

High Speed Simulation and Freeform Optimization of Nanophotonic Devices with Physics-Augmented Deep Learning

Mingkun Chen,[#] Robert Lupoiu,[#] Chenkai Mao,[#] Der-Han Huang,[#] Jiaqi Jiang, Philippe Lalanne, and Jonathan A. Fan^{*}



Cite This: *ACS Photonics* 2022, 9, 3110–3123



Read Online

ACCESS |



Metrics & More



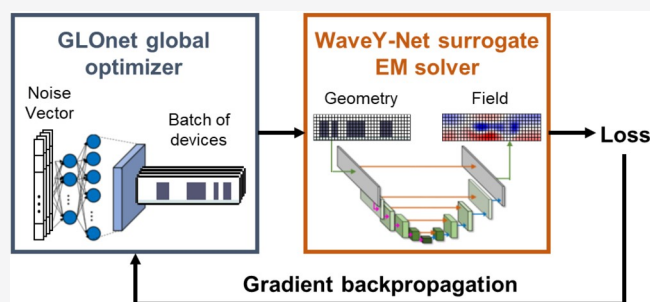
Article Recommendations



Supporting Information

ABSTRACT: We introduce WaveY-Net, a hybrid data- and physics-augmented convolutional neural network that can predict electromagnetic field distributions with ultrafast speeds and high accuracy for entire classes of dielectric nanophotonic structures. This accuracy is achieved by training the neural network to learn only the magnetic near-field distributions of a system and to use a discrete formalism of Maxwell's equations in two ways: to calculate electric fields from the magnetic fields and as physical constraints in the loss function. We show that WaveY-Net can accurately predict the near-fields in periodic, high dielectric contrast nanostructure arrays, and that it can combine with gradient-based algorithms to dramatically accelerate the local and global freeform optimization of diffractive photonic devices by orders of magnitude faster speeds. We anticipate that physics-augmented deep neural networks will transform the practice of nanophotonics simulation and design.

KEYWORDS: *nanophotonics, inverse design, simulation, freeform optimization, neural networks, deep learning*



INTRODUCTION


Maxwell simulators are essential tools for the characterization and design of electromagnetic systems. These systems operate at frequencies spanning the radiowave to X-ray and include a diversity of antennas,^{1–3} diffractive surfaces,^{4–7} metamaterials,^{8–10} and guided wave-based photonic circuits.^{11–13} Among the most popular frequency domain Maxwell solvers are the finite element method (FEM)^{14,15} and finite difference frequency domain (FDFD) algorithms.^{16–19} In both algorithms, the system domain is subdivided into discrete voxels, and the simulator solves for electromagnetic fields by constructing and inverting a sparse matrix with dimensions proportional to the total number of voxels. While these methods can be used to accurately solve general electromagnetics problems, the time and computation cost of matrix inversion serve as practical bottlenecks for the simulation of large domains and for design tasks, where large numbers of electromagnetic simulations are required to perform iterative optimization.

To address this bottleneck, deep neural networks serving as high speed surrogate Maxwell simulators have emerged as promising algorithms that can operate orders-of-magnitude faster than conventional Maxwell simulators.^{20–26} Many initial attempts to use neural networks in this manner were based on the training of fully connected deep neural networks in an end-to-end manner, where the inputs to the network were geometric parameters describing a photonic structure and the output was a physical performance parameter or spectral response.^{27–31}

While these concepts have served as accurate modeling tools for systems described by a handful of geometric parameters, their extension to more complex, higher dimensional problems has proven challenging. First, while there are clear physical relationships between the inputs and outputs of these networks, these relations are complicated and are not bounded by explicit physical constraints. More generally, the networks do not explicitly utilize spatial relationships involving the device geometries or optical response. Second, there is often a large dimensionality mismatch between the input and output data forms in these networks, which adds difficulty to the network training process. Direct scaling of these concepts to more complex systems would simply not work or require an impractical amount of data for training.

To model high-dimensional data forms, such as electromagnetic field distributions within a photonic structure, modifications to these deep learning approaches have been proposed. In one example, dimensionality-reduced forms of the fields were trained in conjunction with a fully connected deep network to map metasurface geometry to field distribution.³² In

Received: June 9, 2022
Revised: August 8, 2022
Accepted: August 8, 2022
Published: August 22, 2022



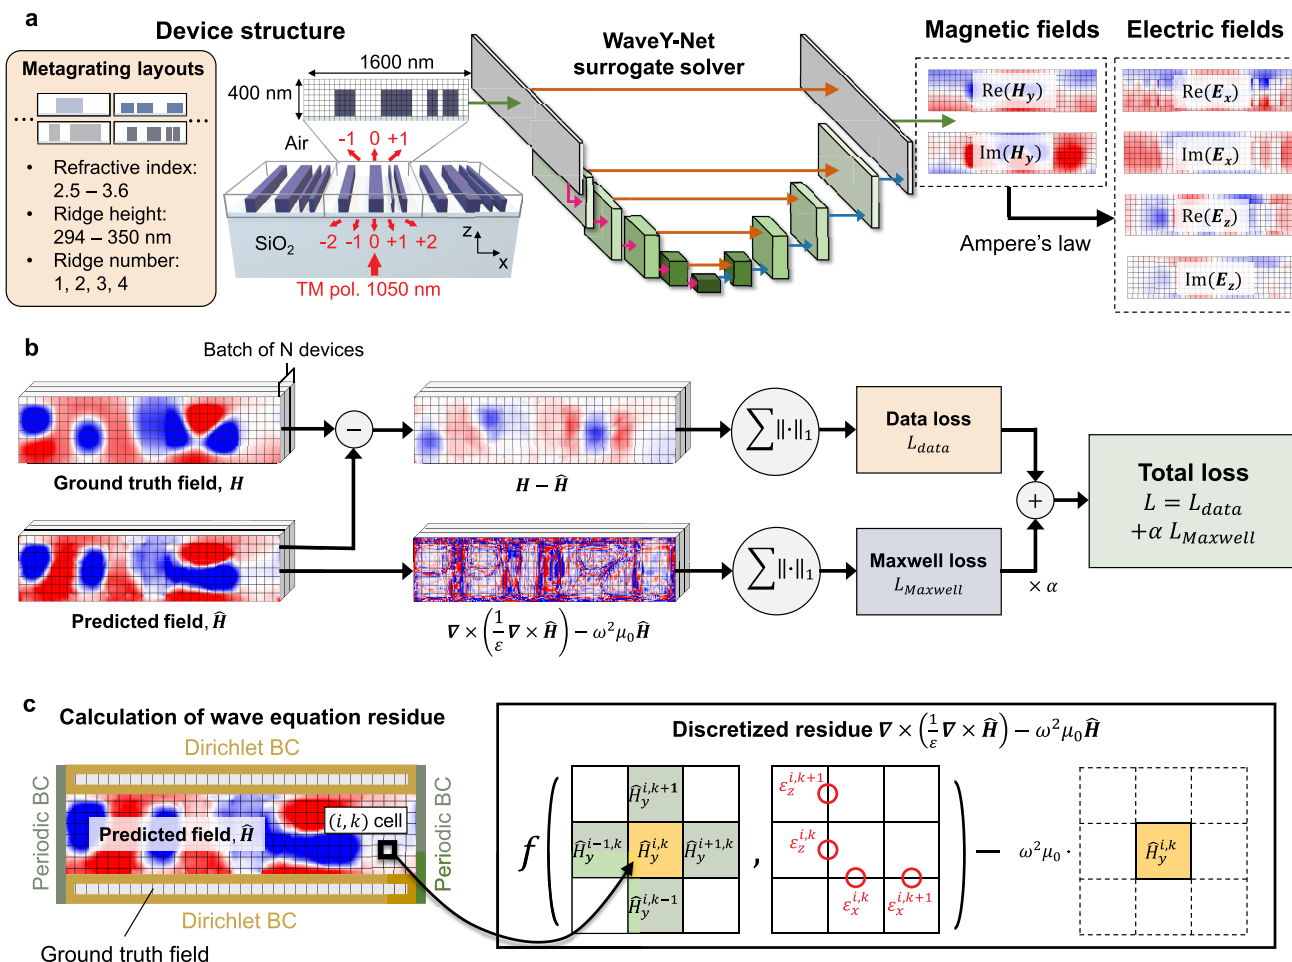


Figure 1. Overview of the WaveY-Net network architecture and training procedure. (a) WaveY-Net trains the UNet architecture to learn the electromagnetic behavior of dielectric-based metagratings (left). The input is an image of a single unit cell of the metagrating (center-right), and the outputs are magnetic field maps (center-right). Electric field maps are calculated from the magnetic field maps using a discrete version of Ampere's law (right). The devices being studied are nanoridge-based freeform metagratings with different grating materials, varying number of nanoridges, and varying ridge heights. Orange arrows: shortcut connections. Pink arrows: periodic convolution and maxpooling operations. Blue arrows: periodic convolution and upsampling operations. (b) Computational graph of the loss function. The loss function comprises two terms: data loss, which quantifies deviations between ground truth and predicted magnetic field maps, and Maxwell loss, which quantifies the deviation of the predicted magnetic field maps from Maxwell's equations. (c) Calculation of Maxwell loss. Periodic boundary conditions along vertical boundaries and Dirichlet boundary conditions along horizontal boundaries enforce a well-posed loss expression. Maxwell loss within the magnetic field maps is calculated using a discretized version of the wave equation residue, which imposes relationships between magnetic field values at a given voxel and those of its nearest neighbors.

another example, convolutional neural networks (CNNs) were shown to be effective at predicting electromagnetic field distribution maps within a photonic structure.³³ While these demonstrations pointed to the potential of neural networks as simulators, they were exclusively trained using data with no knowledge of Maxwell's equations,³⁴ placing limits on their ability to process and learn wavelike electromagnetic phenomena.

To expand on the capabilities of deep neural networks for physical science problems, techniques have been developed to either augment or completely replace the data-driven loss maps used for training with physics-driven loss maps generated using the governing equations of the system. In recent work, the concept of physics-only training for electromagnetics was successfully demonstrated for the modeling of low-contrast refractive optical systems.²⁶ More generally, physics-informed neural networks (PINNs) have been proposed, which explicitly incorporate physical constraints, such as a governing differential

equation, into the loss function.^{35–39} These concepts have been developed in the field of fluid mechanics and have been adapted to fully connected and CNN architectures.^{40–43} More recently, they have been utilized to solve vectorial electromagnetics problems, where an analytic form of Maxwell's equations and boundary conditions were used to constrain the loss function.^{44–46} The explicit incorporation of physics into the network training process produced impressive simulation and inverse design demonstrations. However, these methods that solve differential equation problems through network training are resource intensive to run, and they involve the solving of individual problems and not classes of problems. To be more specific, PINNs tasked with solving PDEs undergo a whole neural network training process to solve a single partial differential equation problem with predefined equation parameters, with the training process taking tens of minutes to hours.⁴² Furthermore, in our case a separate PINN would need to be trained for each device in the training set, as they have

unique sets of positions for the dielectric ridge-air boundaries that can only be learned by a single network. As such, PINNs are less practical for quickly solving classes of differential equation problems.

In this Article, we introduce Maxwell surrogate simulators based on a hybrid data- and physics-augmented training approach. We term these networks WaveY-Nets, and they combine the generalization and high-speed solving capabilities of data-only training with explicit Maxwell equation constraints from physics-informed training. To maximize network accuracy and capacity, we train the networks to output the magnetic field maps of a system and calculate electric field maps from the magnetic fields using Ampere's law. This approach follows FEM and FDTD formalisms, which solve for one field type and use Maxwell's equations to calculate the other field type,^{14,16,19} and it takes advantage of the fact that the electric and magnetic field distributions have the same information content and do not need to be independently learned. We also show that these high speed surrogate solvers can be used in conjunction with established gradient-based optimization algorithms to perform local and global freeform nanophotonic inverse design. Unlike concepts that attempt end-to-end inverse design solely through the training or evaluation of deep neural networks,^{47–52} our method leverages the efficacy of known gradient-based methods to accelerate inverse design in a stable manner. As a model system, we consider surrogate simulators that apply to periodic arrays of dielectric nanoridges comprising variable nanoridge heights, material refractive index, and topology. We show these surrogate simulators can accelerate electromagnetic simulations by orders of magnitude speeds and can be directly used in local and global adjoint optimization algorithms for designing metagratings with selective light diffraction capabilities.

We note that aspects of our concept share superficial similarities with developed ideas in the PINNs literature: PINNs can also utilize a combination of data and physics-based loss to solve differential equations, and they can be configured to solve forward and inverse physical science problems. However, we emphasize that our method is conceptually distinct from PINNs in the following ways. First, PINNs take spatial coordinates as inputs and solve for continuous field representations while WaveY-Nets take the dielectric distribution as inputs and solve for discrete field representations. Second, an individual PINN is trained to solve for fields for a single fixed dielectric distribution configuration. On the other hand, a single WaveY-Net is trained to solve for a class of configurations. That is, WaveY-Nets perform operator learning of Maxwell equations.⁵³ Third, PINNs can take tens of minutes to hours to perform forward and inverse design,⁴² which is impractically long. WaveY-Nets, on the other hand, can perform a forward simulation in milliseconds and be used to perform gradient-based inverse design in seconds, reducing the time for these calculations by orders of magnitude.

RESULTS

Network Architecture and Loss Formulation. The diffractive system captured by our surrogate model consists of dielectric nanoridges illuminated by a normally incident beam with a wavelength of 1050 nm (Figure 1a). The wave has transverse magnetic (TM) polarization, such that the electromagnetic field distribution is fully described by field components H_y , E_x , and E_z . The dielectric nanoridges are situated on a silicon dioxide substrate and have a period of 1600 nm, producing a total of eight diffraction orders. The full simulation window is

defined on a 64×256 pixel grid with a window height of 400 nm, and it covers a single device period with thin substrate and superstrate regions below and above the device, respectively. This discrete pixel representation sets the spatial resolution of the device layout and electromagnetic fields to be 6.25 nm along both the horizontal and vertical directions. The surrogate model applies to freeform diffractive structures comprising varying materials, number of nanoridges per unit cell, and ridge heights. The refractive index of the structures is selected from the continuous range of 2.48 to 3.57, which overlaps with popular high contrast metasurface materials including anatase titanium dioxide ($n = 2.48$),⁵⁴ gallium phosphide ($n = 3.11$),⁵⁵ indium phosphide ($n = 3.29$),⁵⁶ gallium arsenide ($n = 3.45$),⁵⁷ and polycrystalline silicon ($n = 3.57$)⁵⁸ (all the refractive indices are taken at a wavelength of 1050 nm and only the real part of the refractive index is considered as these materials are lossless at the wavelength of interest). The grating period contains from one to four total dielectric ridges and the ridge height ranges from 293.75 to 350 nm with step sizes of 6.25 nm. The device landscape contains all devices in which the dielectric ridges and air spacers have widths of 62.5 nm or greater, and it corresponds to approximately three trillion unique device configurations for a given device refractive index.

Our network scheme is outlined in Figure 1a and Figure S1 and is based on the UNet architecture,⁵⁹ which is a CNN architecture where the input and output data structures are images with the same dimension. Shortcut connections strengthen relationships between the input and output images. UNets were initially popularized in the computer vision community for image processing tasks such as image segmentation,^{59–62} and they are particularly well suited for our application because there exists a strong correspondence between input geometry and output field distribution. These networks were the basis for the data-driven surrogate electromagnetic simulators reported in ref 33. Our network input is a 64×256 pixel image of the simulation window, where the input image pixel values are relative permittivity values of the corresponding material. The output is 64×256 pixel images of the real and imaginary H_y field maps within the structure. The complex E_x and E_z field maps are calculated from H_y using Ampere's law.

We name our network WaveY-Net because it uses Maxwell's equations based on the Yee grid formalism to enforce wavelike field behavior of the output. The Yee grid is an established framework for finite-difference time and frequency domain electromagnetic simulations, and it is formulated to rigorously specify spatial relationships between discrete field components, boundary conditions at dielectric discontinuities, and discrete derivative operations.⁶³ A schematic of the two-dimensional TM Yee grid is in Figure S3. To summarize, the H_y fields are placed at the center of each pixel and parallel electric field components are placed at the pixel boundaries. Permittivity profiles are calculated for each field component to account for the pixel-level spatial offset defining each component. With the Yee grid, the discrete formulation of Ampere's law used to calculate E_x and E_z from H_y is as follows:

$$\frac{H_y^{i,k} - H_y^{i,k-1}}{\Delta z} = -i\omega \epsilon_x^{i,k} E_x^{i,k} \quad (1)$$

$$\frac{H_y^{i,k} - H_y^{i-1,k}}{\Delta x} = i\omega \epsilon_z^{i,k} E_z^{i,k} \quad (2)$$

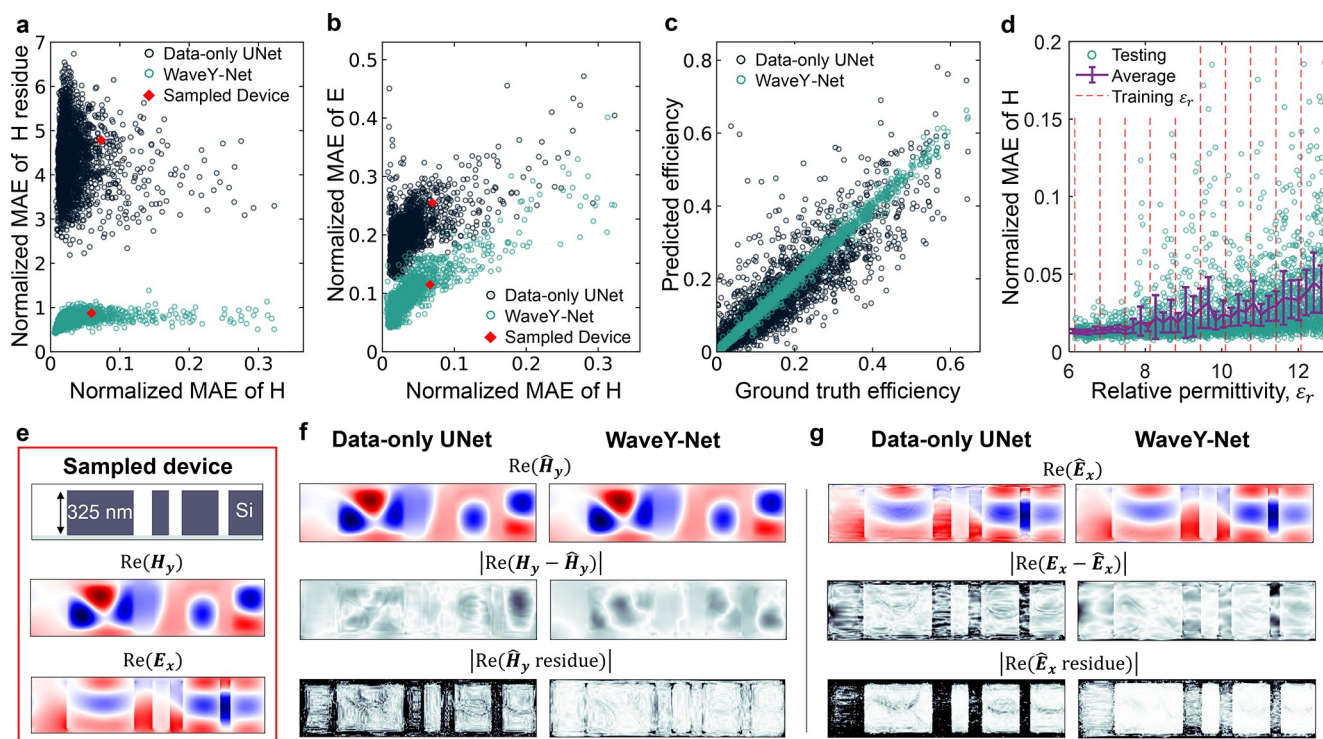


Figure 2. Benchmark comparison between WaveY-Net and a data-only UNet, both which predict magnetic field maps. (a) Scatter plot of magnetic wave equation residue MAE vs magnetic field MAE. The predicted magnetic field maps are from 3000 test devices evaluated by each network. (b) Scatter plot of the electric field MAE versus magnetic field MAE for the same devices. The definition of normalized MAE is in [Methods](#). (c) Scatter plot of the predicted versus ground truth far-field diffraction efficiencies for the same devices, for light diffracting into the transmitted +1 order. (d) Scatter plot of the magnetic field MAE versus relative permittivity of the grating material. Error bar stands for standard deviation. Red dashed lines mark 11 permittivity values seen in the training set. (e) Device layout, H_y field map, and E_x field map of a representative, randomly sampled device from the test set. (f) Magnetic and (g) electric field maps, field MAE maps, and wave equation residue MAE maps for the sampled device evaluated by both networks. H residue: wave equation residue of the magnetic field. E residue: wave equation residue of the electric field. Color scales: H-field [-0.015 (blue), 0.015 (red)], H-field MAE [0 (white), 0.001 (black)], H residue [0 (white), 0.01 (black)], E-field [-3 (blue), 3 (red)], E-field MAE [0 (white), 0.2 (black)], E residue [0 (white), 20 (black)].

where i and k are discrete index labels for the horizontal and vertical pixel positions, respectively.

To train WaveY-Net, the loss function is specified to have two components ([Figure 1b](#)):

$$L = L_{\text{data}} + \alpha L_{\text{Maxwell}} \quad (3)$$

where L_{data} is data-driven loss from which the network attempts to fit the network output fields with a ground truth training set. It takes the form of mean absolute error (MAE):

$$L_{\text{data}} = \frac{1}{N} \sum_{n=1}^N \|\mathbf{H}^{(n)} - \hat{\mathbf{H}}^{(n)}\|_1 \quad (4)$$

where \mathbf{H} represents ground truth field profiles from the training set, $\hat{\mathbf{H}}$ are the field profiles outputted from the network, N is a given batch size, and n is the index of the device within a batch. L_{Maxwell} specifies the compliance of the outputted fields with Maxwell's equations and is the MAE of the magnetic field wave equation residue:

$$L_{\text{Maxwell}} = \frac{1}{N} \sum_{n=1}^N \left\| \nabla \times \left(\frac{1}{\epsilon^{(n)}} \nabla \times \hat{\mathbf{H}}^{(n)} \right) - \omega^2 \mu_0 \hat{\mathbf{H}}^{(n)} \right\|_1 \quad (5)$$

The discrete Yee grid-based formalism of the magnetic field wave equation is [eq S1 in the Supporting Information, Section 2](#). If the calculated Maxwell loss for a given pixel in an outputted field profile is zero, it means that the fields local to that pixel are consistent with Maxwell's equations. α is a hyperparameter that

balances the contributions of data and Maxwell loss and it is dynamically tuned during network training in a manner that stabilizes the training process. More details pertaining to the network architecture and training methodology are in the [Methods Section](#) and [Supporting Information, Sections 1 and 3](#).

While L_{data} and L_{Maxwell} both reduce to zero as the outputted fields converge to ground truth values, each loss term captures different information about the fields. L_{data} captures the accuracy of the fields on a pixel-by-pixel basis, ensuring that the outputted fields are close to ground truth values but without explicitly accounting for spatial relationships between pixels. L_{Maxwell} , on the other hand, utilizes discrete spatial derivatives that explicitly capture physical relationships between neighboring field pixels. While CNN kernels are able to implicitly learn spatial electromagnetic field relationships through data-based training, L_{Maxwell} more directly imposes spatial constraints and relationships between neighboring pixels, as defined explicitly with Maxwell's equations. These spatial correlations can be visualized by the calculation of L_{Maxwell} for a given pixel, shown in [Figure 1c](#), which has the form of a 3×3 pixel kernel. As a kernel, L_{Maxwell} can be efficiently calculated for every pixel in a field map using a graphical processing unit (GPU), which naturally processes convolution operations in a fast and parallel manner. With the complementary nature of these loss functions, it is possible for L_{data} to be low and L_{Maxwell} to be high if the outputted field values are approximately correct but there are large field fluctuations between neighboring pixels. It is also possible for L_{data} to be high

and L_{Maxwell} to be low if the local fields are wavelike but dissimilar to the ground truth values.

To ensure that L_{Maxwell} is well posed, it is essential that the loss function incorporates proper boundary conditions. Without proper boundary conditions, L_{Maxwell} can push the outputted fields to wavelike profiles that locally satisfy Maxwell's equations but are far from ground truth values. At the left- and right-most columns of the field maps, we naturally apply periodic boundary conditions within the UNet framework by applying periodic padding to the convolutional calculations along the x -axis (Figure 1c). Dirichlet boundary conditions at the horizontal field map boundaries are enforced by substituting the top and bottom rows of the predicted field maps with ground truth magnetic field values (see a more detailed explanation in the Supporting Information, Section 4). In this manner, L_{Maxwell} is zero at all pixels only when the fields converge exactly to ground truth values.

WaveY-Net Solver. To evaluate the impact of L_{Maxwell} on UNet simulator accuracy, we train WaveY-Net, which trains using both L_{data} and L_{Maxwell} , and compare it with a data-only UNet trained with only L_{data} . In both cases, the neural networks use 30 000 random device layouts and their associated fields as training data, have outputs consisting of the real and imaginary magnetic field maps, and use eqs 1 and 2 to calculate the electric field maps. Training data are generated using an open source FDFD solver.¹⁷ More details pertaining to the data set preparation are in the Methods Section and Supporting Information, Section 1. A summary of the performance of both networks, compiled from 3000 test data, is presented as scatter plots in Figures 2a-2c and Table 1. We find that WaveY-

the data-only UNet, with an approximately six times difference in L_{Maxwell} MAE between the two networks (Figure 2b). This consistency of the magnetic field maps with the magnetic field wave equation leads to more accurate calculations of the electric fields, with WaveY-Net producing electric fields with average MAE values over twice smaller than those from the data-only UNet and with significantly improved consistency with the electric field wave equation (Table 1 and Supporting Information, Section 13).

To further quantify the utility of L_{Maxwell} , we perform near-to-far-field transformations⁶⁴ on the electric fields produced by each network and calculate diffraction efficiencies into the transmitted +1 order. Accurate far-field amplitudes and phases are required for tasks such as local and global freeform optimization and will be utilized later in this study. More details pertaining to the near-to-far-field calculation are in the Methods Section and Supporting Information, Section 5. Scatter plots of the predicted diffraction efficiencies compared to the ground truth for fields generated by the two networks are presented in Figure 2c. For field plots generated by the data-only UNet, there are clear deviations between the predicted and ground truth diffraction efficiencies, with an average efficiency error of 4.3%. WaveY-Net, on the other hand, outputs fields that produce relatively accurate near-to-far-field efficiency calculations, with an average efficiency error of only 0.68%. This error does not increase for high efficiency devices in spite of the fact that there is a disproportionately low number of high efficiency devices in the network training set. As such, Maxwell regularization is effective at predicting usual figures of merit, for example, efficiency, by enforcing long-range consistency through local field-map constraints.

The generalization ability of WaveY-Net as a surrogate Maxwell solver is noteworthy and captured in its ability to accurately process different nanoridge refractive indices. This analysis is shown in Figure 2d and Supporting Information, Figure S26. While WaveY-Net is trained with data that only contains 11 discrete permittivity values (red dashed line in Figure 2d), the network is able to interpolate and predict consistently accurate electromagnetic fields for devices comprising permittivity values unseen in the training set. In fact, the accuracy of the magnetic field maps outputted by WaveY-Net for interpolated and training set device refractive index values are nearly identical. We note that the field MAE increases as the permittivity value increases, which is as expected because the complexity of the electromagnetic field distributions increases as a function of material index, but that the average field MAE for high index devices is nonetheless only approximately 4%. A more extensive and detailed discussion of the generalization ability of WaveY-Net is provided in Supporting Information, Section 14.

To more clearly visualize the discrepancy between the fields produced by both networks, we examine the electromagnetic fields outputted by each network for a randomly selected test device which has four silicon nanoridges within the unit cell with a ridge height of 325 nm (Figure 2e). Ground truth $\text{Re}(E_x)$ and $\text{Re}(H_y)$ field maps are also shown, with the other field components presented in the Supporting Information, Section 6. The outputted magnetic field maps from both networks are in Figure 2f and display good agreement with the ground truth field map. Their corresponding MAE maps show pixel-level deviations that each feature approximately 3% relative error. However, the MAE map from the data-only UNet has a higher nonphysical spatial frequency noise component. This difference

Table 1. Summary of Average MAE Values and Associated Standard Deviation Values for the Test Data Analyzed in Figure 2a^a

normalized MAE		data-only UNet	WaveY-Net
$H - \hat{H}$	avg	0.032	0.029
	std	0.035	0.034
\hat{H} residue	avg	4.331	0.696
	std	0.687	0.117
$E - \hat{E}$	avg	0.193	0.083
	std	0.038	0.034
\hat{E} residue	avg	275.66	27.20
	std	54.15	5.23

^aHere, " $H - \hat{H}$ " and " $E - \hat{E}$ " refer to the MAE data losses of the predicted H and the subsequently calculated E fields, respectively. " \hat{H} residue" is the Maxwell Loss calculated using the discretized wave equation derived for the H field, as defined in eq 5. " \hat{E} residue" refers to the average of the MAE values of the residue maps calculated for both components of the E field using the discretized wave equations defined by eqs S12 and S13.

Net and the data-only UNet are reasonably accurate magnetic field surrogate solvers, with magnetic field MAE averages of 0.029 and 0.032, respectively (Figure 2a). As such, the addition of L_{Maxwell} produces a modest but not substantial improvement in predicted magnetic field accuracy, as defined on a pixel-by-pixel basis.

However, the fields outputted by WaveY-Net are more self-consistent with the magnetic field wave equation compared to

is attributed to the spatial filtering functionality of the L_{Maxwell} loss term, which enforces spatial derivative constraints between neighboring pixels and pushes the fields to have spatially smooth, wavelike forms. The benefits featured by Maxwell regularization carry over to the calculation of electric field, where the electric field maps produced from WaveY-Net are relatively smooth and low in error while the data-only UNet electric fields have visible high spatial frequency noise and regions featuring as high as 25% error (Figure 2g). Errors are particularly amplified at dielectric discontinuities, where the electric field components have discontinuities themselves, and they carry over to the electric field wave equation residue maps.

The acceleration in computation enabled by WaveY-Net, compared to a conventional full wave solver, is significant due to a combination of software and hardware features. A summary of the computation time required by a conventional FDFD solver¹⁷ and WaveY-Net for different numbers of simulations is shown in Figure 3. A single simulation is defined as the evaluation of the

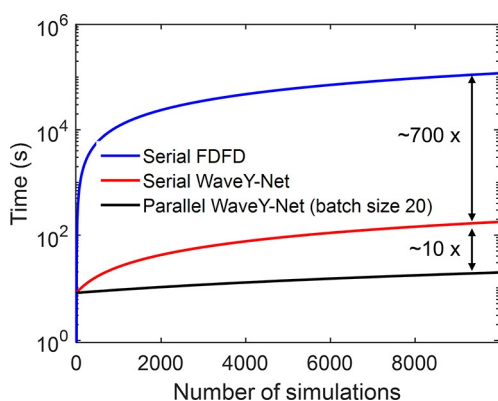


Figure 3. Plot of computation time vs number of simulations for three different simulation methods: a conventional FDFD solver; serial WaveY-Net, which simulates individual devices in series; and parallel WaveY-Net, which simulates devices in batches of 20 in parallel.

complex magnetic field maps for an individual device pattern. Computations for the conventional FDFD solver are performed with four cores of a 2.70 GHz Intel Xeon Gold 6150 CPU Processor with 32 GB of RAM, and those for WaveY-Net are performed with one NVIDIA A100 GPU with 40 GB of VRAM and PCIe connection. The serial WaveYNet solver provides a well over 2 orders of magnitude speedup than the FDFD solver and the parallel WaveYNet provides nearly a 4 orders of magnitude faster speedup in computation. For the evaluation of 7000 devices, the FDFD solver takes approximately 1 day while the serial WaveYNet, which evaluates one device at a time, take 3 min. The parallel WaveY-Net evaluates 20 devices at a time by taking advantage of the parallel computing capability of GPUs and can evaluate 7000 devices within 20 s. A more detailed comparison can be found in Supporting Information, Section 8. It is also noted that these speedups featured by WaveY-Net lead to significant reductions in environmental impact, which is further detailed in the Supporting Information, Section 11.

WaveY-Net-Based Freeform Optimizers. We next turn our attention to the utilization of WaveY-Net in design and optimization algorithms, which are ideal platforms for benchmarking high speed electromagnetic solvers because they require the iterative evaluation of distinct device structures and can require batches of devices to be evaluated at a given time. Our focus will be on the freeform boundary optimization

of dielectric metagratings that selectively diffract normally incident light to the transmitted +1 diffraction order. We consider two methods for optimization: gradient-based local optimization based on the adjoint variables method^{5,65–67} and global optimization based on GLOnets,^{67–69} which is a population-based evolutionary algorithm that performs optimization through the training of a generative neural network. Our approach for device optimization is fundamentally different from prior studies of photonic inverse design with deep discriminative networks, which are trained end-to-end and use methods such as backpropagation to perform optimization. As previously discussed, end-to-end networks have difficulty in generalizing to complex, high dimensional problems due to the lack of explicit regularization methods.^{27–29} Our neural network-based solvers, on the other hand, are much more capable of generalizing because they are explicitly constrained to relate geometries with electromagnetic fields that obey Maxwell’s equations. By combining generalized neural network-based solvers together with the adjoint variables method, our hybrid algorithm combines the accuracy and high speeds of WaveY-Net with the capabilities of freeform gradient calculations for optimization.

A computational graph of the local adjoint-based optimizer is shown in Figure 4a. At the core of this algorithm is the adjoint solver, which takes the device geometry as the input and outputs the diffraction efficiency and adjoint gradient of the device. The gradient is used to iteratively perturb the device geometry in a manner that improves device diffraction efficiency. We formulate a WaveY-Net surrogate adjoint solver that utilizes two separately trained WaveY-Nets. The first is a “forward” simulator that predicts the magnetic fields H_{fwd} given a normally incident plane wave. The second is an “adjoint” simulator that predicts the magnetic fields H_{adj} given an obliquely incident plane wave oriented oppositely to that of the transmitted diffracted beam (Figure 4a). The electric field maps E_{fwd} and E_{adj} are calculated from each corresponding magnetic field map using eqs 1 and 2. The calculated electric field maps from the forward simulator are also used to calculate the far-field amplitude and phase response in the desired diffraction order, using a near-to-far-field transformation, which yield diffraction efficiency and the adjoint phase. The adjoint gradient is calculated by integrating the forward and adjoint electric fields and the adjoint phase. More details are provided in Methods.

To benchmark the adjoint solver, we calculate the adjoint gradient maps for a random device using three types of simulators: a fullwave FDFD simulator, a data-only UNet, and WaveY-Net. The gradient maps are shown in Figure 4b and indicate that the gradients from the WaveY-Net match well with the ground truth values computed with the FDFD solver. On the contrary, the adjoint gradients calculated using the data-only UNets exhibit large errors in magnitude and even sign, indicating these solvers are not sufficiently accurate for use in optimization (more details are provided in the Supporting Information, Section 9).

To evaluate the efficacy of the WaveY-Net adjoint solver to perform optimization, we use this solver and the FDFD-based adjoint solver to perform full gradient optimizations on silicon metagratings with four nanoridges within the unit cell and a ridge height of 325 nm. Figure 4c summarizes the results for representative optimization runs where the same starting device geometry and 100 iterations are used for each solver. The algorithm using the FDFD solver shows a nearly monotonic increase in efficiency over the course of the iterative process,

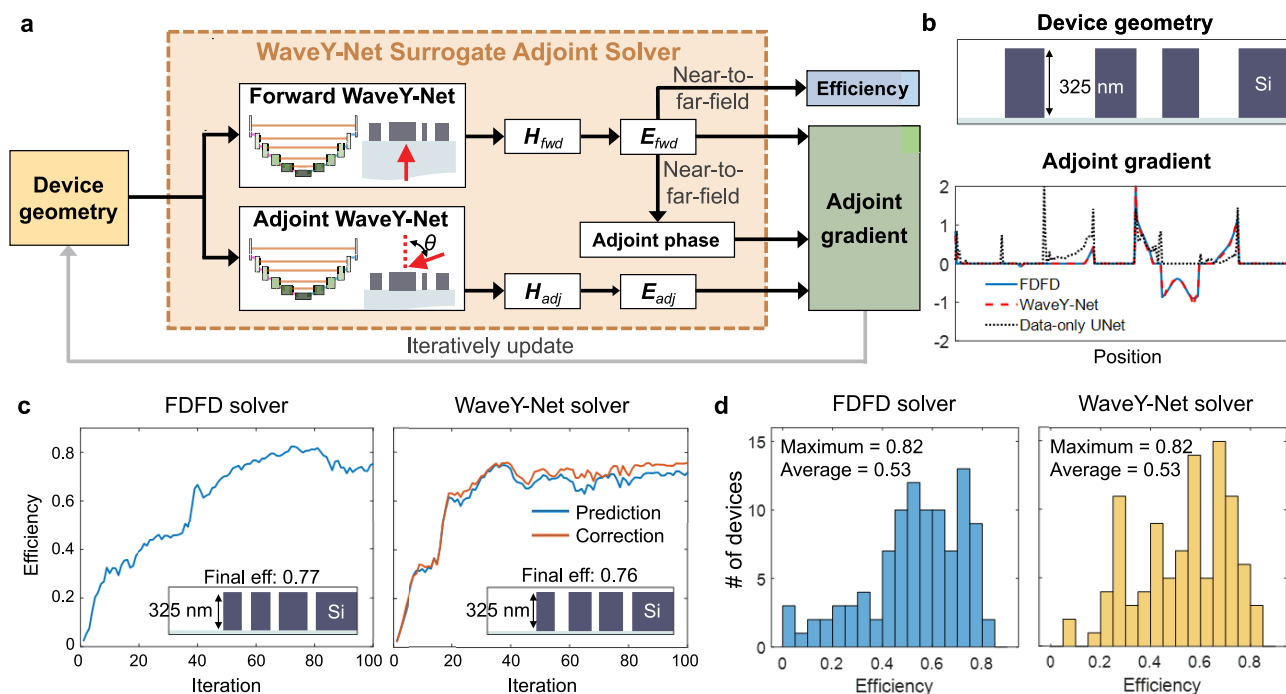


Figure 4. Local freeform metagrating optimization based on a WaveY-Net surrogate adjoint solver. (a) Computational graph of the local adjoint optimizer that maximizes device diffraction efficiency into the transmitted +1 diffraction order. Two WaveY-Nets are used to perform forward and adjoint magnetic field simulations, which are used to calculate diffraction efficiency and the adjoint gradient. (b) A representative, randomly sampled device layout and the corresponding adjoint gradient calculated using three different electromagnetic solvers: an FDFD solver that produces ground truth gradients, WaveY-Net, and a data-only UNet. The gradients calculated using the FDFD solver and WaveY-Net are nearly identical. (c) Optimization trajectories for local adjoint optimizations performed using an FDFD solver and WaveY-Net. The blue efficiency curve in the WaveY-Net plot is predicted by WaveY-Net while the orange efficiency curve is based on FDFD simulations. The insets show the final device layouts and efficiencies. (d) Histograms of 100 locally optimized metagratings with calculations performed using the FDFD solver and WaveY-Net. Devices being optimized have four silicon nanoridges within the unit cell with a ridge height of 325 nm.

with a final device exhibiting a 77% diffraction efficiency. The optimizer based on WaveY-Nets shows similar behavior and produces a final device exhibiting a 76% diffraction efficiency, indicating that the surrogate adjoint solver is sufficiently accurate to perform iterative optimization. The final efficiency value cited here is calculated using a fullwave solver. Ground truth device efficiencies at each optimization iteration (orange curve) nearly match predicted values (blue curve), indicating that our near-to-far-field calculations of WaveY-Net-outputted fields yield consistently accurate diffraction efficiencies. The optimization curves do not match exactly, indicating that small amounts of error in the gradient and efficiency calculations produce small deviations in the optimization trajectory. To further benchmark the consistency of our WaveY-Net adjoint solver, we perform 100 local optimizations with random starting devices using WaveY-Net and the FDFD simulator. Final device efficiency histograms for optimizers using each simulator are summarized in Figure 4d and show both distributions to have similar average and maximum efficiencies. We have also performed benchmark tests on optimization with other grating materials such as titanium dioxide and gallium phosphide and validated the effectiveness our WaveY-Net adjoint solver (Supporting Information, Figure S27).

The WaveY-Net-based adjoint solver can also be directly incorporated into the GLOnets algorithm. The optimization pipeline is shown in Figure 5a and involves three principle parts: the iterative generation of a batch of devices from a generative network, calculating the performance gradients and efficiencies of those devices to evaluate a loss function, and updating the

network weights in the generative network with backpropagation in a manner that minimizes the loss function. The generator G_ϕ contains a single fully connected layer followed by a batch normalization layer, where ϕ are the network weights, and it produces a distribution of devices \mathbf{x} from uniformly distributed noise vectors \mathbf{z} . To enforce a minimum feature size of 62.5 nm throughout the optimization process, we use a reparameterization transformation in which an analytic transformation converts network-outputted latent device representations to physical, constrained devices.⁶⁷ The loss function, formally derived in ref 68, is defined to be

$$L_G(\mathbf{x}, \mathbf{g}, \text{Eff}) = -\frac{1}{K} \sum_{k=1}^K \frac{1}{\sigma} \exp\left(\frac{\text{Eff}^{(k)}}{\sigma}\right) \mathbf{x}^{(k)} \cdot \mathbf{g}^{(k)} \quad (6)$$

The efficiencies, Eff, and adjoint gradients, \mathbf{g} , of the devices are calculated with the WaveY-Net adjoint solver. As the WaveY-Net adjoint solver relies on GPU hardware, it can evaluate full batches of devices in parallel. K is the generator batch size and σ is a hyperparameter for which 0.1 is used for optimal performance.

We run the WaveY-Net-GLOnet algorithm for the same metagrating example discussed above and also run an FDFD-based GLOnet as a ground truth benchmark. A total of 200 optimization iterations, each with a generator batch size of 100 devices, are used in both cases. The results are summarized in Figure 5b and show that WaveY-Net-GLOnet is able to converge to nearly the same optimal device as the FDFD-based GLOnet. The best sampled device has an efficiency of

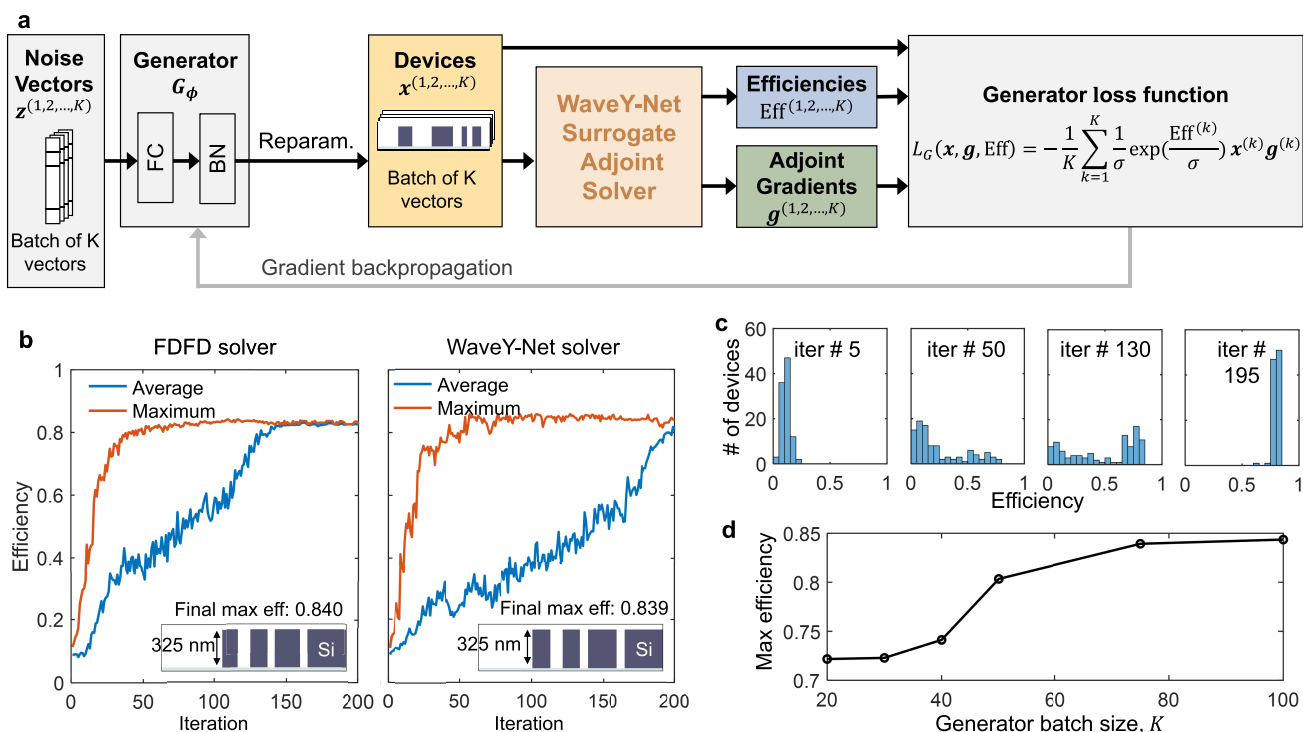


Figure 5. WaveY-Net-GLOnet algorithm for the population-based search of the global optimum. (a) Computational graph of the WaveY-Net-GLOnet algorithm. The WaveY-Net surrogate adjoint solver is the module featured in Figure 4a. To perform optimization, a batch of devices with latent space representations is produced by the generator and transformed into physical devices using reparameterization. Device efficiencies and adjoint gradients, computed by the adjoint solver, are utilized in a custom loss function to push the generated device distribution toward the global optimum. FC: fully connected layer. BN: batch normalization layer. (b) Optimization trajectories of GLOnet runs based on FDFD and WaveY-Net solvers. The final optimal device layouts and efficiency values, shown in the insets, are similar. Batch size is $K = 100$. (c) Histogram of device efficiencies as a function of iteration number from the WaveY-Net-GLOnet run. (d) Plot of final optimal device efficiency vs GLOnet training batch size obtained from WaveY-Net-GLOnet. Devices being optimized have four silicon nanoridges within the unit cell with a ridge height of 325 nm.




Figure 6. WaveY-Net simulator for the modeling and design of volumetric metamaterials. (a) Schematic of a unit cell of a three-layer silicon metamaterial modeled using WaveY-Net. (b) Efficiencies from a set of 11 multilayer devices, designed using WaveY-Net-accelerated local freeform optimization, which can diffract incident light to the +1 and -1 transmitted diffraction orders with tailored amplitudes in each order. (c) Poynting vector (S) distribution of Device 11 from (b), calculated from WaveY-Net field outputs. The curved Poynting vector trajectories are indicative of strong nonlocal light–matter interactions. Color scales: S_x [−0.003 (blue), 0.003 (red)], S_z [−0.004 (blue), 0.004 (red)], S [0 (white), 0.003 (orange)].

0.84, which is higher than the best locally optimized device in Figure 4d and indicative of an effective global search algorithm. Figure 5c shows histograms of generated device efficiencies as a function of iteration number for the WaveY-Net-GLOnet and show that the generator initially has no knowledge of high quality metagrating designs, but over the course of training, the best generated device and the overall device distribution shift toward higher efficiency regimes. A plot of the best device

efficiency obtained using WaveY-Net-GLOnet with different batch sizes (Figure 5d) shows that to effectively search for the global optimum, a batch size of at least 100 devices is ideal. The requirement of large batch sizes plays to the strengths of the UNet simulator approach, where the forward and adjoint simulations of entire batches are computed in parallel. We calculate that WaveY-Net-GLOnet runs approximately 7000

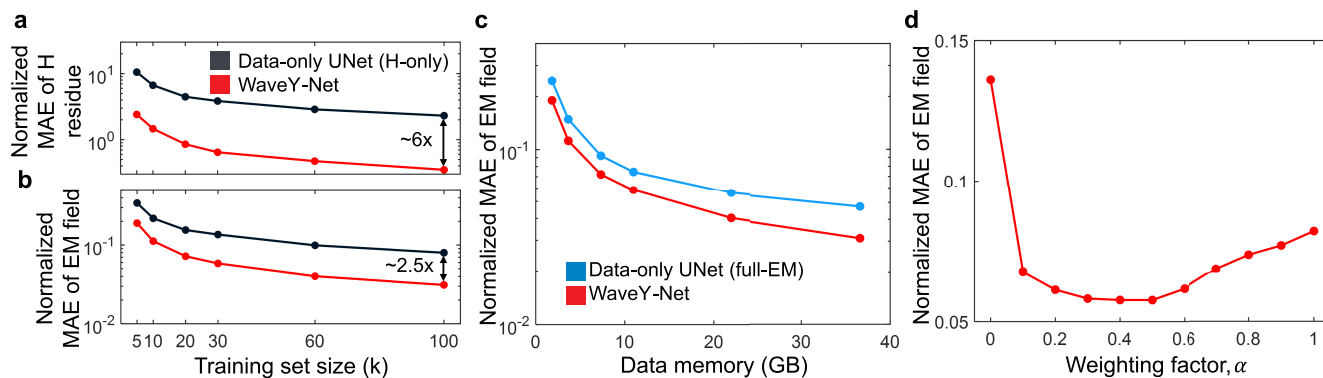


Figure 7. Benchmark numerical experiments with WaveY-Net. (a) Plot of magnetic wave equation residue MAE versus training set batch size, computed for WaveY-Net and a data-only UNet that outputs magnetic fields. (b) Plot of electromagnetic field MAE versus training set batch size, computed for WaveY-Net and a data-only UNet that outputs magnetic fields. (c) Plot of electromagnetic field MAE vs total training set memory size, computed for WaveY-Net and a data-only UNet that outputs the full electromagnetic fields. (d) Plot of electromagnetic field MAE versus the hyperparameter α , computed for WaveY-Net with normalized L_{Maxwell} and L_{data} loss terms.

times faster than FDFD-based GLOnet, which is consistent with the trends shown in Figure 3.

We further show that WaveY-Nets can readily generalize to more complex and higher dimensional photonic systems. For this demonstration, we consider multilayer metamaterials, which are an emergent class of volumetric structured media that utilize nontrivial nonlocal light-matter interactions to perform wavefront engineering tasks.^{70–72} The added degrees of design freedom are ideal for realizing devices exhibiting ultrahigh efficiencies and multifunctional optical responses. A schematic of our simulated system is in Figure 6a and shows a three layer silicon metamaterial with the bottom two layers embedded in glass. Such devices can be fabricated using a series of film deposition, patterning, and planarization steps.^{72,73} We implement a WaveY-Nets surrogate adjoint solver for this system and perform multiobjective freeform optimization of devices that diffract incident light to the +1 and –1 diffraction orders with tailored amplitudes in each order. The performances of a set of these devices, designed using local freeform optimization, are summarized in Figure 6b and show that the desired objective can be achieved with arbitrary choice of amplitudes in each diffraction order. The results from the WaveY-Nets-based optimizer effectively match those from a fullwave FDFD solver (see Supporting Information, Section 15). With the full electric and magnetic field maps from WaveY-Nets, we calculate the Poynting vector distribution for Device 11 in Figure 6b and observe that power flow through the device is highly nonlocal and represented by highly curved pathways within the full device volume (Figure 6c). These WaveY-Net-based Poynting vector maps match well with those calculated from an FDFD solver (Supporting Information, Section 15), demonstrating the capability of WaveY-Nets to produce accurate full field calculations. Remarkably, the WaveY-Nets used for multilayer devices have the exact same network structure and training set size as when it is used for single-layer devices even though the multilayer devices define a significantly higher dimensional problem. We attribute this dimension-robustness to the fact that WaveY-Net is learning the local relationship constrained by Maxwell equation between geometry and field response, which does not depend on the problem dimension.

DISCUSSION

We perform further WaveY-Net benchmarking with two types of data-only UNets: the previously discussed two channel H-only

network that outputs magnetic field maps and a six channel full-EM network that outputs field maps for all electromagnetic field components (see Supporting Information, Section 10 for more details). With the H-only data-only UNets, we previously observed that the lack of Maxwell regularization led to inaccurate electric field maps. In spite of these shortcomings, one might hypothesize that these field inaccuracies could be mitigated by simply increasing the data set size. This “big data” mentality arises from trends observed in computer vision and natural language processing tasks, where it is observed that models generalize better with less overfitting when larger training sets are used.^{74–77} To investigate the impact of training set size on network performance, we train data-only UNet and WaveY-Net models with the same network architecture on data sets with a total of 5k, 10k, 20k, 30k, 60k, and 100k random training devices. For the purposes of this benchmark task, we consider a more specialized training set comprising silicon devices with four nanoridges per unit cell and a height of 325 nm.

Comparisons of the magnetic wave equation residue and full field MAE as a function of training set size for each model is presented in Figure 7a,b. We find that WaveY-Nets produce magnetic wave equation residue values that are approximately 10× lower than those from the data-only UNets, independent of the training set size. Furthermore, the residue value from the data-only UNet trained with 100K data is similar to the residue value from WaveY-Nets trained with only 5K data (Figure 7a). Without explicit Maxwell regularization, data-only UNets are not able to learn wavelike correlations between neighboring pixels, even in the limit of large training sets. An examination of the full field MAE trends in Figure 7b indicates that even with 100K training data, the data-only UNet still has limited accuracy, with normalized field MAEs of approximately 10%. In addition, WaveY-Net requires approximately 10× less data than the data-only UNet to produce total electromagnetic field maps with similar MAE. More discussions are in the Supporting Information, Section 7. This relative reduction in training data preparation is critical for our application, where the generation of training data consumes the vast majority of computational resources used for algorithm development. We also note that the amount of training data and trade-off between training set size and field prediction accuracy are dictated by the final application, and we selected 30K as the training set size for this study as it gives low enough error for the fields as to enable accurate and stable device optimization runs. Using less training

data, such as 20K, does not meet the requirement for optimizers used in this study (see Supporting Information, Figure S14), but it might suffice for other applications that do not require as high accuracy of the field prediction.

A comparison of WaveY-Net with a full-EM data-only UNet shows that WaveY-Net has a performance advantage. When both networks are trained with 30 000 training devices, the normalized MAE of the full fields (see Methods), is 0.054 from WaveY-Net while the MAE values from the data-only network is 0.056. This slight advantage exists in spite of the fact that WaveY-Net uses nearly three times less total training data, as quantified by data memory. The origin of this advantage is clarified when examining the normalized MAE of the network-outputs magnetic fields, which is 0.033 for WaveY-Net and 0.046 for the full-EM data-only UNet. This relatively large difference arises because WaveY-Net learns data relationships for a lower dimensional problem (i.e., two output channels instead of six), allowing the trained network to display improved accuracy and generalizability. These trends suggest that the full WaveY-Net fields can be further improved, in a manner that far surpasses the capabilities of the full-EM data-only UNet, by improving the method for calculating electric fields from magnetic fields.

A more fair comparison, from the point of view of network training, is to benchmark both networks using the same amount of training data as quantified by data memory. Practical limits to maximum usable training set memory are bounded by hardware limitations such as disk capacity and GPU memory. In addition, the speed and computational resources used for network training, which involve data processing in CPUs and data transferring between GPUs and CPUs, directly correlate with training set memory. The results are summarized in Figure 7c and show a clear performance improvement with WaveY-Net. It indicates that for a fixed amount of training data memory, it is always advantageous to have a more diverse data set with less information contained by each data point and to use physical relationships to compute the rest of the information, instead of having the network learn all the information with a reduced data set size.

Finally, we examine the impact of the weighting factor α in eq 3 to further elucidate the network training process and the relative roles of L_{Maxwell} and L_{data} . On one hand, hybrid networks trained with a strong L_{Maxwell} weighting can be treated as physics-informed neural networks that predominantly train by solving differential equations but use data to help with network convergence. On the other hand, hybrid networks that use a strong L_{data} weighting can be treated more as conventional data-based networks that use Maxwell regularization to push the outputted data to be more wavelike. We train a series of WaveY-Nets in which L_{Maxwell} is normalized each iteration to have the same magnitude as L_{data} and α is fixed to a chosen number. The plot of the resulting full field MAE values for α ranging from 0 to 1 is shown in Figure 7d and indicates that the best performing networks use an α between 0.2 and 0.6. As such, WaveY-Net most effectively operates as a data-based network that uses physics to regularize the quality of outputted fields. This biasing toward data-based loss is reflected in our observation that while it is straightforward to effectively train a network only with L_{data} , the network does not properly converge when trained only with L_{Maxwell} (see Supporting Information, Section 3). Training methods that use stronger L_{Maxwell} weighting are of interest because their proper implementation may reduce the reliance of large training data sets. Concepts such as the incorporation of an

active weighting scheme for boundary condition contributions may improve the performance of those networks,⁷⁸ and they will be a topic of future study.

In summary, we show that WaveY-Net, which trains using data and physical constraints, can serve as effective electromagnetic solvers. These surrogate simulators can produce accurate field solutions for classes of freeform devices comprising four silicon nanoridges, and they can be directly used in local and global optimization algorithms that search within this design space. An important feature of our approach is that it is data efficient, training with only a single field type and taking advantage of the explicit relationships between electric and magnetic fields fixed by Maxwell's equations. This feature enhances the generalization capabilities of the network through optimal use of the network capacity. This is a particularly important consideration when adapting UNets to large, three-dimensional systems, where the generation and utilization of large training data sets is extremely computationally intensive. While our network considers devices with fixed topology and material type, we anticipate that ensembles of WaveY-Net solvers can be collectively used to solve more general classes of photonics problems. We also anticipate that for applications that require accuracy convergence quantification, WaveY-Net can serve as a preconditioner for a rigorous Maxwell solver, providing a compromise between acceleration and accuracy.⁷⁹

METHODS

Network Architecture. WaveY-Net is implemented by using a traditional encoder-decoder UNet architecture consisted of six successive residual blocks, with each residual block containing six periodic-convolutional layers followed by batch normalization and a leaky rectifying linear unit (leaky ReLU). The periodic-convolutional layers use zero padding for horizontal boundaries and pads vertical boundaries using columns from the opposite side to account for periodic boundary condition. The number of convolutional kernels doubles after each pooling layer in the encoder, which is mirrored in the decoder. Two residual connections, each across three convolutional layers, are implemented within the residual block, which are proven to be beneficial for efficient optimization, as well as higher accuracy especially for deep networks.⁵⁹ For the first two encoding blocks, nonuniform maxpooling is used where the window size is (1×2) . Correspondingly, for the last two decoding blocks, nonuniform Upsampling is implemented. Shortcut connections are utilized between corresponding encoding and decoding blocks such that the last leaky ReLU layer of the encoder is concatenated to the input layers of the decoder. This has been proven to enhance the reconstruction of finer features.⁵⁹ Finally, the network produces the output fields with two channels (or six channels) for the real and imaginary part of the H_y field (or H_x , E_x , E_z fields). Detailed schematic of UNet structure is shown in the Supporting Information, Section 1.

Data Set Preparation. The neural network used to collect performance statistics and to perform local and global device optimization is trained using 27 000 examples and tested using 3000 examples. Both the training and test data sets are composed of dielectric device structures with varying permittivity of the dielectric metagrating, a variable number of nanoridges per unit cell, and a varying ridge height. The refractive index of the dielectric material of each device in the data set is one of 11 values, ranging from that of anatase titanium dioxide⁵⁴ to that of polycrystal silicon,⁵⁸ with the in-between

values incremented by a constant step-size in dielectric constant. The proportion of devices generated with each of the 11 possible refractive index values was sampled in a manner that is directly proportional to the permittivity constant. The number of nanoridges within the metagrating unit cell is 1, 2, 3, and 4, and the number of devices generated with each of these four values scales with the exponential increase in geometrical complexity with increasing nanoridge number. As such, 50, 2800, 10 150, and 17 000 devices were generated for 1, 2, 3, and 4 nanoridges, respectively. The 10 ridge heights are uniformly sampled from 293.75 to 350 nm with a step size of 6.25 nm.

The H_y^{fwd} and H_y^{adi} field responses of the dielectric metagrating unit cells were generated using the FDFD-based electromagnetic simulator Ceviche.¹⁷ Perfectly matched layers (PMLs)⁸⁰ are utilized in the vertical z direction, and periodic boundary conditions are adopted in the horizontal x direction. The metagrating device is placed at the center of the simulation domain, leaving at least a wavelength of space from the PML to ensure the integrity of its high absorbing performance. The lower half of the simulation domain is set as silicon dioxide, and the upper half is set as air. An infinitely large magnetic current sheet is used for field radiation to simulate a plane wave source. Only the fields and grating pattern shown in the red window in Figure S2 are saved to generate the labels for WaveY-Net training. The window is of size [256, 64] and contains five rows of substrate pixels at the bottom of the window and seven rows of pixels of air at the top. The simulation is set such that the phase of the x component of the incident electric field is zero at the lower edge of the window in forward simulation. For the adjoint simulation, the source is set to ensure the phase of the y component of the incident electric field is zero at the center point of the upper edge of the window. A more detailed discussion regarding the training set generation and Ceviche setup is in Supporting Information, Section 1.

Training Procedure. All the models are trained for 200 epochs in batch sizes of 32. Adaptive Moment Estimation (Adam)⁸¹ optimizer is used with an l_2 -regularization coefficient of 3×10^{-3} . A learning rate of 1×10^{-4} is applied, and an exponentially decaying learning rate scheduler is used with $\gamma = 0.98$. The convolutional layer weights are initialized by default using Kaiming Initialization.⁸² All the models are trained with PyTorch version 1.8.1. We adopt a train-test split ratio of 9:1 for all training processes.

Evaluation Metrics. Mean absolute error (MAE) is mainly used as the evaluation metric for this study. For an individual device, the normalized MAE of a certain field component is calculated by first computing the l_1 -norm of the difference between the predicted field matrix and its corresponding ground truth, and then it is normalized through dividing by the mean absolute magnitude of the ground truth field. Note that since the field is complex (2 channel), the normalized MAE is first computed for both the real part and the imaginary part, and the average value between the two is taken. The normalized MAEs of H_y , E_x , and E_z are all calculated in this manner. The normalized MAE of E is computed by taking the average between E_x and E_z . Similarly, the normalized MAE of the full electromagnetic field is computed by taking the average between all three field components. The normalized MAE of the wave equation residue is calculated in the same way except that the l_1 -norm is taken for the wave equation residue matrix computed from Supporting Information, eqs S1, S2, and S13. The normalized MAE of the entire set (Table 1 and Figure 7) is

evaluated by taking the mean value of the normalized MAE of each individual device within the data set.

Near-to-Far-Field Transformation. Given the near field profile of the electric field within the simulation window, we can take a horizontal slice of the E-field $E(x, z = z_1)$ and compute the far-field of the outgoing plane waves c_q as a function of diffraction channel q with the near- to far-field transformation as below (see Supporting Information, Section 5 for the derivation):

$$c_q = \frac{1}{a} \sum_x E(x, z_1) e^{ik_{x,\text{inc}}x} e^{-iqKx} e^{ik_{z,q}z_1}$$

where a is the grating period, $k_{x,\text{inc}}$ is the x -component of the incident wave vector, K is the spatial frequency of the grating computed as $K = \frac{2\pi}{a}$, and $k_{z,q}$ is the z -component of the outgoing wave vector of diffraction channel q . To calculate diffraction efficiency, the Poynting vector flux of the diffracted wave is integrated and divided by that of the incident field.⁸³

Adjoint Gradient Calculation. To calculate the adjoint gradient used in the adjoint variable method, the adjoint phase is first calculated given the forward electric field. More specifically, the electric field maps from the forward simulator are used to calculate the far-field response in the desired diffraction order, \tilde{E}_{fwd} , using the near- to far-field transformation. This far-field metric is then used to calculate diffraction adjoint phase, which is computed as $\theta_{\text{adj}} = \text{angle}(\tilde{E}_{\text{fwd}}^*)$, where \tilde{E}_{fwd}^* is the complex conjugate of \tilde{E}_{fwd} . The voxel-wise gradient is computed as $\text{Re}(\tilde{E}_{\text{fwd}} \cdot E_{\text{adj}} e^{i\theta_{\text{adj}}})$, which is a 64×256 matrix. It is then turned into a vector with dimension 1×256 by taking summation along the z -axis. It is cropped with a filtering operation where positive gradient is set to zero when the corresponding region is already the high-index material (dielectric ridge in this case), and negative gradient is set to zero when the corresponding region is already the low-index material (air in this case). Lastly, the magnitude is normalized by dividing by one-half of the maximal magnitude. The boundary gradients are calculated according to the methods described in refs 66 and 67.

■ ASSOCIATED CONTENT

SI Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acsphotonics.2c00876>.

Network architecture and training setup; background on the two-dimensional Yee grid formalism; discussing of the Maxwell loss; numerical Dirichlet boundary condition background; near-to-far-field formalism utilized; comparison of the entire electromagnetic field response prediction between the data-only UNet and the WaveY-Net; comparison of the efficiency prediction between the data-only UNet and the WaveY-Net for varying training data set sizes; computation time analysis for different solver types compared to WaveY-Net; optimization results from the data-only UNet; discussion on the direct prediction of the entire electromagnetic field response using the data-only UNet; energy consumption and environmental impact analysis of the WaveY-Net; comparison between WaveY-Net and Rigorous Coupled-Wave Analysis; error analysis on field conversions between H-field and E-field; metagrating generalization performance analysis of WaveY-Net; analysis and discussion on utilizing WaveY-Net for optimizing multilayer metamaterial devices; Figures S1–S27 (PDF)

AUTHOR INFORMATION

Corresponding Author

Jonathan A. Fan – Department of Electrical Engineering, Stanford University, Stanford, California 94305, United States; orcid.org/0000-0001-9816-9979; Email: jonfan@stanford.edu

Authors

Mingkun Chen – Department of Electrical Engineering, Stanford University, Stanford, California 94305, United States; orcid.org/0000-0001-9796-0021

Robert Lupoiu – Department of Electrical Engineering, Stanford University, Stanford, California 94305, United States; orcid.org/0000-0003-1649-2305

Chenkai Mao – Department of Electrical Engineering, Stanford University, Stanford, California 94305, United States

Der-Han Huang – Department of Electrical Engineering, Stanford University, Stanford, California 94305, United States

Jiaqi Jiang – Department of Electrical Engineering, Stanford University, Stanford, California 94305, United States; orcid.org/0000-0001-7502-0872

Philippe Lalanne – Laboratoire Photonique, Institut d'Optique Graduate School, Université Bordeaux, Bordeaux 33400, France; orcid.org/0000-0003-1979-2290

Complete contact information is available at:

<https://pubs.acs.org/10.1021/acsp Photonics.2c00876>

Author Contributions

[#]M.C., R.L., C.M., and D.-H.H. contributed equally to this study.

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

The authors acknowledge support from ARPA-E with Agreement Number DE-AR0001212, AFOSR with Agreement Number FA9550-18-1-0070, NASA with Agreement Number 80NSSC21K0220, National Science Foundation under Award 2103301, and the Packard Fellowship Foundation. Robert Lupoiu is supported by a graduate fellowship award from Knight-Hennessy Scholars at Stanford University. The authors thank Dr. Anton Ovcharenko for insightful discussion on electromagnetic field simulation methods. The data sets generated during and/or analyzed during the current study are available on MetaNet (<http://metanet.stanford.edu/search/waveynet-study/>) and are also available from the corresponding author upon request. The code that has been used to generate the results reported in this paper will be made available on GitHub (<https://github.com/jonfanlab/waveynet>) upon publication of this manuscript and is also available from the corresponding author upon request.

REFERENCES

- (1) Koziel, S.; Ogurtsov, S. *Antenna Design by Simulation-Driven Optimization*; Springer, 2014.
- (2) Hannan, P.; Balfour, M. Simulation of a phased-array antenna in waveguide. *IEEE transactions on Antennas and Propagation* **1965**, *13*, 342–353.
- (3) Sun, J.; Timurdogan, E.; Yaacobi, A.; Hosseini, E. S.; Watts, M. R. Large-scale nanophotonic phased array. *Nature* **2013**, *493*, 195–199.
- (4) Yu, N.; Capasso, F. Flat optics with designer metasurfaces. *Nature materials* **2014**, *13*, 139–150.
- (5) Sell, D.; Yang, J.; Doshay, S.; Yang, R.; Fan, J. A. Large-angle, multifunctional metagratings based on freeform multimode geometries. *Nano Lett.* **2017**, *17*, 3752–3757.
- (6) Phan, T.; Sell, D.; Wang, E. W.; Doshay, S.; Edee, K.; Yang, J.; Fan, J. A. High-efficiency, large-area, topology-optimized metasurfaces. *Light: Science & Applications* **2019**, *8*, 1–9.
- (7) Zhou, M.; Liu, D.; Belling, S. W.; Cheng, H.; Kats, M. A.; Fan, S.; Povinelli, M. L.; Yu, Z. Inverse Design of Metasurfaces Based on Coupled-Mode Theory and Adjoint Optimization. *ACS Photonics* **2021**, *8*, 2265–2273.
- (8) Gansel, J. K.; Thiel, M.; Rill, M. S.; Decker, M.; Bade, K.; Saile, V.; von Freymann, G.; Linden, S.; Wegener, M. Gold helix photonic metamaterial as broadband circular polarizer. *Science* **2009**, *325*, 1513–1515.
- (9) Valentine, J.; Zhang, S.; Zentgraf, T.; Ulin-Avila, E.; Genov, D. A.; Bartal, G.; Zhang, X. Three-dimensional optical metamaterial with a negative refractive index. *nature* **2008**, *455*, 376–379.
- (10) Chen, H.-T.; Padilla, W. J.; Zide, J. M.; Gossard, A. C.; Taylor, A. J.; Averitt, R. D. Active terahertz metamaterial devices. *Nature* **2006**, *444*, 597–600.
- (11) Fan, S.; Johnson, S. G.; Joannopoulos, J.; Manolatos, C.; Haus, H. Waveguide branches in photonic crystals. *JOSA B* **2001**, *18*, 162–165.
- (12) Piggott, A. Y.; Lu, J.; Lagoudakis, K. G.; Petykiewicz, J.; Babinec, T. M.; Vučković, J. Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer. *Nat. Photonics* **2015**, *9*, 374–377.
- (13) Roeloffzen, C. G.; Zhuang, L.; Taddei, C.; Leinse, A.; Heideman, R. G.; van Dijk, P. W.; Oldenbeuving, R. M.; Marpaung, D. A.; Burla, M.; Boller, K.-J. Silicon nitride microwave photonic circuits. *Opt. Express* **2013**, *21*, 22937–22961.
- (14) Jin, J.-M. *The Finite Element Method in Electromagnetics*; John Wiley & Sons, 2015.
- (15) Silvester, P. P.; Ferrari, R. L. *Finite Elements for Electrical Engineers*; Cambridge University Press, 1996.
- (16) Lalanne, P.; Hugonin, J.-P. Numerical performance of finite-difference modal methods for the electromagnetic analysis of one-dimensional lamellar gratings. *JOSA A* **2000**, *17*, 1033–1042.
- (17) Hughes, T. W.; Williamson, I. A.; Minkov, M.; Fan, S. Forward-Mode Differentiation of Maxwell's Equations. *ACS Photonics* **2019**, *6*, 3010–3016.
- (18) Alkan, E.; Demir, V.; Elsherbeni, A.; Arvas, E. Double-Grid Finite-Difference Frequency-Domain (DG-FDFD) Method for Scattering from Chiral Objects. *Synthesis Lectures on Computational Electromagnetics* **2013**, *8*, 1–129.
- (19) Rumpf, R. C. Simple implementation of arbitrarily shaped total-field/scattered-field regions in finite-difference frequency-domain. *Progress In Electromagnetics Research* **2012**, *36*, 221–248.
- (20) Jiang, J.; Chen, M.; Fan, J. A. Deep neural networks for the evaluation and design of photonic devices. *Nature Reviews Materials* **2021**, *6*, 679–700.
- (21) So, S.; Badloe, T.; Noh, J.; Bravo-Abad, J.; Rho, J. Deep learning enabled inverse design in nanophotonics. *Nanophotonics* **2020**, *9*, 1041–1057.
- (22) Ma, W.; Liu, Z.; Kudyshev, Z. A.; Boltasseva, A.; Cai, W.; Liu, Y. Deep learning for the design of photonic structures. *Nat. Photonics* **2021**, *15*, 77–90.
- (23) An, S.; Zheng, B.; Shalaginov, M. Y.; Tang, H.; Li, H.; Zhou, L.; Ding, J.; Agarwal, A. M.; Rivero-Baleine, C.; Kang, M.; et al. Deep learning modeling approach for metasurfaces with high degrees of freedom. *Opt. Express* **2020**, *28*, 31932–31942.
- (24) Yeung, C.; Tsai, J.-M.; King, B.; Pham, B.; Ho, D.; Liang, J.; Knight, M. W.; Raman, A. P. Multiplexed supercell metasurface design and optimization with tandem residual networks. *Nanophotonics* **2021**, *10*, 1133–1143.
- (25) Lin, Z.; Roques-Carnes, C.; Pestourie, R.; Soljačić, M.; Majumdar, A.; Johnson, S. G. End-to-end nanophotonic inverse design for imaging and polarimetry. *Nanophotonics* **2021**, *10*, 1177–1187.

- (26) Lim, J.; Psaltis, D. MaxwellNet: Physics-driven deep neural network training based on Maxwell equations. *APL Photonics* **2022**, *7*, 011301.
- (27) Peurifoy, J.; Shen, Y.; Jing, L.; Yang, Y.; Cano-Renteria, F.; DeLacy, B. G.; Joannopoulos, J. D.; Tegmark, M.; Soljačić, M. Nanophotonic particle simulation and inverse design using artificial neural networks. *Sci. Adv.* **2018**, *4*, eaar4206.
- (28) Malkiel, I.; Mrejen, M.; Nagler, A.; Arieli, U.; Wolf, L.; Suchowski, H. Plasmonic nanostructure design and characterization via deep learning. *Light: Science & Applications* **2018**, *7*, 1–8.
- (29) Ma, W.; Cheng, F.; Liu, Y. Deep-learning-enabled on-demand design of chiral metamaterials. *ACS Nano* **2018**, *12*, 6326–6334.
- (30) Liu, D.; Tan, Y.; Khoram, E.; Yu, Z. Training deep neural networks for the inverse design of nanophotonic structures. *ACS Photonics* **2018**, *5*, 1365–1369.
- (31) An, S.; Fowler, C.; Zheng, B.; Shalaginov, M. Y.; Tang, H.; Li, H.; Zhou, L.; Ding, J.; Agarwal, A. M.; Rivero-Baleine, C.; et al. A deep learning approach for objective-driven all-dielectric metasurface design. *ACS Photonics* **2019**, *6*, 3196–3207.
- (32) Zhelyeznyakov, M. V.; Brunton, S.; Majumdar, A. Deep learning to accelerate scatterer-to-field mapping for inverse design of dielectric metasurfaces. *ACS Photonics* **2021**, *8*, 481–488.
- (33) Wiecha, P. R.; Muskens, O. L. Deep learning meets nanophotonics: a generalized accurate predictor for near fields and far fields of arbitrary 3D nanostructures. *Nano Lett.* **2020**, *20*, 329–338.
- (34) Maxwell, J. C., VIII. A dynamical theory of the electromagnetic field. *Philosophical Transactions of the Royal Society of London* **1865**, *155*, 459–512.
- (35) Raissi, M.; Perdikaris, P.; Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **2019**, *378*, 686–707.
- (36) Goswami, S.; Anitescu, C.; Chakraborty, S.; Rabczuk, T. Transfer learning enhanced physics informed neural network for phase-field modeling of fracture. *Theoretical and Applied Fracture Mechanics* **2020**, *106*, 102447.
- (37) Meng, X.; Li, Z.; Zhang, D.; Karniadakis, G. E. PPINN: Parareal physics-informed neural network for time-dependent PDEs. *Computer Methods in Applied Mechanics and Engineering* **2020**, *370*, 113250.
- (38) Kharazmi, E.; Zhang, Z.; Karniadakis, G. E. Variational Physics-Informed Neural Networks For Solving Partial Differential Equations. *arXiv:1912.00873 [cs.NE]* **2019**, na.
- (39) Jagtap, A. D.; Kawaguchi, K.; Karniadakis, G. E. Adaptive activation functions accelerate convergence in deep and physics-informed neural networks. *J. Comput. Phys.* **2020**, *404*, 109136.
- (40) Raissi, M.; Yazdani, A.; Karniadakis, G. E. Hidden fluid mechanics: A Navier-Stokes informed deep learning framework for assimilating flow visualization data. *arXiv:1808.04327 [cs.CE]* **2018**, na.
- (41) Zhu, Y.; Zabarav, N.; Koutsourelakis, P.-S.; Perdikaris, P. Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data. *J. Comput. Phys.* **2019**, *394*, 56–81.
- (42) Gao, H.; Sun, L.; Wang, J.-X. PhyGeoNet: physics-informed geometry-adaptive convolutional neural networks for solving parameterized steady-state PDEs on irregular domain. *J. Comput. Phys.* **2021**, *428*, 110079.
- (43) Fang, Z. A High-Efficient Hybrid Physics-Informed Neural Networks Based on Convolutional Neural Network. *IEEE Transactions on Neural Networks and Learning Systems* **2021**, 1–13.
- (44) Chen, Y.; Lu, L.; Karniadakis, G. E.; Dal Negro, L. Physics-informed neural networks for inverse problems in nano-optics and metamaterials. *Opt. Express* **2020**, *28*, 11618–11633.
- (45) Lu, L.; Pestourie, R.; Yao, W.; Wang, Z.; Verdugo, F.; Johnson, S. G. Physics-Informed Neural Networks with Hard Constraints for Inverse Design. *SIAM Journal on Scientific Computing* **2021**, *43*, B1105–B1132.
- (46) Fang, Z.; Zhan, J. Deep physical informed neural networks for metamaterial design. *IEEE Access* **2020**, *8*, 24506–24513.
- (47) Ma, W.; Xu, Y.; Xiong, B.; Deng, L.; Peng, R. W.; Wang, M.; Liu, Y. Pushing the Limits of Functionality-Multiplexing Capability in Metasurface Design Based on Statistical Machine Learning. *Adv. Mater.* **2022**, *34*, 2110022.
- (48) Liu, Z.; Zhu, D.; Rodrigues, S. P.; Lee, K. T.; Cai, W. Generative Model for the Inverse Design of Metasurfaces. *Nano Lett.* **2018**, *18*, 6570–6576.
- (49) Ma, W.; Cheng, F.; Xu, Y.; Wen, Q.; Liu, Y. Probabilistic Representation and Inverse Design of Metamaterials Based on a Deep Generative Model with Semi-Supervised Learning Strategy. *Adv. Mater.* **2019**, *31*, 1901111.
- (50) Ren, H.; Shao, W.; Li, Y.; Salim, F.; Gu, M. Three-dimensional vectorial holography based on machine learning inverse design. *Science Advances* **2020**, *6*, eaaz4261.
- (51) Kudyshev, Z. A.; Kildishev, A. V.; Shalae, V. M.; Boltasseva, A. Machine-learning-assisted metasurface design for high-efficiency thermal emitter optimization. *Applied Physics Reviews* **2020**, *7*, 021407.
- (52) Ma, W.; Liu, Y. A data-efficient self-supervised deep learning model for design and characterization of nanophotonic structures. *Science China: Physics, Mechanics and Astronomy* **2020**, *63*, 284212.
- (53) Lu, L.; Jin, P.; Pang, G.; Zhang, Z.; Karniadakis, G. E. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence* **2021**, *3*, 218–229.
- (54) DeVore, J. R. Refractive indices of rutile and sphalerite. *JOSA* **1951**, *41*, 416–419.
- (55) Bond, W. Measurement of the refractive indices of several crystals. *J. Appl. Phys.* **1965**, *36*, 1674–1677.
- (56) Pettit, G.; Turner, W. Refractive index of InP. *J. Appl. Phys.* **1965**, *36*, 2081–2081.
- (57) Papatryfonos, K.; Angelova, T.; Brimont, A.; Reid, B.; Guldin, S.; Smith, P. R.; Tang, M.; Li, K.; Seeds, A. J.; Liu, H.; et al. Refractive indices of MBE-grown Al_xGa_(1-x)As ternary alloys in the transparent wavelength region. *AIP Advances* **2021**, *11*, 025327.
- (58) Green, M. A. Self-consistent optical parameters of intrinsic silicon at 300 K including temperature coefficients. *Sol. Energy Mater. Sol. Cells* **2008**, *92*, 1305–1310.
- (59) Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. *LNCS* **2015**, *9351*, 234–241.
- (60) Li, X.; Chen, H.; Qi, X.; Dou, Q.; Fu, C.-W.; Heng, P.-A. H-DenseUNet: hybrid densely connected UNet for liver and tumor segmentation from CT volumes. *IEEE transactions on medical imaging* **2018**, *37*, 2663–2674.
- (61) Weng, Y.; Zhou, T.; Li, Y.; Qiu, X. Nas-unet: Neural architecture search for medical image segmentation. *IEEE Access* **2019**, *7*, 44247–44257.
- (62) Zeng, Z.; Xie, W.; Zhang, Y.; Lu, Y. RIC-Unet: An improved neural network based on Unet for nuclei segmentation in histology images. *Ieee Access* **2019**, *7*, 21420–21428.
- (63) Yee, K. Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media. *IEEE Transactions on Antennas and Propagation* **1966**, *14*, 302–307.
- (64) Snyder, A. W.; Love, J. *Optical Waveguide Theory*; Springer Science & Business Media, 2012.
- (65) Sigmund, O.; Maute, K. Topology optimization approaches. *Structural and Multidisciplinary Optimization* **2013**, *48*, 1031–1055.
- (66) Miller, O. D. *Photonic design: From fundamental solar cell physics to computational inverse design*; University of California: Berkeley, 2012.
- (67) Chen, M.; Jiang, J.; Fan, J. A. Design space reparameterization enforces hard geometric constraints in inverse-designed nanophotonic devices. *ACS Photonics* **2020**, *7*, 3141–3151.
- (68) Jiang, J.; Fan, J. A. Simulator-based training of generative neural networks for the inverse design of metasurfaces. *Nanophotonics* **2019**, *9*, 1059–1069.
- (69) Jiang, J.; Fan, J. A. Global optimization of dielectric metasurfaces using a physics-driven neural network. *Nano Lett.* **2019**, *19*, 5366–5372.

- (70) Sell, D.; Yang, J. et al. Device components formed of geometric structures. U.S. Patent US10,725,290, 2020.
- (71) Lin, Z.; Groever, B.; Capasso, F.; Rodriguez, A. W.; Lončar, M. Topology-optimized multilayered metaoptics. *Physical Review Applied* **2018**, *9*, 044030.
- (72) Mansouree, M.; Kwon, H.; Arbabi, E.; McClung, A.; Faraon, A.; Arbabi, A. Multifunctional 2.5 D metastructures enabled by adjoint optimization. *Optica* **2020**, *7*, 77–84.
- (73) McGuinness, C.; Byer, R.; Colby, E.; Cowan, B.; England, R.; Noble, R.; Plettner, T.; Sears, C.; Siemann, R.; Spencer, J.; et al. Woodpile structure fabrication for photonic crystal laser acceleration. *AIP Conf. Proc.* **2009**, *1086*, 544–549.
- (74) Barbedo, J. G. A. Impact of dataset size and variety on the effectiveness of deep learning and transfer learning for plant disease classification. *Computers and electronics in agriculture* **2018**, *153*, 46–53.
- (75) Catal, C.; Diri, B. Investigating the effect of dataset size, metrics sets, and feature selection techniques on software fault prediction problem. *Information Sciences* **2009**, *179*, 1040–1058.
- (76) Shorten, C.; Khoshgoftaar, T. M. A survey on image data augmentation for deep learning. *Journal of Big Data* **2019**, *6*, 1–48.
- (77) Fadaee, M.; Bisazza, A.; Monz, C. Data augmentation for low-resource neural machine translation. *arXiv:1705.00440 [cs.CL]* **2017**, na.
- (78) Wang, S.; Wang, H.; Perdikaris, P. Learning the solution operator of parametric partial differential equations with physics-informed DeepOnets. *arXiv:2103.10974 [cs.LG]* **2021**, na.
- (79) Trivedi, R.; Su, L.; Lu, J.; Schubert, M. F.; Vuckovic, J. Data-driven acceleration of photonic simulations. *Sci. Rep.* **2019**, *9*, 1–7.
- (80) Berenger, J.-P. A perfectly matched layer for the absorption of electromagnetic waves. *J. Comput. Phys.* **1994**, *114*, 185–200.
- (81) Kingma, D. P.; Ba, J. Adam: A method for stochastic optimization. *arXiv:1412.6980 [cs.LG]* **2014**, na.
- (82) He, K.; Zhang, X.; Ren, S.; Sun, J. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *Proceedings of the IEEE international conference on computer vision* **2015**, 1026–1034.
- (83) Hugonin, J. P.; Lalanne, P. RETICOLO software for grating analysis. *arXiv:2101.00901 [physics.optics]* **2021**, na.

NOTE ADDED AFTER ASAP PUBLICATION

This paper was published on August 22, 2022. Missing labels were added to Figure 4a, and the corrected version was reposted on August 29, 2022.

Recommended by ACS

Reparameterization Approach to Gradient-Based Inverse Design of Three-Dimensional Nanophotonic Devices

Erez Gershnel, Jonathan A. Fan, *et al.*

OCTOBER 17, 2022

ACS PHOTONICS

READ 

VarRCWA: An Adaptive High-Order Rigorous Coupled Wave Analysis Method

Ziwei Zhu and Changxi Zheng

SEPTEMBER 14, 2022

ACS PHOTONICS

READ 

Constraining Continuous Topology Optimizations to Discrete Solutions for Photonic Applications

Conner Ballew, Andrei Faraon, *et al.*

JANUARY 09, 2023

ACS PHOTONICS

READ 

Efficient Inverse Design of Large-Area Metasurfaces for Incoherent Light

Raphaël Pestourie, Steven G. Johnson, *et al.*

SEPTEMBER 19, 2022

ACS PHOTONICS

READ 

Get More Suggestions >