Research Article | Vol. 30, No. 25 / 5 Dec 2022 / Optics Express 45365
Optics EXPRESS

Exploiting geometric biases in inverse nano-optical problems using artificial neural networks

TAAVI REPÄN,[1,*] YANNICK AUGENSTEIN,[2] AND CARSTEN ROCKSTUHL[2,3]

[1]Institute of Physics, University of Tartu, W. Ostwald St 1, Tartu 50411, Estonia
[2]Institute of Theoretical Solid State Physics, Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany
[3]Institute of Nanotechnology, Karlsruhe Institute of Technology, 76344 Eggenstein-Leopoldshafen, Germany
*taavi.repan@ut.ee

Abstract: Solving the inverse problem is a major challenge in contemporary nano-optics. However, frequently not just a possible solution needs to be found but rather the solution that accommodates constraints imposed by the problem at hand. To select the most plausible solution for a nano-optical inverse problem additional information can be used in general, but how to specifically formulate it frequently remains unclear. Here, while studying the reconstruction of the shape of an object using the electromagnetic field in its proximity, we show how to take advantage of artificial neural networks (ANNs) to produce solutions consistent with prior assumptions concerning the structures. By preparing suitable datasets where the specific shapes of possible scatterers are defined, the ANNs learn the underlying scatterer present in the datasets. This helps to find a plausible solution to the otherwise non-unique inverse problem. We show that topology optimization, in contrast, can fail to recover the scatterer geometry meaningfully but a hybrid approach that is based on both, ANNs and a topology optimization, eventually leads to the most promising performance. Our work has direct implications in fields such as optical metrology.

1. Introduction

In experimental nanophotonics, given the growing importance and complexity of nanostructures [1–8], the characterization of fabricated structures is of crucial importance. As the structures are realized, they will inevitably look different from the design considered in prior numerical simulations. Then, establishing how the structures exactly ended up looking is the first step in iterating towards a design with properties that were initially envisioned. Of course, there is a spectrum of different characterization methods available [9,10]. However, the most comprehensive methods tend up to be destructive and time-consuming. Therefore, probing the optical response from a given structure with light and reconstructing from that information the actual structure is a prime task for the field of metrology. For example, solving the inverse problem to characterize grating lines fabricated using electron-beam lithography by examining non-destructive X-ray diffraction patterns [11]. Also, it is crucial to constantly monitor the quality and geometry of fabricated samples in a production line. This occurs, for example, in the semiconductor industry. Here again, observing the samples under high-resolution electron microscopy is an option. But this is time-consuming and potentially destroys the samples. Then, again, probing the samples with light and measuring the optical response from which one can conclude on the details of

#474260 https://doi.org/10.1364/OE.474260
Journal © 2022 Received 14 Sep 2022; revised 4 Nov 2022; accepted 7 Nov 2022; published 30 Nov 2022

Optics EXPRESS

fabricated samples is a prime goal. This is the general problem we consider here: how to obtain the geometry of the object based on observing scattered fields from that object.

Such problems fall under the context of inverse scattering problems (ISP), where—in the general case—the shape and composition of the scatterer is to be determined based on measuring scattered fields from a number of sources by a number of detectors. Traditionally, different iterative solution methods have been developed [12], while recently, there has been a series of results based on artificial neural networks (ANNs) [13–16]. However, here we look at a slightly different formulation, where we consider a single plane wave illumination and look at the electric fields known over a specific area. Although this is a more restrictive formulation of the problem, it should be stressed that, in principle, the method can be extended to a wider range of ISP schemes since the phase information can be obtained from intensity measurements [17,18] and thus the spatial distributions of the field in real space can be reconstructed [12]. Note that the problem is non-unique: there are many possible shapes that produce the same scattered fields. So, one should seek to use prior assumptions about the expected structure of the solution to limit the solution space. The novel contribution of our work is to consider biased training sets for the neural network. This way, knowledge about the expected structure of the scatterers can be easily integrated into the inverse problem, yielding more plausible solutions.

Potentially, the state of the art method for solving inverse design problems in nanophotonics is topology optimization (TO) [19–22]. With the recent growing interest in machine learning, deep learning methods have proven useful as well [23–30]. It is well known that the quality of ANN predictions depends on the size and variety of the training dataset. Thus, care must be taken to use large representative datasets and tools to improve training efficiency [31,32]. For the inverse problem considered here, however, this weakness can be exploited as a strength instead. In particular, we note that by using biased datasets to train ANNs to solve the inverse problem, we can easily embed assumptions about expected solutions into the solver. In contrast, such assumptions for TO solver need to be implemented in a differentiable manner, which is not always straightforward to do. For example, these might take the form of additional partial differential equations [33,34]. It has been shown [35,36] that in inverse design problems, it can be beneficial to refine ANN predictions by using them as an initial guess for a TO solver. Indeed, we also show that the quality of ANN predictions can be improved, if needed, by a TO refinement step.

Technically speaking, we consider an inverse problem, where we have (partially) known electric fields, and we want to know the structure that produced these fields, as illustrated in Fig. 1. Here, we consider a 2D case in the x-y-plane and look at the TE polarization, i.e., we consider E_z fields. Throughout the paper, we consider a 5 μm by 5 μm domain, discretized by a regular 128×128 grid. The structures consists of an isotropic material with a permittivity of 2.25 and are illuminated with a plane wave (wavelength 1 μm). For calculating the fields, we use frequency-domain finite-difference (FDFD) method [37]. Note that we only supply a fraction of the spatial distributions of the fields to the ANN or TO code. This sets a somewhat more realistic challenge, where not all the fields are known. Notably, only the fields outside the structure are known in this case, rendering the retrieval task much more challenging. Specifically, we exclude the central 64×64 pixels from our consideration and simply put them to zero, thus depriving the ANN of the information of the fields inside and near the unknown object. That is, of course, reasonable, as measuring the field inside an object is out of reach, if not to say impossible. For this, we have generated our datasets such that all scatterers always fit within the central 64×64 pixels, so the fields given for the solver do not contain internal fields of the scatterer. Although we present here the results based on complete complex fields, i.e., the real and imaginary components of E_z are available, the methods here work (with minor tweaks) also for the case where only field intensity $|E_z|^2$ is known.

For our networks, we use the UNet architecture [38–42], which is a convolutional neural network consisting of downscaling blocks followed by upscaling blocks (see Fig. 1). In our case,

Optics EXPRESS

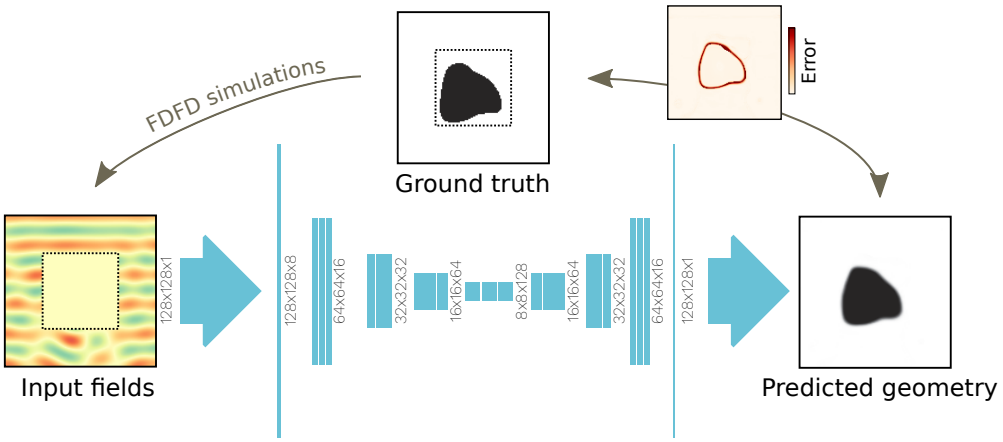

Fig. 1. Sketch of the problem and the used ANN architecture. Here, we use FDFD simulations to calculate the fields of the object under plane wave illumination. The ANN then predicts from the fields in its proximity the supposed scatterer geometry, which we then compare against the ground truth during training. The ANN takes an image with a size of 128x128 of the fields (with central region masked out) and predicts the geometry corresponding to the fields. The ANN consists of downscaling and upscaling blocks, each having three convolutional layers.

we use four blocks for each (unless specified otherwise). That is, we scale the initial 128×128 input image down to 64×64, and so on until 8×8. With each downscaling, we double the number of channels. Before the first downscaling block, we use a single convolution layer to expand the single-channel input image to eight channels. At the output, we end up with again an 128×128 image, with one channel containing the predicted permittivity. Each upscaling and downscaling block consists of three 2D convolutional layers with Leaky ReLU activations, followed by MaxPool layer (only for downscaling blocks).

2. Learning dataset structure

To highlight the benefits of our ANN-based approach, we will first consider rather specific structures with small (but well-defined) features. The dataset consists of shapes formed out of three circles connected with small beams with a small hole in the middle each of the beams [see Fig. 2(a)]. Since the finest features of these samples are strongly subwavelength, there is insufficient information in the fields outside the scattering object to restore these features uniquely. This suggests that for successfully recovering the features, the ANN has to learn the underlying structure of the samples in the dataset.

To set the stage, we exploit at first topology optimization (TO) [43–52] as the reference to benchmark our method, which is state of the art for such inverse problems (see Methods for details). TO is able to find solutions that reproduce the known fields within the spatial domain that was provided as information, i.e., excluding the fields in the central region. However, in the results both, the fields in that central domain and the restored material distribution, have minimal if not to say negligible similarity to the ground truth. Figure 2(c) to 2(f) shows the four best solutions from topology optimization, i.e., with the lowest error in reproducing the known fields. None of the four solutions resembles the ground truth geometry, although all of them have converged to nearly zero in terms of the optimization objective. Looking at the field error distribution in detail shows that in the spatial domain considered for the optimization, the error

Optics EXPRESS

is nearly zero (as expected), while in the central domain, the fields differ significantly from the ground truth.

That disagreement in the field just mirrors the fact that the structure in this central domain was not retrieved correctly, hence the large deviation to the ground truth in all these samples. This is very illustrative for the non-uniqueness of this kind of inverse problem with only partial information available. The difficulty with TO is that even if there are strong constraints on the plausible solutions, it is not always straightforward to implement them in a differentiable manner, which is a requirement for TO. In this work, we formulate the topology optimization problem with the commonly used filter-and-project method [19,33,53]. Except for the feature size, such parametrization places few restrictions on the design space, a property that is beneficial for topology optimization in the general case. However, the results in Fig. 2 show that this parametrization is not suitable to recover the correct structures in our problem.

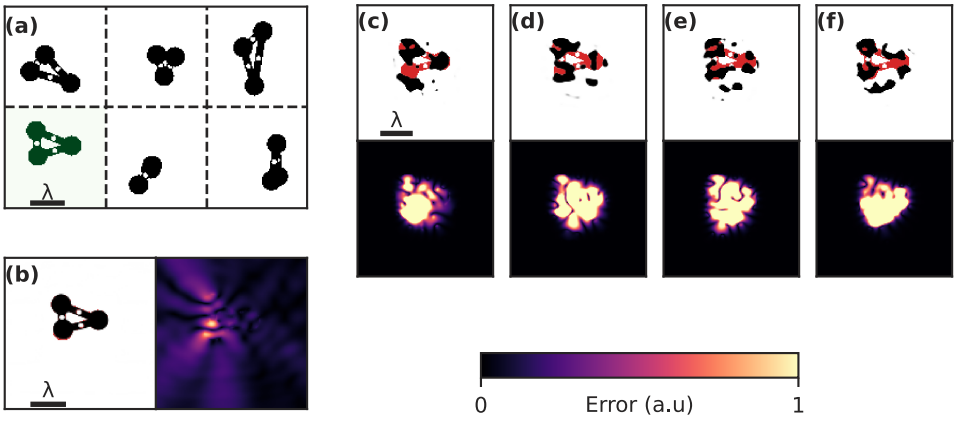

Fig. 2. (a) Samples from the validation dataset. The structure shown in green is used as an example for the inverse problem in the subfigures (b-f) of this figure. (b) The result from the ANN prediction, with errors in reconstructed fields shown. The ground truth structure is underlaid with red color. We stress that barely a difference can be seen. (c-f) Top four results from topology optimization approach, with corresponding errors in reconstructed fields shown below. Largely different structures can generate a very similar field as considered in the solution of the inverse problem.

In contrast, the ANN-based approach works well here. Notably, the inherent structure in the dataset is learned by the ANN without the need to implement any assumptions in the code. Figure 2(b) shows the prediction for the same input as used for topology optimization, showing that the ANN prediction is much closer to the ground truth geometry. Crucially, the geometry is now relatively well recovered, as only minimal deviations from the ground truth are visible in the Fig. 2(b). Nevertheless, the predicted geometry is very slightly deformed compared to the ground truth. This leads to a small deviation of the scattered fields (compared to the ground truth) over the whole domain. However, these errors are significantly lower than corresponding errors of fields from the TO solutions. The ANN shown here was trained on 16 000 samples for around 60 epochs. On GTX 1080 GPU, this took around 30 minutes. Data generation took around 0.1–0.2 seconds per sample on one CPU core, thus generating the whole dataset on our 24-core machine took just a few minutes.

ANN only performs a single-shot prediction, unlike TO, which iteratively improves the solution until the optimization objective (field error in the outside domain) has converged. Thus, it is not surprising that the error in the scattered fields is much higher for the solution from the ANN. However, in terms of actual geometry, the ANN solution is nearly perfect, while the TO falls

Optics EXPRESS

victim to the non-uniqueness of the problem. Although this is a toy example, it highlights the versatility of the ANN approach, as the ANNs are able to learn this highly specific geometry. Importantly, no changes to the code are required to handle different kinds of structures: training a new ANN for a new dataset is sufficient.

3. General datasets

We now look at more general shapes to show that (1) ANNs offer reasonably good predictions even for datasets with more diverse shapes and (2) ANNs learn the underlying structure in the dataset. For this, we consider two different datasets: "blobs" [Fig. 3(a)] and "triangles" [Fig. 3(b)]. Each sample in the "blobs" dataset is generated by sampling 128×128 points from a uniform random distribution with values between 0 and 1. A convolution with a two-dimensional Gaussian kernel with a size of 10×10 pixels is then applied to this grid, which is subsequently thresholded at a value of 0.5. This procedure yields smooth, random shapes with no discernible underlying structure and can include multiple scatterers. Before thresholding we set outer pixels to zero, so that for the given filtering and thresholding parameters the generated shapes will always be limited to the center 64×64 region. The "triangles" dataset is constructed out of variable number of fixed size triangles, leading to shapes with a more obvious underlying structure. The shapes are generated in a way to ensure that all structures are contained within the center 64×64 region. Note that neither of these datasets have samples with small subwavelength features, and thus the non-uniqueness of the inverse problem will not play a crucial role here. Then we trained two ANNs: one was trained exclusively on the "blobs" dataset, while the other was trained on the "triangles" dataset. For both trainings we used 16 000 samples and ran the training for 50–80 epochs (around 50 minutes). For "blobs" dataset, five downscaling and five upscaling blocks were used in the UNet, instead four as for the other nets.

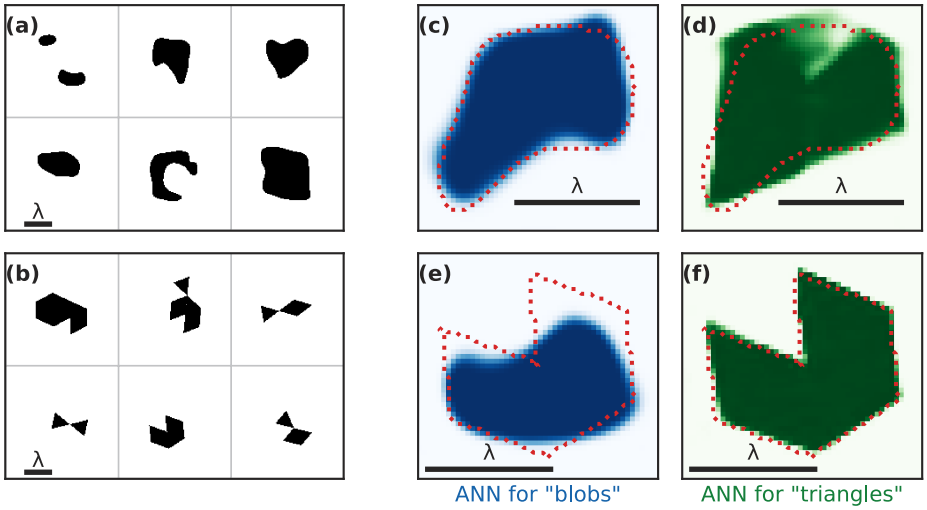

Fig. 3. (a) Examples from the "blobs" dataset. (b) Examples from the "triangles" dataset. (c,d) Predictions for a "blob" structure from ANN trained on blobs (c) or triangular structures (d). (e,f) Predictions for a "triangular" structure from ANN trained on blobs (e) or triangular structures (f). The dotted red lines in (c-f) indicate the ground truth.

Figure 3(c) and 3(d) show predictions from both ANNs for a validation sample from the "blobs" dataset. Similarly, Fig. 3(e) and 3(f) compare the ANNs for an input from the "triangles" dataset. These results show that the ANNs learn the characteristics of structures in the dataset, and it tends to replicate those features even on input data that is not produced by such data. As exemplified

Optics EXPRESS

above (Fig. 2), this allows—with little effort—to use prior assumptions to retrieve more plausible solutions. In contrast, implementing such assumptions in the topology optimization procedure is not straightforward.

For more statistical comparison, we generated additional 100 testing samples (i.e., not used during the training) for both datasets. We then compared ANN performance against topology optimization on these samples. To compare the results, we consider the error in the reconstructed parameters

$$\Delta\varepsilon = \sum_i \left[\varepsilon\left(x_i, y_i\right) - \varepsilon^{\mathrm{ref}}\left(x_i, y_i\right)\right]^2, \tag{1}$$

where $\varepsilon^{\mathrm{ref}}\left(x_i, y_i\right)$ is the permittivity distribution from which the input fields are calculated (the ground truth). Similarly, we look at the error in the reconstructed fields

$$\Delta E_z = \sum_i \left[E_z\left(x_i, y_i\right) - E_z^{\mathrm{ref}}\left(x_i, y_i\right)\right]^2, \tag{2}$$

where $E_z\left(x_i, y_i\right)$ are electric fields calculated using the predicted permittivity distribution $\varepsilon\left(x_i, y_i\right)$, while $E_z^{\mathrm{ref}}\left(x_i, y_i\right)$ are the input fields corresponding to $\varepsilon^{\mathrm{ref}}\left(x_i, y_i\right)$. Results are shown in Fig. 4, both for the blobs [Fig. 4(a)] and the triangles dataset [Fig. 4(b)]. There, we show a scatterplot of the error in predicted structures vs. error in corresponding fields. The errors are compared to errors obtained using a TO as well for comparison. In the TO, the problem was solved for 13 different initial conditions, and the best result from these 13 different runs is considered as the result.

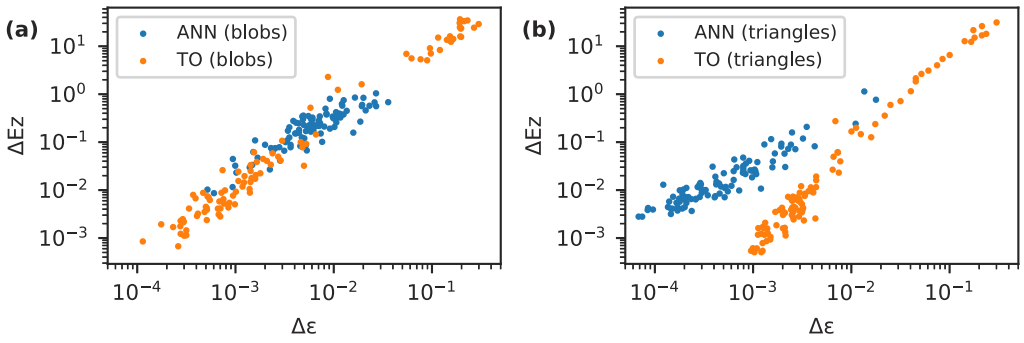

Fig. 4. Scatterplot of the error in the predicted structure ($\Delta\epsilon$) vs. error in corresponding fields (ΔE_z) comparing results from TO and ANN for (a) the "blobs" dataset and (b) the "triangles" dataset.

As expected, by tendency the TO approach yields lower ΔE_z, as this corresponds to the optimization objective. However, we note that the ANN results are competitive in reconstructing the scatterer geometry. While in the case of the blobs dataset, the TO outperforms the ANNs, for the triangles dataset, the ANN approach actually manages to recover the scatterer geometry better. This is because the TO parametrization works naturally for the blobs dataset, but due to the Gaussian filtering in the filter-and-project method, the TO cannot produce such sharp features. Another consideration is the dataset complexity: while the triangles dataset itself is also rather diverse, it is still less diverse compared to the blobs dataset. We used 16 000 training samples for training both ANNs for blobs and triangles. However, we expect that for the blobs dataset, more samples would improve the prediction performance, if so is needed.

It must be noted that topology optimization is sensitive to initial conditions. In the results shown, we have performed 13 different optimization runs for each sample, with slightly different initial conditions (see Methods), to ensure that a good solution would be found. It often converged

Optics EXPRESS

to suboptimal solutions or even failed to find a good solution at all. Results of all optimization runs (grey dots) are shown in Fig. 5(a), where we see that only a small amount of optimization runs converges to a suitable solution (red dots). To more clearly highlight the proportion of failed optimization runs, the figure also shows kernel density estimate of the scatter plot data. For example, for the sample shown in Fig. 5(b), only one out of the 13 optimization runs succeeded in finding a good solution. Furthermore, every optimization run consists of 500 iterations, with two full-wave simulations per iteration. These are significant costs, at times even comparable to the costs of generating data for the ANN training. This clearly highlights the data-effectiveness of ANNs over TO for the present problem setting.

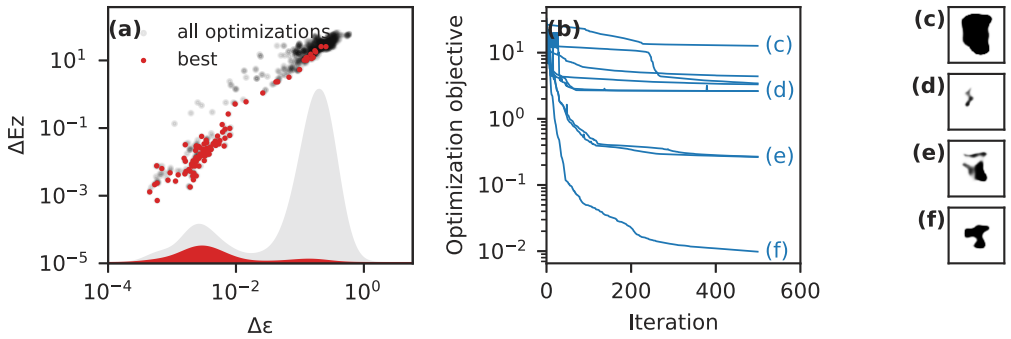

Fig. 5. (a) Scatterplot of error in predicted structures vs. error in corresponding fields for different optimizations using TO but with different initial conditions. The results from all optimization runs are shown in grey (100 different inputs and 13 different initial conditions), while the best run for each of the 100 inputs are highlighted in red. The filled contours indicate kernel density estimate, showing the large amount of failed optimization runs. (b) Plots of optimization objective of all optimization runs for one particular input sample from the dataset. (c-f) Selected optimization results.

The results of this section show that even on more general datasets the ANNs offer sufficiently accurate predictions. However, at this point, it is worth reemphasizing that the motivation behind using ANNs is not replacing TO in general. Instead, the aim is to exploit highly specific training datasets to solve the inverse problem in a restricted solution space.

4. Refinement of ANN predictions

As shown above, the ANNs offer reasonably good accuracy considering the low numerical costs needed for evaluating the ANNs (in comparison to costs of topology optimization). However, there might be cases where higher levels of accuracy are required. For this, it is straightforward to use ANN predictions as an initial guess for topology optimization, which seeds the optimization with a reasonably good geometry that tends to follow the typical shapes provided by the dataset. This (mostly) avoids issues related to poor local minima in topology optimization, and one optimization run is sufficient. Results of this approach for the "blobs" dataset are shown in Fig. 6(a). We show the results from the ANN alone and from the TO alone (best of 13 TO runs). In the combined approach, we use the predictions from the ANN and let the topology optimization run again for 500 iterations. As clearly seen in the results, thanks to the good starting point, this hybrid approach provides more accurate reconstructions than the pure topology optimization approach. For example, Figure 6(b) shows one of the less successful ANN predictions [e.g. compare against TO result in Fig. 6(d)]. However, after running a refinement optimization, the hybrid approach [Fig. 6(c)] converges to the level of accuracy of pure TO, even exceeding it. For practical purposes, we consider them to be equally converged, though, as the small differences in retrieved permittivity (on the order of 10^{-3}) are unlikely to be practically relevant. Moreover,

Optics EXPRESS

the final error for the TO approach depends on the number of iterations, which we have fixed at 500 here. It must be emphasized here that the refinement of the ANN prediction is significantly cheaper than topology optimization alone, as here, only a single topology optimization run was necessary to obtain good results. Of course, for the results shown here, the ANN training costs (up to an hour) dominate the time requirements of a single TO optimization run (two to three minutes), so typically, for ANN-based methods, the trained ANN should be reused in order to amortize the training costs.

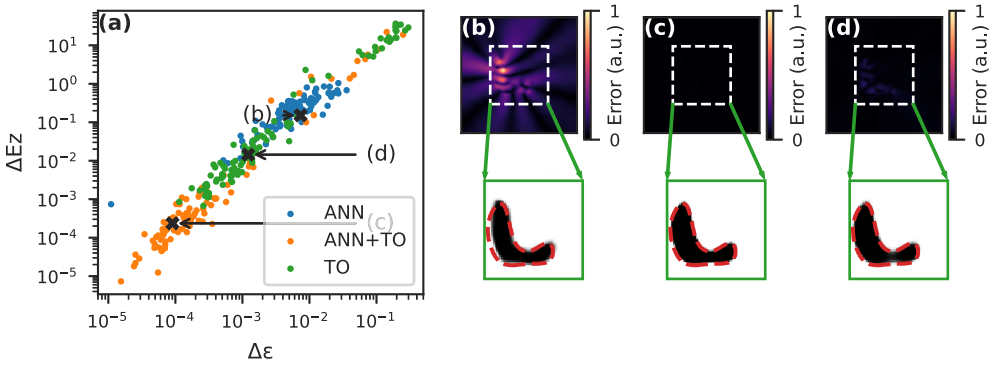

Fig. 6. (a) Scatterplot of MSE of permittivity and field errors. We compare here the results from the ANN (blue dots) and best TO result out of 13 runs (green dots). Moreover, we use a single TO run using the ANN prediction as the initial guess (orange dots). The superiority of that approach is obvious. (b) ANN prediction for one particular input field. (c) The same prediction, after refinement with a single TO run. (d) Pure TO result for the same input, best of 13 optimization runs. Red dashed line shows the ground truth in (b-d). Error plots are showing errors in the predicted fields to the input fields.

5. Conclusions

To summarize, we propose using ANNs to retrieve the scatterer geometry from partially known fields. Specifically, we consider working in a restricted solution space, i.e., the allowed scatterer geometries are limited. Implementing such restrictions in the state of the art method, topology optimization, can be exceedingly difficult, depending on the restrictions sought. In contrast, it is much easier to implemenent such restrictions with ANNs: by training the ANN on specific datasets, it learns the features of the structures and reproduces them in the predicted geometries. This is the key benefit of using ANNs for this problem.

Interestingly, we saw other benefits in using ANNs aswell. Recovering the structures from only partially known fields has multiple solutions (i.e., it is a non-unique problem). For topology optimization, this means that the optimization landscape suffers from local minima, requiring multiple optimization attempts with different initial conditions to find a good solution. On the other hand, ANN predicted structures are generally relatively close to the ground truth.

Furthermore, ANN predictions serve well as an initial guess to the topology optimization, combining best of the two methods: ANN will provide a good initial guess with guess conforming to the structure of the training dataset, and topology optimization will then refine the structure to minimize the error between the predicted and input fields.

These findings are important in metrology, but also in the design of photonic nanostructures that should cope with specific geometrical constraints hard to express in alternative optimization techniques.

Optics EXPRESS

6. Methods

FDFD simulations For simulations we use a Python FDFD code based on https://github.com/fancompute/fdfdpy. The simulation domain is terminated by perfectly-matched layers (with thickness $\lambda/2 = 0.5$ μm).

Topology optimization In TO, the design domain is discretized into material elements with continuous permittivities. An objective function is then defined and evaluated via a forward simulation. The optimization objective is mean square error between the input fields (from the unknown object) and fields produced by the structure under optimization. Through the adjoint method, it is possible to derive the gradients of the objective function with respect to each element's permittivity by performing an additional so-called "adjoint" simulation. The forward and adjoint simulations are carried out by the FDFD method described above. The permittivities are then iteratively updated using a gradient-based optimization algorithm (e.g., BFGS, MMA) until sufficient convergence of the error is achieved. The design density ρ is usually initialized with a constant value of 0.5. In order to have different optimization runs (in case where the optimization fails to converge a good solution), the additional optimization runs initialized with a random constant value between 0.4 and 0.6. A Gaussian filter is applied to the design density ρ, which regularizes the optimization problem and places a constraint on the minimum feature size in the optimized design. The variables are then passed through a "projection"—a soft-thresholding function that helps to binarize the design, resulting in the permittivity distribution of the optimized structure.

ANN training The ANNs are implemented using PyTorch framework and trained on a single Nvidia GPU. Adam optimizer is used, with a fixed learning rate of 0.001, with no learning rate scheduling. Batch size for the training is 20.

Funding. Eesti Teadusagentuur (PSG716); Helmholtz Association; Deutsche Forschungsgemeinschaft (390761711); Carl-Zeiss-Stiftung.

Acknowledgments. We acknowledge support by the German Research Foundation within the Excellence Cluster 3D Matter Made to Order (EXC 2082/1 under project number - 390761711), by the Carl Zeiss Foundation, and by the Helmholtz Association via the Helmholtz program "Materials Systems Engineering" (MSE). This work was supported by the Estonian Research Council grant (PSG716). This work was performed on the HoreKa supercomputer funded by the Ministry of Science, Research and the Arts Baden-Württemberg and by the Federal Ministry of Education and Research.

Disclosures. The authors declare no conflicts of interest.

Data availability. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

References

1. M. Rajaei, J. Zeng, M. Albooyeh, M. Kamandi, M. Hanifeh, F. Capolino, and H. K. Wickramasinghe, "Giant Circular Dichroism at Visible Frequencies Enabled by Plasmonic Ramp-Shaped Nanostructures," ACS Photonics **6**(4), 924–931 (2019).
2. S. Nocentini, F. Riboli, M. Burresi, D. Martella, C. Parmeggiani, and D. S. Wiersma, "Three-Dimensional Photonic Circuits in Rigid and Soft Polymers Tunable by Light," ACS Photonics **5**(8), 3222–3230 (2018).
3. P. Camayd-Muñoz, C. Ballew, G. Roberts, and A. Faraon, "Multifunctional volumetric meta-optics for color and polarization image sensors," Optica **7**(4), 280 (2020).
4. V. Hahn, P. Kiefer, T. Frenzel, J. Qu, E. Blasco, C. Barner-Kowollik, and M. Wegener, "Rapid Assembly of Small Materials Building Blocks (Voxels) into Large Functional 3D Metamaterials," Adv. Funct. Mater. **30**(26), 1907795 (2020).
5. L. Su, A. Y. Piggott, N. V. Sapra, J. Petykiewicz, and J. Vuckovic, "Inverse Design and Demonstration of a Compact on-Chip Narrowband Three-Channel Wavelength Demultiplexer," ACS Photonics **5**(2), 301–305 (2018).
6. C. Wang, Z. Li, R. Pan, W. Liu, H. Cheng, J. Li, W. Zhou, J. Tian, and S. Chen, "Giant Intrinsic Chirality in Curled Metasurfaces," ACS Photonics **7**(12), 3415–3422 (2020).
7. Y. Chen, Y. Hu, J. Zhao, Y. Deng, Z. Wang, X. Cheng, D. Lei, Y. Deng, and H. Duan, "Topology Optimization-Based Inverse Design of Plasmonic Nanodimer with Maximum Near-Field Enhancement," Adv. Funct. Mater. **30**(23), 2000642 (2020).
8. J. He, H. Chen, J. Hu, J. Zhou, Y. Zhang, A. Kovach, C. Sideris, M. C. Harrison, Y. Zhao, and A. M. Armani, "Nonlinear nanophotonic devices in the ultraviolet to visible wavelength range," Nanophotonics **9**(12), 3781–3804 (2020).

Optics EXPRESS

9. S. Myhra and J. C. Rivière, *Characterization of nanostructures* (Taylor & Francis, Boca Raton, 2013).
10. S. Amelinckx, D. van Dyck, J. van Landuyt, and G. van Tendeloo, *Electron Microscopy Principles and Fundamentals* (2008). OCLC: 890037233.
11. V. Soltwisch, A. Fernández Herrero, M. Pflüger, A. Haase, J. Probst, C. Laubis, M. Krumrey, and F. Scholze, "Reconstructing detailed line profiles of lamellar gratings from GISAXS patterns with a Maxwell solver," J. Appl. Crystallogr. **50**(5), 1524–1532 (2017).
12. D. Colton and R. Kress, *Inverse Acoustic and Electromagnetic Scattering Theory*, no. 93 in Applied Mathematical Sciences (Springer New York Springer e-books Imprint: Springer New York, NY, 2013), 3rd ed.
13. X. Chen, Z. Wei, M. Li, and P. Rocca, "A Review of Deep Learning Approaches for Inverse Scattering Problems," Prog. Electromagn. Res. **167**, 67–81 (2020).
14. Y. Khoo and L. Ying, "SwitchNet: A Neural Network Model for Forward and Inverse Scattering Problems," SIAM J. Sci. Comput. **41**(5), A3182–A3201 (2019).
15. L. Li, L. G. Wang, F. L. Teixeira, C. Liu, A. Nehorai, and T. J. Cui, "DeepNIS: Deep Neural Network for Nonlinear Electromagnetic Inverse Scattering," IEEE Trans. Antennas Propag. **67**(3), 1819–1825 (2019).
16. Y. Sanghvi, Y. Kalepu, and U. K. Khankhoje, "Embedding Deep Learning in Inverse Scattering Problems," IEEE Trans. Comput. Imaging **6**, 46–56 (2020).
17. J. R. Fienup, "Phase retrieval algorithms: a comparison," Appl. Opt. **21**(15), 2758 (1982).
18. K. Creath, "V Phase-Measurement Interferometry Techniques," in *Progress in Optics*, vol. 26 (Elsevier, 1988), pp. 349–393.
19. J. Jensen and O. Sigmund, "Topology optimization for nano-photonics," Laser Photonics Rev. **5**(2), 308–321 (2011).
20. S. Molesky, Z. Lin, A. Y. Piggott, W. Jin, J. Vuckovic, and A. W. Rodriguez, "Inverse design in nanophotonics," Nat. Photonics **12**(11), 659–670 (2018).
21. J. Park, S. Kim, D. W. Nam, H. Chung, C. Y. Park, and M. S. Jang, "Free-form optimization of nanophotonic devices: from classical methods to deep learning," Nanophotonics **11**(9), 1809–1845 (2022).
22. H. Chung and O. D. Miller, "High-NA achromatic metalenses by inverse design," Opt. Express **28**(5), 6945 (2020).
23. O. Khatib, S. Ren, J. Malof, and W. J. Padilla, "Deep Learning the Electromagnetic Properties of Metamaterials—A Comprehensive Review," Adv. Funct. Mater. **31**(31), 2101748 (2021).
24. P. Wiecha, A. Arbouet, C. Girard, and O. Muskens, "Deep learning in nano-photonics: inverse design and beyond," Photonics Research (2021).
25. S. So, T. Badloe, J. Noh, J. Bravo-Abad, and J. Rho, "Deep learning enabled inverse design in nanophotonics," Nanophotonics **9**(5), 1041–1057 (2020).
26. W. Ma, Z. Liu, Z. A. Kudyshev, A. Boltasseva, W. Cai, and Y. Liu, "Deep learning for the design of photonic structures," Nature Photonics (2020).
27. R. S. Hegde, "Deep learning: a new tool for photonic nanostructure design," Nanoscale Adv. **2**(3), 1007–1023 (2020).
28. 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 (2021).
29. Y. Chen and L. Dal Negro, "Physics-informed neural networks for imaging and parameter retrieval of photonic nanostructures from near-field data," APL Photonics **7**(1), 010802 (2022).
30. S. Ren, A. Mahendra, O. Khatib, Y. Deng, W. J. Padilla, and J. M. Malof, "Inverse deep learning methods and benchmarks for artificial electromagnetic material design," Nanoscale **14**(10), 3958–3969 (2022).
31. W. Ma and Y. Liu, "A data-efficient self-supervised deep learning model for design and characterization of nanophotonic structures," Sci. China Phys. Mech. Astron. **63**(8), 284212 (2020).
32. C. Qiu, X. Wu, Z. Luo, H. Yang, G. He, and B. Huang, "Nanophotonic inverse design with deep neural networks based on knowledge transfer using imbalanced datasets," Opt. Express **29**(18), 28406 (2021).
33. Y. Augenstein and C. Rockstuhl, "Inverse Design of Nanophotonic Devices with Structural Integrity," ACS Photonics **7**(8), 2190–2196 (2020).
34. Q. Li, W. Chen, S. Liu, and L. Tong, "Structural topology optimization considering connectivity constraint," Struct. Multidisc. Optim. **54**(4), 971–984 (2016).
35. Z. A. Kudyshev, A. V. Kildishev, V. M. Shalaev, and A. Boltasseva, "Machine-learning-assisted metasurface design for high-efficiency thermal emitter optimization," Appl. Phys. Rev. **7**(2), 021407 (2020).
36. J. Jiang, D. Sell, S. Hoyer, J. Hickey, J. Yang, and J. A. Fan, "Free-Form Diffractive Metagrating Design Based on Generative Adversarial Networks," ACS Nano **13**(8), 8872–8878 (2019).
37. W. Shin and S. Fan, "Choice of the perfectly matched layer boundary condition for frequency-domain Maxwell's equations solvers," J. Comput. Phys. **231**(8), 3406–3431 (2012).
38. O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, vol. 9351 N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, eds. (Springer International Publishing, Cham, 2015), pp. 234–241. Series Title: Lecture Notes in Computer Science.
39. 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).
40. Y. Li, Y. Wang, S. Qi, Q. Ren, L. Kang, S. D. Campbell, P. L. Werner, and D. H. Werner, "Predicting Scattering From Complex Nano-Structures via Deep Learning," IEEE Access **8**, 139983–139993 (2020).

Optics EXPRESS

41. Z. Wei and X. Chen, "Deep-Learning Schemes for Full-Wave Nonlinear Inverse Scattering Problems," IEEE Trans. Geosci. Remote Sensing **57**(4), 1849–1860 (2019).
42. C. Majorel, C. Girard, A. Arbouet, O. L. Muskens, and P. R. Wiecha, "Deep Learning Enabled Strategies for Modeling of Complex Aperiodic Plasmonic Metasurfaces of Arbitrary Size," ACS Photonics **9**(2), 575–585 (2022).
43. M. P. Bendsøe and N. Kikuchi, "Generating optimal topologies in structural design using a homogenization method," Comput. Methods Appl. Mech. Eng. **71**(2), 197–224 (1988). Publisher: Elsevier BV.
44. M. P. Bendsøe and O. Sigmund, *Topology optimization: theory, methods, and applications* (Springer, Berlin ; New York, 2003).
45. J. S. Jensen and O. Sigmund, "Topology optimization of photonic crystal structures: a high-bandwidth low-loss T-junction waveguide," J. Opt. Soc. Am. B **22**(6), 1191 (2005).
46. C. M. Lalau-Keraly, S. Bhargava, O. D. Miller, and E. Yablonovitch, "Adjoint shape optimization applied to electromagnetic design," Opt. Express **21**(18), 21693 (2013).
47. T. W. Hughes, M. Minkov, I. A. D. Williamson, and S. Fan, "Adjoint Method and Inverse Design for Nonlinear Nanophotonic Devices," ACS Photonics **5**(12), 4781–4787 (2018).
48. M. Zhou, D. Liu, S. W. Belling, H. Cheng, M. A. Kats, S. Fan, M. L. Povinelli, and Z. Yu, "Inverse Design of Metasurfaces Based on Coupled-Mode Theory and Adjoint Optimization," ACS Photonics **8**(8), 2265–2273 (2021).
49. D. Vercruysse, N. V. Sapra, K. Y. Yang, and J. Vuckovic, "Inverse-Designed Photonic Crystal Circuits for Optical Beam Steering," ACS Photonics **8**(10), 3085–3093 (2021).
50. L. Raju, K.-T. Lee, Z. Liu, D. Zhu, M. Zhu, E. Poutrina, A. Urbas, and W. Cai, "Maximized Frequency Doubling through the Inverse Design of Nonlinear Metamaterials," ACS Nano **16**(3), 3926–3933 (2022).
51. Z. A. Kudyshev, A. V. Kildishev, V. M. Shalaev, and A. Boltasseva, "Optimizing Startshot Lightsail Design: A Generative Network-Based Approach," ACS Photonics **9**(1), 190–196 (2022).
52. E. Hassan and A. C. Lesina, "Topology optimization of dispersive plasmonic nanostructures in the time-domain," arXiv, arXiv:2203.01462 (2022) [physics.optics].
53. O. Sigmund and K. Maute, "Topology optimization approaches: A comparative review," Struct. Multidisc. Optim. **48**(6), 1031–1055 (2013).