



Accurate and scalable deep Maxwell solvers

Chenkai Mao^a and Jonathan A. Fan^{a,1}

Edited by Naomi J. Halas, Rice University, Houston, TX; received November 5, 2025; accepted March 26, 2026

Neural networks have promise as surrogate partial differential equation (PDE) solvers, but it remains a challenge to use these concepts to solve problems with high accuracy and scalability. In this work, we show that neural network surrogates can combine with iterative algorithms to accurately solve PDE problems featuring different scales, resolutions, and boundary conditions. We develop a subdomain neural operator model that supports arbitrary Robin-type boundary condition inputs, and we show that it can be utilized as a flexible preconditioner to iteratively solve subdomain problems with bounded accuracy. We further show that our subdomain models can facilitate the construction of global coarse spaces to enable accelerated, large scale PDE problem solving based on iterative multilevel domain decomposition. With two-dimensional Maxwell's equations as a model system, we train a single network to simulate large scale problems with different sizes, resolutions, wavelengths, and dielectric media distribution. We further demonstrate the utility of our platform in performing the accurate inverse design of multiwavelength nanophotonic devices. Our work presents a promising path to building accurate and scalable multiphysics surrogate solvers for large practical problems.

neural operator | surrogate models | iterative methods | GMRES | domain decomposition

Artificial neural networks have emerged as powerful tools for accelerating scientific computing, particularly in solving partial differential equations (PDEs) that govern complex physical processes including thermal transport (1, 2), structural mechanics (3, 4), fluid dynamics (5, 6), and electromagnetics (7–9). As specified by the Universal Approximation Theorem (10), these networks can serve as direct solution approximators and are theoretically capable of outputting arbitrary PDE solutions with an accuracy level determined by the training dataset and model size. Practically, however, it has been observed that neural network surrogate solvers are difficult to scale for reasons beyond model and training dataset size limitations. Most models are limited to solving toy problems with fixed domain sizes, fixed or implicit boundary conditions, and a narrow distribution of parameters (i.e., constitutive parameters, source types, etc.), which prevents them from generalizing beyond the training settings. Neural network solvers also become increasingly difficult to train and are less accurate as the problem size and parameter distribution grow.

To overcome these limitations, there has been interest to investigate how iterative PDE solving methods can be incorporated into machine learning workflows to enhance their capabilities and scalability. Iterative methods are established methods in conventional PDE solving that were developed to efficiently solve large scale problems, and they encompass different concepts. One class of concepts involves iterative neural network-enhanced algebraic preconditioners, where the idea is to train deep networks that take PDE residuals as inputs and that output the corresponding error corrections. These errors can then be used directly with basic iterative methods (11) or couple with Krylov subspace methods such as conjugate gradient or generalized minimal residual (GMRES) methods (12–14). These concepts have been adapted using all major classes of neural networks, including physics informed neural networks (PINNs) (15), neural operators (10, 16, 17) and graph neural networks (18, 19).

A second class of concepts pertains to the iterative solving of large scale PDE problems and include multigrid (20, 21), domain decomposition (22–27), and multilevel Schwarz methods (28). With multigrid methods, global problems are iteratively solved using a hierarchy of spatial resolutions, and neural network models have been utilized to predict low spatial frequency errors (20, 21) and to more generally solve PDE problems at different resolutions (29). With domain decomposition, global problems are subdivided into a set of subdomain problems that are iteratively evaluated until global self-consistency is obtained. Many machine learning efforts have focused on the solving of subdomain problems, where the online training of PINN models and the use of neural operator-based surrogate solvers (25) has been demonstrated at the subdomain level. For multilevel

Significance

This work introduces a model-agnostic framework for training and inference to enable accurate partial differential equation solving (down to double precision) for problems with arbitrary sizes and parameters, bridging the gap between problem-specific surrogate solvers and general-purpose accurate solvers. Specifically, we demonstrate that the generalized minimal residual method can be used in both training and inference to obtain neural operators that solve problems with arbitrary robin-type boundary conditions. Combined with the domain decomposition method and a coarse space constructed with the same neural operator, large problems can be solved accurately and efficiently. We use 2D Maxwell's equations as a model system and demonstrate broadband simulation and inverse design of problems larger than 200 wavelengths based on deep neural networks.

Author affiliations: ^aDepartment of Electrical Engineering, Stanford University, Stanford, CA 94305

Author contributions: C.M. and J.A.F. designed research; C.M. performed research; C.M. contributed new reagents/analytic tools; C.M. analyzed data; and C.M. and J.A.F. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2026 the Author(s). Published by PNAS. This article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹To whom correspondence may be addressed. Email: jonfan@stanford.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2530330123/-/DCSupplemental>.

Published April 27, 2026.

domain decomposition methods, many works focus on using neural networks to reduce the number of subdomain eigenvalue problems to be solved (30–32).

While these approaches show promise in certain capacities, it remains challenging to adapt them to general problems featuring a wide range of boundary conditions and domain sizes. It is particularly difficult to apply these techniques to indefinite PDEs such as Maxwell's equations, which is the focus here, where high spatial resolution is required and where the system matrix is indefinite and becomes ill-conditioned for large scale heterogeneous global domains. To date, efforts have been made to solve Maxwell's equations with neural networks integrated with iterative preconditioning (8, 33), multigrid (29, 34, 35), and domain decomposition methods (25, 36, 37), with the intent to extend the accuracy and scalability of deep solvers. While these efforts have demonstrated improved capabilities in certain limits, they still display shortcomings in key metrics pertaining to accuracy, scalability, dielectric media contrast, wavelength range, and boundary condition types. For instance, concepts based on iterative preconditioners and multigrid methods have been limited to fixed domain sizes and do not generalize to arbitrary boundary conditions. Iterative neural domain decomposition methods show promise in adapting to a wide range of global domain dimensions, but the limited accuracy of the surrogate subdomain solvers limits global domain simulation accuracy and scalability.

We propose an approach to neural network-enhanced iterative algorithms that supports the solving of PDE problems with scalability and bounded accuracy. As a model system, we consider 2D Maxwell's equations for TE polarization (i.e., E_z , H_x , and H_y field components) in the frequency domain, and we focus on the simulation of dielectric media with subwavelength-scale structuring, which is typical in the fields of nanophotonics and metamaterials. Our concept is outlined in Fig. 1 and showcases the use of iterative algorithms at both the global domain and subdomain levels to achieve scalable and accurate solving capabilities. At the subdomain level, we train a generalized neural operator that maps residuals to errors and that works for subdomain systems featuring different parameter maps, boundary conditions, resolutions, and wavelengths. We use the neural network as a general preconditioner in iterative residual minimization schemes to solve subdomain problems with bounded accuracy. At the global domain level, we use domain decomposition with a coarse-space correction to iteratively update subdomain boundary values until the global solution is self-consistent.

Our approach to iterative PDE solving combines multiple innovations that collectively enable PDE problem solving with unprecedented capabilities. First, we generalize our neural operator subdomain preconditioner to work for arbitrary Robin-type boundary conditions, which enables the adaptation of these subdomain solvers with advanced methods in domain decomposition such as the overlapping Schwarz method. We additionally generalize our solver to work for a continuous range of wavelengths, enabling multiwavelength and broadband evaluation. Second, we use the trained subdomain networks to efficiently solve subdomain eigenvalue problems with robin-to-robin maps. The solutions to these eigenvalue problems enable the construction of efficient coarse spaces in the global domain, achieving near-optimal scaling in which the iterations required remain constant or grow very slowly with increased problem size. Third, we combine iterative methods at both the subdomain and global domain level, which is a synergistic and critical strategy to global domain scaling because domain decomposition methods

require highly accurate subdomain solutions to prevent global-scale error propagation. To demonstrate the capabilities of our platform, we show the utilization of a deep learning platform that can accurately and efficiently optimize large scale chip-based and free space multiwavelength nanophotonic devices.

1. Results

1.1. Subdomain Neural Operators as General Preconditioners.

We first focus on the solving of subdomain problems, which can each be framed as a boundary value PDE problem with robin-type boundary conditions, g_i :

$$\nabla \times \nabla \times \mathbf{E}(\mathbf{r}) - \frac{\omega^2}{c^2} \epsilon(\mathbf{r}) \mathbf{E}(\mathbf{r}) = -i\omega\mu_0 \mathbf{J}(\mathbf{r}), \quad [1]$$

$$jk(\mathbf{r})\mathbf{E}_i(\mathbf{r}) - \frac{\partial \mathbf{E}_i}{\partial \mathbf{n}} = g_i, \quad [2]$$

Here, ω is the angular frequency, $\mathbf{E}$ is the electric field, $\epsilon(\mathbf{r})$ is the heterogeneous relative permittivity, $\mathbf{J}(\mathbf{r})$ is the electric current density, $k(\mathbf{r}) = \frac{\omega}{c} \sqrt{\epsilon(\mathbf{r})}$ is the wavevector on the boundary, and $\mathbf{n}$ is the outward normal vector. Using finite difference frequency domain (FDFD) discretization, the system can be described as a linear equation $Ax = b$, in which A is a sparse matrix that involves the differential operator and boundary operator on the left-hand sides of Eqs. 1 and 2, and b contains the source term and boundary values from the right-hand sides (Fig. 1 B, Left).

To solve these problems, our approach is to train a Fourier neural operator that takes auxiliary data (i.e., dielectric distribution, source maps, PML maps) and a residual as inputs and outputs the corresponding error (Fig. 1 B, Center). For a given approximate solution $\hat{x}$, the error can be quantified as $e = x^* - \hat{x}$ and the corresponding residual is $r = A(x^* - \hat{x}) = Ae$ (Fig. 1 B, Left). In this manner, the network approximates the A^{-1} operator on broad residual inputs, serving as a general preconditioner for iterative methods, in our case flexible GMRES (F-GMRES), whereby the error are used to construct basis directions for the Krylov subspace. The algorithmic details are shown in Section 1. As with all preconditioning methods, simulation accuracy can be bounded by evaluating the residual over the course of F-GMRES iterations and stopping the solver once a desired residual threshold has been met (Fig. 1 B, Right).

To implement our neural operator, we utilize a modified Fourier neural operator architecture (Fig. 2) together with a training scheme that is adapted for photonics problems. Conventionally, the Fourier neural operator architecture utilizes a linear transformation with weight R in the Fourier space, which has learned static weights that apply the same transformation to different input data. In this manner, the size of R can quickly become a storage bottleneck as the diversity of input data increases and the network scales up. To boost network expressiveness for a diverse input data distribution, we multiply the tensor R with a modulation tensor obtained from a modulation path that applies convolutional, pooling, and fully connected layers to auxiliary inputs. We also adopt a low-rank parameterization of R , where a linear layer expands a low-dimensional tensor to construct the full-sized tensor. The design of the modulation path is improved from previous work (25), and similar Fourier-space augmentations have been reported in related works (26).

To train the network, we construct a training set that is representative of the nanophotonics design space. The training set comprises 1M subdomains and their associated fields, and they include a wide range of freeform dielectric structures,




Fig. 1. PDE solving framework combining a deep subdomain solver with nested inner and outer iterative loops. (A) Iterative global domain solving using domain decomposition. The global domain is partitioned into overlapping subdomains that are iteratively solved using the overlapping Schwarz method until self-consistency. The subdomain solver also facilitates coarse space construction. Multiwavelength simulations are achieved by performing multiple global domain simulations at different wavelengths in parallel. (B) Iterative subdomain solving using neural operator-preconditioned F-GMRES. The neural network model uses the residual and auxiliary data as inputs, and it predicts the corresponding error, which serve as Krylov subspace basis directions in F-GMRES.

fields, sources, and Robin-type boundary conditions. The dielectric structures consist of discrete and continuous dielectric distributions with permittivity ranging from 1 to 8, and they are obtained by cropping subdomains from large scale FDFD simulations. To perform training, each training data sample

is recurrently processed in a F-GMRES loop preconditioned with our model and gradients are backpropagated to minimize both data loss and physical residual loss. Importantly, network outputs are subsequently used to build a Krylov subspace via F-GMRES, yielding new residuals that are recurrently fed back




Fig. 2. The model architecture, which contains a series of Fourier neural operator (FNO) blocks. Auxiliary inputs are processed with convolution, pooling, and fully connected layers to form a tensor that modulates linearly transformed lower Fourier modes in each FNO block. During training, a dynamic training set of residuals is constructed by iteratively performing F-GMRES on the training data with the model, with the output error vectors being used as Krylov bases in F-GMRES. The same F-GMRES loop is used in inference. $\mathcal{F}$: Fourier transform. σ : leaky-ReLU activation function. $\mathbf{W}$: channel mixer via 1-by-1 kernel convolution. $\mathbf{R}$: linear transform on the lower Fourier modes. $\mathbf{R}^*$: modulated linear transform through an element-wise multiplication with the modulation tensor.

into the network training process. We define the number of total F-GMRES iterations used in the training process as the “training data recurrence.” For the ease of implementation, we found out that a simple fixed point iteration training scheme can produce comparable results, but a problem-specific clipping is required for stable training (Algorithms 1 and 2).

With this approach, network training is explicitly streamlined with the F-GMRES process. In addition, a large ensemble of residuals are created from a single structure-field pair from the training set, amplifying the effective size of the training set without the need for additional simulated data. We further note that this recurrent training scheme has the capability to learn PDE hyperparameters implicitly. For example, our trained network is capable of solving broadband subdomain boundary value problems without explicit wavelength or resolution inputs, and the scaled information is implicitly learned from the encoding of residuals throughout the iterations. This provides a simple yet elegant solution to enable broadband model training without explicit feature or architecture engineering. Finally, the same F-GMRES loop can be utilized in both training and inference, naturally positioning the network as a general preconditioner. More details pertaining to the training process are provided in *SI Appendix*.

To evaluate the efficacy of the trained subdomain models, we utilize them as preconditioners for F-GMRES and solve problems based on test data comprising arbitrary dielectric material and source distributions. Fig. 3A shows examples of subdomain boundary value problems with heterogeneous parameter maps and random sources, PML distributions, wavelengths, and resolutions that are accurately solved using F-GMRES preconditioned by the trained model. A comparison of our approach with efficient implementations of CPU-based GMRES and BiCGSTAB using the PETSc library(38) indicates that our implementation requires 20 to 50 times fewer iterations to reach a given residual threshold (Fig. 3B), showcasing the advantage of neuro-preconditioning over conventional methods. These implementations consider a batch size of 1 for the time benchmark. In addition to these CPU baselines, we consider problems of increased size where GPU-acceleration of large matrix-vector multiplication becomes necessary. We include benchmarks with GPU-accelerated PETSc baselines in *SI Appendix*, section S2. Compared with conventional preconditioners including incomplete LU factorizations, the additive Schwarz method, and successive overrelaxation, our model-based preconditioning remains favorable with more than a 50× savings in number of iterations and a 5× reduction in runtime.

Interestingly, we observe that even if the ground truth data and model weights are in single precision with 32-bit floating point numbers, the model is able to accurately solve problems with double precision when used in a double precision F-GMRES algorithm with 64-bit floating point numbers. We further find that our network can also be used as an effective preconditioner for BiCGSTAB, but due to the sensitivity of BiCGSTAB to precision, it is reliable only up to single precision.

We further consider models that are trained with different training data recurrence (1 to 16) within the training F-GMRES loop, to quantify the impact of our iterative training scheme on preconditioner performance. The residual and field error as a function of model preconditioner iteration are shown in Fig. 3C. We observe that as the training data recurrence in F-GMRES increases, the preconditioner also improves with faster convergence rate. Note that the training epochs for models with different training data recurrence are adjusted such that the total number of updates and training time are kept the same. This

is an indication that the recurrent training scheme is effective in creating more diverse residual-error pairs to better train the model from a fixed-size dataset.

Our neural operator preconditioner can exhibit improved performance as the network gets larger, at the expense of inference time. This naturally raises the question of what the best model size is for a given problem: small models may take too many iterations to reach a certain error threshold while large models take too much time for each iteration. To understand and optimize this tradeoff, we parametrically sweep the number of model weights by varying the number of Fourier modes in the linear transform and number of hidden channels, and we quantify the performance and run time of each model. The results are summarized in Fig. 3D and confirm the tradeoff between model size, performance, and run time. In our experiments, models with approximately 1.5M weights are optimal. We leave further improvements in model performance via the meta-optimization of model architecture, size, and training trajectories to future research.

An examination of how well our preconditioner performs with out-of-distribution data is summarized in Fig. 3E, which shows the final relative residual after 100 iterations of F-GMRES conducted on 2,000 samples with different scales (i.e., grid points per wavelength) and dielectric distributions, including those featuring complex dielectric constants. For in-distribution cases that follow the parameter distribution of the training set, the neural network-preconditioned F-GMRES is effective in minimizing the residual with double precision. For distributions outside of the training range, the effectiveness of the model will decrease, but instead of producing completely erroneous results or causing stagnation, we observe the model-preconditioned F-GMRES can still provide accurate solutions in extended ranges, at the cost of more time and iterations. For subdomain systems featuring larger and complex-valued dielectric constants beyond the training set distribution, our solver still converges as long as the spatial frequencies in the subdomain are within the range of the training set. However, when the model attempts to minimize residuals with higher or lower spatial frequencies outside the range of the training set, it will fail eventually, as the learned weights in the convolution and Fourier layers are targeted for a given range of spatial frequencies. In general, we suggest a thoughtful construction of the training dataset to include the full range of parameters encountered during inference, rather than rely on GMRES convergence to solve out-of-distribution problems. While further enlarging the training set or network capacity can help address these limitations, global-domain partitioning, discussed in the following section, presents a more scalable approach.

1.2. Global Domain Solving with Domain Decomposition. We now show that our subdomain preconditioner models can be utilized in domain decomposition algorithms to solve large scale electromagnetics problems. We first show that our subdomain preconditioner can combine with one-level overlapping Schwarz methods to solve global problems with machine precision. We further quantify how subdomain model accuracy affects global solution accuracy as the problem size increases. To demonstrate, we solve a range of problems featuring random dielectric distributions, source distributions, boundary conditions, wavelengths, and sizes, all using the same subdomain preconditioner. To solve these problems, we use the approach outlined in Fig. 1: the global problem is solved using an outer loop of domain decomposition iterations to update the subdomain boundary conditions




Fig. 3. PDE subdomain solving with network-enhanced preconditioned F-GMRES. (A) Examples of the subdomain boundary value problems showcasing heterogeneous material maps, Robin-type boundary conditions, sources (cyan), use of PML (yellow), wavelengths, and spatial resolutions. The outputs are produced with NN-preconditioned F-GMRES. (B) Benchmarks of network-preconditioned iterative methods with conventional iterative methods. The shaded width for each line represents one SD of variation as computed using 30 test data. (C) Comparison between models with the same number of weights but trained under different numbers of training data recurrence with GMRES. The shaded region represents one SD of variation computed using 500 test data. (D) Comparison between models with the same training setup and different number of weights. The green boxes indicate the best model size for achieving each field error threshold within the given time budget. (E) Benchmark on out-of-distribution samples. Each dot represents the final residual of a test device after running 100 iterations of NN-preconditioned F-GMRES algorithm.

until global convergence, and for each domain decomposition iteration, an inner loop of network-preconditioned F-GMRES subdomain iterations are performed in parallel. The inner loop stops after the F-GMRES reaches a certain predefined residual threshold. The results are summarized in Fig. 4 for different subdomain residual thresholds.

While all setups are able to achieve convergence, the final global solution converges closer to the ground truth solution when we impose a tighter subdomain residual threshold, at the expense of needing more iterations. In addition, as the global problem size increases, more accurate subdomain solutions are required in order to obtain a globally accurate solution. In this regard, we envision that fully convergent subdomain solvers like F-GMRES are required for true scalability to large global problems. We also want to point out that even though we stop the subdomain F-GMRES iterations when it reaches a certain residual threshold in this benchmark, in practice we found out that the global solution can also converge even if the subdomain is not converged at each iteration. For example, even if we only perform 1 GMRES iteration for the inner loop, the global outer loop can still converge, often with less total time. However, the stability is worsened and breakdowns in convergence can occur. We leave the meta-optimization of subdomain and global residual minimization, as well as their robustness, to future work.

While the one-level Schwarz method converges when using accurate subdomain solutions, it ultimately requires very large numbers of domain decomposition iterations for large global domain sizes. This scaling trend is intrinsic to the one-level

Schwarz method where subdomain boundary conditions propagate and update only for nearest neighboring subdomains in an iteration, thereby requiring increasing numbers of iterations for information to propagate across the global domain as the domain size increases. To overcome these limits, a means of global information exchange is required. One typical approach involves the construction of a global coarse space, which is a low-dimensional subspace spanning the global domain that is obtained by solving a dimensionality-reduced global problem (39). Coarse space solutions can be incorporated with the one-level Schwarz method to yield the two-level overlapping Schwarz method framework. In a given iteration, the coarse space correction is efficiently computed to globally correct the subdomain boundary mismatch errors from the one-level Schwarz solution (Fig. 5A), accelerating domain decomposition convergence.

Traditional spectral coarse space construction relies on solving eigenvalue problems for each subdomain followed by solving a dimensionality-reduced global problem (40, 41), and large computational overhead is required to solve these eigenvalue problems from scratch (39). We propose to use our subdomain models to accelerate the solving of these subdomain-scale eigenvalue problems. The concept is illustrated in Fig. 5B. The eigenvalue problems for each subdomain are defined by the linear operator $\mathbf{B}$, which acts on a given Robin-type boundary value vector x and produces the boundary values for its neighbors y . Eigenvalue solutions to $\mathbf{B}$ comprise eigenvectors u and associated eigenvalues η , and they represent electromagnetic modal solutions with




Fig. 4. Maxwell PDE solving with the one-level Overlapping Schwarz Method. Three problems with different wavelengths, resolutions, and sizes are simulated. The problems comprise grids of 5×5 , 12×12 , and 25×25 subdomains. The global solutions converge with increasing accuracy as we increase the accuracy of the subdomain model by imposing stricter residual thresholds. The final field error plots are shown in a scale spanning -1% to 1% the maximum ground truth field value. The simulations are conducted with single precision.


boundaries u and fields that can be readily evaluated using our neural operator. Modes with the largest eigenvalues η decay the slowest, and for the purposes of coarse space construction, we consider coarse space solutions to comprise linear combinations of a finite number of slowly converging modes with a lower eigenvalue threshold. To calculate u and η for a given subdomain, we evaluate (x, y) pairs using our subdomain neural operator, where x are heuristically chosen as pseudoplane waves, and we use the Rayleigh–Ritz method to obtain the approximated u and η , with details in Section 1. Calculation of the modes only needs to be performed once for a given global domain layout. This process enables automatic coarse space basis function construction, even for subdomains with highly heterogeneous dielectric distributions. The method is also accelerated with our subdomain neural operator because our subdomain models can perform fast batched evaluation of (x, y) pairs for multiple subdomains in parallel.

Our process for performing coarse space correction for a given domain decomposition iteration is standard: boundary error from the one-level Schwarz method is evaluated and linear combinations of modes that can minimize this error are determined (Fig. 5C). For n subdomains and k_{coarse} modes per subdomain, this coarse problem reduces to solving a $n \times k_{\text{coarse}}$ dimensional linear problem, which in most cases are very small and can be efficiently and quickly solved using direct methods. The fields from the one-level Schwarz method and coarse space correction are combined and used for the subsequent domain decomposition iteration. To study the scaling trend of this coarse space problem, we evaluate the two-level overlapping Schwarz Method on two subdomain model choices: one with 64 by 64 subdomain size and the other with 256 by 256 subdomain size. We perform both one-level and two-level overlapping Schwarz Methods on problems with different global domain sizes, and the results are summarized in Tables 1 and 2.


Both models are able to solve arbitrary-sized problems with different wavelengths and resolutions up to grid size of about $3,000$ by $3,000$, and physical size of up to 200 wavelengths, while the two-level Schwarz method enables a significant reduction in domain decomposition iterations. In fact, when the coarse space is well constructed, the number of iterations to converge either stays constant or grows very slowly as the global problem size increases (42). The term k_{coarse} indicates the number of basis functions that we evaluate in each subdomain, and it is effectively equivalent to running an additional k_{coarse} global domain decomposition iterations compared to the one-level Schwarz method. This one-time cost for a specific dielectric distribution is especially efficient in cases where multiple source configurations are evaluated on the same dielectric distribution. We emphasize that the effectiveness of the coarse space is crucial and remains an active research area for systems including Maxwell’s equations with heterogeneous medium. In our experiments, highly resonant subdomains with discrete material boundaries worsen the effectiveness of the coarse space, and with larger subdomain size and higher frequency (fewer grids per λ), the required size of the coarse space also grows accordingly. When the heterogeneous dielectric map is grayscale and varies smoothly, the coarse space appear to be very effective, as shown in Tables 1 and 2. We have included one visual example of the convergence history in SI Appendix.

To further demonstrate the difficulty imposed by discrete material boundaries, we have included an example of the simulation of a piecewise linear dielectric and its variants in SI Appendix. While solving the piecewise linear dielectric system with the two-level Schwarz method took 27 steps to converge, it took only nine steps to converge when the same dielectric distribution was smoothed using a Gaussian filter. Interestingly, if we modify the piecewise linear dielectric to have uniform dielectric values at the subdomain boundaries, only 11 iterations are required for convergence. This indicates that noise in the robin-type boundary

A Two-level overlapping Schwarz method



B NN-assisted subdomain eigenvalue problem



C Basis function as eigenvectors




Fig. 5. Two-level overlapping Schwarz Method and coarse space. (A) Exemplary problem with 3 by 3 subdomains. Within one iteration, the global coarse space correction amends the boundary mismatches resulting from the local solve, resulting in an accurate combined solution. (B) The linear operator and eigenvalue problem in the subdomain. (C) Examples of basis functions as eigenvectors, for each subdomain with heterogeneous parameter map. The basis functions are then distributed globally as a basis function in the coarse space.

conditions due to dielectric discontinuities can impact global domain convergence. We anticipate that modifications to the finite difference description of material discontinuities at boundaries, modifications to the boundary conditions themselves, and better techniques for domain decomposition that minimize the presence of discrete material boundaries at subdomain boundaries can collectively improve the convergence of simulations involving heterogeneous materials.

Finally, we demonstrate that our general purpose multilevel solver can be used for accurate and scalable inverse design using gradient-based optimization. Freeform density-based optimization is performed using the adjoint variables method (AVM) coupled with gradient-based optimization. We begin with a homogeneous grayscale dielectric representing our photonic device, and we use Adam-based gradient descent to iteratively push the dielectric values at each voxel in a manner that maximizes our figure of Merit. Differentiable filtering and projection operations

can be used to push the dielectric distribution toward increasingly binary values over the course of optimization. Gradients to the dielectric constant at each voxel are calculated using forward and adjoint simulations with our solver, and in cases where the gradients from AVM are part of a larger computational graph, a custom Jacobian vector product (JVP) ensures the full computational graph is differentiable. We consider implicit differentiation over automatic differentiation because the computational graph quickly becomes intractable for large problems, rendering automatic differentiation infeasible due to memory constraint issues.

Results for the freeform design of a wavelength division multiplexer, multiwavelength volumetric coupler, and multifunctional metalens are shown in Fig. 6. Note that even though the topology of the final binary and optimized devices are clearly out of the training distribution, the dielectric constants and wavenumbers are still within the training range, thus the neural preconditioner

Table 1. Comparison between one-level and two-level Schwarz methods, for 64 by 64 grid subdomain preconditioner model

subdomain grid 64 by 64		grids per $\lambda = 80$ $k_{coarse} = 16$		grids per $\lambda = 44$ $k_{coarse} = 24$		grids per $\lambda = 24$ $k_{coarse} = 36$	
# of sub-domains	grid size	N_λ	iterations 1L / 2L	N_λ	iterations 1L / 2L	N_λ	iterations 1L / 2L
4^2	240^2	3.1	22 / 7	5.7	26 / 5	10.7	24 / 4
8^2	480^2	5.8	48 / 9	10.7	69 / 7	20.1	67 / 5
12^2	720^2	8.8	75 / 10	16.2	107 / 8	30.5	117 / 6
16^2	960^2	12.3	102 / 10	22.7	163 / 8	42.6	155 / 6
24^2	1440^2	18.4	181 / 19	33.9	284 / 10	63.6	271 / 7

N_λ : simulation size in terms of number of wavelengths within the global domain.

1L / 2L: short for one-level and two-level.

Table 2. Comparison between one-level and two-level Schwawrz methods, using a 256 by 256 grid subdomain preconditioner model

subdomain size 256 by 256		grids per $\lambda = 28$ $k_{coarse} = 100$		grids per $\lambda = 22$ $k_{coarse} = 120$		grids per $\lambda = 15$ $k_{coarse} = 160$	
# of sub-domains	grid size	N_λ	iterations 1L / 2L	N_λ	iterations 1L / 2L	$/ N_\lambda$	iterations 1L / 2L
4^2	994^2	35.4	44 / 4	46.5	40 / 5	67.9	50 / 5
6^2	1486^2	53.8	91 / 4	70.7	75 / 6	103	79 / 9
8^2	1978^2	70.9	170 / 7	93.0	174 / 9	136	138 / 12
10^2	2470^2	89.5	262 / 8	117	224 / 11	172	174 / 16
12^2	2962^2	106	421 / 15	140	350 / 19	203	251 / 23

remains effective in each subdomain. The multiplexer and metalens are designed to have binary dielectric values while the coupler is designed to have grayscale dielectric values. The subdomain solutions are specified to have high accuracy (i.e., residual $<10^{-6}$). In all cases, the devices exhibit high efficiencies typical of freeform volumetric photonic media. A comparison of the multiplexer optimization trajectory produced using our solver and a ground truth FDFD solver, which uses automatic differentiation to calculate gradients, are nearly identical. The final devices optimized from both methods have nearly identical layouts and performance, indicating that our solver performs nearly identically to a ground truth FDFD solver. These results represent demonstrations of photonic inverse design for large scale domains ($>100\lambda$) with a neural network surrogate solver.

2. Discussion

In this work, we demonstrate a scalable and accurate approach to building general purpose PDE surrogate solvers using neural networks. The key to our approach is to build a subdomain surrogate solver that is capable of solving a general boundary value problem for a diverse range of wavelengths, source distributions, and auxiliary parameters. Utilization of this subdomain solver as a subdomain preconditioner for F-GMRES and as an eigenvalue mode solver for course grid construction enables streamlining with multilevel domain decomposition methods and the solving of subdomain and global domain problems with machine precision. Our framework targets the building of scalable and accurate models for repeated application to a broad range of problems as shown in Fig. 6, such that the upfront training cost can be amortized to essentially negligible. These methods can be directly integrated with reparameterization algorithms to regularize device layouts (43–45), generative deep networks to perform advanced freeform optimization (46–49), and agentic workflows (50) to automate computer aided design processes.

While our approach points out a promising path forward in neural PDE solving, there are many interesting directions for future research and opportunities to improve our concept. First, there are many opportunities to explore and understand how the model architecture, model size, and its implementation in domain decomposition workflows can be optimized to accelerate PDE solving. New architecture concepts beyond our modified FNO have the potential to utilize fewer network weights without a reduction in performance, pushing subdomain capabilities beyond the networks optimized in Fig. 3D. The optimization of subdomain size in the context of multilevel PDE solving is important to navigate the tradeoff between subdomain size, model complexity, coarse space construction speed, and domain decomposition convergence. There also are opportunities to

modulate the accuracy of subdomain and global solutions over the course of global iterative solving in a manner that reduces computational cost. For example, an intuitive strategy involves the gradual increase of subdomain accuracy over the course of global iterative solving such that the subdomain solutions in early iterations are approximate but sufficiently accurate to facilitate global boundary updating.

Second, a detailed computational cost and speed analysis is required to understand where our neural network solvers have an opportunity to outperform conventional solvers. The problems that we consider in this study are the two-dimensional problems with up to several million voxels, for which efficient direct solvers based on sparse matrix inversion can apply. In this limit, when benchmarked against a highly efficient direct sparse matrix solver [PARDISO (51)] running on a workstation with 80 CPU cores, our solver (running on 1 NVIDIA RTX A6000 GPU) is approximately twice as slow. We anticipate that for problems with larger sizes where direct solvers cannot apply (i.e., >10 million grids in 2D and >1 million voxels in 3D), our methods will have a decisive advantage, and proper benchmarking with fully optimized workflows will be required to understand these computational scaling trends.

Third, research is required to understand how to stably scale these methods to domain sizes much larger than 200λ , particularly for noncoercive and indefinite PDE problems. Although domain decomposition methods were originally developed for coercive elliptic problems, their robust and scalable extension to indefinite wave problems, including Maxwell’s equations, remains an active area of research. While we were able to obtain convergent results in this study, we did also observe that as the problem size grew, convergence became more and more sensitive to hyperparameters, such as the amount of overlap between subdomains and the momentum term in the boundary update step. The effectiveness of the coarse space also slowly degraded as the problem size increased. We anticipate that neural networks specifically trained to perform coarse space computation can potentially help with domain size scaling. We also anticipate that other types of domain decomposition methods, including finite element tearing and Interconnecting and Multigrid methods, could potentially offer certain benefits in terms of efficiency and stability.

3. Methods

3.1. Training Data Generation. For the model with 64 by 64 subdomain problem size, ground truth simulations are conducted with an FDFD solver (52). We generate diverse dielectric distributions using Gaussian random field and Voronoi map masked by free-from shapes, to create continuous and discrete material distributions encountered in nanophotonics inverse designs.




Fig. 6. Freeform gradient-based inverse design with our multilevel solver. Topology optimization is performed on three types of photonic devices using the adjoint variables method and implicit differentiation: a visible light wavelength division multiplexer (*Left*), a multiwavelength radio frequency coupler (*Center*), and a multifunctional volumetric near-infrared metalens (*Right*). The cyan region indicates the forward source region and the red regions indicate the figure-of-merit windows, which is also the location for adjoint sources. The optimization trajectories and final freeform devices based on our neural network solver are nearly identical to the ground truth. A single subdomain model of input size 256×256 is used in all three examples.

The Gaussian random fields and Voronoi maps have feature sizes from 1 to 50 pixels to cover structures from subwavelength to several-wavelength large. We run 4,000 simulations with grid size 600 by 600 on these dielectric data, with randomly placed line sources with diverse sinusoidal profiles, under wavelength from 2 to 4 cm, resolution from 0.2 to 0.8 mm, and permittivity from 1 to 8. The generated field maps, dielectric distributions, source and PML maps are cropped into a total of 1M training samples with size of 64 by 64. The Robin-type boundary values are prepared using the cropped dielectric maps and the fields, which are used to compute the residual maps used during training.

For the model with 256 by 256 subdomain problem size, the same methods of Gaussian random field and Voronoi map are used to generate dielectric distributions. The source maps are used as pure Gaussian noise that span the whole domain, and the PMLs are randomly placed in both x and y directions to cover scenarios encountered in DDM subdomains. The Robin-type boundary values are also generated as Gaussian noise with different length scales (SD). The source and boundary values are heuristically scaled down within the PML regions for more stable training.

In both cases, a uniaxial-PML (UPML) is adopted. Which has the formulation:

$$\nabla \times \mathbf{E} = -i\omega\mu[\mathbf{s}]\mathbf{H}, \quad [3]$$

$$\nabla \times \mathbf{H} = \mathbf{J} + i\omega\epsilon[\mathbf{s}]\mathbf{E}, \quad [4]$$

in which the tensor $[\mathbf{s}]$ has form:

$$[\mathbf{s}] = \begin{bmatrix} \frac{s_y s_z}{s_x}, & 0, & 0 \\ 0, & \frac{s_z s_x}{s_y}, & 0 \\ 0, & 0, & \frac{s_x s_y}{s_z} \end{bmatrix} \quad [5]$$

with

$$s_\zeta(\zeta) = 1 + \frac{\sigma_\zeta(\zeta)}{i\omega\epsilon_0}, \quad \sigma_\zeta(\zeta) = \sigma_{\zeta, \max} \cdot \left(\frac{\zeta}{d}\right)^m \quad [6]$$

$$\sigma_{\zeta, \max} = -\frac{(m+1) \ln R}{2\eta_0 d \zeta}, \quad \zeta = x, y, z, \quad [7]$$

in which d is the width of the PML, R is the ideal reflectivity (chosen as e^{-30}), and η_0 is the impedance of free space, and m is a coefficient chosen as 3. The s matrices (in our case s_x and s_y) are prepared and cropped as auxiliary inputs to the model.

3.2. Subdomain Training and Inference Algorithms. Here we present the NN-preconditioned Flexible GMRES algorithm for training and inference. Training is performed under single precision and both single and double precision can be used for inference. Both physics and data loss functions can be utilized to train the network. The network gradients are updated in each GMRES iteration. We want to mention that backpropagation through the entire GMRES iterations is also possible, but in our practice we observe more stable training with backpropagation at each step.

For the ease of implementation, we also present a training method based on fixed point iterations, which gives near identical result compared to training with GMRES. One difference is that fix-point-iterations are divergent for indefinite systems, especially at the beginning of the training where network produces noisy outputs. We found out that a clipping strategy is effective to obtain stable training. The exact clipping threshold is dependent on the data recurrence N , as well as the PDE.

3.3. Two-Level Schwarz Method with Coarse Space. All of the one-level and two-level overlapping Schwarz methods used in this work are implemented with an abstraction which we call the "partitioner." For a given global problem configuration, the partitioner takes care of partitioning global domain data into

Algorithm 1 NN-preconditioned Flexible GMRES (NN-FGMRES)

Require: Model **M**, right hand side **b**, auxiliary input **aux**, operator **A**, residual threshold r_{th} or maximum data recurrence **N**

Require: (in training) MSE loss function **L₂**, optimizer **opt**, weight for data and residual loss **c_d**, **c_r**, ground truth field **y**

- 1: $\mathbf{x}_0 = \mathbf{0}$
- 2: $\mathbf{r}_0 = \mathbf{b} - \mathbf{A}(\mathbf{x}_0)$
- 3: $\beta = \|\mathbf{r}_0\|_2$
- 4: $\mathbf{v}_1 = \mathbf{r}_0 / \beta$
- 5: for $j = 1, \dots, N$
- 6: $\text{opt.zero_grad}()$ ▷ Training only
- 7: $\mathbf{z}_j = \mathbf{M}(\mathbf{v}_j, \text{aux})$ ▷ Inference: convert input to 32 bit, output back to 64 bit
- 8: $\mathbf{w} = \mathbf{A}(\mathbf{z}_j)$
- 9: For $i = 1, \dots, j$
- 10: $h_{ij} = (\mathbf{v}_i, \mathbf{w})$ ▷ complex inner product
- 11: $\mathbf{w} = \mathbf{w} - h_{ij}\mathbf{v}_j$
- 12: $h_{j+1,j} = \|\mathbf{w}\|_2$
- 13: $\mathbf{v}_{j+1} = \mathbf{w} / h_{j+1,j}$
- 14: $\mathbf{Z}_j := [\mathbf{z}_1, \dots, \mathbf{z}_j]$, $\bar{\mathbf{H}}_j = \{h_{ik}\}_{1 \leq i \leq j; 1 \leq k \leq j}$
- 15: $\mathbf{y}_j = \text{argmin}_{\mathbf{y}} \|\beta \mathbf{e}_1 - \bar{\mathbf{H}}_j \mathbf{y}\|_2$
- 16: $\mathbf{x}_j = \mathbf{x}_0 + \mathbf{Z}_j \mathbf{y}_j$
- 17: $\mathbf{r}_j = \mathbf{b} - \mathbf{A}(\mathbf{x}_j)$
- 18: $\text{loss} = c_r \mathbf{L}_2(\mathbf{r}_j) + c_d \mathbf{L}_2(\mathbf{x} - \mathbf{y})$ ▷ Training only
- 19: $\text{loss.backward}(); \text{opt.step}()$ ▷ Training only
- 20: $\text{stop_gradient}(\mathbf{v}_j, \mathbf{Z}_j, \bar{\mathbf{H}}_j)$ ▷ Training only
- 21: if $r_j < r_{th}$ ▷ Inference only
- 22: break ▷ Inference only

overlapping (or nonoverlapping) subdomain data, and assembling them back to the global domain. The subdomains all have a fixed grid size (so they can be efficiently processed as a batch), but they can in principle have different grid resolutions and thus different physical sizes (as shown in Fig. 14). For example, a natural strategy can be a quad-tree partition of the global domain, and whenever the physical size of a quadrant is larger than a chosen threshold (e.g., 5λ), it will be decomposed into four smaller regions. The subdomains can be arbitrarily arranged and share different overlap with its neighbors, as long as their union area covers the whole global domain. In Figs. 4–6 and Tables 1 and 2, the subdomains are arranged in a rectangular grid, with each row and column have the same overlapping configurations. Once a specific partition geometry is chosen, at each domain decomposition iteration, the partitioner will take care of the boundary update for each subdomain from its neighbors,

Algorithm 2 Fixed-point-iteration training of NN-preconditioner

Require: Model **M**, right hand side **b**, auxiliary input **aux**, operator **A**, maximum data recurrence **N**, clipping threshold th_{clip} , MSE loss function **L₂**, optimizer **opt**, weight for data and residual loss **c_d**, **c_r**, ground truth field **y**

- 1: $\mathbf{x} = \mathbf{0}$
- 2: $\mathbf{r} = \mathbf{b} - \mathbf{A}(\mathbf{x})$
- 3: for $j = 1, \dots, N$
- 4: $\text{opt.zero_grad}()$
- 5: $\mathbf{e} = \mathbf{M}(\mathbf{r}, \text{aux})$
- 6: $\mathbf{x} = \mathbf{x} + \mathbf{e}$
- 7: $\mathbf{r} = \mathbf{b} - \mathbf{A}(\mathbf{x})$
- 8: $\text{loss} = c_r \mathbf{L}_2(\mathbf{r}) + c_d \mathbf{L}_2(\mathbf{x} - \mathbf{y})$
- 9: $\text{loss.backward}(); \text{opt.step}()$
- 10: $\mathbf{r} = \text{stop_gradient}(\text{clip}(\mathbf{r}, th_{clip}))$

Algorithm 3 Two-level Schwarz method with coarse space

Require: Dielectric distribution ϵ , source distribution **j**, PML distribution **S**, subdomain partition configuration **config**, number of basis functions per subdomain **k_{coarse}**, subdomain solver **M**, subdomain residual threshold th_{sub} , global residual threshold th_{global} , momentum **m**

- 1: $\mathbf{P} = \text{init_partitioner}(\epsilon, \text{config})$
- 2: $\text{eps_batch}, \text{src_batch}, \text{PML_batch} = \mathbf{P}.\text{partition}(\epsilon, \mathbf{j}, \mathbf{S})$
- 3: $\mathbf{M}.\text{setup}(\text{eps_batch}, \text{PML_batch})$ ▷ zero source
- 4: ifuse_coarse_space ▷ One-time precompute for coarse space
- 5: $\mathbf{x} = \text{prepare_pseudo_plane_wave_bc}(k_{coarse})$
- 6: $\mathbf{x} = \text{gram_schmidt}(\mathbf{x})$
- 7: $\mathbf{fields} = \mathbf{M}.\text{solve}(\mathbf{x}, th_{sub})$
- 8: $\mathbf{y} = \mathbf{P}.\text{extract_neighboring_bc}(\mathbf{fields})$
- 9: $\mathbf{H} = (\mathbf{x}, \mathbf{y})$ ▷ complex inner product
- 10: $\eta, \mathbf{u} = \text{eig}(\mathbf{H})$
- 11: $\mathbf{P}.\text{assemble_BQTBQ}(\mathbf{u})$
- 12: $\mathbf{x}, \mathbf{r} = \mathbf{0}, \mathbf{0}$
- 13: $bc = \mathbf{P}.\text{get_bc}(\mathbf{x_batch})$
- 14: $\mathbf{M}.\text{setup}(\text{eps_batch}, \text{src_batch}, \text{PML_batch})$
- 15: while $r > th_{global}$ ▷ Domain decomposition loop
- 16: $\mathbf{x}_{new} = \mathbf{M}.\text{solve}(bc, th_{sub})$
- 17: $bc_{new} = \mathbf{P}.\text{get_bc}(\mathbf{x}_{new})$
- 18: $bc = m \cdot bc + (1 - m) \cdot bc_{new}$
- 19: $r = \text{get_residual}(\mathbf{P}.\text{assemble}(\mathbf{x}_{new}))$
- 20: ifuse_coarse_space
- 21: $e_{bc} = \mathbf{P}.\text{get_bc_mismatch}(\mathbf{x}_{new})$
- 22: $\mathbf{z}_{coarse} = \mathbf{P}.\text{solve_coarse_space}(e_{bc})$ ▷ Solve for equation 9
- 23: $bc_{coarse} = \mathbf{P}.\text{coarse_space_correction}(\mathbf{z}_{coarse})$
- 24: $bc = bc + bc_{coarse}$
- 25: **Return** $\mathbf{P}.\text{assemble}(\mathbf{x}_{new})$

as well as a coarse space correction in two-level methods. The subdomain data will be batched and iteratively solved with the NN-preconditioned F-GMRES algorithm.

For the coarse space construction, it is well known that for subdomains with homogeneous medium, the natural basis are plane waves traveling in different directions spanning the unit circle (42). We use the Rayleigh-Ritz method and heuristically choose a set of pseudoplane wave trial basis as probe boundary conditions to construct a subspace that approximates the Robin-to-Robin map. The pseudoplane wave boundary conditions are constructed by first compute the wave vector inside the medium for a chosen plane wave direction, and then accumulate the phase around the boundary. We have also experimented with other trial boundary conditions such as random Gaussian or Fourier modes at the boundary, which produces acceptable but worse results. Note that when solving for each boundary value problem for these Rtr maps, homogeneous equations without the source term are used. We have also tried the singular value decomposition method to get the basis vectors, and it yields similar quality coarse space compared to solving for the eigenvectors. At each domain decomposition iteration, we solve the coarse space correction as a linear system:

$$BQz_{coarse} = e_{bc}, \quad [8]$$

in which **B** is the Robin-to-Robin map, each column of **Q** is the boundary values for a basis vector, z_{coarse} is the linear coefficients of the basis we want to solve, which minimizes the boundary mismatch error e_{bc} . The standard practice is to

solve for its normal equation:

$$(\overline{BQ})^T BQz_{\text{coarse}} = (\overline{BQ})^T e_{bc}, \quad [9]$$

for which we can assemble (and factor if needed) the small $(\overline{BQ})^T BQ$ matrix beforehand for fast solve at run time. The process is summarized in Algorithm 3.

1. S. Cai, Z. Wang, S. Wang, P. Perdikaris, G. E. Karniadakis, Physics-informed neural networks for heat transfer problems. *J. Heat Transf.* **143**, 060801 (2021).
2. Y. J. Wu, L. Fang, Y. Xu, Predicting interfacial thermal resistance by machine learning. *npj Comput. Mater.* **5**, 56 (2019).
3. W. He, J. Li, X. Kong, L. Deng, Multi-level physics informed deep learning for solving partial differential equations in computational structural mechanics. *Commun. Eng.* **3**, 151 (2024).
4. X. Li *et al.*, Mechanics-informed autoencoder enables automated detection and localization of unforeseen structural damage. *Nat. Commun.* **15**, 9229 (2024).
5. B. Font, F. Alcántara-Ávila, J. Rabault, R. Vinuesa, O. Lehmkuhl, Deep reinforcement learning for active flow control in a turbulent separation bubble. *Nat. Commun.* **16**, 1422 (2025).
6. T. Li, L. Biferale, F. Bonaccorso, M. A. Scarpolini, M. Buzicotti, Synthetic Lagrangian turbulence by generative diffusion models. *Nat. Mach. Intell.* **6**, 393–403 (2024).
7. E. Zhu *et al.*, Frequency transfer and inverse design for metasurface under multi-physics coupling by Euler latent dynamic and data-analytical regularizations. *Nat. Commun.* **16**, 2251 (2025).
8. R. Trivedi, L. Su, J. Lu, M. F. Schubert, J. Vuckovic, Data-driven acceleration of photonic simulations. *Sci. Rep.* **9**, 19728 (2019).
9. J. Jiang, M. Chen, J. A. Fan, Deep neural networks for the evaluation and design of photonic devices. *Nat. Rev. Mater.* **6**, 679–700 (2021).
10. L. Lu, P. Jin, G. E. Karniadakis, DeepONet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv [Preprint]* (2019). <https://doi.org/10.48550/arXiv.1910.03193> (Accessed 12 April 2026).
11. J. T. Hsieh, S. Zhao, S. Eismann, L. Mirabella, S. Ermon, Learning neural PDE solvers with convergence guarantees. *arXiv [Preprint]* (2019). <https://doi.org/10.48550/arXiv.1906.01200> (Accessed 12 April 2026).
12. A. Kopaničková, G. E. Karniadakis, DeepONet based preconditioning strategies for solving parametric linear systems of equations. *SIAM J. Sci. Comput.* **47**, C151–C181 (2025).
13. J. Luo *et al.*, Neural Krylov iteration for accelerating linear system solving. *Adv. Neural Inf. Process. Syst.* **37**, 128636–128667 (2024).
14. V. Trifonov *et al.*, Learning from linear algebra: A graph neural network approach to preconditioner design for conjugate gradient solvers. *arXiv [Preprint]* (2024). <https://doi.org/10.48550/arXiv.2405.15557> (Accessed 12 April 2026).
15. M. Raissi, P. Perdikaris, G. E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019).
16. Z. Li *et al.*, Fourier neural operator for parametric partial differential equations. *arXiv [Preprint]* (2020). <https://doi.org/10.48550/arXiv.2010.08895> (Accessed 12 April 2026).
17. Q. Cao, S. Goswami, G. E. Karniadakis, Laplace neural operator for solving differential equations. *Nat. Mach. Intell.* **6**, 631–640 (2024).
18. A. Sanchez-Gonzalez *et al.*, “Learning to simulate complex physics with graph networks” in *International Conference on Machine Learning* (PMLR, 2020), pp. 8459–8468.
19. A. Bryutkin *et al.*, Hamlet: Graph transformer neural operator for partial differential equations. *arXiv [Preprint]* (2024). <https://doi.org/10.48550/arXiv.2402.03541> (Accessed 12 April 2026).
20. D. Dong, W. Suo, J. Kou, W. Zhang, PINN-MG: A multigrid-inspired hybrid framework combining iterative method and physics-informed neural networks. *arXiv [Preprint]* (2024). <https://doi.org/10.48550/arXiv.2410.05744> (Accessed 12 April 2026).
21. E. Zhang *et al.*, Blending neural operators and relaxation methods in PDE numerical solvers. *Nat. Mach. Intell.* **6**, 1303–1313 (2024).
22. A. D. Jagtap, G. E. Karniadakis, Extended physics-informed neural networks (XPINNs): A generalized space-time domain decomposition based deep learning framework for nonlinear partial differential equations. *Commun. Comput. Phys.* **28**, 2002–2041 (2020).
23. S. Basir, I. Senocak, A generalized Schwarz-type non-overlapping domain decomposition method using physics-constrained neural networks. *arXiv [Preprint]* (2023). <https://doi.org/10.48550/arXiv.2307.12435> (Accessed 12 April 2026).
24. V. Dolean, A. Heinlein, S. Mishra, B. Moseley, Multilevel domain decomposition-based architectures for physics-informed neural networks. *Comput. Methods Appl. Mech. Eng.* **429**, 117116 (2024).
25. C. Mao, R. Lupo, T. Dai, M. Chen, J. Fan, “Towards general neural surrogate solvers with specialized neural accelerators” in *Proceedings of the 41st International Conference on Machine Learning*, R. Salakhutdinov *et al.*, Eds. (PMLR, 2024), vol. 235, pp. 34693–34711.
26. Z. Li *et al.*, Scale-consistent learning for partial differential equations. *arXiv [Preprint]* (2025). <https://doi.org/10.48550/arXiv.2507.18813> (Accessed 12 April 2026).
27. R. Wu, N. Kovachki, B. Liu, A learning-based domain decomposition method. *arXiv [Preprint]* (2025). <https://doi.org/10.48550/arXiv.2507.17328> (Accessed 12 April 2026).

Data, Materials, and Software Availability. All study data are included in the article and/or *SI Appendix*. The code of the project can be found at: <https://github.com/ChenkaiMao97/Deep-Maxwell-Solvers> (53).

ACKNOWLEDGMENTS. This work is funded by NSF under Award No. 2103301 and the Air Force Office of Scientific Research under Award No. FA9550-25-1-0373.

28. I. Luz, M. Galun, H. Maron, R. Basri, I. Yavneh, “Learning algebraic multigrid using graph neural networks” in *International Conference on Machine Learning* (PMLR, 2020), pp. 6489–6499.
29. B. Lerer, I. Ben-Yair, E. Treister, Multigrid-augmented deep learning preconditioners for the Helmholtz equation using compact implicit layers. *SIAM J. Sci. Comput.* **46**, S123–S144 (2024).
30. A. Heinlein, A. Klawonn, M. Lanser, J. Weber, Combining machine learning and adaptive coarse spaces—A hybrid approach for robust FETI-DP methods in three dimensions. *SIAM J. Sci. Comput.* **43**, S816–S838 (2021).
31. A. Heinlein, A. Klawonn, M. Lanser, J. Weber, Machine learning in adaptive domain decomposition methods—Predicting the geometric location of constraints. *SIAM J. Sci. Comput.* **41**, A3887–A3912 (2019).
32. A. Heinlein, A. Klawonn, M. Lanser, J. Weber, “Predicting the geometric location of critical edges in adaptive GDSW overlapping domain decomposition methods using deep learning” in *Domain Decomposition Methods in Science and Engineering XXVI*, S. C. Brenner *et al.*, Eds. (Springer, 2023), pp. 307–315.
33. Y. Azulay, E. Treister, Multigrid-augmented deep learning preconditioners for the Helmholtz equation. *SIAM J. Sci. Comput.* **45**, S127–S151 (2022).
34. C. Cui, K. Jiang, S. Shu, A neural multigrid solver for Helmholtz equations with high wavenumber and heterogeneous media. *arXiv [Preprint]* (2024). <https://doi.org/10.48550/arXiv.2404.02493> (Accessed 12 April 2026).
35. Y. Xie, M. Lv, C. S. Zhang, “MGCFNN: A neural multigrid solver with novel Fourier neural network for high wave number Helmholtz equations” in *Thirteenth International Conference on Learning Representations* (2025).
36. T. Knoke *et al.*, Domain decomposition with neural network interface approximations for time-harmonic Maxwell’s equations with different wave numbers. *arXiv [Preprint]* (2023). <https://doi.org/10.48550/arXiv.2303.02590> (Accessed 12 April 2026).
37. S. Piao, H. Gu, A. Wang, P. Qin, A domain-adaptive physics-informed neural network for inverse problems of Maxwell’s equations in heterogeneous media. *IEEE Antennas Wirel. Propag. Lett.* **23**, 2905–2909 (2024).
38. L. D. Dalcin, R. R. Paz, P. A. Kler, A. Cosimo, Parallel distributed computing using Python. *Adv. Water Resour.* **34**, 1124–1139 (2011).
39. V. Dolean, P. Jolivet, F. Nataf, *An Introduction to Domain Decomposition Methods: Algorithms, Theory, and Parallel Implementation* (SIAM, 2015).
40. F. Nataf, H. Xiang, V. Dolean, A two level domain decomposition preconditioner based on local Dirichlet-to-Neumann maps. *C. R. Math.* **348**, 1163–1167 (2010).
41. R. Haferss, P. Jolivet, F. Nataf, A robust coarse space for optimized Schwarz methods: SORAS-GenEO-2. *C. R. Math.* **353**, 959–963 (2015).
42. C. Farhat, A. Macedo, M. Lesoinne, A two-level domain decomposition method for the iterative solution of high frequency exterior Helmholtz problems. *Numer. Math.* **85**, 283–308 (2000).
43. M. Chen, J. Jiang, J. A. Fan, Design space reparameterization enforces hard geometric constraints in inverse-designed nanophotonic devices. *ACS Photon.* **7**, 3141–3151 (2020).
44. E. Gershnel *et al.*, Reparameterization approach to gradient-based inverse design of three-dimensional nanophotonic devices. *ACS Photon.* **10**, 815–823 (2023).
45. T. Dai *et al.*, Shaping freeform nanophotonic devices with geometric neural parameterization. *npj Comput. Mater.* **11**, 259 (2025).
46. J. Jiang, J. A. Fan, Global optimization of dielectric metasurfaces using a physics-driven neural network. *Nano Lett.* **19**, 5366–5372 (2019).
47. J. Jiang, J. A. Fan, Multiobjective and categorical global optimization of photonic structures based on ResNet generative neural networks. *Nanophotonics* **10**, 361–369 (2021).
48. J. Jiang, J. A. Fan, Simulator-based training of generative neural networks for the inverse design of metasurfaces. *Nanophotonics* **9**, 1059–1069 (2020).
49. J. Jiang *et al.*, Free-form diffractive metagrating design based on generative adversarial networks. *ACS Nano* **13**, 8872–8878 (2019).
50. R. Lupo *et al.*, A multi-agentic framework for real-time, autonomous freeform metasurface design. *Sci. Adv.* **11**, ead8006 (2025).
51. O. Schenk, K. Gärtner, W. Fichtner, A. Stricker, PARDISO: A high-performance serial and parallel sparse linear solver in semiconductor device simulation. *Fut. Gener. Comput. Syst.* **18**, 69–78 (2001).
52. T. W. Hughes, I. A. Williamson, M. Minkov, S. Fan, Forward-mode differentiation of maxwell’s equations. *ACS Photon.* **6**, 3010–3016 (2019).
53. C. Mao, Deep-Maxwell-Solvers. GitHub. <https://github.com/ChenkaiMao97/Deep-Maxwell-Solvers>. Deposited 13 April 2026.