9

Research Article

Ronald P. Jenkins*, Sawyer D. Campbell and Douglas H. Werner

Establishing exhaustive metasurface robustness against fabrication uncertainties through deep learning

https://doi.org/10.1515/nanoph-2021-0428 Received August 5, 2021; accepted October 24, 2021; published online November 2, 2021

Abstract: Photonic engineered materials have benefitted in recent years from exciting developments in computational electromagnetics and inverse-design tools. However, a commonly encountered issue is that highly performant and structurally complex functional materials found through inverse-design can lose significant performance upon being fabricated. This work introduces a method using deep learning (DL) to exhaustively analyze how structural issues affect the robustness of metasurface supercells, and we show how systems can be designed to guarantee significantly better performance. Moreover, we show that an exhaustive study of structural error is required to make strong guarantees about the performance of engineered materials. The introduction of DL into the inverse-design process makes this problem tractable, enabling optimization runtimes to be measurable in days rather than months and allowing designers to establish exhaustive metasurface robustness guarantees.

Keywords: deep learning; fabrication; robustness; supercell; tolerance.

1 Introduction

Optical metasurfaces are a class of engineered materials (EnMats) that use large numbers of patterned subwavelength structures to modify the electromagnetic

*Corresponding author: Ronald P. Jenkins, Department of Electrical Engineering, The Pennsylvania State University, 121 Electrical Engineering East Building, University Park, PA 16802-1503, USA, E-mail: rpj1@psu.edu. https://orcid.org/0000-0002-4173-5754

Sawyer D. Campbell and Douglas H. Werner, Department of Electrical Engineering, The Pennsylvania State University, 121 Electrical Engineering East Building, University Park, PA 16802-1503, USA. https://orcid.org/0000-0002-3973-2730 (S.D. Campbell)

relationship of two interfacing media. By intelligently designing a metasurface, engineers can functionalize the surface and use it to manipulate light, producing a wide variety of useful phenomena like polarization change, anomalous reflection/refraction, and birefringence among many others [1, 2]. In order to produce such phenomena, the geometrical features of these metasurfaces must be small relative to the operating wavelength λ . For optical metasurfaces, this entails fabricating structures on the nanoscale. Consequently, controlling light with subwavelength structures demands very fine control over the design and fabrication of nanometer-scale features, which presents numerous challenges for the practical realization of metasurface designs.

For several decades, the go-to method for research nanofabrication has been electron beam lithography (EBL) because of its comparatively low cost and high accuracy [3]. EBL uses a focused electron beam to bombard and expose resist materials for chemical etching. Over the years this technique has been refined to the point where researchers can successfully produce structures with feature sizes as small as 2–4 nm using certain resists like hydrogen silsesquioxane (HSQ) and polymethyl methacrylate (PMMA) [4].

In step with the development of improved fabrication techniques, designers have found new and improved methods for discovering useful metasurface geometries. For example, unit cells (also referred to as meta-atoms), which are the building blocks of metasurfaces, had for many years been designed using conventional shapes or patterns (e.g., rectangles and ellipses). However, metasurface designers have found that unconventional unit cells composed of more complex shapes arising organically from the application of genetic algorithms, evolutionary computation, and topology optimization (TO) can lead to significantly higher-performing metasurfaces compared to those based on canonical geometries [5-13]. Unit cell designs often rely on electromagnetic resonances in order to achieve the aforementioned properties, and, unfortunately, that makes them quite sensitive to structural defects.

The structural details of unit cells can be extremely influential in forming their final electromagnetic behavior, so an understanding on the part of designers of how these structural details can change upon fabrication is extremely important. For example, the minimum feature size (MFS), also known as the critical dimension, of a nanostructure is determined by many factors such as depth of focus, temperature, dose, exposure time, etc. Fabrication within $\pm 10\%$ of the MFS is a standard target for quality fabrication [3]. Thus, introducing tolerance analysis into the inversedesign process for these metasurfaces is not only highly desirable but represents an essential tool for the successful creation of robust nanophotonic devices. However, because metasurface designs must be evaluated with fullwave solvers to measure their properties, the computational expense of these solvers is a major hurdle for any kind of exhaustive tolerancing study. As a result, tolerancing is usually neglected in the inverse-design process entirely or only included in the form of a local sensitivity measurement that cannot exhaustively guarantee true performance bounds.

The objective of this article is to provide a method for exhaustively establishing the interaction of unit-cell-based metasurface optical functionality and structural error resulting from manufacturing uncertainty. Our approach to overcoming this challenge is to introduce a deep learning (DL) step into the inverse-design problem. DL has made a significant impact in nanophotonics and propagation in recent years owing to its versatility and unparalleled model generalization [14-17]. Numerous recent contributions have applied DL to metasurfaces specifically, with usecases such as all-dielectric metasurfaces [18-21], chiral metasurfaces [22], diffraction gratings [23, 24], and absorbers [25], among many others. DL has also been applied to metasurfaces in many different ways such as electric field modeling [26], and inverse-design of many kinds of metasurfaces through generative models and latent space engineering [27–30].

In this work, we present a tractable method for exhaustive tolerance modeling and analysis of nanophotonic structures by applying established DL techniques. For our application, we employ a composite deep neural network (DNN) model to predict the transmission diffraction efficiencies of a dielectric supercell from silica to air at a single frequency. Through this approach, we are able to show how metasurface optimizations that perform exhaustive tolerance analysis can provide strong performance guarantees across a wide range of edge deviations. This knowledge and the accompanying DL technique can better inform the development of robust nanophotonic devices.

This paper shows how using our DL augmented optimization method one can increase the guaranteed firstorder diffraction performance of a supercell by more than 35% absolute efficiency over standard optimization methods. We believe this approach will significantly impact the current state-of-the-art in metasurface inversedesign, with application to other functional materials more broadly.

The remainder of the paper is divided into three sections as follows. Section 2 covers the supercell formulation, our chosen method for estimating tolerance, and the DNN model. Section 3 provides the results of the study, including the training of the networks, optimization, edge deviation study, and finally a performance comparison. Section 4 concludes the paper with a discussion of the results, their impact, and a look forward to the next steps.

2 Problem definition and approach

Metasurface performance can degrade when going from design to manufacturing due to inevitable variations in the desired geometry arising from fabrication process imperfections. This problem is especially noticeable with highperformance metasurfaces due to their tendency to rely on complex subwavelength structures.

Some geometric deviations which affect the MFS are partially or entirely beyond the control of the manufacturing process (e.g., edge roughness). On the other hand, flaws such as structure erosion or dilation or undesirable sidewall angles can arise as a result of over/underdosing and/or over/under-etching. These steps in the manufacturing process can be controlled, so they are more systematic in nature. Structures with very tight tolerances can often be fabricated with the necessary precision, but there is an associated cost. This cost can be measured in terms of the process window—the acceptable ranges of manufacturing parameters (i.e., dose, temperature, time, etc.) that will produce a successful product [31]. Structures, where the performance is less sensitive to variations in the process, are desirable because they provide a larger window to the manufacturer [32]. This can reduce costs and increase yield-important aspects in ensuring the successful transition of a product to market and maintaining profitability.

In an effort to accommodate smaller and smaller MFSs for nanophotonic structures, recent work has also focused on analyzing and counteracting sources of error on the fabrication side. Sophisticated process analyses like threedimensional (3D) modeling of line edge roughness through a molecular dynamics simulation or machine-learningassisted analysis of stochastic lithography defects can lead designers toward more realistic modeling of their metadevices [33, 34]. Others have tested the performance consequences of process variation in other varieties of photonic devices or used proximity effect correction to eliminate process-specific defects [35, 36]. Yet another consideration is that of scale. For example, recent work has demonstrated that deep UV immersion lithography can be used to produce an a-Si metasurface at the wafer-scale with an MFS under 100 nm while nevertheless maintaining an MFS error of <10% [37]. The path forward for optical metasurfaces which use complex subwavelength features from EBL to more scalable nanofabrication techniques demands either a relaxation of the MFS or increased tolerance of the aforementioned manufacturing errors.

As discussed earlier, a number of different geometrical deviations can arise during EBL fabrication. When it comes to managing the process window, finding the acceptable ranges of dosing and etching are two major factors. The performance consequences of over/under dosing/etching on nanophotonic structures have been studied for some time now [8-10, 38]. For at least a decade, managing performance sensitivity due to process variation has been explored in the context of TO. Topology optimized structures have enough geometric flexibility that, when successfully optimized, they have the potential to reach higher performance values than conventional unit cell-based supercell structures. However, this high performance is often accompanied by an associated increase in geometrical sensitivity, which is why researchers have sought out methods to improve their stability. As a compromise between achieving some increased tolerance, while maintaining computational efficiency, the current state-ofthe-art in TO sensitivity analysis has been to introduce additional gradients in one form or another. One such approach has been to combine these gradients using a weighted sum. At each iteration, in addition to evaluating the nominal design, two other designs which represent over and underdosing/etching are also evaluated. The gradients computed for each of these are weighted and then combined to form a single gradient. This approach has been shown to be effective for increasing the robustness of photonic crystal waveguides [9], contiguous supercell metasurfaces [10], and nanolenses [39].

Another approach has handled the additional gradients as separate steps in the optimization. In this case, each iteration consists of at least two steps-a local improvement in the nominal performance, and then some local optimization to reduce uncertainty and improve robustness. This method has been used to develop robust photonic waveguide components [40] and 3D photonic crystal band-gap structures [41].

While the above techniques have yielded significant improvements in nanophotonic robustness, the results they produce cannot be said to be exhaustive, nor do they intend to be. Indeed, the necessity of performing a fullwave simulation at each step for each of these methods places a significant limitation on the scope of the robustness with which they can successfully imbue a design.

2.1 Metasurface parameterization

Here, we step away from TO in order to address the question of how to make an exhaustive guarantee of metasurface robustness. Because our designs are not parameterized per pixel, we exchange the TO limitation in exhaustive robustness for a limitation in structural flexibility.

Whereas a traditional TO supercell has the potential to produce entirely contiguous structures, we chose to divide our supercell into discrete unit cells (although our method is flexible enough to handle contiguous supercells). This partitioning helps to preserve the analogy with the more traditional library-sourced supercells in the literature, as developing a library of individual unit cells is necessary to scale metasurfaces up for applications such as metalenses [42-44]. Traditional metasurface designs select an ordering of unit cells with different phase responses to tile into the metasurface. This causes a transmitted wave to be deflected anomalously into a chosen diffraction order m. A metasurface can be made to focus by engineering a phase gradient that changes linearly from its center to the exterior. The core of the focusing metasurface is the development of supercells. Each supercell is composed of several unit cells, with the combination of unit cells providing the phase gradient needed. The diffraction angle is chosen with Equation (1):

$$\theta = \sin^{-1}(m\lambda/d) \tag{1}$$

where d is the supercell width. This applies to all diffraction orders m where $\theta \in \mathbb{R}$.

In [45], the authors demonstrated how supercell construction is complicated by coupling effects between adjacent unit cells which interfere with the supercell's diffraction efficiency. Indeed, constructing a supercell with maximized efficiency from individually optimized unit cells is often nontrivial. In the case of this study, such an approach could also serve to obscure the source of losses between fabrication error and coupling effects. As our goal is to clearly identify losses in efficiency that are directly due

to fabrication error, we simulate and optimize full supercells rather than individual unit cells.

For this work, we chose a freeform geometrical parameterization for our unit cells which has been used successfully in [46]. As shown in Figure 1, unit cell masks are created using a fourfold symmetric spline surface function defined in terms of several control points.

The spline surface function interpolates between the control points, producing a smooth freeform surface that can be manipulated by varying the control point heights. In this work, each unit cell is defined by a 3×3 uniformly spaced grid of control points in one quadrant which is then translated to the other quadrants to enforce fourfold symmetry (although twofold symmetry and completely arbitrary geometries are also trivial to generate with this approach). To produce a mask, the function is rasterized and then a thresholding value is used to create the final binary mask. Each unit cell is given a buffer of pixels

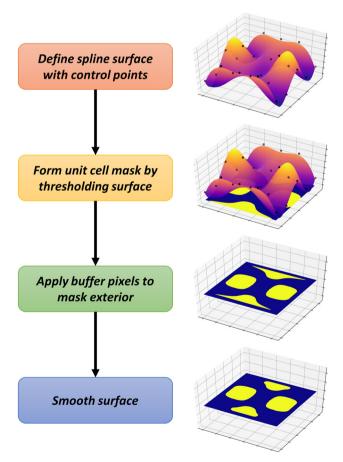


Figure 1: Process for creating unit masks using thresholding of spline surfaces, buffering, and MFS filtering. Each unit cell is assumed to be fourfold symmetric. After constructing these masks, the unit cells are combined together in sets of four to create supercells.

around its exterior to enforce discontinuity between the unit cells. To ensure that unit cell structures have a realistic MFS, each cell is further refined using a series of image processing techniques that enforce the desired MFS. This process is performed for each unit cell before they are combined together to form the final four-element supercell. The spline-based structural parameterization used in this study serves as a proof-of-concept, but the techniques that follow can be extended to individual unit cells, contiguous supercells, and metadevices based on true 3D geometries [47].

For a chosen operational wavelength of $\lambda = 1.55 \,\mu\text{m}$, the supercell width was selected to be $4\lambda/3$ (2067 nm). Unit cells had equal side lengths of $\lambda/3$ (516.7 nm). Silicon was chosen for the patterned layer and it is assumed to be on top of an infinite SiO₂ substrate. The Si pattern height was chosen to be $\lambda/2$ (775 nm). Transmission through the supercell was assumed in the +z direction from the substrate into the air. This arrangement produces three possible diffraction orders, -1, 0, and +1 at angles of -48.6°, 0°, and 48.6°, respectively. Each supercell is evaluated using a rigorous coupled-wave analysis (RCWA) code, and diffraction orders were calculated based on the fields exiting the structure [48]. The structures were rasterized at a resolution of 256 px by 1024 px, leading to a pixel side length of approximately 2 nm.

2.2 Tolerance analysis

There are many ways to model over/under dosing/etching numerically, ranging from highly physics informed to simpler image processing methods [9, 10, 49, 50]. Nevertheless, they all produce similar geometrical erosion or dilation of a structure's mask, also known as edge deviation.

Certain types of designs are less prone to sensitivity than others. However, as will be shown, even designs with relatively relaxed goals can be overly sensitive despite being optimized for robustness in a nonexhaustive manner. Capturing edge deviation with a granularity relevant to a thorough sensitivity analysis produces a demanding resolution requirement. Whereas previous work has carried out inverse-design using only nominal and extreme erosion/dilations, we found that relevant fluctuations in the performance can occur due to very small edge deviation increments between the extrema. Even at the maximum resolution feasible due to our hardware constraints, pixel-by-pixel changes ($\pm 0.0013\lambda$ or ± 2 nm for our problem) could still produce noticeable changes in a design's performance. Combining this level of detail with a

wide range of target edge deviations can cause a dramatic increase in the number of simulations required to test a given design. Beyond the nominal performance and performances at the extreme positive or negative edge deviations, an exhaustive approach must evaluate a significant number of designs in-between.

For a proof-of-concept demonstration, we implemented binary mask erosion and dilation using standard image processing algorithms. For n pixels of erosion or dilation, the masks were iteratively modified using a single-pixel square structuring element n times. Since our goal was to measure changes on a per-pixel level, this method made it easy to track edge deviation length.

2.3 Deep learning

To properly stress test the proposed process, we resolved to test a very wide range of possible edge deviations, out to ± 20 nm ($\pm 0.013\lambda$). At our resolution, this corresponds to a total of 21 variations of the base design which must be simulated per supercell. Despite using a fast RCWA solver to evaluate designs and their variants, the resolution and variation count requirements of this problem present a significant computational challenge, especially when contextualized inside an optimization. A consequence of more complex shape parameterizations than canonical shapes such as rectangles and disks is that, in addition to

the improved performance opportunity [46], they are simultaneously more challenging to optimize due to the increased degrees of freedom. This supercell configuration specifically has 36 degrees of freedom (nine control points per unit cell) which dictates the supercell mask geometry. We found that the optimizer NSGA-II converged after more than 150,000 samples, a substantial number. Indeed, this may be the lower end of how many function evaluations may actually be needed to exhaustively optimize a problem in 36 dimensions such as the aforementioned.

On the one hand, the scale of a supercell optimization where each sample requires 21 different simulations anticipates a total number of simulations in excess of 3,000,000; a number that would be utterly intractable with commercial finite element (FEM) or finite difference time domain (FDTD) solvers. However, the total number of samples needed to train a neural network to approximate a supercell simulation, even with edge deviations, is only a fraction of that amount; on the order of \sim 65,000 simulations. This is the core motivation for introducing DL into this problem.

As has already been discussed, there are many different potential DL models available for estimating electromagnetic phenomena, and even metasurface behavior specifically. This work employs a model similar to others found in the literature and relies specifically on a combination of the U-Net and convolutional neural network (CNN) topologies [51, 52]. Figure 2 shows the full

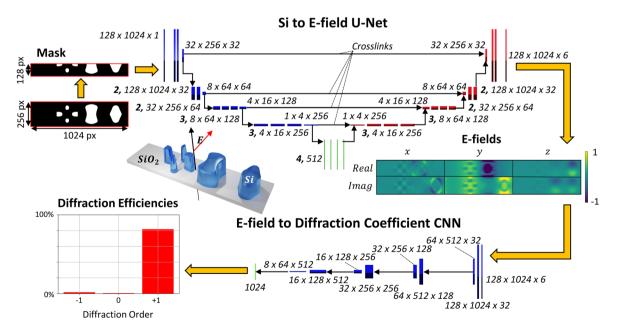


Figure 2: Diagram of network topology, composed of two subnetworks. The first subnetwork (top) converts a binary supercell mask to electric fields via a U-Net. Crosslinks from the U-net help ensure that the geometry of the structure is preserved in the E-fields. The second subnetwork (bottom) converts electric fields to diffraction coefficients using a standard multi-layered CNN. Both subnetworks exploit mirror/anti-mirror *y*-symmetry of the mask and fields to cut the size of the network by half. Each layer or set of layers is labeled to show its shape: *y-pixels* by *x-pixels* by # of features.

structure and details of the model, with subnetworks labeled.

As shown in the figure, the first half of the network (Si to E-field) converts a supercell binary mask to E-fields within the Si pattern and is based on the U-Net architecture. The U-Net architecture shares some structural similarities with an autoencoder in that it compresses an input image down to a latent space (or set of features) before decompressing it back to an image. However, a key innovation of the U-Net is the introduction of crosslinks which copy values from the compression half of the U-Net to the decompression half. Developed for image segmentation and tagging purposes, the U-Net is useful for preserving sharp geometrical features despite the compression process, which makes it ideal for converting a binary mask to E-fields.

The second half of the network (E-fields to Diffraction Coefficients) takes the interior E-fields and predicts diffraction coefficients for the transmitted fields. This subnetwork is a standard CNN that uses a series of convolution layers to extract features from the image before performing a regression on the resulting feature set.

A unique attribute of this work is the use of an intermediate E-field representation that precedes the final prediction of diffraction coefficients. This structure was selected after experimentation with a wide variety of alternate structures due to its comparatively low error. It is worth noting that a very low error in the diffraction coefficient is necessary as computing the efficiency will tend to exaggerate any discrepancies due to its squaring of the coefficient's magnitude.

3 Results

3.1 Data and training

Successfully training a DNN requires the selection of the right terms, which dictate the learning process, known as hyperparameters. While there are some general guidelines available for how to choose these (i.e., learning rate, minibatch size, dataset size, etc.), ultimately the right choice can vary from problem to problem.

For image-based computational electromagnetics DL models, recent work has typically used training sets on the order of 20,000 to 50,000 samples [18-20, 22, 23, 25]. For our problem, we used RCWA to sample a total of 81,408 random supercell designs. 65,024 of these samples were used for training while 16,384 were reserved for validation; an 80:20 training to validation set ratio.

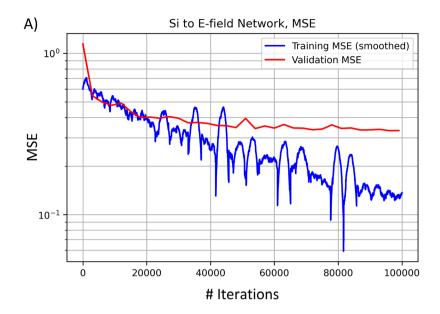
Table 1: Choices are shown for each hyperparameter. These terms dictate the progression of the training process. The same hyperparameters were used for training both halves of the network.

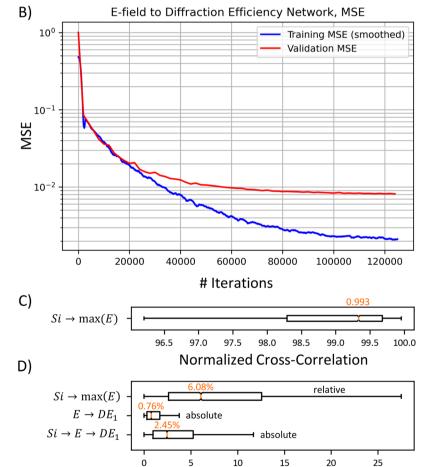
Hyperparameter	Values
Training set size	65,024
Validation set size	16,384
Initial learning rate	1e-4
Learning rate decay	0.9925
Learning rate period	200 iterations
Minibatch size	16

Beyond the 36 variables defining the supercell mask control points, an additional variable representing the erosion/dilation was randomly sampled as well. Each data point was a set of three terms: a mask, E-fields (real and imaginary), and diffraction coefficients (real and imaginary). Because the network is constructed as a pair, both subnetworks needed to be trained separately. The hyperparameters governing the learning can be seen in Table 1. As is typical with DL models, their applicability from one problem to the next is highly sensitive to the proper selection of hyperparameters. These were selected through a process of trial-and-error, a reality of DL development that is itself a point of much research.

Some resources provided critical guidance in the selection of our hyperparameters. For example, for some time now it has been known that DL models can generalize better when trained with smaller minibatch sizes [53]. Owing to the high-resolution requirements of fine-grained tolerancing, the pair of models took up substantial memory during training. Thus, we used a relatively small minibatch size of 16 designs for both the Si to E-field and E-field to Diffraction Coefficient networks. Another way GPU memory overuse was ameliorated was to exploit mirror symmetry in the y-plane of the supercell. This allowed us to cut the data size in half and was achieved by introducing special boundary conditions in the convolution layers which mixed circular and reflection padding.

Figure 3A and B show the training progression of the two networks for both the training and validation datasets. In both cases, the learning rate was initialized to 1e-4 and then made to decay with a scalar of 0.9925 at an interval of 200 iterations. This decay helped both the training and validation mean-squared errors (MSEs) to converge to smaller values at the conclusion of training. Before training, all data were normalized to have a mean of 0 and a variance of 1. The model was developed and trained using the PyTorch library on a machine with four Nvidia RTX





Error (%)

Figure 3: MSE curves show the training of both halves of the network.

(A) Si to E-field U-Net training convergence curves. (B) E-field to diffraction efficiency CNN training convergence curves. MSE for training and validation sets are shown in both figures. Training curves have been smoothed across multiple minibatches to emphasize convergence trends. (C) Normalized cross-correlation of the maximum E-field predicted by the Si to E-field network for the validation dataset, with a median value of 0.993 shown. (D) Percent errors for both the Si to E-field network and E-field to diffraction efficiency network tested with the validation set. For the Si to E-field network, the relative error is reported between prediction and ground truth max E-field yielding a median value of 6.08%. Diffraction efficiency errors are also reported but as an absolute error. The E-field to diffraction efficiency network alone has a median absolute percent error of 0.76%, while the full network chain has a median absolute percent error of 2.45%.

2080 ti graphics cards, each providing 11 GB of RAM for a total of 44 GB of GPU RAM.

To evaluate the accuracy of the networks, we employed several error metrics. Additionally, the networks

had to be tested individually and in series. For the first network which outputs E-fields, we evaluated two error metrics for each design in the validation set: normalized cross-correlation (NCC) and relative percent error. Both metrics were evaluated by comparing the max E-field within the structures. Figure 3C shows the NCC of the validation set E-fields. Since an ideal value of NCC is 1 and any value over 0.8 is considered good, a median NCC of 0.993 indicates that the network is a good predictor. Because the interior E-fields are unbounded, a relative error metric is more helpful in understanding the network's accuracy as opposed to an absolute error metric. Figure 3D shows the validation set's median relative percent error of 6.08%, which is in good agreement with similar networks reported elsewhere in the literature [26]. For the second network, which outputs diffraction efficiencies, we compute an absolute percent error for each design. Figure 3D shows this metric for both the second network by itself and when put in series with the first network. These arrangements both indicate very good agreement, achieving median absolute errors of 0.76% and 2.45%, respectively.

As will be discussed in the next section on optimization, training our network to have a low absolute DE error was crucial to making this solution work. Even under simpler conditions, applying a state-of-the-art evolutionary optimizer to a DNN can often be met with challenges. However, our task was made significantly more difficult than is typical in nanophotonic DL because the high-resolution requirement (necessitated by tolerance analysis) scaled up the network size dramatically. This made training slower and more difficult. Nevertheless, our training regimen was successful in producing a suitable network for use in optimization.

3.2 Optimization and filtering

With the networks trained, we approached the inverse design of exhaustively robust supercells using a multiobjective evolutionary algorithm (MOEA). MOEAs are a category of optimizers designed to characterize the tradeoff between two or more objective functions at the same time [54]. Rather than identifying a single solution as is typical for a single objective optimizer, the goal of a multiobjective optimizer is to produce a family of solutions, called a Pareto set, that represent the best tradeoff between the objectives. Like their single objective counterparts, MOEAs are typically global optimizers. This study relies on this global characteristic to overcome the inherent multimodality which arises when using more complex supercell parameterizations and which challenges TO-based inverse-design. For this work, the NSGA-II optimizer was selected due to its popularity and proven applicability to a wide range of problems [54].

For this study, two cost functions were used to explore the space of possible supercells, summarized in Equation (2). First, the nominal diffraction efficiency DE_0 is used to search for designs that are performant in the base case. Second, a worst-case change is computed by taking the difference of the diffraction efficiencies DE; of all possible variations of the structure with its nominal diffraction efficiency. By optimizing against these two costs, we can clearly show a tradeoff between nominal performance and guaranteed performance under edge deviation.

$$cost_1 = -DE_0$$

$$cost_2 = max_{i>0}|DE_0 - DE_i|$$
(2)

However, combining optimization with DL is not without its challenges. Because DL models can only be used to produce good approximations of an objective function, there are often multiple subdomains of inputs to the neural network which will produce artificially attractive cost values. Therefore, pairing neural networks with optimization can regularly lead to situations where inaccuracies of the network are exploited by the optimizer. Fortunately, there are different methods to overcome this problem. First, neural networks can be integrated into MOEAs to form a surrogate-assisted optimization. This has the advantage of adapting the neural network to better fit the objective function where it matters. Another option the one chosen for this work—is to filter the cost values as they are computed. By comparing the network's prediction with the full-wave solution at just the nominal case, designs which have inherently high error can be filtered from the set. Because there are so many more variations to be tested than just the nominal performance for any given supercell, there is still a substantial speedup to be had from this approach. If a design's nominal performance exceeds some maximum error value (2% in our case), then that cost can be modified according to Equation (3):

$$cost_{1, filt} = -DE_0 + fail * 1e6$$
 (3)

This method exposes and disincentivizes those input subdomains of the objective function which have large MSE in a way that will guide the optimizer away from them and toward well-fit solutions. It is of note that this issue can arise even with networks that have relatively high accuracy (e.g., <5% error in absolute DE). Furthermore, by testing nominal designs and culling those with high error, it is possible to short-circuit the exhaustive test with the DNN on such samples, which can afford further time savings.

3.3 Edge deviation study

Multiobjective optimization of the spline-cut supercells for nominal diffraction efficiency and stability reveals very important robustness characteristics of optimized supercells. Figure 4 summarizes the differences that can result from selecting supercell designs with or without an exhaustive performance guarantee. Within the Pareto front, there is a clear trend where the highest efficiency designs are unstable against perturbations, but lower efficiency designs can be more stable. While this difference would be expected to a certain extent, we can now explicitly quantify this tradeoff. Indeed, as shown in Figure 4E, the improvement in guaranteed performance can be significant, on the order of >35% absolute diffraction efficiency. The degree of the improvement depends on the extent of perturbation considered.

In Figure 4A, all the samples from optimization are plotted with respect to their Nominal and Guaranteed +1 DE, as predicted by the DNN. The Pareto front, which is the optimal tradeoff between these two objectives, is marked with a dashed green line. The extrema for each objective, marked a and c, demonstrate how much difference there is between a design selected for performance-only and robustness, respectively.

Also included is point b, selected as the maximum of a different objective. Figure 4B recontextualizes the samples from the optimization on a new y-axis which shows a weighted sum of the nominal +1 DE with the extreme perturbation (± 20 nm) +1 DEs. While design b does ultimately achieve a higher guaranteed +1 DE than design a, it is clear that there are even more robust solutions like design c.

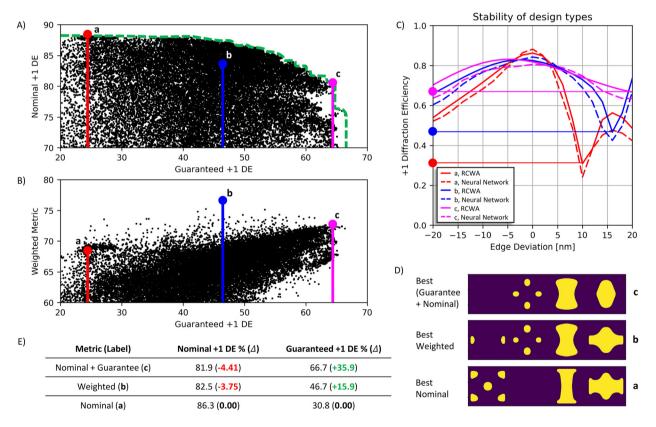


Figure 4: Robustness tradeoff analysis.

(A) Multiobjective study of supercell nominal performance versus loss under ±20 nm of perturbation. The nominal performance of each design is the *y*-axis, and the guaranteed performance is the *x*-axis. The Pareto front is marked with a dashed green line—an ideal design would lie in the upper-right corner of the plot. Three designs marked a, b, and c show the best selection against different metrics: Best nominal, best weighted, and best guarantee + nominal, respectively. (B) Replotting the samples from (A) with a new *y*-axis now showing the weighted metric. (C) Comparison of edge deviation curves for options a, b, and c, including prediction made by DNN and validation from RCWA solver. Paying a loss of 4.5% in absolute diffraction efficiency for the nominal case can lead to an improvement of 35% absolute diffraction efficiency in the guarantee. (D) Supercell masks for the three designs shown. (E) Table comparing the performance difference between the three designs in terms of nominal +1 DE and guaranteed +1 DE. Design found using the neural-network-predicted guarantee (c) achieves a substantial boost in guaranteed performance over a traditionally optimized design (a).

Figure 4C is an edge deviation plot that shows how these three designs compare to each other, as predicted by the DNN and validated with RCWA. For example, with just ±5 nm of perturbation, design a's performance degrades to such an extent as to make it interchangeable with designs b or c in terms of guaranteed performance. Indeed, the full curve showing the falloff trends for the two designs indicates just how much better design c does across the board, culminating in a >35% absolute performance gain over at 10 nm of perturbation. All that for <4.5% loss in nominal performance makes design c a potential preference in terms of overall performance.

This study affirms what has been shown in other studies, that a naïve single objective optimization of nominal performance will not only yield unstable results but significantly worse guaranteed performance as well. However, this work expands upon this result by showing, for the first time, that establishing exhaustive metasurface robustness requires fine-grained sampling of variations. Considering a weighted metric of the nominal case and extreme perturbations provides some advantage over purely nominal-performance focused optimization. However, failures in the edge deviation curve can occur at scales missed by this kind of analysis. Therefore, designers who want to guarantee performance over a range of edge deviations can take advantage of the proposed DL-augmented technique to affect optimizations that achieve exhaustive tolerancing.

3.4 Speedup performance summary

A thorough analysis of the performance of this approach is presented in Table 2. The expected time that would be taken

by a purely RCWA version of our study is listed in one column, with timing for our DL augmented study shown next to it. A speedup column further clarifies the improvement we get over just RCWA.

In this analysis, we include the initial time investment required to collect data and train the networks as part of the total time attributed to the DL augmented approach. This extra time is often not included nor discussed in other DL papers involving optimization in nanophotonics, but it makes a significant difference in realistically understanding the overall utility of the technique. For example, the speedup for a single optimization without accounting for the initial time investment is 14.8 times faster than RCWA alone. However, when the real-time investment is included, this amount is reduced to a 4.37 times speedup.

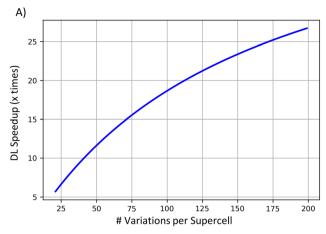
In reality, oftentimes a single optimization is not enough to solve a problem. In the course of collecting results for the studies in this paper, for example, we ran no fewer than 6 DL augmented optimizations. From this point of view, our full study's actual speedup was 10.6 times over a purely RCWA-based study. We envision this technique providing a more powerful platform for studying robustness in general than a purely full-wave approach.

Additionally, the speedup reported is the lower bound of what is achievable using our technique, even including the startup time cost. Our choice of RCWA represents a best-case choice for full-wave solver speed. If this solver were replaced with FEM or FDTD techniques, the speedup would significantly increase.

This DL augmented approach to metasurface design opens the doors to more complex studies in the future. As shown in Figure 5, further increases to optimization count or a number of variation tests per supercell make the problem even harder to solve, which increasingly favors

Table 2: Performance comparison of robustness study with and without DL augmentation. The columns show timings for different parts of the study. The field description column explains what each step is. The RCWA and DL columns each record thetime taken for the different approaches at each step. Finally, the speedup column reports the improvement in time taken by using the DL method for each step. Even when factoring in training time, the DL augmented approach affords a substantial speedup over the standard method. Moreover, as shown in rows G and H, the performance advantage DL affords increases substantially with further optimizations and increases in the number of variations tested for each design.

	Field description	RCWA	DL	Speedup
A	Single sample time (serialized)	1.61 s	0.034 s	47.4
В	Data collection sample time (includes internal fields)	N/A	3.35 s	N/A
C	Data collection total time	N/A	3.16 days = 81,408 samples * B[DL]	N/A
D	Training time	N/A	3.15 days	N/A
Ε	Full supercell evaluation (including variants)	33.9 s = (21 * A[RCWA])	2.29 s = (A[RCWA] + 20 * A[DL])	14.8
F	Optimization (>100,000 function evaluations)	>5.60 weeks	>2.65 days	14.8
G	Total time (one optimization)	>5.60 weeks = F[RCWA]	>8.96 days = C[DL] + D[DL] + F[DL]	4.37
Н	Total time (six optimizations)	>8 months = 6 * F[RCWA]	>3 weeks = C[DL] + D[DL] + 6 * F[DL]	10.60



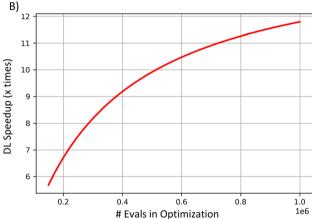


Figure 5: Speedup curves using DL augmentation versus pure fullwave solvers.

(A) Speedup when increasing the number of variation tests required by each supercell. (B) Speedup when increasing the number of total evaluations performed in an optimization.

our DL method. With finer resolved structures and/or additional tolerance considerations (e.g., edge roughness and sidewall angle), the number of variation tests required per unit cell would increase combinatorically, offering speedups on the order of 20–30 times over a purely fullwave approach.

4 Conclusions

We have successfully implemented an approach for establishing exhaustive metasurface performance guarantees which uses DL to overcome the fundamental computational challenges associated with rigorous tolerancing of nanoscale structurally functionalized materials with specific application to metasurface design. Ensuring that a metasurface is robust to edge deviation is a challenging problem because of the fine-grained perturbation

study required to assess each design. The performance of these structures can be sensitive to variation even at the finest scale of EBL edge deviations (e.g., 2 nm), which in turn demands a consummately small sampling period to avoid aliasing over drops in performance.

This challenge is exacerbated by optimization. Performance at the nominal and extreme values of edge deviation has a higher upper bound when the performance between them is allowed to dip. Thus, to find designs that can make firm guarantees about performance within a range of edge deviations, many samples must be taken for each design.

By frontloading the problem (i.e., sampling for a dataset and training a model), it is possible to nevertheless perform the necessary amount of sampling in order to make these guarantees and thereby render optimizing for guaranteed performance viable. Training a model to achieve a low MSE is also critical to making the optimization successful, as high model error can be a major limitation for optimizing using DL.

We showed that by including guaranteed performance as an objective in an MOEA optimization, it is possible to extract a Pareto front that characterizes the tradeoff between designs with the highest nominal performance and those with the best guarantee. We demonstrate that an exhaustively robust design selected using guaranteed performance as part of the metric yielded a substantial increase of more than 35% in the guaranteed absolute performance over a design selected only for having the best nominal performance. Moreover, this robust design suffered a minimal loss in nominal performance of <4.5%.

Compared to a purely full-wave study, we achieved a more than 10 times speedup, bringing the study down from what would be more than an eight-month effort to a little over three weeks. We believe this development provides a path forward for improving the state of the art in engineered nanomaterial robustness, and we anticipate that the method will be adapted beyond the scope of this manuscript to other nanophotonic devices and EnMats as well.

4.1 Future work

Many possible next steps exist to extend this work further. To overcome the necessity to filter designs during optimization, the learning process could be integrated directly into the optimizer. As a now surrogate-assisted optimizer, the model may converge to even higher performances than previously discovered. Other improvements might include developing a smaller DNN model, applying a more physically realistic erosion/dilation model to the data generation,

extending the complexity of the supercell parameterization to include more control points per unit cell, and training the model to generalize across more types of fabrication errors (edge roughness, sidewall angle, etc.).

Author contributions: All the authors have accepted responsibility for the entire content of this submitted manuscript and approved submission.

Research funding: This research was supported in part by DARPA EXTREME (contract HR00111720032).

Conflict of interest statement: The authors declare no conflicts of interest regarding this article.

References

- [1] H. T. Chen, A. J. Taylor, and N. Yu, "A review of metasurfaces: physics and applications," Rep. Prog. Phys., vol. 79, p. 076401, 2016.
- [2] D. H. Werner, S. D. Campbell, and L. Kang, Nanoantennas and Plasmonics Modelling, Design and Fabrication, Raleigh, NC, USA, SciTech Publishing Inc., 2020.
- [3] Z. Cui, Nanofabrication, Basel, Switzerland, Springer International Publishing, 2017.
- [4] Y. Chen, "Nanofabrication by electron beam lithography and its applications: a review," Microelectron. Eng., vol. 135, pp. 57-72,
- [5] S. D. Campbell, D. Sell, R. P. Jenkins, E. B. Whiting, J. A. Fan, and D. H. Werner, "Review of numerical optimization techniques for metadevice design [Invited]," Opt. Mater. Express, vol. 9, pp. 1842-1863, 2019.
- [6] Y. Chen, S. Zhou, and Q. Li, "Multiobjective topology optimization for finite periodic structures," Comput. Struct., vol. 88, pp. 806-811, 2010.
- [7] A. R. Diaz and O. Sigmund, "A topology optimization method for design of negative permeability metamaterials," Struct. Multidiscip. Optim., vol. 41, pp. 163-177, 2010.
- [8] H. W. Dong, Y. S. Wang, T. X. Ma, and X. X. Su, "Topology optimization of simultaneous photonic and phononic bandgaps and highly effective phoxonic cavity," JOSA B, vol. 31, pp. 2946-2955, 2014.
- [9] F. Wang, J. S. Jensen, and O. Sigmund, "Robust topology optimization of photonic crystal waveguides with tailored dispersion properties," JOSA B, vol. 28, pp. 387-397, 2011.
- [10] E. W. Wang, D. Sell, T. Phan, and J. A. Fan, "Robust design of topology-optimized metasurfaces," Opt. Mater. Express, vol. 9, pp. 469-482, 2019.
- [11] G. Yi and B. D. A. Youn, "Comprehensive survey on topology optimization of phononic crystals," Struct. Multidiscip. Optim., vol. 54, pp. 1315-1344, 2016.
- [12] J. A. Fan, "Freeform metasurface design based on topology optimization," MRS Bull., vol. 45, pp. 196-201, 2020.
- [13] M. Zhou, B. S. Lazarov, and O. Sigmund, "Topology optimization for optical projection lithography with manufacturing uncertainties," Appl. Opt., vol. 53, pp. 2720-2729, 2014.
- [14] S. D. Campbell, R. P. Jenkins, P. J. O'Connor, and D. H. Werner, "The explosion of artificial intelligence in antennas and propagation: how deep learning is advancing our state of the art," IEEE Antenn. Propag. Mag., vol. 63, pp. 16-27, 2020.

- [15] A. Massa, D. Marcantonio, X. Chen, M. Li, and M. Salucci, "DNNs as applied to electromagnetics, antennas, and propagation-a review," IEEE Antenn. Wireless Propag. Lett., vol. 18, pp. 2225-2229, 2019.
- [16] O. Khatib, S. Ren, J. Malof, and W. J. Padilla, "Deep learning the electromagnetic properties of metamaterials—a comprehensive review," Adv. Funct. Mater., vol. 2021, p. 2101748.
- [17] W. Ma, A. Liu, Z. A. Kudyshev, A. Boltasseva, W. Cai, and Y. Liu, "Deep learning for the design of photonic structures," Nat. Photonics, vol. 15, pp. 77-90, 2021.
- [18] S. An, C. Fowler, B. Zheng, et al., "A deep learning approach for objective-driven all-dielectric metasurface design," ACS Photonics, vol. 6, pp. 3196-3207, 2019.
- [19] S. An, B. Zheng, M. Y. Shalaginov, et al., "A freeform dielectric metasurface modeling approach based on deep neural networks," ArXiv Prepr, 2020, ArXiv 200100121.
- [20] S. An, B. Zheng, M. Y. Shalaginov, et al., "Deep learning modeling approach for metasurfaces with high degrees of freedom," Opt. Express, vol. 28, p. 31932, 2020.
- [21] C. C. Nadell, B. Huang, J. M. Malof, and W. J. Padilla, "Deep learning for accelerated all-dielectric metasurface design," Opt. Express, vol. 27, pp. 27523-27535, 2019.
- [22] W. Ma, F. Cheng, and Y. Liu, "Deep-learning-enabled on-demand design of chiral metamaterials," ACS Nano, vol. 12, pp. 6326-6334, 2018.
- [23] S. Inampudi and H. Mosallaei, "Neural network based design of metagratings," Appl. Phys. Lett., vol. 112, p. 241102, 2018.
- [24] J. Jiang and J. A. Fan, "Global optimization of dielectric metasurfaces using a physics-driven neural network," Nano Lett., vol. 19, pp. 5366-5372, 2019.
- [25] T. Qiu, X. Shi, J. Wang, et al., "Deep learning: a rapid and efficient route to automatic metasurface design," Adv. Sci., vol. 6, p. 1900128, 2019.
- [26] 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., vol. 20, pp. 329-338, 2020.
- [27] W. Ma, F. Cheng, Y. Xu, Q. Wen, and Y. Liu, "Probabilistic representation and inverse design of metamaterials based on a deep generative model with semi-supervised learning strategy," Adv. Mater., vol. 31, p. 1901111, 2019.
- [28] Z. A. Kudyshev, A. V. Kildishev, V. M. Shalaev, and A. Boltasseva, "Machine-learning-assisted metasurface design for highefficiency thermal emitter optimization," Appl. Phys. Rev., vol. 7, p. 021407, 2020.
- [29] D. Zhu, Z. Liu, L. Raju, A. S. Kim, and W. Cai, "Building multifunctional metasystems via algorithmic construction," ACS Nano, vol. 15, pp. 2318-2326, 2021.
- [30] 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., vol. 63, p. 284212, 2020.
- [31] K. Keil, K. H. Choi, C. Hohle, et al., "Determination of best focus and optimum dose for variable shaped e-beam systems by applying the isofocal dose method," Microelectron. Eng., vol. 85, pp. 778-781, 2008.
- [32] J. W. Bossung, "Projection printing characterization," in Developments in Semiconductor Microlithography II, J. W. Giffin, Ed., vol. 0100, International Society for Optics and Photonics, 1977, pp. 80-85.

- [33] S. Pinge, Y. Qiu, V. Monreal, D. Baskaran, A. Ravirajan, and Y. Lak Joo, "Three-dimensional line edge roughness in pre- and post-dry etch line and space patterns of block copolymer lithography," Phys. Chem. Chem. Phys., vol. 22, pp. 478-488, 2020.
- [34] K. Azumagawa and T. Kozawa, "Application of machine learning to stochastic effect analysis of chemically amplified resists used for extreme ultraviolet lithography," Jpn. J. Appl. Phys., vol. 60, p. SCCC02, 2021.
- [35] X. Mu, Z. Chen, L. Cheng, et al., "Effects of fabrication deviations and fiber misalignments on a fork-shape edge coupler based on subwavelength gratings," Opt. Commun., vol. 482, p. 126562, 2021.
- [36] M. Eissa, T. Mitarai, T. Amemiya, Y. Miyamoto, and N. Nishiyama, "Fabrication of Si photonic waveguides by electron beam lithography using improved proximity effect correction," *Ipn*. J. Appl. Phys., vol. 59, p. 126502, 2020.
- [37] T. Hu, C. K. Tseng, Y. H. Fu, et al., "Demonstration of color display metasurfaces via immersion lithography on a 12-inch silicon wafer," Opt. Express, vol. 26, pp. 19548-19554, 2018.
- [38] F. B. Arango, R. Thijssen, B. Brenny, T. Coenen, and A. F. Koenderink, "Robustness of plasmon phased array nanoantennas to disorder," Sci. Rep., vol. 5, pp. 1-9, 2015.
- [39] Y. Augenstein and C. Rockstuhl, "Inverse design of nanophotonic devices with structural integrity," ACS Photonics, vol. 7, pp. 2190-2196, 2020.
- [40] N. Lebbe, C. Dapogny, E. Oudet, K. Hassan, and A. Gliere, "Robust shape and topology optimization of nanophotonic devices using the level set method," J. Comput. Phys., vol. 395, pp. 710-746, 2019.
- [41] H. Men, K. Y. K. Lee, R. M. Freund, J. Peraire, and S. G. Johnson, "Robust topology optimization of three-dimensional photoniccrystal band-gap structures," Opt. Express, vol. 22, pp. 22632-22648, 2014.
- [42] M. Khorasaninejad, Z. Shi, A. Y. Zhu, et al., "Achromatic metalens over 60 nm bandwidth in the visible and metalens with reverse chromatic dispersion," Nano Lett., vol. 17, pp. 1819-1824, 2017.
- [43] W. T. Chen, A. Y. Zhu, V. Sanjeev, et al., "A broadband achromatic metalens for focusing and imaging in the visible," Nat. Nanotechnol., vol. 13, pp. 220-226, 2018.

- [44] J. Nagar, S. D. Campbell, and D. H. Werner, "Apochromatic singlets enabled by metasurface-augmented GRIN lenses," Optica, vol. 5, pp. 99-102, 2018.
- [45] S. An, B. Zheng, M. Y. Shalaginov, et al., "Deep Convolutional Neural Networks to Predict Mutual Coupling Effects in Metasurfaces," Arxiv 2021, arXiv 210201761.
- [46] E. B. Whiting, S. D. Campbell, L. Kang, and D. H. Werner, "Metaatom library generation via an efficient multi-objective shape optimization method," Opt. Express, vol. 28, pp. 24229-24242, 2020.
- [47] D. Z. Zhu, E. B. Whiting, S. D. Campbell, D. B. Burckel, and D. H. Werner, "Optimal high efficiency 3D plasmonic metasurface elements revealed by lazy ants," ACS Photonics, vol. 6, pp. 2741-2748, 2019.
- [48] W. Jin, W. Li, M. Orenstein, and S. Fan, "Inverse design of lightweight broadband reflector for relativistic lightsail propulsion," ACS Photonics, vol. 7, pp. 2350-2355, 2020.
- [49] R. J. Hawryluk, "Exposure and development models used in electron beam lithography," J. Vac. Sci. Technol., vol. 19, pp. 1-17, 1981.
- [50] P. Hudek and D. Beyer, "Exposure optimization in highresolution e-beam lithography," Microelectron. Eng., vol. 83, pp. 780-783, 2006.
- [51] 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, Lecture Notes in Computer Science, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds., Switzerland, Springer International Publishing, 2015, pp. 234-241.
- [52] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," Commun. ACM, vol. 60, pp. 84-90, 2017.
- [53] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang, "On large-batch training for deep learning: generalization gap and sharp minima," ArXiv Cs Math, 2017, ArXiv 160904836.
- [54] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," IEEE Trans. Evol. Comput., vol. 6, pp. 182-197, 2002.