



PAPER • OPEN ACCESS

A multi-objective permittivity optimization for object classification at the speed of light

To cite this article: Ergun Simsek and Sumya H Oishe 2025 *Mach. Learn.: Sci. Technol.* **6** 045014

View the [article online](#) for updates and enhancements.

You may also like

- [Design and analysis of underwater capture device for high sea conditions](#)
Xiong Deng, Xia Yang, Xiaomei Li et al.
- [Outline of optical design and viewing geometry for divertor Thomson scattering on MAST upgrade](#)
J Hawke, R Scannell, J Harrison et al.
- [Synergistic design method for MRE isolator combining extended-variable tri-phase topology and multi-objective parameter optimization](#)
Mi Zhu, Honglin Zhang, Wang Li et al.



PAPER

OPEN ACCESS

RECEIVED
15 April 2025REVISED
25 September 2025ACCEPTED FOR PUBLICATION
15 October 2025PUBLISHED
24 October 2025

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



A multi-objective permittivity optimization for object classification at the speed of light

Ergun Simsek* and Sumya H Oishe

Department of Computer Science and Electrical Engineering, University of Maryland Baltimore County, Baltimore, MD 21250, United States of America

* Author to whom any correspondence should be addressed.

E-mail: simsek@umbc.edu**Keywords:** Inverse design, machine learning, adjoint method, classification

Abstract

This paper presents a numerical demonstration of the use of the adjoint method for permittivity optimization to design a dielectric medium capable of object classification at the speed of light. In a two-dimensional setup, the system comprises an input waveguide, a design region, and three output ports made of a lossless dielectric material. The design medium is optimized to guide light into specific output ports based on the type and variation of scatterers placed between the input waveguide and the design region. For proof of concept, scatterers derived from the MNIST data-set's digits 0, 1, and 2 are used to represent different object classes with varying shapes and sizes. The optimization process dynamically adjusts the material distribution within the design region to maximize classification performance. The final structure achieved a classification accuracy of 96.3%, with light successfully directed to the correct output port corresponding to each scatterer class. This work demonstrates the potential of permittivity optimization for developing advanced photonic devices capable of ultrafast object recognition, paving the way for future research in three-dimensional designs and more complex classification tasks.

1. Introduction

Inverse photonic design involves leveraging advanced computational techniques to create photonic structures with the desired optical properties. This process begins with the desired outcome and works backwards to identify the optimal structure [1–20]. Unlike traditional design methods, inverse design utilizes algorithms to explore complex design spaces and discover nonintuitive geometries that meet specific performance goals. Among various techniques, permittivity optimization has emerged as a powerful approach in the past two decades [1–20]. Permittivity optimization, often employing the adjoint method for efficient gradient calculations of performance metrics with respect to design parameters, aims to determine the optimal distribution of materials within a design domain to achieve the desired photonic behavior. Originated in fluid dynamics and meteorology for sensitivity analysis [21], the adjoint method significantly accelerates the optimization process.

Historically, Young-Seek *et al*'s and Georgieva *et al*'s works [22, 23] can be considered the first demonstrations of the feasibility of adjoint sensitivity technique for electromagnetic design optimization, laying the groundwork for its application in photonics. Jensen and Sigmund's review paper [2] summarizes the rapid developments in the following decade of permittivity optimization in nanophotonics, utilizing the adjoint method to design photonic crystals, waveguides, and resonators. This pixel-like parameterization of geometry allowed for significant design freedom and efficiency and in the following years, the researchers came up with novel designs for optical cloaking [3], mode conversion [4] and splitting (based on wavelength [4, 6], polarization [4], or power [4]), beam splitters [5], fiber couplers [4], and photonic fibers [9]. Lin *et al* further demonstrated the power of adjoint optimization in nonlinear photonics through cavity-enhanced second-harmonic generation, optimizing the nonlinear overlap to achieve superior efficiency [7]. The integration of machine learning with the adjoint method, as




Figure 1. Schematic illustration of object classification using light. (a) The setup consists of four waveguides: one input waveguide, where the light source is positioned, and three output waveguides. The dark gray box represents the design region. The permittivity of each pixel of the design region will be determined using permittivity optimization. When different types of objects are placed between the input waveguide and design medium, the light is expected to exit through different output ports, as shown in (b), (c), and (d).

explored by Tahersima *et al* [13], introduced deep neural networks for the inverse design of integrated photonic power splitters, combining data-driven approaches with physics-based optimization. Recent studies showcased the freeform metasurface design using topology optimization, illustrating the method's adaptability to complex metasurface geometries [9, 14, 15]. Morrison *et al* introduced a physics-agnostic inverse design framework by integrating the adjoint method with the transfer matrix method, broadening its applicability across various physical domains [17]. Substrate and coating optimization have also benefited from the adjoint method [18–20]. For example, we recently applied it to broadband substrate optimization using layered medium Green's functions, achieving significant performance improvements in quantum efficiency of the two-dimensional (2D) material-based phototransistors [20]. In short, the adjoint method has made a significant impact on inverse photonic design by providing a computationally efficient framework for optimizing complex photonic structures.

In this purely numerical work, we employed the adjoint method for permittivity optimization to design a medium capable of achieving object classification at the speed of light. To demonstrate the feasibility of this concept, we consider a 2D scenario. Assume that our setup involves three distinct object types: rock, paper, and scissors, which we aim to classify using light. Our computation domain includes one input and three output ports made from a lossless dielectric material and a design medium, as illustrated in figure 1(a), with light and dark gray rectangles. The light source is positioned within the input waveguide on the left of the design medium. The object to be classified is placed between the input waveguide and the design medium. If we can achieve our goal, the light should exit from different ports for different objects, i.e. if it is a rock, the light will exit from the lower port; if it is a paper, from the middle port; and if it is a pair of scissors, then from the upper port, as illustrated in figures 1(b)–(d). However, this classification must be achieved for objects with slight variations in size and shape. To accomplish this, we need to use scatterers from different classes, and for each class, we should have similar objects with varying sizes and shapes during the adjoint optimization. In other words, unlike all previous implementations aforementioned, training needs to be done dynamically here. To our knowledge, this is the first time multi-objective permittivity optimization has been carried out with dynamic training.

The remainder of this paper is structured as follows: In section 2, we detail our methodology, beginning with the construction of scatterers from the MNIST dataset [24]. We then describe the computational domain and the finite-difference frequency-domain (FDFD) simulation setup. We explain the adjoint-based permittivity optimization process, including the cost function, the parameterization of the design region, and the use of automatic differentiation. In section 3, we present the numerical results, demonstrating how strategic weighting in the cost function improved classification accuracy, and analyze the resulting field distributions. We also discuss the system's robustness and its inherent linearity. Following a discussion on experimental validation, practical challenges, and the potential for 3D extension in section 4, we conclude.

2. Method

2.1. Scatterers

To validate our hypothesis, we require a dataset comprising three distinct groups of objects, with each group containing numerous similar yet non-identical objects. Therefore, we decided to create scatterers using the digits 0, 1, and 2 from the MNIST dataset [24]—a widely used collection of 60 000 handwritten digit images (0–9) for training and testing image classification algorithms. Its simplicity and standardized format (each image is 28×28 pixels with integer values ranging from 0 to 255, as illustrated in the top row of figure 2) have established it as a foundational benchmark for evaluating and comparing machine learning models in computer vision [24, 25]. Following the approach in [25], we first apply the following expression to convert MNIST images into scatterers with dielectric permittivity values ranging from $\varepsilon_r^{\min}$ to $\varepsilon_r^{\max}$

$$\varepsilon_r(u, v) = \varepsilon_r^{\min} + (\varepsilon_r^{\max} - \varepsilon_r^{\min}) \frac{D_{u,v}}{255}, \quad (1)$$

where $D_{u,v}$ represents the value of the pixel at (u, v) and $1 \leq u, v \leq 28$. We then apply a 2D interpolation algorithm (RegularGridInterpolator module from the SciPy library) to upscale the images from 28×28 to 70×70 pixels, as shown in the bottom row of figure 2. The rationale behind this enlargement will be explained in the subsequent section. Notably, the two ‘digit-2’ images/scatterers in the last two columns of figure 2 exhibit significant variations, which will pose challenges later in the analysis.

2.2. Computation domain

Figure 3(a) illustrates the 10λ by 10λ computation domain used in this work, where λ is the excitation wavelength. To ensure high accuracy, the mesh sampling density along the x (horizontal) and y (vertical) directions is set to $\lambda/60$, resulting in a computational domain composed of 600 by 600 pixels. All the waveguides are 50 pixels wide. The output waveguides are centered at pixel positions 200, 300, and 400 along the y -axis. The design region measures 300 pixels by 400 pixels. The target region—where the scatterer will be placed—is 70 pixels by 70 pixels. This is why we enlarge the MNIST images in the first step of this study. Both the design and target regions are centered at pixel 300 along the y -axis. The source and observation ports are positioned at the centers of the input and output ports, indicated by red and light green dashed lines. The background is assumed to be a vacuum with a relative electrical permittivity of 1, while the waveguides have a relative electrical permittivity of 2.96. The mustard-colored regions in figure 3(a) represent perfectly matched layers, each 40 pixels thick. Figure 3(b) lists the critical coordinates of each component within the computational domain. The objective of the permittivity optimization is to determine the relative electrical permittivity of each pixel in the design region so that light exits through one of the three output waveguides based on the scatterer type (0, 1, or 2), as illustrated in figures 1(b)–(d).

To simulate the wave propagation along the computation domain, we utilize a freely available 2D electromagnetic FDFD simulation tool called Ceviche [26]. The `fdd_ez` module in Ceviche specifically solves for 2D transverse magnetic (TM) modes, where the electric field has only a z -component (E_z) and the magnetic fields have x - and y -components (H_x and H_y). Briefly, this FDFD algorithm solves the following form of Maxwell’s equation:

$$\nabla \times \nabla \times E_z(x, y) - \omega^2 \mu_0 \varepsilon_0 \varepsilon_r(x, y) E_z(x, y) = 0, \quad (2)$$

where $\omega = 2\pi f$, f is the frequency of the excitation, μ_0 and ε_0 are the magnetic permeability and electrical permittivity of vacuum, and ε_r is the relative electrical permittivity. The total field E_z is expressed as $E_z = E_{\text{inc}} + E_{\text{scat}}$, where E_{inc} and E_{scat} are the incident field due to mode injection and scattered field due to the structures inside the computation domain, respectively. Hence, we first determine the fundamental TM mode in the input waveguide. Then, this excitation is applied as a boundary condition by injecting a known modal field E_{inc} . Once the linear system is solved, we compute the electric field intensity at the output ports, as indicated by the green dashed lines in figure 3(a). Integrating this intensity along the prob yields the mode overlap $O_{s,p}$, where s is 0, 1, or 2, representing the scatterer object. The p is either 1, 2, or 3, representing the output ports, which are centered at pixel positions 200, 300, and 400, respectively, along the y -direction.




Figure 2. Top row: Sample digits from the MNIST dataset, where the color of each pixel represents an integer between 0 and 255. Bottom row: Sample scatterers created from the MNIST dataset using 2D interpolation, where the color of each pixel represents the local relative electrical permittivity value ranging from 1 to 4.




Figure 3. (a) Schematic illustration of the computation domain: black and dark gray rectangles are the waveguides and design region, light gray and mustard-colored regions are vacuum and perfectly matched layers (PML), respectively, dashed green square shows the area where the scatterer will be placed; (b) critical pixel numbers along the x and y directions for all the components of the computation domain.

2.3. Permittivity optimization

Permittivity optimization is an iterative process, and for each iteration, we run three simulations to calculate mode overlaps at each port, assuming the scatterer is made from a digit-0 first, then from a digit-1, and finally from a digit-2. We use the following expression as our cost function C

$$C = \sum_{s=0}^2 \left(\frac{O_{s,s+1}}{O_{s,1} + O_{s,2} + O_{s,3}} \right)^{\alpha_s}, \quad (3)$$

where α_s is a positive number that determines the priority given to digit- s . For example, if $\alpha_0 = \alpha_1 = \alpha_2 = 1$, each calculation contributes equally during the permittivity update. Maximizing this cost function should lead to the following scenarios: if the scatterer is created from a digit-0, the strongest mode should appear in the lower output port; if the scatterer is created from a digit-1, the strongest mode should appear in the middle output port; and if the scatterer is created from a digit-2, the strongest mode should appear in the top output port.

In our optimization framework, we employ a gradient-based topology optimization approach utilizing reverse-mode automatic differentiation combined with the Adam optimizer to design photonic structures capable of routing optical signals to specified output ports based on input digit patterns. The design variable ρ representing a density distribution undergoes a sophisticated parameterization process before defining the physical permittivity distribution. This process begins with a blurring operation implemented via 2D convolution using a circular kernel with a radius of 2 pixels, applied once to ensure feature smoothness and mitigate discretization artifacts. Subsequently, a projection operation employing a hyperbolic tangent function enforces binarization through the transformation $\rho = \tanh(\beta\eta) + \tanh(\beta(\rho - \eta)) / \tanh(\beta\eta) + \tanh(\beta(1 - \eta))$, where β controls the projection strength and




Figure 4. Evolution of the cost function C versus epoch number for optimization iterations 1 (blue), 2 (red), 10 (yellow), and 20 (purple).

$\eta = 0.5$ represents the transition midpoint. The projection strength β follows an adaptive scheduling protocol to facilitate gradual change: during iterations 1–39, $\beta = 30$ provides gentle projection; iterations 40–60 employ $\beta = 50$ for increased permittivity change; iterations 61–89 utilize $\beta = 75$ for stronger enforcement; and iterations 90 and beyond apply $\beta = 100$ for permittivity finalization.

Gradient computation leverages reverse-mode automatic differentiation via the adjoint method, providing computational efficiency independent of the number of design parameters. The Adam optimizer [27] operates with a step size of 5×10^{-3} for maximization over 10 steps per digit pattern, resulting in 200 epochs across 20 digits. Figure 4 illustrates the evolution of the cost function across training epochs for different optimization iterations. At iteration 1, the cost starts near unity and increases gradually with epoch number, reflecting the optimizer’s attempt to align scatterer-port mappings with the design objectives. By iteration 2, the cost rises more sharply, reaching values above 2.5 by epoch 6 before stabilizing, indicating improved-but still fluctuating-mode routing performance. At iteration 10, the cost baseline shifts upward to around 3, showing cumulative improvement from earlier iterations, while maintaining smoother growth. By iteration 20, the cost stabilizes above 4, demonstrating that the optimizer has progressively refined the permittivity distribution to direct optical energy toward the desired output ports with higher fidelity. The monotonic upward trend in the cost across iterations confirms that the gradient-based adjoint optimization, combined with adaptive projection and the Adam optimizer, effectively maximizes the overlap of the desired output modes for each digit scatterer. Small fluctuations within individual iterations reflect the nonconvex nature of the optimization landscape and the stochasticity inherent in automatic differentiation-based gradient updates.

To summarize, we compute the gradient of the cost function C using automatic differentiation, specifically reverse-mode differentiation (also known as backpropagation). At each iteration of the optimization, we update the permittivity distribution based on the gradient of the objective function with respect to the design parameters with the following steps. We first run three full-wave simulations to compute the electromagnetic field distributions corresponding to each possible scatterer configuration. Then we evaluate the mode overlap integrals to determine how well the scattered fields couple into the desired output ports. Using automatic differentiation, we compute the gradient of the objective function with respect to the permittivity values in the design region, and we update the permittivity distribution using the Adam optimization algorithm, which adaptively adjusts the step size based on the computed gradients. By employing reverse-mode differentiation, we efficiently compute the necessary gradients without manually deriving and implementing them. The updated permittivity matrix is then used to re-evaluate the objective function, and the process repeats for N_{it} iterations.

Table 1. Confusion matrix showing the classification results for the final optimized design on the test set of 100 samples per class. Rows represent the true class, columns represent the predicted class based on the output port with the strongest signal.

True Class	Predicted '0'	Predicted '1'	Predicted '2'
Digit '0' (True)	100	0	0
Digit '1' (True)	0	100	0
Digit '2' (True)	9	2	89

3. Results

All the codes and datasets used to generate these results are publicly available, please see the 'Data availability statement' below. The frequency of excitation is 6 μm . The input and output ports are assigned a constant relative permittivity value of 2.96. The design region is initialized with a relative permittivity value of 2 and is allowed to vary between 1 and 4 during optimization. For training, we use 20 different sets of scatterers. For each set, we iterate 10 times ($N_{it} = 10$.) In other words, the total number of epochs is 200. The initial learning rate is 5×10^{-3} . For testing, we used 100 sets of scatterers that were not included in the training.

In our initial attempt, we set $\alpha_0 = \alpha_1 = \alpha_2 = 1$ and obtained our first permittivity-optimized medium. Although we correctly identified all the zeros and ones with 100% accuracy, the success rate for identifying scatterers created from digit 2 s was very low at 61%. The main reason for this low accuracy is the non-uniformity of the scatterers created from digit-2 s. The last two columns of figure 2 show two examples. Upon examining the digit-2 images, we realized that people write the number two in many different ways. For our second attempt, we set $\alpha_0 = \alpha_1 = 1$, and $\alpha_2 = 1.2$, giving higher weight (priority) to the permittivity updates from scatterer-2 calculations. In this case, the accuracy for identifying scatterers created from digit 2 s increased to 73%. In our third and last attempt, we set $\alpha_0 = \alpha_1 = 1$, and $\alpha_2 = 2$, and we achieved an accuracy of 89% for identifying scatterers created from digit 2 s. To be more precise, for all the 100 test cases with a 0-shaped scatterer, the strongest mode occurred in the lower port; for all the 100 test cases with a 1-shaped scatterer, the strongest mode occurred in the middle port, and 89 of the 100 test cases with a 2-shaped scatterer, the strongest mode occurred in the upper port. Nine 2-shaped scatterers produced the strongest signal in the lower port and were mislabeled as scatterer-0, while two 2-shaped scatterers produced the strongest signal in the middle port and were mislabeled as scatterer-1. In other words, for 289 test cases in total, the strongest mode was in the correct port, resulting in an overall classification accuracy of 96.33%. For clarity, the classification performance is summarized in the confusion matrix (table 1).

Figure 5 shows the relative permittivity (ϵ_r) distribution of the final version of the computation domain. We observe that ϵ_r changes smoothly between 1 and 4 over the 30 μm by 40 μm design domain.

Figures 6(a)–(c) present the electric field distributions $|E_z|$ across the entire computational domain for scatterers formed by digits 0, 1, and 2, respectively. In contrast to the earlier results shown in figures 1(b)–(d), where only one waveguide is bright while the other two remain dark, here all three output waveguides receive a non-negligible portion of the transmitted light. Nevertheless, the dominant mode coupling occurs into the lower, middle, and upper ports when the scatterer corresponds to digits 0, 1, and 2, respectively. As illustrated in figure 6(a), the designed medium after the scatterer '0' produces a wavefront that interferes constructively towards the lower port, while destructive interference suppresses the excitation of the middle and upper ports. A similar effect is observed for the scatterer '2' in figure 6(c), where constructive interference directs energy predominantly into the upper port.

To better understand the intermediate case of the scatterer '1', we turn to the phase distributions shown in figure 7. Unlike the vertical phase patterns observed for scatterers '0' and '2', the phase pattern for the scatterer '1' exhibits a strong horizontal variation within the design region. This phase modulation indicates that the scatterer modifies both the amplitude and phase of the propagating wave in such a way that efficient coupling occurs primarily into the fundamental mode of the middle port. Overall, the design region acts as a physical analog of a linear kernel: it transforms the scattered fields into a representation where the output classes (digits) are distinguished by the port with maximum transmitted power. This demonstrates how the scatterer geometry effectively encodes information into spatial field distributions that can be mapped onto separable output channels.

To verify the accuracy of these simulations, we re-simulated all the test cases using Tidy3D, a commercial finite-difference time-domain solver. In all cases, identical numerical results were obtained.




Figure 5. The distribution of relative permittivity in the computation domain at the end of the optimization.




Figure 6. Electric field distribution $|E_z|$ over the entire computation domain for scatterers created from (a) digit-0, (b) digit-1, and (c) digit-2. The contour of the permittivity distribution is overlaid on the field distribution for more intuitive visualization.




Figure 7. Phase of the electric field over the entire computation domain for scatterers created from (a) digit-0, (b) digit-1, and (c) digit-2. The contour of the permittivity distribution is overlaid for more intuitive visualization.

To assess the practical viability of our classifier, we evaluated its robustness against two common perturbations: additive noise and small misalignments. For the first robustness study, we introduced a Gaussian noise with a standard deviation of 5% of the maximum permittivity value into the scatterer

profiles of the test set. When tested with these noisy scatterers, the overall classification accuracy dropped slightly from 96.33% to 94.33%. The accuracy for digit '2' decreased to 83%, while digits '0' and '1' remained perfectly classified. This indicates a reasonable level of resilience to fabrication imperfections or measurement noise. For the second robustness study, we tested the system's sensitivity to misalignment by shifting the scatterer's position within the target area by up to ± 5 pixels (approximately $\lambda/12$) in both x and y directions. This perturbation caused a more significant performance drop, reducing the overall accuracy to 88%. When we randomly change the orientation of the scatterers, the overall classification accuracy dropped to 57 %. Considering we only use three waveports for the classification task, such a low classification accuracy for randomly oriented objects is inevitable. We recently addressed this challenge in a separate study [28] using an attention-based neural network. Since this was more of a post-processing solution, rather than direct classification with light, we do not delve into the details of that study. We conclude that for passive classification systems such as the one examined in this work, the object should be centered between the input waveport and design region, and it should be aligned parallel to the waveport.

We acknowledge that the classification mechanism demonstrated here is inherently linear, as it relies on light propagation governed by Maxwell's equations within a passive, linear dielectric medium. The optimization procedure tailors the spatial permittivity distribution, but the resulting classification corresponds to a linear projection of the input field onto the output ports. This represents a fundamental limitation, since linear classifiers are known to struggle with data that is not linearly separable. Nevertheless, linear systems remain powerful for certain classes of problems and can serve as energy-efficient front-end processors. To overcome this limitation, future extensions could integrate optical nonlinearities into the design region, for example, by incorporating Kerr-type or saturable absorber materials, which would enable intensity-dependent responses and thereby increase the expressive power of the classifier. Alternatively, hybrid optical–electronic architectures could be envisioned, where the optical module performs fast linear transformations while subsequent electronic or optoelectronic layers implement nonlinear activation. Such approaches would combine the energy efficiency and parallelism of optical processing with the representational power of nonlinear computation, thereby addressing the theoretical constraints of purely linear classifiers.

4. Discussion

4.1. Experimental validation plan

To further support the numerical demonstrations, we outline here a potential pathway toward experimental validation. The proposed setup would consist of a coherent mid-infrared source, such as a quantum cascade laser capable of generating monochromatic light near $6\text{ }\mu\text{m}$. The input and output ports can be realized as single-mode ridge waveguides with grating couplers for efficient excitation and readout [29]. The light would then be incident on the target object, which is placed over the designed region. The optical transmission at each output port would be recorded either with an infrared camera for full-field imaging or using discrete photodetectors (e.g. the Hamamatsu P16113-011 M, an InAsSb photovoltaic detector) connected to an optical power meter for quantitative measurements. By sequentially testing scatterer patterns corresponding to different digit classes, classification accuracy could be directly benchmarked against numerical predictions, following procedures similar to those adopted in recent demonstrations of optical neural networks [30, 31].

4.2. Practical implementation challenges

This was a simple proof-of-concept study, and we are aware of the fabrication-related challenges associated with the permittivity-optimized designs. Small deviations in scatterer geometry can alter interference conditions and reduce classification fidelity, highlighting the need for robust design strategies. Precise alignment of the scatterers and stable coupling of the input laser beam are critical to minimize systematic measurement errors, while the sensitivity of photodetectors will ultimately limit the precision of port intensity readout.

Beyond fabrication, several system-level challenges must be addressed for real-world deployment. In high-speed optical communication systems, for example, precise signal synchronization between the laser source, the scatterer, and the detection module is essential, as even minor temporal drifts can compromise classification performance. Furthermore, noise resilience is a critical concern, as environmental fluctuations (e.g. thermal drift, vibrations, or laser intensity noise) can perturb the scattered fields and obscure the detected signal. These challenges can be effectively mitigated by employing a lock-in amplification technique, where the laser source is amplitude-modulated at a stable reference frequency and the photodetector outputs are measured using synchronous detection. By doing so, our system can reject

uncorrelated noise and ensure precise, drift-resistant measurement of the port intensities. For ultimate robustness, such signal recovery strategies can be combined with scatterers designed for inherent stability and digital error-correction protocols. Additionally, stability and long-term reliability are crucial considerations for optical sensing applications. Over time, material degradation, photodetector aging, or surface contamination can alter the effective permittivity distribution, thereby affecting reproducibility. Strategies such as protective coatings, effective heat sinks, self-calibration protocols, and redundant sensor arrays may help mitigate these effects.

Nevertheless, rapid progress in nanofabrication technologies offers promising solutions. For instance, nanocomposite materials [32, 33] enable the realization of graded permittivity profiles by embedding high-permittivity nanoparticles (e.g. titanium dioxide [32] or barium titanate [33]) into a low-permittivity polymer matrix (e.g. silicon dioxide or PMMA [34]). By spatially varying the nanoparticle concentration, the local permittivity can be precisely tuned to match the optimized design. Advanced fabrication techniques, such as two-photon polymerization [35] or direct ink writing (for book-size objects, using GHz-level electromagnetic antennas and receivers), can then provide the necessary resolution for the accurate replication of permittivity-optimized structures.

4.3. Extension to three-dimensional designs

While the present work is implemented in 2D, the broader applicability of this approach motivates considering extensions to three-dimensional (3D) configurations. A transition to 3D would enlarge the design space and could enable higher-capacity classification schemes, more compact device layouts, and improved tolerance to certain fabrication imperfections. It is important to clarify a common misconception: the optimization formulation does not fundamentally change when moving from 2D to 3D. The adjoint-based gradient computation retains the same mathematical structure and requires only $M + 1$ full-wave solutions of Maxwell's equations per gradient evaluation (where M is the number of exit ports). Consequently, although increasing the number of output ports (to support more classes) increases the number of full-wave solves per iteration, the algorithmic complexity of the optimization loop remains the same. That said, practical resource requirements do grow in 3D: the volumetric mesh contains many more degrees of freedom, the number of permittivity voxels to be updated increases, and each full-wave solve becomes more memory- and time-intensive. In 3D, one must also account for all three vector components of the electric field (E_x, E_y, E_z) in both the forward and adjoint evaluations; these are naturally incorporated into the gradient expressions without altering the optimization formulation, but they do increase per-solve computational effort. Despite these practical cost increases, we expect the number of iteration steps required for convergence to remain modest—typically a few tens of iterations, as in the 2D case—because the adjoint method provides efficient exact gradients regardless of dimensionality. Fabrication of the design medium remains more challenging in 3D (requiring techniques such as multiphoton lithography [36], focused-ion-beam approaches [35], or high-resolution additive manufacturing [37]) at high frequencies, but these are separate engineering hurdles: they do not change the adjoint-based optimization workflow and can be addressed with the same material- and process-level strategies aforementioned before. Also note that in this work, we restrict our analysis to real positive permittivity values due to solver limitations, but more sophisticated 3D solvers such as MEEP and COMSOL can directly handle lossy materials, thereby introducing an additional degree of freedom that could enhance both the accuracy and capacity of 3D classification; in fact, by placing photodetectors on all five output facets (with excitation from the sixth side), one could, in principle, classify a substantially larger number of object classes.

4.4. Efficiency of adjoint method and ultra-fast classification with light

The computational efficiency of the adjoint method is a critical advantage that enables the practical design of large-scale, wavelength-scale devices, such as the one presented here. Unlike other modern optimization techniques, such as particle swarm optimization [38] or genetic algorithms [39], which are often derivative-free and rely on stochastic exploration of the design space, the adjoint method is a gradient-based approach. Its power stems from a physics-based derivation that allows it to compute the sensitivity of the figure of merit (e.g. field intensities at target locations) to design parameters (e.g. the permittivity of each pixel in the structure) in just M additional full-wave simulations per iteration, regardless of the number of parameters (pixels), where M is the number of ports. This stands in stark contrast to brute-force finite-difference methods, which would require $N + 1$ simulations for an N -parameter problem, becoming prohibitively expensive for the thousands of pixels in a typical permittivity optimization. Consequently, while methods like genetic algorithms or surrogate modeling [38] may require hundreds or even thousands of iterations to converge—each involving a computationally expensive electromagnetic simulation—the adjoint method typically converges to a high-quality solution in only a

few tens of iterations. This reduction by an order of magnitude in the number of required simulations makes the optimization of complex, large-area photonic structures computationally feasible.

Last but not least, please remember that our computation domain is 10λ wide. The light propagation from one end to another takes only $\sim 10\lambda\sqrt{\epsilon_r^{\max}}/c_0$ seconds, where c_0 is the speed of light in vacuum. This process takes 0.40 picoseconds in this work. An implementation at 1 GHz would take 20 nanoseconds. In other words, object classification via a permittivity-optimized medium happens at the speed of light.

5. Conclusion

In conclusion, this work demonstrates a successful application of the adjoint method for permittivity optimization in designing a photonic medium capable of high-speed object classification. By leveraging a dynamically evolving material distribution, our design effectively routes light through specific output ports based on the shape and size variations of scatterers derived from the MNIST dataset. The optimization process, guided by the Adam algorithm and tailored cost function, resulted in a photonic structure that achieved a classification accuracy of 96.33% across diverse scatterer samples. Notably, increasing the weighting factors for more challenging scatterer classes substantially improved performance, highlighting the method's adaptability. Future research may explore extending this framework to 3D systems and integrating more complex datasets to broaden the scope of optical computing applications.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://github.com/simsekergun/PermittivityOptimization>.

Acknowledgments

The authors gratefully acknowledge the support of the College of Engineering and Information Technology (COEIT) at the University of Maryland, Baltimore County (UMBC) for providing internal funding that made this research possible.

Conflict of interest

The authors declare that they have no conflict of interest.

Ethical statement

This research did not involve any studies with human participants or animals performed by any of the authors.

Author contributions

Ergun Simsek  0000-0001-9075-7071

Conceptualization (equal), Data curation (equal), Formal analysis (equal), Investigation (equal), Methodology (equal), Project administration (equal), Supervision (lead), Visualization (equal), Writing – original draft (equal), Writing – review & editing (equal)

Sumya H Oishe  0009-0001-7239-6580

Data curation (equal), Investigation (equal)

References

- [1] Chen M *et al* 2024 *J. Opt. Soc. Am. B* **41** A161–76
- [2] Jensen J S and Sigmund O 2011 *Laser Photon. Rev.* **5** 308–21
- [3] Miller O D 2012 Photonic design: from fundamental solar cell physics to computational inverse design *PhD Thesis* University of California Berkeley doctoral dissertation
- [4] Lu J and Vučković J 2013 *Opt. Express* **21** 13351–67
- [5] Lalau-Keraly C M, Bhargava S, Miller O D and Yablonovitch E 2013 *Opt. Express* **21** 21693–701
- [6] Piggott A Y, Lu J, Lagoudakis K G, Petykiewicz J P, Babiniec T M and Vučković J 2015 *Nat. Photon.* **9** 374–7
- [7] Lin Z, Liang X, Lončar M, Johnson S G and Rodriguez A W 2016 *Optica* **3** 233–9
- [8] Sell D, Yang J, Doshay S, Yang R and Fan J A 2017 *Nano Lett.* **17** 3752–7

- [9] Sitawarin C, Jin W, Lin Z and Rodriguez A W 2018 *Photon. Res.* **6** B82–B88
- [10] Hughes T W, Minkov M, Williamson I A D and Fan S 2018 *ACS Photon.* **5** 4781–7
- [11] Wang J, Shi Y, Hughes T, Zhao Z and Fan S 2018 *Opt. Express* **26** 3236–45
- [12] Bradley A M 2024 PDE-constrained optimization and the adjoint method *Technical Report* Standord University detailed tutorial on PDE-constrained optimization
- [13] Tahersima M H, Kojima K, Koike-Akino T, Jha D, Wang W, Lin X, Yiğit A, Katsube Y and Oikawa Y 2019 *Sci. Rep.* **9** 1–10
- [14] Fan J A 2020 *MRS Bull.* **45** 196–201
- [15] Mansouree M, McClung A, Samudrala S and Arbabi A 2021 *ACS Photon.* **8** 455–63
- [16] Ma W, Liu Z, Kudyshev Z A, Boltasseva A, Cai W and Liu Y 2021 *Nat. Photon.* **15** 77–90
- [17] Morrison N, Pan S and Ma E Y 2024 *APL Mach. Learn.* **2** 016115
- [18] Kiziltas G, Psychoudakis D, Volakis J L and Kikuchi N 2003 *IEEE Trans. Antennas Propag.* **51** 2732–43
- [19] Koucheh A B, Kecebas M A and Sendur K 2024 *J. Opt. Soc. Am. B* **41** A98–A107
- [20] Simsek E, Islam R, Oishe S H and Menyuk C R 2024 *J. Opt. Soc. Am. B* **41** 2259–65
- [21] Errico R M 1997 *Bull. Am. Meteorol. Soc.* **78** 2539–58
- [22] Chung Y S, Cheon C, Park I H and Hahn S Y 2000 *IEEE Trans. Microw. Theory Tech.* **48** 2289–96
- [23] Georgieva N K, Glavic S, Bakr M H and Bandler J W 2002 *IEEE Trans. Microw. Theory Tech.* **50** 2751–8
- [24] Deng L 2012 *IEEE Signal Process. Mag.* **29** 141–2
- [25] Simsek E and Manyam H R 2024 *IET Microw. Antennas Propag.* **18** 898–910
- [26] Hughes T W, Williamson I A D, Minkov M and Fan S 2019 *ACS Photon.* **6** 3010–6
- [27] Kingma D P and Ba J 2017 Adam: a method for stochastic optimization (arXiv:1412.6980)
- [28] Simsek E and Manyam H R 2025 Improving object classification accuracy from electromagnetic data using attention mechanisms *26th Int. Conf. on Electromagnetics in Advanced Applications (ICEAA 2025) (Palermo, Italy)*
- [29] Taillaert D, Chong H, Borel P I, Frandsen L H, De La Rue R M and Baets R 2004 *IEEE Photonics Technol. Lett.* **15** 1249–51
- [30] Lin X, Rivenson Y, Yardimci N T, Veli M, Luo Y, Jarrahi M and Ozcan A 2018 *Science* **361** 1004–8
- [31] Zhang Y *et al* 2021 *Nat. Commun.* **12** 457
- [32] Zhou Y, Hu J, Dang B and He J 2017 *IEEE Trans. Dielectr. Electr. Insul.* **24** 1380–9
- [33] Saini P, Arora M, Gupta G, Gupta B K, Singh V N and Choudhary V 2013 *Nanoscale* **5** 4330–6
- [34] Sengwa R J and Dhatarwal P 2021 *Opt. Mater.* **113** 110837
- [35] Baldacchini T 2015 *Three-Dimensional Microfabrication Using two-Photon Polymerization: Fundamentals, Technology and Applications* (William Andrew)
- [36] Gonzalez-Hernandez D, Varapnickas S, Bertoncini A, Liberale C and Malinauskas M 2023 *Adv. Opt. Mater.* **11** 2201701
- [37] Loterie D, Delrot P and Moser C 2020 *Nat. Commun.* **11** 852
- [38] Anjum I M, Simsek E, Mahabadi S E J, Carruthers T F, Menyuk C R, Campbell J C, Tulchinsky D A and Williams K J 2023 *J. Lightwave Technol.* **41** 7285–91
- [39] Simsek E and Menyuk C R 2023 Designing silicon-germanium photodetectors with numerical optimization: the tradeoffs and limits *Proc. SPIE* **12647** 126470L