$

$$\text{such that } \phi^*(\theta) \in \arg \min_{\phi} \mathbb{E}_{z \sim \mathbb{P}_z(\theta)} [\|G^\phi(z) - G^\dagger(z)\|_2^2]. \quad (16b)$$

This methodology is developed in depth in Akyildiz et al. (2025), Vadeboncoeur et al. (2025b), Zhang et al. (2026a,b) with various implementation strategies and approximations involving physics residuals or non-differentiable “black-box” models. If one were interested in using surrogate acceleration in the periodic case with neural operators, it is straightforward to combine (Bhattacharya et al., 2024) and Vadeboncoeur et al. (2025b), Akyildiz et al. (2025). This work will explore a different surrogate modelling strategy in the stochastic homogenization case, using simple feedforward multilayer perceptron neural networks.

3. Methodology

This section brings together the ideas from the preceding section to define the proposed methodology for distributional inverse homogenization. We make the idea encapsulated in (1) concrete. The implications of the proposed approach, these being mathematical, numerical, and practical, will be explored in later sections.

We start by assuming that there exists a data-generating distribution on material microstructure $\mathbb{P}_A^\dagger$. We observe, by way of macroscopic measurement, a collection of bulk material properties from a generating model

$$\bar{A}^{(n)} = H^\dagger(A^{(n)}) + \xi^{(n)}, \quad (17a)$$

$$A^{(n)} \sim \mathbb{P}_A^\dagger, \quad (17b)$$

$$\xi^{(n)} \sim \mathbb{P}_\xi. \quad (17c)$$

Here $H^\dagger$ is the homogenization map taking a physical microstructure of a specimen to a macroscopically measurable bulk material property. The probability measure $\mathbb{P}_\xi$ describes noise entering the measurements. As described in abstract in (12), we collect data $n = 1, \dots, N$ into an empirical distribution

$$\mathbb{P}_A^N = \frac{1}{N} \sum_{n=1}^N \delta_{\bar{A}^{(n)}}. \quad (18)$$

We are interested in inferring aspects of the distribution on microstructure $\mathbb{P}_A^\dagger$ by way of parametric approximation.

Distributional Inverse Periodic Homogenization Recall map $F^\dagger$ given by (7). As explained in the general setting in Section 2.2, we directly posit a map $\theta \mapsto \mathbb{P}_A(\theta)$ such that $A \sim \mathbb{P}_A(\theta)$, and minimize

$$J_p(\theta) = D\left(\mathbb{P}_A^N, \mathbb{P}_\xi * (F^\dagger)_\# \mathbb{P}_A(\theta)\right). \quad (\text{Per. Inv. Hom. Objective})$$

For example in our 2D periodic homogenization examples we take $\mathbb{P}_A(\theta)$ to be a generative model for Voronoi diagrams with random seed locations where each Voronoi seed is deterministically assigned a material property. Furthermore, the Voronoi diagram may have a material dependent additive weighting. The goal of distributional inversion is to determine the statistics of the Voronoi diagrams. This idea is explored in Section 5. If we now choose D to be the sliced-Wasserstein metric (14), then we have defined a specific instance of (1).

Distributional Inverse Stochastic Homogenization Similarly to the periodic case, we posit a parametric statistical model for $A_n^{(n)}$ which is a map $\varpi^{(n)} \mapsto \mathbb{P}_\omega(\varpi^{(n)})$; the distribution $\mathbb{P}_\omega(\varpi^{(n)})$ subsumes all information of $A^{(n)}$, which according to exact stochastic homogenization is defined on an unbounded domain. We then define a map $\theta \mapsto \mathbb{P}_\omega(\theta)$ such that $\varpi^{(n)} \sim \mathbb{P}_\omega(\theta)$. This way, $\{A_n^{(n,i)}\}_{i=1}^\infty$ for $A_n^{(n,i)} \sim \mathbb{P}_\omega(\varpi^{(n)})$ summarizes all information in $A^{(n)}$, yet $A_n^{(n,i)}$ is defined on a bounded domain. We are now in a position to pose our objective: minimize

$$J_s(\theta) = D\left(\mathbb{P}_A^N, \mathbb{P}_\xi * (G^\dagger)_\# \mathbb{P}_\omega(\theta)\right). \quad (\text{Stoch. Inv. Hom. Objective})$$

In this paper, when studying 2D stochastic homogenization, $\mathbb{P}_\omega(\varpi^{(n)})$ is the generative model from which we can sample random Voronoi diagrams, with random seeding locations; each Voronoi cell is assigned a material property at random. The variables $\varpi^{(n)}$ are themselves drawn from a distribution with the parameters θ we seek to infer; ϖ will represent a description of volume fraction per material type as well as the associated constituent material property. If we again choose D to be the sliced-Wasserstein metric (14), then we have defined a specific instance of (1).

To implement (Per. Inv. Hom. Objective) and (Stoch. Inv. Hom. Objective), we use the approximate Sliced-Wasserstein distance (15) as our divergence D , rather than the exact Sliced-Wasserstein distance (14). The methodology presented in this current form is defined for observational data $\{A^{(n)}\}_{n=1}^N$ which are i.i.d. Of course the homogenized coefficients are symmetric, so we only compute the sliced-Wasserstein distance with the three unique coefficient elements. In this work we also explore a novel surrogate modelling direction specifically for stochastic homogenization, by concurrently learning an approximation to $G^\dagger$ given by (10), which we will call G^ϕ . More details on this may be found in Section 6. In later sections, we further extend our stochastic homogenization methodology to consider physically realistic scenarios where the data originates from a single large specimen queried at various locations, and is only approximately i.i.d. Here, homogenization approximately holds due to local conditions. Fig. 2 is graphical summary of that particular workflow in the stochastic setting. More on this in later sections.

4. 1D homogenization

We now explore certain mathematical properties, in one spatial dimension, of the problem posed in Section 3. We take a perspective that subsumes both periodic and stochastic homogenization into a single framework, in this 1D setting. Section 4.1 constructs an explicit example and a nontrivial inverse homogenization problem concerning distributions on volume fractions of M materials, which we prove to be identifiable. Section 4.2 explores a more practically relevant class of distributional inverse homogenization problems using Dirichlet distributed volume fractions. We prove identifiability in a specific limiting regime, of both material properties and parameters of the distribution and, in Section 4.3, provide numerical evidence demonstrating that this theorem extends beyond the limiting case that our theory covers. This addresses Contribution (C1).

In one spatial dimension, periodic homogenization of a periodic field, $a_{\text{per}}(\cdot)$, has closed form (Pavliotis and Stuart, 2008)

$$\bar{a}_{\text{per}} = \left(\int_{\mathbb{T}} a_{\text{per}}(y)^{-1} dy \right)^{-1}. \quad (19)$$

Stochastic homogenization of a stationary-ergodic field, $a_{\text{sc}}(\cdot)$, also has closed form (Kozlov, 1980) given by

$$\bar{a}_{\text{sc}} = \lim_{L \rightarrow \infty} \left(\frac{1}{2L} \int_{-L}^L a_{\text{sc}}(y, \omega)^{-1} dy \right)^{-1} \text{ for } \mathbb{P}_{\omega} - a.e. \omega. \quad (20)$$

This limiting procedure in L is what we approximate with the large (n, J_{avg}) limit in (10).

In the periodic homogenization case, assume the piecewise constant structure

$$a_{\text{per}}(y) = \sum_{m=1}^M a_j \mathbb{1}_{Q_m}(y), \quad (21)$$

where $\{Q_m\}_{m=1}^M$ is a partition of $\mathbb{T}$ with $a_1 > \dots > a_M > 0$. In the stochastic homogenization case, assume

$$a_{\text{sc}}(y, \omega) = \sum_{m=1}^M a_j \mathbb{1}_{U_m(\omega)}(y), \quad (22)$$

where $\{U_m(\omega)\}_{m=1}^M$ is a partition of $\mathbb{R}$ with $a_1 > \dots > a_M > 0$. Now, in both cases assume $\rho = (\rho_1, \dots, \rho_M) \in \Delta_{M-1}$ is the volume fraction (of Q_m or U_m) associated to the material coefficient a_m . We define $a = (a_1, \dots, a_M)$ and note that we may then write the following simple formula for the homogenized coefficient $\bar{a}$, valid in both periodic and stochastic homogenization:

$$\bar{a} = \langle a^{-1}, \rho \rangle^{-1}. \quad (23)$$

That is to say, in both periodic and stochastic 1D homogenization the homogenized coefficient is the volume fraction weighted harmonic mean. In the rest of this section we use the 1D simple-coefficient homogenization map

$$F(\rho; a) = \left(\sum_{m=1}^M a_m^{-1} \rho_m \right)^{-1}. \quad (24)$$

4.1. Volume fractions and a first distribution on the simplex

Recall Fig. 1 and in particular recall that the left-hand side illustrates the fact, apparent from (23), that in the 1D setting multiple materials have the same homogenized response – if the volume fractions are fixed then the harmonic average is the same regardless of how the materials are arranged, as formula (23) shows. We develop here our first construction that illustrates the right-hand side of Fig. 1. We employ a distribution on the simplex of volume fractions and show that it is possible to recover parameters of this distribution.

Theorem 4.1. *Let $M = 2$, fix $a \in (0, \infty) \times (0, \infty)$ and assume that $\rho_1 \sim \text{Unif}(b-l, b+l)$ with $b \geq l > 0$ and $b+l \leq 1$; necessarily $\rho_2 = 1 - \rho_1$ to ensure the volume fractions add to unity. Then b, l are determined uniquely by the distribution of homogenized material property $F(\rho; a)$.*

Thus we can recover the distribution of the volume fractions of the microstructure from the distribution of the bulk properties. Theorem 4.1 can be extended to $l = 0$ where it is simply the statement that in one dimension the volume fractions are identifiable from a single measurement of bulk properties if the constituent materials are known. Theorem 4.1 is a special case of a more general result, applying for any integer $M \geq 2$, and given in Appendix A.

4.2. The dirichlet distribution

The construction of a probability measure on Δ_{M-1} used in the preceding subsection, and in Appendix A, is straightforward as it uses only uniform distributions; but it is somewhat unnatural. The canonical distribution on the simplex Δ_{M-1} is the Dirichlet distribution, which we study in this subsection. We describe a setting in which the parameters of this distribution and the microstructural material properties are jointly identifiable. The following subsection provides numerical evidence beyond the particular parameter limit in which we prove the result.

The Dirichlet distribution is a multivariate generalization of the beta distribution, and may also be derived as a normalized multivariate version of the Gamma distribution. The Gamma distribution ensures positivity and the normalization enforces that draws are in the simplex. The density of the Dirichlet distribution, over variable $\rho \in \Delta_{M-1}$ and parameterized by Euclidean vector $\alpha = (\alpha_1, \dots, \alpha_M) \in (0, \infty)^M$, is given by

$$\text{Dir}(\rho; \alpha) = \frac{1}{B(\alpha)} \prod_{m=1}^M \rho_m^{\alpha_m-1}, \quad B(\alpha) = \frac{\prod_{m=1}^M \Gamma(\alpha_m)}{\Gamma\left(\sum_{m=1}^M \alpha_m\right)}. \quad (25)$$

Here Γ is the Gamma function. We use $\text{Dir}(\rho; \alpha)$ as a distribution on material volume fractions; combined with material constants a , and because 1D homogenization is invariant to material layout, (23), it is a generative model for 1D material microstructure with parameters (α, a) . Note that since the α_m are positive and finite in number we may write $\alpha := \beta\gamma$ for a scalar parameter β and vector $\gamma \in \Delta_{M-1}$. We introduce the notation $\text{Dir}(\rho; \gamma, \beta) := \text{Dir}(\rho; \gamma\beta)$ (and often drop argument ρ). Our second claim about identifiability of γ concerns the limit $\beta \rightarrow 0$. It is proved in Appendix B. Appendix B also contains Lemma B.1 which characterizes the behaviour of $\text{Dir}(\gamma, \beta)$ in the limits $\beta \rightarrow 0$ and $\beta \rightarrow \infty$.

Theorem 4.2. *Consider the distributional inverse homogenization problem in the restricted case where $\beta \rightarrow 0$. Then this problem is jointly identifiable for the parameter of the Dirichlet distribution, $\gamma \in \Delta_{M-1}$ and the material properties $a \in (0, \infty)^M$ for $a_j \neq a_i$, $i \neq j$, and $1 \leq i, j \leq M$.*




Fig. 3. (a) shows the histogram of the observed homogenized coefficient dataset, (b) the decreasing Objective, (c) the estimated volume fraction parameters of the microstructural generative model during minimization Objective, (d) the estimated constituent material parameters of the microstructural generative model.

4.3. Numerical experiments

We now present 1D numerical experiments demonstrating identifiability in the Dirichlet setting. These suggest identifiability beyond the narrow setting of [Theorem 4.2](#). We first demonstrate convergence plots for the distributional inversion problem, then we explore the parameter recovery as a function of the magnitude of $\beta > 0$.

Data and Inference: [Fig. 3a](#) shows the dataset histogram of the of homogenized coefficients for Dirichlet distributed volume fractions. Here, the number of materials considered, $M = 3$, is fixed, and the number of data points in the empirical data distribution is $N = 10^4$. The values for $a^\dagger = (0.5, 1.75, 3.)$ and $\alpha^\dagger = (0.46, 0.30, 0.24)$.

Results: [Fig. 3\(b-d\)](#) show the decreasing ([Per. Inv. Hom. Objective](#)), and the convergence of Dirichlet parameters and material properties.

These results provide further evidence for the well-posedness of distributional inversion problems of this form, beyond the setting of the somewhat restrictive claims. In the next two sections we develop our numerical studies further to explore 2D material fields described by Voronoi diagrams.

Remark 4.3. We briefly discuss the role of optimization in parameter identification as the problems considered are typically non-convex. In [Fig. 4](#) we test the accuracy of recovery of the Dirichlet parameters, parametrized with $\beta\gamma = \alpha$, and the material coefficients. In test where the ground truth value of β is small we observe occasional convergence to local minima; to mitigate this effect we use ten random restarts and select the run with least final value of ([Per. Inv. Hom. Objective](#)) for reporting. We observe the relative error in recovering β increases when its magnitude decreases. The choice of metric or divergence explored to define our objection will also have bearing on the optimization; however we do not explore alternatives to sliced-Wasserstein in this paper.

5. 2D periodic homogenization

In this section, we apply distributional inverse homogenization to problems in periodic homogenization. We first look at a class of i.i.d. random microstructures generated via Voronoi tessellation in periodic homogenization. We then explore a practically relevant




Fig. 4. Errors in estimating β, γ, a , as compared to ground truth, from distributional inversion, 1D periodic.

locally periodic Voronoi construction based on statistical copulas; we infer distributional information from spatially decorrelated measurements of bulk material properties. This section, together with Section 6, addresses Contributions (C2) and (C4). In what follows we use the terms “Voronoi seed” and “nucleation sites” interchangeably.

Remark 5.1. In conducting the experiments that follow, in this Section 5 and in Section 6, nonconvexity is evidenced by the presence of occasional local minima. Studying the properties and reasons for the existence of such local minima is a problem of interest in its own right, but not one we focus on in this paper. To address nonconvexity, we perform five random restarts and return the outcome giving the lowest objective value (best data-fit) for the last 100 iterations.

Remark 5.2. The off-diagonal element of the constant homogenized material field is typically on a different scale to the diagonal elements. Hence, when computing the data-fit objective we scale each component by the standard deviation of that component in the data set to equally weigh contributions from the three elements. This is done both for (Per. Inv. Hom. Objective) and (Stoch. Inv. Hom. Objective).

5.1. Periodic i.i.d. microstructures

We approach this section by first discussing the model setup, followed by the synthetic data, and finally the results.

Setup: In this section we consider i.i.d. data and an i.i.d. data generating model. Building on the previous section where we studied 1D materials, we seek to jointly infer information pertaining to volume fractions and material properties. As we are in 2D, we apply this methodology to material property fields described by Voronoi diagrams, as these are widely used in generative modelling for materials.

For the Voronoi description of a material field, it is not obvious how to directly specify the volume fractions in a cell. One could resort to optimizing Voronoi seed centers to tune the volume fractions shared by the M material types, however we choose a different approach. We cast the problem as the task of inferring additive Voronoi seed weights, w , which can be used to model shifts in crystal nucleation start times. This construction is also called a Laguerre tessellation (Quey and Renversade, 2018). We define a partition $\bigcup_{m=1}^M Q_m = Q$, where

$$Q_m(v; w) = \{y \in Q : d_{\text{per}}(y, v_m)^2 - w_m \leq d_{\text{per}}(y, v_j)^2 - w_j, \forall j \neq m\}, \quad (26a)$$




Fig. 5. Randomly generated periodic microstructure for (w, a) randomizing over Voronoi seed locations in the cell, Q .

$$d_{\text{per}}(y, y') = \sqrt{\sum_{i=1}^d (y_i - y'_i)^2 \wedge (1 - |y_i - y'_i|)^2}, \quad (26b)$$

where $\wedge$ is a binary function returning the least of the two elements being compared. Thus d_{per} is a metric on the torus – it measures distance accounting for periodicity. We sample a candidate material field from our microstructure generative model by assigning scalar material constant a_m to the m^{th} region of a random Voronoi tessellation, Q_m . Specifically,

$$v^{(m)} \sim \text{Unif}(Q), \quad (27a)$$

$$A(y; a, w) = \mathbb{I}_2 \sum_{m=1}^M a_m \mathbb{1}_{Q_m(v^{(1:M)}; w)}(y), \quad (27b)$$

where $v^{(m)}$ are uniformly distributed random nucleation sites. We refer to generative model (27) as the periodic i.i.d. model. Note that the size of Q_m , and so the volume fraction associated to material m , will be affected by the weight w_m , the relation to the weights of associated with cells defined by other nearby nucleation sites, and the location of randomly generated nucleation sites. Here, parameter w plays the role of ρ in Section 4. We seek to infer w and the material properties a for fixed distribution on the location of the Voronoi nucleation sites, (27a). Thus we define $\theta = (w, a)$ and note that we have a map $\theta \mapsto \mathbb{P}_{\bar{A}}$. The source of stochasticity yielding an information rich distribution on $\bar{A}$ comes from the randomization over the Voronoi seeds $v^{(1:M)}$ for each $\bar{A}$. In this setup, there is an additive lack of identifiability with regards to the start times; the Voronoi diagram is invariant to constant shifts in the collection of lag times (physically, it does not matter if all nucleations are delayed by the same amount). And so, for inference we fix $w_1 = 0$ and infer the other $\{w_m\}_{m=2}^M$.

Remark 5.3. In the Voronoi construction (27), for any material nucleation time shifts $0 \leq w_1 < \dots < w_m < \dots < w_M$, there is positive probability that a randomly generated diagram yields a volume fraction, for any material $m < M$, of zero. In practice we notice that if the difference in shift times is small enough with respect to the cell size, so that zero volume fractions are rare occurrence, we maintain identifiability of (w, a) .

Remark 5.4. In homogenization, Hashin-Shtrikman bounds are of general interest. These specify the upper and lower extremal values of isotropic homogenized coefficients that can be obtained with any geometric configuration of material phases for given material volume fractions and constituent material property. Under the weighted Voronoi periodic construction outlined in this section, and if $a_M = \max(a^\dagger)$ with w_M the largest Voronoi weight, an upper bound on the distribution of homogenized coefficients can be obtained (it is achieved when all nucleation sites are sufficiently close to $v^{(M)}$). Then, the periodic cell will contain only material a_M and hence the homogenized coefficient with be $\bar{A} = \mathbb{I}_2 a_M$. Fig. 6 (d)-(f) show this upper bound is approximately achieved by a random sample, noting the data generating $a^\dagger = (1, 2, 4)$ which implies a homogenized coefficient upper bound $a^+ = 4$. Obtaining general upper and lower bounds on the distribution of homogenized coefficients is less straightforward and we leave this line of investigation to future work.

We will generate from the i.i.d. model and we will attempt to estimate the crystal growth rates and material properties (w, a) by minimizing (Per. Inv. Hom. Objective). To deploy gradient based minimization on this objective function requires computation of gradients of J_p , and hence of $\bar{A}$, with respect to $w \in \theta$. As it is, the generative model (27) is not differentiable with respect to w . Thus we consider a continuous relaxation of (27b). To do this, specify τ , a relaxation parameter and define the softmax function acting on an M -dimensional vector $sm : \mathbb{R}^M \times \mathbb{R}_+ \rightarrow \mathbb{R}_+^M$ is

$$sm(z; \tau)_i = \frac{\exp(z_i/\tau)}{\sum_{j=1}^M \exp(z_j/\tau)}. \quad (28)$$

We call the relaxed random field $\tilde{A}$, it is sampled via

$$v^{(m)} \sim \text{Unif}(Q), \quad (29a)$$




Fig. 6. Histograms and scatter plots of the observed data-distribution. In (d)-(f) the color is representative the local density of particles computed via the mean 10-neighbor distance to nearby particles.




Fig. 7. Convergence of (Per. Inv. Hom. Objective) and (w, a) for the i.i.d. periodic microstructure model.

$$\tilde{A}(y; w, a, \tau) = \mathbb{I}_2 \left\langle a, \text{sm}(\text{d}_{\text{per}}(y, v^{(1:M)})^2 - w; \tau) \right\rangle. \quad (29b)$$

The softmax function ensures that for every sub-Voronoi-cell some amount of information from each material parameter, (w, a) , is present; this allows automatic differentiation through this model to estimate (w, a) . As $\tau \rightarrow 0$, $\tilde{A}(y; w, a, \tau) \rightarrow A(y; w, a)$ and $\nabla_{(w,a)} A(y; w, a, \tau)$ is well behaved for $\tau > 0$.

Data and Inference: We use $N = 5 \cdot 10^3$ data samples. We set $a^\dagger = (1, 2, 4)$, $w^\dagger = (0, 0.05, 0.1)$. To generate data we set τ to 10^{-20} to have sharp numerically discontinuous Voronoi cells. We do not include noise in this experiment. Fig. 5 shows three randomly generated periodic Voronoi diagrams for the parameter settings $(w^\dagger, a^\dagger)$. Fig. 6 shows the histograms of the data distribution $\mathbb{P}_{\bar{A}}^N$ as scatter plots of combinations of the three data components. We fix the computational mesh to be 100×100 when representing $\bar{A}(y)$ and in solving for the corrector function χ with finite elements in $F^\dagger$.

Results: For inference we set τ to 10^{-3} to allow gradients to flow through the Voronoi construction and estimate w . Fig. 7 shows the convergence of the objective, and (w, a) against $(w^\dagger, a^\dagger)$. We run 5 random restarts with (w, a) independently drawn from $\text{Unif}(0.01, 0.5)$ and $\text{Unif}(0.5, 3)$ respectively, and independently componentwise. The relative error on the recovered a is 0.40% and for w it is 2.88%.

Remark 5.5. In practice, it may not be feasible to measure all three homogenized coefficient values $\bar{A}_{11}, \bar{A}_{22}, \bar{A}_{12}$, but only some subset of these. With the same setup as described in this subsection, Table 1 reports the relative errors in the recovery of a and w for various combinations of parameters observations, listed on the left. From these results, it can be inferred that in the prescribed setting, not every component marginal contains the same amount of information about a, w and in practice, one may wish to prioritize

Table 1

Relative estimation errors for w and a based on observed homogenization components.

Observed Component	Rel. Error in w	Rel. Error in a
$\bar{A}_{11}, \bar{A}_{22}, \bar{A}_{12}$	2.88%	0.40%
$\bar{A}_{11}, \bar{A}_{22}$	3.15%	0.836%
$\bar{A}_{11}, \bar{A}_{12}$	3.96%	0.996%
$\bar{A}_{11}$	9.07%	3.16%
$\bar{A}_{12}$	91.07%	21.41%




Fig. 8. Locally periodic Voronoi construction (a-b) 10×10 Voronoi cells, (c) 8000×8000 locally periodic Voronoi cells (plot is 40 fold cell-level sub-sampled and 5-fold in-cell sub-sampled for display – making spatial correlation appear to be on a shortened scale). The total specimen is $2 \cdot 10^4 \times 2 \cdot 10^3$ cells.




Fig. 9. Homogenized locally periodic Voronoi construction, $\bar{A} : S \rightarrow M_2$ corresponding to applying periodic homogenization to every cell of Fig. 8c.

observing certain components to maximize data efficiency. Nevertheless, observing more components yields better recovery of the ground truth.

5.2. Locally periodic microstructures

In this section we explore a more physically realistic data-acquisition setup. We imagine a single large specimen with a spatially varying locally periodic microstructure. That is, $A^\epsilon(x/\epsilon, x)$ is assumed to vary both on a small scale, of $\mathcal{O}(\epsilon)$, and the macroscopic scale of D , of $\mathcal{O}(1)$. In this setting, we show how distributional inversion can leverage natural spatial variations in microstructure for statistical characterization. To employ distributional inverse homogenization simply requires that the measurements of bulk material properties are at points distant enough in the specimen to be approximately i.i.d.. The generative model used, and in particular its dependence on the material properties a and the additive Voronoi weights w , remain unchanged from Section 5.1.

Here, in the way we generate the observed dataset, $\{\bar{A}^{(n)}\}_{n=1}^N$, the Voronoi seeds are not randomly sampled on each cell, but rather, we generate the seeds in such a way such that contiguous cells have Voronoi centers in close but not identical layout. This results in a locally periodic microstructure. We will attempt to recover microstructural information based on homogenized measurements at locations in the large specimen which are far apart enough so that the data becomes close to i.i.d..

The construction of the locally periodic microstructure follows a *copula* construction on the Voronoi seeds, where one can independently specify one- and two -point statistics: spatial correlation and point-wise marginals. It is straightforward to construct Gaussian random fields with specified one and two-point statistics (Sklar, 1959; Nelsen, 2006) and such a field may be used to define a spatially correlated distribution for the Voronoi seeds.

We start by describing the Gaussian random field construction. Let $S = [0, L]^2 = \bigcup_{j \in I} S_j$ a finite partition with each S_j a square of width w with local origin at p_j . Draw $2M$ i.i.d. random functions $s \in C(S; \mathbb{R})$ via

$$s^{(\cdot)} \sim \mathcal{N}(0, (\ell^{-2} \mathbf{I} - \Delta)^{-\nu}), \quad k = 1, 2, \quad m = 1, \dots, M. \quad (30)$$

Notation $s^{(m,k)}$ denotes a specific element from this collection of random fields. Here $\mathcal{N}$ denotes the normal distribution and Δ is the Laplacian operator defined on an extended domain $S' \supset S$ equipped with homogeneous Neumann boundary conditions; it has eigenfunctions $\varphi_{i,o}(x) = \cos(i\pi x_1) \cos(o\pi x_2)$ for i, o positive integers. This function-space definition of a Gaussian measure can be approximately sampled via the truncated Karhunen-Loève expansion (Dashti and Stuart, 2011)

$$s^{(\cdot)}(x) = \sum_{i,o=1}^{J_{\text{KL}}} ((i^2 + o^2)\pi^2 + \ell^{-2})^{-\nu/2} \xi^{(i,o)} \varphi_{i,o}(x); \quad \xi^{(i,o)} \sim \mathcal{N}(0, 1) \text{ i.i.d.} \quad (31)$$

This Gaussian random field has a Gaussian distribution at every point $x \in S$. We now invert to find a spatially varying copula field which is uniformly distributed at every point $x \in S$. The coupla is realized via $s^{(\cdot)}$ as

$$\zeta^{(\cdot)}(x) = F_{\mathbf{N}}(s^{(\cdot)}(x)) \quad (32)$$

where $F_{\mathbf{N}}$ is the scale-normalized CDF of the one-dimensional Gaussian, where the scaling is determined by the variance of the one-dimensional Gaussian. Thus, $\zeta^{(\cdot)}$ will have, ignoring nonstationary boundary effects of (30), point-wise distribution $\text{Unif}(0, 1)$, and will be spatially correlated with two-point statistics controlled by ℓ, ν .

We now use the copula field to construct a slowly varying Voronoi tessellation. To this end we write $\zeta^{(\cdot)} \sim \Xi(\ell, \nu; S)$, where S is included as an argument to emphasize the restriction of the copula to $S \subset S'$. Now specify the Voronoi seeds, for $m = 1 \dots, M$ and $k \in \{1, 2\}$,

$$\zeta^{(m,k)}(x) \sim \Xi(\ell, \nu; S) \quad (33a)$$

$$v^{(j,m,k)} = p_j + w \zeta^{(m,k)}(p_j). \quad (33b)$$

Recall that each S_j a square of width w with local origin at p_j so that the construction centers and scales the seeds locations with respect to the S_j . Index $k \in \{1, 2\}$ allows construction of a point in $S_j \subset \mathbb{R}^2$, j indexes the different cells and m indexes different random Voronoi seeds. The special case where $s^{(m,k)}(x)$ is constant with respect to $x \in S$, rather than a Gaussian random field, yields a periodic Voronoi diagram on S .

We have just used the copula construction to specify Voronoi seed locations. Given the seeds, we use the Laguerre-Voronoi diagrams of Section 5.1, and hence with $v^{(n,1:M,1:2)}$, n in the index set of partition cells I , we can compute $A^{(n)}$ and $\tilde{A}^{(n)}$ the same way as in (27b) and (29b). Fig. 8 shows the locally periodic construction: the first two plots is the zoomed in Voronoi diagram at two different locations showing the local periodicity, the third plot is a larger window of the Voronoi construction where the spatial variations become visible. Fig. 9 shows the three components of the periodically homogenized coefficient in each cell in Fig. 8c, exposing the spatial variability of the copula construction.

To mirror a physically testable setup, we assume the local $\tilde{A}^{(n)}$ are noisily measured at sparse distant locations with the noise distribution $\mathbb{P}_{\xi} := \mathcal{N}(0, \sigma_{\xi}^2 \mathbf{I}_3)$. We attempt to estimate (w, a) from sparsely observed homogenized coefficients using (Per. Inv. Hom. Objective).

Data and Inference: The correlation lengthscale is set to $\ell = 1/100$ and $\nu = 4$ (this information is not inferred from data or needed in any way for inference, the only assumption is that measurement locations are distant enough to be approximately i.i.d.) and we use 100 KL expansion terms to sample from $\Xi(\ell, \nu; S)$, the cell size Q is set to $5 \cdot 10^{-5}$, meaning the full Voronoi diagram contains $4 \cdot 10^8$ cells. The noise $\sigma_{\xi} = 10^{-2}$. We set $a^{\dagger} = (0.5, 2.25, 4)$, $w^{\dagger} = (0, 1.38 \cdot 10^{-10}, 2.78 \cdot 10^{-9})$. To generate data we set $\tau = 10^{-20}$ to have sharp numerically discontinuous Voronoi cells. We use $N = 5 \cdot 10^3$ coefficient observations. We use $N = 5 \cdot 10^3$ coefficient observations. Fig. 10 shows the histograms of the data distribution $\mathbb{P}_{\tilde{A}}^N$ as scatter plots of combinations of the three data components.

Results: For inference we set τ to 10^{-3} to allow gradients to flow through the Voronoi construction and estimate w . Fig. 11 shows the convergence of the objective, and (w, a) against $(a^{\dagger}, w^{\dagger})$. The inference ran in 5×9.33 hours. The relative error on the recovered w is 1.76%, and for a is 0.35%. The relative error on a is 0.383% and the relative error on w is 0.97%.

6. 2D stochastic homogenization

In this section we explore applications of distributional inverse homogenization in 2D stochastic homogenization. As with the previous section, we focus on the Voronoi microstructure model. We first consider an i.i.d. data generating mode, followed by a locally stationary-ergodic Voronoi construction, again based on statistical copulas, supporting Contributions (C2) and (C4). To efficiently implement our inference scheme we propose an adaptive surrogate learning scheme, answering Contribution (C3).




Fig. 10. Histograms and scatter plots of the observed data-distribution (before adding noise for inference). In (d)-(f) the color is representative the local density of particles computed via the mean 10-neighbor distance to nearby particles.




Fig. 11. Convergence of (Per. Inv. Hom. Objective) and (w, a) for the locally periodic microstructure.

6.1. Stationary-ergodic i.i.d. microstructures

We first introduce the setup for the i.i.d. microstructure model, followed by the exposition of the surrogate modelling method, and we finish with details on the dataset and results.

Setup: In this subsection we revisit Dirichlet-distributed volume fractions, ρ ; we make the identification $\varpi = (\rho, a)$ in (Stoch. Inv. Hom. Objective) and attempt to invert the map $(\alpha, a) \mapsto \mathbb{P}_{\bar{A}}$. Prescribing the exact volume fraction of a material type in a Voronoi construction requires direct manipulation of Voronoi seeds,¹ we do not attempt this. Instead, we approximately prescribe volume fractions by randomly assigning cell material-types by sampling from a categorical distribution with category weights sampled from $\text{Dir}(\alpha)$. This use of the Dirichlet distribution to specify the volume fractions mirrors Section 4.2. In the following, for $\rho \in \Delta_{M-1}$, we have $r \sim \text{Cat}(\rho)$ where $r = (0, \dots, 1, \dots, 0)$ is a M -dimensional vector of zeros with a 1 in position k with probability ρ_k . We define i.i.d. generative model for the material microstructure in the n -cell Q via

$$\rho \sim \text{Dir}(\alpha), \quad (34a)$$

$$v^{(j)} \sim \text{Unif}(Q), \quad (34b)$$

$$r^{(j)} \sim \text{Cat}(\rho), \quad (34c)$$

¹ If keeping other parameter such as Voronoi-weights fixed.




Fig. 12. Randomly generated stationary-ergodic (S.E.) microstructure for (a, ρ) randomizing over Voronoi seed locations in Q and $\rho \sim \text{Dir}(\alpha)$. These are realizations of the i.i.d. stationary-ergodic material model.




Fig. 13. Histograms and scatter plots of the observed data-distribution for the i.i.d. stationary-ergodic material model. In (d)–(f) the color is representative the local density of particles computed via the mean 10-neighbor distance to nearby particles.

$$A(x; a, \rho, v^{(1:J_{\text{vor}})}) = \mathbb{I}_2 \sum_{j=1}^{J_{\text{vor}}} \langle a, r^{(j)} \rangle \mathbb{1}_{Q_j(v^{(1:J_{\text{vor}})})(x)}, \quad (34d)$$

where we define

$$Q_j(v) = \{x \in Q : \|x - v^{(j)}\| \leq \|x - v^{(k)}\|; k, j \in \{1, \dots, J_{\text{vor}}\}; k \neq j\}. \quad (35)$$

This procedure can be explained informally as follows: (34a) sample a volume fraction vector for a Dirichlet distribution; (34b) sample J_{vor} Voronoi seeds for the cell Q ; (34c) draw a categorical vector of all zeros expect for one 1 in position m with probabilities given by ρ ; (34d) assemble the microstructure field by constructing the Voronoi diagram by (35) and assigning material value to each cell via the inner product $\langle a, r^{(j)} \rangle$ (notice $r^{(j)} = e^m$ simply picks out the value a_m from a). Fig. 12 shows component $A_{11}(y)$ for two independent random coefficient fields as sampled per (34) with $J_{\text{vor}} = 3 \cdot 10^4$,

Surrogate Training: The objective is to invert the map $(\alpha, a) \mapsto \mathbb{P}_{\bar{A}}$. We are now presented with two challenges for numerical implementation: i) the forward map is very computationally expensive as it requires a Monte-Carlo expectation over Voronoi seeds; ii) we require gradients w.r.t to α for a computation which passes via the discontinuous categorical distribution. Challenge ii) could

be resolved by smoothing out discontinuities, similarly to [Section 5](#), using ([Jang et al., 2017](#); [Figurnov et al., 2018](#)), but then we are still faced with i). The solution we explore in this section is to concurrently learn a neural network directly approximating $\mathcal{G}^\dagger$ of (2.1.2) with $\mathcal{G}^{\phi^*} : (\rho, a) \mapsto \bar{A}$ with adaptive surrogate training as outlined in [Section 2.2.2](#). In estimating θ we run a gradient-flow on a pseudo-time t , representing the steps in the optimization routine. At a given step t we update our surrogate model $\mathcal{G}^{\phi_t^*}$ with

$$\phi_t^* = \arg \min_{\phi} \mathbb{E}_{(\rho, a), \bar{A} \sim \mathbb{P}^t} \left[\|\mathcal{G}^\phi(\rho, a) - \bar{A}\|^2 \right], \quad (36a)$$

$$\mathbb{P}^t = \frac{1}{t + N_{\text{pre}}} \sum_{i=1-N_{\text{pre}}}^t \delta_{((\rho, a)^{(i)}, \mathcal{G}^\dagger((\rho, a)^{(i)}))}, \begin{cases} (\rho^{(i)}, a^{(i)}); \rho^{(i)} \sim \text{Dir}(\alpha^0), & \text{for } i \leq 0, \\ (\rho^{(i)}, a^{(i)}); \rho^{(i)} \sim \text{Dir}(\alpha^i), & \text{for } 0 < i \leq t. \end{cases} \quad (36b)$$

In (36a) we select ϕ^* at iteration t to minimize the input-output error for pairs of training data drawn from the adaptive dataset which is defined via (36b). Notice, we do not use the exact $\phi^*(\theta)$ in (16a) as we wish to avoid differentiating $\mathcal{G}^\dagger$, hence the pseudo-time t formulation. This particular time stepping scheme is proposed, in a related context, in [Vadeboncoeur et al. \(2025b\)](#). In implementing this scheme we withhold updating $\mathbb{P}^t$ until 100 new samples $(\rho^{(i)}, a^{(i)})$ are ready for evaluation under $\mathcal{G}^\dagger$, to make the most of hardware parallelization. The surrogate model, $\mathcal{G}^\phi$, is a 4-layer, 32-neuron, feedforward network with “swish” activation functions ([Ramachandran et al., 2017](#)) and a layer-wise Lipschitz number constrained to be under 50 ([Gouk et al., 2021](#)).

Remark 6.1. In a similar spirit to [Remark 5.4](#), we comment on bounds on the distribution of homogenized coefficients. In this section, volume fractions are random vectors drawn from a Dirichlet distribution; for any parameter setting of the Dirichlet distribution, it is always possible to draw a vector of volume fractions which represents a “pure material” (volume fraction of the form $(1, 0, 0)$, $(0, 1, 0)$, or $(0, 0, 1)$). Hence, the lower and upper bounds over the distribution of bulk material properties in this setting are simply $a^- = \min(a^\dagger)$ and $a^+ = \max(a^\dagger)$. This is readily visible from [Fig. 13](#) (d)-(f) by noting that the data-generating $a^\dagger = (1, 2, 3)$.

Data and Inference: We generate $5 \cdot 10^3$ data samples to be used for inference. We do not include noise in this experiment. We set $\alpha^\dagger = (0.66, 1.32, 2)$ and $a^\dagger = (1, 2, 3)$. The Voronoi construction is fully discontinuous. [Fig. 13](#) shows the data distribution in terms of histograms and scatter plots. We set $J_{\text{vor}} = 3 \cdot 10^4$ seeds per cell Q_n . The data is generated using single-sample Monte-Carlo estimate of $\bar{A}^{(n)}$, in contrast to this, the i.i.d. model used for inference uses $J_{\text{avg}} = 10$ Monte-Carlo samples.

Results: The surrogate modelling dataset is limited to a size 2500 evaluations of $\mathcal{G}^\dagger$. Every random restart includes acquiring the training dataset of evaluations of $\mathcal{G}^\dagger$ for accumulation in $\mathbb{P}^t$ along with training the surrogate $\mathcal{G}^\phi$. For initialization, (w, a) are drawn independently from $\text{Unif}(0.2, 2)$ and $\text{Unif}(0.5, 4)$ respectively, and i.i.d. with respect to each component of w and of a . This is on the same order of compute time as periodic homogenization case – which did not use a surrogate model – despite: i) the computational mesh used in $F^\dagger$ being 500×500 for stochastic homogenization instead of 100×100 in the periodic case; ii) each compute of $\bar{A}$ requiring 10 evaluations of $F^\dagger$ in the Monte-Carlo stochastic estimation. [Fig. 14](#) shows the convergence of the algorithm on (α, a) . By using the surrogate construction with an adaptively acquired training set we only compute the expensive $\mathcal{G}^\dagger$ map 2500 times and instead, in ([Stoch. Inv. Hom. Objective](#)) evaluate the inexpensive $\mathcal{G}^\phi$ map $2.5 \cdot 10^7$ times. The relative error on a is 2.98% and the relative error on w is 0.533%.

6.2. Locally stationary-ergodic microstructures

In this, our final example, we consider a material field with two approximately stationary-ergodic levels: a microscopic stationary-ergodic property which allows us to apply stochastic homogenization in a small cell locally, and a macroscopic stationary-ergodic property which allows us to estimate relevant statistics over the entire large specimen. In this context, distributional inverse homogenization can infer statistical properties of the microstructure, without needing to estimate or assume specific spatial correlation information, only that the measured bulk properties are separated well enough so as to be approximately i.i.d.. We now introduce this construction.

Using a copula, $\Xi(\ell, \nu; S)$, defined in [Section 5.2](#), we consider $\rho(x)$, a spatially correlated random vector field describing local material volume fractions with $\text{Dir}(\alpha)$ marginals. Such a field is constructed via

$$\zeta^{(m)}(x) \sim \Xi(\ell, \nu; S) \quad (37a)$$

$$Y^{(m)}(x) = F_{\text{Gamma}(\alpha_m, 1)}^{-1}(\zeta^{(m)}(x)), \quad (37b)$$

$$\rho_{(m)}(x) = \frac{Y^{(m)}(x)}{\sum_{j=1}^M Y^{(j)}(x)}, \quad m = 1 \dots, M. \quad (37c)$$

We abbreviate this procedure as

$$\rho \sim \Lambda(\alpha; \nu, \ell, S). \quad (38)$$

We can now summarize this procedure informally as follows: (37a) sample M independent uniform copula fields; (37b) for copula field value at location x , evaluate the inverse CDF of the Gamma distribution, this yields M Gamma distributed values at x ; (37c) normalize each M Gamma distributed values by their sum, the resulting values are Dirichlet distributed and, for varying x , spatially correlated. Following this construction, $\rho(x) \sim \text{Dir}(\alpha) \forall x \in S$, and $\text{Cov}(\rho(x), \rho(x'))$ is proportional to (ℓ, ν) . We note, the inverse CDF of the Gamma




Fig. 14. Convergence of Objectives and (α, a) for the i.i.d. stationary-ergodic material model.

distribution at a value y , $F_{\text{Gamma}(\alpha_m, 1)}^{-1}(y)$, is the root of $y - F_{\text{Gamma}(\alpha_m, 1)}(x)$, where $F_{\text{Gamma}(\alpha_m, 1)}$ is the CDF; we can approximately compute the inverse CDF via bisection iterations. Now to construct an approximately locally and globally stationary-ergodic field on a specimen S and a subdomain S_p , sample

$$\rho \sim \Lambda(\alpha, \ell, \nu; S) \quad (39a)$$

$$v^{(j)} \sim \text{Unif}(S_p), \quad (39b)$$

$$r^{(j)} \sim \text{Cat}(\rho(v^{(j)})), \quad (39c)$$

$$A(x; a, \rho, v^{(1:J_{\text{vor}})}) = I_2 \sum_{j=1}^{J_{\text{vor}}} \langle a, r^{(j)} \rangle \mathbb{1}_{Q_j}(x). \quad (39d)$$

We can now summarize this procedure informally as follows: (39a) sample a random vector valued Dirichlet Copula field; (39a) for a select sub-specimen S_p sample J_{vor} seeds; (39c) at each seed location evaluate the Dirichlet copula field, this return a vector of entries which sum to 1, this the local volume fraction, draw a random categorical vector for this volume fraction vector; (39c) assemble the microstructure field using the seeds to specify the Voronoi diagram and the categorical variable to assign material property to each Voronoi cell. In Fig. 15 we show the three components of $\rho \sim \Lambda(\alpha, \ell, \nu)$ on a $1/20 \times 1/20$ subset of S , notice the sum of the three fields is equal 1 at all spatial locations. Collections of evaluations of $\rho(x)$ at sparse locations will be Dirichlet distributed. Fig. 16 shows the locally stochastic homogenization of random Voronoi diagrams where the volume fractions are specified by the field in Fig. 15.

Data and Inference: The number of Voronoi seeds per sub-specimen S_p is $3 \cdot 10^4$. The inter-seed distance in the Voronoi construction with respect the Specimen size is roughly 1 to $6 \cdot 10^5$. We set the number of observed homogenized coefficients to $N = 5 \cdot 10^3$, the noise standard deviation $\sigma_\xi = 5 \cdot 10^{-3}$, $(\ell, \nu) = (1/500, 4)$, $J_{\text{vor}} = 3 \cdot 10^4$ for S_p in S . N_{srgt} . Each observed data $\bar{A}^{(n)}$ is generated from a single $\bar{A}_n$ defined on S_p , hence it is a single-sample Monte-Carlo estimate of $\bar{A}^{(n)}$; in contrast to this, the i.i.d. model used for inference uses $J_{\text{avg}} = 10$ Monte-Carlo samples.

Results: The surrogate modelling dataset is limited to a size 2500 evaluations of $\mathcal{G}^\dagger$. We use 5 random restarts for different initializations of (α, a) , these are drawn independently from $\text{Unif}(0.2, 2)$ and $\text{Unif}(0.5, 4)$ respectively. We report the run achieving the




Fig. 15. Dirichlet marginal copula, $\Lambda(\alpha, \ell, \nu; S)$, components on a small $1/20 \times 1/20$ subset of S . Notice, $\sum_{m=1}^M \rho_m(x) = 1$. This sub-specimen contains 100×100 sub-specimen cells, Q_n , which appear in Fig. 12.




Fig. 16. The independent components of the homogenized coefficient field on a subset of a $1/20 \times 1/20$ subset of S where the volume fractions are specified by the random field in Fig. 15. This sub-specimen contains 100×100 computational cells Q_n .

lowest value of (Stoch. Inv. Hom. Objective) over the last 100 iterations (Fig. 17). Here, again, the expensive map $\mathcal{G}^\dagger$ is computed 2500 times and $\mathcal{G}^\phi$ is evaluated $2.5 \cdot 10^7$ times. The relative error on a is 0.46% and the relative error on w is 1.32%.

7. Elasticity

In this section we demonstrate that the ideas developed in the paper can be applied to the complex, and computationally demanding, setting of elasticity. We apply distributional inversion to the case of elastic homogenization on an i.i.d. periodic microstructure similar to that of Section 5. We begin with a brief introduction to homogenization for elasticity, followed by a description of the inference task, a discussion of the results and a summary.

The elasticity problem is an elliptic vector-valued counterpart of the scalar elliptic (3). Given Dirichlet boundary conditions, it has the form

$$-\nabla \cdot (C^\epsilon \nabla^{\text{sym}} u^\epsilon) = f, \text{ for } x \in D \subset \mathbb{R}^2, \quad (40a)$$

$$u^\epsilon(x) = 0, \text{ for } x \in \partial D. \quad (40b)$$

where $u = (u_1, u_2)^\top$ is two-dimensional displacement, and the symmetric gradient $\nabla^{\text{sym}} u = \frac{1}{2}(\nabla u + (\nabla u)^\top)$. Here $C^\epsilon : D \rightarrow \mathcal{V}^4$ is a fourth order tensor field. Analogously to the scalar elliptic case, $C^\epsilon(x) = C(x/\epsilon)$ for small ϵ . Also analogously to the scalar elliptic case, the homogenization task is to identify a constant 4-tensor $\bar{C}$ for which the solution to (40), with $\bar{C}$ in place of C^ϵ , approximates (40). In this work we will focus on the plane strain 2D variant of elasticity (Cioranescu and Donato, 1999). For a given periodic 4-tensor material field, C , one first solves the vector-valued corrector problems for $\chi^{kl} \in H_{\text{per}}^1(\mathbb{T}^2; \mathbb{R}^2)$, $kl \in \{1, 2\}$ which weakly satisfy

$$-\frac{\partial}{\partial y_j} \left(c_{ijop}(y) \frac{\partial \chi_o^{kl}}{\partial y_p} \right) = \frac{\partial c_{ijkl}}{\partial y_j}, \quad (41a)$$

$$\int_{\mathbb{T}^2} \chi^{kl}(y) dy = 0. \quad (41b)$$

The homogenized coefficient is then computed from the formula

$$\bar{C}_{ijkl} = \int_{\mathbb{T}^2} \left(c_{ijkl}(y) + c_{ijop}(y) \frac{\partial \chi_o^{kl}}{\partial y_p} \right) dy. \quad (42)$$




Fig. 17. Convergence of objectives and (α, a) for the locally stationary-ergodic microstructure model.

The symmetries of $\bar{\mathbf{C}}$ imply 6 independent tensor entries.

We specify the distributional inverse homogenization problem as follows. The constituent material properties (a_1, a_2, a_3) and Voronoi weights (w_1, w_2, w_3) parameterise a spatially varying Young's modulus

$$E(y; a, w) = \sum_{m=1}^3 a_m \mathbb{1}_{Q_m(v^{(1:M)}; w)}(y), \quad (43)$$

where the Voronoi cells, Q , are defined as in (26), (27). The periodic field $E(y)$ and constant Poisson ratio (assumed known and fixed at 0.3 for data generation and inference) parametrise the periodic 4-tensor field $\mathbf{C}(y)$. To use automatic differentiation in minimizing (Per. Inv. Hom. Objective), we apply the continuous relaxation of (29b) to (43). The inference task is to identify (a, w) from collections of homogenized coefficients $\bar{\mathbf{C}}$ generated via randomisation of the Voronoi seeds.

Data and Inference: We use $N = 5 \cdot 10^3$ data samples. We set $a^\dagger = (1, 2, 4)$, $w^\dagger = (0, 0.05, 0.1)$. To generate data we set τ to 10^{-20} to have sharp numerically discontinuous Voronoi cells. We do not include noise in this experiment. We fix the computational mesh to be 100×100 when representing $\mathbf{C}(y)$ and in solving for the vector-valued corrector functions with finite elements in (41).

Results: For inference we set τ to 10^{-3} to allow gradients to flow through the Voronoi construction and estimate w . Fig. 18 shows the convergence of the objective, and (w, a) against $(w^\dagger, a^\dagger)$. We run 5 random restarts with (w, a) independently drawn from $\text{Unif}(0.01, 0.5)$ and $\text{Unif}(0.5, 3)$ respectively, and independently component-wise, and extend the optimization of the initialization achieving lowest Objective value by an additional 1000 iterations. The relative error on the recovered a is 2.57% and for w it is 6.23%.

Summary The work in this section shows that the ideas developed for scalar thermo-mechanical problems earlier in the paper do extend to the mechanical problems of elasticity. The concept that it is possible to invert for the statistics of microstructure, and that it is feasible to do so by means of distributional matching, have been shown to apply. The increased computational burden for the vector-valued elasticity problems suggests that use of surrogate modelling, to evaluate the microstructure to macroscopic response map, has great potential in this area; the paper (Akyildiz et al., 2025) provides a methodology that is well-adapted to this task; and the paper (Bhattacharya et al., 2024) provides theory concerning the ability to approximate the microstructure to macroscopic response map using neural operators.




Fig. 18. Convergence of (Per. Inv. Hom. Objective) and (w, a) for the distributional inverse elastic homogenization.

8. Conclusions and future work

In this paper we present a new approach for investigating the microscopic properties of materials by way of measurement of large collections of bulk material properties. Using this approach it is possible to leverage indirect, homogenized, observed quantities to characterize microstructure at a distributional level, both in the case of periodic and stochastic homogenization. In practice one may think of collecting the bulk material property data by making measures at different, decorrelated, points in the same specimen. Algorithmically we model this by using spatial copula. Interesting directions for future work include generalization to two and three dimensional elasticity; a starting point would be to first consider similar scalar-valued microstructures which results in 4-tensor homogenized coefficients. We anticipate that going from measured symmetric 2×2 homogenized coefficient, $\bar{A}$, which has three independent components, to measuring up to 21 independent components for fourth order elasticity tensors, $\bar{C}$, may provide the help needed to tackle this broader class of problems. The direction we describe and deploy in order to accelerate computations, using surrogate modelling, is also applicable to elasticity.

CRedit authorship contribution statement

Arnaud Vadeboncoeur: Writing – review & editing, Writing – original draft, Software, Methodology, Investigation, Conceptualization; **Mark Girolami:** Writing – review & editing, Funding acquisition; **Kaushik Bhattacharya:** Writing – review & editing, Investigation, Conceptualization; **Andrew M. Stuart:** Writing – review & editing, Methodology, Investigation, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Arnaud Vadeboncoeur, Mark Girolami reports financial support was provided by Engineering and Physical Sciences Research Council. Mark Girolami reports financial support was provided by Royal Academy of Engineering. Kaushik Bhattacharya reports financial support was provided by Office of Naval Research. Andrew M. Stuart reports financial support was provided by US Department of Defense. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

The authors are thankful to John R Willis for conversation on the topic of homogenization which helped improve this paper.

Data availability

Code is shared in attached file step.

Appendix A. A first probability distribution on the simplex

In this section we prove [Theorem 4.1](#), and also extend it to the case $M > 2$. In 1D it is more convenient to work with reciprocals of material properties and so we define $c_m = a_m^{-1}$, $\bar{c} = (\bar{a})^{-1}$ and write $c = (c_1, \dots, c_M)$. Then [\(23\)](#) may be written

$$\bar{c} = \langle c, \rho \rangle. \quad (\text{A.1})$$

Using the fact that $\rho_M = 1 - \sum_{m=1}^M \rho_m$ we can write

$$\bar{c} = c_M \sum_{m=1}^{M-1} (c_M^{-1} c_m - 1) \rho_m + c_M. \quad (\text{A.2})$$

Thus we have established the explicit dependence of $\bar{a}$ (equivalently $\bar{c}$) on the M material parameters c and the $M - 1$ volume fraction parameters $(\rho_1, \dots, \rho_{M-1})$. We now explore a volume-fraction distribution on the simplex and the map taking its parameters to the distribution of homogenized coefficients. We need a preliminary technical result. To state the result, note that any uniform distribution on an interval may be written as $\text{Unif}(b - l, b + l)$ and we refer to b as the *shift* and $2l$ as the *width* of the probability distribution.

Lemma A.1. *The convolution of n uniform distributions uniquely identifies the set of all n widths and the sum of the n shifts of the individual distributions.*

Proof. The Fourier transform (characteristic function) of $f := \text{Unif}(b - l, b + l)$ is given by Bradley and Gupta (2002)²

$$\hat{f}(\xi) = \int_{b-l}^{b+l} \frac{1}{2l} e^{-i2\pi\xi t} dt = e^{-i2\pi b\xi} \text{sinc}(2\pi l\xi),$$

Now let $f_j := \text{Unif}(b_j - l_j, b_j + l_j)$. Then

$$(f_1 * \dots * f_n)(\xi) = \prod_{j=1}^n e^{-i2\pi b_j \xi} \text{sinc}(2\pi l_j \xi).$$

On $L^2(\mathbb{R})$, the Fourier transform is bijective, and as $e^{-i2\pi\xi}$ has no complex zeros,

$$\prod_{j=1}^n e^{-i2\pi b_j \xi} \text{sinc}(2\pi l_j \xi) = \prod_{j=1}^n e^{-i2\pi b'_j \xi} \text{sinc}(2\pi l'_j \xi) \quad (\text{A.3a})$$

$$\implies \left\{ \xi \in \mathbb{R} : \prod_{j=1}^n \text{sinc}(2\pi l_j \xi) = 0 \right\} = \left\{ \xi \in \mathbb{R} : \prod_{j=1}^n \text{sinc}(2\pi l'_j \xi) = 0 \right\}. \quad (\text{A.3b})$$

The zeros of $\prod_{j=1}^n \text{sinc}(2\pi l_j \xi)$ are $\cup_{j=1}^n \{k/2l_j\}_{k \in \mathbb{N} \setminus \{0\}}$. These uniquely identify $\cup_{j=1}^n \{l'_j\} = \cup_{j=1}^n \{l_j\}$.

Since $l_j = l'_j$, the identity (A.3a) implies that

$$\prod_{j=1}^n e^{-i2\pi b_j \xi} = \prod_{j=1}^n e^{-i2\pi b'_j \xi}.$$

Hence, $\exp(-i2\pi\xi \sum_{j=1}^n b_j) = \exp(-i2\pi\xi \sum_{j=1}^n b'_j)$, $\forall \xi \in \mathbb{R}$. Thus we can identify $\sum_{j=1}^n b_j$. $\square$

Now we address Theorem 4.1, and generalization to $M > 2$. Define, for $m = 1, \dots, M - 1$,

$$\mathbb{P}_{\rho_m} = \text{Unif}(b_j - l_j, b_j + l_j).$$

We assume that the $l_j > 0$ so that the uniform distribution has support with positive Lebesgue measure. To ensure positivity of the ρ_j drawn from this distribution we require $b_j - l_j \geq 0$. We now construct a probability measure by taking the independent product of these measures, to obtain

$$\mathbb{P}_{\rho_{1:M-1}} = \prod_{j=1}^{M-1} \mathbb{P}_{\rho_j}.$$

It remains to specify the choice of ρ_M in a manner which ensure $\rho \in \triangle_{M-1}$, the probability simplex. To this end we add the assumption $\sum_{m=1}^{M-1} b_j + l_j \leq 1$ and define

$$f(\rho_{1:M-1}) = 1 - \sum_{j=1}^{M-1} \rho_j.$$

Then we set $\mathbb{P}_\rho = (\text{Id}(\cdot), f(\cdot))_{\#} \mathbb{P}_{\rho_{1:M-1}}$, thereby creating a distribution on $\triangle_{m-1}$. Define $\theta = (\sum_{m=1}^{M-1} b_m, \{l_m\}_{m=1}^{M-1})$. Recall formula (24) for the homogenized diffusion coefficient. The distribution over these coefficients implied by $\mathbb{P}_\rho$ is thus

$$\mathbb{P}_{\bar{a}} = F(\cdot; a)_{\#} \mathbb{P}_\rho(\theta).$$

Although the homogenization map, $\rho \mapsto F(\rho; a)$, is not injective, we have the following form of distributional injectivity:

Theorem A.3. *Fix $a \in (0, \infty)^m$. Then map $\rho \mapsto (F(\cdot; a))_{\#} \mathbb{P}_\rho(\theta)$ is injective.*

² Within this proof we use the symbol ξ to denote wave number. This is not to be confused with its use as the noise variable in the main body of the paper.

Proof. The map $\bar{a} \mapsto \bar{c}$ is bijective, so it suffices to study pushforward of $\mathbb{P}_\rho(\theta)$ onto $\bar{c}$. Note from (A.2), $\bar{c}c_M^{-1} - 1 = \sum_{m=1}^{M-1} (c_M^{-1}c_m - 1)\rho_m$. Hence, we can re-write $\mathbb{P}_{\bar{c}} = (\langle \cdot, c \rangle)_\# \mathbb{P}_\rho(\theta)$ as a convolution of scaled uniforms on $\rho_{1:M-1}$. Using Lemma A.1, the distribution of $\bar{c}$ identifies the widths of all uniforms and the sum of centers of the uniforms on $\rho_{1:M-1}$. Thus, the distributions of $\bar{a}$ identifies θ . The distribution of ρ_M is given by the distribution of $\sum_{m=1}^{M-1} \rho_m - 1$. $\square$

Appendix B. Theorem 4.2: Dirichlet limit identifiability

We find it useful to think of the M -dimensional Dirichlet parameters α as $\alpha := \beta\gamma$ where $\gamma \in \Delta_{M-1}$ and $\beta \in (0, \infty)$, noting that the condition $\sum_{m=1}^M \gamma_m = 1$ implies that we still only have M degrees-of-freedom. Recall the standard notation $\text{Dir}(\alpha)$ for a Dirichlet distribution, introduced in (25), and the definition $\text{Dir}(\gamma, \beta) := \text{Dir}(\beta\gamma)$, for $\gamma \in \Delta_{M-1}$, in the text that follows (25). For our purposes it is useful to extend the definition of the Dirichlet distribution to $\beta \in [0, \infty]$. To do this, we prove the existence of the limits $\beta = 0, \beta = \infty$.

Lemma B.1. Let $\gamma \in \Delta_{M-1}$ and $\beta \in (0, \infty)$. Define $\text{Dir}_0(\gamma) := \sum_{m=1}^M \gamma_m \delta_{e^m}$ and $\text{Dir}_\infty(\gamma) := \delta_\gamma$. Then, $\text{Dir}(\gamma, \beta) \xrightarrow{\beta \rightarrow 0} \text{Dir}_0(\gamma)$ and $\text{Dir}(\gamma, \beta) \xrightarrow{\beta \rightarrow \infty} \text{Dir}_\infty(\gamma)$.

Proof. Let $\alpha = \beta\gamma$. We start with the first limit. Let $\rho \sim \text{Dir}(\alpha)$. Then $\mathbb{E}_{\text{Dir}(\alpha)}[\rho_k] = \alpha_k / \sum_{m=1}^M \alpha_m$. Thus i): $\mathbb{E}_{\text{Dir}(\beta\gamma)}[\rho_k] = \gamma_k, \forall \beta > 0$. Also, ii): from (25) we can see that for any $A \subset \Delta_{M-1}$ such that $e^m \notin A$ for any $m \in \{1, \dots, M\}$, then,

$$\lim_{\beta \rightarrow 0} \int_A \text{Dir}(\rho; \gamma, \beta) d\rho = 0.$$

Thus, probability concentrates on the vertices of the simplex as $\beta \rightarrow 0$. Points i) and ii) can be used to conclude convergence in distribution: $\text{Dir}(\beta\gamma) \rightarrow \sum_{m=1}^M \gamma_m \delta_{e^m}$ as β goes to zero.

We now prove the second limit. To sample from $\text{Dir}(\beta\gamma), \beta > 0$, first sample $x_j \sim \text{Gamma}(\beta\gamma_j, 1)$ where

$$\text{Gamma}(x; \beta\gamma_j, 1) = \frac{x^{\beta\gamma_j-1} e^{-x}}{\Gamma(\beta\gamma_j)}.$$

Next set

$$\rho_j = \frac{x_j}{\sum_{k=1}^M x_k},$$

thus producing $\rho \sim \text{Dir}(\beta\gamma)$. For $\beta \geq 1$ the density $\text{Gamma}(x; \beta\gamma_j, 1)$ has unique mode at $x^* = (\beta\gamma_j - 1)$ and approaches a Gaussian with mean and variance $\beta\gamma_j$; this may be shown using the central limit theorem. As a consequence, as β increases, the Dirichlet density behaves like

$$\rho_j = \frac{\beta\gamma_j + \mathcal{O}(\sqrt{\beta})}{\sum_{k=1}^M \beta\gamma_k + \mathcal{O}(\sqrt{\beta})}.$$

Applying l'Hôpital reveals, for $\beta \rightarrow \infty, \rho_j = \gamma_j$. These calculations can be used to establish convergence in distribution: $\text{Dir}(\beta\gamma) \rightarrow \delta_\gamma$. $\square$

We now define the 1D distributional homogenization map

$$\mathbb{P}_{\bar{a}}(\gamma, \beta, a) = F(\cdot; a)_\# \text{Dir}(\gamma, \beta).$$

The following proves Theorem 4.2.

Lemma B.2. The map $(\gamma, \beta, a) \mapsto F(\cdot; a)_\# \text{Dir}(\gamma, \beta)$ is injective at $\beta = 0$, and is not injective at $\beta = \infty$.

Proof. We start with $\beta = 0$. We work with $\bar{a}^{-1} = \bar{c}$ and $a_m^{-1} = c_m$. Note that the distribution $(\langle c, \cdot \rangle)_\# \text{Dir}_0(\gamma) = \sum_{m=1}^M \gamma_m \delta_{c_m e^m}$. This distribution identifies (γ, β, c) which in turn identifies (γ, β, a) .

For $\beta = \infty, (\langle c, \cdot \rangle)_\# \text{Dir}_\infty(\gamma) = \delta_{(c, \gamma)}$, which does not identify (γ, c) and so does not identify (γ, β, a) . $\square$

References

- Akyildiz, O.D., Girolami, M., Stuart, A.M., Vadeboncoeur, A., 2025. Efficient prior calibration from indirect data. *SIAM J. Sci. Comput.* 47 (4), C932–C958.
- Al-Ostaz, A., Diwakar, A., Alzebedh, K.I., 2007. Statistical model for characterizing random microstructure of inclusion–matrix composites. *J. Mater. Sci.* 42 (16), 7016–7030.
- Bensoussan, A., Lions, J.-L., Papanicolaou, G., 2011. *Asymptotic Analysis for Periodic Structures*. Vol. 374. American Mathematical Soc.
- Bernton, E., Jacob, P.E., Gerber, M., Robert, C.P., 2019. On parameter estimation with the wasserstein distance. *Inf. Inference: J. IMA* 8 (4), 657–676.
- Bhattacharya, K., Kovachki, N.B., Rajan, A., Stuart, A.M., Trautner, M., 2024. Learning homogenization for elliptic operators. *SIAM J. Numer. Anal.* 62 (4), 1844–1873.
- Blanc, X., Le Bris, C., Legoll, F., 2016. Some variance reduction methods for numerical stochastic homogenization. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* 374 (2066), 20150168.
- Bond-Taylor, S., Leach, A., Long, Y., Willcocks, C.G., 2021. Deep generative modelling: a comparative review of vaes, gans, normalizing flows, energy-based and autoregressive models. *IEEE Trans. Pattern Anal. Mach. Intell.* 44 (11), 7327–7347.
- Bonneel, N., Rabin, J., Peyré, G., Pfister, H., 2015. Sliced and radon wasserstein barycenters of measures. *J. Math. Imaging Vis.* 51 (1), 22–45.
- Bradley, D.M., Gupta, R.C., 2002. On the distribution of the sum of n non-identically distributed uniform random variables. *Ann. Inst. Stat. Math.* 54 (3), 689–700.
- Catalanotti, G., 2016. On the generation of RVE-based models of composites reinforced with long fibres or spherical particles. *Compos. Struct.* 138, 84–95.

- Chatzopoulos, M., Koutsourelakis, P.-S., 2024. Physics-aware neural implicit solvers for multiscale, parametric PDEs with applications in heterogeneous media. *Comput. Methods Appl. Mech. Eng.* 432, 117342.
- Chen, X., Yang, Y., Li, Y., 2020. Augmented sliced wasserstein distances. In: *International Conference on Learning Representations*.
- Cherkaev, E., 2001. Inverse homogenization for evaluation of effective properties of a mixture. *Inverse Probl.* 17 (4), 1203–1218.
- Cherkaev, E., Golden, K.M., 1998. Inverse bounds for microstructural parameters of composite media derived from complex permittivity measurements. *Waves Random Media* 8 (4), 437–450.
- Cioranescu, D., Donato, P., 1999. *An Introduction to Homogenization*. Oxford university press.
- Dashti, M., Stuart, A.M., 2011. Uncertainty quantification and weak approximation of an elliptic inverse problem. *SIAM J. Numer. Anal.* 49 (6), 2524–2542.
- Deshpande, I., Hu, Y.-T., Sun, R., Pyros, A., Siddiqui, N., Koyejo, S., Zhao, Z., Forsyth, D., Schwing, A.G., 2019. Max-sliced wasserstein distance and its use for gans. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10648–10656.
- Figurnov, M., Mohamed, S., Mnih, A., 2018. Implicit reparameterization gradients. *Adv. Neural Inf. Process. Syst.* 31.
- Flamary, R., Courty, N., Gramfort, A., Alaya, M.Z., Boisbunon, A., Chambon, S., Chapel, L., Corenflos, A., Fatras, K., Fournier, N., Gautheron, L., Gayraud, N. T.H., Janati, H., Rakotomamonjy, A., Redko, I., Rolet, A., Schutz, A., Seguy, V., Sutherland, D.J., Tavenard, R., Tong, A., Vayer, T., 2021. Pot: python optimal transport. *J. Mach. Learn. Res.* 22 (78), 1–8.
- Flamary, R., Vincent-Cuaz, C., Courty, N., Gramfort, A., Kachaiev, O., Quang Tran, H., David, L., Bonet, C., Cassereau, N., Gnassounou, T., Tanguy, E., Delon, J., Collas, A., Mazalet, S., Chapel, L., Kerdoncuff, T., Yu, X., Feickert, M., Krzakala, P., Liu, T., Fernandes Montesuma, E., Neike, N., Genest, B., Coeurjolly, D., Germain, T., O'Shea, S., Corneli, M., Genans, F., 2026. Pot python optimal transport (version 0.9.7). <https://doi.org/10.5281/zenodo.17161062>
- Fratzl, P., Weinkamer, R., 2007. Nature's hierarchical materials. *Prog. Mater. Sci.* 52 (8), 1263–1334.
- Gao, H., Kaltenbach, S., Koumoutsakos, P., 2024. Generative learning for forecasting the dynamics of high-dimensional complex systems. *Nat. Commun.* 15 (1), 8904.
- Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B., 1995. *Bayesian Data Analysis*. Chapman and Hall/CRC.
- Gent, A.N., 1970. Hypothetical mechanism of crazing in glassy plastics. *J. Mater. Sci.* 5 (11), 925–932.
- Gibson, L.J., 2003. Cellular solids. *MRS Bull.* 28 (4), 270–274.
- Gloria, A., 2012. Numerical approximation of effective coefficients in stochastic homogenization of discrete elliptic equations. *ESAIM: Math. Model. Numer. Anal.* 46 (1), 1–38.
- Gloria, A., Otto, F., 2011. An optimal variance estimate in stochastic homogenization of discrete elliptic equations. *Ann. Probab.* 39 (3), 779–856.
- Gloria, A., Otto, F., 2012. An optimal error estimate in stochastic homogenization of discrete elliptic equations. *Ann. Appl. Probab.* 22 (1), 1–28.
- Glyn-Davies, A., Vadeboncoeur, A., Akyildiz, O.D., Kazlauskaitė, I., Girolami, M., 2025. A primer on variational inference for physics-informed deep generative modelling. *Philos. Trans. A* 383 (2299), 20240324.
- Goh, H., Sherifdeen, S., Wittmer, J., Bui-Thanh, T., 2019. Solving bayesian inverse problems via variational autoencoders. *arXiv preprint arXiv:1912.04212*.
- Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2014. Generative adversarial nets. In: *Advances in Neural Information Processing Systems* vol. 27.
- Gouk, H., Frank, E., Pfahringer, B., Cree, M.J., 2021. Regularisation of neural networks by enforcing lipschitz continuity. *Mach. Learn.* 110 (2), 393–416.
- Gretton, A., Borgwardt, K.M., Rasch, M.J., Schölkopf, B., Smola, A., 2012. A kernel two-sample test. *J. Mach. Learn. Res.* 13 (1), 723–773.
- Jang, E., Gu, S., Poole, B., 2017. Categorical reparameterization with gumbel-softmax. In: *International Conference on Learning Representations*.
- Kingma, D.P., Ba, J., 2014. Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kolouri, S., Nadjahi, K., Simsekli, U., Badeau, R., Rohde, G., 2019. Generalized sliced wasserstein distances. *newblock in: Advances in Neural Information Processing Systems* vol. 32.
- Kozlov, S.M., 1980. Averaging of random operators. *Sb. Math.* 37 (2), 167–180.
- Li, Q., Oprea, M., Wang, L., Yang, Y., 2024. Stochastic inverse problem: stability, regularization and wasserstein gradient flow. *arXiv preprint arXiv:2410.00229*.
- Li, Q., Oprea, M., Wang, L., Yang, Y., 2025a. Inverse problems over probability measure space. *arXiv preprint arXiv:2504.18999*.
- Li, Q., Wang, L., Yang, Y., 2025b. Least-squares problem over probability measure space. *arXiv preprint arXiv:2501.09097*.
- Liu, R., Kumar, A., Chen, Z., Agrawal, A., Sundararaghavan, V., Choudhary, A., 2015. A predictive machine learning approach for microstructure optimization and materials design. *Sci. Rep.* 5 (1), 11551.
- Martin-Pérez, 1999. *Service life modelling of RC highway structures exposed to chlorides*. University of Toronto, Toronto, ON, Canada, Doctoral dissertation.
- Martin-Pérez, B., Thomas, M., 2001. Numerical solution of mass transport equations in concrete structures. *Comput. Struct.* 79 (13), 1251–1264.
- Martin-Pérez, B., Zibara, H., Hooton, R.D., Thomas, M., 2000. A study of the effect of chloride binding on service life predictions. *Cem. Concr. Res.* 30 (8), 1215–1223.
- Nelsen, R.B., 2006. *An Introduction to Copulas*. Springer.
- Nguyen, K., Ho, N., 2023. Energy-based sliced wasserstein distance. *Adv. Neural Inf. Process. Syst.* 36, 18046–18075.
- Nguyen, K., Ho, N., Pham, T., Bui, H., 2021. Distributional sliced-wasserstein and applications to generative modeling. In: *International Conference on Learning Representations*.
- Panaretos, V.M., Zemel, Y., 2019. Statistical aspects of wasserstein distances. *Annu. Rev. Stat. Appl.* 6 (1), 405–431.
- Pavliotis, G., Stuart, A., 2008. *Multiscale Methods: Averaging and Homogenization*. Springer Science & Business Media.
- Quey, R., Dawson, P.R., Barbe, F., 2011. Large-scale 3d random polycrystals for the finite element method: generation, meshing and remeshing. *Comput. Methods Appl. Mech. Eng.* 200 (17–20), 1729–1745.
- Quey, R., Kasemer, M., 2022. The neper/fepx project: free/open-source polycrystal generation, deformation simulation, and post-processing. In: *IOP Conference Series: Materials Science and Engineering*. Vol. 1249. IOP Publishing, p. 012021.
- Quey, R., Renversade, L., 2018. Optimal polyhedral description of 3d polycrystals: method and application to statistical and synchrotron x-ray diffraction data. *Comput. Methods Appl. Mech. Eng.* 330, 308–333.
- Ramachandran, P., Zoph, B., Le, Q.V., 2017. Searching for activation functions. *arXiv preprint arXiv:1710.05941*.
- Rixner, M., Koutsourelakis, P.-S., 2022. Self-supervised optimization of random material microstructures in the small-data regime. *npj Comput. Mater.* 8 (1), 46.
- Robbins, H.E., 1992. An empirical bayes approach to statistics. In: *Breakthroughs in Statistics: Foundations and Basic Theory*. Springer, pp. 388–394.
- Rollett, A.D., Lee, S.-B., Campman, R., Rohrer, G.S., 2007. Three-dimensional characterization of microstructure by electron back-scatter diffraction. *Annu. Rev. Mater. Res.* 37 (1), 627–658.
- Sejdinovic, D., Sriperumbudur, B., Gretton, A., Fukumizu, K., 2013. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *Ann. Stat.* 41, 2263–2291.
- Senthilnathan, A., Saseendran, V., Acar, P., Yamamoto, N., Sundararaghavan, V., 2024. Comparison and validation of stochastic microstructure characterization and reconstruction: machine learning vs. deep learning methodologies. *Acta Mater.* 278, 120220.
- Sigmund, O., 1994. Materials with prescribed constitutive parameters: an inverse homogenization problem. *Int. J. Solids Struct.* 31 (17), 2313–2329.
- Sklar, M., 1959. Fonctions de répartition à n dimensions et leurs marges. In: *Annales De L'ISUP*. Vol. 8, pp. 229–231.
- Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B., 2020. Score-based generative modeling through stochastic differential equations. In: *International Conference on Learning Representations*.
- Sternfels, R., Koutsourelakis, P.-S., 2011. Stochastic design and control in random heterogeneous materials. *Int. J. Multiscale Comput. Eng.* 9 (4).
- Sundararaghavan, V., Zabarab, N., 2005. Classification and reconstruction of three-dimensional microstructures using support vector machines. *Comput. Mater. Sci.* 32 (2), 223–239.
- Sundararaghavan, V., Zabarab, N., 2006. Design of microstructure-sensitive properties in elasto-viscoplastic polycrystals using multi-scale homogenization. *Int. J. Plast.* 22 (10), 1799–1824.
- Székely, G.J., Rizzo, M.L., 2013. Energy statistics: a class of statistics based on distances. *J. Stat. Plan. Inference* 143 (8), 1249–1272.
- Vadeboncoeur, A., Duthé, G., Girolami, M., Chatzi, E., 2025a. Geometric autoencoder priors for bayesian inversion: learn first observe later. *arXiv preprint arXiv:2509.19929*.

- Vadeboncoeur, A., Girolami, M., Stuart, A.M., 2025b. Efficient deconvolution in populational inverse problems. [arXiv preprint arXiv:2505.19841](#).
- Yi, J., Bessa, M.A., 2025. Cooperative Bayesian and variance networks disentangle aleatoric and epistemic uncertainties. [arXiv preprint arXiv:2505.02743](#).
- Zang, Y., Koutsourelakis, P.-S., 2026. Design-genNO: a physics-informed generative model with neural operators for inverse microstructure design. *Comput. Methods Appl. Mech. Eng.* 450, 118597.
- Zhang, P., Vadeboncoeur, A., Glyn-Davies, A., Girolami, M., 2026a. Hierarchical inference and closure learning via adaptive surrogates for ODEs and PDEs. [arXiv preprint arXiv:2603.03922](#).
- Zhang, R.Z., Miles, C.E., Xie, X., Lowengrub, J.S., 2026b. Bilo: bilevel local operator learning for PDE inverse problems. *J. Comput. Phys.*, 114679.