Check for updates

Nanophotonics

WILEY

RESEARCH ARTICLE OPEN ACCESS

Neural Adjoint Inverse Design of Three-Dimensional Nanophotonic Devices

Chanik Kang[1] | Hyewon Suk[2] | Ikbeom Jang[3] | Haejun Chung[1,2,4]

[1]Department of Artificial Intelligence, Hanyang University, Seoul, Korea | [2]Department of Artificial Intelligence Semiconductor Engineering, Hanyang University, Seoul, Korea | [3]Department of Computer Engineering, Hankuk University of Foreign Studies, Yongin, Korea | [4]Department of Electronic Engineering, Hanyang University, Seoul, Korea

Correspondence: Ikbeom Jang (ijang@hufs.ac.kr) | Haejun Chung (haejun@hanyang.ac.kr)

Received: 19 May 2026 | **Revised:** 6 October 2026 | **Accepted:** 9 October 2026

Keywords: adjoint method | fourier neural operator | inverse design | meta-optics | scientific machine learning | surrogate modeling

ABSTRACT

Volumetric meta-optics have emerged as a compact high-performance platform for advanced imaging, color routing in image sensors, and on-chip wavefront control. Realizing such devices depends on solving a high-dimensional nonconvex inverse-design problem governed by Maxwell's equations. Adjoint optimization is the standard route to gradients, yet broadband three-dimensional (3D) design typically requires thousands of full-wave forward/adjoint simulations, making industrial-scale optimization prohibitively slow. To address this bottleneck, a neural adjoint method is introduced, in which a solver-supervised Fourier neural operator (FNO) directly predicts dense, per-voxel 3D adjoint gradient fields from a voxelized permittivity volume. To preserve the sharp, resonance-driven sensitivity peaks that single-pass operator networks tend to oversmooth, a stage-wise FNO is proposed, in which successive stages refine residual errors with progressively higher-frequency emphasis. A domain-specific dataset is curated from paired forward/adjoint finite-difference time-domain (FDTD) simulations spanning three volumetric tasks: spectral sorting (color routers), achromatic focusing (metalenses), and waveguide mode conversion. Across all benchmarks, the surrogate-driven optimization attains figures of merit within 86%–99% of rigorous FDTD-based adjoint optimization whilst reducing end-to-end design time from hours to seconds, with speedups exceeding $10^4\times$. These results establish a pathway toward fast large-scale volumetric meta-optical design enabled by neural-operator surrogates.

1 | Introduction

Optical-system miniaturization has accelerated the commercialization of meta-optics, marking a major transition from academic research to industrial-scale manufacturing [1–4]. In particular, the display and semiconductor industries are increasingly adopting 3D volumetric meta-optics, with color routers for complementary metal-oxide-semiconductor (CMOS) image sensors being a prominent example, to address the fundamental efficiency constraints of conventional dye-based color filters [5–7]. Unlike thin diffractive components, volumetric meta-optical devices leverage complex multiple scattering within micrometer-scale volumes to enable precise control of wavefronts and spectra [8–10]. For deployment in consumer electronics, these devices must meet stringent performance requirements across broad wavelength ranges, necessitating an unprecedented degree of precision in the design and optimization process [11].

However, industrial deployment of these devices remains limited by substantial computational overhead. The inverse design of 3D volumetric meta-optics is typically a high-dimensional nonconvex optimization problem constrained by Maxwell's equations [12]. The adjoint method is the standard approach for computing gradients, requiring two full-wave

This is an open access article under the terms of the Creative Commons Attribution 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.

simulations per iteration (one forward and one adjoint) [13]. Although adjoint optimization is computationally efficient in terms of gradient evaluation, the overall cost of rigorous 3D solvers such as finite-difference time-domain (FDTD) grows rapidly with problem size [14]. In addition, commercial designs must be optimized over the full visible band, which typically requires separate simulations for each wavelength or color channel [15]. As a result, a single optimization run can involve tens of thousands of full-wave simulations and take days to weeks to complete. This turnaround time poses a major obstacle for high-throughput industrial research and development, where rapid prototyping and iterative refinement are critical.

To mitigate this cost, deep learning has emerged as a promising alternative. Generative models (GANs [16–18], diffusion [19–22]) offer rapid one-shot designs, but they suffer from a precision-speed trade-off [23]. Conversely, numerical optimization guarantees physical optimality but is constrained by the extreme slowness of the solver [24].

To resolve this precision-speed trade-off, we propose the neural adjoint method, a surrogate framework that retains the precision of iterative gradient-based optimization whilst drastically reducing its solver cost. Our approach retains the conventional gradient-based optimization loop and replaces only the per-iteration gradient evaluation, namely the paired forward/adjoint full-wave solves, with a Fourier neural operator (FNO)-based surrogate [25, 26]. In contrast to end-to-end generative approaches, which substitute the entire iterative design pipeline with a learned generator, this decouples the per-iteration cost of optimization from the cost of full-wave simulation. The 3D volumetric gradient field is predicted within a second, supporting a solver-reduced optimization loop that drastically shortens the design turnaround time.

The proposed framework addresses the challenges of volumetric inverse design through three key contributions. First, inverse design is formulated around dense 3D adjoint sensitivities, and a neural operator is trained to predict full volumetric adjoint gradient fields directly from voxelized devices, enabling a single surrogate architecture, trained separately for each task, to serve heterogeneous meta-optical objectives. Second, to counteract the low-frequency bias of single-pass operator networks, a stage-wise Fourier neural operator (SW-FNO) is introduced, in which successive stages refine residual errors with progressively higher-frequency emphasis, thereby improving the recovery of localized high-magnitude sensitivity peaks that are critical for topology refinement. Third, the inferred gradients are integrated into a gradient-based optimization loop, enabling rapid design iteration without per-iteration full-wave Maxwell solves. The framework is validated on three representative 3D inverse-design tasks: RGB color routing [7, 27, 28], achromatic metalens design [29–31], and waveguide mode conversion [32–35]. Across all tasks, the neural adjoint method achieves performance comparable to rigorous FDTD-based adjoint optimization whilst reducing end-to-end design time by several orders of magnitude, with speedups of up to $\sim 10^4$–$10^5\times$. By mitigating the bottleneck of repeated full-wave simulations, these results suggest a practical pathway toward rapid large-scale 3D meta-optical design enabled by AI-accelerated scientific computing.

The present work builds on several lines of inquiry. Adjoint-based optimization is widely used for large-scale photonic and meta-optical inverse design, enabling efficient gradient computation with only two full-wave simulations per objective evaluation [13, 36–41]; however, the overall runtime remains dominated by repeated full-wave simulations across many iterations and wavelengths, motivating efforts to reduce the cost of solver-in-the-loop optimization. Neural operators, including FNO, have emerged as mesh-agnostic surrogates for PDE-governed systems, such as electromagnetics [42, 43], with most prior work focused on predicting fields or spectra to reduce the cost of repeated simulations [12, 44–48]; more recently, gradient learning has been investigated as an alternative, in which adjoint sensitivities required by gradient-based optimizers are directly regressed [49–52]. A complementary line of work introduces fabrication awareness directly into the objective itself through density filtering, projection schemes, length-scale control, and three-field robust formulations [53–56]; because these formulations modify the objective rather than the gradient evaluation routine, they compose orthogonally with the present approach, which focuses on accelerating the gradient evaluation itself and leaves the choice of objective formulation to the user.

Building on gradient learning with neural operators, the present work focuses on 3D volumetric inverse design across diverse physics-based objectives, including focusing, spectral routing, and mode conversion. Extending neural adjoint surrogates from low-dimensional or 2D problems to 3D introduces practical challenges related to representational capacity, preservation of high-spatial-frequency sensitivity, and scalability. These challenges are addressed through a volumetric multi-objective surrogate formulation that predicts 3D gradient fields on voxel grids and generalizes across multiple classes of objectives, a stage-wise refinement scheme with progressive frequency emphasis that better preserves sharp sensitivity peaks critical for faster convergence, and a controllable synthetic data-generation pipeline that produces diverse, smoothly varying topologies together with solver-generated gradients for supervised training.

2 | Background and Problem Formulation

This section formulates the volumetric inverse design problem as a constrained optimization task governed by Maxwell's equations, outlines the dominant computational cost, and contrasts the present work with prior efforts on adjoint optimization and neural surrogates.

2.1 | Inverse Design for Diverse Physical Objectives

Inverse design in meta-optics and nanophotonics seeks a device parameterization (e.g., a 2D/3D permittivity distribution) that maximizes or minimizes a physics-based objective function [14]. A common approach is gradient-based adjoint optimization: given an objective functional $\mathcal{J}(\mathbf{x})$ evaluated by solving Maxwell's equations, the adjoint method computes $\nabla_{\mathbf{x}}\mathcal{J}$ using only two full-wave simulations, a forward solve to obtain the electromagnetic fields and an adjoint solve to back-propagate sensitivity, enabling scalable updates even when $\mathbf{x}$ is high-

21928614, 2026, 20, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/nap2.70334 by Hanyang University Library, Wiley Online Library on [08/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

dimensional [13]. In this work, this conventional formulation is used as the reference pipeline, and the focus is on accelerating the expensive field/gradient evaluations.

As illustrated in Figure 1, volumetric inverse design is considered within a three-dimensional physical domain $\Omega \subset \mathbb{R}^3$. The design variable is the spatial permittivity distribution $\epsilon(\mathbf{r})$, which is discretized as a voxel grid $\mathbf{x} \in \mathbb{R}^N$ with $N = N_x N_y N_z$ degrees of freedom. For a given design $\mathbf{x}$ and wavelength λ, a full-wave Maxwell solver computes the complex electric field $\mathbf{E}(\mathbf{r}; \mathbf{x}, \lambda)$, from which task-specific figures of merit (FoMs) are defined.

Three representative classes of meta-optics/nanophotonics objectives are studied as visualized in Figure 2. For the metalens task, the field intensity within a prescribed focal volume $V_f \subset \Omega$ is maximized across K target wavelengths $\{\lambda_k\}_{k=1}^K$:

$$\mathcal{J}_{\text{focus}}(\mathbf{x}) = \sum_{k=1}^{K} \int_{V_f} \|\mathbf{E}(\mathbf{r}; \mathbf{x}, \lambda_k)\|^2 \, d\mathbf{r}. \tag{1}$$

For the color router task, K target wavelengths $\{\lambda_k\}_{k=1}^K$ are spatially sorted and the field intensity integrated over wavelength-specific target regions $\{R_k\}_{k=1}^K$ is maximized:

$$\mathcal{J}_{\text{sort}}(\mathbf{x}) = \sum_{k=1}^{K} \int_{R_k} \|\mathbf{E}(\mathbf{r}; \mathbf{x}, \lambda_k)\|^2 \, d\mathbf{r}. \tag{2}$$

For the waveguide mode converter, the modal overlap between the output field on the cross-section S_{out} and a target mode profile $\mathbf{\Psi}_{\text{tgt}}$ is maximized:

$$\mathcal{J}_{\text{mode}}(\mathbf{x}) = \left| \int_{S_{\text{out}}} \mathbf{E}(\mathbf{r}; \mathbf{x}, \lambda) \cdot \mathbf{\Psi}^*_{\text{tgt}}(\mathbf{r}) \, d\mathbf{r} \right|^2. \tag{3}$$

Although the objectives differ across tasks, they are all governed by the same frequency-domain Maxwell equation:

$$\nabla \times \nabla \times \mathbf{E}(\mathbf{r}) - k_0^2 \epsilon(\mathbf{r}) \mathbf{E}(\mathbf{r}) = -i\omega\mu_0 \mathbf{J}_{\text{src}}(\mathbf{r}), \tag{4}$$

where $k_0 = 2\pi/\lambda$.

2.2 | The Adjoint-Solver Bottleneck

To optimize the objectives above, gradient-based inverse design typically employs the adjoint method (also known as the adjoint-variable method) to compute the sensitivity of the FoM with respect to the design parameters [13, 57]. Let $\mathbf{g} = \nabla_{\mathbf{x}} \mathcal{J}$

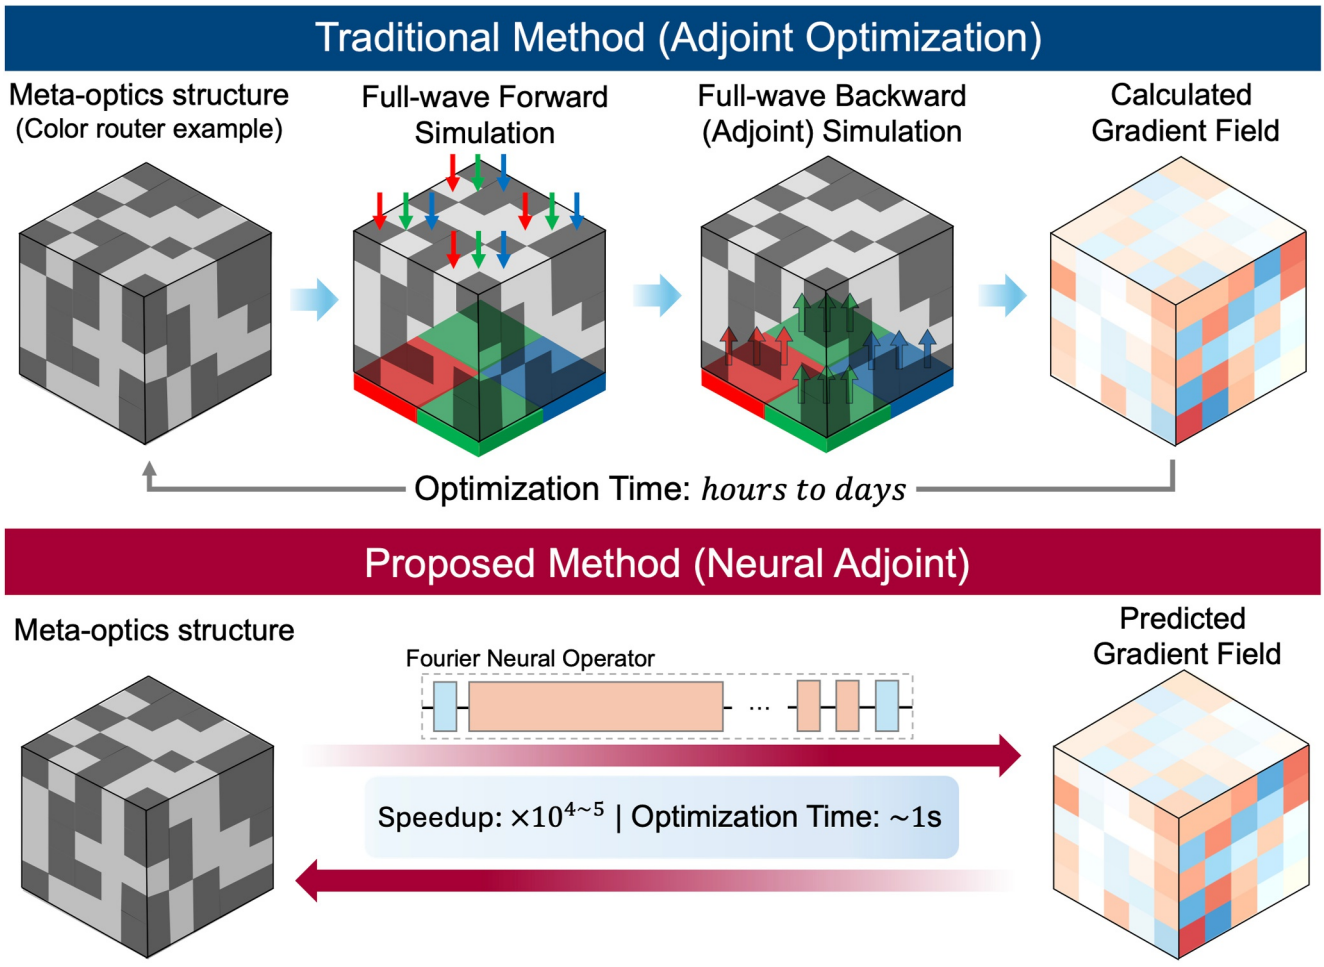

FIGURE 1 | Breaking the simulation bottleneck in 3D meta-optics. (Top) Conventional adjoint optimization relies on iterative full-wave high-fidelity simulations (Maxwell solvers) to compute gradients, requiring hours to days per design. (Bottom) The proposed neural adjoint method replaces the solver with a Fourier neural operator, inferring volumetric gradient fields in seconds.

21928614, 2026, 20, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/nap2.70334 by Hanyang University Library, Wiley Online Library on [08/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

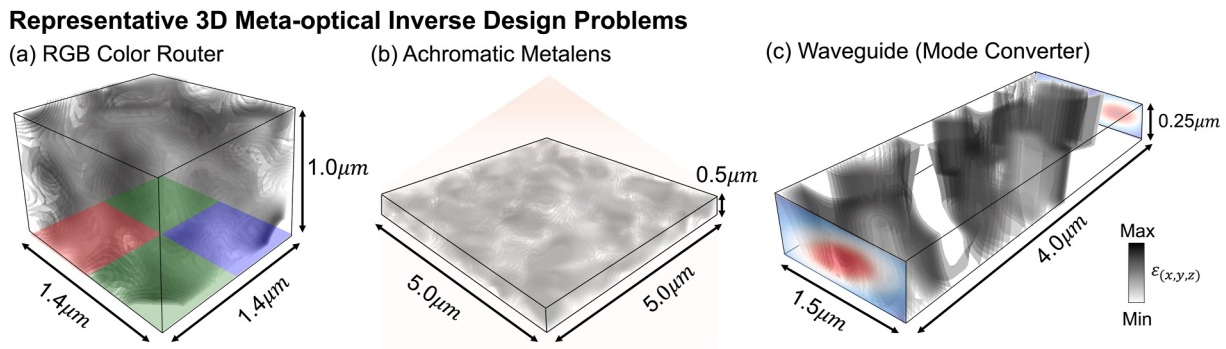

FIGURE 2 | 3D meta-optical inverse-design tasks considered in this work: (a) spectral sorting (color router), (b) light focusing (metalens), and (c) guided-wave mode conversion (waveguide mode converter).

denote the gradient. In electromagnetic adjoint formulations, $\mathbf{g}$ can be expressed as a point-wise dot product between the forward field $\mathbf{E}_{\text{fwd}}$ and the adjoint field $\mathbf{E}_{\text{adj}}$:

$$\mathbf{g}(\mathbf{r}) \propto \text{Re}\big[\mathbf{E}_{\text{fwd}}(\mathbf{r}) \cdot \mathbf{E}_{\text{adj}}(\mathbf{r})\big], \tag{5}$$

where the adjoint source term is determined by the particular FoM.

Importantly, independent of the chosen FoM, the dominant cost in adjoint-based optimization arises from repeatedly solving the full-wave Maxwell system in Equation (4) to obtain the forward and adjoint fields. In 3D, high-resolution, and broadband settings, these repeated full-wave computations become the primary bottleneck, rendering conventional iterative optimization prohibitively expensive in practice.

2.3 | Neural Surrogate for Adjoint Gradients

The objective is to approximate the adjoint-to-gradient mapping $\mathcal{G} : \mathbf{x} \mapsto \nabla_{\mathbf{x}}\mathcal{J}$ using a neural surrogate f_θ. Although the scalar FoM $\mathcal{J}$ differs across applications, the volumetric gradient field $\mathbf{g} = \nabla_{\mathbf{x}}\mathcal{J}$ encodes the local sensitivity information required by gradient-based optimizers. Accordingly, the problem is formulated as dense field prediction:

$$\hat{\mathbf{g}} = f_\theta(\mathbf{x}) \approx \nabla_{\mathbf{x}}\mathcal{J} \in \mathbb{R}^{N_x \times N_y \times N_z}. \tag{6}$$

Each task uses a single task-specific surrogate, and wavelength is not an explicit network input. Instead, wavelength information enters through the design objective. For the multi-wavelength FoMs of Equations (1) and (2), the supervision label is the gradient of the aggregated objective, $\mathbf{g} = \nabla_{\mathbf{x}}\mathcal{J} = \sum_{k=1}^{K} \nabla_{\mathbf{x}}\mathcal{J}_k$, where $\mathcal{J}_k$ denotes the single-wavelength summand at λ_k. By linearity of differentiation, this aggregation is exact and coincides with the field required by the design-update rule. In the remainder of this paper, f_θ is treated as a drop-in replacement for repeated adjoint solves within iterative optimization. The model architecture and training procedure are described in the Methodology section.

3 | Methodology

This section presents the neural adjoint method for accelerating volumetric meta-optical inverse design. As illustrated in Figure 3, a stage-wise Fourier neural operator (SW-FNO) is introduced that predicts 3D adjoint gradient fields via coarse-to-fine residual refinement, improving the recovery of localized high-sensitivity structures. A solver-reduced optimization loop is then described that updates the design using the inferred gradients, enabling rapid iterative refinement without repeatedly invoking full-wave forward/adjoint simulations.

3.1 | Stage-Wise Fourier Neural Operator (SW-FNO)

To map a voxelized permittivity distribution $\mathbf{x} \in \mathbb{R}^{N_x \times N_y \times N_z}$ to a dense adjoint gradient field $\mathbf{g}$, the surrogate must represent electromagnetic phenomena, including wave propagation, scattering, and resonance, with the latter producing sharp high-spatial-frequency sensitivity peaks. In practice, a single-stage FNO often under-represents these high-frequency components, resulting in over-smoothed gradient fields and, consequently, degraded optimization performance [58].

To address this limitation, a stage-wise refinement strategy is adopted that better matches the multiscale nature of 3D volumetric scattering [59]. The proposed SW-FNO performs progressive refinement: it first captures coarse, global trends, and then successively adds residual corrections to recover fine-scale high-spatial-frequency components [60]. Formally,

$$\hat{\mathbf{g}} = \hat{\mathbf{g}}_1 + \sum_{s=2}^{S} \Delta\hat{\mathbf{g}}_s, \quad \hat{\mathbf{g}}_1 = \mathcal{N}_1(\mathbf{x}),\ \Delta\hat{\mathbf{g}}_s = \mathcal{N}_s(\mathbf{x}) \approx \mathbf{r}_s, \tag{7}$$

where S is the number of stages and $\Delta\hat{\mathbf{g}}_s$ denotes the residual correction predicted at stage s. This additive coarse-to-fine decomposition encourages later stages to focus on what earlier stages miss, enabling sharper recovery of high-frequency gradient features whilst preserving global consistency of the predicted field. The three stages use nested, progressively enlarged Fourier-mode cutoffs rather than architecturally disjoint band-pass filters. The low/mid/high terminology refers to each stage's effective emphasis, which is induced by residual training. Because the frozen earlier stages already capture the content representable within their smaller cutoffs, the learning signal of each later stage concentrates on the modes newly admitted between the previous and the current cutoffs, and each stage's correction is concentrated in, though not strictly limited to, this newly admitted range. The enlarged cutoff also retains the lower modes, which allows later stages to refine residual

21928614, 2026, 20, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/nap2.70334 by Hanyang University Library, Wiley Online Library on [08/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

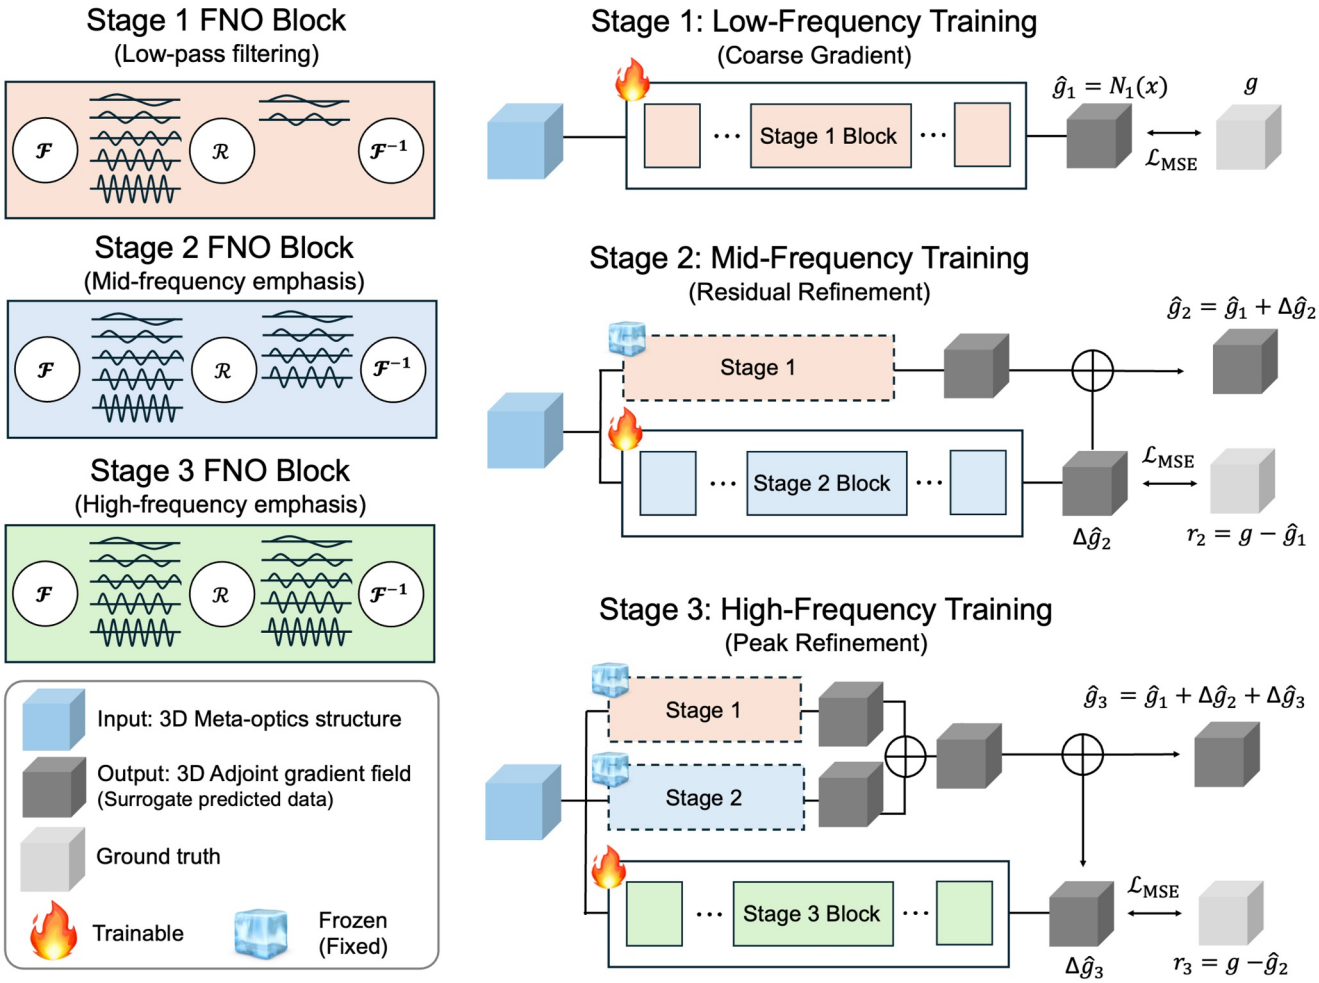

FIGURE 3 | Stage-wise Fourier neural operator (SW-FNO). Overview of the stage-wise training scheme for predicting dense 3D adjoint gradient fields from a voxelized meta-optics structure. Each stage employs an FNO block operating in the Fourier domain, with frequency emphasis progressively shifted from low to mid to high bands (left: simplified nested-cutoff FNO blocks). Training proceeds sequentially: Stage 1 learns a coarse low-frequency prediction; subsequent stages are trained on residual targets to refine missing components, whereas earlier stages are frozen. The final prediction is obtained by summing the stage outputs, enabling accurate reconstruction across the full frequency spectrum whilst preserving sharp high-frequency sensitivity peaks. Each stage retains all Fourier modes up to its per-stage cutoff; the low/mid/high labels denote each stage's effective frequency emphasis induced by residual training.

low-frequency errors left by earlier stages. The cutoffs follow a simple rule anchored to the single-stage baseline. The mid-stage cutoff equals the truncation cutoff of the single-stage FNO baseline. The Stage 1 cutoff retains roughly 60%–78% of the baseline cutoff per axis, so the first stage models only the coarse global envelope. The Stage 3 cutoff extends 20%–33% beyond the baseline, approaching the largest mode count representable on the voxel grid. For the color router, for example, the cutoffs are refs [6, 8, 10] and refs. [10, 12] modes per axis, and the per-task values are listed in Supporting Information S1: Table S1.

3.1.1 | *Stage 1: Global Physics Learning*

The first stage uses a 3D FNO to capture the dominant low-spatial-frequency global structure of the adjoint gradient field. These components reflect long-range interactions, wave propagation effects, and coarse sensitivity patterns throughout the volume. For the l-th Fourier layer, the feature update is defined as

$$\mathbf{v}_{l+1}(\mathbf{r}) = \sigma\big(W_l\,\mathbf{v}_l(\mathbf{r}) + \mathcal{F}^{-1}(R_l(\mathbf{k}) \odot \mathcal{F}(\mathbf{v}_l)(\mathbf{k}))(\mathbf{r})\big), \tag{8}$$

where $\mathcal{F}$ and $\mathcal{F}^{-1}$ denote the fast Fourier transform (FFT) and inverse FFT, W_l is a pointwise linear projection, and R_l is a complex-valued spectral weight tensor defined on the truncated set of Fourier modes (up to k_{base}). Operating in the Fourier domain yields a global convolution that captures nonlocal physics-driven correlations.

The Stage 1 predictor produces a coarse global estimate of the gradient field,

$$\hat{\mathbf{g}}_1 = \mathcal{N}_1(\mathbf{x}), \tag{9}$$

which provides a globally consistent initialization for subsequent residual stages and stabilizes training by avoiding the need to learn fine-scale corrections from scratch. With a limited mode budget (k_{base}) and the spectral bias of operator networks, $\hat{\mathbf{g}}_1$ may inevitably smooth out boundary-aligned variations and localized resonant features that require higher spectral resolution.

21928614, 2026, 20, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/nap2.70334 by Hanyang University Library, Wiley Online Library on [08/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

3.1.2 | *Stage 2: Mid-Frequency Residual Refinement*

Although Stage 1 captures dominant global sensitivity patterns, its prediction $\hat{\mathbf{g}}_1$ can under-represent mid-frequency structures around material interfaces, boundaries, and transition regions due to low-mode truncation and spectral bias. Therefore, Stage 2 predicts a residual correction on top of the frozen Stage 1 estimate, following the residual-learning principle [61]. The Stage 1 residual is defined as $\mathbf{r}_2 = \mathbf{g} - \hat{\mathbf{g}}_1$, where $\mathbf{g}$ denotes the ground-truth adjoint gradient field obtained from the FDTD solver. Stage 2 then learns this residual via

$$\Delta\hat{\mathbf{g}}_2 = \mathcal{N}_2(\mathbf{x}), \tag{10}$$

and the refined estimate is obtained as

$$\hat{\mathbf{g}}_2 = \hat{\mathbf{g}}_1 + \Delta\hat{\mathbf{g}}_2. \tag{11}$$

Because the frozen Stage 1 already captures the content representable within its smaller cutoff, the learning signal of $\mathcal{N}_2$ concentrates on the modes admitted between the Stage 1 and Stage 2 cutoffs (Supporting Information S1: Table S1).

3.1.3 | *Stage 3: High-Frequency Peak Refinement*

Even after Stage 2 recovers boundary-aligned mid-frequency residuals, the target adjoint gradient field may still exhibit sharp highly localized sensitivity peaks at higher spatial frequencies. These peaks often arise from abrupt spatial variations (e.g., resonance-induced confinement and interface discontinuities) and indicate fine-scale regions where topology updates must remain responsive.

Standard FNOs exhibit a low-frequency bias and, under finite mode truncation, tend to smooth spiky high-frequency components, which can underestimate gradient magnitudes near localized maxima. To better preserve these high-frequency extrema, Stage 3 predicts an additional residual correction on top of the frozen Stage 2 estimate. The Stage 2 residual is $\mathbf{r}_3 = \mathbf{g} - \hat{\mathbf{g}}_2$, and Stage 3 learns this residual as

$$\Delta\hat{\mathbf{g}}_3 = \mathcal{N}_3(\mathbf{x}), \tag{12}$$

yielding the final prediction

$$\hat{\mathbf{g}}_{\text{final}} = \hat{\mathbf{g}}_2 + \Delta\hat{\mathbf{g}}_3. \tag{13}$$

By explicitly modeling the remaining residual, Stage 3 improves the recovery of high-magnitude localized peaks that are critical for fine-scale topology refinement, helping to avoid attenuated sensitivities and to preserve reliable update directions. Similarly, the learning signal of $\mathcal{N}_3$ concentrates on the modes admitted between the Stage 2 and Stage 3 cutoffs (Supporting Information S1: Table S1).

3.2 | Learning Adjoint Gradients From Full-Wave Simulations

The surrogate f_θ is trained on paired samples of voxelized device representations and their corresponding adjoint gradients, $\{(\mathbf{x}_i, \mathbf{g}_i)\}$. Each $\mathbf{g}_i$ is obtained from a conventional adjoint formulation using the open-source FDTD package Meep [62]: for each design $\mathbf{x}_i$, a forward simulation is run to evaluate the objective and an adjoint simulation is run to compute the volumetric sensitivity field. The learning problem is formulated as dense regression. Given $\mathbf{x}$, the model predicts the full 3D gradient field $\hat{\mathbf{g}} = f_\theta(\mathbf{x})$ and is optimized to match the solver-generated $\mathbf{g}$. Dataset sizes are task-dependent, with 600 waveguide, 1100 metalens, and 3300 color-router pairs used for training and validation. Each training set was grown until the validation gradient fidelity entered the high-fidelity regime and further samples yielded diminishing gains (Supporting Information S1: Section 5.1, Table S3).

4 | Results and Discussion

The neural adjoint method is empirically evaluated on three volumetric meta-optical inverse-design tasks. The experimental study is organized around three questions. The first concerns generalization: whether a single neural framework can be applied to physically distinct inverse-design problems using a consistent architecture, without task-specific modifications. The second concerns efficiency: how much speedup the proposed method provides relative to conventional FDTD-based adjoint optimization. The third concerns component analysis: whether multistage refinement and high-frequency boosting improve gradient fidelity and final design performance, and what is the trade-off with model size and inference latency.

4.1 | Baselines and Evaluation Metrics

The neural adjoint method is compared against four baselines covering both physics-grounded optimization and learning-based surrogates. The first is the adjoint method (FDTD reference), a conventional adjoint optimization pipeline that computes gradients via paired forward/adjoint full-wave simulations and updates the design with a gradient-based optimizer; this provides the physically accurate benchmark FoM for comparison. The surrogate-driven pipeline shares this optimization loop, including the optimizer and all update settings, and differs only in the source of the gradient (Supporting Information S1: Section 6, Table S4). The second is the one-stage FNO, a vanilla Fourier neural operator that predicts the full 3D adjoint gradient field in a single pass without stage-wise refinement. The third is the triple-stage FNO (SW-FNO, the proposed model), which cascades three independent FNO blocks with progressively larger Fourier-mode budgets to enable coarse-to-fine refinement from low-to-high-frequency content. The fourth is a 3D U-Net, a volumetric encoder–decoder architecture with skip connections trained to predict the dense adjoint gradient field. For

TABLE 1 | Comparison of gradient prediction performance (cosine similarity, %).

	Color router		Metalens		Waveguide	
Model	Mean (std)	Max	Mean (std)	Max	Mean (std)	Max
U-net	61.4 (13.7)	86.4	53.6 (6.5)	69.1	32.6 (23.9)	71.9
FNO	66.3 (13.1)	89.7	73.5 (10.3)	90.0	46.4 (25.2)	80.5
Ours	82.3 (13.7)	99.8	87.6 (7.5)	98.2	87.9 (30.3)	99.4

21928614, 2026, 20, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/nap2.70334 by Hanyang University Library, Wiley Online Library on [08/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

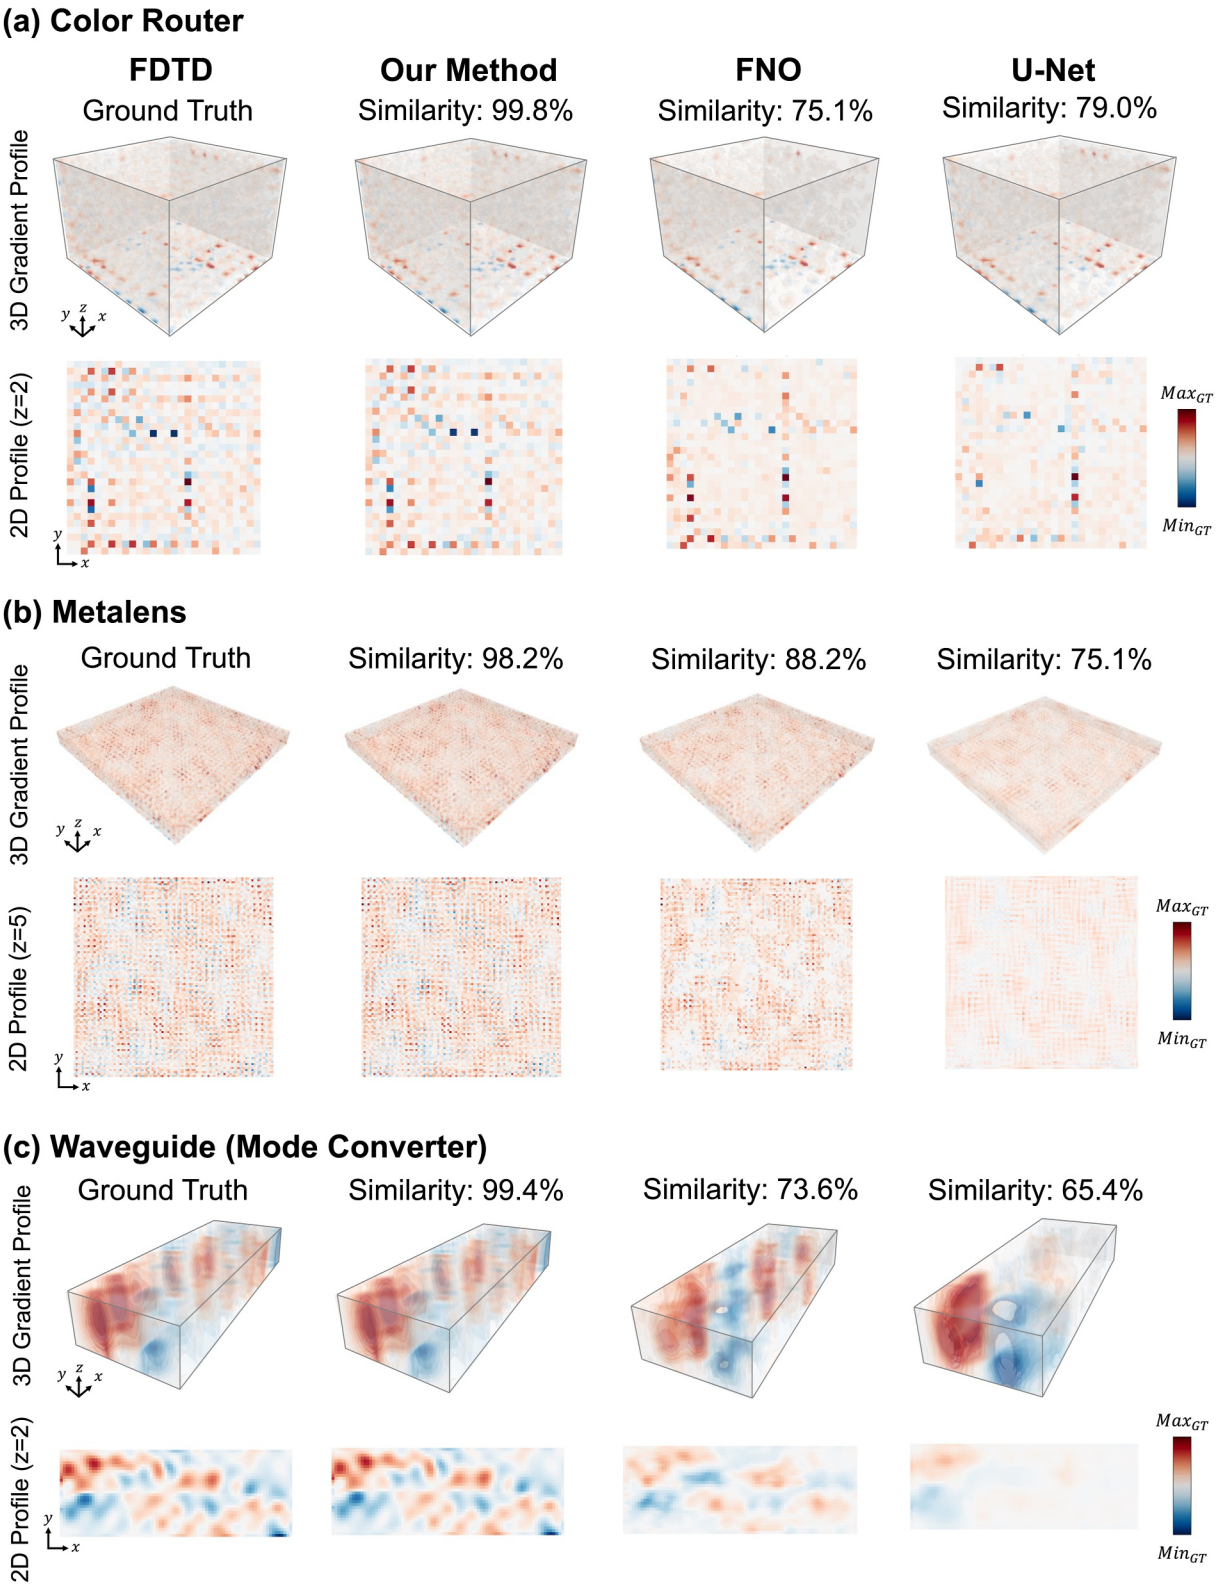

FIGURE 4 | Qualitative comparison of volumetric adjoint-gradient predictions. Ground-truth gradient fields are visualized alongside predictions from the proposed method, a standard FNO, and a 3D U-Net across three tasks: (a) RGB color router, (b) metalens, and (c) waveguide mode converter. Each panel displays a 3D volumetric rendering and a corresponding 2D cross-sectional slice at the indicated z-depth. Although baseline models (FNO and U-Net) tend to produce over-smoothed fields that miss critical high-frequency components, the proposed method accurately reconstructs sharp gradient peaks and fine-grained spatial features, achieving maximum cosine similarity scores exceeding 98% across all benchmarks.

21928614, 2026, 20, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/nap2.70334 by Hanyang University Library, Wiley Online Library on [08/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

evaluation, two complementary metrics are reported: functional fidelity, defined as the ratio of the final FoM achieved by the surrogate-driven optimization to the FDTD benchmark; and structural fidelity, defined as the cosine similarity between the predicted and ground-truth gradient fields (for gradient inference) or between the final voxelized designs (for end-to-end optimization).

4.2 | Adjoint Gradient Inference Across Diverse Benchmarks

The proposed framework is evaluated on three representative high-dimensional inverse-design problems: an RGB color router, an achromatic metalens, and a waveguide mode converter. Both structural fidelity and functional performance are analyzed.

To quantify gradient fidelity, the cosine similarity between the surrogate-predicted gradient fields and the rigorous FDTD adjoint gradients is computed, directly measuring the alignment of update directions used by gradient-based optimizers. A global rescaling of the gradient is absorbed by the optimizer step size, whereas directional errors directly alter the update, so cosine similarity is the decisive measure for optimization performance. Magnitude fidelity is quantified separately in the Supporting Information. As summarized in Table 1, the proposed method consistently achieves the highest cosine similarity across all three tasks, outperforming both the standard one-stage FNO and 3D U-Net baselines. Qualitatively, Figure 4 supports this conclusion: baseline models tend to produce over-smoothed gradient fields that attenuate localized extrema, whereas the proposed method reconstructs sharper peak structures and fine-grained spatial features across all benchmarks.

4.3 | Surrogate-Driven Inverse Design

To assess end-to-end utility for inverse design, two complementary outcomes are reported: (i) the final design performance achieved by surrogate-driven optimization and (ii) the total wall-clock design time required to reach the final design. Performance is reported as relative optimality (FoM and structure similarity), defined as the ratio of the achieved FoM/structure similarity to the FDTD baseline. Results across tasks are summarized in Table 2, and a qualitative example of an optimized color router is shown in Figure 5.

Across all three benchmarks, the proposed method achieves high relative optimality whilst substantially reducing design time. The resulting designs remain close to the FDTD-optimized solutions in terms of final performance, despite eliminating repeated full-wave forward/adjoint solves during optimization. This indicates that the neural adjoint surrogate supports a shared optimization pipeline that transfers across disparate meta-optical objectives. Notably, the color-router and waveguide tasks reach $\gtrsim$95% relative FoM, whilst structural similarity remains lower, consistent with the known nonuniqueness of inverse-designed photonic solutions: different topologies can satisfy the same physical objective. A spectral-coverage analysis of the surrogate-driven trajectory further shows that the iterates remain within the training envelope through the structure-forming phase and progressively exceed it as binarization proceeds, whereas the closely matched trajectories of Figure 5a confirm that this late-phase mismatch does not degrade optimization fidelity (Supporting Information S1: Section 7).

TABLE 2 | Relative optimality and design speed.

	Rel. opt. (%)		Design time		
Task	**FoM**	**Struct.**	**FDTD**	**Ours**	**Speedup**
Color router	95.9	72.1	3.4 h	1.0 s	×12, 240
Metalens	86.0	83.6	40.1 h	1.8 s	×80, 200
Waveguide	98.8	91.2	11.5 h	0.7 s	×59, 143

Note: Surrogate-driven designs are compared against the FDTD benchmark. All methods are run for a fixed number of optimization iterations: 50 (color router), 100 (metalens), and 30 (waveguide).

4.4 | Computational Efficiency

The primary advantage of the proposed method is the elimination of the simulation bottleneck. Compared to conventional solver-in-the-loop adjoint optimization, the surrogate yields orders-of-magnitude speedups on all tasks (Table 2 and Figure 6), reducing design time from hours to seconds. Although a single iteration of 3D FDTD simulation for the color router takes approximately minutes to hours on a workstation (CPU: AMD Ryzen Threadripper PRO 9995WX), neural adjoint inference takes only milliseconds on a single GPU (NVIDIA H200).

4.5 | Effect of Multi-Stage Refinement

Table 3 shows that the substantial gain emerges at three stages, with two-stage refinement (L →H) alone yielding no improvement over the single-stage baseline. This nonmonotonic pattern indicates that direct low-to-high residual learning is insufficient: the wide spectral gap between coarse global features and sharp resonance-driven peaks requires an intermediate mid-frequency stage that bridges the two scales. With a three-stage cascade (L → M → H), the model recovers localized high-magnitude extrema that the single-pass and two-stage variants attenuate, consistent with the multiscale nature of 3D volumetric scattering. This behavior is also consistent with the spectral-bias and frequency-principle literature on neural networks [58, 63]. The rule used to select the per-stage cutoffs is stated in Section 3.1 and verified per task in Supporting Information S1: Section 2.1.

Notably, moving from three to four stages increases relative FoM by only 1.1 percentage points, whilst substantially increasing both the number of parameters and inference latency. The four-stage variant also covers the same spectral range with different interior cutoffs, and its similar relative FoM (88.7% vs. 87.6%) indicates that the final performance is not hyper-sensitive to the exact placement of the interior cutoffs. Given this trade-off, the three-stage configuration is adopted as the default, providing a favorable balance between accuracy and computational overhead.

4.6 | Manufacturability Considerations and Outlook

The designs in this work are freeform volumetric structures, and fabrication-constrained design is beyond the scope of this proof-

21928614, 2026, 20, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/nap2.70334 by Hanyang University Library, Wiley Online Library on [08/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

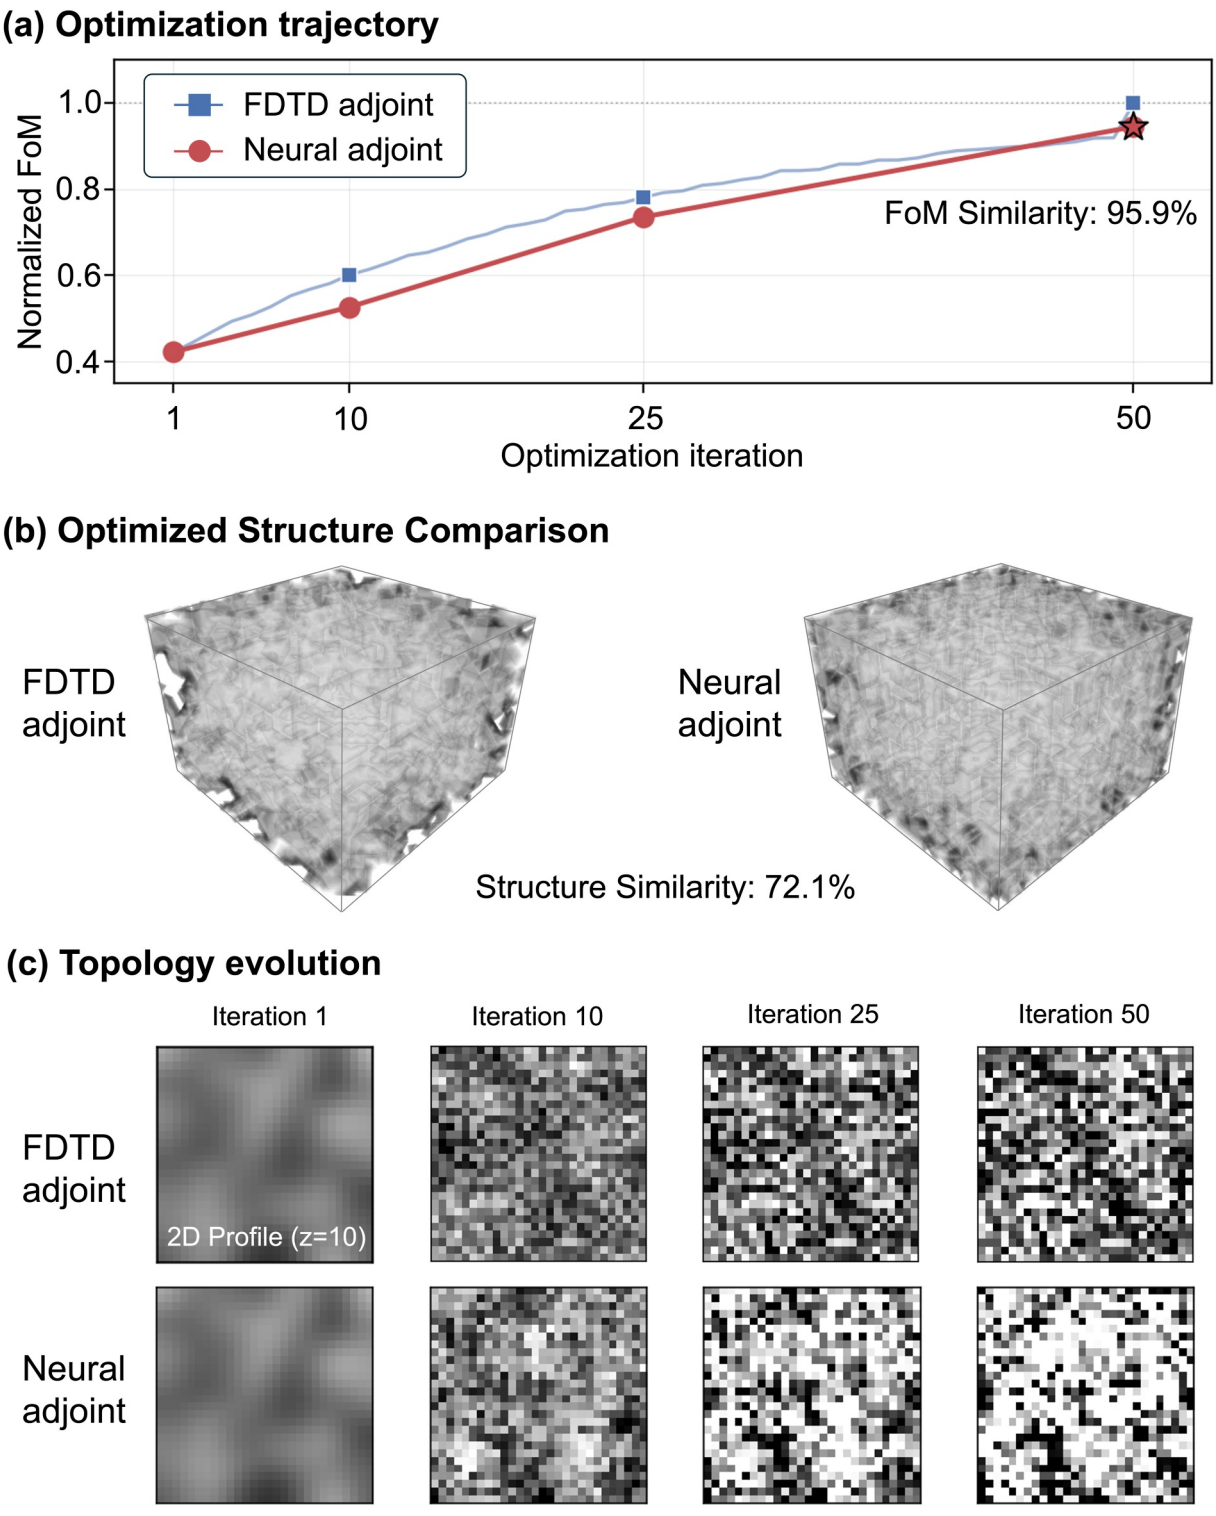

FIGURE 5 | Surrogate-driven inverse design on the RGB color router. (a) Normalized figure-of-merit (FoM) trajectory over 50 optimization iterations for FDTD-based adjoint optimization (blue squares) and the proposed Neural Adjoint method (red circles), both initialized from an identical random permittivity distribution. The two pipelines follow closely matched convergence curves, and the neural adjoint final design attains 95.9% of the FDTD-optimized FoM. (b) Volumetric renderings of the final optimized permittivity distributions. Although the resulting geometries are not identical (structure similarity 72.1%), they achieve comparable functional performance, consistent with the known nonuniqueness of photonic inverse-design solutions. (c) Topology evolution shown as $z = 10$ cross-sectional slices at iterations 1, 10, 25, and 50 for both methods. The two trajectories exhibit similar coarse-to-fine refinement patterns, indicating that the surrogate-driven update directions traverse a comparable optimization landscape to the rigorous FDTD reference, whilst reducing per-iteration wall-clock cost by roughly four orders of magnitude.

21928614, 2026, 20, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/nap2.70334 by Hanyang University Library, Wiley Online Library on [08/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

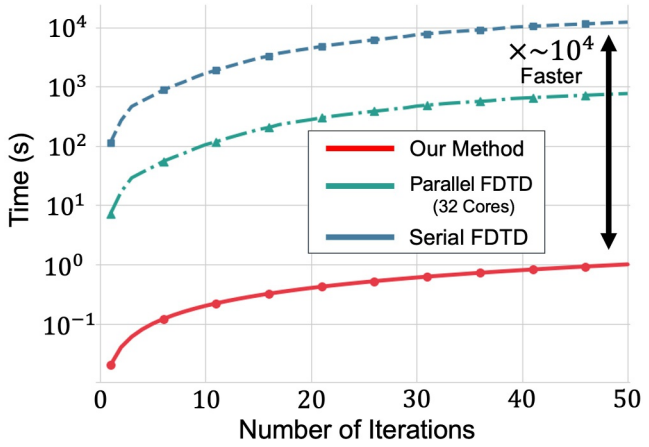

FIGURE 6 | Optimization time comparison. Wall-clock time versus the number of optimization iterations for the neural adjoint method and FDTD-based adjoint optimization (serial and 32-core parallel implementations). Across iterations, the proposed method remains orders of magnitude faster, achieving up to $\sim 10^4\times$ speedup over serial FDTD.

TABLE 3 | Ablation of the number of refinement stages in SW-FNO (1–4) on the metalens task.

#Stages	Stage emphasis	#Params (*M*) ↓	Latency (ms) ↓	Rel. opt. FoM (%) ↑
1	L	7.09	4.8	73.5
2	L → H	14.71	7.8	73.1
3 (Ours)	L → M → H	21.80	10.9	87.6
4	L → LM → HM → H	42.53	14.0	88.7

Note: The Fourier-mode emphasis of each stage is denoted by L (low), M (mid), and H (high); intermediate emphasis levels are abbreviated as LM (low–mid) and HM (high–mid).

of-concept study. Freeform inverse-designed structures have been used as unconstrained performance references because their large design freedom reduces the geometric bias introduced by predefined or fabrication-constrained parameterizations [64, 65]. Such optimized structures are not rigorous upper bounds by themselves. Their performance can instead be interpreted through comparison with independently derived physical bounds [66, 67], and bound-guided inverse design can explicitly target performance near such limits [65]. Accordingly, the present freeform designs should be interpreted as empirical performance benchmarks for a specified design volume, material system, and objective, against which future fabrication-constrained implementations can be evaluated.

Fabrication routes for this class of meta-optics are developing rapidly. Metasurface color routers have been fabricated and characterized [7, 10, 28]. Industrial-scale manufacturing has been demonstrated using roll-to-roll [1] and wafer-scale [2] processes. Two-photon polymerization has been used to produce 3D-patterned inverse-designed meta-optics [68] and 3D-printed multilayer metalenses [31], and its capabilities for imaging optics have been reviewed in Ref. [69]. Layer-by-layer stacking also underlies the volumetric platforms demonstrated in Refs. [8, 39].

Relative to a reduced two-dimensional model, which treats one direction as invariant, a three-dimensional model can represent fabrication deviations throughout the volume, including interface displacements, thickness variations, layer misalignment, and shrinkage. Shrinkage and feature-size mismatch have been identified as likely contributors to the reduced experimental contrast of 3D-printed meta-optics [68]. The effects of these deviations depend on the material distribution, operating wavelengths, and device function, so dimensionality alone does not determine whether a design is more or less tolerant.

Neural adjoint acceleration is also compatible in principle with fabrication-constrained design. Training pairs could be generated from process-relevant geometries, and constraints, such as prescribed layer thicknesses, material choices, and minimum feature sizes [55, 56], would be imposed through the design parameterization or update rule. The predicted gradients would then be propagated to the design variables by the chain rule as is already done for the threshold projection in the present pipeline (Supporting Information S1: Section 6).

Perturbation-insensitive solutions can be targeted by jointly optimizing nominal, eroded, and dilated realizations, together with process-specific perturbations, under an average or worst-case objective [53–55]. Each scenario requires its own gradient evaluation, so such formulations multiply the per-iteration cost of FDTD-based adjoint optimization, which already requires 3.4–40.1 h per design in the present benchmarks (Table 2). A surrogate trained and validated on the perturbed geometries could perform these evaluations in milliseconds, and the gradient fidelity reported in Section 4.2 is a prerequisite for this acceleration. Finite-error performance would still need to be confirmed by forward electromagnetic simulations, and application to a specified process would require validating both the surrogate and the device under the relevant perturbations.

5 | Conclusion

The neural adjoint method has been introduced as a neural-operator-based approach that targets a bottleneck in computational meta-optics: the repeated full-wave forward/adjoint simulations required for 3D adjoint optimization. Instead of predicting only scalar objectives, the method learns to directly infer dense 3D adjoint gradient fields on voxelized device domains, enabling rapid iterative refinement whilst preserving the physically meaningful update directions that drive topology optimization.

A stage-wise Fourier neural operator has been proposed that performs coarse-to-fine residual refinement with progressively higher-frequency emphasis, improving the recovery of localized high-magnitude sensitivity peaks that strongly influence topology updates. To support rigorous analysis and reproducibility, a domain-specific dataset has been constructed from conventional adjoint formulations using full-wave FDTD simulations, pairing voxelized designs with solver-generated volumetric adjoint gradients. Across three physically distinct inverse-design benchmarks (RGB color router, achromatic metalens, and waveguide mode converter), surrogate-driven optimization reaches solutions that remain close to FDTD adjoint optimization in final FoM whilst reducing design time from hours to seconds.

21928614, 2026, 20, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/nap2.70334 by Hanyang University Library, Wiley Online Library on [08/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

These results suggest a practical route toward interactive volumetric meta-optics design, bridging recent advances in neural operators with the stringent fidelity requirements of photonic inverse design.

Author Contributions

Chanik Kang: conceptualization, methodology, software, data curation, investigation, validation, formal analysis, visualization, writing – original draft, writing – review and editing, project administration. **Hyewon Suk:** investigation, data curation, visualization, writing – review and editing, software. **Ikbeom Jang:** writing – review and editing, supervision, project administration. **Haejun Chung:** resources, project administration, writing – review and editing, supervision, funding acquisition, visualization, conceptualization.

Acknowledgments

This work was supported by National Research Foundation of Korea (NRF) grants funded by the Korean government (Ministry of Science and ICT, MSIT) (RS-2024-00338048, RS-2024-00414119, RS-2025-25463760, RS-2026-25487796, and RS-2026-25620621). This work was also supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grants (RS-2025-02283217, RS-2026-25507992, RS-2020-II201373 [Hanyang University], IITP-(2026)-RS-2023-00253914). Additional support was provided by the Culture, Sports and Tourism R&D Program through the Korea Creative Content Agency (KOCCA), funded by the Ministry of Culture, Sports and Tourism in 2024 (RS-2024-00332210). 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).

Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

References

1. T. Hoang, Y. Park, J. Kim, et al., "300-Unit-Per-Second Roll-to-Roll Manufacturing of Visible Metalenses," *Nature* 652, no. 8112 (2026): 1188–1194, https://doi.org/10.1038/s41586-026-10369-y.

2. C. Park, E. Lee, J. Kim, et al., "12″Wafer-Scale Mass-Manufactured Metal–Insulator–Metal Reflective Metaholograms by Nanotransfer Printing," *ACS Applied Materials & Interfaces* 17, no. 2 (2025): 3749–3756, https://doi.org/10.1021/acsami.4c16216.

3. 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/acsphotonics.3c00457.

4. D. E. Jung, J. Amann, V. J. Einck, et al., "All-Inorganic TiO2 Nanoparticle-Based Metalenses Manufactured by Direct Nanoimprint Lithography for High Energy Applications: Femtosecond Laser-Induced Damage Threshold Testing," *Advanced Optical Materials* 14, no. 2 (2026): e01769, https://doi.org/10.1002/adom.202501769.

5. J. Jeon, C. Park, D. Heo, H. Chung, and M. S. Jang, "Inverse Design of Nanophotonic Color Router Robust to Oblique Incidence," *Advanced Optical Materials* 14, no. 4 (2026): e01697, https://doi.org/10.1002/adom.202501697.

6. 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.

7. 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.

8. P. Camayd-Muñoz, C. Ballew, G. Roberts, and A. Faraon, "Multifunctional Volumetric Meta-Optics for Color and Polarization Image Sensors," *Optica* 7, no. 4 (2020): 280–283, https://doi.org/10.1364/optica.384228.

9. Y. Shi, S. Wan, Z. Wang, R. Rao, and Z. Li, "On-Chip Nonlocal Metasurface for Color Router: Conquering Efficiency-Loss From Spatial-Multiplexing," *Light: Science & Applications* 15, no. 1 (2026): 66, https://doi.org/10.1038/s41377-025-02146-9.

10. M. Miyata, K. Shikama, H. Takehara, J. Ohta, and T. Hashimoto, "High-Sensitivity RGB–NIR Image Sensor With Dispersion-Engineered Meta-Optics," *ACS Photonics* 13, no. 3 (2026): 849–860, https://doi.org/10.1021/acsphotonics.5c02801.

11. T. Gu, H. J. Kim, C. Rivero-Baleine, and J. Hu, "Reconfigurable Metasurfaces Towards Commercial Success," *Nature Photonics* 17, no. 1 (2023): 48–58, https://doi.org/10.1038/s41566-022-01099-4.

12. 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.

13. O. D. Miller, "Photonic Design: From Fundamental Solar Cell Physics to Computational Inverse Design," Ph.D. Thesis (University of California, 2012).

14. C. Kang, C. Park, M. Lee, J. Kang, M. S. Jang, and H. Chung, "Large-Scale Photonic Inverse Design: Computational Challenges and Breakthroughs," *Nanophotonics* 13, no. 20 (2024): 3765–3792, https://doi.org/10.1515/nanoph-2024-0127.

15. Y. Wang, Q. Chen, W. Yang, et al., "High-Efficiency Broadband Achromatic Metalens for Near-IR Biological Imaging Window," *Nature Communications* 12, no. 1 (2021): 5560, https://doi.org/10.1038/s41467-021-25797-9.

16. J. Jiang, D. Sell, S. Hoyer, J. Hickey, J. Yang, and J. A. Fan, "Free-Form Diffractive Metagrating Design Based on Generative Adversarial Networks," *ACS Nano* 13, no. 8 (2019): 8872–8878, https://doi.org/10.1021/acsnano.9b02371.

17. S. So and J. Rho, "Designing Nanophotonic Structures Using Conditional Deep Convolutional Generative Adversarial Networks," *Nanophotonics* 8, no. 7 (2019): 1255–1261, https://doi.org/10.1515/nanoph-2019-0117.

18. S. An, B. Zheng, H. Tang, et al., "Multifunctional Metasurface Design With a Generative Adversarial Network," *Advanced Optical Materials* 9, no. 5 (2021): 2001433, https://doi.org/10.1002/adom.202001433.

19. Z. Zhang, C. Yang, Y. Qin, H. Feng, J. Feng, and H. Li, "Diffusion Probabilistic Model Based Accurate and High-Degree-of-Freedom Metasurface Inverse Design," *Nanophotonics* 12, no. 20 (2023): 3871–3881, https://doi.org/10.1515/nanoph-2023-0292.

20. R. Lin, C. Zhang, W. Mao, J. Ge, H. Dong, and L. Zhang, "Diffusion Model-Based Inverse Design of Photonic Crystals for Customized Refraction," *Nanophotonics* 14, no. 27 (2025): 5537–5544, https://doi.org/10.1515/nanoph-2025-0499.

21. H. Wang, Z. Du, F. Feng, Z. Kang, S. Tang, and X. Guo, "DiffMat: Data-Driven Inverse Design of Energy-Absorbing Metamaterials Using Diffusion Model," *Computer Methods in Applied Mechanics and Engineering* 432 (2024): 117440, https://doi.org/10.1016/j.cma.2024.117440.

22. C. Mao, S. S. Leuthold, and J. A. Fan, "Metamaterial Inverse Design Using Diffusion Models," in *AI and Optical Data Sciences VI* (SPIE, 2025): PC1337516.

21928614, 2026, 20, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/nap2.70334 by Hanyang University Library, Wiley Online Library on [08/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

23. C. Kang, D. Seo, S. V. Boriskina, and H. Chung, "Adjoint Method in Machine Learning: A Pathway to Efficient Inverse Design of Photonic Devices," *Materials & Design* 239 (2024): 112737, https://doi.org/10.1016/j.matdes.2024.112737.

24. M. Choi, J. Park, J. Shin, et al., "Realization of High-Performance Optical Metasurfaces Over a Large Area: A Review From a Design Perspective," *npj Nanophotonics* 1, no. 1 (2024): 31, https://doi.org/10.1038/s44310-024-00029-2.

25. Z. Li, N. Kovachki, K. Azizzadenesheli, et al., "Fourier Neural Operator for Parametric Partial Differential Equations," *Preprint arXiv: 2010.08895* (2020).

26. V. Duruisseaux, J. Kossaifi, and A. Anandkumar, "Fourier Neural Operators Explained: A Practical Perspective," *Preprint arXiv: 2512.01421* (2025).

27. C. Liu, J. Zhang, Y. Zhao, and B. Ai, "Inverse Design of Plasmonic Nanohole Arrays by Combing Spectra and Structural Color in Deep Learning," *Advanced Intelligent Systems* 5, no. 10 (2023): 2300121, https://doi.org/10.1002/aisy.202300121.

28. C. Kim, J. Hong, J. Jang, et al., "Freeform Metasurface Color Router for Deep Submicron Pixel Image Sensors," *Science Advances* 10, no. 22 (2024): eadn9000, https://doi.org/10.1126/sciadv.adn9000.

29. R. Pestourie, C. Pérez-Arancibia, Z. Lin, W. Shin, F. Capasso, and S. G. Johnson, "Inverse Design of Large-Area Metasurfaces," *Optics Express* 26, no. 26 (2018): 33732–33747, https://doi.org/10.1364/oe.26.033732.

30. Q. Jiang, J. Liu, J. Li, et al., "Multiwavelength Achromatic Metalens in Visible by Inverse Design," *Advanced Optical Materials* 11, no. 15 (2023): 2300077, https://doi.org/10.1002/adom.202300077.

31. C.-F. Pan, H. Wang, H. Wang, et al., "3D-Printed Multilayer Structures for High–Numerical Aperture Achromatic Metalenses," *Science Advances* 9, no. 51 (2023): eadj9262, https://doi.org/10.1126/sciadv.adj9262.

32. C. Shang, J. Yang, A. M. Hammond, et al., "Inverse-Designed Lithium Niobate Nanophotonics," *ACS Photonics* 10, no. 4 (2023): 1019–1026, https://doi.org/10.1021/acsphotonics.3c00040.

33. W. Jiang, L. Zhang, Z. Han, et al., "Deep Learning Enabled Inverse Design of Ultracompact Silicon Metadevices for Mode Conversion," *Optics Express* 34, no. 2 (2026): 1294–1306, https://doi.org/10.1364/oe.577558.

34. R. E. Christiansen, "Inverse Design of Optical Mode Converters by Topology Optimization: Tutorial," *Journal of Optics* 25, no. 8 (2023): 083501, https://doi.org/10.1088/2040-8986/acdbdd.

35. 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.

36. R. Marzban, A. Adibi, and R. Pestourie, "Inverse Design in Nanophotonics via Representation Learning," *Advanced Optical Materials* 14, no. 1 (2026): e02062, https://doi.org/10.1002/adom.202502062.

37. H. Zhou, H. Yang, H. Ren, J. Matres, and J. Gu, "End-To-End Physical Design Automation Flow for Yield-Optimized Inverse-Designed Large-Scale Electronic-Photonic Integrated Circuits," *Preprint arXiv:2604.15493* (2026).

38. T. Dai, Y. Shao, C. Mao, et al., "Shaping Freeform Nanophotonic Devices With Geometric Neural Parameterization," *npj Computational Materials* 11, no. 1 (2025): 259, https://doi.org/10.1038/s41524-025-01752-w.

39. P. Pearson, G. Roberts, and A. Faraon, "Inverse-Designed Metasurfaces for Multifunctional Spatial Frequency Filtering," *Optica* 12, no. 7 (2025): 1090–1099, https://doi.org/10.1364/optica.560985.

40. B. M. d. A. Jokisch, A. Cerjan, R. E. Christiansen, J. Mørk, O. Sigmund, and S. G. Johnson, "Efficient First-Principles Inverse Design of Nanolasers," *Laser & Photonics Reviews* 20, no. 14 (2026): e01614, https://doi.org/10.1002/lpor.202501614.

41. T. Bi, S. Zhang, E. Bostan, et al., "Inverse-Designed Silicon Nitride Nanophotonics,"2025, https://arxiv.org/abs/2505.13383.

42. K. Azizzadenesheli, N. Kovachki, Z. Li, M. Liu-Schiaffini, J. Kossaifi, and A. Anandkumar, "Neural Operators for Accelerating Scientific Simulations and Design," *Nature Reviews Physics* 6, no. 5 (2024): 320–328, https://doi.org/10.1038/s42254-024-00712-5.

43. 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.

44. J. Seo, C. Kang, D. Seo, and H. Chung, "Wave Interpolation Neural Operator: Interpolated Prediction of Electric Fields Across Untrained Wavelengths," in *NeurIPS 2024 Workshop on Data-Driven and Differentiable Simulations, Surrogates, and Solvers* (2024).

45. J. Seo, C. Kang, D. Seo, and H. Chung, "Broadband Electromagnetic Field Prediction at Unseen Wavelengths via Physics-Guided Neural Operators," *Nanophotonics* 15, no. 15 (2026): e70242, https://doi.org/10.1002/nap2.70242.

46. R. Pestourie, Y. Mroueh, T. V. Nguyen, P. Das, and S. G. Johnson, "Active Learning of Deep Surrogates for PDEs: Application to Metasurface Design," *npj Computational Materials* 6, no. 1 (2020): 164, https://doi.org/10.1038/s41524-020-00431-2.

47. S. Nam, C. Y. Park, and M. S. Jang, "Data-Efficient Electromagnetic Surrogate Solver Through Dissipative Relaxation Transfer Learning," *Preprint arXiv:2601.18235* (2026).

48. O. Furat, V. C. Gogineni, H. Bindslev, and E. S. Nadimi, "Physics-Informed Neural Operators for Predicting 3D Electromagnetic Fields Transformed by Metasurfaces," *Preprint arXiv:2512.15694* (2025).

49. J. Jiang and J. A. Fan, "Global Optimization of Dielectric Metasurfaces Using a Physics-Driven Neural Network," *Nano Letters* 19, no. 8 (2019): 5366–5372, https://doi.org/10.1021/acs.nanolett.9b01857.

50. 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.

51. 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).

52. R. Trivedi, L. Su, J. Lu, M. F. Schubert, and J. Vuckovic, "Data-Driven Acceleration of Photonic Simulations," *Scientific Reports* 9, no. 1 (2019): 19728, https://doi.org/10.1038/s41598-019-56212-5.

53. F. Wang, B. S. Lazarov, and O. Sigmund, "On Projection Methods, Convergence and Robust Formulations in Topology Optimization," *Structural and Multidisciplinary Optimization* 43, no. 6 (2011): 767–784, https://doi.org/10.1007/s00158-010-0602-y.

54. R. E. Christiansen and O. Sigmund, "Inverse Design in Photonics by Topology Optimization: Tutorial," *Journal of the Optical Society of America B* 38, no. 2 (2021): 496–509, https://doi.org/10.1364/josab.406048.

55. A. M. Hammond, A. Oskooi, S. G. Johnson, and S. E. Ralph, "Photonic Topology Optimization With Semiconductor-Foundry Design-Rule Constraints," *Optics Express* 29, no. 15 (2021): 23916–23938, https://doi.org/10.1364/oe.431188.

56. D. Vercruysse, N. V. Sapra, L. Su, R. Trivedi, and J. Vučković, "Analytical Level Set Fabrication Constraints for Inverse Design,"

21928614, 2026, 20, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/nap2.70334 by Hanyang University Library, Wiley Online Library on [08/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Scientific Reports 9, no. 1 (2019): 8999, https://doi.org/10.1038/s41598-019-45026-0.

57. J. Wang, Y. Shi, T. Hughes, Z. Zhao, and S. Fan, "Adjoint-Based Optimization of Active Nanophotonic Devices," *Optics Express* 26, no. 3 (2018): 3236–3248, https://doi.org/10.1364/oe.26.003236.

58. N. Rahaman, A. Baratin, D. Arpit, et al., "On the Spectral Bias of Neural Networks," in *International Conference on Machine Learning* (PMLR, 2019), 5301–5310.

59. S. Qin, F. Lyu, W. Peng, et al., "Toward a Better Understanding of Fourier Neural Operators From a Spectral Perspective," *Preprint arXiv: 2404.07200* (2024).

60. X. Liu, B. Xu, S. Cao, and L. Zhang, "Mitigating Spectral Bias for the Multiscale Operator Learning," *Journal of Computational Physics* 506 (2024): 112944, https://doi.org/10.1016/j.jcp.2024.112944.

61. K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2016), 770–778.

62. A. F. Oskooi, D. Roundy, M. Ibanescu, P. Bermel, J. Joannopoulos, and S. G. Johnson, "Meep: A Flexible Free-Software Package for Electromagnetic Simulations by the FDTD Method," *Computer Physics Communications* 181, no. 3 (2010): 687–702, https://doi.org/10.1016/j.cpc.2009.11.008.

63. Z.-Q. J. Xu, Y. Zhang, T. Luo, Y. Xiao, and Z. Ma, "Frequency Principle: Fourier Analysis Sheds Light on Deep Neural Networks," *Preprint arXiv:1901.06523* (2019).

64. H. Chung and O. D. Miller, "High-NA Achromatic Metalenses by Inverse Design," *Optics Express* 28, no. 5 (2020): 6945–6965, https://doi.org/10.1364/oe.385440.

65. S. Li, H.-C. Lin, and C. W. Hsu, "High-Efficiency High-Numerical-Aperture Metalens Designed by Maximizing the Efficiency Limit," *Optica* 11, no. 4 (2024): 454–459, https://doi.org/10.1364/optica.514907.

66. H. Zhang, C. W. Hsu, and O. D. Miller, "Scattering Concentration Bounds: Brightness Theorems for Waves," *Optica* 6, no. 10 (2019): 1321–1327, https://doi.org/10.1364/optica.6.001321.

67. M. I. Abdelrahman and F. Monticone, "How Thin and Efficient Can a Metasurface Reflector Be? Universal Bounds on Reflection for Any Direction and Polarization," *Advanced Optical Materials* 11, no. 4 (2023): 2201782, https://doi.org/10.1002/adom.202201782.

68. 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.

69. H. Wang, C.-F. Pan, C. Li, et al., "Two-Photon Polymerization Lithography for Imaging Optics," *International Journal of Extreme Manufacturing* 6, no. 4 (2024): 042002, https://doi.org/10.1088/2631-7990/ad35fe.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.
Supporting Information S1: nap270334-sup-0001-suppl-data.pdf.

21928614, 2026, 20, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/nap2.70334 by Hanyang University Library, Wiley Online Library on [08/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License