

RESEARCH ARTICLE OPEN ACCESS

Broadband Electromagnetic Field Prediction at Unseen Wavelengths via Physics-Guided Neural Operators

Joonhyuk Seo¹  | Chanik Kang²  | Dongjin Seo³  | Haejun Chung² 

¹University of California Irvine, Irvine, California, USA | ²Hanyang University, Seoul, South Korea | ³Yale University, New Haven, Connecticut, USA

Correspondence: Haejun Chung (haejun@hanyang.ac.kr)

Received: 11 November 2025 | **Revised:** 19 April 2026 | **Accepted:** 30 July 2026

Keywords: Maxwell's equations | metasurface | neural operators | surrogate solver

ABSTRACT

Surrogate solvers that deliver full-wave accuracy across broad spectral ranges can unlock the rapid design and analysis of nanophotonic devices. We present a physics-guided neural-operator surrogate solver that predicts electromagnetic field distributions throughout the visible band from a permittivity map and a query wavelength, including wavelengths unseen during training. Our approach is built on two key ideas: (i) spectral consistency, which formalizes the intrinsic relationship between wavelength-dependent field variations and spatial frequency, and (ii) a conditional embedding framework that comprises a refined wave prior. We validate the approach on three representative nanophotonic platforms—a single-layer metasurface (metalenses), a five-layer volumetric metasurface (spectrum-splitters), and a freeform waveguide for photonic integrated circuits—achieving full-wave fidelity while reducing mean field-prediction error relative to prior operator-learning baselines by up to 80.3% (single-layer), 56.8% (multilayer), and 32.3% (waveguide). The solver remains compact (0.43 M parameters; 86.9% fewer than FNO) and accelerates simulation by $\sim 42\times$ compared with a finite-difference frequency-domain solver while yielding robust broadband generalization. By providing continuous-wavelength field predictions with full-wave accuracy, this surrogate solver eliminates a key bottleneck in modeling and inverse design of metalenses, spectrum splitters, and emerging large-area meta-optics, offering a drop-in engine for system-level co-optimization in nanophotonics.

1 | Introduction

Recent advances at the intersection of deep learning (DL) and photonics have led to significant breakthroughs across various domains, including the inverse design of photonic structures [1–5], high-resolution imaging and computational microscopy [6–8], optical signal processing and communication systems [9–11], and the realization of optical neural networks [12–14]. Specifically, DL-based approaches have demonstrated remarkable capabilities in rapidly predicting electromagnetic phenomena traditionally analyzed through numerical methods [2, 15–17]. By training neural networks on datasets derived from simulations or experiments, DL approaches achieve orders-of-

magnitude acceleration in predicting electromagnetic fields, device performance, and structural optimizations, all while maintaining accuracy comparable to traditional numerical solvers [2, 15–17].

One important application of DL in photonics is the development of surrogate solvers that approximate electromagnetic simulations. These surrogate solvers predominantly utilize deep neural networks due to their rapid inference and notable learning capability. Traditional data-driven neural network-based surrogate solvers [15, 16, 18–20] learn end-to-end mappings from structure/excitation to fields or responses, approximating the input-output operator of Maxwell's equations.

Joonhyuk Seo and Chanik Kang contributed equally to this work.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 The Author(s). *Nanophotonics* published by Wiley Periodicals LLC.

However, these approaches frequently encounter limitations in accurately capturing the underlying physical characteristics of Maxwell's equations [2].

To address these shortcomings, physics-informed neural networks (PINNs) [21–23] and physics-regularized learning methods [2, 24] have been widely explored. PINNs incorporate Maxwell's equations into the training objective by penalizing PDE residuals and boundary/source conditions, enabling training without labeled simulation data. Physics-regularized methods, by contrast, augment supervised/data losses with physics-based penalties, reducing but not eliminating the need for labeled data. Despite these advantages, standard PINNs are largely instance-specific: unless the network is explicitly conditioned on simulation parameters (e.g., wavelength and relative permittivity ϵ_r) and trained across their ranges, they typically require retraining for each parameter setting. Even with parameter conditioning, accuracy often degrades at parameter values not covered during training, leading to reduced field-prediction fidelity on unseen simulation settings.

Recently, surrogate solvers utilizing neural operators have gained attention, as these methods effectively learn mappings between infinite-dimensional function spaces, reflecting the intrinsic characteristics of partial differential equations (PDEs) [25, 26]. Notable progress in operator learning includes transformer-based neural operators [27–30], convolutional neural operators [31], and Fourier neural operators (FNO) [17, 26, 32–35]. Among these, FNO has demonstrated strong inference performance across various electromagnetic applications [17, 36, 37], making it particularly promising for accelerating electromagnetic simulations.

Nevertheless, recent surrogate solver studies, including FNO-based methods, exhibit limited prediction accuracy under the untrained simulation conditions, thereby failing to demonstrate the generalization capability of these models. The notable decline in predictive accuracy at intermediate wavelengths between discretely sampled training points raises fundamental concerns about whether these models adequately capture the underlying physics governing electromagnetic phenomena.

In this work, we present a surrogate solver that delivers highly accurate, full-wave solutions of Maxwell's equations across a continuous, broadband spectrum, even when training data are available only at sparsely sampled wavelengths (Figure 1). Our model maintains robust accuracy on untrained intermediate wavelengths, overcoming the generalization limits of existing surrogate solvers. Moreover, by leveraging the intrinsic scale invariance of Maxwell's equations [38], our demonstrations satisfy the solutions of Maxwell's equations at any target frequency band by multiplying a scaling factor.

2 | Broadband Maxwell Surrogate Solver

Similar to a prior work [25], our model has an iterative process to map between infinite-dimensional function spaces Figure 2.

$$u = \mathcal{G}_\theta(a) = (\mathcal{Q} \circ \mathcal{L}^{(L)} \circ \dots \circ \mathcal{L}^{(1)} \circ \mathcal{P})(a) \quad (1)$$

where $\mathcal{L}^{(l)}$ is a l -th operator layer, L indicates the number of layers, and $\mathcal{P}$, $\mathcal{Q}$ are lifting and projection operators, respectively. We define the layers $\mathcal{L}^{(l)}$ as follows:

$$\mathcal{L}^{(l)}(z^{(l)}) = z^{(l)} + W_2^{(l)} \sigma \left(W_1^{(l)} \tilde{\mathcal{K}}^{(l)}(z^{(l)}) + b_1^{(l)} \right) + b_2^{(l)} \quad (2)$$

where $\tilde{\mathcal{K}}^{(l)}$ is our kernel integral operator based on the Fourier transform [32].

Following prior work [33, 39], we employ a residual block consisting of two linear transformations separated by a nonlinear activation, which improves representational capacity and facilitates optimization. To preserve discretization invariance, we use pointwise 1×1 convolutions in place of spatial 3×3 convolutions, in contrast to ref. [17].

2.1 | Wave Prior

As illustrated in Figure 3, the polarization-dependent wave prior $\mathcal{W} = [\mathcal{W}_x, \mathcal{W}_y]$ is a synthetic wave pattern derived from homogeneous-medium solutions of the wave equation. It is defined as

$$\begin{aligned} \mathcal{W}_x &= \exp \left(i \frac{2\pi}{\lambda} \sqrt{\epsilon_r} \mathbf{x} \mathbf{1}^\top \Delta \ell_x \right), \\ \mathcal{W}_y &= \exp \left(i \frac{2\pi}{\lambda} \sqrt{\epsilon_r} \mathbf{1} \mathbf{y}^\top \Delta \ell_y \right), \end{aligned}$$

where $\epsilon_r \in \mathbb{C}^{H \times W}$ is the spatial map of relative permittivity, λ is the wavelength, $\Delta \ell_x$ and $\Delta \ell_y$ are the grid spacings, $\mathbf{x}$ and $\mathbf{y}$ denote the coordinate index vectors along x and y , and $\mathbf{1}$ is a vector of ones used to tile coordinates over the 2D grid. Because it originates from a homogeneous-medium solution, $\mathcal{W}$ reproduces the wavelength-dependent oscillation patterns observed in simulator-generated fields within uniform regions.

At material interfaces, the fields satisfy electromagnetic boundary conditions. For transverse-electric (TE) polarization, the tangential electric field is continuous, $E_{t,1} = E_{t,2}$. For transverse-magnetic (TM) polarization, the normal component of the electric displacement is continuous, $\epsilon_1 E_{n,1} = \epsilon_2 E_{n,2}$, where ϵ_1 and ϵ_2 are the relative permittivities of the adjoining media.

However, the wave prior used in this work is derived solely from the macroscopic permittivity map. It thus does not account for other complex electromagnetic phenomena in sub-wavelength structures, such as wavelength-dependent resonances and evanescent fields. Consequently, it exhibits sharp and physically implausible features in regions where the permittivity changes.

2.2 | Spectral Consistency

As discussed earlier, complex optical effects such as scattering and resonance induce nonlinear, wavelength-dependent variations in electric field profiles, which can impede accurate predictions at wavelengths not encountered during training. To address this challenge, the model must encode the relevant physics at arbitrary wavelengths, including unseen ones, by (i) constructing *conditional embeddings* that reliably capture




FIGURE 1 | Overview of the approach. (a) *Training*: the network is supervised on paired permittivity profiles (maps) and electric-field distributions over the simulation domain. Field data are provided at discrete wavelengths uniformly spaced every 20 nm across the visible spectrum. (b) *Inference*: given a permittivity profile and a query wavelength, the trained model predicts accurate full-wave field distributions, including at unseen wavelengths. (c) *Example geometry*: schematic of a 2D subwavelength grating and its electric-field distribution under normal-incidence plane-wave illumination.




FIGURE 2 | Model architecture and WIME layer. (a) *Network architecture*: the polarization-dependent wave prior (conditional embedding) and the permittivity profile (input) are lifted with 1×1 convolutions (blue) and then passed through a stack of wave-informed multiplicative encoding (WIME) layers (pink). The latent refined wave-prior embedding Z_{wp} is injected at every layer. (b) *Single WIME layer*: features are first updated by a kernel integral operator (green), followed by 1×1 convolutions (blue) and a nonlinear activation (red); the projected wave-prior embedding modulates the feature map via element-wise multiplication ($\odot$). (c) *Kernel integral operator schematic*.

wavelength-dependent features and (ii) designing conditioning mechanisms that effectively integrate these embeddings throughout the network.

As illustrated in Figure 4, shorter wavelengths generally yield field patterns with higher spatial frequencies, whereas longer wavelengths produce sparser, lower-frequency structures. This




FIGURE 3 | TM and TE wave priors. The TM wave prior $\mathcal{W}_x = e^{\frac{2\pi\sqrt{\epsilon_r}}{\lambda} \cdot x1^T \Delta k_x}$ incorporates the spatially varying relative permittivity ϵ_r ; the pronounced vertical seam marks interfaces where ϵ_r changes and the normal electric field is discontinuous. In contrast, the TE wave prior $\widetilde{\mathcal{W}}_x = e^{\frac{2\pi}{\lambda} \cdot x1^T \Delta k_x}$ omits the ϵ_r factor, yielding a smooth, plane-wave-like pattern.




FIGURE 4 | Spectral consistency. Across wavelengths, the latent wave prior (top) and the corresponding predicted electric-field distributions (bottom) exhibit similar spatial-frequency content, indicating that their spectra remain aligned under wavelength variation.

inverse relationship is inherent to solutions of Maxwell's equations and follows from $\nu = 1/\lambda$, where λ and ν denote the wavelength and spatial frequency, respectively [40]. We define *spectral consistency* as the requirement that wavelength-dependent conditional embeddings reflect this inverse relation and that their spatial-frequency content aligns with that of the desired electromagnetic-field distributions. Embeddings that satisfy spectral consistency provide an informed prior on the expected field complexity at target wavelengths, including those absent from the training set, thereby facilitating robust broadband generalization in surrogate models.

In this setting, conditional information that satisfies spectral consistency encodes the expected field complexity and spatial-frequency distribution at target wavelengths, including those not seen during training. Enforcing spectral consistency, therefore, substantially enhances the surrogate model's broadband generalization.

3 | Result and Discussion

In the following, we compare the field prediction performance for the continuous-spectrum interpolation within the targeted bandwidth, including untrained intermediate wavelengths, in various simulation environments. We select three representative optical structures to investigate the generalization performance of our model under diverse electromagnetic phenomena.

3.1 | Single Layer Metasurface: Metalens

Metasurfaces enable precise wavefront engineering by imposing arbitrary, spatially varying phases and amplitude profiles on transmitted or reflected waves. This versatility has been exploited in a wide range of applications, including beam steering [41–43], high-resolution imaging [6–8], holography [44–46], vortex beam generation [47, 48], and communications [9–11]. Owing to their subwavelength thickness, metasurfaces provide a compact alternative to bulky optics and can be manufactured at scale using CMOS-compatible lithographic processes [49].

Meta-atoms in metasurfaces are individual scatterers and resonators that produce a prescribed macroscopic wavefront. Due to the resonance property of the meta-atoms, the meta-atoms project Maxwell's solutions nonlinearly over the varying wavelengths. This hinders large-scale metasurface simulations; therefore, most of the metasurface studies assume a local periodic assumption for domain decomposition. This assumption often greatly degrades metasurface efficiency when the required phase and amplitude vary rapidly over the design region [50]. Therefore, it is very crucial to develop full-scale metasurface simulators. Here, we apply our model to predict broadband solutions of Maxwell's equations in a single metasurface example, which can be applied to a metalens design.

Metasurface unit cells, often referred to as meta-atoms, serve as subwavelength scatterers and resonators that collectively synthesize a prescribed macroscopic wavefront. Because each meta-atom's resonant response varies nonlinearly with wavelength, simulating an entire metasurface at full scale becomes computationally expensive. To mitigate this, most designs invoke a local periodicity assumption for domain decomposition [51], but this approximation may not hold whenever the required phase and amplitude profiles change rapidly across the surface [50]. Such limitations impose the necessity for full-scale metasurface simulators. Here, we apply our broadband surrogate model to predict full-wave Maxwell solutions for a single-layer metasurface [52, 53], illustrating its potential applicability to metalens design.

3.1.1 | Problem Setup

In this study, we utilize the open-source FDFD simulation tool Ceviche [54] to generate electromagnetic field data for photonic structures. All simulations employ a uniform grid resolution of 40 points per micrometer and include a 1- μm -thick perfectly matched layer at the boundaries. For all cases, electric field distributions are generated under normally incident plane-wave illumination for wavelengths spanning 400–700 nm. The single-layer metalens is simulated over a $5.2\ \mu\text{m} \times 6.57\ \mu\text{m}$ domain using SiO_2 (refractive index 1.46) and air in 25 nm grid patches. Boundary conditions include a 1- μm PML on all sides Figure 5.

In addition, we compare the field prediction performance of UNet [55, 56], FNO2d [32], and F-FNO2d [33], commonly used in surrogate solver study. Moreover, we evaluate our model against NeurOLight [17], the state-of-the-art model for field prediction of a narrow wavelength range in electromagnetic simulations. For several baseline models (UNet, FNO2d, and F-FNO2d), the conditional information was incorporated by simply concatenating the wave prior to the input. NeurOLight similarly utilized the wave prior concatenation, following the methodology proposed in the original study.

3.1.2 | Result

Figure 6 compares the proposed model's prediction with the full-wave FDFD reference for the single-layer metasurface and shows the corresponding absolute error. The left and middle panels show the electric-field magnitude $|E|/E_{\text{max}}$ (shared color scale), whereas the right panel displays the absolute error $|E_{\text{pred}} - E_{\text{gt}}|$ (separate scale). The predicted field closely reproduces the interference pattern and lobe structure observed in the reference solution, with small residual errors largely confined to regions of rapid spatial variation near the patterned surface and along fine fringes.

Figure 7 visualizes absolute error maps $|E_{\text{pred}} - E_{\text{gt}}|$ at the unseen wavelength $\lambda = 410\text{ nm}$ for the single-layer metasurface. The five panels compare prior baselines (Refs. [17, 32, 33, 55, 56]) with the proposed model under a shared color scale (right; larger values indicate higher error). The orange dashed line marks the boundary of the design region (labeled “Design space”). Relative to the baselines, the proposed model exhibits uniformly lower errors across the domain with residuals largely confined near the device and fine interference fringes,

indicating stronger generalization at wavelengths not used for training.

Figure 8 plots the normalized mean squared error (NMSE) versus wavelength for the single-layer metasurface, comparing the proposed method with prior baselines (refs. [17, 32, 33, 55, 56]) over the full simulation domain. Green dashed vertical lines mark the wavelengths included during training; all other wavelengths are evaluated without retraining. Curves show the mean NMSE and shaded regions indicate ± 2 standard deviations. Across the 400–700 nm band, the proposed model maintains consistently low error on both trained and untrained wavelengths, whereas baselines exhibit larger fluctuations and pronounced error spikes at certain wavelengths, evidencing stronger broadband generalization of our approach.

Table 1 compares single-layer metasurface field prediction across five models. Entries show the average error (lower is better) over two spatial regions, the entire field and the structure (design) region, and for two evaluation regimes: wavelengths used during training (“Trained”) and untrained wavelengths. Our method attains the lowest error on both untrained settings (0.002035 for the entire field and 0.003812 for the structure field), reducing error by $\sim 39\%$ relative to the best baseline in each column (UNet) and by $\sim 80\%$ relative to NeurOLight. On trained wavelengths, our model remains competitive (0.001774 entire field; 0.003929 structure field), within 2%–4% of UNet while outperforming FNO/F-FNO and NeurOLight by 10%–13%, showing superior generalization capability of our approach.

3.2 | Multi-Layer Metasurface: Spectrum Splitter

Although metasurfaces offer arbitrary wavefront control, their meta-atoms exhibit inherently frequency-dependent resonances that reduce efficiency outside their design wavelength [57]. This bandwidth limitation restricts critical imaging applications, such as achromatic metalenses [50] and spectral beam splitters [58], to narrow operating ranges. One promising solution is a volumetric, multilayer metasurface architecture, where stacking successive patterned layers increases the effective design volume [59, 60], enabling tailored phase and amplitude responses over a wider spectrum and bringing performance closer to that of conventional bulky optics.

Although multilayer metasurfaces remain uncommon in the infrared and visible regimes, primarily due to the fabrication




FIGURE 5 | Single-layer metasurface: geometry and simulation setup. (a) Schematic of the single-layer metasurface. (b) Simulation setup for the single-layer metasurface. (c) Detailed view of the design region within the simulation domain.




FIGURE 6 | Single-layer metasurface. Side-by-side comparison of the full-wave (ground-truth) electric-field distribution and the field predicted by the proposed model for the single-layer metasurface configuration.




FIGURE 7 | Absolute error at an unseen wavelength. Absolute error maps for the single-layer metasurface at $\lambda = 410$ nm, an unseen wavelength that exhibits the largest prediction error across the evaluated band (computed as $|E_{\text{pred}} - E_{\text{gt}}|$).




FIGURE 8 | Broadband evaluation on the single-layer metasurface. Comparison of all methods across the continuous broadband spectrum within the targeted bandwidth, evaluated over the full simulation domain. Shaded regions denote ± 2 standard deviations for visual clarity, and green dashed lines mark the wavelengths used during training.

complexity of volumetric architectures, recent advances in two-photon polymerization lithography [61, 62] have enabled their realization. To demonstrate the versatility of our approach, we apply our broadband surrogate model to a representative multilayer metasurface, thereby validating its applicability to complex optical structures.

3.2.1 | Problem Setup

Figure 9 summarizes the multilayer metasurface problem and simulation configuration. The computational domain is $5.2\mu\text{m}$

wide, $3.5\mu\text{m}$ tall, whereas the design region (blue) is embedded with a SiO_2 substrate (green), a normally incident source placed $1.0\mu\text{m}$ below the design region (red), and perfectly matched layers (gray) on the boundaries. This multilayer metasurface simulation environment is directly applicable to real-world applications, such as spectrum splitters [58, 63], diffraction gratings [64, 65], LiDAR sensors [43], and many others. The design region includes five layers, each $0.7\mu\text{m}$ thick, patterned across a $3.2\mu\text{m}$ lateral span using alternating ϵ_{low} , ϵ_{high} materials; the minimum design unit (pixel) has a width of 25 nm. This setup

specifies the geometry, materials, and boundary conditions used in the multilayer metasurface experiments.

3.2.2 | Result

Figure 10 compares the proposed model's prediction against the full-wave FDFD result for the multilayer metasurface and visualizes the resulting absolute error. The left and middle panels show the normalized electric-field magnitude $|E|/E_{\max}$ under a shared color scale, whereas the right panel reports the absolute error $|E_{\text{pred}} - E_{\text{gt}}|$ with an independent scale. The prediction closely reproduces the interference fringes and scattering features present in the reference solution. The overall low error indicates high-fidelity agreement between the surrogate and the full-wave solver.

Figure 11 visualizes absolute error maps $|E_{\text{pred}} - E_{\text{gt}}|$ for the multilayer metasurface at the untrained wavelength $\lambda = 410$ nm, which yields the largest error across the evaluated band. The panels compare prior baselines (refs. [17, 32, 33, 55, 56]) with the proposed model using a shared color scale (right); the orange box indicates the design region. Relative to the baselines, our method shows consistently lower errors both inside the device and in the surrounding field with residuals mainly confined to fine interference fringes, evidence of stronger generalization to unseen wavelengths.

Figure 12 plots NMSE versus wavelength for the multilayer metasurface, comparing the proposed method with prior baselines over the full simulation domain. Green dashed vertical lines mark wavelengths included during training; all other wavelengths are evaluated without retraining. Solid curves show

mean NMSE across test structures, and shaded bands indicate ± 2 standard deviations. Across the 400–700 nm range, our model sustains the lowest errors on both trained and unseen wavelengths and exhibits reduced variance, whereas baselines display larger fluctuations and pronounced error peaks, indicating a stronger broadband generalization of the proposed approach.

Table 2 shows NMSE (lower is better) for multilayer metasurface field prediction, averaged over the evaluated band. Results are shown for two spatial regions, the entire field and the structure (design) region, and for two evaluation regimes: wavelengths included during training (“Trained”) and untrained wavelengths. Our method achieves the lowest error in all four settings (Entire: 0.1360 trained/0.1329 untrained; Structure: 0.1549 trained/0.1515 untrained), outperforming UNet, FNO, F-FNO, and NeuOLight. Notably, although some baselines degrade on unseen wavelengths (e.g., NeuOLight), our model maintains nearly unchanged error between trained and untrained evaluations, indicating stronger broadband generalization. In conclusion, this demonstrates that our method can generalize reliably across wavelengths, even in environments where the spectral consistency assumption is significantly weakened.

3.3 | Freeform Waveguide: Photonic-Integrated Circuit

Integrated photonics [66, 67] has emerged as a leading platform for next-generation high-performance computing, leveraging light's intrinsic advantages, low latency, immense bandwidth, and multiplexing capabilities. Advances in photonic integrated circuits (PICs), which integrate a diverse array of optoelectronic components on a single chip, have paved the way for photonic–electronic co-integrated systems. These hybrid circuits hold great promise as hardware accelerators for both conventional high-performance computing and emerging AI workloads. Recent advances in photonic integrated circuits have demonstrated on-chip optical tensor convolution [14, 68], enabling fully optical neural networks that can dramatically accelerate convolutional operations.

Adjoint-based inverse design has become one of the prevailing standards in PIC optimization because it efficiently computes gradients of performance metrics. However, each optimization

TABLE 1 | Comparisons of single-layer metasurface field prediction.

Model	Entire field (absolute error)		Structure field (absolute error)	
	Trained	Untrained	Trained	Untrained
UNet	0.001732	0.003332	0.003775	0.006190
FNO	0.007719	0.013300	0.025067	0.028339
F-FNO	0.022040	0.033090	0.037091	0.043035
NeuOLight	0.001973	0.010730	0.004526	0.019503
Ours	0.001774	0.002035	0.003929	0.003812




FIGURE 9 | Multilayer metasurface: geometry and simulation setup. (a) Schematic of the multilayer metasurface. (b) Simulation configuration for the multilayer device. (c) Detailed view of the design region and its parameterization.




FIGURE 10 | Multilayer metasurface. Side-by-side comparison of the full-wave (ground-truth) electric-field distribution and the field predicted by the proposed model in the multilayer metasurface simulation.




FIGURE 11 | Absolute error at an unseen wavelength. Absolute error maps for the multilayer metasurface at $\lambda = 410$ nm—a wavelength not used in training that exhibits the largest prediction error over the evaluated band (computed as $|E_{\text{pred}} - E_{\text{gt}}|$).




FIGURE 12 | Broadband evaluation on the multilayer metasurface. Comparison of all methods across the continuous broadband spectrum within the targeted bandwidth, evaluated over the full simulation domain. Shaded regions denote ± 2 standard deviations for visual clarity, and green dashed lines mark the wavelengths used during training.

step still requires full-wave Maxwell solutions, so a single device often demands hundreds of electromagnetic simulations. As a result, adjoint methods have been largely confined to individual components of PICs (e.g., wavelength demultiplexers [66], mode

converters [69, 70], or grating couplers [71, 72]). Here, we leverage our broadband surrogate model to predict the field profile of the freeform waveguide, achieving grid-size precision while significantly reducing the computational cost.

3.3.1 | Problem Setup

Figure 13 summarizes the simulation setup of the freeform waveguide problem. It has a $5.2 \mu\text{m} \times 8.85 \mu\text{m}$ simulation domain. Polymer regions ($\epsilon_r = 2.25$) and air are arranged in 25 nm-scale pixels within an unconstrained design region.

3.3.2 | Result

Figure 14 compares the surrogate's prediction with the full-wave FDFD result for the freeform waveguide and visualizes the residual error. Although this simulation setting includes a relatively smaller permittivity contrast compared to the spectrum splitter simulation, the unconstrained geometric design causes complex electromagnetic phenomena. The left and middle panels display the normalized electric-field magnitude $|E|/E_{\text{max}}$ under a shared color scale, whereas the right panel shows the absolute error $|E_{\text{pred}} - E_{\text{gt}}|$ with its own scale. The prediction closely reproduces the interference fringes and scattering features present in the reference. The uniformly low error indicates high-fidelity agreement between the surrogate and the full-wave solver.

Figure 15 visualizes absolute error maps $|E_{\text{pred}} - E_{\text{gt}}|$ for the freeform waveguide at the untrained wavelength $\lambda = 410 \text{ nm}$, which yields the largest error across the evaluated band. The panels compare prior baselines (refs. [17, 32, 33, 55, 56]) with the proposed model under a shared color scale (right); the orange rectangle delineates the design region ("Design space"). Relative to the baselines, our method exhibits visibly lower errors both inside the device and in the surrounding field.

Figure 16 plots NMSE versus wavelength for the freeform waveguide, comparing the proposed method with prior

baselines over the full simulation domain. Green dashed vertical lines mark wavelengths included during training; all other wavelengths are evaluated without retraining. Solid curves report the mean NMSE across test structures, and shaded bands indicate ± 2 standard deviations. Across the 400–700 nm range, our model maintains the lowest errors on both trained and unseen wavelengths and exhibits reduced variance, whereas baselines show larger fluctuations and pronounced error peaks, evidencing stronger broadband generalization.

Table 3 shows NMSE for freeform waveguide field prediction across two spatial regions, the entire field and the design region, and two evaluation settings: wavelengths seen during training ("Trained") and untrained wavelengths. Our method attains the lowest error in all four cases (Entire: 0.1305 trained/0.1273 untrained; design: 0.1528 trained/0.1479 untrained), improving over the strongest baseline, NeurOLight, by $\approx 23\%$ – 33% depending on the setting. The minimal trained–untrained gap—indeed a slight decrease on unseen wavelengths—indicates robust broadband generalization, whereas UNet, FNO, and F-FNO exhibit substantially higher errors, particularly within the design region.

Despite the increased complexity of the freeform structure, the proposed model still achieves the lowest average NMSE and error variance across the simulation region. In particular, the error increase for untrained wavelengths remains minimal, further validating the model's generalization ability. In contrast, the other models exhibit both higher average errors and broader error distributions. The results also demonstrate that our method enables robust electromagnetic field prediction even

TABLE 2 | Comparisons of multilayer metasurface field prediction.

	Entire field (absolute error)		Structure field (absolute error)	
	Trained	Untrained	Trained	Untrained
Unet	0.2182	0.2167	0.2013	0.1987
FNO	0.4712	0.4962	0.5055	0.5184
F-FNO	0.4579	0.4579	0.4730	0.5143
NeurOLight	0.1683	0.3160	0.1949	0.2785
Ours	0.1360	0.1329	0.1549	0.1515




FIGURE 14 | Freeform waveguide. Comparison of the full-wave (ground-truth) electric-field distribution and the field predicted by the proposed model for the freeform waveguide configuration.




FIGURE 13 | Freeform waveguide: geometry and simulation setup. (a) Schematic of the freeform waveguide. (b) Simulation configuration for the freeform device. (c) Detailed view of the design region and its parameterization.




FIGURE 15 | Absolute error at an unseen wavelength. Absolute error maps for the freeform waveguide at $\lambda = 410$ nm—a wavelength not used in training that exhibits the largest prediction error across the evaluated band (computed as $|\mathbf{E}_{\text{pred}} - \mathbf{E}_{\text{gt}}|$).




FIGURE 16 | Broadband evaluation on the freeform waveguide. Comparison of all methods across the continuous broadband spectrum within the targeted bandwidth, evaluated over the full simulation domain. Shaded regions denote ± 2 standard deviations for visual clarity, and green dashed lines mark the wavelengths used during training.

TABLE 3 | Comparisons of freeform waveguide field prediction.

	Entire field (absolute error)		Design field (absolute error)	
	Trained	Untrained	Trained	Untrained
Unet	0.2886	0.3723	0.3294	0.4211
FNO	0.3681	0.3735	0.4296	0.4313
F-FNO	0.3529	0.3650	0.4140	0.4193
NeurOLight	0.1696	0.1892	0.1996	0.2124
Ours	0.1305	0.1273	0.1528	0.1479

under the weakened spectral consistency assumption by a highly complex electromagnetic interplay from structurally unconstrained design conditions.

3.4 | Computational Efficiency

3.4.1 | Evaluating Performance Across Varying Dataset Sizes

To analyze the dependency between dataset size and model performance, we conduct additional experiments evaluating prediction accuracy across various training dataset sizes. As shown in Figure 17, our method consistently outperforms the baseline models for untrained wavelengths across all dataset

sizes, including significantly reduced size. These results demonstrate the data efficiency of our approach. Such data-efficient performance has practical implications for applications where data collection is expensive or limited.

3.4.2 | Computational Efficiency and Model Compactness

Our method also delivers substantial acceleration relative to a conventional numerical solver. To assess computational efficiency and model compactness, we compare per-sample runtime between the FDFD simulator (Ceviche [54]) and our surrogate solver (Table 4). With a batch size of 1, our model achieves $\sim 16\times$ speedup; under the default batch size of 32, the speedup increases to $\sim 42\times$, representing an order-of-magnitude improvement in runtime efficiency over the numerical solver.

Beyond runtime, we analyze parameter complexity (Table 5). In 2D, a standard FNO layer requires $C^2 M_v M_h$ parameters, where M_v and M_h are the numbers of retained frequency modes in the vertical and horizontal directions, respectively, and C is the channel count, leading to large models. In contrast, our surrogate uses only 0.43 million parameters, $\approx 86.9\%$ reduction relative to FNO-2D [32]'s 3.29 million and $\approx 96.33\%$ reduction compared with UNet [55, 56]'s 11.60 million. This reduction




FIGURE 17 | Dataset size ablation on the single-layer metalens. Field-prediction accuracy as a function of training-set size, reported for wavelengths included during training (Trained) and for untrained wavelengths (Unseen). Lower values indicate better performance. Whiskers denote the standard deviation scaled by 1/5 for visual clarity.

TABLE 4 | Runtime per sample.

	Ours (batch size 1)	Ours (batch size 32)	Numerical simulation
Time (s)	0.0172	0.0067	0.2786

Note: Computational time per sample for the proposed surrogate compared with a conventional finite-difference frequency-domain (FDFD) solver (Ceviche). Field prediction with our model was performed on a single RTX 6000 Ada GPU, whereas the FDFD simulation ran on a 64-core AMD 5995WX CPU.

TABLE 5 | Comparison of the number of parameters across models.

Model	Parameters (M)
U-Net	11.60
FNO	3.29
F-FNO	1.89
NeurOLight	1.65
Ours	0.43

stems from our *Fourier Group Convolutional Shuffling*, which decomposes each Fourier layer into separate horizontal and vertical transforms and applies group-wise weight sharing with channel shuffling, compressing the spectral weight tensor from $C^2 M_v M_h$ to $C^2 (M_v + M_h) / (4G)$, where G is the number of groups. Despite the substantially smaller parameter count, the model attains state-of-the-art accuracy in both the design and propagation regions.

3.5 | Limitations and Future Directions

Although the proposed physics-guided neural operator demonstrates robust performance and a highly compact architecture across diverse nanophotonic platforms, certain limitations present clear opportunities for future research. A primary

simplification in our current baseline test cases is the assumption of constant, nondispersive refractive indices. In physical broadband applications, material dispersion and loss are dominant effects that dictate the lineshape and spectral position of optical resonances. Although our extended analyses indicate that the model's predictive capability remains stable across a wide range of arbitrary permittivities, incorporating wavelength-dependent complex refractive indices into the conditional embedding is an essential next step.

Furthermore, it is important to clarify that the current study demonstrates high-fidelity continuous-spectrum interpolation within a targeted operational bandwidth (400–700 nm), rather than unbounded spectral extrapolation outside the training range. Although interpolating across this wide fractional bandwidth is a substantial achievement, predicting electromagnetic phenomena entirely outside the trained domain remains an active challenge for data-driven solvers. Nevertheless, by leveraging the intrinsic scale invariance of Maxwell's equations, the interpolative capabilities demonstrated here can be mathematically mapped to different physical bandwidths using scaling factors, thereby broadening the model's applicability without requiring direct spectral extrapolation.

Another important consideration is the training data sampling strategy. The current model relies predominantly on fixed uniform spectral sampling (20 nm gaps) to systematically validate the architectural contributions of the proposed physics-guided priors. However, the ability to perfectly resolve highly abrupt spectral features, such as sharp resonances, is intrinsically tied to this sampling distribution. Future studies investigate optimal data acquisition strategies, including sparse, nonuniform, random, or Chebyshev sampling, to determine the most data-efficient methods for capturing complex spectral landscapes while avoiding the harmonic blind spots inherent to uniform grids [73].

Finally, although the proposed model exhibits orders-of-magnitude acceleration compared to single-instance FDFD calculations, a more comprehensive computational benchmark is necessary for a holistic assessment. In practical broadband nanophotonics design, conventional workflows often utilize single-shot FDTD simulations or highly parallelized FDFD runs distributed across multi-node clusters. Future work will benchmark the proposed neural operator against these massively parallel and inherently broadband numerical environments to definitively quantify its computational advantages in true industrial and academic workflows.

4 | Conclusion

In this paper, we present an electromagnetic surrogate solver that generates highly accurate, full-wave solutions to Maxwell's equations over a continuous, broadband spectrum, even when the training set contains only sparsely sampled wavelengths. Our model inherently maintains strong predictive accuracy at intermediate wavelengths not seen during training, thereby overcoming the generalization limitations of existing surrogate solvers in electromagnetism. Moreover, by exploiting the intrinsic scale invariance of Maxwell's equations [38], the solutions obtained at one frequency can be uniformly extended to any other target frequency band through simple multiplication by an appropriate scaling factor.

To ensure broad-spectrum generalizability, we introduce two foundational concepts. First, spectral consistency, which formalizes the inherent link between wavelength-dependent field variations and spatial frequency as dictated by Maxwell's equations. Second, a conditional embedding framework supported by an efficient conditioning mechanism. Concretely, we use the polarization-dependent wave prior and a wave-informed element-wise multiplicative encoding, each explicitly crafted to enforce spectral consistency. By uniting these components, our solver produces reliable, high-fidelity predictions at wavelengths beyond those seen during training, thereby overcoming the shortcomings of prior surrogate approaches and markedly advancing the generalization capabilities of deep-learning-based electromagnetic simulators. We evaluate our surrogate solver on three representative structures, namely a single layer metasurface, a multilayer metasurface, and a freeform waveguide, each with distinct operating mechanisms and geometric complexity. Across their target bandwidths, our model reduces average electric-field prediction error relative to the state of the art [17] by 80.3%, 56.8%, and 32.3% for three examples, and lowers model complexity by 86.9% compared with FNO2d [32]. Extension of our framework to macroscopic length scales is attainable through integration with emerging fast-solver techniques [74–76]. Although enlarging the simulation domain naturally raises concerns about model complexity and training efficiency, contemporary deep-learning architectures comprising hundreds of millions—or even billions—of parameters can be trained effectively on modern hardware. Consequently, developing a large-scale, three-dimensional extension of our surrogate solver represents a significant advance toward replacing conventional numerical methods for solving Maxwell's equations.

Author Contributions

Joonhyuk Seo: conceptualization, methodology, software, data curation, validation, visualization, writing – review and editing, writing – original draft, investigation, formal analysis. **Chanik Kang:** conceptualization, methodology, software, data curation, validation, visualization, writing – original draft, writing – review and editing, investigation, formal analysis. **Dongjin Seo:** conceptualization. **Haejun Chung:** supervision, funding acquisition, visualization, writing – review and editing, project administration, validation, writing – original draft.

Acknowledgments

The authors thank Jiwoo Kim and Dongyoung Kim (Hanyang University) for providing simulation data across varying permittivity ranges used in the analysis of this work. This work was supported by National Research Foundation of Korea (NRF) grants funded by the Korean government (MSIT) (RS-2024-00338048, RS-2024-00414119, RS-2026-25620621, RS-2026-25519864, RS-2025-02218723); by the Global Research Support Program in the Digital Field (RS-2024-00412644), funded by the Republic of Korea's MSIT and supervised by the Institute for Information and Communications Technology Promotion (IITP); by the Culture, Sports, and Tourism R&D Program through the Korea Creative Content Agency, funded by the Ministry of Culture, Sports, and Tourism in 2024 (RS-2024-00332210); by the Artificial Intelligence Graduate School Program (No. RS-2020-II201373, Hanyang University), supervised by the IITP; by the artificial intelligence semiconductor support program to nurture the best talents (IITP-(2026)-RS-2023-00253914), funded by the Korean government; and by the MSIT (RS-2025-02283217). This research was supported by the ANCHOR through the Seoul ANCHOR Center, funded by the Ministry of Education (MOE) and the Seoul Metropolitan Government (2026-ANCHOR-01-027-01).

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

All data supporting the findings of this study are available within the article and its Supporting Information S1. Source data are provided with this paper. And the photonic structure datasets are available at <https://drive.google.com/drive/folders/1YRuzbOaIv3q6AjBgva2H2cizaP1Rh?usp=sharing>.

The code used for the experiments presented in this paper is available at https://github.com/yhy258/WINO_interp.

References

1. F. Wen, J. Jiang, and J. A. Fan, “Progressive-Growing of Generative Adversarial Networks for Metasurface Optimization,” preprint arXiv:1911.13029 (2019).
2. M. Chen, R. Lupoiu, C. Mao, et al., “High Speed Simulation and Freeform Optimization of Nanophotonic Devices With Physics-Augmented Deep Learning,” *ACS Photonics* 9, no. 9 (2022): 3110–3123, <https://doi.org/10.1021/acsphotonics.2c00876>.
3. D. Seo, S. Um, S. Lee, J. C. Ye, and H. Chung, “Physics-Guided and Fabrication-Aware Inverse Design of Photonic Devices Using Diffusion Models,” preprint arXiv:2504.17077 (2025).
4. J. Seo, C. Kang, D. Seo, and H. Chung, “High-Speed Multiwavelength Adjoint Optimization With Surrogate Solver,” in *2024 Conference on Lasers and Electro-Optics Pacific Rim (CLEO-PR)* (IEEE, 2024), 1–2.

5. C. Kang, J. Seo, I. Jang, and H. Chung, "Adjoint Method-Based Fourier Neural Operator Surrogate Solver for Wavefront Shaping in Tunable Metasurfaces," *iScience* 28, no. 1 (2025): 111545, <https://doi.org/10.1016/j.isci.2024.111545>.
6. J. Seo, J. Jo, J. Kim, et al., "Deep-Learning-Driven End-to-End Metasens Imaging," *Advanced Photonics* 6, no. 6 (2023): 066002, <https://doi.org/10.1117/1.ap.6.6.066002>.
7. E. Tseng, S. Colburn, J. Whitehead, et al., "Neural Nano-Optics for High-Quality Thin Lens Imaging," *Nature Communications* 12, no. 1 (2021): 6493, <https://doi.org/10.1038/s41467-021-26443-0>.
8. D. Deb, Z. Jiao, R. Sims, et al., "Fouriernet Enable the Design of Highly Non-Local Optical Encoders for Computational Imaging," *Advances in Neural Information Processing Systems* 35 (2022): 25224–25236, <https://doi.org/10.52202/068431-1829>.
9. M. Tan, X. Xu, A. Boes, et al., "Photonic Signal Processor Based on a Kerr Microcomb for Real-Time Video Image Processing," *Communications Engineering* 2, no. 1 (2023): 94, <https://doi.org/10.1038/s44172-023-00135-7>.
10. Q. Fan, G. Zhou, T. Gui, C. Lu, and A. P. T. Lau, "Advancing Theoretical Understanding and Practical Performance of Signal Processing for Nonlinear Optical Communications Through Machine Learning," *Nature Communications* 11, no. 1 (2020): 3694, <https://doi.org/10.1038/s41467-020-17516-7>.
11. Z. Huang, P. Wang, J. Liu, et al., "All-Optical Signal Processing of Vortex Beams With Diffractive Deep Neural Networks," *Physical Review Applied* 15, no. 1 (2021): 014037, <https://doi.org/10.1103/physrevapplied.15.014037>.
12. X. Lin, Y. Rivenson, N. T. Yardimci, et al., "All-Optical Machine Learning Using Diffractive Deep Neural Networks," *Science* 361, no. 6406 (2018): 1004–1008, <https://doi.org/10.1126/science.aat8084>.
13. X. Xu, M. Tan, B. Corcoran, et al., "11 Tops Photonic Convolutional Accelerator for Optical Neural Networks," *Nature* 589, no. 7840 (2021): 44–51, <https://doi.org/10.1038/s41586-020-03063-0>.
14. T. Fu, J. Zhang, R. Sun, et al., "Optical Neural Networks: Progress and Challenges," *Light: Science & Applications* 13, no. 1 (2024): 263, <https://doi.org/10.1038/s41377-024-01590-3>.
15. J. Jiang, M. Chen, and J. A. Fan, "Deep Neural Networks for the Evaluation and Design of Photonic Devices," *Nature Reviews Materials* 6, no. 8 (2021): 679–700, <https://doi.org/10.1038/s41578-020-00260-1>.
16. S. So, T. Badloe, J. Noh, J. Bravo-Abad, and J. Rho, "Deep Learning Enabled Inverse Design in Nanophotonics," *Nanophotonics* 9, no. 5 (2020): 1041–1057, <https://doi.org/10.1515/nanoph-2019-0474>.
17. J. Gu, Z. Gao, C. Feng, et al., "NeuroLight: A Physics-Agnostic Neural Operator Enabling Parametric Photonic Device Simulation," *Advances in Neural Information Processing Systems* 35 (2022): 14623–14636, <https://doi.org/10.52202/068431-1063>.
18. I. Malkiel, M. Mrejen, A. Nagler, U. Arieli, L. Wolf, and H. Suchowski, "Plasmonic Nanostructure Design and Characterization via Deep Learning," *Light: Science & Applications* 7, no. 1 (2018): 60, <https://doi.org/10.1038/s41377-018-0060-7>.
19. J. Peurifoy, Y. Shen, Li Jing, et al., "Nanophotonic Particle Simulation and Inverse Design Using Artificial Neural Networks," *Science Advances* 4, no. 6 (2018): eaar4206, <https://doi.org/10.1126/sciadv.aar4206>.
20. I. Sajedian, J. Kim, and J. Rho, "Finding the Optical Properties of Plasmonic Structures by Image Processing Using a Combination of Convolutional Neural Networks and Recurrent Neural Networks," *Microsystems & Nanoengineering* 5, no. 1 (2019): 27, <https://doi.org/10.1038/s41378-019-0069-y>.
21. Lu Lu, R. Pestourie, W. Yao, Z. Wang, F. Verdugo, and S. G. Johnson, "Physics-Informed Neural Networks With Hard Constraints for Inverse Design," *SIAM Journal on Scientific Computing* 43, no. 6 (2021): B1105–B1132, <https://doi.org/10.1137/21m1397908>.
22. M. Zhelyeznyakov, J. Fröch, A. Wirth-Singh, et al., "Large Area Optimization of Meta-Lens via Data-Free Machine Learning," *Communications Engineering* 2, no. 1 (2023): 60, <https://doi.org/10.1038/s44172-023-00107-x>.
23. S. Piao, H. Gu, A. Wang, and P. Qin, "A Domain-Adaptive Physics-Informed Neural Network for Inverse Problems of Maxwell's Equations in Heterogeneous Media," *IEEE Antennas and Wireless Propagation Letters* 23, no. 10 (2024): 2905–2909, <https://doi.org/10.1109/LAWP.2024.3413851>.
24. J. Lim and D. Psaltis, "Maxwellnet: Physics-Driven Deep Neural Network Training Based on Maxwell's Equations," *APL Photonics* 7, no. 1 (2022): 011301, <https://doi.org/10.1063/5.0071616>.
25. N. Kovachki, Z. Li, B. Liu, et al., "Neural Operator: Learning Maps Between Function Spaces With Applications to PDES," *Journal of Machine Learning Research* 24, no. 89 (2023): 1–97, <https://jmlr.org/papers/v24/21-1524.html>.
26. K. Azizzadenesheli, N. Kovachki, Z. Li, M. Liu-Schiaffini, J. Kos-saifi, and A. Anandkumar, "Neural Operators for Accelerating Scientific Simulations and Design," *Nature Reviews Physics* 6, no. 5 (2024): 1–9, <https://doi.org/10.1038/s42254-024-00712-5>.
27. S. Cao, "Choose a Transformer: Fourier or Galerkin," *Advances in Neural Information Processing Systems* 34 (2021): 24924–24940.
28. H. Zhongkai, Z. Wang, H. Su, et al., "Gnot: A General Neural Operator Transformer for Operator Learning," in *International Conference on Machine Learning* (PMLR, 2023), 12556–12569.
29. Z. Li, K. Meidani, and A. Barati Farimani, "Transformer for Partial Differential Equations' Operator Learning," preprint arXiv:2205.13671 (2022).
30. Z. Li, D. Shu, and A. B. Farimani, "Scalable Transformer for PDE Surrogate Modeling," *Advances in Neural Information Processing Systems* 36 (2024), <https://doi.org/10.52202/075280-1216>.
31. B. Raonic, R. Molinaro, T. Rohner, S. Mishra, and E. de Bezenac, "Convolutional Neural Operators," in *ICLR 2023 Workshop on Physics for Machine Learning* (2023).
32. Z. Li, N. Kovachki, K. Azizzadenesheli, et al., "Fourier Neural Operator for Parametric Partial Differential Equations," preprint arXiv:2010.08895 (2020).
33. A. Tran, A. Mathews, L. Xie, and C. Soon Ong, "Factorized Fourier Neural Operators," preprint arXiv:2111.13802 (2021).
34. Z. Li, D. Z. Huang, B. Liu, and A. Anandkumar, "Fourier Neural Operator With Learned Deformations for PDES on General Geometries," *Journal of Machine Learning Research* 24, no. 388 (2023): 1–26, <https://www.jmlr.org/papers/v24/23-0064.html>.
35. Z. Li, N. Kovachki, C. Choy, et al., "Geometry-Informed Neural Operator for Large-Scale 3d PDES," *Advances in Neural Information Processing Systems* 36 (2024), <https://doi.org/10.52202/075280-1556>.
36. Y. Augenstein, T. Repan, and C. Rockstuhl, "Neural Operator-Based Surrogate Solver for Free-Form Electromagnetic Inverse Design," *ACS Photonics* 10, no. 5 (2023): 1547–1557, <https://doi.org/10.1021/acsp Photonics.3c00156>.
37. C. Mao, R. Lupoiu, T. Dai, M. Chen, and J. A. Fan, "Towards General Neural Surrogate Solvers With Specialized Neural Accelerators," preprint arXiv:2405.02351 (2024).
38. J. D. Jackson, *Classical Electrodynamics* (John Wiley & Sons, 2012).

39. A. Vaswani, N. Shazeer, N. Parmar, et al., "Attention Is all You Need," *Advances in Neural Information Processing Systems* 30 (2017), https://proceedings.neurips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.
40. B. D. Guenther, *Modern Optics* (Oxford University Press, 2018).
41. X. Zhang, K. Kwon, J. Henriksson, J. Luo, and M. C. Wu, "A Large-Scale Microelectromechanical-Systems-Based Silicon Photonics Lidar," *Nature* 603, no. 7900 (2022): 253–258, <https://doi.org/10.1038/s41586-022-04415-8>.
42. A. Lukashchuk, H. K. Yildirim, A. Bancora, et al., "Photonic-Electronic Integrated Circuit-Based Coherent LiDAR Engine," *Nature Communications* 15, no. 1 (2024): 3134, <https://doi.org/10.1038/s41467-024-47478-z>.
43. H. Chung and O. D. Miller, "Tunable Metasurface Inverse Design for 80% Switching Efficiencies and 144 Angular Deflection," *Acs Photonics* 7, no. 8 (2020): 2236–2243, <https://doi.org/10.1021/acsp Photonics.0c00787>.
44. S. Choi, C. Jang, L. Douglas, and G. Wetzstein, "Synthetic Aperture Waveguide Holography for Compact Mixed-Reality Displays With Large Étendue," *Nature Photonics* 19, no. 8 (2025): 854–863, <https://doi.org/10.1038/s41566-025-01718-w>.
45. S. Moon, S. Kim, J. Kim, C.-K. Lee, and J. Rho, "Single-Layer Waveguide Displays Using Achromatic Metagratings for Full-Colour Augmented Reality," *Nature Nanotechnology* 20, no. 6 (2025): 1–8, <https://doi.org/10.1038/s41565-025-01887-3>.
46. M. Gopakumar, G.-Y. Lee, S. Choi, et al., "Full-Colour 3d Holographic Augmented-Reality Displays With Metasurface Waveguides," *Nature* 629, no. 8013 (2024): 791–797, <https://doi.org/10.1038/s41586-024-07386-0>.
47. J. Jo, S. Byun, M. Bae, J. Wang, H. Chung, and S. Kim, "Spin-Orbit Coupling for Optical Vortex Generation in van der Waals Materials," preprint arXiv:2408.12061 (2024).
48. A. D. White, L. Su, D. I. Shahrar, et al., "Inverse Design of Optical Vortex Beam Emitters," *Acs Photonics* 10, no. 4 (2022): 803–807, <https://doi.org/10.1021/acsp Photonics.2c01007>.
49. A. I. Kuznetsov, M. L. Brongersma, J. Yao, et al., "Roadmap for Optical Metasurfaces," *ACS photonics* 11, no. 3 (2024): 816–865, <https://doi.org/10.1021/acsp Photonics.3c00457>.
50. H. Chung and D. M. Owen, "High-Na Achromatic Metalenses by Inverse Design," *Optics Express* 28, no. 5 (2020): 6945–6965, <https://doi.org/10.1364/oe.385440>.
51. W. T. Chen, A. Y. Zhu, and F. Capasso, "Flat Optics With Dispersion-Engineered Metasurfaces," *Nature Reviews Materials* 5, no. 8 (2020): 604–620, <https://doi.org/10.1038/s41578-020-0203-3>.
52. J.-S. Park, S. Zhang, A. She, et al., "All-Glass, Large Metalens at Visible Wavelength Using Deep-Ultraviolet Projection Lithography," *Nano Letters* 19, no. 12 (2019): 8673–8682, <https://doi.org/10.1021/acs.nanolett.9b03333>.
53. E. Bayati, R. Pestourie, S. Colburn, Z. Lin, S. G. Johnson, and A. Majumdar, "Inverse Designed Metalenses With Extended Depth of Focus," *ACS photonics* 7, no. 4 (2020): 873–878, <https://doi.org/10.1021/acsp Photonics.9b01703>.
54. W. H. Tyler, I. A. D. Williamson, M. Minkov, and S. Fan, "Forward-Mode Differentiation of Maxwell's Equations," *ACS Photonics* 6, no. 11 (2019): 3010–3016, <https://doi.org/10.1021/acsp Photonics.9b01238>.
55. J. Ho, A. Jain, and P. Abbeel, "Denoising Diffusion Probabilistic Models," *Advances in Neural Information Processing Systems* 33 (2020): 6840–6851.
56. J. K. Gupta and J. Brandstetter, "Towards Multi-Spatiotemporal-Scale Generalized PDE Modeling," preprint arXiv:2209.15616 (2022).
57. P. Lalanne and P. Chavel, "Metalenses at Visible Wavelengths: Past, Present, Perspectives," *Laser & Photonics Reviews* 11, no. 3 (2017): 1600295, <https://doi.org/10.1002/lpor.201600295>.
58. X. Zou, Y. Zhang, R. Lin, et al., "Pixel-Level Bayer-Type Colour Router Based on Metasurfaces," *Nature Communications* 13, no. 1 (2022): 3288, <https://doi.org/10.1038/s41467-022-31019-7>.
59. F. Presutti and F. Monticone, "Focusing on Bandwidth: Achromatic Metalens Limits," *Optica* 7, no. 6 (2020): 624–631, <https://doi.org/10.1364/optica.389404>.
60. A. Arbabi and A. Faraon, "Fundamental Limits of Ultrathin Metasurfaces," *Scientific Reports* 7, no. 1 (2017): 43722, <https://doi.org/10.1038/srep43722>.
61. G. Roberts, C. Ballew, T. Zheng, et al., "3d-Patterned Inverse-Designed Mid-Infrared Metaoptics," *Nature Communications* 14, no. 1 (2023): 2768, <https://doi.org/10.1038/s41467-023-38258-2>.
62. H. Wang, W. Zhang, D. Ladika, et al., "Two-Photon Polymerization Lithography for Optics and Photonics: Fundamentals, Materials, Technologies, and Applications," *Advanced Functional Materials* 33, no. 39 (2023): 2214211, <https://doi.org/10.1002/adfm.202214211>.
63. S. Lee, J. Hong, J. Kang, et al., "Inverse Design of Color Routers in CMOS Image Sensors: Toward Minimizing Interpixel Crosstalk," *Nanophotonics* 13, no. 20 (2024): 3895–3914, <https://doi.org/10.1515/nanoph-2024-0269>.
64. J. Jo, M. Lee, S. Lee, M. Bae, C. Kang, and H. Chung, "Inverse-Designed Metasurfaces for Wavefront Restoration in Under-Display Camera Systems," *Nanophotonics* 14, no. 17 (2025): 2963–2978, <https://doi.org/10.1515/nanoph-2025-0242>.
65. D. Seo, D. W. Nam, J. Park, Y. P. Chan, and M. S. Jang, "Structural Optimization of a One-Dimensional Freeform Metagrating Deflector via Deep Reinforcement Learning," *Acs Photonics* 9, no. 2 (2021): 452–458, <https://doi.org/10.1021/acsp Photonics.1c00839>.
66. A. Y. Piggott, J. Lu, K. G. Lagoudakis, J. Petykiewicz, T. M. Babinec, and J. Vučković, "Inverse Design and Demonstration of a Compact and Broadband on-Chip Wavelength Demultiplexer," *Nature Photonics* 9, no. 6 (2015): 374–377, <https://doi.org/10.1038/nphoton.2015.69>.
67. S. Ning, H. Zhu, C. Feng, et al., "Photonic-Electronic Integrated Circuits for High-Performance Computing and AI Accelerators," *Journal of Lightwave Technology* 42, no. 22 (2024): 7834–7859, <https://doi.org/10.1109/JLT.2024.3427716>.
68. S. Chen, Y. Zheng, Y. Xu, et al., "High-Bit-Efficiency Topological Optical Tensor Convolutional Accelerator Using Microcombs," *Laser & Photonics Reviews* 19, no. 9 (2025): 2401975, <https://doi.org/10.1002/lpor.202401975>.
69. K. Kwon, H.J. Heo, D. Lee, et al., "Photon-Pair Generation Using Inverse-Designed Thin-Film Lithium Niobate Mode Converters," *APL Photonics* 9, no. 5 (2024): 056108, <https://doi.org/10.1063/5.0192026>.
70. C. Kim, M. Bae, M. Choi, et al., "Freeform Thin-Film Lithium Niobate Mode Converter for Photon-Pair Generation," *Nanophotonics* 14, no. 11 (2025): 1949–1960, <https://doi.org/10.1515/nanoph-2024-0515>.
71. X. He, D. Sun, J. Chen, and Y. Shi, "Inverse Designed Grating Coupler With Low Loss and High Bandwidth on LNOI Platform," *IEEE Photonics Journal* 16, no. 1 (2024): 1–5, <https://doi.org/10.1109/jphot.2024.3351199>.
72. Y. Xie, M. Nie, and S.-W. Huang, "Inverse-Designed Broadband Low-Loss Grating Coupler on Thick Lithium-Niobate-on-Insulator Platform," *Applied Physics Letters* 124, no. 5 (2024): 051108, <https://doi.org/10.1063/5.0184413>.
73. R. Pestourie, Y. Mroueh, C. Rackauckas, P. Das, and S. G. Johnson, "Physics-Enhanced Deep Surrogates for Partial Differential Equations,"

Nature Machine Intelligence 5, no. 12 (2023): 1458–1465, <https://doi.org/10.1038/s42256-023-00761-y>.

74. W. H. Tyler, M. Minkov, V. Liu, Z. Yu, and S. Fan, “A Perspective on the Pathway Toward Full Wave Simulation of Large Area Metalenses,” *Applied Physics Letters* 119, no. 15 (2021): 150502, <https://doi.org/10.1063/5.0071245>.

75. Y. Mahlau, F. Schubert, K. Bethmann, et al., “A Flexible Framework for Large-Scale FDTD Simulations: Open-Source Inverse Design for 3D Nanostructures,” *Photonic and Phononic Properties of Engineered Nanostructures XV* 13377 (2025): 40–52, <https://doi.org/10.1117/12.3052639>.

76. W. Xue, H. Zhang, A. Gopal, V. Rokhlin, and D. M. Owen, “Full-wave Design of cm-scale Cylindrical Metasurfaces via Fast Direct Solvers,” preprint arXiv:2308.08569 (2023).

Supporting Information

Additional supporting information can be found online in the Supporting Information section.

Supporting Information S1: nap270242-sup-0001-suppl-data.pdf.