Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Mie Optical Computing

Mihail Petrov

trisha.petrov@gmail.com

ITMO University https://orcid.org/0000-0001-8155-9778

Vsevolod Kleshchenko

ITMO University

Vladimir Igoshin

ITMO University

Costantino De Angelis

University of Brescia https://orcid.org/0000-0001-8029-179X

Article

Keywords:

Posted Date: September 11th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10975200/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: There is **NO** Competing Interest.

Mie Optical Computing

Vsevolod Kleshchenko[1], Vladimir Igoshin[1], Costantino De Angelis[2], and Mihail Petrov [1]

[1]*Department of Physics and Engineering, ITMO University, 191002, Lomonosova St. 9, St. Petersburg, Russia and*

[2]*Department of Information Engineering, University of Brescia, Via Branze 38, 25123 Brescia, Italy*

Abstract

Optical computing is emerging as a promising paradigm for next-generation information processing. Diffractive optical processors rely on spatially distributed trainable degrees of freedom, leading to extended architectures. Here, we propose a compact neuromorphic optical-computing approach where the entire trainable transformation is implemented by a single Mie scatterer. By formulating computation in vector spherical harmonics basis, trainable modal couplings can be concentrated within a finite object through its T-matrix. Since available T-matrix parameters scale as the fourth power of maximal multipole order, this architecture can overcome trainable-parameter-density limitations of conventional spatially distributed diffractive processors. We demonstrate classification of phase-encoded MNIST images using scattered-field intensity. At a particle size parameter of $ka = 15$, the trained T-matrix reaches approximately 90% test accuracy, comparable to a single-layer artificial neural network. Similar performance can be achieved using near-fields, enabling on-chip integration. We show how reciprocity, passivity, and particle symmetry constrain performance: passivity reduces the accessible operator space while improving robustness, whereas symmetry reduces the number of independent parameters. Finally, we inverse-design a non-absorbing dielectric scatterer that realizes the classification task with 84% accuracy. These results demonstrate that nontrivial neuromorphic transformations can be encoded within the multipolar response of a single compact scatterer.

I. INTRODUCTION

Free-space optical neural networks have been rapidly developing over the last several years as one of the most promising approaches for analog optical computing as an alternative to electronic architectures [1], although whether this translates into a genuine computational advantage over digital electronics – and what role metamaterial and metasurface platforms specifically can play in achieving it – remains an open and actively debated question [2]. Such optical systems offer information processing at the speed of light with intrinsically low latency, massive parallelism, and potentially energy saving [3–5]. While the basic principles of optical computing have been established for several decades [6, 7] the rapid progress in the diffractive optics and metaoptics technologies has made possible wider spread of optical neural networks and formation of the modern architecture of diffractive optical neural

networks (DNN) based on cascaded passive or active diffractive layers [8]. The spatial diffraction and scattering channels, with the latter potentially outnumbering the open diffraction channels [9], play the role of communication channels and define the connectivity between the neurons [3, 10], while the trainable parameters are encoded in the local amplitude and phase response of the diffractive layers. This architecture has been successfully applied to a broad range of computational tasks, including image classification [8, 11, 12], object recognition [13], turbid media imaging [14], polarization [15] and spectral [4] processing. Additional momentum to the field was given by the development of metaoptics and metasurfaces, which enabled subwavelength control of optical states for image processing improving the performance of diffractive optical neural networks [16, 17] and making them more compact.

Despite these advances, most existing optical neural-network architectures remain fundamentally distributed systems in which increasing computational complexity requires additional diffractive layers, extending free-space propagation distances, cascaded Fourier components, or multilayer metasurface stacks [4]. More generally, diffraction and the overlapping nonlocality of an optical transformation impose a minimum transverse size and, for sufficiently nonlocal operations, a minimum system thickness [18, 19]. On the way of downscaling the physical footprint of optical computation unit a few approaches have been suggested such as using a single scatterer for integral transformation [20], multiple linear [21] and nonlinear [22] scattering in disordered microstructures, as well as scattering based nonlinear computing [23] and programmable multimode propagation [24]. However, the question of how to implement complex trainable transformations in a single compact optical element while keeping the required connectivity remains largely unexplored.

In this work, we explore an alternative approach to optical neural computation in which the entire trainable transformation is implemented by a single scatterer compared to a wavelength and, thus, operating in the Mie regime [25]. Multipolar nanophotonics enables control of the optical response through the excitation and interference of electromagnetic multipoles [26]. Instead of distributing computation across multiple diffractive layers, we represent the optical processor by a trainable electromagnetic T-matrix [27] acting in the basis of vector spherical multipoles (VSH). The basic principle of this approach are illustrated in Fig. 1 (a) , where the input and output optical signals are represented as structured far-field illumination and scattered radiation, respectively, while the trainable transformation

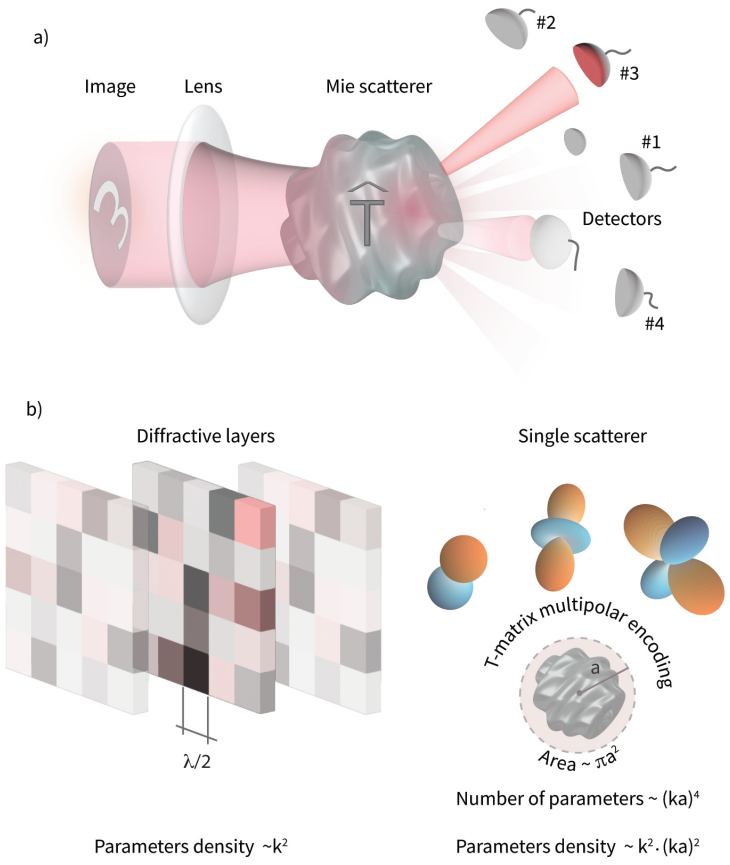


Figure 1. The concept of optical computing with a single Mie scatterer with a trainable multipolar transformation matrix for optical inference. (a) A phase-encoded image is focused onto a Mie scatterer whose T-matrix couples incident and scattered multipolar channels. Spatially separated detectors collect the scattered power and assign the predicted class. (b) Comparison of trainable-parameter densities. A conventional diffractive layer provides locally tunable pixels with a density of order $1/\lambda^2 \sim k^2$. A scatterer of size a supports approximately $(ka)^2$ multipolar channels and $(ka)^4$ pairwise T-matrix couplings, corresponding to a parameter density increase of the order of $(ka)^2$.

is implemented a single compact scatterer within the T-matrix formalism. The proposed approach encodes the computational operation into the modal scattering response of a compact object, where the channel mixing in the multipolar basis can occur within a single scatterer implied by the symmetry and geometry of the structure [28, 29].

The recent work has highlighted that the performance and scalability of optical computing systems are fundamentally constrained by the physics of wave propagation and scattering. In particular, limits associated with spatial complexity and communication-channel overlap

[10, 18], as well as optical connectivity [19], constrain the size of a system implementing a given optical transformation or inference task. In this prospect, the single scatterer has much higher density of trainable parameters compared to the diffractive optical neural networks. Indeed, the DNN is characterized by a single "pixel" of typical size as small as $\sim \lambda/2$, which provides the density of trainable parameters $\sim 1/\lambda^2$ per unit area (see Fig. 1 (b)). In contrast, the number of free parameters of a finite scatterer in the VSH basis depends on the number of free parameters on the T-matrix which scale as $\sim n_{\max}^4$, where $n_{\max}$ is the multipole truncation number. The $n_{\max} \propto ka$ [30, 31], where a is the typical size of the scatterer, and $k = 2\pi/\lambda$ is the wavenumber (see Fig. 1 (b)). Physically, $n_{\max}^2 \sim (ka)^2$ determines the number of accessible multipolar channels, whereas the $n_{\max}^4$ scaling of the T-matrix originates from pairwise couplings between the input and output channel spaces [32, 33]. Thus, the ratio between the number of free parameters in layered DNN and single scatterer scale as $\sim (ka)^2$ (see more details in Methods Sec. IV B) that gives already two order of increase for particle comparable to the wavelength $a \sim \lambda/2$. That opens a promising way towards new DNN architectures based on compact scatterers which can implement nontrivial neuromorphic transformations with a high density of trainable parameters.

In the following, we consider the optical classification task of phase-encoded MNIST images [34] using the detected scattered-field intensity in the far- and near-fields. We demonstrate that a single scatterer can be trained to perform nontrivial classification tasks, achieving test accuracy comparable to a single-layer artificial neural network. We further investigate how physical constraints, such as reciprocity, passivity, and particle symmetry, affect the achievable performance of the T-matrix-based optical classifier. Finally, we present an inverse-designed non-absorbing dielectric scatterer that realizes the classification task with high accuracy, demonstrating the practical feasibility of implementing complex transformations within a single compact optical element.

II. RESULTS AND DISCUSSION

A. Spherical multipole representation

We consider a single scatterer as an optical computation unit which performs a trainable linear transformation of the incident structured electromagnetic fields in the basis of vector

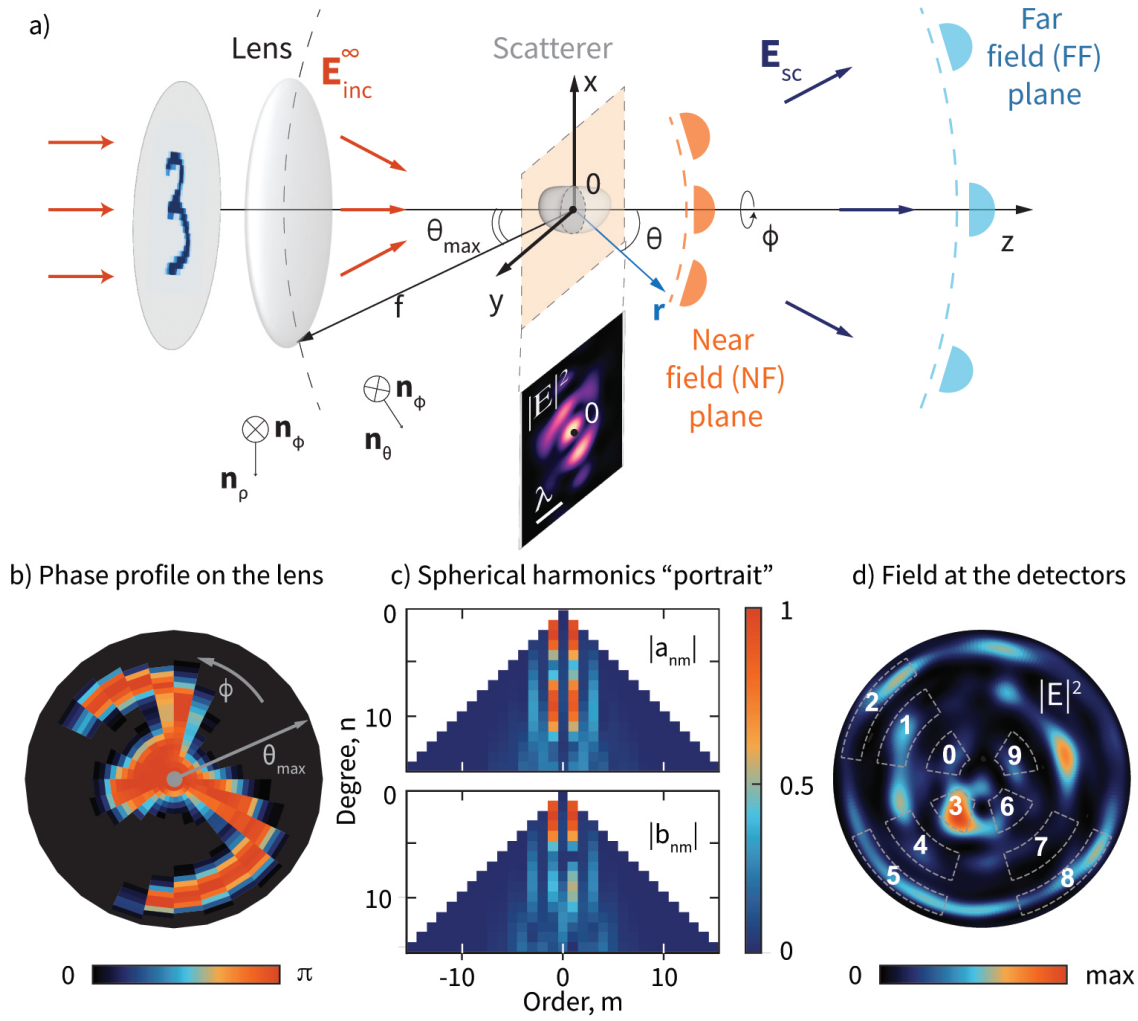


Figure 2. (a) Near-field and far-field formulation of the T-matrix-based optical computation scheme. (b) Example of an MNIST digit encoded in the spherical-harmonic modal basis. The input image is encoded in phase of x-component of electric field projected on a lens. (c) Amplitudes of the multipole decomposition of the input field in (b). (d) Intensity pattern of scattered fields in the far-field plane. Detectors regions are denoted with dashed lines.

spherical harmonics. The general architecture of the system is illustrated in Fig. 2 (a). In this system, input information is encoded in the electromagnetic wave that passes through the infinitely thin lens and focused onto the particle. The scatterer performs processing of the input and the output of the system is collected by the set of detectors. In our work we demonstrate the working principle of the computational unit on the MNIST handwritten digits classification task, where the image of the digit serves as input. Detectors collect the scattered-field intensity in the near-field zone ($kr \sim n_{\max}$) or far-field ($kr \gg 1$) of the

scatterer. The number of detectors is equal to the number of classes K to make one-to-one correspondence between detector number and digit class.

Here, we consider monochromatic fields with the wavelength λ ($k = 2\pi/\lambda$). The lens is characterized by its numerical aperture (NA) which defines the maximal angle of the incident light $\theta_{\max} = \arcsin(\mathrm{NA})$. The NA of the detector grid is equal to the NA of the lens for definiteness. The maximal illumination angle determines how many transverse spatial frequencies of the input field can reach the scatterer. Larger $\theta_{\max}$, or equivalently larger NA, excites more rapidly varying angular structure and therefore requires a larger set of multipole channels for an accurate representation. At the same time, the number of physically relevant multipoles is limited by the particle size parameter ka: a scatterer with larger ka can support higher-order electric and magnetic multipoles. Thus, the truncation order $n_{\max}$ must be chosen large enough to resolve both the angular content imposed by the lens and the highest multipolar orders supported by the particle, with the natural scaling $n_{\max} \sim ka$ up to a numerical safety margin.

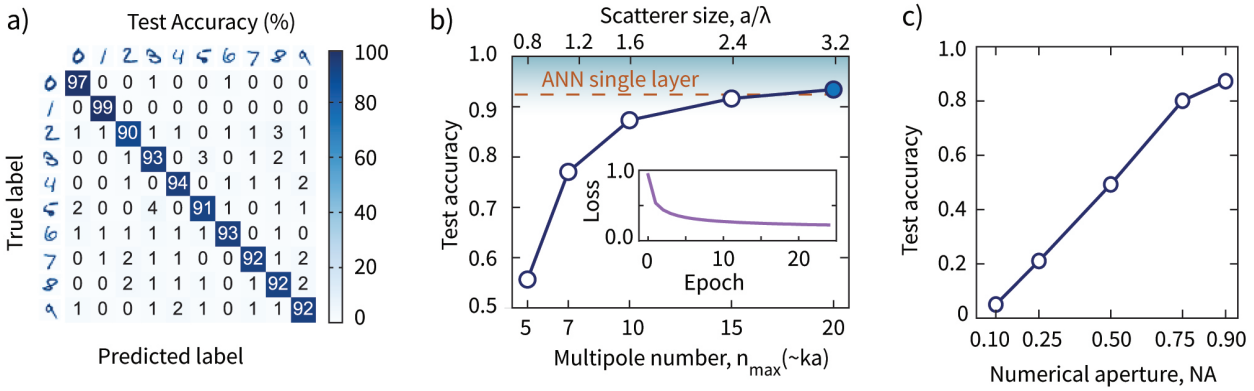


Figure 3. (a) Confusion matrix for the MNIST test set, with rows and columns denoting true and predicted labels, respectively ($n_{\max} = 20$, NA$= 0.9$, corresponds to blue point in (b)). (b) MNIST test accuracy as a function of the maximum retained multipole order $n_{\max}$ for the T-matrix-based optical classifier (NA$= 0.9$). The upper axis shows the corresponding scatterer size parameter a/λ, the dashed line indicates the reference accuracy of a single-layer artificial neural network, and the inset shows the training-loss evolution for $n_{\max} = 20$. (c) Test accuracy as a function of the numerical aperture ($n_{\max} = 10$).

The incident and scattered fields are expanded in the basis of vector spherical harmonics as

$$\mathbf{E}_{\text{inc, sc}}(\mathbf{r}) = \sum_{n=1}^{n_{\max}} \sum_{m=-n}^{n} \left[a_{nm}(p_{nm})\, \mathbf{M}_{nm}^{(1),(3)}(k\mathbf{r}) + b_{nm}(q_{nm})\, \mathbf{N}_{nm}^{(1),(3)}(k\mathbf{r}) \right], \tag{1}$$

where "inc" and "sc" subscripts denote the incident and scattered fields, respectively, and the superscripts (1) and (3) correspond to regular and outgoing electric ($\mathbf{N}_{nm}$) and magnetic ($\mathbf{M}_{nm}$) vector spherical wave functions [35]. The coefficients $a_{nm}(p_{nm})$ and $b_{nm}(q_{nm})$ represent the multipolar content of the incident (scattered) field. The total field is given by the sum of the incident and scattered fields:

$$\mathbf{E}(\mathbf{r}) = \mathbf{E}_{\text{inc}}(\mathbf{r}) + \mathbf{E}_{\text{sc}}(\mathbf{r}). \tag{2}$$

In our formulation, the input object is a far-field image encoded as an angular field distribution incident on the particle after lens transmission. The input image can be represented in the angular domain by a complex vector field defined over the numerical-aperture cone Ω_{NA} right before the lens with $\mathbf{E}_{\text{inp}}^{\infty}(\theta, \phi)$ and right after the lens with $\mathbf{E}_{\text{inc}}^{\infty}(\theta, \phi)$ at the radius $r = f$, where f is the focal length. The angular field is then projected onto the far-field angular parts $\boldsymbol{\mathcal{A}}_{nm}(\theta, \phi)$ and $\boldsymbol{\mathcal{B}}_{nm}(\theta, \phi)$ associated with vector spherical harmonics (see details in Methods Sec. IV A)

$$a_{nm} = \int_{\Omega_{\text{NA}}} \mathbf{E}_{\text{inc}}^{\infty}(\theta, \phi) \cdot \boldsymbol{\mathcal{A}}_{nm}^{*}(\theta, \phi)\, d\Omega, \tag{3}$$

$$b_{nm} = \int_{\Omega_{\text{NA}}} \mathbf{E}_{\text{inc}}^{\infty}(\theta, \phi) \cdot \boldsymbol{\mathcal{B}}_{nm}^{*}(\theta, \phi)\, d\Omega, \tag{4}$$

where $d\Omega = \sin\theta\, d\theta\, d\phi$. These coefficients define the modal representation of the input image. The input image is scalar and linearly polarized, the angular field may be written as

$$\mathbf{E}_{\text{inp}}^{\infty}(\theta, \phi) = u(\theta, \phi)\, \hat{\mathbf{e}}_{\text{pol}}(\theta, \phi), \tag{5}$$

where $u(\theta, \phi)$ is the complex image amplitude and $\hat{\mathbf{e}}_{\text{pol}}$ is the local polarization unit vector. The modal coefficients are then computed by substituting the corresponding field after lens transmission $\mathbf{E}_{\text{inc}}^{\infty}$ into the projection formulas (3)-(4) above.

An example of the MNIST digit that is encoded in the far field as a linearly polarized angular field distribution and projected onto the lens reference sphere via a gnomonic projection is shown in the Figure 2(b). The image is encoded in the phase of the x-component of the electric field on a grid of 28×28 points in θ and ϕ coordinates, i.e $\hat{\mathbf{e}}_{\text{pol}} = \mathbf{e}_x$ and $u = \exp\{i\pi \cdot \text{image}_{\text{proj}}\}$, where $\text{image}_{\text{proj}}$ encodes the gnomonic projection of the input image (see Methods Sec. IV A). The representation of the input image in the multipolar basis

(spherical harmonics "portrait") is shown in Fig. 2(c) for magnetic and electric amplitude coefficients. Evaluated coefficients for all the samples form the MNIST dataset in multipole represenation (see Supplementary Materials Sec. C). One can notice that the input image is not fully captured by the first few multipole coefficients and requires modes with higher n. At the same time, linearly polarized excitation efficiently populates mainly modes with small $|m|$, which is inefficient in terms of input-feature dimensionality, but determines the potential capacity of the spherical multipole representation for the input data.

B. Optical classification procedure

In the proposed architecture, the trainable transformation is implemented by a single compact scatterer within the T-matrix formalism. The T-matrix describes the linear mapping between the incident and scattered multipole coefficients [35], effectively encoding the computational operation of the optical processor into the modal scattering response of a compact object:

$$\begin{pmatrix} \mathbf{p} \\ \mathbf{q} \end{pmatrix} = \begin{pmatrix} \mathbf{T}^{11} & \mathbf{T}^{12} \\ \mathbf{T}^{21} & \mathbf{T}^{22} \end{pmatrix} \begin{pmatrix} \mathbf{a} \\ \mathbf{b} \end{pmatrix}. \tag{6}$$

Here, $\mathbf{a}$, $\mathbf{b}$, $\mathbf{p}$, and $\mathbf{q}$ are column vectors obtained by stacking the coefficients a_{nm}, b_{nm}, p_{nm}, and q_{nm}, respectively. Since the multipole expansion is truncated at $n_{\max}$, the T-matrix generally is a finite-dimensional matrix of size $2N \times 2N$, where $N = n_{\max}(n_{\max}+2)$ is the number of multipole modes retained in the expansion. Therefore, $n_{\max}$ determines not only the quality of the multipole representation of the input data but also the number of trainable parameters in the T-matrix, which scales approximately as $\sim n_{\max}^4$. This quartic scaling reflects dense pairwise coupling between two modal spaces whose individual dimensionalities scale as $n_{\max}^2$. The relation Eq. (6) defines the central linear operation of the optical network (see Supplementary Materials Sec. D) and the weight matrix that can be trained to perform the classification task.

Once the scattered coefficients are obtained, the scattered electromagnetic field is reconstructed from the multipole expansion. In the far field, the scattered field takes the asymptotic form $\mathbf{E}_{\text{sca}}^{\infty}(\theta,\phi)$ based on the asymptotic behavior of the outgoing vector spherical wave functions $\mathbf{M}_{nm}$ and $\mathbf{N}_{nm}$ [35].

The classification signal is obtained by integrating the scattered intensity over a set of

detector regions Ω_k at a radial distance r_{det}, corresponding to either the near- or far-field detection regime:

$$I_k = \int_{\Omega_k} |\mathbf{E}_{\mathrm{sc}}(r_{\mathrm{det}}; \theta, \phi)|^2 \, d\Omega, \tag{7}$$

where $k = 1, \ldots, K$ labels the detectors and K is the number of classes, and, in the case of MNIST dataset, $K = 10$. Each detector therefore collects optical power within a prescribed angular sector. Detectors positions and the resulting intensity in detection plane are illustrated in Figure 2(d) for the far-field formulation where "3" digit is classified correctly with the highest intensity among all the detectors.

Alternatively, the classification can be performed in the *near-field zone*, where the scattered field is evaluated at a distance λ from the surface of the implied scatterer and the intensity is integrated over detector regions Ω_k on a spherical surface with $kr = n_{max} + 2\pi$.

C. Training procedure

The detector outputs $\{I_k\}$ are interpreted as class scores. For a sample with ground-truth label y, we define the classification loss by the softmax cross-entropy

$$\mathcal{L}_{\mathrm{cls}} = -\log \frac{\exp(I_y)}{\sum_{k=1}^{K} \exp(I_k)}. \tag{8}$$

The trainable parameters are the admissible elements of the T-matrix. In our work, all the constraints were implemented using parameterization and regularization of the T-matrix, so the loss function was not explicitly modified to include additional penalty terms. Details of the optimization procedure and the implementation of physical constraints are given in Methods Sec. IV D.

D. Results of training and classification performance

Figure 3 summarizes the performance of the T-matrix-based optical classifier on the MNIST dataset. In panel Fig. 3 (a), the confusion matrix is shown for the case NA = 0.9 and $n_{\mathrm{max}} = 20$, corresponding to the filled blue marker in Fig. 3(b). It confirms that the trained T-matrix achieves strong discrimination across most digit classes, with only a few off-diagonal misclassifications remaining. Fig. 3(b) shows the test accuracy as a function of the retained multipole order n_{max} (NA = 0.9), showing that the classification performance

improves steadily with increasing modal capacity and begins to saturate once enough multipole channels are included. The dashed reference line indicates the accuracy of a single-layer artificial neural network (ANN, see details in Methods Sec. IV D), while the inset shows the training-loss evolution during optimization for $n_{\max} = 20, \mathrm{NA} = 0.9$, indicated by the filled blue marker. Similar results were obtained for the total field formulation (see Supplementary Materials Sec. A). The number of free parameters in ANN is approximately $7.8 \cdot 10^3$ that matches well with number of free parameters in the T-matrix with $n_{\max} = 15$ (approximately $5.9 \cdot 10^3$ with assumed reciprocity, passivity and axial symmetry - see Sec. II E). Thus, a single compact scatterer described by an arbitrary T-matrix with a sufficiently rich multipolar response can perform nontrivial classification tasks on structured optical inputs.

Panel Fig. 3 (c) reports the classification accuracy as a function of the numerical aperture for fixed multipole order $n_{\max} = 10$. The required number of multipoles is directly related to the NA of the system. Indeed, small NA will require significantly larger number of angular frequencies to capture the features within small $\theta_{\max}$ and processed by the scatterer. Thus, simultaneous increase of the NA and the number of multipole modes leads to improved classification performance, while insufficient modal capacity or limited NA can suppress the ability of the system to discriminate between different classes effectively.

Together, the results in Fig. 3 indicate that both the number of retained multipole modes and the numerical aperture are key factors in achieving reliable optical classification with the T-matrix-based architecture.

E. Physical constraints

One of the central questions in the design of optical neural networks is how physical constraints on wave propagation and scattering affect the achievable performance and scalability of the system [10, 19]. In the case of a single scatterer, unlike an unconstrained neural-network weight matrix, the T-matrix is restricted by electromagnetic principles [36]. There are several key constraints that must be considered when designing and training the T-matrix-based optical classifier, including reciprocity, passivity, energy conservation, and symmetry relations. While the results shown in Fig. 3 were obtained by enforcing these constraints without explicitly discussing them, it is important to understand how they can be incorporated into the T-matrix and how they affect the classification performance.

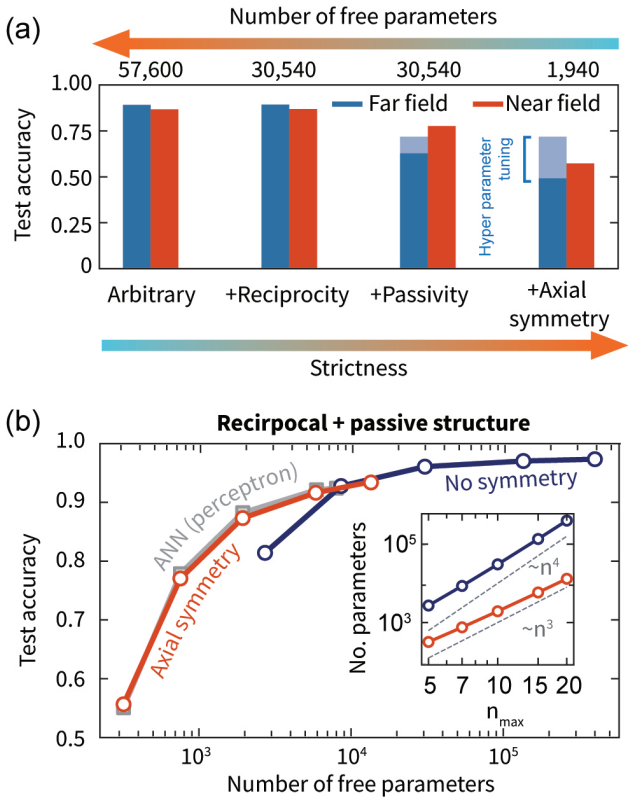


Figure 4. (a) Classification accuracy for reciprocal, passive, and axially symmetric T-matrices trained at $n_{\max} = 10$ and NA $= 0.5$ in both near- and far-field formulations. The panel shows how reciprocity, passivity, and axial symmetry affect the achievable performance under electromagnetic constraints. (b) Comparison of the constrained T-matrix classifiers with a single-layer ANN of similar parameter count, demonstrating that the physically constrained multipolar scatterer can achieve competitive accuracy.

Reciprocity. While there are many theoretical and experimental examples of non-reciprocal single scatterers [37–39], in the context of passive linear optical systems, reciprocity remains a fundamental constraint in the T-matrix formalism. Reciprocity imposes symmetry relations between coupling channels [35]:

$$T^{kl}_{n,m,n',m'} = (-1)^{m+m'} T^{lk}_{n',-m',n,-m}, \tag{9}$$

where $k, l = 1, 2$ denote the electric and magnetic multipole channels, respectively, and n, m are the multipole indices. When reciprocity is imposed, the number of free parameters is reduced by approximately a factor of two. For sufficiently large $n_{\max}$, the resulting accuracy can remain comparable to that of an unconstrained T-matrix because the model still retains

enough trainable degrees of freedom.

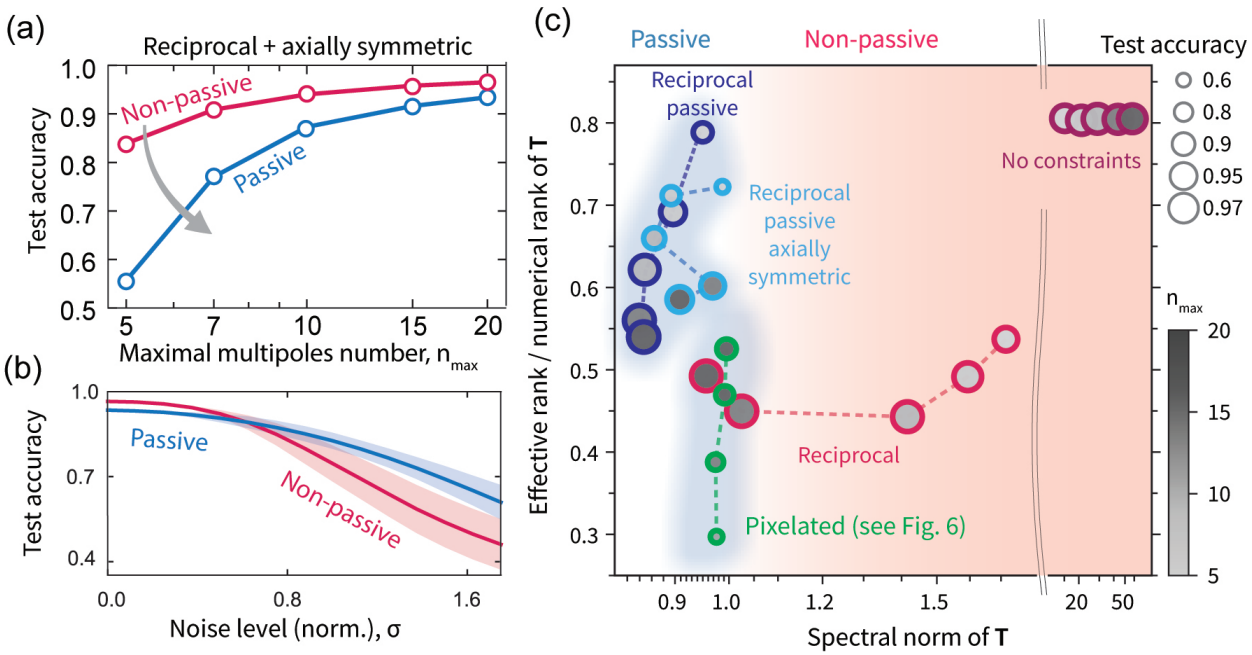


Figure 5. (a) Accuracy as a function of retained multipole order for the far-field formulation at NA = 0.9 with symmetrical reciprocal T-matrices, showing that higher modal capacity reduces the performance gap introduced by passivity. (b) Classification accuracy of the T-matrix-based classifier as a function of the noise magnitude σ for the symmetrical reciprocal passive T-matrix and symmetrical reciprocal T-matrix for $n_{\max} = 20$ and NA = 0.9. (c) Distribution of the optimized T-matrices and accuracy in the coordinate space defined by the spectral norm of the scattering matrix $\|\mathbf{T}\|_2$ and the spectral entropy rank of the T-matrix.

Passivity. Further constrains the admissible parameter space by requiring that the scattered power does not exceed the incident power. In the T-matrix formalism, equality of scattered and incident power imposes a generalized unitarity (optical theorem) condition on the transition operator. In operator form, this can be written as

$$\mathbf{T} + \mathbf{T}^{\dagger} + 2\,\mathbf{T}^{\dagger}\mathbf{T} = 0. \tag{10}$$

In terms of S-matrix, where $\mathbf{S} = \mathbf{E} + 2\mathbf{T}$, this condition can be written as $\mathbf{S}^{\dagger}\mathbf{S} = \mathbf{E}$ with $\|\mathbf{S}\|_2 = 1$. The required passivity condition can be formulated as

$$\|\mathbf{S}\|_2 \leq 1, \tag{11}$$

which ensures that the extracted power (extinction) equals the sum of scattered and absorbed power and that no net energy is generated by the system. Equivalently, in terms of modal coefficients, the passivity condition implies

$$\|\mathbf{p}\|^2 + \|\mathbf{q}\|^2 \leq \|\mathbf{a}\|^2 + \|\mathbf{b}\|^2, \tag{12}$$

for any incident field amplitudes $(\mathbf{a}, \mathbf{b})$, where $\|\cdot\|$ denotes the appropriate power-normalized norm of multipole coefficients. The passivity constraint does not reduce the nominal number of parameters, but it restricts optimization to the manifold of passive scatterers. Therefore, passivity can reduce the achievable accuracy even when the number of trainable coefficients is unchanged.

Axial symmetry of the scatterer can be additionally imposed by requiring that the T-matrix elements vanish for $m \neq m'$. This constraint, on the contrary to previously discussed limitations, substantially reduces the number of free parameters because the T-matrix becomes block-diagonal in m and couples only modes with identical azimuthal index:

$$T^{kl}_{n,m,n',m'} = \delta_{mm'}\, T^{kl}_{n,n',m,m}. \tag{13}$$

This symmetry of the T-matrix leads to a reduction in the number of independent parameters and also results in $\sim n_{\max}^3$ scaling. In the following, we will rely on axially symmetric scatterer designs in order to reduce the computational complexity during the scatterer design process and to reduce the number of trainable parameters.

Thus, the nominally large modal-coupling space of the T-matrix is progressively restricted by physical constraints: reciprocity and symmetry reduce number of independent parameters, while passivity restricts the admissible operator spectrum. More generally, not every desired set of input-output transformations can be realized by a single physical structure [36].

The results of the simulations with the physically constrained T-matrix are summarized in Figure 4. Panel (a) compares classification accuracy under different physical constraints, starting from an arbitrary T-matrix and then sequentially imposing reciprocity, passivity, and axial symmetry at $n_{\max} = 10$ and NA $= 0.5$ in both near- and far-field formulations. The corresponding number of free parameters is shown above each case. Reciprocity alone has only a mild impact on performance, whereas passivity suddenly causes a clear accuracy decrease despite leaving the nominal parameter count unchanged. This indicates that the main effect of passivity is not parameter reduction, but the spectral restriction imposed on

the scattering operator: the classifier can no longer use arbitrary amplification or non-passive mode mixing to separate classes.

The scaling of the accuracy with the number of trainable parameters is shown in Figure 4 (b) for the case of reciprocal and passive structure (see results for the other constraints combinations in Supplementary Materials Sec. B). The blue line corresponds to the reciprocal passive system, while the orange line corresponds to the axially symmetric case. In both cases, the multipole order $n_{\max}$ is varied, so that the number of free parameters changes together with the modal capacity of the classifier. Because axial symmetry conserves the azimuthal index m, the axially symmetric T-matrix contains approximately one order of magnitude fewer trainable parameters for the same $n_{\max}$. Nevertheless, its accuracy scales comparably to that of the ANN baseline and can remain competitive even with a substantially smaller parameter count. Here, the single-layer ANN was implemented as a linear classifier and trained similarly using softmax cross-entropy loss and masking out a random subset of weights to investigate dependence on the parameters number. The inset of Figure 4 (b) shows the corresponding scaling of the T-matrix number of parameters with $n_{\max}$.

The effect of accuracy drop due to passivity, in our opinion, deserves a more detailed discussion. Figure 5 (a) shows the accuracy as a function of the retained multipole order in the far-field formulation for reciprocal structures with axial symmetry. The comparison highlights that passivity leads to a significant drop in accuracy, although the number of trainable parameters is the same as in the reciprocal non-passive case. This behavior highlights the role of passivity as a spectral constraint and indicate that the performance loss cannot be explained only by the parameter count.

In the non-passive model, training can exploit large singular values of the T-matrix to amplify class-separating directions in the multipolar feature space. In the passive model, the associated scattering matrix must satisfy $\|\mathbf{S}\|_2 \leq 1$, so the classifier cannot increase the optical power in arbitrary outgoing channels. This distinction also suggests that the passive model may be more sensitive to the details of the training procedure. Moreover, this bound limits the scale of the detector intensities and, consequently, the classification margins for a fixed incident field. This is particularly relevant for physical neural networks, where the squared field amplitudes at the detectors can be small. In this case, the resulting bound on the cross-entropy loss and the reduced gradients can make it difficult to separate samples across the decision boundary during finite-time optimization (see Supplementary Materi-

als Sec. E). Related studies on Lipschitz-constrained networks with spectral normalization report competitive accuracy and, in some cases, improved robustness [40]. In practice, however, achieving such performance requires careful training design, including appropriate dataset normalization to produce logits with a suitable scale for cross-entropy optimization. Alternatively, this scale can be controlled by introducing a softmax temperature τ [41]. The loss depends on the classification margins m through m/τ, so choosing $\tau < 1$ increases the effective margin and strengthens the training signal from a given physical intensity difference. Temperature can therefore partially compensate for the optimization penalty imposed by bounded optical gain.

At the same time, bounded gain suppresses the amplification of perturbations. To test this regularization effect directly, we add complex Gaussian perturbations to the trained T-matrix during inference,

$$\mathbf{T}_{\text{noisy}} = \mathbf{T} + \sigma \cdot \gamma \cdot \mathbf{N}, \quad \mathbf{N} \sim \mathcal{CN}(0, 1)$$

where σ is the dimensionless noise magnitude and $\mathbf{N}$ is a matrix of independent and identically distributed complex Gaussian samples. The scale factor $\gamma = \frac{\|T\|_F}{\sqrt{N}}$ is chosen from the normalized Frobenius norm of the trained matrix, so that the perturbation amplitude is proportional to the typical magnitude of the T-matrix elements. Figure 5 (b) compares axially symmetric reciprocal T-matrices trained with and without passivity at $n_{\max} = 20$ and NA $= 0.9$. It is observed that the passive T-matrix model exhibits stronger robustness to noise compared to the unconstrained model, which can be attributed to the regularizing effect of the passivity constraint that limits the spectral norm of the T-matrix and prevents excessive amplification of perturbations.

Finally, in Fig. 5 (c) we summarized all optimized T-matrix solutions in the coordinates of the spectral norm $\|\mathbf{T}\|_2$ and the normalized effective rank

$$r_{\text{eff}} = \frac{\exp\left(-\sum_i \tilde{s}_i \log \tilde{s}_i\right)}{\text{rank}(\mathbf{T})}, \qquad \tilde{s}_i = \frac{s_i}{\sum_j s_j}, \tag{14}$$

with s_i being the singular values of $\mathbf{T}$. The quantity r_{eff} estimates the fraction of singular scattering channels that are effectively used by the trained operator. Values close to unity correspond to a broad singular-value spectrum, where many channels contribute with comparable weights, while lower values indicate that the transformation is dominated by a smaller number of relatively high-gain channels. This quantity shows the distinction between nominal parameter count and effectively utilized scattering channels demonstrating

that large T-matrix does not by itself imply a proportionally large information capacity [42]. The spectral norm of $\mathbf{T}$ is not itself a complete passivity criterion, because passivity is imposed on $\mathbf{S} = \mathbf{I} + 2\mathbf{T}$ through $\|\mathbf{S}\|_2 \leq 1$. Nevertheless, $\|\mathbf{T}\|_2 > 1$ is incompatible with passivity, while $\|\mathbf{T}\|_2 \leq 1$ is only a necessary, not sufficient, indicator of passive behavior.

In all the types of solutions accuracy (shown with the circle size in Fig. 5(c)) grows with number of multipoles from $n_{\max} = 5$ to $n_{\max} = 20$ (shown with the grey filling color), as was already discussed previously, and reaches high accuracy even with small number of multipole modes used for the unconstrained and reciprocal non-passive solutions. These solutions occupy the larger $\|\mathbf{T}\|_2$-values, while passive solutions are confined to the low-norm region. The effective-rank trend shows that non-constrained models use all the available channels, while passive and axially symmetric passive classifiers tend to use relatively fewer effective singular channels as the modal basis grows.

To recapitulate, Fig. 5 identifies two distinct mechanisms: non-constrained models improve classification by channels amplification, whereas passive models trade part of this accuracy for bounded gain, physical admissibility, and enhanced robustness to parameter noise.

F. Designing a physical scatterer

The results above establish that reciprocity, passivity, and axial symmetry still leave sufficient modal capacity for nontrivial optical classification. However, these constraints are only necessary conditions for physical realizability: an arbitrary T-matrix satisfying them does not necessarily correspond to a scatterer with a realizable spatial distribution of material parameters. In particular, a large number of T-matrix coefficients does not imply an equal number of independently controllable material parameters, which simultaneously affect all modal couplings. We therefore proceed from the operator-level description to a direct design of an inhomogeneous dielectric particle and examine how closely such a structure can approach the performance of the optimized physically constrained T-matrix.

The particle design and the corresponding classification results are summarized in Fig. 6. Here we restricted the particle geometry to be axially symmetric that reduces the number of trainable parameters and simplifies the optimization. As shown in Fig. 6 (a), the axially symmetric particle is discretized into pixels in the radial and polar directions. Each pixel is

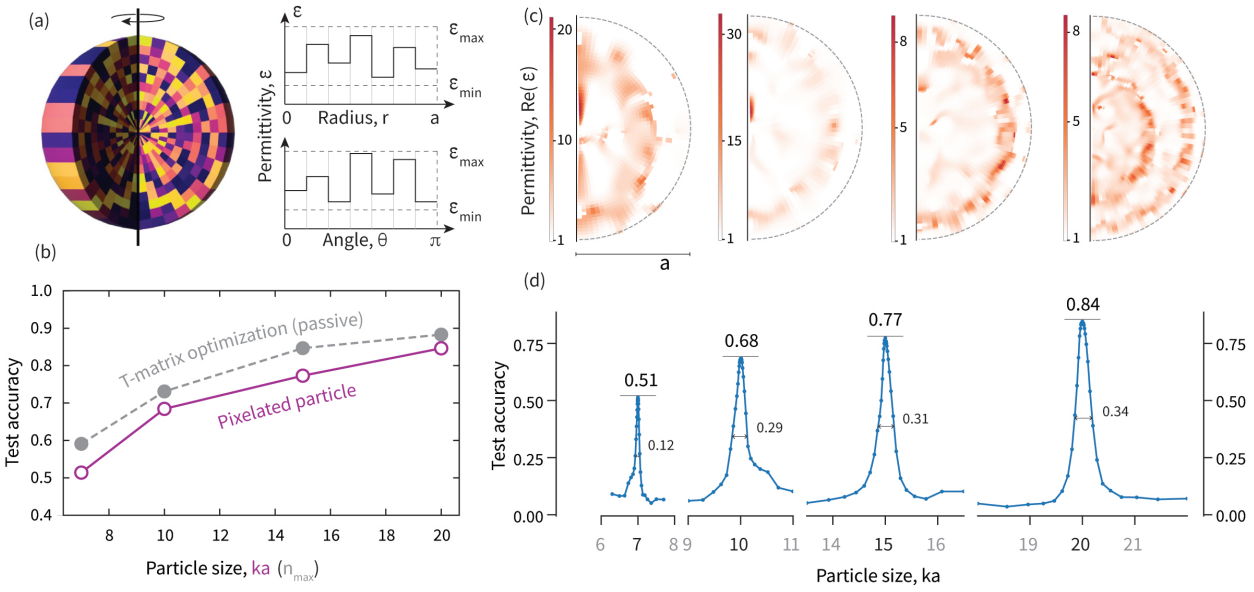


Figure 6. (a) Parameterization of an axially symmetric pixelated particle with a trainable relative-permittivity distribution discretized along the radial coordinate r and polar angle θ, with each pixel constrained to $\varepsilon_{\min} \leq \varepsilon \leq \varepsilon_{\max}$. (b) Classification accuracy of the optimized pixelated particle compared with the reciprocal, passive, and axially symmetric T-matrix as a function of the particle size parameter ka and the associated multipole truncation order $n_{\max}$. (c) Optimized permittivity distributions obtained for representative particle sizes. (d) Test accuracy as a function of ka in the vicinity of the selected resonant solutions. The horizontal markers and numerical values indicate the maximum accuracy achieved near $ka = 7$, 10, 15, and 20, while the horizontal arrows and the corresponding numerical values denote the full width at half maximum (FWHM) of each resonance.

assigned a trainable relative permittivity bounded by $\varepsilon_{\min}$ and $\varepsilon_{\max}$. The absorption losses are set to zero for this particular case so $\mathrm{Im}\,\varepsilon = 0$ and $\mathrm{Re}\,\varepsilon \geq 1$. Details of the optimization are given in Methods Sec. IV D.

The T-matrix of the pixelated particle is computed using the invariant imbedding T-matrix method (IITMM), with the computational procedure and relevant equations described in Methods Sec. IV C. IITMM is an approach for calculating the T-matrices of inhomogeneous particles. The method is particularly efficient for axially symmetric structures, providing a significant computational speedup in such cases[31, 43, 44].

Figure 6 (b) compares the classification accuracy of the optimized pixelated particle with

that of the arbitrary T-matrix with reciprocal, passive, and axially symmetry restrictions (axis of symmetry is aligned with the z-axis). The horizontal axis relates the particle size parameter ka to the corresponding multipole truncation order $n_{\max}$ ($n_{\max} \approx ka$). The pixelated realization follows the same general trend as the ideal constrained operator: the accuracy increases as the particle becomes larger and more multipolar channels become available. Although the directly optimized material distribution performs slightly below the T-matrix model over the considered range, the relatively small gap demonstrates that a substantial fraction of the operator-level performance can be retained in a spatially realizable scatterer. In these simulations the number of pixels was set to $N_r = ka \cdot 5$ and $N_\theta = 60$ in the radial and polar directions, respectively, which was found to be sufficient to achieve convergence of the classification accuracy. For each value of ka, the T-matrix was computed with a truncation order of $n_{\max} = ka + 5$ multipoles, which provides sufficient evaluation accuracy and accounts for the exponentially decaying amplitudes of the matrix elements.

Examples of the optimized permittivity distributions for several particle sizes are presented in Fig. 6 (c). The resulting structures are strongly inhomogeneous and exploit both radial and angular variations of the permittivity, indicating that classification relies on coordinated coupling among multiple resonant channels rather than on a single conventional Mie resonance. Figure 6 (d) further shows that the test accuracy depends nonmonotonically on the particle size and exhibits narrow resonant maxima near $ka = 7$, 10, 15, and 20 and the accuracy increases from approximately 0.51 to 0.84 across these parameters. At the same time, high classification accuracy is obtained only in the very narrow spectral range of around $\delta\lambda/\lambda \approx 2.1 \cdot 10^{-2}$ when the material distribution and the size parameter jointly realize a favorable multipolar response. Out of this band, the accuracy drops to insufficient values.

Although the finely pixelated geometry considered here would be difficult to fabricate and implement directly at optical or microwave frequencies, it is not intended as a fabrication-ready design. Rather, it serves as a proof of principle that a T-matrix with the functionality required for classification can arise from a finite passive object with realistic material parameters. Advances in two-photon polymerization have enabled the fabrication of complex three-dimensional photonic structures with subwavelength scale features [45, 46], including inverse-designed volumetric scattering systems [47]. In this context, the target modal response could, in principle, be approximated using more fabrication-oriented geometries, in-

cluding composite scatterer arrangements optimized within the T-matrix framework [48, 49], multilayer and core-shell particles [50] and inverse-designed resonators with tailored multipolar responses [51, 52] or voxels [53].

III. CONCLUSIONS AND OUTLOOK

An advantage of the proposed architecture is the concentration of a dense many-to-many modal transformation within a single finite scatterer with the density of free parameters overcoming the conventional layered diffractive elements by the factor of $\sim (ka)^2$ which in the case of a scatterer comparable with wavelength can give increase in two orders of magnitude. While limited density of parameters layered structures can be potentially overcome by non-locally coupled metasurfaces [33], the single scatterer immediately provides an efficient solution. This does not resolve the issue of the limited communication channel as the channel space scales as $\sim (ka)^2$ [32]. Nevertheless, one of the big advantage of the proposed geometry is that all the computation can be performed within a near-field domain, where the number of channels can be significantly larger than in the far-field. That opens a route towards the optical computation at the interface with the integrated photonic systems allowing to directly couple preprocessed optical data to on-chip photonic circuitry.

Along the same line, one should note that the concentration of the trainable transformation in a single scatterer does not remove the longitudinal-space requirement of the complete optical system. The phase-encoded image must first be prepared over a finite aperture and focused onto the particle, so propagation between the image plane or spatial light modulator and the scatterer remains necessary. Likewise, far-field readout requires collection optics or a propagation distance, although the near-field formulation can replace this output path by direct coupling to detectors or an integrated photonic circuit. Thus, the present architecture primarily reduces the volume occupied by the trainable transformation itself; it does not evade the diffraction- and communication-based thickness bounds that apply to the entire input–processor–output system [18].

To conclude, in this work we have demonstrated that a single compact scatterer can be trained to perform nontrivial optical classification tasks on structured inputs. The T-matrix formalism provides a natural framework for describing the trainable transformation and for enforcing physical constraints, including reciprocity, passivity, and symmetry. The results

show that even under these constraints, the T-matrix retains sufficient modal capacity to achieve competitive classification accuracy. Furthermore, we have shown that a physically realizable inhomogeneous dielectric particle can approximate the performance of the optimized T-matrix, demonstrating the feasibility of implementing such optical classifiers in practice. This work opens new avenues for designing compact and efficient optical neural networks based on single-scatterer architectures.

IV. METHODS

A. Incident multipole coefficients derivation

Transmitted through the ideal lens field can be expressed [54] as

$$\mathbf{E}^{\infty}_{\text{inc}}(\theta,\phi) = \left[(\mathbf{E}^{\infty}_{\text{inp}}\cdot\mathbf{n}_{\phi})\mathbf{n}_{\phi} + (\mathbf{E}^{\infty}_{\text{inp}}\cdot\mathbf{n}_{\rho})\mathbf{n}_{\theta}\right](\cos\theta)^{1/2}, \tag{15}$$

where $\mathbf{n}_{\phi} = (-\sin\phi, \cos\phi, 0)$, $\mathbf{n}_{\rho} = (\cos\phi, \sin\phi, 0)$, and $\mathbf{E}^{\infty}_{\text{inp}}$ is the encoded field in the input plane, that is in our case is the gnomonic projection of the image on the lens (see Figure 2) encoded in the phase of only nonzero x-component of the field:

$$\text{image}_{\text{proj}}(\theta,\phi) = \text{image}\left(\frac{\tan(\pi-\theta)}{\tan\theta_{\max}}\begin{pmatrix}-\sin\phi\\ \cos\phi\end{pmatrix}\right),$$

where image is the input image given in cartesian coordinates, and $\Theta_{\max}$ is the maximum angle defined by the numerical aperture of the lens, and:

$$\mathbf{E}^{\infty}_{\text{inp}} = \mathbf{e}_x u(\theta,\phi), \quad u(\theta,\phi) = \exp\{i\pi\text{image}_{\text{proj}}\}.$$

The angular spectrum representation of the focal field is given by the angular representation of the far field on the lens reference sphere [54]:

$$\mathbf{E}_{\text{inc}} = \gamma\int_0^{\theta_{\max}}\int_0^{2\pi}\mathbf{E}^{\infty}_{\text{inc}}(\theta,\phi)\text{PW}(\theta,\phi)d\Omega,$$

where $d\Omega = \sin\theta\, d\phi\, d\theta$, θ_{max} is the maximum angle defined by the numerical aperture of the lens, $\gamma = -\frac{ikfe^{-ikf}}{2\pi}$, where f is focal length and $\text{PW}(\theta,\phi) = e^{ikz\cos\theta + ik\rho\sin\theta\cos(\phi-\varphi)} = e^{ik_x x + ik_y y + ik_z z}$.

Since the multipole coefficients $[a/b]^{\text{PW}}_{mn}(\mathbf{E}^{\infty}_{\text{inc}},\theta,\phi)$ are known for plane waves [35]:

$$a^{\text{PW}}_{mn}(\theta,\phi) = i\beta_{m,n}\mathbf{E}^{\infty}_{\text{inc}}(\theta,\phi)\cdot\mathbf{C}^{*}_{mn}(\theta)e^{-im\phi},$$
$$b^{\text{PW}}_{mn}(\theta,\phi) = \beta_{m,n}\mathbf{E}^{\infty}_{\text{inc}}(\theta,\phi)\cdot\mathbf{B}^{*}_{mn}(\theta)e^{-im\phi},$$

with $\beta_{m,n} = 4\pi(-1)^m i^{n-1}\sqrt{(2n+1)/(4\pi n(n+1))}$ and $\mathbf{C}_{mn}(\theta)$ and $\mathbf{B}_{mn}(\theta)$ are the angular parts of the magnetic and electric multipoles, respectively [35]. the incident field can be expressed through the multipole expansion

$$\mathbf{E}_{\text{inc}} = \gamma\iint\sum_{m,n}[a^{\text{PW}}_{mn}\mathbf{M}^{(1)}_{mn} + b^{\text{PW}}_{mn}\mathbf{N}^{(1)}_{mn}]d\Omega,$$

where $\mathbf{M}_{mn}^{(1)}$ and $\mathbf{N}_{mn}^{(1)}$ are the regular vector spherical wave functions [35] Finally, changing the order of summation and integration gives multipole decomposition of focused incident field

$$\mathbf{E}_{\text{inc}} = \sum_{n=1}^{\infty} \sum_{m=-n}^{n} \left[a_{mn} \mathbf{M}_{mn}^{(1)}(k\mathbf{r}) + b_{mn} \mathbf{N}_{mn}^{(1)}(k\mathbf{r}) \right]$$

with integral expressions for the coefficients $[a/b]_{mn}$ given by

$$[a/b]_{mn} = \gamma \int_0^{\theta_{\max}} \int_0^{2\pi} [a/b]_{mn}^{\text{PW}}(\theta, \phi) d\Omega.$$

This yields the final expressions for Eqs. (3) and (4):

$$\boldsymbol{\mathcal{A}}_{mn}^{*}(\theta, \phi) = i\gamma\beta_{m,n} \mathbf{C}_{mn}^{*}(\theta) e^{-im\phi},$$

$$\boldsymbol{\mathcal{B}}_{mn}^{*}(\theta, \phi) = \gamma\beta_{m,n} \mathbf{B}_{mn}^{*}(\theta) e^{-im\phi}.$$

B. Density of parameters

For a fixed reference area A, taken as the projected area of the scatterer and the physical area of the diffractive layer, we estimate the scaling of the number and density of parameters for a T-matrix representation and a diffractive layer. Here, a is the particle radius, k is the wavenumber, and λ is the wavelength.

1. Parameter estimate for the T-matrix

The number of parameters of an unconstrained T-matrix (modal coupling coefficients) is determined by the matrix dimension, which is set by the maximum retained multipole order $n_{\max}$:

$$N_T \sim [2n_{\max}(n_{\max} + 2)]^2 \sim 4n_{\max}^4$$

The relevant characteristic length scale is the particle radius a, with the minimal scaling requirement:

$$n_{\max} \sim ka$$

The parameter count therefore has the following asymptotic scaling with ka:

$$N_T \sim 4(ka)^4$$

For a fixed area $A = \pi a^2$, the T-matrix parameter count can be written as:

$$N_T \sim 4\left(\frac{2\pi a}{\lambda}\right)^4 \sim 64\pi^4\left(\frac{a^2}{\lambda^2}\right)^2 \sim 64\pi^2\frac{A^2}{\lambda^4}$$

The corresponding volumetric parameter density of the T-matrix representation is:

$$\rho_T \sim \frac{N_T}{(4/3)\pi a^3} \sim \frac{3}{\pi a^3}(ka)^4 \sim \frac{48\pi^3 a}{\lambda^4}$$

2. Parameter estimate for a diffractive layer

The characteristic length scale of a diffractive layer is its period d, taken as:

$$d \sim \frac{\lambda}{2}\cdot\frac{1}{\alpha}$$

Here, α is a dimensionless factor introduced to describe metasurfaces for which d can be smaller than $\lambda/2$. The number of parameters of the diffractive layer is then:

$$N_{dif} \sim \frac{A}{d^2} \sim \frac{4A}{\lambda^2}\alpha^2$$

In terms of ka (using $A = \pi a^2$), the parameter count has the asymptotic scaling:

$$N_{dif} \sim \frac{4\pi a^2}{\lambda^2}\alpha^2 \sim \frac{\alpha^2}{\pi}(ka)^2$$

The corresponding areal parameter density of the diffractive layer is:

$$\rho_{dif} \sim \frac{N_{dif}}{A} \sim \frac{4}{\lambda^2}\alpha^2$$

C. IITMM

We calculate the T-matrix of the inhomogeneous particle using the invariant imbedding T-matrix method (IITMM) [44]. In the IITMM formulation, the radial evolution of the T-matrix is determined by the matrix $U_{mm'nn'}(r)$, whose elements contain angular integrals over the particle permittivity distribution at each radius r. For a pixelated particle with piecewise-constant angular permittivity, these angular integrals can be evaluated analytically, avoiding their numerical evaluation at every radial step. Following the approach of Ref. [55], we use a Legendre expansion of the angular permittivity profile.

Specifically, we write

$$u(r,\theta) = \varepsilon(r,\theta) - 1 = \sum_{p=0}^{\infty} u_p(r) P_p(\cos\theta),$$

$$\frac{u(r,\theta)}{\varepsilon(r,\theta)} = \sum_{p=0}^{\infty} \widetilde{u}_p(r) P_p(\cos\theta),$$

where P_p denotes the Legendre polynomial of order p.

For an axially symmetric particle, the resulting matrix has the block form

$$U_{mm'nn'}(r) = \frac{k^2 r^2}{2} \delta_{mm'} \begin{pmatrix} U^{11}_{mnn'}(r) & -U^{12}_{mnn'}(r) & 0 \\ U^{12}_{mnn'}(r) & U^{11}_{mnn'}(r) & 0 \\ 0 & 0 & U^{33}_{mnn'}(r) \end{pmatrix}.$$

Its nonzero components can be evaluated directly from the Legendre coefficients as

$$U^{11}_{mnn'} = -2(-1)^m \gamma_{nn'} \sum_{p=0}^{\infty} u_p \left[1 - \big(n + n' + p \bmod 2\big)\right] C^{npn'}_{m\,0\,-m} C^{npn'}_{1\,0\,-1},$$

$$U^{12}_{mnn'} = -2i(-1)^m \gamma_{nn'} \sum_{p=0}^{\infty} u_p \big(n + n' + p \bmod 2\big) C^{npn'}_{m\,0\,-m} C^{npn'}_{1\,0\,-1},$$

$$U^{33}_{mnn'} = 2(-1)^m \gamma_{nn'} \sum_{p=0}^{\infty} \widetilde{u}_p C^{npn'}_{0\,0\,0} C^{npn'}_{m\,0\,-m}.$$

where $\gamma_{nn'} = \sqrt{(2n+1)(2n'+1)}$ and $C^{j_1 j_2 j_3}_{m_1 m_2 m_3}$ denotes the corresponding Wigner $3j$-symbol.

For a single angular pixel occupying $\theta_1 < \theta < \theta_2$ and having relative permittivity ε, its contribution to the Legendre coefficients is

$$u_p = (\varepsilon - 1) w_p, \qquad \widetilde{u}_p = \left(1 - \frac{1}{\varepsilon}\right) w_p, \tag{16}$$

where

$$w_p = \begin{cases} \dfrac{1}{2}\left[P_0(\cos\theta_1) - P_0(\cos\theta_2)\right], & p = 0, \\ \dfrac{1}{2}\left[P_{p-1}(\cos\theta_2) - P_{p+1}(\cos\theta_2) - P_{p-1}(\cos\theta_1) + P_{p+1}(\cos\theta_1)\right], & p \geq 1. \end{cases} \tag{17}$$

Consequently, the angular integrals required by the IITMM can be obtained analytically for each pixel and assembled linearly for an arbitrary piecewise-constant angular permittivity profile.

D. Optimization procedure

T-matrix optimization. The T-matrix is optimized directly with respect to the classification objective of Eq. (8), starting from a random complex-valued initialization. For each batch, the incident multipole coefficients are transformed $(\mathbf{a}, \mathbf{b})$ by the current T-matrix according to Eq. (6), and the resulting detector powers are used as class scores. The cross-entropy loss is differentiated with respect to the trainable complex-valued T-matrix parameters, which are updated using AdamW gradient-based optimizer. The standard MNIST training set is divided into 55 000 training and 5 000 validation samples, while the official 10 000-sample test set is used only for the final evaluation.

Different physical constraints are incorporated into the same optimization procedure by restricting the admissible T-matrix rather than modifying the classification loss. For the unconstrained model, the independent complex valued matrix elements are optimized directly. Reciprocity and geometrical symmetries are imposed through the parameterization of the T-matrix, such that only the independent matrix elements are treated as trainable variables and the remaining elements are reconstructed from the corresponding symmetry relations. Passivity is enforced during optimization by projecting the reconstructed scattering matrix $\mathbf{S}$ onto the unit spectral norm ball satisfying $\|\mathbf{S}\|_2 \leq 1$. The projection is performed by singular-value decomposition of $\mathbf{S}$ and restricting singular values that exceed unity.

The T-matrix models were optimized for approximately 10^2 epochs, using batch sizes of order 10^2 and learning rates in the range 10^{-4} to 10^{-3}.

Direct optimization of the pixelated particle. For the physically realizable particle, the optimization is performed directly in the space of material distributions rather than by first optimizing a T-matrix and subsequently fitting a particle to it. The trainable variables are the relative-permittivity values $\varepsilon(r, \theta)$ of the pixels forming the axially symmetric particle. For every forward evaluation, the T-matrix corresponding to the current permittivity distribution is calculated using the IITMM described in Sec. IV C. This T-matrix is then used in exactly the same optical classification pipeline as above to obtain the detector signals and the cross-entropy loss. Since the T-matrix calculation is differentiable with respect to the material parameters, gradients of the classification loss are propagated through the IITMM calculation directly to the permittivity pixels. For the non-absorbing particles considered here, $\operatorname{Im}\varepsilon = 0$ and $\operatorname{Re}\varepsilon$ is constrained to the allowed range $[1; +\infty)$.

The direct particle optimization is computationally more demanding because the particle T-matrix is recomputed during every forward pass. Accordingly, these optimizations were performed for order of 10^2 epochs using batches of order 10^3 samples.

Artificial neural network. For ANN training optimization, we use the Adam optimizer [56] with a learning rate of order 10^{-4} and a batch size of 256 for 50 epochs. Similar splitting was used for training as for the T-matrix and pixel particle optimization.

ACKNOWLEDGEMENTS

The authors thank Andrey Bogdanov and Alexey Shcherbakov.

FUNDING

The work was supported by the Federal Academic Leadership Program Priority 2030.

AUTHOR CONTRIBUTIONS

V.K. contributed to methodology, software, formal analysis, investigation, visualization, conceptualization and writing of the original draft. V.I. contributed to methodology, software, formal analysis, investigation, visualization, conceptualization, and writing of the original draft. C.D.A. contributed to conceptualization, writing – review and editing. M.P. contributed to conceptualization, methodology, software, formal analysis, visualization, writing of the original draft, and supervision. All authors discussed the results and reviewed the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

DATA AVAILABILITY

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

CODE AVAILABILITY

The code that supports the findings of this study is available from the corresponding author upon request.

---

[1] Ali Momeni, Babak Rahmani, Benjamin Scellier, Logan G Wright, Peter L McMahon, Clara C Wanjura, Yuhang Li, Anas Skalli, Natalia G Berloff, Tatsuhiro Onodera, et al. Training of physical neural networks. *Nature*, 645(8079):53–61, 2025.

[2] Yandong Li and Francesco Monticone. Exploring the role of metamaterials in achieving advantage in optical computing. *Nature Computational Science*, 4(8):545–548, 2024.

[3] Alexandre Silva, Francesco Monticone, Giuseppe Castaldi, Vincenzo Galdi, Andrea Alu, and Nader Engheta. Performing Mathematical Operations with Metamaterials. *Science*, 343(6167):160–163, 2014.

[4] Jingtian Hu, Deniz Mengu, Dimitrios C. Tzarouchis, Brian Edwards, Nader Engheta, and Aydogan Ozcan. Diffractive optical computing in free space. *Nature Communications*, 15(1):1525, 2024.

[5] Bhavin J Shastri, Alexander N Tait, Thomas Ferreira de Lima, Wolfram HP Pernice, Harish Bhaskaran, C David Wright, and Paul R Prucnal. Photonics for artificial intelligence and neuromorphic computing. *Nature Photonics*, 15(2):102–114, 2021.

[6] Joseph W. Goodman. *Introduction to Fourier optics*. McGraw-Hill physical and quantum electronics series. McGraw-Hill, New York, reissued edition, 1988.

[7] Kelvin Wagner and Demetri Psaltis. Multilayer optical learning networks. *Applied Optics*, 26(23):5061, 1987.

[8] Xing Lin, Yair Rivenson, Nezih T. Yardimci, Muhammed Veli, Yi Luo, Mona Jarrahi, and Aydogan Ozcan. All-optical machine learning using diffractive deep neural networks. *Science*, 361(6406):1004–1008, 2018.

[9] Thanh Xuan Hoang, Ayan Nussupbekov, Jie Ji, Daniel Leykam, Jaime Gómez Rivas, and Yuri Kivshar. Mie-tronics supermodes and symmetry breaking in nonlocal metasurfaces. *Physical Review Research*, 8(3):033076, 2026.

[10] David A. B. Miller. Communicating with waves between volumes: evaluating orthogonal spatial channels and limits on coupling strengths. *Applied Optics*, 39(11):1681, April 2000.

[11] Deniz Mengu, Yi Luo, Yair Rivenson, and Aydogan Ozcan. Analysis of Diffractive Optical Neural Networks and Their Integration With Electronic Neural Networks. *IEEE Journal of Selected Topics in Quantum Electronics*, 26(1):1–14, January 2020.

[12] Yichen Sun, Mingli Dong, Mingxin Yu, Xiaolin Liu, and Lianqing Zhu. Review of diffractive deep neural networks. *Journal of the Optical Society of America B*, 40(11):2951, November 2023.

[13] Tiankuang Zhou, Xing Lin, Jiamin Wu, Yitong Chen, Hao Xie, Yipeng Li, Jingtao Fan, Huaqiang Wu, Lu Fang, and Qionghai Dai. Large-scale neuromorphic optoelectronic computing with a reconfigurable diffractive processing unit. *Nature Photonics*, 15(5):367–373, May 2021.

[14] Yi Luo, Yifan Zhao, Jingxi Li, Ege Çetintaş, Yair Rivenson, Mona Jarrahi, and Aydogan Ozcan. Computational imaging without a computer: seeing through random diffusers at the speed of light. *eLight*, 2(1):4, December 2022.

[15] Yuhang Li, Jingxi Li, Yifan Zhao, Tianyi Gan, Jingtian Hu, Mona Jarrahi, and Aydogan Ozcan. Universal Polarization Transformations: Spatial Programming of Polarization Scattering Matrices Using a Deep Learning-Designed Diffractive Polarization Transformer. *Advanced Materials*, 35(51):2303395, December 2023.

[16] Yueqiang Hu, Hanbin Chi, and Huigao Duan. Metaoptics merging computational optics and optical computing toward intelligent visual perception. *Science Advances*, 12(1):eaea8941, January 2026.

[17] Yetao Shu, Laixi Sun, Yuhai Li, Tong Fu, Yubin Zhang, Rong Wu, Wenhua Gu, Li Li, Fang Wang, Xiaodong Yuan, and Junwei Zhang. All-optical object classification using an edge-detecting spin-differential diffractive network. *Optics Express*, 34(3):3494, February 2026.

[18] David A. B. Miller. Why Optics Needs Thickness. *Science*, 379(6627):41–45, 2023.

[19] Yandong Li and Francesco Monticone. The spatial complexity of optical computing: toward space-efficient design. *Nature Communications*, 16(1):8588, September 2025.

[20] Heedong Goh and Andrea Alu. Nonlocal Scatterer for Compact Wave-Based Analog Computing. *Physical Review Letters*, 128(7):073201, 2022.

[21] Xuhao Luo, Qiang Song, Weiwei Cai, Lei Chen, Enbo Yang, Hao Wang, Zhipei Sun, Yueqiang Hu, Joel K W Yang, and Huigao Duan. Volumetric Optical Scattering Neural Networks. *arXiv preprint arXiv:2605.13177*, 2026.

[22] Hao Wang, Jianqi Hu, Andrea Morandi, Alfonso Nardi, Fei Xia, Xuanchen Li, Romolo Savo, Qiang Liu, Rachel Grange, and Sylvain Gigan. Large-scale photonic computing with nonlinear disordered media. *Nature Computational Science*, 4(6):429–439, 2024.

[23] Clara C Wanjura and Florian Marquardt. Fully nonlinear neuromorphic computing with linear wave scattering. *Nature Physics*, 20(9):1434–1440, 2024.

[24] Tatsuhiro Onodera, Martin M Stein, Benjamin A Ash, Mandar M Sohoni, Melissa Bosch, Ryotatsu Yanagimoto, Marc Jankowski, Timothy P McKenna, Tianyu Wang, Gennady Shvets, et al. Arbitrary control over multimode wave propagation for machine learning. *Nature Physics*, 22(1):164–171, 2026.

[25] Gustav Mie. Beiträge zur Optik trüber Medien, speziell kolloidaler Metallösungen. *Annalen der Physik*, 330(3):377–445, 1908.

[26] Wei Liu and Yuri S Kivshar. Multipolar interference effects in nanophotonics. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, 375(2090):20160317, 2017.

[27] Michael I. Mishchenko, Larry D. Travis, and Daniel W. Mackowski. T-matrix computations of light scattering by nonspherical particles: a review. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 55(5):535–575, 1996.

[28] Sergey Gladishev, Kristina Frizyuk, and Andrey Bogdanov. Symmetry analysis and multipole classification of eigenmodes in electromagnetic resonators for engineering their optical properties. *Physical Review B*, 102(7):75103, 2020.

[29] Maria Poleva, Kristina Frizyuk, Kseniia Baryshnikova, Andrey Evlyukhin, Mihail Petrov, and Andrey Bogdanov. Multipolar theory of bianisotropic response of meta-atoms. *Physical Review B*, 107(4):L041304, 2023.

[30] W. J. Wiscombe. Improved Mie scattering algorithms. *Applied Optics*, 19(9):1505, May 1980.

[31] Yuheng Zhang, Jiachen Ding, Ping Yang, and R. Lee Panetta. Vector spherical wave function truncation in the invariant imbedding T-matrix method. *Optics Express*, 30(17):30020, August 2022.

[32] David AB Miller, Zeyu Kuang, and Owen D Miller. Tunnelling escape of waves. *Nature Photonics*, 19(3):284–290, 2025.

[33] Adam Overvig and Francesco Monticone. Nonlocal metasurfaces: universal modal maps governed by a nonlocal generalized Snell's law. *Nanophotonics*, 14(23):3851–3860, 2025.

[34] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[35] Michael I Mishchenko, Larry D Travis, and Andrew A Lacis. *Scattering, absorption, and emission of light by small particles*. Cambridge university press, 2002.

[36] Sean Molesky, Pengning Chao, Jewel Mohajan, Wesley Reinhart, Heng Chi, and Alejandro W Rodriguez. T-operator limits on optical communication: metaoptics, computation, and input-output transformations. *Physical Review Research*, 4(1):013020, 2022.

[37] Heedong Goh, Alex Krasnok, and Andrea Alù. Nonreciprocal Scattering and Unidirectional Cloaking in Nonlinear Nanoantennas. *Nanophotonics*, 13(18):3347–3353, 2024.

[38] Dimitrios L. Sounas, Christophe Caloz, and Andrea Alù. Giant Non-reciprocity at the Sub-wavelength Scale Using Angular Momentum-Biased Metamaterials. *Nature Communications*, 4:2407, 2013.

[39] Christophe Caloz, Andrea Alù, Sergei Tretyakov, Dimitrios Sounas, Karim Achouri, and Zoé-Lise Deck-Léger. Electromagnetic Nonreciprocity. *Physical Review Applied*, 10(4):047001, 2018.

[40] Louis Béthune, Thibaut Boissin, Mathieu Serrurier, Franck Mamalet, Corentin Friedrich, and Alberto Gonzalez Sanz. Pay attention to your loss: understanding misconceptions about lipschitz neural networks. *Advances in Neural Information Processing Systems*, 35:20077–20091, 2022.

[41] Atish Agarwala, Jeffrey Pennington, Yann Dauphin, and Sam Schoenholz. Temperature check: theory and practice for training models with softmax-cross-entropy losses. *arXiv preprint arXiv:2010.07344*, 2020.

[42] Alessio Amaolo, Pengning Chao, Benjamin Strekha, Stefan Clarke, Jewel Mohajan, Francis J Chen, Sean Molesky, and Alejandro W Rodriguez. Maximum Shannon capacity of photonic structures. *npj Nanophotonics*, 3(1):14, 2026.

[43] B. R. Johnson. Invariant imbedding T matrix approach to electromagnetic scattering. *Applied Optics*, 27(23):4861, December 1988.

[44] Bingqiang Sun, George W. Kattawar, Lei Bi, Michael Kahnert, and Ping Yang. *Invariant imbedding T-matrix method for light scattering by nonspherical and inhomogeneous particles*. Elsevier, Amsterdam, 2020.

[45] Niyazi Ulas Dinc, Amirhossein Saba, Jorge Madrid-Wolff, Carlo Gigli, Antoine Boniface, Christophe Moser, and Demetri Psaltis. From 3d to 2d and back again. *Nanophotonics*, 12(5):777–793, 2023.

[46] Kanhaya Sharma, Adrià Grabulosa, Erik Jung, and Daniel Brunner. 3d photonic integration leveraging hybrid-confinement circuits. *arXiv preprint arXiv:2605.24990*, 2026.

[47] Xuhao Luo, Qiang Song, Weiwei Cai, Lei Chen, Enbo Yang, Hao Wang, Zhipei Sun, Yueqiang Hu, Joel KW Yang, and Huigao Duan. Volumetric optical scattering neural networks. *arXiv preprint arXiv:2605.13177*, 2026.

[48] Nigar Asadova, Jan David Fischbach, Renaud Vallée, Oliver Kuster, Yannick Augenstein, Dmytro Vovchuk, Anton Kharchevskii, Pavel Ginzburg, and Carsten Rockstuhl. Gradient-based optimization of scatterer arrangements based on the T-matrix method. *APL Photonics*, 11(5), 2026.

[49] Nigar Asadova, Kaoutar Boussaoud, Jörg Meyer, Frank Tristram, and Carsten Rockstuhl. T-matrix database to promote information-driven research in nanophotonics. *Optical Materials Express*, 16(6):1534–1550, 2026.

[50] John Peurifoy, Yichen Shen, Li Jing, Yi Yang, Fidel Cano-Renteria, Brendan G DeLacy, John D Joannopoulos, Max Tegmark, and Marin Soljačić. Nanophotonic particle simulation and inverse design using artificial neural networks. *Science advances*, 4(6):eaar4206, 2018.

[51] Zhanyuan Zhang, Wanchun Zhang, Jiayi Yang, Tuqiang Pan, Yi Xu, and Yuwen Qin. Spherical-harmonic-based inverse design of arbitrary Mie scatterers mediated by deep learning. *Physical Review Applied*, 24(6):064002, 2025.

[52] Sadeq Bahmani, Andrey B Evlyukhin, Emadeldeen Hassan, and Antonio Calà Lesina. Topology optimization of optical nanoantennas with desired multipoles. *Optics express*, 33(9):19418–19441, 2025.

[53] Joel Sved, Shijie Song, Liwei Li, George Li, Debin Meng, and Xiaoke Yi. Inverse-designed nanophotonic neural network accelerators for ultra-compact optical computing. *Nature Communications*, 17(1):1059, 2026.

[54] Lukas Novotny and Bert Hecht. *Principles of Nano-Optics*. Cambridge University Press, September 2012.

[55] Alexey A. Shcherbakov. Calculation of the electromagnetic scattering by non-spherical particles based on the volume integral equation in the spherical wave function basis. *J. Quant. Spectrosc. Radiat. Transfer*, 231:102–114, July 2019.

[56] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations (ICLR 2015)*, 2015.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- MieopticalcomputingSupplsubmit.pdf