



HAL
open science

Gradient-based optimization of core-shell particles with discrete materials for directional scattering

Dalin Soun, Antoine Azéma, Lucien Roach, Glenna L Drisko, Peter R Wiecha

► **To cite this version:**

Dalin Soun, Antoine Azéma, Lucien Roach, Glenna L Drisko, Peter R Wiecha. Gradient-based optimization of core-shell particles with discrete materials for directional scattering. *Optics Express*, 2025, 33 (12), pp.25945-25958. <10.1364/OE.561657>. <hal-05236796>

HAL Id: hal-05236796

<https://hal.science/hal-05236796v1>

Submitted on 3 Sep 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.



L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License



Gradient-based optimization of core-shell particles with discrete materials for directional scattering

DALIN SOUN,¹ ANTOINE AZÉMA,^{1,2} LUCIEN ROACH,³  GLENNA L. DRISKO,³ AND PETER R. WIECHA^{1,*} 

¹LAAS-CNRS, Université de Toulouse, 31000 Toulouse, France

²CEMES-CNRS, Université de Toulouse, 31000 Toulouse, France

³Laboratoire de Chimie, CNRS, ENS de Lyon, 69364 Lyon, France

*pwiecha@laas.fr

Abstract: Designing nanophotonic structures traditionally grapples with the complexities of discrete parameters, such as real materials, often resorting to costly global optimization methods. This paper introduces an approach that leverages generative deep learning to map discrete parameter sets into a continuous latent space, enabling direct gradient-based optimization. For scenarios with non-differentiable physics evaluation functions, a neural network is employed as a differentiable surrogate model. The efficacy of this methodology is demonstrated by optimizing the directional scattering properties of core-shell nanoparticles composed of a selection of realistic materials. We derive suggestions for core-shell geometries with strong forward scattering and minimized backscattering. Our findings reveal significant improvements in computational efficiency and performance when compared to global optimization techniques. Beyond nanophotonics design problems, this framework holds promise for broad applications across all types of inverse problems constrained by discrete variables.

Published by Optica Publishing Group under the terms of the [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

1. Introduction

From the transmission of radio signals to the ongoing exploration of quantum computing – the manipulation of electromagnetic waves has been a cornerstone of advancing technology ever since James Clerk Maxwell's groundbreaking work on electromagnetic fields [1]. In recent years, the field of nano-photonics has emerged as an exciting arena where the interplay of light and matter at the nanoscale can be controlled to an unprecedented extent [2,3]. This control has powerful implications for applications ranging from information processing over sensing to subwavelength imaging. However, the design of nanophotonic structures that can precisely dictate light behavior remains a challenging inverse problem, constrained by the limitations of available materials and intricate fabrication processes [4–6].

Traditional methods to address these design challenges have relied on direct search and global optimization algorithms. These are computationally expensive and often converge very slowly due to the complex, high-dimensional search spaces involved [7–9]. Recent advancements in deep learning, particularly generative models, suggest a way forward by extracting complex correlations from large banks of data, and can also be trained to solve discrete inverse problems [10–15]. Such data driven methods however often suffer from insufficient dataset sizes or problematic data quality [16].

Gradient-based optimization, either with adjoint formulations like in topology optimization, based on automatic differentiation, or in combination with deep learning surrogate models, has often been found to be the superior approach [17–22]. In combination with data-based

physics predictors however, this requires very careful constraint handling, due to a high risk of erroneous extrapolation [15,23]. Furthermore, these techniques generally do not work with discrete parametrizations, such as fixed sets of real materials, since gradients are then not defined. Complicated hybrid optimization procedures [24,25], customized interpolation schemes [26,27], or costly global optimization are typical fallback approaches [4,6]. Recently, generative deep learning has been used to re-parametrize image representations of structures to enable gradient-based optimization [28,29], and it has been proposed to use the expressiveness of untrained neural networks by passing the original parametrization of a problem through a neural network and optimizing this network instead of the parameters [30,31].

In this paper, we propose an approach in a similar spirit, merging the strengths of automatic differentiation with deep generative models to address the design optimization problem of nanophotonic structures with discrete parameters. We use a Wasserstein Generative Adversarial Network with Gradient Penalty (WGAN-GP) to map a discrete parametrization into a continuous latent space, enabling gradient-based techniques to probe the design space efficiently. For situations where direct gradient evaluation is not feasible, a neural network can be used as a differentiable surrogate model [23], ensuring that the evaluation loop remains smoothly integrable [15].

Core-shell nano-spheres are a promising class of photonic structures, holding potential for applications in low-loss metasurfaces, sensing, quantum emitters or medicine [32–37]. Core-shell particles can be synthesized through bottom-up approaches, and are promising for scalable and cost-effective mass production [38,39]. It is therefore compelling to design such structures for applications in nanophotonics. However, designing core-shell particles from a discrete selection of real materials, poses a challenge due to the lack of gradients for the discrete material parameters, which is why former approaches used simplifications like idealized permittivity [23,40], less accurate direct deep learning methods [41], or complex optimization schemes to treat the discrete part separately [25].

We demonstrate our approach of gradient-based discrete parameter optimization with the design of core-shell nano-spheres for optimal and strong forward light scattering.

In this manuscript, we first introduce the details of our approach and its implementation for the optimization of core-shell nanoparticles. We describe the dataset generation, preprocessing techniques, and the architecture of the models involved. Subsequently, a statistical evaluation demonstrates the superior performance against traditional global optimization. Finally, we apply the method on the specific task of maximizing forward scattering whilst minimizing backscattering.

2. Gradient optimization with discrete parameters

2.1. Problem: design of core-shell nanospheres with discrete materials

Throughout this work, we illustrate our approach for gradient-based optimization of discrete problems by an example in nano-photonics design optimization. Our goal is to design spherical core-shell particles for light scattering. Specifically, we want to find particle geometries that maximize forward scattering while backscattering is minimized, which occurs at the so-called Kerker condition. This is highly interesting for applications in bottom-up fabricated, low-loss metasurfaces (so-called “Huygens” metasurfaces) [42–45]. Our geometry is described by four parameters: core material, shell material, core radius, and shell thickness. While the sizes of core and shell are continuous parameters, the respective materials are chosen from a list of available materials and represent discrete parameters.

The optical scattering of spherical particles can be analytically solved by Mie theory [46]. The forward and backward scattering cross sections can thereby be expressed as function of the n th

order Mie coefficients a_n and b_n : [47]

$$\sigma_{\text{fwd}} = \frac{\pi}{k^2} \left| \sum_{n=1}^{\infty} (2n+1)(-1)^n (a_n - b_n) \right|^2 \quad (1)$$

and:

$$\sigma_{\text{back}} = \frac{\pi}{k^2} \left| \sum_{n=1}^{\infty} (2n+1)(a_n + b_n) \right|^2 \quad (2)$$

The Mie coefficients are obtained by matching the fields at the spherical interfaces between the different materials [48]. Finally, the scattering efficiencies Q_{fwd} and Q_{back} used here, are simply the cross sections divided by the geometric cross section $\sigma_{\text{geo}} = \pi r^2$ of the sphere with outer radius r . To calculate the directional scattering, we use a homemade, freely available Python / NumPy implementation “pymiecs” [49].

The inverse problem however, is an ill-posed problem and generally cannot be solved directly [11,50]. We want to use gradient-based optimization to find solutions to the design problem. In addition to the discrete parametrization not being differentiable, also the Mie toolkit is not differentiable. Deriving symbolic derivatives of core-shell Mie theory is very tedious. Implementation in an automatic differentiation (AD) framework furthermore requires AD versions of spherical Bessel functions and their derivatives, which are so far not available in the most popular AD toolkits so far. Finally, in order to ensure proper gradient flow, non-discrete parameter passing to the Mie solver would need to be implemented, e.g. using interpolated mixing of the actual material properties. To avoid these challenges, we choose to replace the Mie solver by a feed-forward neural network, which is described in the following.

2.2. Data generation and preprocessing

We build a dataset using core radii between 1 nm and 100 nm and shell thickness between 1 nm and 100 nm. The available materials are illustrated in Fig. 1, and comprise the plasmonic metals gold (Au) and silver (Ag), as well as the dielectric materials silicon (Si), silica (SiO₂), silicon nitride (Si₃N₄), zirconium dioxide (ZrO₂), and titanium dioxide (TiO₂). In addition to the radii and materials of core and shell, the scattering depends on the environment and on the illumination wavelength. For the environment, we use vacuum ($n_{\text{env}} = 1$). The illumination is at visible wavelengths from $\lambda_0 = 400$ nm to $\lambda_0 = 800$ nm, sampled with 64 equally spaced points. For every configuration, we calculate the scattering efficiencies of core-shell nanospheres in forward Q_{fwd} and backward direction Q_{back} via Mie theory. In this way, we generate a total of 122500 data points.

It is common in deep learning to preprocess and normalize the data before training the model, to ensure it is in a suitable format, and in a numerical range similar to a unit normal distribution [51]. The categorical features (the core and shell materials) are represented in one-hot encoding (vectors of N zeros, with the i -th element, corresponding to the represented category, being set to 1). The numerical features (the core and shell radii), are normalized using a min-max standard scaler [52], which re-scales the data into the range $[-1, 1]$. The numerical target variables (the forward and backward scattering efficiencies Q_{fwd} and Q_{back}) are preprocessed with a logarithm transformation (“log1p”), to reduce the strong skewness of the raw data distribution, followed by applying a min-max scaler. The inverse transformation is performed on the predictions of the feed-forward Mie neural network. Note, that we verified each step of the preprocessing procedure to improve training convergence. Specifically we found that “log1p” preprocessing does not only increase the global performance of the Mie predictor model cutting the residual average error roughly by half, it improves in particular weak scattering cases, in which cases relative errors improve significantly.

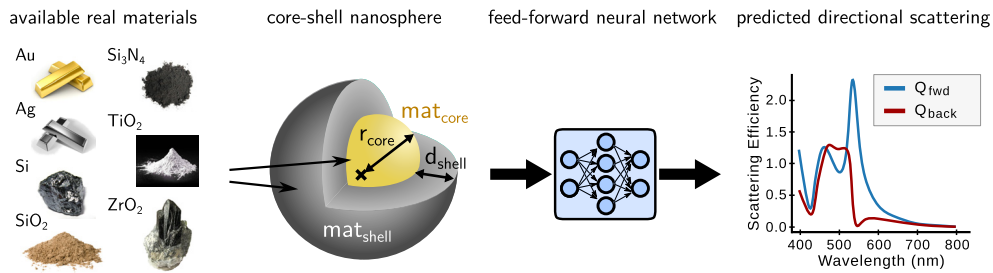


Fig. 1. Illustration of the example problem of core shell nanosphere design with a discrete set of realistic materials. A feed-forward neural network model predicts the nanosphere's directional scattering response in form of the forward and backward scattering efficiencies Q_{fwd} and Q_{back} across the visible spectrum ($\lambda_0 = 400\text{--}800\text{ nm}$). The inputs to the model are core material, shell material, core radius, and shell thickness.

2.3. Scattering modeling (the direct problem)

As mentioned above, gradient-based optimization requires all parts in the computational chain of scattering evaluation to be differentiable. To this end, we replace the Mie toolkit in our demonstration by a learned feed-forward model.

Architecture of the scattering prediction model. Our neural networks are implemented with Tensorflow [53]. The neural network takes the core-shell geometry as input and returns predictions for the scattering response (Q_{fwd}) and (Q_{back}). This is shown schematically in Fig. 1. In our implementation we use a 1D residual network (ResNet), consisting of two parts: a first dense part takes as input the geometry and reshapes it into a 2D array for compatibility with 1D convolutional layers (Conv1D). The second part consists of multiple 1D convolutional residual blocks and upsampling stages. Each residual block includes two Conv1D layers with batch normalization, leaky rectifier linear unit activations (leaky ReLU), and a skip connection. Sequences of four residual blocks are followed by a factor 2 upsampling layer, to gradually increase the data dimension towards the final output of 64 wavelengths per spectrum. A final Conv1D layer consists of two output channels (one channel for Q_{fwd} , the second for Q_{back}), followed by a linear activation function. A detailed illustration of the scattering prediction model is shown in the Supplement 1 figure S2.

Note that any differentiable model could be used instead of the ResNet. We found for example, that a simple multilayer perceptron (MLP), which evaluates a single wavelength instead of the full spectrum, also performs well, yet with errors approximately twice as large. The main advantage of an MLP would be, that it is not limited to the specific wavelengths defined by the dataset.

Training. We use 108,000 samples of our dataset for training and 12,000 samples for validation and 2,500 for testing. The model is trained using an MSE loss with the Adam optimizer [54] over 255 epochs, ended through early stopping after convergence (see Supplement 1 figure S3). We use a hybrid approach [55] of increasing batch size up to 256 and a simultaneous learning rate decay schedule, which halves the learning rate every 15 epochs. The model converges well and training and validation loss curves show no sign of overfitting (see Sup.Info. Fig. S3).

Performance. The model performance is finally evaluated using a relative mean absolute error (rMAE) metric on the test set predictions. We find rMAE percentages as low as 0.80% and 0.95% for, respectively, Q_{fwd} and Q_{back} . Histograms of the error distributions are shown in Fig. 2.

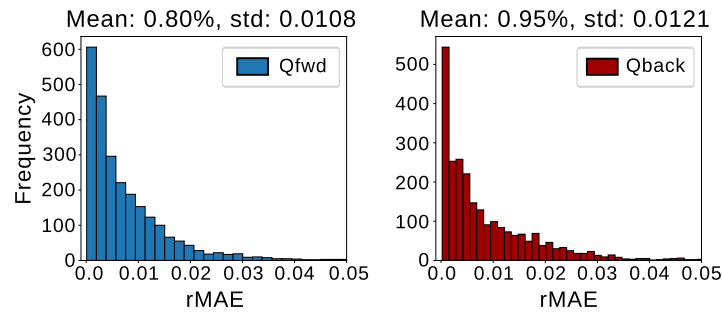


Fig. 2. Relative Mean Absolute Error (rMAE) distribution for the Q_{fwd} and Q_{back} predictions from the feed-forward neural network on the test set. Note, that the histogram bars have different sizes because of different data ranges for forward and backward scattering.

2.4. Re-parametrization of discrete materials

Our goal is gradient-based optimization of the core-shell particles. Therefore, in addition to the differentiable surrogate for the scattering evaluation, we also need to convert the discrete material parameters of the geometry into a continuous representation.

To do so, we employ a WGAN-GP [56,57], which is an adaptation of the original GAN [58] with greatly improved training stability. GAN models such as the WGAN-GP learn a meaningful, smooth, and compact mapping of the original training data into a convex latent representation that follows a predefined distribution (we chose a standard normal distribution) [59]. The WGAN-GP has two main components. First the generator network (red in Fig. 3), which generates fake samples (here: core-shell geometries) from a latent vector input (z). And second, the discriminator (green in Fig. 3). Its task is to distinguish whether data is fake (from the generator) or real (from the dataset), it is trained in alternance with the generator, for which it acts as a dynamically learned loss function. A detailed illustration of the WGAN-GP architecture we use, is shown in the [Supplement 1](#) figure S2.

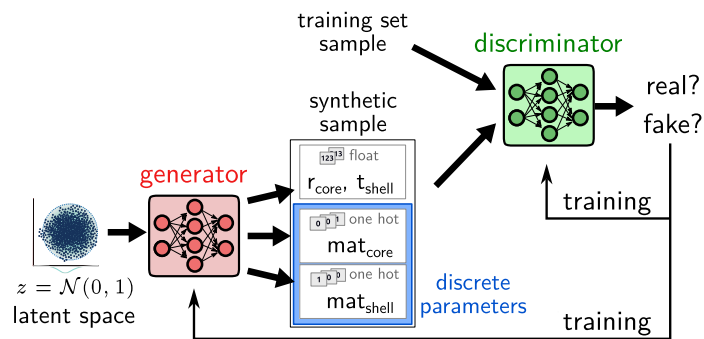


Fig. 3. Global architecture of WGAN-GP Model. The generator network (red) takes inputs from a Gaussian distribution and produces two numerical values (core radius, shell thickness), and two categorical outputs (core material, shell material), the latter are returned as one-hot encoded vectors. The discriminator network (green) is trained on distinguishing synthetic from real samples, and it is used as an optimized loss function for the generator training.

Our WGAN-GP model is trained exclusively on the design parameters (sizes and materials of the core and shell), which are preprocessed in the same way as for the scattering prediction model. The architecture is based on two MLPs for generator and discriminator, in which each

fully connected layer is followed by a leaky ReLU activation. The generator uses an input latent dimension of 128, and has three distinct output layers: The first output layer consists of two neurons and “tanh” activation, yielding the numerical values for the core and shell radii. The second and third output layer use “softmax” activation, each with seven neurons (as here we consider seven distinct materials), yielding the materials for the core and the shell. This configuration guarantees that the WGAN-GP cannot generate unphysical geometries. It enforces the synthetic samples to lie within the data normalization range, hence within the user-defined parameters boundaries. Note that we deliberately use a large latent dimension, because information compression is not our goal. We require the smoothness and compactness properties of the latent space, together with the best possible generation quality.

As training hyperparameters we found that a batch size of 256, and two different learning rates, 5×10^{-5} for the generator, and 5×10^{-4} for the discriminator optimizer (both Adam) give best results. The other training parameters follow the original WGAN-GP paper [56]. We train the WGAN-GP over 24 epochs.

To assess the model fidelity, after the training, we generate 10,000 samples and compare the distributions of the synthetic data with the uniform distribution of the training data, as shown in figure S4. The distributions of the synthetic data’s material parameters are almost uniform, meaning that the WGAN-GP model generates synthetic samples with the same statistics as the categorical features in the training set. The synthetic data distributions of the two size parameters are visually a little distorted, compared to the training set distributions. Yet, the statistics are still very similar and, most importantly, the entire training data parameter range is covered by the synthetic data. Therefore, an optimizer algorithm should be able to reach the full parameter range.

2.5. Gradient-enabled design optimization

The actual design process consists in optimizing the core-shell nanostructures using gradient-descent methods to find a geometry that matches as closely as possible a target optical scattering response. To this end, we define a fitness function that quantifies the deviation from the design target scattering response. This fitness function is either the mean square error with respect to a specific target scattering response, or simply the numerical value to be minimized or maximized. For maximization we use the negative of the value, as typical algorithms perform minimization by default.

We connect the output of the pre-trained generator from the WGAN-GP to the input of the pre-trained Mie scattering prediction model. The output of the scattering predictor is used for calculation of the design fitness. The gradient of the design’s fitness L with respect to the generator’s latent vector z is calculated using automatic differentiation. z is updated accordingly via a gradient-based optimizer (here we simply use again Adam). The full workflow is depicted in Fig. 4. We use a learning rate of 0.01, which we found to be a robust and fast choice in a series of convergence test (see Supplement 1 figure S7). Note that convergence speed could certainly be further optimized by employing a learning rate schedule and tuning the optimizer hyperparameters. The latent representation z is iteratively adjusted until convergence of the fitness.

A final problem needs to be resolved: gradient-based optimization converges towards the closest *local* minimum. To nevertheless find solutions close to the global optimum, we initiate the optimization process with a large number of different, random latent vectors, which we then all optimize concurrently. Thanks to the high degree of optimization of modern deep learning frameworks, this can be done with essentially no technical effort and batch optimization is extremely efficient. On our system with an NVIDIA GeForce RTX 4090 GPU, optimization of 500 latent vectors in parallel for 100 iterations takes around 9 s. For the here shown core-shell problem, 100 parallel optimizations were generally sufficient, we use 500 as it is not time-critical.

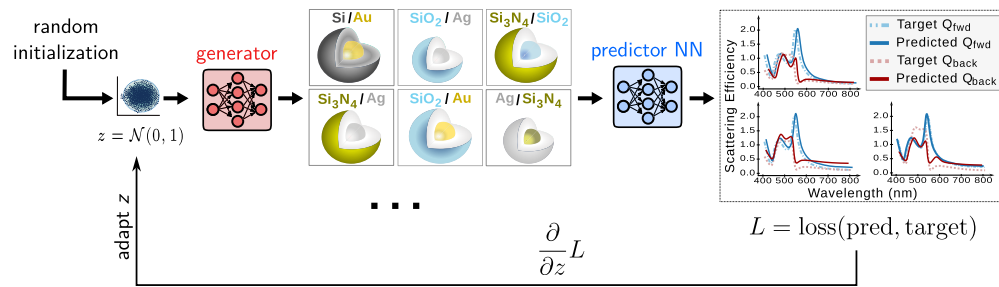


Fig. 4. Our framework uses the pre-trained generator (red) and the feed-forward neural network (blue) for gradient calculation. The generator produces synthetic core-shell geometries from a set of randomly initialized latent vectors z . These geometries are then passed to the feed-forward model to predict their Q_{fwd} and Q_{back} spectra. A loss function L is defined to quantify the difference between predicted scattering and the target (this may for example be the mean square error between a full target reference spectrum and the predicted spectrum). The gradients of the loss L with respect to the latent vectors z are calculated via automatic differentiation and are used to adjust the latent vectors to better match the design target.

The required amount of initial guesses is typically problem specific. For problems with large numbers of local optima, gradient-free global pre-optimization with a gradient-based second step may be an efficient strategy [15].

3. Results

3.1. Examples and statistics from the test-set

We first test the performance of our optimization technique on the task of finding geometries from the test dataset based on the known optical spectra. The fitness function is the mean square error between the target spectrum and the predicted spectrum. Figure 5 shows a selection of random examples, where the spectrum of the optimized geometry is recalculated with exact Mie theory. Remarkably, and despite the discrete material parameters, the gradient-based optimization reliably finds geometries that reproduce almost perfectly the desired solution. In Fig. 5(d) we also see that, as expected, ambiguous solutions (e.g. homogeneous particles with identical core- and shell-material) are not identically reproduced. The excellent performance of the optimization scheme is confirmed statistically on the entire test dataset (Fig. 6(a)). We find a deviation of only around one percent between optimization results and the existing solutions.

To put this performance in context with conventional methods, we compare our gradient-based method to global optimization. For this comparison, we use the evolutionary optimization algorithm “Differential Evolution” (DE) from the “NeverGrad” toolbox to optimize the entire test dataset [60,61]. DE is known to be a robust global algorithm that performs well on a large variety of problems [6].

100 iterations do not suffice to get reasonable results with global optimization. Therefore we let the global optimization run for 250 iterations, on a population of 100 individuals. The average error with respect to the target is still around 10 times higher than with gradient-based optimization (mean absolute errors of around 1.12% for gradient-based vs. around 11.88% for global optimization). Please note also, that the latter uses an error-prone forward neural network for scattering predictions, while the global optimization uses exact Mie theory for fitness evaluation. We finally note, that global optimization with our vectorized Mie solver is similarly fast compared to the forward neural network framework. However, the feed-forward neural network would be expected to have a drastic runtime advantage for computationally more difficult

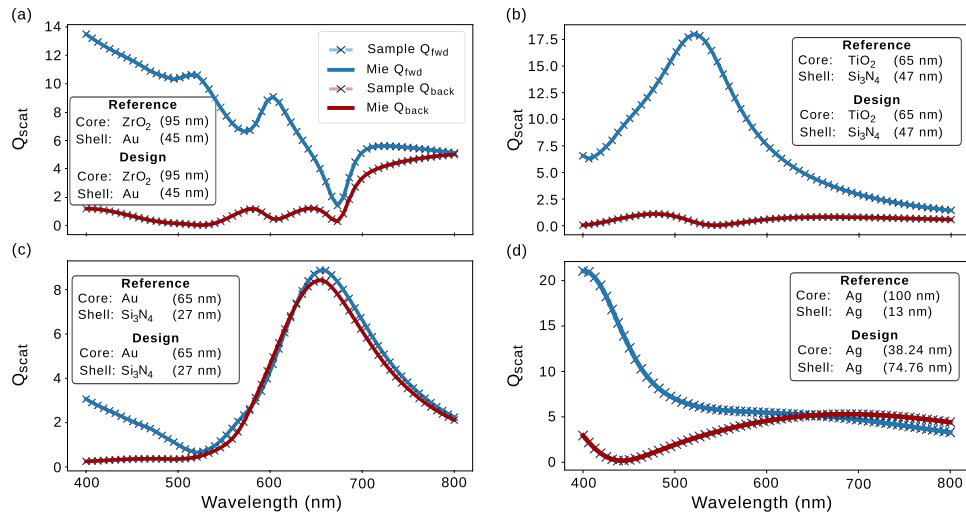


Fig. 5. Random optimized scattering spectra vs the design targets. Each subplot (a-d) presents a different sample, showing both Q_{fwd} and Q_{back} for the design target (solid lines) and of the optimized geometries (cross markers). The indicated numbers are core radius and shell thickness.

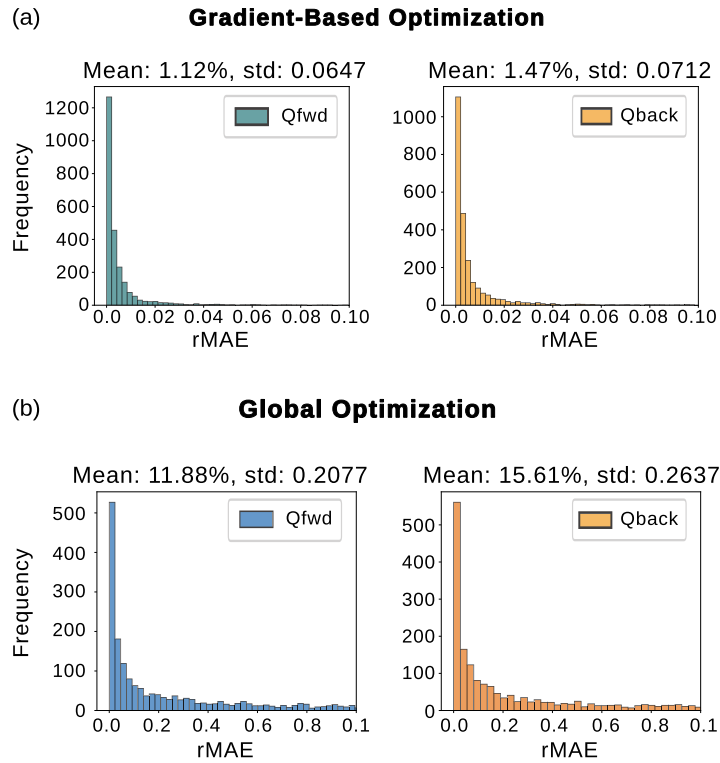


Fig. 6. Comparison of relative mean absolute errors (rMAE) for Q_{fwd} (left) and Q_{back} (right). (a) Gradient-based optimization. (b) Global optimization.

problems (such as free-form geometries), in addition to the strongly improved convergence rate [62].

3.2. Maximize forward scattering and minimize backscattering

Now we want to use gradient-based design optimization to find core-shell particles that achieve maximum forward scattering while minimizing backward scattering at specific wavelengths. This problem is more complicated than the above full spectrum matching, as it is effectively a multi-objective optimization. In earlier work we found that the quotient of forward over backscattering tends to mostly minimize backscattering [63]. Therefore, to solve for strong forward scattering while minimizing backscattering, we define the following directional scattering fitness function:

$$L_{\text{dir}} = -Q_{\text{fwd}} + wQ_{\text{back}}, \quad (3)$$

which contains two terms that correspond to maximize forward and minimize backward scattering. To balance the two objectives, we add a weight w to the backscattering term. We performed various series of random optimizations for different weights w . These are described in the supporting information (see [Supplement 1](#) figure S5). It turns out that a weight of $w = 1$ is adequate for the goal of maximizing forward scattering with as little as possible backscattering, while not being too restrictive on the overall scattering efficiency. Note, that for each multi-objective optimization problem the fitness function should be chosen individually and, if applicable, its weight parameters should be optimized. We want to mention also, that there exist different ways to approach such multi-objective optimization (MOO) problems. The most popular alternative to a combined fitness function such as Eq. (3) would be Pareto optimization, that searches all solutions that cannot be further improved without worsening at least one of the objectives [64]. While Pareto-optimization historically uses gradient-free heuristics [65], currently research interests are emerging to develop gradient based MOO algorithms for machine learning applications [66]. Our work may be of interest also in this context.

Having set up the fitness function, we now use the possibility to perform very efficiently numerous optimizations with different initialization to globally explore the locally optimum regions in the design space [23], here solutions that maximize Q_{fwd} , while minimizing Q_{back} . Thanks to the WGAN-GP reparameterization, this is now possible also with discrete input parameters. We perform 5000 optimizations with random initializations, using a backscattering weight of $w = 1$ and working wavelength at $\lambda_0 = 705$ nm. During the optimization, we fix silicon as the core material, simply by replacing the WGAN-GP's output core material. After the optimization, we remove all non-silicon core samples, after this 686 solutions remain (which is approximately 1/7, matching the 7 different available materials). Alternatively, a new WGAN-GP could have been trained for a fixed core material.

We indicate in Fig. 7(a) the fitness of the results by the color of the markers in a scatter plot. The color bar indicates the loss values after the optimization process, with yellow corresponding to the best (lowest) loss. The x -axis indicates the shell material, and the y -axis corresponds to the shell thickness. The core radius is also subject to optimization but is not shown in this figure. A detailed figure including core sizes is given in the [Supplement 1](#) figure S6.

Every dot corresponds to a local minimum in the optimization landscape. Therefore, "gap" regions with no solution indicate locally non-optimum parameter regions. This analysis reveals that two different types of locally optimal forward-scattering silicon core-shell nanospheres exist: (1) Structures with lower-index dielectric shell (SiO_2 , Si_3N_4 , ZrO_2 , TiO_2) ideally have a large shell thickness while (2) plasmonic shells (Au or Ag) have an optimal thickness below 10 nm. In the case of silver, the shell should be even a bit thinner than for gold, presumably due to its better plasmonic properties and lower dissipation. The directional scattering spectra for the lowest loss solution of each shell material is shown in Fig. 7(b). These findings are in agreement with recent studies about such types of core-shell nanostructures [32,67,68]

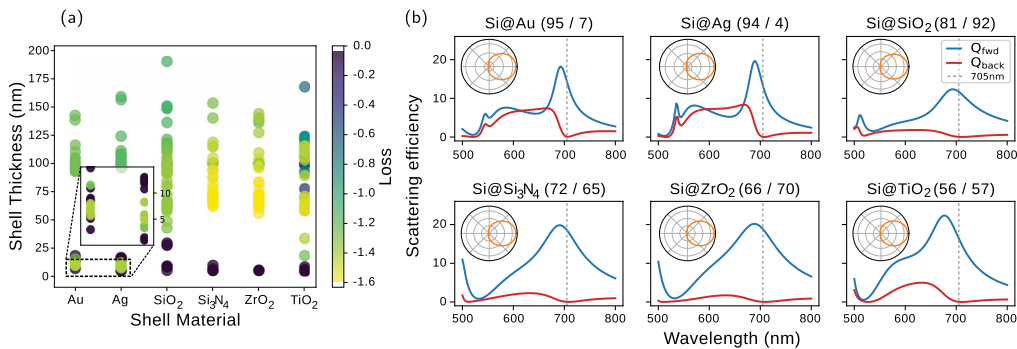


Fig. 7. (a) Analysis of locally optimal regions in the design space for silicon core material nano-spheres. Scatter plots of 686 independent directional optimizations with weight $w = 1$. The y-axis represents the shell thickness, the x-axis the shell material. The marker color corresponds to the optimization loss (lower is better). The inset shows a zoom on thin shell thicknesses for plasmonic shells. (b) Best design per shell material (fixed silicon core). The silicon core radius and shell thickness is indicated in parentheses above each subplot (in units of nm). Insets show the polar radiation patterns of the scattering at the design wavelength ($\lambda_0 = 705$ nm).

4. Discussion

The optimized core-shell geometries predicted by our approach were chosen because of their feasible realization. Crystalline silicon cores in this size range have already been produced [69–71]. The creation of thin plasmonic shells made from gold (or even thinner silver) on a silicon core has not yet been produced. There are currently two reports of Si@Au core-shells in the literature [72,73], neither of which produced shells of suitable thickness or continuity to replicate the results predicted here.

However, thin plasmonic shells have been achieved to varying degrees of success on silica (SiO₂) [34,67,74]. Silicon particles are typically coated by 2-3 nm native silica layer, hence the surface chemistry is identical to a silica particle, allowing all techniques used on silica to be transposed to silicon. A much thicker silica coating around a silicon core has already been achieved via thermal annealing in air [68] and through the co-decomposition of precursors [32]. Alternatively, it would be possible to deposit a silica shell around the silicon core using well-known sol-gel techniques [75,76]. Metal oxide cores have been coated with titania [77,78] and zirconia [79,80] shells using a variety of sol-gel protocols, rendering shells of comparable thickness to those studied here. Thus, we selected core-shell structures not only that theoretically are optimized in terms of their forward/backward scattering ratio, but hold promise in terms of realizability.

In the context of chemical feasibility, we can conclude from the above results that low-index dielectric materials seem the most interesting choice for fabrication of Huygens' sources with optimized scattering efficiency. Silicon nitride and zirconium oxide seem to be particularly interesting, as in addition to strong directional scattering at the Kerker wavelength, the directionality is also very broadband. The range of “optimality” over a wide span of shell thicknesses (Fig. 7(a)) indicates also a certain robustness against fabrication size polydispersity.

Alternatively, thin plasmonic shells made from gold or silver, seem to be capable of increasing the scattering efficiency at which the Kerker effect occurs. However, the experimental challenge of the fabrication of thin, yet crystalline and fully covering shells stills needs to be overcome [67].

Please note, that all the solutions shown in Fig. 7 are locally optimal with respect to L_{dir} . This means they dependent on the choice of the backscattering weight w in the optimization loss L_{dir} and therefore do not necessarily represent “physical” feasibility.

5. Conclusions

In conclusion, we propose a WGAN-GP based re-parametrization of discrete problem parameters, to use gradient-based optimization on (partially) discrete problems. To avoid local solutions, the (very fast) optimization can simply be run many times in parallel. Specifically we demonstrate the gradient-based design of directional optical scattering properties of core-shell nanospheres, specifically aiming to maximize forward scattering with as little backscattering as possible. In comparison with global optimization, the gradient-based approach converges significantly faster. With a comparable compute budget, the error of the final solutions compared to the target, are in the order of 10% for global optimization, while our gradient-based approach converges down to around 1% residual error. Finally, we demonstrated by a practical example, how discrete parameter landscapes can be systematically analyzed for locally optimal solutions. We found that silicon core spheres with dielectric shells seem to be very interesting candidates for broadband, strong directional scattering. As a perspective, we foresee that gradient-based optimization will be very interesting for the design and study of nanophotonic multi-material structures and devices. The method is of course not limited to nanophotonics, it could be applied to any other discrete problem, possible examples may be synthesis protocol optimization in chemistry or process automation with binary triggers as system parameters.

Funding. Agence Nationale de la Recherche (ANR-22-CE24-0002, ANR-23-CE09-0011); CALMIP (p20010); Institut Quantique Occitan (Q-META).

Acknowledgments. We thank Aurélien Cuche for fruitful discussions. This work was supported by the French Agence Nationale de la Recherche (ANR) under grants ANR-22-CE24-0002 (project NAINOS) and ANR-23-CE09-0011 (project AIM), and by the Toulouse HPC CALMIP (grant p20010). A.A. acknowledges funding by the Institute of Quantum Technology in Occitanie IQO (project Q-META).

Disclosures. The authors declare no conflicts of interest

Data availability. All codes required to reproduce the results are available at [81].

Supplemental document. See [Supplement 1](#) for supporting content.

References

1. J. C. Maxwell, “A Dynamical Theory of the Electromagnetic Field,” *Philos. Trans. R. Soc. London* **155**, 459–512 (1865).
2. L. Novotny and N. van Hulst, “Antennas for light,” *Nat. Photonics* **5**(2), 83–90 (2011).
3. A. I. Kuznetsov, A. E. Miroshnichenko, M. L. Brongersma, *et al.*, “Optically resonant dielectric nanostructures,” *Science* **354**(6314), 1 (2016).
4. M. M. R. Elsayy, S. Lanteri, R. Duvigneau, *et al.*, “Numerical Optimization Methods for Metasurfaces,” *Laser Photonics Rev.* **14**(10), 1900445 (2020).
5. W. Ma, Z. Liu, Z. A. Kudyshev, *et al.*, “Deep learning for the design of photonic structures,” *Nat. Photonics* **15**(2), 77–90 (2021).
6. P. Bennet, D. Langevin, C. Essoual, *et al.*, “Illustrated tutorial on global optimization in nanophotonics,” *J. Opt. Soc. Am. B* **41**(2), A126–A145 (2024).
7. B. Wu, S. Xue, Z. Zhang, *et al.*, “Invisible devices with natural materials designed by evolutionary optimization,” *Phys. Rev. E* **106**(5), 055312 (2022).
8. O. Teytaud, P. Bennet, and A. Moreau, “Discrete global optimization algorithms for the inverse design of silicon photonics devices,” *Photonics Nanostructures - Fundam. Appl.* **52**, 101072 (2022).
9. C. Majorel, A. Loucif, E. Marinov, *et al.*, “Bio-inspired flat optics for directional 3D light detection and ranging,” *npj Nanophotonics* **1**(1), 18 (2024).
10. D. Liu, Y. Tan, E. Khoram, *et al.*, “Training Deep Neural Networks for the Inverse Design of Nanophotonic Structures,” *ACS Photonics* **5**(4), 1365–1369 (2018).
11. P. R. Wiecha, A. Arbouet, C. Girard, *et al.*, “Deep learning in nano-photonics: Inverse design and beyond,” *Photonics Res.* **9**(5), B182–B200 (2021).
12. J. Jiang, M. Chen, and J. A. Fan, “Deep neural networks for the evaluation and design of photonic devices,” *Nat. Rev. Mater.* **6**(8), 679–700 (2020).

13. A. Estrada-Real, A. Khairah-Walieh, B. Urbaszek, *et al.*, “Inverse design with flexible design targets via deep learning: Tailoring of electric and magnetic multipole scattering from nano-spheres,” *Photonics Nanostructures - Fundam. Appl.* **52**, 101066 (2022).
14. P. Dai, K. Sun, X. Yan, *et al.*, “Inverse design of structural color: Finding multiple solutions via conditional generative adversarial networks,” *Nanophotonics* **11**(13), 3057–3069 (2022).
15. A. Khairah-Walieh, D. Langevin, P. Bennet, *et al.*, “A newcomer’s guide to deep learning for inverse design in nano-photonics,” *Nanophotonics* **12**(24), 4387–4414 (2023).
16. P. R. Wiecha, “Deep learning for nano-photonic materials – The solution to everything!?” *Curr. Opin. Solid State Mater. Sci.* **28**, 101129 (2024).
17. J. S. Jensen and O. Sigmund, “Topology optimization for nano-photonics,” *Laser Photonics Rev.* **5**(2), 308–321 (2011).
18. T. Ma, M. Tobah, H. Wang, *et al.*, “Benchmarking deep learning-based models on nanophotonic inverse design problems,” *Opto-Electronic Sci.* **1**(1), 210012 (2022).
19. A. G. Baydin, B. A. Pearlmutter, A. A. Radul, *et al.*, “Automatic Differentiation in Machine Learning: A Survey,” *J. Mach. Learn. Res.* **18**(1), 5595–5637 (2018).
20. T. W. Hughes, I. A. D. Williamson, M. Minkov, *et al.*, “Forward-Mode Differentiation of Maxwell’s Equations,” *ACS Photonics* **6**(11), 3010–3016 (2019).
21. M. Minkov, I. A. D. Williamson, L. C. Andreani, *et al.*, “Inverse Design of Photonic Crystals through Automatic Differentiation,” *ACS Photonics* **7**(7), 1729–1741 (2020).
22. S. Colburn and A. Majumdar, “Inverse design and flexible parameterization of meta-optics using algorithmic differentiation,” *Commun. Phys.* **4**(1), 65 (2021).
23. Y. Deng, S. Ren, K. Fan, *et al.*, “Neural-adjoint method for the inverse design of all-dielectric metasurfaces,” *Opt. Express* **29**(5), 7526–7534 (2021).
24. R. S. Hegde, “Photonics Inverse Design: Pairing Deep Neural Networks With Evolutionary Algorithms,” *IEEE J. Sel. Top. Quantum Electron.* **26**(1), 1–8 (2020).
25. L. Kuhn, T. Repän, and C. Rockstuhl, “Inverse design of core-shell particles with discrete material classes using neural networks,” *Sci. Rep.* **12**(1), 19019 (2022).
26. W. Zuo and K. Saitou, “Multi-material topology optimization using ordered SIMP interpolation,” *Struct. Multidiscip. Optim.* **55**(2), 477–491 (2017).
27. D. Li and I. Y. Kim, “Multi-material topology optimization for practical lightweight design,” *Struct. Multidiscip. Optim.* **58**(3), 1081–1094 (2018).
28. Y. Augenstein, T. Repän, and C. Rockstuhl, “Neural Operator-Based Surrogate Solver for Free-Form Electromagnetic Inverse Design,” *ACS Photonics* **10**(5), 1547–1557 (2023).
29. T. W. Radford, P. R. Wiecha, A. Politi, *et al.*, “Inverse Design of Unitary Transmission Matrices in Silicon Photonic Coupled Waveguide Arrays Using a Neural Adjoint Model,” *ACS Photonics* **12**(3), 1480–1493 (2025).
30. K. Monakhova, V. Tran, G. Kuo, *et al.*, “Untrained networks for compressive lensless photography,” *Opt. Express* **29**(13), 20913–20929 (2021).
31. A. Chandrasekhar and K. Suresh, “Multi-Material Topology Optimization Using Neural Networks,” *Comput. Des.* **136**, 103017 (2021).
32. M. L. De Marco, T. Jiang, J. Fang, *et al.*, “Broadband Forward Light Scattering by Architectural Design of Core-Shell Silicon Particles,” *Adv. Funct. Mater.* **31**(26), 2100915 (2021).
33. L. Lermusiaux, L. Roach, M. Lehtihet, *et al.*, “Silver Nanoshells with Optimized Infrared Optical Response: Synthesis for Thin-Shell Formation, and Optical/Thermal Properties after Embedding in Polymeric Films,” *Nanomaterials* **13**(3), 614 (2023).
34. B. P. Shah, N. Pasquale, G. De, *et al.*, “Core-Shell Nanoparticle-Based Peptide Therapeutics and Combined Hyperthermia for Enhanced Cancer Cell Apoptosis,” *ACS Nano* **8**(9), 9379–9387 (2014).
35. S. Sun, I. L. Rasskazov, P. S. Carney, *et al.*, “Critical Role of Shell in Enhanced Fluorescence of Metal-Dielectric Core-Shell Nanoparticles,” *The J. Phys. Chem. C* **124**(24), 13365–13373 (2020).
36. P. K. Kalambate, Dhanjai, Z. Huang, *et al.*, “Core@shell nanomaterials based sensing devices: A review,” *TrAC Trends Anal. Chem.* **115**, 147–161 (2019).
37. G.-X. Liu, J.-F. Liu, W.-J. Zhou, *et al.*, “Inverse design in quantum nanophotonics: Combining local-density-of-states and deep learning,” *Nanophotonics* **12**(11), 1943–1955 (2023).
38. R. Ghosh Chaudhuri and S. Paria, “Core/Shell Nanoparticles: Classes, Properties, Synthesis Mechanisms, Characterization, and Applications,” *Chem. Rev.* **112**(4), 2373–2433 (2012).
39. A. M. El-Toni, M. A. Habila, J. P. Labis, *et al.*, “Design, synthesis and applications of core-shell, hollow core, and nanorattle multifunctional nanostructures,” *Nanoscale* **8**(5), 2510–2531 (2016).
40. J. Peurifoy, Y. Shen, L. Jing, *et al.*, “Nanophotonic particle simulation and inverse design using artificial neural networks,” *Sci. Adv.* **4**(6), eaar4206 (2018).
41. S. So, J. Mun, and J. Rho, “Simultaneous Inverse Design of Materials and Structures via Deep Learning: Demonstration of Dipole Resonance Engineering Using Core-Shell Nanoparticles,” *ACS Appl. Mater. Interfaces* **11**(27), 24264–24268 (2019).
42. M. Kerker, D.-S. Wang, and C. L. Giles, “Electromagnetic scattering by magnetic spheres,” *J. Opt. Soc. Am.* **73**(6), 765 (1983).

43. M. Decker, I. Staude, M. Falkner, *et al.*, “High-Efficiency Dielectric Huygens’ Surfaces,” *Adv. Opt. Mater.* **3**(6), 813–820 (2015).
44. A. Rahimzadegan, D. Arslan, D. Dams, *et al.*, “Beyond dipolar Huygens’ metasurfaces for full-phase coverage and unity transmittance,” *Nanophotonics* **9**(1), 75–82 (2020).
45. C. Gigli, Q. Li, P. Chavel, *et al.*, “Fundamental Limitations of Huygens’ Metasurfaces for Optical Beam Shaping,” *Laser Photonics Rev.* **15**(8), 2000448 (2021).
46. G. Mie, “Beiträge zur Optik trüber Medien, speziell kolloidaler Metallösungen,” *Ann. Phys.* **330**(3), 377–445 (1908).
47. W. Liu and Y. S. Kivshar, “Generalized Kerker effects in nanophotonics and meta-optics,” *Opt. Express* **26**(10), 13085–13105 (2018).
48. C. F. Bohren and D. R. Huffman, *Absorption and Scattering of Light by Small Particles* (Wiley, 1998).
49. P. Wiecha, “Pymiecs - a simple python Mie toolbox,” GitLab, 2024, <https://gitlab.com/wiechapeter/pymiecs>.
50. T. W. Odom, E.-A. You, and C. M. Sweeney, “Multiscale Plasmonic Nanoparticles and the Inverse Problem,” *J. Phys. Chem. Lett.* **3**(18), 2611–2616 (2012).
51. U. Qamar and M. S. Raza, “Data Preprocessing,” in *Data Science Concepts and Techniques with Applications*, U. Qamar and M. S. Raza, eds. (Springer, Singapore, 2020), pp. 63–86.
52. I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016).
53. M. Abadi, A. Agarwal, P. Barham, *et al.*, “TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems,” <https://www.tensorflow.org/> (2015).
54. D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” *arXiv* (2014).
55. S. L. Smith, P.-J. Kindermans, C. Ying, *et al.*, “Don’t Decay the Learning Rate, Increase the Batch Size,” *arXiv* (2018).
56. I. Gulrajani, F. Ahmed, M. Arjovsky, *et al.*, “Improved Training of Wasserstein GANs,” (2017).
57. S. Raghavendra, P. Shirzadian, and K. Zhang, “A New Robust Partial p-Wasserstein-Based Metric for Comparing Distributions,” in *Proceedings of the 41st International Conference on Machine Learning*, (PMLR, 2024), pp. 41867–41885.
58. I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, *et al.*, “Generative Adversarial Networks,” *arXiv* (2014).
59. T. Sainburg, M. Thielk, B. Theilman, *et al.*, “Generative adversarial interpolative autoencoding: Adversarial training on latent space interpolations encourage convex latent distributions,” (2019).
60. D. Simon, *Evolutionary Optimization Algorithms* (John Wiley & Sons, 2013).
61. P. Bennet, C. Doerr, A. Moreau, *et al.*, “Nevergrad: Black-box optimization platform,” *ACM SIGEVOlution* **14**(1), 8–15 (2021).
62. P. R. Wiecha and O. L. Muskens, “Deep Learning Meets Nanophotonics: A Generalized Accurate Predictor for Near Fields and Far Fields of Arbitrary 3D Nanostructures,” *Nano Lett.* **20**(1), 329–338 (2020).
63. P. R. Wiecha, C. Majorel, C. Girard, *et al.*, “Design of plasmonic directional antennas via evolutionary optimization,” *Opt. Express* **27**(20), 29069–29081 (2019).
64. P. R. Wiecha, A. Arbouet, C. Girard, *et al.*, “Evolutionary multi-objective optimization of colour pixels based on dielectric nanoantennas,” *Nat. Nanotechnol.* **12**(2), 163–169 (2017).
65. K. Deb, *Multi-Objective Optimization Using Evolutionary Algorithms*, vol. 16 (Wiley, 2001).
66. W. Chen, X. Zhang, B. Lin, *et al.*, “Gradient-Based Multi-Objective Deep Learning: Algorithms, Theories, Applications, and Beyond,” (2025).
67. S. J. Oldenburg, R. D. Averitt, S. L. Westcott, *et al.*, “Nanoengineering of optical resonances,” *Chem. Phys. Lett.* **288**(2–4), 243–247 (1998).
68. Y. Tsuchimoto, T.-a. Yano, T. Hayashi, *et al.*, “Fano resonant all-dielectric core/shell nanoparticles with ultrahigh scattering directionality in the visible region,” *Opt. Express* **24**(13), 14451–14462 (2016).
69. W. Chaâbani, J. Proust, A. Movsesyan, *et al.*, “Large-Scale and Low-Cost Fabrication of Silicon Mie Resonators,” *ACS Nano* **13**(4), 4199–4208 (2019).
70. M. A. Eslamisaray, P. R. Wray, Y. Lee, *et al.*, “A Single-Step Bottom-up Approach for Synthesis of Highly Uniform Mie-Resonant Crystalline Semiconductor Particles at Visible Wavelengths,” *Nano Lett.* **23**(5), 1930–1937 (2023).
71. M. A. Parker, S. Khaddad, N. Fares, *et al.*, “Unveiling the Potential of Redox Chemistry to Form Size-Tunable, High-Index Silicon Particles,” *Chem. Mater.* **36**(22), 10986–10993 (2024).
72. W. Chaâbani, J. Proust, S. Ouellet, *et al.*, “Si@Au Core–Shell Nanostructures: Toward a New Platform for Controlling Optical Properties at the Nanoscale,” *J. Phys. Chem. C* **125**(37), 20606–20616 (2021).
73. H. Sugimoto, T. Hinamoto, Y. Kazuoka, *et al.*, “Mode Hybridization in Silicon Core–Gold Shell Nanosphere,” *Small* **18**(45), 2204890 (2022).
74. L. Lermusiaux, M. Plissonneau, L. Bertry, *et al.*, “Seeded growth of ultrathin gold nanoshells using polymer additives and microwave radiation,” *Sci. Rep.* **11**(1), 17831 (2021).
75. A. Désert, I. Chaduc, S. Fouilloux, *et al.*, “High-yield preparation of polystyrene/silica clusters of controlled morphology,” *Polym. Chem.* **3**(5), 1130–1132 (2021).
76. H. Blas, M. Save, P. Pasetto, *et al.*, “Elaboration of Monodisperse Spherical Hollow Particles with Ordered Mesoporous Silica Shells via Dual Latex/Surfactant Templating: Radial Orientation of Mesopore Channels,” *Langmuir* **24**(22), 13132–13137 (2008).
77. W. Li, A. Elzatahry, D. Aldhayan, *et al.*, “Core–shell structured titanium dioxide nanomaterials for solar energy utilization,” *Chem. Soc. Rev.* **47**(22), 8203–8237 (2018).
78. A. Imhof, “Preparation and Characterization of Titania-Coated Polystyrene Spheres and Hollow Titania Shells,” *Langmuir* **17**(12), 3579–3585 (2001).

79. P. M. Arnal, C. Weidenthaler, and F. Schüth, “Highly Monodisperse Zirconia-Coated Silica Spheres and Zirconia/Silica Hollow Spheres with Remarkable Textural Properties,” *Chem. Mater.* **18**(11), 2733–2739 (2006).
80. T. Kondratowicz, M. Drozdek, A. Rokicińska, *et al.*, “Novel CuO-containing catalysts based on ZrO₂ hollow spheres for total oxidation of toluene,” *Microporous Mesoporous Mater.* **279**, 446–455 (2019).
81. S. Dalin, “Gradient-based optimization of core-shell particles with discrete materials for directional scattering: code,” GitHub, 2025, https://github.com/S-Dalin/gradient_optimization_with_discrete_materials.