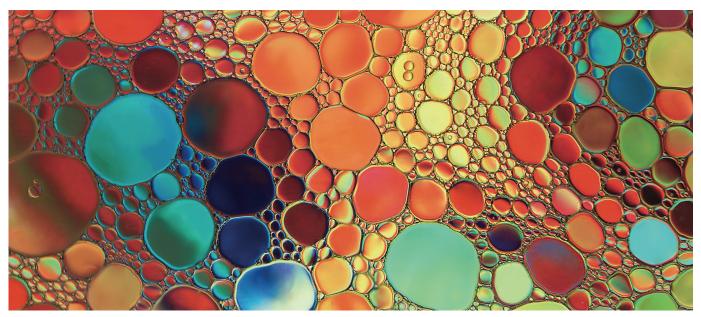
Low-Rank Approximation for Multiscale PDEs



Ke Chen, Shi Chen, Qin Li, Jianfeng Lu, and Stephen J. Wright

1. Introduction

Multiscale phenomena are ubiquitous, with applications in many physical sciences and engineering fields: aerospace, material sciences, geological structure analysis, and many others. The different scales often have

Ke Chen is a postdoctoral researcher in the Department of Mathematics at University of Texas at Austin. His email address is ke.chen@austin.utexas.edu.

Shi Chen is a graduate student in the Department of Mathematics at University of Wisconsin, Madison. His email address is schen636@wisc.edu.

Qin Li is an associate professor in the Department of Mathematics and the Wisconsin Institutes for Discovery at University of Wisconsin, Madison. Her email address is qinli@math.wisc.edu.

Jianfeng Lu is a professor in the Department of Mathematics, the Department of Physics, and the Department of Chemistry at Duke University. His email address is jianfeng@math.duke.edu.

Stephen J. Wright is a professor in the Department of Computer Sciences at University of Wisconsin, Madison. His email address is swright@cs.wisc

Communicated by Notices Associate Editor Reza Malek-Madani.

For permission to reprint this article, please contact:

 ${\tt reprint-permission@ams.org}.\\$

DOI: https://doi.org/10.1090/noti2488

different physics, which entangle to produce complicated nonlinearities. Partial differential equations (PDEs) are often used to model these problems, with different scales captured in the coefficients and functions that define the PDE. These PDE models are challenging to compute directly, so analysis and algorithms specifically targeted to multiscale problems have been developed and investigated. Following convention, we focus in this review on problems with two distinct scales, with a small positive parameter ϵ capturing the ratio between the small and large scales.

Though modern multiscale analysis dates back to asymptotic PDE analysis that was seen already in Hilbert and Poincaré expansions early last century (see review in [PS08]), the impetus for computations involving multiscale PDEs came largely from the US Department of Energy (DOE) National Labs within the ASCI (Advanced Strategic Computing Initiative) [Hor09]. Since that time, analysis and computation in multiscale PDEs have taken different paths. Analysis has tended to follow a single "universal" strategy, passed down from tradition. The equation is decomposed into several levels according to

asymptotic expansions involving the scale parameters, with the subequation at each level representing physics at a single scale, and subequations at the finer level feeding information to those at the coarser level. This analytical machinery has been used to treat multiscale PDEs arising from such varied backgrounds as kinetic theory, semiclassical quantum systems, and homogenization of composite materials, among others [PS08, E11].

On the computational side, strategies for handling multiscale PDEs are more varied. Problems are usually handled by specifically designed solvers. One class of solvers called *asymptotic-preserving* schemes [HJL17] are designed to preserve asymptotic limits of kinetic equations. These schemes usually contain some component of macro-solvers and micro-solvers, integrated in a clever way to reveal different structures in different regimes. Another class of solvers called *numerical homogenization* methods [E11, OS19, EE03] usually target elliptic and parabolic equations in which the coefficients that represent media have oscillatory elements. These methods usually consist of offline and online stages, with either the homogenized media or the representative basis functions being prepared in the offline stage.

Why are most numerical schemes for multiscale PDEs equation-specific despite the analytical tools being largely unified? This intriguing question has motivated our investigations into devising a universal numerical strategy for solving multiscale PDEs. While the approach is yet to be developed fully, we believe that our progress on this issue is of wide interest, and this article surveys our progress to date. Crucially, our approach exploits the low-rank structure present in discretizations of multiscale PDEs.

To demonstrate the fundamental idea, we consider the following problem:

$$\mathcal{L}^{\epsilon}u^{\epsilon} = f, \tag{1}$$

where $\mathcal{L}^{\varepsilon}$ is a linear partial differential operator that depends explicitly on the small parameter ε , while f represents the boundary condition or the source term, which is assumed to have no dependence on ε . Multiscale problems that can be formulated in this way include elliptic equations with highly oscillating media and the neutron transport equation with small Knudsen number. Due to the ε -dependence of the operator, the solution u^{ε} inherits structures at both fine and coarse scales.

An asymptotic limit is revealed by multiscale analysis using asymptotic expansions as $\varepsilon \to 0$. In this limit, the oscillation at the fine scale is fast and the detailed oscillation pattern no longer matters — only macroscopic quantities are relevant. Formally, writing the homogenization limit as

$$\mathcal{L}^* u^* = f, \tag{2}$$

we have

$$||u^{\epsilon} - u^*|| \to 0 \text{ as } \epsilon \to 0.$$
 (3)

The norm of the approximation error depends heavily on the particular equation at hand.

The numerical challenge in solving (1) is that many degrees of freedom may be needed. Naïve finite element or finite difference methods would require mesh size $h \ll \epsilon$ to resolve fine-scale structure of the solution at the ϵ level. For a problem on \mathbb{R}^d , the discretized system would therefore have $O(\epsilon^{-d})$ degrees of freedom, leading to prohibitive computational and memory cost for small ϵ . From an application perspective, it often suffices to characterize the solutions on the macroscopic level, where oscillations at the ϵ scale are largely absent. This property raises the question of whether we can obtain an approximate solution of this type using only O(1) degrees of freedom. If we know how to derive (2), we can simply solve for u^* , which has the required macroscopic properties, and typically requires a discretization with O(1) degrees of freedom. Often, though, the limiting equation (2) and its solution u^* are difficult to find explicitly, even when it is possible to establish their existence. These difficulties have led researchers to propose problem-specific solutions.

We believe that a universal approach for finding the large-scale solution can be devised, and that exploitation of the low-rank structure of the solution space is the key to developing such an approach. As suggested above, the Green's matrix G^{ε} (the discretized Green's function on fine grids) for the multiscale system (1) requires dimension $O(\varepsilon^{-d})$ to represent the underlying Green's function accurately. However, if a limiting system such as (2) exists, this limiting system can be well-represented numerically by G^* , a Green's matrix with dimension only O(1). This phenomenon suggests the system can largely be "compressed" and hence is of low rank; see illustration in Figure 1. In the language of numerical linear algebra, this transition amounts to performing a truncated singular value decomposition (SVD) of G^{ε} to obtain G^* .

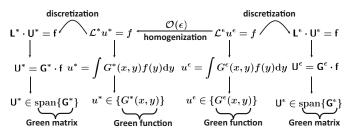


Figure 1. PDEs with small parameters have homogenized limits, meaning the solutions to the original PDEs can be well-approximated by the solutions to the limiting equations. While analytically the two solution spaces are "close," the original equation requires many more degrees of freedom to solve numerically than its limiting counterpart. The numerical Green's matrix is intrinsically low rank.

If we obtain the truncated SVD of the matrix G^{ε} by starting with a full SVD, the resulting algorithms would be impractical because of the large dimension of the matrix and the expense of preparing and storing the full matrix G^{ε} and computing its SVD. Several new linear algebra solvers take a quite different approach. Instead of accessing the full matrix, these new solvers merely require computation of matrix-vector products, involving the target matrix and several randomly selected vectors (typically vectors with Gaussian i.i.d. entries). Translated to the PDE solver setting, these matrix-vector multiplications amount to computing numerical solutions to PDEs with some random source terms, a task that may be practical if the number of such operations required is modest. The randomized SVD (rSVD) approach is one method of this type. It is equipped with a thorough analysis and achieves optimality in terms of computational efficiency. We make use of this method in the techniques described in the remainder of this article.

The main theme of our article, then, is the use of randomized SVD solvers to exploit the low-rank features of multiscale PDEs. We will describe two strategies both of which are divided into "offline" and "online" stages. The offline stage sees the preparation of either the solution space or the boundary-to-boundary map used in the domain decomposition, while the online stage singles out the specific solution for the given source f. The two strategies are described in Section 4.1 and 4.2, respectively. In Section 5, we present the nonlinear extension utilizing manifold learning algorithms for reconstructing the lowrank features of the solution manifold. Prior to these discussions, we describe in Section 2 two algorithm classes — asymptotic preserving and numerical homogenization — for identifying the asymptotic limits of multiscale problems. As examples, we use the multiscale radiative transfer equation (RTE) and the elliptic equation with rough media. Section 3 explores the two main elements of our approaches: the numerical low-rank feature of multiscale PDEs and the randomized SVD solver for efficient reconstruction of low-rank operator/spaces. We conclude with a discussion of future work in Section 6.

2. Examples

Kinetic equations and elliptic equations with oscillating media are two examples of multiscale PDEs, for which computational schemes were developed separately. The specific features of these problems were incorporated into the design of asymptotic preserving schemes and numerical homogenization methods, respectively. We review these techniques and highlight the shared low-rank property of these two problems.

2.1. Kinetic equations and asymptotic preserving methods. Kinetic equations, which originate from statistical mechanics, describe the evolution of probability density

for identical particles in phase space. A model equation, the radiative transfer equation (RTE), characterizes the evolution of photon density. In the steady state, this equation is

$$-v \cdot \nabla_x u^{\varepsilon} + \mathsf{S}^{\varepsilon}[u^{\varepsilon}] = f(x, v), \quad (x, v) \in \mathcal{K} \times \mathbb{V}, \quad (4)$$

where f(x, v) is the light source, and the linear collision operator S^{ϵ} describes the interaction of photons with the optical media. The small parameter ϵ is encoded in this operator.

The operator S^{ε} defines several distinct regimes. In the optically thick regime, it is defined by

$$S^{\epsilon}u(x,v) = \frac{1}{\epsilon} \int_{\mathbb{V}} k(x,v,v')u(x,v')dv' - \frac{1}{\epsilon} \int_{\mathbb{V}} k(x,v',v)u(x,v)dv'.$$
 (5)

In this case, k(x, v, v') is the scattering coefficient that describes the possibility of a photon located at x changing its velocity from v to v', and the parameter ϵ is called the *Knud*sen number, standing for the ratio of the mean free path to the typical domain length. When the medium is optically thick, the mean free path is small, with $\epsilon \ll 1$. This means the photon particles are scattered fairly often, and the system statistically achieves the equilibrium state, which can itself be characterized mathematically. One example is to observe light in atmosphere, where the average mean free path is about 10 m, and the observation is conducted at the scale of 10 km, leading to $\varepsilon \sim 10^{-3}$. By performing asymptotic expansion in terms of ϵ , the inhomogeneity in the velocity domain vanishes, and one can show that $u^{\varepsilon}(x,v)$ asymptotically approximates $u^*(x)$, a function without dependence on v that solves a diffusion equation. We have the following result from [BLP11].

Theorem 1. Suppose that u^{ε} solves (4) with collision term S being isotropic, that is, $k(x, v, v') = \sigma(x)$ for some function σ . Let $\mathcal{K} \subset \mathbb{R}^d$ be bounded with smooth boundary, and $V = \mathbb{S}^{d-1}$. Assume that the boundary condition is

$$u^{\epsilon}(x,v) = \phi(x,v)$$
 on $x \in \partial \mathcal{K}$, $v \cdot n_x < 0$. (6)

Then

$$||u^{\varepsilon} - u^*||_{L_2(\mathrm{d}x\mathrm{d}v)} \to 0, \tag{7}$$

where $u^* = u^*(x)$ solves

$$\nabla_x \cdot \left(\frac{1}{\sigma(x)} \nabla_x u^*(x)\right) = g(x), \quad x \in \mathcal{K},$$
 (8)

with the boundary condition

$$u^*(x) = \xi_{\phi}(x)$$
, on $x \in \partial \mathcal{K}$,

where $\xi_{\phi}(x)$ solves a proper boundary layer equation and g can be obtained from $\int f(x, v) dv$.

This result indicates that the homogenized operator as $\varepsilon \to 0$ is $\mathcal{L}^* \propto \nabla_x \cdot ((1/\sigma)\nabla_x)$. Similar results, when k(x,v,v') fails to have the form of $\sigma(x)$ in the anisotropic optical media, are still available, but the explicit form of \mathcal{L}^* is no longer available.

A second regime of interest for S^{ε} is one in which the media is highly heterogeneous [DG00]:

$$S^{\epsilon}u(x,v) = \int_{\mathbb{N}} k\left(\frac{x}{\epsilon}, v, v'\right) u(x, v') dv' - \int_{\mathbb{N}} k\left(\frac{x}{\epsilon}, v', v\right) u(x, v) dv'.$$
(9)

In this case, the photons go through the media that oscillates at a small scale: For example, sunlight passing through a heavy cloud with a large number of small droplets or a laser beam passing through crystals. The amplitude of k determines the photon scattering frequency. Since k oscillates rapidly, photons also change rapidly between the high- and low-scattering regimes. On a large scale, the photons can be viewed approximately as scattering with an averaged frequency. A mathematical result is as follows [DG00].

Theorem 2. Let the conditions from Theorem 1 hold, and suppose that the collision term S^{ε} is defined in (9). Then

$$||u^{\epsilon} - u^*||_{L_2(\mathrm{d}x\mathrm{d}v)} \to 0, \tag{10}$$

where $u^*(x, v)$ solves

$$-v \cdot \nabla_{\mathbf{x}} u^* + \mathsf{S}^*[u^*] = f(\mathbf{x}, v), \quad (\mathbf{x}, v) \in \mathcal{K} \times \mathbb{V}, \quad (11)$$

where $S^*u(x,v) = \sigma^*(x) \int_{\mathbb{V}} u(x,v') - u(x,v) dv'$ for some $\sigma^*(x)$. Furthermore, if $k(x,v,v') = \sigma(x)$ is periodic in x with period $[0,1]^d$, then $\sigma^* = \int_{[0,1]^d} \sigma(x) dx$.

In special cases, such as under periodic or random ergodic conditions, the function σ^* can be computed explicitly. (There are also works that investigate the asymptotic limit of the RTE when the system is both highly oscillatory and in diffusion regime; see [GM01].)

In both limiting regimes, the limiting equations (8) and (11) can be solved much more efficiently than the original equation (4). The discretization of (4) is constrained strongly by ϵ , due either to stability (as in (5)) or accuracy (as in (9)). By contrast, the solution u^* varies smoothly, containing no ϵ -scale effects, so it can be obtained accurately by applying a discretization with mesh width O(1) to the asymptotic limiting equation. If the latter equation is available, computation of u^* by this means is the recommended methodology.

Methods for kinetic equations are termed "asymptotic preserving" (AP) if they can relax the requirement that the discretization width h satisfies $h = o(\varepsilon)$ yet still capture the asymptotic limits. Many different AP approaches have been proposed. For linear equations, existing AP methods rely on even-odd or micro-macro decomposition. For

nonlinear equations, knowledge of the specific forms of the limits is usually required, and this knowledge is built into the solvers [HJL17]. As mentioned above, these specific forms are often not available, so many AP methods cannot be applied to a large set of multiscale kinetic equations. This observation begs the question: Knowing the existence of the limit, but not its particular form, can we still devise efficient methods for solving kinetic equations? 2.2. Elliptic equations and numerical homogenization. Another class of multiscale equations that has been investigated deeply is elliptic equations with highly oscillatory coefficient. These problems have the form

$$-\nabla_{x} \cdot (a^{\epsilon}(x) \nabla_{x} u^{\epsilon}) = f, \qquad (12)$$

where $\epsilon \ll 1$ is the scale on which the media oscillates. (The source term f has no small-scale contribution.)

This equation is a model problem from petroleum engineering where it is crucial to precompute the underground flow before expensive construction of infrastructure takes place [BL11]. The problem is typically solved on kilometerscale domains, but the heterogeneities in the media can scale at centimeters. Certain forms of this equation can be approximated effectively by an equation that can be solved efficiently. Suppose the media coefficient $a^{\epsilon}(x)$ has the form $a(x, x/\epsilon)$, that is, it varies on two scales (1 and ϵ), and moreover is periodic with respect to the fast variable (the second argument in $a(x, x/\epsilon)$). Then in the limiting regime as $\epsilon \to 0$, the solution u^{ϵ} converges to that of a homogenized equation, with the media "smoothed-out," as described in the following result [All92].

Theorem 3. Let u^{ε} solve (12) in the domain $x \in \mathcal{K}$ with zero boundary condition. Suppose $a(x, x/\varepsilon)$ is periodic with respect to the second argument. Then

$$||u^{\epsilon} - u^*||_{L^2} \lesssim \epsilon ||u^*||_{H^2} , \qquad (13)$$

where u^* solves the following effective equation with zero boundary condition:

$$-\nabla_x \cdot (a^*(x)\nabla_x u^*) = f, \quad x \in \mathcal{K}, \tag{14}$$

where a^* , the effective media, can be computed from a cell problem (See Definition 2.1 in [All92]).

As in the previous section, when a limiting equation can be derived explicitly, the best course for obtaining a useful solution it to solve this equation directly, as the mesh width in the discretization scheme can be much larger than ϵ . See [EH09] for a discussion of a reduced number of basis functions and [E11] for computation of the effective media.

However, the validity and the specific form of the effective limit are known only in special cases like the one described in Theorem 3. In other cases, we seek a solver that relies on as little analytical knowledge as possible. An

approach known as *numerical homogenization* has been investigated extensively. This approach is founded on two principles: a discretization scheme independent of ε , and a numerical solution scheme that captures the true limiting behavior of the solution on the discrete level. Variants of numerical homogenization include application of the \mathcal{H} -matrix, a purely algebraic technique [Hac15]; and a Bayesian approach that views the source f, and hence the solution u^{ε} , as Gaussian fields [Owh15], which further translates to game theory [OS19]. All these methods are successful, but they all implicitly rely on properties of the underlying elliptic equation. Can we devise an approach that applies to general problems with oscillatory media that exploits the low-rank property in the solution space, without using analytical structure explicitly?

3. A Unified Framework for Multiscale PDEs Based on Random Sampling

We have given several examples of multiscale models that arise in applications, and mentioned several algorithmic approaches that make use of the limiting equations, when available. We describe next the foundations of a unified scheme that captures asymptotic limiting behavior automatically, even when the asymptotic limits are unavailable. Our method exploits low-rank structure and uses random sampling to discover this structure. We describe the low-rank property in Section 3.1 and the randomized SVD method for revealing this structure in Section 3.2.

3.1. Numerical rank. We consider a bounded linear operator \mathcal{A} , which maps $f \in \mathcal{X}$ to a space \mathcal{Y} , that is

$$\begin{array}{cccc} \mathcal{A} : & \mathcal{X} & \rightarrow & \mathcal{Y} \\ & f & \mapsto & u. \end{array}$$

In the PDE setting, \mathcal{A} is the solution operator that maps the boundary conditions and/or source term f to the solution u. The numerical rank of such an operator is defined as follows.

Definition 1 (Numerical rank). The numerical τ -rank of \mathcal{A} is the rank of the lowest-rank operator within the τ -neighborhood of \mathcal{A} , that is,

$$k_{\tau}(\mathcal{A}) := \min \{ \dim \operatorname{ran} \tilde{\mathcal{A}} : \tilde{\mathcal{A}} \in \mathcal{L}(\mathcal{X}, \mathcal{Y}), ||\tilde{\mathcal{A}} - \mathcal{A}|| \leq \tau \}.$$

In other words, $k_{\tau}(\mathcal{A})$ is this smallest dimension of the range among all the operators within distance τ of \mathcal{A} .

When \mathcal{A} is the PDE solution map, then $\tilde{\mathcal{A}}$ with low rank is also a linear map with a finite dimensionality. It can be viewed as the discrete version (or a matrix of dimension $k_{\tau}(\mathcal{A})$) that approximates \mathcal{A} within τ accuracy. The definition suggests that if $\tilde{\mathcal{A}}$ can be found, it is optimal in the sense of numerical efficiency. The concept is rather similar to the Kolmogorov N-width, defined as follows.

Definition 2 (Kolmogorov *N*-width). Given the linear operator $\mathcal{A}: \mathcal{X} \to \mathcal{Y}$, the Kolmogorov *N*-width $d_N(\mathcal{A})$ is the shortest distance from its range to all *N*-dimensional space, that is,

$$d_{N}(\mathcal{A}) := \min_{S: \dim S = N} d(\mathcal{A}, S)$$

$$= \min_{S: \dim S = N} \sup_{f} \min_{v \in S} \frac{\|\mathcal{A}f - v\|_{\mathcal{Y}}}{\|f\|_{\mathcal{X}}}.$$
(15)

Indeed, the Kolmogorov *N*-width and numerical rank are related by the following result [CLLW20a].

Proposition 1. For any linear operator $\mathcal{A}: \mathcal{X} \to \mathcal{Y}$, we have the following.

- (a) If the numerical τ -rank is N, then $d_N(\mathcal{A}) \leq \tau$.
- (b) If $d_N(\mathcal{A}) \leq \tau < d_{N-1}(\mathcal{A})$, then the numerical τ -rank is N.

For the three examples presented in Section 2, the numerical ranks can be calculated from their limiting equations. For one-dimensional RTE in the diffusion regime, if we denote by \mathcal{A}^{ϵ} and \mathcal{A}^{*} the solution operators of (4) and (8), respectively, then noting that A^* can be approximated using $1/\sqrt{\tau}$ grid points to achieve τ accuracy, when $\epsilon < \tau$, the numerical rank is naturally $k_{\tau}(\mathcal{A}^{\epsilon}) \lesssim 1/\sqrt{\tau - \epsilon}$. Without employing the knowledge of the existence of the limit, however, a brute-force discretization naturally requires $O(1/\epsilon \tau^{\alpha+1})$ degrees of freedom: $O(1/\epsilon \tau)$ for the upwind discretization in *x* and $O(1/\tau^{\alpha})$ for the discretization in v, where α depends on the particular numerical integral accuracy. Translating into Green's-matrix language, this observation means that G^{ϵ} is represented by $O(1/\epsilon \tau^{\alpha+1})$ degrees of freedom but its range can be captured by a compressed Green's matrix G^* with just O(1) column vectors.

The same argument applies to the elliptic equation on a two-dimension domain with high oscillations. When second-order linear finite elements are used, with no knowledge of the limiting system, $O(1/\epsilon^2\tau)$ degrees of freedom are required, dropping to $O(1/\tau)$ when the the limiting system is known. In other words, the full Green's matrix G^{ϵ} requiring $O(1/\epsilon^2)$ degrees of freedom can be well-represented using just O(1) column vectors.

In all these cases, the degrees of freedom for a given numerical method are substantially larger than the numerical rank of the problem. Thus, much of the information in these full-blown representations is redundant and compressible. A low-rank representation exists and yields a much more economical representation.

3.2. Random sampling in numerical linear algebra. Knowing the existence of the low-rank structure and finding such a structure are very different goals. The Kolmogorov *N*-width is a concept developed in numerical PDEs, but it has made little impact in numerical PDEs for a simple reason: Traditional PDE solvers require a

predetermined set of basis functions, while the Kolmogorov N-width looks for "optimal" basis functions. How can an optimal basis be found without first forming the full basis? Translated to linear algebra, this question is about finding the dominant singular vectors in a matrix without forming the whole matrix. Specifically, if $A \in \mathbb{R}^{m \times n}$ is known to be approximately low rank, meaning that there exists U_r , a $m \times r$ matrix with orthonormal columns with $r \ll \min(m, n)$ and

$$||A - A_r|| = ||A - U_r U_r^T A|| \ll ||A||,$$

can we find U_r without forming the full matrix A?

In linear algebra, it is well-known that U_r is simply the collection of the first r singular vectors of A. Writing

$$A = U\Sigma V^{\top} = \sum_{i=1}^{n} \sigma_i u_i v_i^{\top}, \qquad (16)$$

where $U = [u_1, u_2, ..., u_n] \in \mathbb{R}^{m \times n}$ and $V = [v_1, v_2, ..., v_n] \in \mathbb{R}^{n \times n}$ contain the left/right singular vectors and $\Sigma = \operatorname{diag}(\sigma_1, \sigma_2, ..., \sigma_n)$ contains the singular values, then U_r is the first r columns in U.

The standard method for computing the SVD requires A to be stored and computed with. But the celebrated randomized SVD (rSVD) method [HMT11] captures the range of a given matrix by means of random sampling of its column space, which requires only computation of matrix-vector products involving A and random vectors — operations that can be performed without full storage or knowledge of A. Implementation of rSVD is easy and its performance is robust.

The idea behind the algorithm is simple. If matrix $A \in \mathbb{R}^{m \times n}$ has approximate low rank $r \ll \min\{m, n\}$, the matrix maps an n-dimensional sphere to an m-dimensional ellipsoid that is "thin:" r of its axes are significantly larger than the rest. With high probability, vectors that are randomly sampled on the n-dimensional sphere are mapped to vectors that lie mostly in a r-dimensional subspace of \mathbb{R}^m — the range of A. An approximation to A_r can be obtained by projecting onto this subspace.

A precise statement of the performance of randomized SVD is as follows [HMT11].

Theorem 4. Let A be an $m \times n$ matrix. Define

$$Y = A\Omega, \qquad (17)$$

where $\Omega = [\omega_1, ..., \omega_{r+p}]$ is a matrix of size $n \times (r+p)$ with its entries randomly drawn from an i.i.d. normal distribution, where p is an oversampling parameter. If $\sigma_{r+1} \ll \sigma_1 = O(1)$, then the projection of A onto the space spanned by Y, defined by

$$P_{V}(A) = Y(Y^{T}Y)^{-1}Y^{T}A,$$

yields that $\|A-P_Y(A)\|\ll \sigma_1$ with high probability, and

$$\mathbb{E} \left\| \mathbf{A} - \mathbf{P}_{\mathsf{Y}}(\mathbf{A}) \right\| \lesssim \frac{r}{p-1} \sigma_{r+1} \ll \sigma_1.$$

The result reconstructs the range of A in a nearly optimal way. It is optimal in efficiency because to capture a rank-r matrix, only r+p matrix-vector products involving A are required for the calculation of Y, and the oversampling parameter p is typically quite modest. (p=5 is a typical value.) The result is nearly optimal in accuracy as well. The error bound relies only on σ_{r+1} , which is expected to be smaller than σ_1 . The decay profile of singular values do not affect the approximation accuracy.

If a low-rank approximation to the matrix A is required, and not just an approximation of its range, another step involving multiplications with its transpose is needed. The full method is shown in Algorithm 1.

Algorithm 1 Randomized SVD.

- 1: Given an $m \times n$ matrix A, target rank r and oversampling parameter p;
- 2: Set k = r + p;
- 3: Stage A:
- 4: Generate an $n \times k$ Gaussian test matrix Ω :
- 5: Form $Y = A\Omega$;
- 6: Perform the QR-decomposition of Y: Y = QR
- 7: Stage B:
- 8: Form $B = A^TQ$;
- 9: Compute the SVD of the $k \times n$ matrix $B^{T} = \widetilde{U}\Sigma V^{T}$;
- 10: Set $U = Q\widetilde{U}$;
- 11: **Return:** U, Σ, V .

4. Random Sampling for Multiscale Computation

Here we describe how rSVD can be incorporated into multiscale PDE solvers to exploit the low-rank structure of these equations. Our procedure is composed of both offline and online stages. Low-rank structure is learned in the offline stage, while in the online stage, the solution for the given source / boundary term f in (1) is extracted.

We consider in particular the following boundary value problem:

$$\begin{cases} (\mathcal{L}^{\varepsilon}u^{\varepsilon})(x) = 0, & x \in \mathcal{K}, \\ \mathcal{B}u(x) = \phi(x), & x \in \partial \mathcal{K}, \end{cases}$$
 (18)

where \mathcal{B} is the boundary condition operator, $\partial \mathcal{K}$ the boundary associated with domain \mathcal{K} , and we now denote the source term (the boundary data) by ϕ . Our fundamental goal is to construct the low-rank approximation to the Green's operator G^{ε} for (18). With this operator in hand, the solution can be computed for any value of the boundary conditions ϕ at a relatively small incremental cost.

If we apply rSVD to approximate G^{ε} directly, we need to compute products of this operator with random vectors. For problem (18), this operation corresponds to solving

the problem with ϕ replaced by random boundary conditions. Solving even one such problem efficiently is a computationally challenging task. We use the domain decomposition framework.

We start by partitioning the domain \mathcal{K} into subdomains as follows:

$$\mathcal{K} = \bigcup_{m=1}^{M} \mathcal{K}_m, \tag{19}$$

where the patches \mathcal{K}_m overlap, in general. We denote by $\partial \mathcal{K}_m$ the boundary associated with \mathcal{K}_m . Furthermore, we identify the subregions that intersect with \mathcal{K}_m as follows:

$$\mathcal{I}_m = \{ n \in \mathbb{N} : 1 \le n \le M, \mathcal{K}_m \cap \mathcal{K}_n \ne \emptyset \},\$$

and define the interior of the patch to be

$$\widetilde{\mathcal{K}}_m = \mathcal{K}_m \backslash \left(\bigcup_{n \in \mathcal{I}_m} \mathcal{K}_n \right).$$

For this particular partition of the domain, we define the partition-of-unity functions χ_m , m = 1, 2, ..., M to have the following properties:

$$\sum_{m=1}^{M} \chi_m(x) = 1, \quad \forall x \in \mathcal{K},$$
with
$$\begin{cases} 0 \le \chi_m(x) \le 1, & x \in \mathcal{K}_m \\ \chi_m(x) = 0, & x \in \mathcal{K} \setminus \mathcal{K}_m. \end{cases}$$
(20)

We choose a discretization that resolves the small scales in the solution, defining a mesh width $h \ll \varepsilon$. (The number of subdomains M is independent of ε .) A typical decomposition is illustrated in Figure 2.

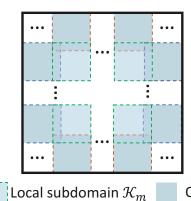


Figure 2. Illustration of domain decomposition of a rectangular domain \mathcal{K} in 2D with overlap.

How do we design the offline stage to "learn" the lowrank approximation? We propose two approaches that lead to two different kinds of algorithms. In the first approach we learn the optimal basis functions within each subdomain, while the second algorithm employs Schwarz iteration, preparing the boundary-to-boundary map in the offline stage. Other PDE solvers that utilize randomness can also be found in [CEGL16, BS18, Mar20]. In particular in [CEGL16] the authors studied, specifically for elliptic type equations, the generalized eigenvalue problem of the stiffness and mass matrices, and give an error bound using the largest eigenvalue obtained offline.

4.1. **Learning basis functions**. In standard domain decomposition, the local discretized Green's matrix G_m is assembled from a "full" collection of basis functions in the patch \mathcal{K}_m . The global solution to (18), confined to each \mathcal{K}_m , is a linear combination of the columns G_m . The coefficients of these combinations are chosen so that that the continuity conditions across patches and the exterior boundary condition are all satisfied. The complete process can be outlined as follows.

1. Offline stage: For m = 1, 2, ..., M, find

$$G_m = [b_{m,1}, b_{m,2}...],$$

where each local function $b_{m,n}$ is a solution to (18) restricted to the subdomain \mathcal{K}_m , with fine grid $h \ll \varepsilon$ and delta-function boundary conditions. That is,

$$\begin{cases} \mathcal{L}^{\epsilon} b_{m,n} = 0, & x \in \mathcal{K}_m \\ b_{m,n} = \delta_{m,n}, & x \in \partial \mathcal{K}_m, \end{cases}$$
 (21)

where $\delta_{m,n}$ is the Kronecker delta that singles out the n-th grid point on the boundary $\partial \mathcal{K}_m$.

2. Online stage: The global solution is

$$u = \sum_{m=1}^{M} u_m \chi_m = \sum_{m=1}^{M} \chi_m \mathsf{G}_m c_m,$$

with the support of each $u_m = G_m c_m$ confined to \mathcal{K}_m , where c_m is a vector of coefficients determined by the boundary conditions ϕ and continuity conditions across the patches.

The complete basis represented by G_m has a low-rank structure that can be revealed using randomized SVD. Instead of using delta functions as the boundary conditions, we propose to obtain basis functions by setting random values on $\partial \mathcal{K}_m$, as follows:

$$\begin{cases} \mathcal{L}^{\epsilon} r_{m,n} = 0, & x \in \mathcal{K}_m, \\ r_{m,n} = \omega_{m,n}, & x \in \partial \mathcal{K}_m, \end{cases}$$
 (22)

where $\omega_{m,n}$ is defined to have a random value drawn i.i.d. from a normal distribution at each grid point in $\partial \mathcal{K}_m$. Denoting $\mathsf{G}_m^{\mathrm{r}} = \{r_{m,1}, r_{m,2}, \cdots\}$, we have from linearity of the equation that

$$G_m^r = G_m \Omega$$
,

where Ω is a random i.i.d. matrix with entries $\omega_{m,n}$. This $\mathsf{G}_m^{\mathrm{r}}$ is used in the online stage, as an accurate surrogate of G_m , see Algorithm 2.

Although we do not apply full-blown rSVD here, the homogenizable and low-rank property of the local solution

Algorithm 2 A general framework for multiscale PDE $\mathcal{L}^{\varepsilon}u^{\varepsilon} = 0$ over \mathcal{K} with $\mathcal{B}u^{\varepsilon} = f$ on $\partial \mathcal{K}$.

1: Domain Decomposition

2: Partition domain according to (19).

3: Offline Stage:

- 4: Prepare i.i.d. Gaussian vectors $\omega_{m,i}$, $i = 1, ..., k_m$ on each $\partial \mathcal{K}_m$.
- 5: Solve the basis function $r_{m,i}$ in (22) on each Ω_m , and collect the local basis in G_m .

6: Online Stage:

7: Use continuity condition and global boundary data ϕ to determine coefficient vectors c_1 , c_2 , ..., c_M , and set

$$u = \sum_{m=1}^{M} \chi_m \mathsf{G}_m c_m \tag{23}$$

8: **Return:** approximate global solution u.

space implies that G_m^r and G_m share similar range with the number of basis functions k_m in G_m^r being much smaller than n_m , the number of grid points on $\partial \mathcal{K}_m$, and independent of ε . In Figure 3 we plot the angles between G_m^r and G_m for two of the equations discussed in Section 2. In both cases, and for small ε , the approximated Green's matrix quickly recovers the true Green's matrix as the number of samples k_m increases, and thus captures the local solution space. We should note that if G_m does not have low-rank structure, in the sense that $k_m \sim n_m$, then solving (22) would be equally expensive as solving (21), hence the random sampling technique does not gain any computational efficiency when the system is not homogenizable.

In Figure 4 we showcase the basis functions on a patch for elliptic equation with media coefficient $a(x_1, x_2) = 1 + 1000 \, \mathbf{1}_S(x_1, x_2)$, with $S = \{(x_1, x_2) \in [0, 1]^2 : (x_1 \cos(100\sqrt{(x_1 - 0.5)^2 + (x_2 - 0.5)^2})) \le x_2 - 0.5\}$. For this non-conventional media without any periodic structure, the traditional multiscale methods are no longer valid, but our method still quickly captures the optimal basis.

For particular boundary conditions ϕ , the global solution is assembled from the local basis functions in the online stage. Two numerical examples are shown in Figure 5. In both examples, there is little visible difference between the reference solution and the approximated one computed from the reduced random basis. Only 8.3% and 62.5%, respectively, of the degrees of freedom required by the full basis are needed to represent these solutions using a random basis; see details in [CLLW20a].

Since we do not have access to the full set of basis functions, the condition that u defined in (23) is continuous across subdomains can be satisfied only in a least-squares sense; see [CLLW20a] for details.

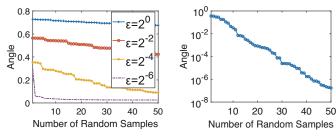


Figure 3. Angle between the true Green's matrix G_m and the approximate version G_m^r , confined on the interior of a patch $\widetilde{\mathcal{K}}_m$ for some m, as the number of random samples increases. Left plot: Angle for 1D RTE (4) with diffusive kernel (5) and various values of ϵ . Right plot: Angle for elliptic equation (12) on a rectangular local patch with $\epsilon = 2^{-4}$. Data from [CLLW20a].

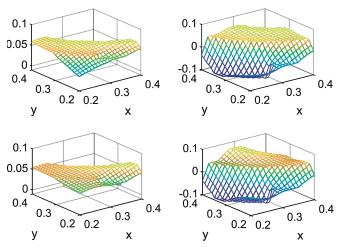


Figure 4. Optimal basis functions and their projections onto the approximate spaces. First row: First two singular vectors of the Green's matrix on a local patch. Second row: Projections onto the space spanned by 6 random sampled basis functions. Note that the small random sample captures well the leading eigenvectors of the true Green's operator. Data from [CLLW20a].

4.2. A low-rank Schwarz method. Our second approach for exploiting the low-rank property in multiscale computations is based on Schwarz iteration. The Schwarz method is a standard iteration algorithm within the domain decomposition framework, in which boundary-value problems are solved on the patches, with neighboring patches subsequently exchanging information and re-solving until consistency is attained. The exchange of boundary information between neighboring patches is known as the boundary-to-boundary (BtB) map. The map has an exploitable low-rank property.

To develop the approach, we write the solution of (18) as

$$u^{\epsilon}(x) = \sum_{m=1}^{N} \chi_m(x) u_m^{\epsilon}(x), \qquad (24)$$

where the partition-of-unity functions χ_m are defined

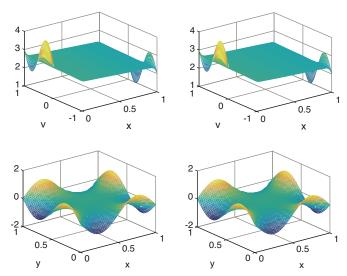


Figure 5. First row: Solutions for 1D RTE (4) with Henyey-Greenstein kernel in (5) with $\epsilon = 2^{-6}$. Second row: Solutions for the elliptic equation (12) with Dirichlet boundary condition and highly oscillatory medium $a(x, x/\epsilon) =$ $2 + \sin(2\pi x_1)\cos(2\pi x_2) + \frac{2 + 1.8\sin(2\pi x_1/\epsilon)}{2 + 1.8\cos(2\pi x_2/\epsilon)} + \frac{2 + \sin(2\pi x_2/\epsilon)}{2 + 1.8\cos(2\pi x_1/\epsilon)}$ $\varepsilon=2^{-4}.$ The left column shows reference solutions while the right column is obtained from randomized reduced bases. Data from [CLLW20a].

in (20). The solution u_m^{ϵ} on patch m is uniquely determined by $f_{m'}$ its local boundary condition, according to the equation

$$\begin{cases} \mathcal{L}^{\varepsilon} u_{m}^{\varepsilon} = 0, & x \in \mathcal{K}_{m} \\ u_{m}^{\varepsilon}(x) = f_{m}(x), & x \in \partial \mathcal{K}_{m}. \end{cases}$$
 (25)

The Schwarz method starts with initial guesses to the local boundary conditions $f_m = f_m^0$, then on iteration t, it solves the subproblems (25) with $f_m = f_m^t$ to obtain all local solutions u_m^{ϵ} . By confining u_m^{ϵ} on the boundaries of adjacent patches, one updates the boundary conditions for surrounding patches:

$$f_m^t(x) \xrightarrow{\mathcal{S}_m} u_m^{\epsilon}(x)|_{\widetilde{\mathcal{K}}_m} \xrightarrow{\mathcal{P}_m} f_n^{t+1}(x), \forall n \in \mathcal{I}_m.$$
 (26)

Here S_m denote the solution to (25) confined in the interior of $\mathcal{K}_{m'}$ and \mathcal{P}_m takes the trace of the solution on the neighboring boundaries $\mathcal{K}_m \cap \partial \mathcal{K}_n$ for $n \in \mathcal{I}_{m'}$ for the updated boundary condition.

Define the BtB map by $\mathcal{A}_m := \mathcal{P}_m \circ \mathcal{S}_m$, and define \mathcal{A} and f^t to be the aggregation of \mathcal{A}_m and f_m^t , respectively, over i = 1, 2, ..., M. We can then write the updating procedure

$$f^{t+1} = \mathcal{A} f^t \,.$$

The overall method is summarized in Algorithm 3.

Most of the computation in the Schwarz method during the iteration comes from solving the boundary-value PDEs on the patches, to implement the map A. Since the PDE is homogenizable, the solution space on each patch is

Algorithm 3 Schwarz method for multiscale PDE $\mathcal{L}^{\epsilon}u^{\epsilon} = 0$ over \mathcal{K} with $\mathcal{B}u^{\epsilon} = f$ on $\partial \mathcal{K}$.

```
1: Given total iterations T;
```

Domain Decomposition

Partition domain according to (19).

```
Schwarz Iteration:
```

11:

```
Initialize f_m^0 for each \partial \mathcal{K}_m and set t=0.
 5:
          While |f_n^t - f_n^{t-1}| > \text{TOL}
 6:
               Solve (25) for u_m^t using f_m^t for each m.
 7:
               Update f_n^{t+1} = u_m^t on \mathcal{K}_m \cap \partial \mathcal{K}_n, n \in \mathcal{I}_m.
 8:
 9:
10:
          Solve (25) for u_m^t using f_m^t for each m.
```

Assemble global solution $u = \sum_{m=1}^{N} \chi_m u_m^t$. 12:

13: **Return:** approximated global solution u^T .

approximately low rank, and the map A can be expected to inherit this property. If we can "learn" this operator in an offline stage, and simply apply a low-rank approximation repeatedly in the online stage, the online part of Algorithm 3 can be made much more efficient. In our approach, Algorithm 1 is used to compress the map A.

This approach is quite different from the one described in Section 4.1, in the sense that it is not only the range of the solution space, but the whole operator that is being approximated. To apply Algorithm 1, we need to define the "adjoint operator" for \mathcal{A} on the PDE level. This operator is composed of the adjoints S_m^* for the local solution operators \mathcal{S}_m of (25) on each domain \mathcal{K}_m . The form of \mathcal{S}_m^* is specific to the PDE; we use the elliptic equation as an example. Defining $\mathcal{L}^{\epsilon} = \nabla \cdot (a(x, x/\epsilon)\nabla)$, \mathcal{S}_{m}^{*} is defined in the following result.

Theorem 5. Let S_m be the confined solution operator for the elliptic equation with Dirichlet boundary condition on patch \mathcal{K}_m . Given any function g supported on $\widetilde{\mathcal{K}}_m$, the adjoint operator S_m^* acting on g is given by:

$$\mathcal{S}_{m}^{*}g = a \frac{\partial h}{\partial n} \Big|_{\partial \mathcal{K}_{m}}, \tag{27}$$

where $\frac{\partial h}{\partial n}$ is the outer normal derivative on $\partial \mathcal{K}_m$ and h solves the following sourced elliptic equation:

$$\begin{cases} \nabla \cdot \left(a\left(x, \frac{x}{\epsilon} \right) \nabla h(x) \right) = g, & x \in \mathcal{K}_m \\ h(x) = 0, & x \in \partial \mathcal{K}_m. \end{cases}$$
 (28)

We also describe calculation of the adjoint operator S_m^* for the RTE (4).

Theorem 6. Let S_m be the confined solution operator for RTE (4) and the conditions in Theorem 1 hold. Given any function g supported on $\widetilde{\mathcal{K}}_m imes \mathbb{V}$, the adjoint operator \mathcal{S}_m^* is

defined as follows:

$$S_m^* g(x, v) = h(x, v), \quad x \in \partial \mathcal{K}_m, \quad v \cdot n_x < 0, \quad (29)$$

where h solves the adjoint RTE over \mathcal{K}_m , which is

$$-v \cdot \nabla_x h - \mathcal{S}^{\epsilon}[h] = g(x, v), \quad (x, v) \in \mathcal{K}_m \times \mathbb{V}, \quad (30)$$

with outgoing boundary condition h(x,v) = 0 on $x \in \partial \mathcal{K}_m$ and $v \cdot n_x > 0$.

The specific form of the adjoint operator S_m^* allows us to adapt the randomized SVD algorithm to compress the confined solution map S_m ; see Algorithm 4. This method requires only k solves of local PDE (25) and sourced adjoint PDE (28) (or (30)), together with a QR factorization and SVD of relatively small matrices. The overall low-rank Schwarz iteration is then summarized in Algorithm 5.

Algorithm 4 Randomized SVD for S_m .

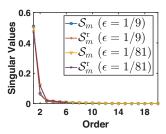
- 1: Given target rank *r* and oversampling parameter *p*;
- 2: Set k = r + p;
- 3: Stage A:
- Generate k random boundary conditions ξ_i on 4:
- Solve (25) using ξ_i as boundary conditions and restrict the solution over $\widetilde{\mathcal{K}}_m$ to obtain u_i^{ε} .
- Find orthonormal basis $Q = [q_1, ..., q_k]$ of $\widetilde{U} =$ $[u_1, ..., u_k].$
- Construct zero extension of q_k over \mathcal{K}_m , denoted by \tilde{q}_k .
- 9: Solve (28) or (30) for h_i using \tilde{q}_k as source.
- Compute b_i using h_i by flux (27) or restriction on 10: incoming boundary (29).
- Assemble all fluxes $B = [b_1, ..., b_k]$. 11:
- Compute SVD of $B^* = \widetilde{\mathsf{U}}_k \Sigma_k \mathsf{V}_{\nu}^*$. 12:
- Compute $U_k = Q\widetilde{U}_k$. 13:
- 14: **Return** U_k , Σ_k , V_k .

In Figure 6 we present the confined solution operator S_m and its low-rank approximation S_m^r . For the radiative transfer equation, the map is a 2880-by-40 matrix, with the size of each patch being $0.2 \times [-1,1]$, with $\Delta x = 1/360$ and $\Delta v = 1/40$. The random sampling procedure reconstructs it with just 6 samples. For the elliptic equation, the map is a 1600-by-160 matrix, and the size of each patch is $1 \times [0,1]$ with $\Delta x = 1/40$. The random sampling approximates it well with 60 samples. The compression rates for these examples are thus 6.7 and 2.7, respectively. See details in [CLLW21].

In Figure 7, we show numerical examples for the global solutions of two problems obtained from the approach of this section. The reference solution (obtained with a fine mesh) is well captured by the approximation that uses the low-rank BtB map as the surrogate in the Schwarz iteration. These two cases use just 15% and 43%, respectively, of the number of local solves needed to capture the BtB map at fine scale. While the relative error of the reduced Schwarz method decays as fast as the standard Schwarz iteration, as shown in Figure 8, the cost is much reduced. See Table 1 for a comparison of computation times.

Algorithm 5 Reduced Schwarz method for multiscale PDE $\mathcal{L}^{\epsilon}u^{\epsilon} = 0$ over \mathcal{K} with $\mathcal{B}u^{\epsilon} = \phi$ on $\partial \mathcal{K}$.

- 1: Given rank k, total iterations T.
- 2: Domain Decomposition
- Partition domain according to (19).
- 4: Offline Stage:
- For all m, use Algorithm 4 to find the rank-k approximation to S_m , denoted by $\bigcup_k^m \sum_k^m \bigvee_k^{m,*}$.
- 6: Online:
- Initiate $f_m^0(x)$ for each $\partial \mathcal{K}_m$, and set t = 0. 7:
- While $|f_n^t f_n^{t-1}| > \text{TOL}$
- 9:
- Evaluate $u_m^t = \bigcup_k^m \sum_k^m \bigvee_k^{m,*} f_m^t$ for each m. Update $f_n^{t+1} = u_m^t$ on $\mathcal{K}_m \cap \partial \mathcal{K}_n$, $n \in \mathcal{I}_m$. 10:
- 11:
- 12:
- Solve (25) for u_m^t using f_m^t for each m. 13:
- Assemble global solution $u = \sum_{n=1}^{N} \chi_m u_m^t$. 14:
- 15: Return u(x).



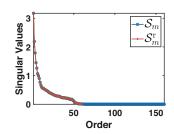


Figure 6. The singular decay of the restricted local solution operator S_m and its low-rank approximation $S_m^{\rm r}$ for the RTE (left) and the elliptic equation (right). In RTE (4) we use heterogeneous collision kernel (9)

 $k(x/\epsilon,v,v')=\sigma^\epsilon(x)=rac{1}{81}rac{1.1+\cos(4\pi x)}{1.1+\sin(2\pi x/\epsilon)}$, and in elliptic equation we use $a(x, x/\epsilon) = \frac{2+1.8 \sin(\pi x_1/\epsilon)}{2+1.8 \cos(\pi x_2/\epsilon)} + \frac{2+\sin(\pi x_2/\epsilon)}{2+1.8 \sin(\pi x_1)}$ with $\epsilon = 2^{-4}$. Data from [CLLW21, CLW20].

5. Manifold Learning and Nonlinear Multiscale **Problems**

It is not straightforward to extend the techniques of the previous section to nonlinear PDEs. Despite low-rank properties still holding due to the existence of the limiting

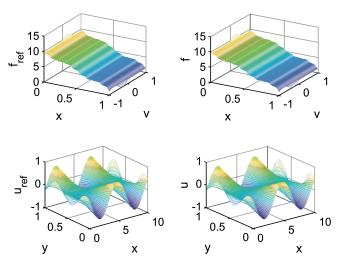


Figure 7. The comparison between the reference solutions (left column) and the approximation using reduced Schwarz method (right column). The first and second rows are for the RTE and the elliptic equation, respectively. In RTE (4) we use heterogeneous collision kernel (9)

 $k(x/\epsilon,v,v')=\sigma^\epsilon(x)=\frac{1}{\epsilon_1}\frac{1.1+\cos(4\pi x)}{1.1+\sin(2\pi x/\epsilon_2)}, \text{ with } (\epsilon_1,\epsilon_2)=(1/81,1/9),$ and in elliptic equation we use $a(x,x/\epsilon)=\frac{2+1.8\sin(\pi x_1/\epsilon)}{2+1.8\cos(\pi x_2/\epsilon)}+\frac{2+\sin(\pi x_2/\epsilon)}{2+1.8\sin(\pi x_1)} \text{ with } \epsilon=2^{-4}. \text{ Data from [CLLW21, CLW20]}.$

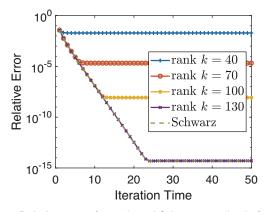


Figure 8. Relative error for reduced Schwarz methods for various ranks k for elliptic equation with oscillatory medium that is the same as Figure 7. Data from [CLW20].

equation, the argument based on compressing the Green's matrix no longer holds. The collection of solutions for different source / boundary terms is not longer a linear subspace, but a solution manifold.

We consider the general nonlinear multiscale problem in the following form

$$\mathcal{N}^{\epsilon}u^{\epsilon} = f, \tag{31}$$

where $\mathcal{N}^{\varepsilon}$ is a nonlinear differential operator that depends explicitly on the small parameter ε . The term f can be the source term, boundary conditions or initial conditions.

RTE	offline (s)	online (s)
Rank = 5	227.99	0.0029
Rank = 6	268.46	0.0162
Full rank	706.99	0.0148
Schwarz		1027.40
elliptic	offline (s)	online (s)
Rank = 40	49.7	0.049
Rank = 70	87.3	0.061
Schwarz		31.4

Table 1. Run time comparison between vanilla Schwarz method and the reduced Schwarz method. k=5,6 for the RTE and k=40,70 for the elliptic equation with Dirichlet boundary condition. The configuration of the media is the same as those in Figure 7. Data from [CLLW21, CLW20].

Assume further that the equation has an asymptotic limit

$$\mathcal{N}^* u^* = f \tag{32}$$

as $\epsilon \to 0$, that is, $\|u^{\epsilon} - u^*\| \to 0$ as $\epsilon \to 0$. The argument for the linear problem is still applicable: The degrees of freedom required by the classical numerical method for solving (31) grows rapidly as $\epsilon \to 0$, while the existence of the homogenized equation (32) indicates that only O(1) degrees of freedom should be needed to resolve macroscale features. From a manifold perspective, the solutions to (31) vary in a high-dimensional space as f changes, but this manifold is approximated to within distance $O(\epsilon)$ by another manifold whose dimension is O(1).

Suppose a manifold in a high-dimensional space is approximately low-dimensional, can we quickly learn it without paying the high-dimensional cost? We turn to manifold learning for answers to this question. We are particularly interested in adopting the ideas from the local linear embedding and multi-scale SVD approaches that learn the manifold from observed point clouds, and interpolate the local solution manifold using multiple tangent-space patches, see references in [CLLW20b].

We denote the nonlinear solution map of (31) by \mathcal{S}^{ϵ} : $f \in \mathcal{X} \to u^{\epsilon} \in \mathcal{Y}$, which maps the source term or initial / boundary conditions f(x) to the solution of the equation. In the offline stage, we randomly sample a large number of configurations f_i in X, and compute the associated solutions $u_i^{\varepsilon} = \mathcal{S}^{\varepsilon} f_i \in \mathcal{Y}$ on fine grids. These solutions form a point cloud in a high-dimensional space \mathcal{Y} . The $\{f_i\}$'s are then subdivided into a number of small neighborhoods, and we construct tangential approximations to the mapping S^{ϵ} on each neighborhood. In the online stage, given a new configuration f, we identify the small neighborhood to which it belongs and find the corresponding solution by performing linear interpolation. The overall offlineonline strategy in summarized in Algorithm 6. We stress that some modifications are needed to reduce the cost of implementation. For example, the algorithm should be

combined with domain decomposition (for example, Schwarz iteration) to further confine the computation to local domains, to save computational cost.

In Figure 9 we plot the local low-dimensional solution manifold for a nonlinear RTE (specifically, a linear RTE nonlinearly coupled with a temperature term). The solution manifold appears to have a local two-dimensional structure; the point clouds lie near on a two-dimensional plane. We refer to [CLLW20b] for more details of the implementation and numerical results.

Algorithm 6 Manifold learning algorithm for solving $\mathcal{N}^{\epsilon}u^{\epsilon} = f$.

- 1: Offline
- 2: Randomly sample $f_i(x)$, i = 1, ..., N, and find solutions $u_i^{\epsilon} = S^{\epsilon} f_i$.
- 3: **Online:** Given f(x):
- 4: **Step 1**: Identify the *k*-nearest neighbors of f(x), call them f_{i_j} , j = 1, 2, ..., k, with f_{i_1} being the nearest neighbor;
- 5: **Step 2**: Compute $\mathcal{S}^{\epsilon} \phi \approx u_{i_1}^{\epsilon} + \mathsf{U} \cdot c$ with

$$\mathsf{U} = \left[\begin{matrix} u_{i_2}^\varepsilon - u_{i_1}^\varepsilon & \dots & u_{i_k}^\varepsilon - u_{i_1}^\varepsilon \\ | & | & \end{matrix} \right],$$

where c is a set of coefficient that fits $f - f_{i_1}$ with a linear combination of $f_{i_j} - f_{i_1}$, for j = 2, 3, ..., k.

6: Return $u^{\epsilon} = \mathcal{S}^{\epsilon} \phi$.

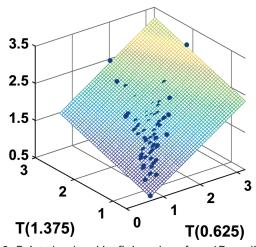


Figure 9. Point cloud and its fitting plane for a 1D nonlinear RTE with Knudsen number $\epsilon=2^{-6}$, when confined in a small patch containing an interval $[0.625\,,1.375]$. The solution profile is approximately determined by the temperature T at two grid points x=0.625 and x=1.375. The z-axis shows the value of $u^\epsilon(x=1)$. The dependence is clearly linear and two-dimensional. Data from [CLLW20b].

6. Looking Forward

We have seen a vast literature addressing all aspects of the computation of multiscale problems. Over the years, the research has been drifting gradually away from its origin, where solvers were influenced by analytical understanding, specifically of the limiting behavior of the specific PDE. Machine learning algorithms have shown more and more power in sketching the solution profile with a much reduced numerical cost. In particular, the existence of the homogenized limit suggests there are low-rank features in the discrete system, and that random linear algebra techniques and manifold learning methods, when utilized properly, can identify these features for a compressed representation of the PDE solutions.

We have reviewed two methods, both of which make use of the domain decomposition framework. They compress either basis functions or the boundary-to-boundary map in an offline learning stage. This review article serves as a showcase of the power of random solvers in numerical PDEs. For time-dependent problems, and homogenization problems that have weak-convergence (instead of strong) such as quantum systems in the semi-classical regime, further development of the approaches is needed. Incorporation of time and weak-limit in the algorithm-design lies at the core of future challenges.

ACKNOWLEDGMENT. The authors acknowledge generous support from NSF, ONR, DOE, and AFOSR. Due to restrictions on the allowed number of references, many important contributions are omitted. The reader is invited to consult the bibliographies in the cited references for a more extensive view of the literature.

References

[All92] Grégoire Allaire, Homogenization and two-scale convergence, SIAM J. Math. Anal. 23 (1992), no. 6, 1482–1518, DOI 10.1137/0523084. MR1185639

[BL11] Ivo Babuska and Robert Lipton, *Optimal local approximation spaces for generalized finite element methods with application to multiscale problems*, Multiscale Model. Simul. 9 (2011), no. 1, 373–406, DOI 10.1137/100791051 MR2801210

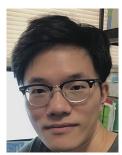
[BLP11] A. Bensoussan, J.-L. Lions, and G. Papanicolaou, *Asymptotic analysis for periodic structures*, AMS Chelsea Publishing, Providence, RI, 2011. Corrected reprint of the 1978 original [MR0503330], DOI 10.1090/chel/374. MR2839402

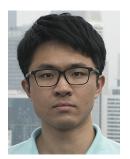
[BS18] Andreas Buhr and Kathrin Smetana, *Randomized local model order reduction*, SIAM J. Sci. Comput. **40** (2018), no. 4, A2120–A2151, DOI 10.1137/17M1138480 MR3824169

[CEGL16] Victor M. Calo, Yalchin Efendiev, Juan Galvis, and Guanglian Li, Randomized oversampling for generalized multiscale finite element methods, Multiscale Model. Simul. 14 (2016), no. 1, 482–501, DOI 10.1137/140988826. MR3477310

- [CLLW20a] Ke Chen, Qin Li, Jianfeng Lu, and Stephen J. Wright, Random sampling and efficient algorithms for multiscale PDEs, SIAM J. Sci. Comput. 42 (2020), no. 5, A2974–A3005. MR4155236
- [CLLW21] Ke Chen, Qin Li, Jianfeng Lu, and Stephen J. Wright, A Low-Rank Schwarz Method for Radiative Transfer Equation With Heterogeneous Scattering Coefficient, Multiscale Model. Simul. 19 (2021), no. 2, 775–801. MR4252068
- [CLW20] Ke Chen, Qin Li, and Stephen J. Wright, Schwarz iteration method for elliptic equation with rough media based on random sampling, Proceedings of the International Consortium of Chinese Mathematicians 2018, Int. Press, Boston, MA, 2020, pp. 163–178. MR4251142
- [CLLW20b] Shi Chen, Qin Li, Jianfeng Lu, and Stephen J Wright, Manifold learning and nonlinear homogenization, Preprint, arXiv: 2011.00568, 2020.
- [DG00] Laurent Dumas and François Golse, *Homogenization of transport equations*, SIAM J. Appl. Math. **60** (2000), no. 4, 1447–1470. MR1760042
- [E11] Weinan E, *Principles of multiscale modeling*, Cambridge University Press, Cambridge, 2011. MR2830582
- [EE03] Weinan E and Bjorn Engquist, *The heterogeneous multiscale methods*, Commun. Math. Sci. 1 (2003), no. 1, 87–132. MR1979846
- [EH09] Yalchin Efendiev and Thomas Y. Hou, *Multiscale finite element methods: Theory and applications*, Surveys and Tutorials in the Applied Mathematical Sciences, vol. 4, Springer, New York, 2009. MR2477579
- [GM01] Thierry Goudon and Antoine Mellet, *Diffusion approximation in heterogeneous media*, Asymptot. Anal. 28 (2001), no. 3-4, 331–358. MR1878799
- [Hac15] Wolfgang Hackbusch, *Hierarchical matrices: algorithms and analysis*, Springer Series in Computational Mathematics, vol. 49, Springer, Heidelberg, 2015, DOI 10.1007/978-3-662-47324-5. MR3445676
- [HMT11] N. Halko, P. G. Martinsson, and J. A. Tropp, Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions, SIAM Rev. 53 (2011), no. 2, 217–288, DOI 10.1137/090771806. MR2806637
- [Hor09] Mark F. Horstemeyer, *Multiscale modeling: a review*, Practical aspects of computational chemistry (2009), 87–135.
- [HJL17] J. Hu, S. Jin, and Q. Li, Asymptotic-preserving schemes for multiscale hyperbolic and kinetic equations, Handbook of numerical methods for hyperbolic problems, Handb. Numer. Anal., vol. 18, Elsevier/North-Holland, Amsterdam, 2017, pp. 103–129. [MR3645390]
- [Mar20] Per-Gunnar Martinsson, Fast direct solvers for elliptic PDEs, CBMS-NSF Regional Conference Series in Applied Mathematics, vol. 96, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2020.
 MR4050504
- [Owh15] Houman Owhadi, *Bayesian numerical homogenization*, Multiscale Model. Simul. **13** (2015), no. 3, 812–828, DOI [0.1137/140974596] [MR3369060]

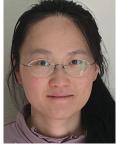
- [OS19] Houman Owhadi and Clint Scovel, Operator-adapted wavelets, fast solvers, and numerical homogenization: From a game theoretic approach to numerical approximation and algorithm design, Cambridge Monographs on Applied and Computational Mathematics, vol. 35, Cambridge University Press, Cambridge, 2019. MR3971243
- [PS08] Grigorios A. Pavliotis and Andrew M. Stuart, Multiscale methods: Averaging and homogenization, Texts in Applied Mathematics, vol. 53, Springer, New York, 2008. MR2382139

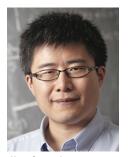




Ke Chen

Shi Chen





Qin Li

Jianfeng Lu



Stephen J. Wright

Credits

The opening photo is courtesy of Eureka_89 via Getty. Figure 1 is courtesy of the Society for Industrial and Applied Mathematics.

All other figures are courtesy of the authors.

All photos of the authors are courtesy of the authors.