

Research article

Jiaqi Jiang and Jonathan A. Fan*

Multiobjective and categorical global optimization of photonic structures based on ResNet generative neural networks

<https://doi.org/10.1515/nanoph-2020-0407>

Received July 20, 2020; accepted September 7, 2020;

published online September 22, 2020

Abstract: We show that deep generative neural networks, based on global optimization networks (GLOnets), can be configured to perform the multiobjective and categorical global optimization of photonic devices. A residual network scheme enables GLOnets to evolve from a deep architecture, which is required to properly search the full design space early in the optimization process, to a shallow network that generates a narrow distribution of globally optimal devices. As a proof-of-concept demonstration, we adapt our method to design thin-film stacks consisting of multiple material types. Benchmarks with known globally optimized antireflection structures indicate that GLOnets can find the global optimum with orders of magnitude faster speeds compared to conventional algorithms. We also demonstrate the utility of our method in complex design tasks with its application to incandescent light filters. These results indicate that advanced concepts in deep learning can push the capabilities of inverse design algorithms for photonics.

Keywords: categorical optimization; global optimization; multiobjective optimization; neural networks; thin-film stack.

1 Introduction

Inverse algorithms are among the most effective methods for designing efficient, multifunctional photonic devices [1–3]. It remains an open question how to select and

implement a design algorithm, and over the last few years, much research has been focused on deep neural networks as inverse design tools [4–6]. Many of these demonstrations are based on the generation of a training set, consisting of device geometries and their optical responses, and modeling these data using discriminative [7, 8] or generative [9–12] neural networks. These methods have proven to be capable of producing high-speed surrogate solvers and can perform inference-type tasks with training data. When the training data are curated using advanced gradient-based optimization methods, such as the adjoint variables [13–17] or objective-first methods [18], the networks can learn to generate high-performing, free-form photonic structures.

To perform global optimization, alternative approaches are required that do not depend on interpolation from a training set. The reason is because the design space is nonconvex and contains multiple local optima, and even devices based on advanced gradient-based optimization methods cannot help a neural network search for the global optimum. In this vein, global optimization networks (GLOnets) have been developed to perform the nonconvex global optimization of free-form photonic devices [19, 20]. GLOnets are gradient-based optimizers that do not use a training set but instead combine a generative neural network with an electromagnetic simulator to perform population-based optimization. The evolution of the generated device distribution is driven by both figure-of-merit values (i.e., efficiencies) and gradients for devices sampled from the generative network. Initial implementations of GLOnets were configured for single-objective problems with binary design variables, such as the maximization of deflection efficiency for a normally incident beam in a metagrating comprising silicon nanostructures. “Single-objective” refers to the optimization of a system operating with one conditional parameter, in this case a system with fixed incidence beam angle, and “binary” refers to silicon and air as our design materials.

A more general formulation of the problem that captures the design space of many photonic technologies is

*Corresponding author: Jonathan A. Fan, Department of Electrical Engineering, Stanford University, 348 Via Pueblo, Stanford, CA 94305, USA, E-mail: jonfan@stanford.edu

Jiaqi Jiang, Department of Electrical Engineering, Stanford University, 348 Via Pueblo, Stanford, CA 94305, USA. <https://orcid.org/0000-0001-7502-0872>

multiobjective, categorical optimization with more than two design materials. “Multiobjective” refers to the optimization of a system operating involving more than one objective function to be optimized simultaneously, such as a metagrating operating over a range of incident beam angles, and “categorical” refers to design variables that have two or more categories without intrinsic ordering, such as multiple material types. In this study, we show that GLOnets can be configured as a multiobjective, categorical global optimizer, and we adapt GLOnets to optimize thin-film stacks to demonstrate the capabilities of our algorithms. Thin-film stacks are an ideal model system for multiple reasons. First, the design problem is multi-objective as devices are typically configured for a range of incident wavelengths, angles, and polarizations. Second, the design problem is categorical as individual layer materials are chosen from a library of materials. Third, thin-film stacks are a well-established technology, and there are a number of pre-existing studies that enable proper benchmarking of algorithm performance [21–23].

Thin-film stacks have been widely used in many optical systems including passive radiative coolers [24], efficient solar cells [25, 26], broadband spectral filtering [27, 28], thermal emitters [29], and spatial multiplexing filters [30]. The materials and thicknesses of thin-film layers have to be carefully optimized to achieve the desired transmission and reflection properties across a broad wavelength and angular bandwidth. Design methods based on physical intuition result in limited performance, and they are generally difficult to scale to aperiodic thin-film stacks comprising many layers. To address these limitations, various global optimization approaches have been explored, including the Monte Carlo approach [31], particle swarm optimization [32], needle optimization [33–35], and the memetic algorithm [21]. These methods are all derivative-free global optimization algorithms that search the design space through the evaluation of a batch of samples without any gradient calculations, limiting their ability to reliably solve for the global optimum.

2 Method

We consider the design of N -layer thin-film stacks each comprising an isotropic material specified from a material library (Figure 1). The refractive indices of the total stack are denoted as a vector $\mathbf{n}(\lambda) = (n_1(\lambda), n_2(\lambda), \dots, n_N(\lambda))^T$, where each index term is a function of wavelength to account for dispersion, and the values can be real or complex valued without loss of generality. The thin-film stack thicknesses are $\mathbf{t} = (t_1, t_2, \dots, t_N)^T$. The material library consists of M material types, and their refractive indices are represented as $\{m_1(\lambda), m_2(\lambda), \dots, m_M(\lambda)\}$.




Figure 1: Schematic of the N -layer thin-film stack system. The refractive index and thickness of each layer are optimized to produce a desired reflection profile, and the composition of each layer is constrained to index values specified in a material library.

The optimization problem is posed as finding the proper \mathbf{n} and \mathbf{t} that produces the desired reflection characteristics over a given wavelength bandwidth, incident angle range, and incident polarization:

$$\{\mathbf{n}^*, \mathbf{t}^*\} = \arg \min_{\{\mathbf{n}, \mathbf{t}\}} \sum_{\lambda, \theta, \text{pol}} (\mathcal{R}(\mathbf{n}, \mathbf{t} | \lambda, \theta, \text{pol}) - \mathcal{R}^*(\lambda, \theta, \text{pol}))^2 \quad (1)$$

The desired reflection spectrum is denoted as $\mathcal{R}^*(\lambda, \theta, \text{pol})$, and $\{\mathbf{n}^*, \mathbf{t}^*\}$ are the corresponding global optimal refractive indices and thicknesses. This optimization problem can be readily cast as the minimization of the objective function: $O(\mathbf{n}, \mathbf{t}) = \sum_{\lambda, \theta, \text{pol}} (\mathcal{R}(\mathbf{n}, \mathbf{t} | \lambda, \theta, \text{pol}) - \mathcal{R}^*(\lambda, \theta, \text{pol}))^2$. \mathbf{n} are categorical variables because the index values are chosen from a material database, while \mathbf{t} can span a continuous set of values and is a continuous variable.

2.1 Transfer matrix method solver

A principle requirement of any gradient-based optimizer is a method to calculate local gradients. For thin-film stacks, these gradients indicate how perturbations to the refractive indices and thicknesses of the device can best reduce the objective function. In prior implementations of GLOnets, local gradients were calculated using the adjoint variable method, in which forward and adjoint simulations are calculated using a conventional electromagnetic solver [19, 20].

While the adjoint variables method provides a general formalism to calculating local gradients using any conventional solver, we pursue an alternative approach based on the transfer matrix method (TMM), which is a fully analytic and high-speed solver for thin-film systems. In particular, we program a TMM solver within the automatic differentiation framework in PyTorch [36], which allows gradients to be directly calculated using the chain rule. Automatic differentiation is the basis for calculating gradients during backpropagation in neural network training, and it generally applies to any algorithm that can be described by a differentiable computational graph. Recently, it was

implemented in finite-difference time domain (FDTD) and finite-difference frequency domain (FDFD) simulators [37, 38]. Compared to generalized differentiable electromagnetic solvers, such as these FDTD and FDFD implementations, our analytic TMM-based algorithms are faster without loss of accuracy because the thin films are described as layers instead of voxels.

2.2 Res-GLOnet algorithm

A schematic of GLOnets configured for our thin-film stack system is outlined in Figure 2a. We term this GLOnet variant as Res-GLOnets because the generator has a residual network architecture that includes skip connections between layers (blue box inset), which will be discussed in a later section. First, a generative neural network G with trainable weights ϕ produces a distribution of thin-film stack configurations. The input to the generator is a uniformly distributed random vector $\mathbf{z} \sim U(0,1)$, so that the generator can be regarded as a function that maps the uniform distribution to a complex distribution of thin-film stack configurations, $G_\phi: U(0,1) \rightarrow P_\phi(\mathbf{n}, \mathbf{t})$. Different samplings of the input random variable $\mathbf{z}^{(k)}$ map onto different device refractive index and thickness configurations within $P_\phi(\mathbf{n}, \mathbf{t})$, denoted as $\{\mathbf{n}^{(k)}, \mathbf{t}^{(k)}\} = G_\phi(\mathbf{z}^{(k)})$. The generated \mathbf{n} from the network do not take categorical values from the material library but are relaxed to be continuous variables, to stabilize the optimization process. These \mathbf{n} are further processed using a probability matrix to enforce the categorical value constraint, which is discussed in the next section. After processing, the reflection

spectra of the generated devices, $\mathcal{R}(\mathbf{n}^{(k)}, \mathbf{t}^{(k)} | \lambda, \theta, \text{pol})$, are calculated using the TMM solver.

The optimization objective, or the loss function, for GLOnet is defined as:

$$L = \mathbb{E} \left[\exp \left(- \frac{O(\mathbf{n}, \mathbf{t})}{\sigma} \right) \right] \quad (2)$$

$$= \int \exp \left(- \frac{O(\mathbf{n}, \mathbf{t})}{\sigma} \right) P_\phi(\mathbf{n}, \mathbf{t}) d\mathbf{n} d\mathbf{t} \quad (3)$$

$$= \int \exp \left(- \frac{O(G_\phi(\mathbf{z}))}{\sigma} \right) P(\mathbf{z}) d\mathbf{z} \quad (4)$$

$$\approx \sum_{k=1}^K \exp \left(- \frac{O(\mathbf{n}^{(k)}, \mathbf{t}^{(k)})}{\sigma} \right) \quad (5)$$

σ is a hyperparameter. These equations follow the derivation of the GLOnet formalism described in the study by Jiang and Fan [20]. To train the generative network and update its weights in a manner that improves the mapping of \mathbf{z} to devices, the gradient of the loss function with respect to the neuron weights, $\nabla_\phi L$, is calculated by backpropagation.

A schematic of the evolution of the generative network over the course of network training is outlined in Figure 2b. Initially, the generator has no knowledge about the design space and outputs a broad distribution of devices spanning the full design space. Over the course of network training, the distribution of generated devices narrows and gets biased toward design space regions that feature relatively small objective function values. Upon the completion of network training, the distribution of generated thin-film stack configurations converges to a narrow distribution centered around the global optimum.

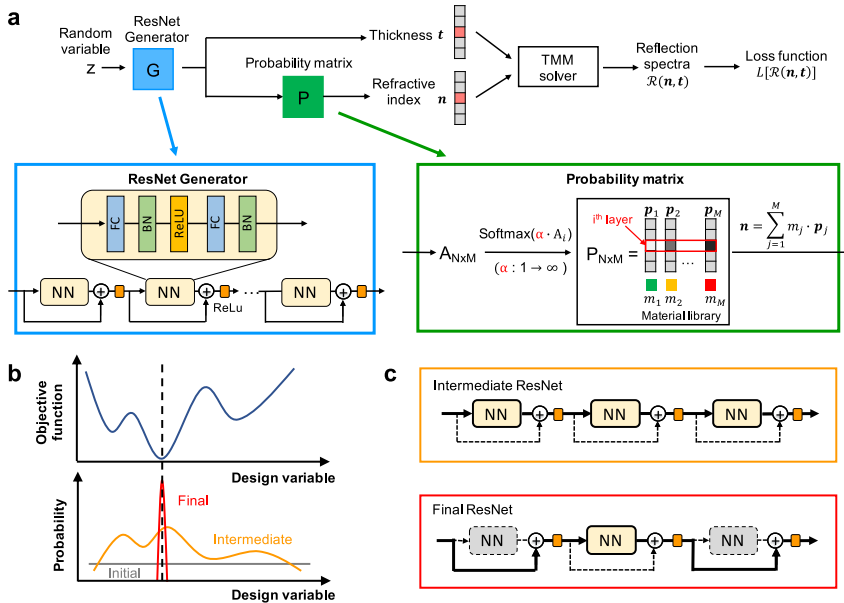


Figure 2: Thin-film global optimization with the Res-GLOnet. (a) Schematic of the Res-GLOnet. A ResNet generator maps a uniformly distributed random variable to a distribution of devices, which are then evaluated using a transfer matrix method solver and used to evaluate the loss function. A probability matrix pushes the continuous generated device indices \mathbf{n} to discrete values. (b) Evolution of the generated device distribution over the course of network training. The network initially samples the full design space and converges to a narrow distribution centered around the global minimum of the objective function. (c) During training, the network operates as a deep architecture with little impact from the skip connections (Intermediate ResNet). Near training completion, the network evolves to a shallow architecture with large impact from the skip connections (Final ResNet). Bold and dashed lines indicate large and small contributions to the network architecture, respectively. TMM, transfer matrix method; GLOnet, global optimization network.

2.3 Enforcing categorical constraints

To update the weights in the generative network during back-propagation, the chain rule is applied to the entire computation graph of the Res-GLOnet algorithm. One required step is the calculation of the gradient of the reflection spectrum with respect to the refractive indices, $\frac{dR}{dn}$. If the refractive indices of thin-film stacks outputted by the generator are directly treated as categorical variables, \mathbf{n} is not a continuous function and the gradient term above cannot be calculated.

To overcome this difficulty, we propose a reparameterization scheme in which the generated \mathbf{n} are relaxed to take continuous values and are then processed in a manner that supports convergence to categorical variable values. The concept is outlined in the green box inset in Figure 2a. The network first maps the random vector \mathbf{z} onto an N -by- M matrix A . These values can vary continuously and take any real number value. A softmax function is then applied to each row of A to generate a probability matrix P :

$$P_{ij} = \frac{\exp(\alpha \cdot A_{ij})}{\sum_{j=1}^M \exp(\alpha \cdot A_{ij})} \quad (6)$$

The i th row of matrix P is a $1 \times M$ vector and represents the probability distribution that the i th thin-film layer takes on a particular material choice within the material library. We use the softmax function because it produces a properly normalized probability distribution and is commonly used in other related tasks, such as classification tasks [39]. The expected refractive index of the i th layer given by this distribution, calculated as $n_i(\lambda) = \sum_{j=1}^M m_j(\lambda) \cdot P_{ij}$, is used to define the thin-film stack in subsequent TMM calculations in Res-GLOnet. All functions in this algorithm can be expanded into a differentiable computational graph, meaning that the loss function gradient with respect to the refractive index is able to backpropagate through the probability matrix P and to the network weights ϕ .

α is a hyperparameter that tunes the sharpness of the softmax function. Initially, α is set to be one, and the expected refractive index of the i th layer has contributions from many different materials in the material library. Over the course of network training, α is linearly increased as a function of the training iteration number until the probability distribution of the i th thin-film layer is effectively a delta function that has converged to a single material. These concepts build on a similar scheme previously used for image sensor multiplexing design [40].

2.4 ResNet generator

Our optimization problem involves searching within a highly complex, nonconvex design space and is made particularly challenging by device requirements spanning a wide range of incident wavelengths and angles. In the early and intermediate stages of network training, a deep neural network is required to properly generate a complex distribution of devices spanning large regions of the design space. However, toward the latter stages of network training, the distribution of the generated devices should ideally converge to a simple and narrow distribution centered around the global optimum, which is more ideally modeled using a shallow network. GLOnet schemes that train using a fixed network architecture do not have the flexibility to capture these trends: deep architectures have general difficulty in training owing to the well-known vanishing gradient problem, while

shallow architectures have the issue of underfitting the design space and are ineffective during the early and intermediate stages of network training [41].

To address these issues, we utilize deep residual networks for the generator architecture, which reformulates our algorithm as Res-GLOnets. Residual networks [41] were developed in the computer vision community to stably process images in very deep networks and overcome the vanishing gradient problem, with the insight that the use of skip connections can enable the depth of the network to be effectively and implicitly tuned over the course of training. A schematic of our Res-GLOnet architecture is shown in the blue box inset in Figure 2a and comprises a series of 16 residual blocks. Each block contains a fully connected layer, a batch normalization layer, and a leaky ReLU nonlinear activation layer. The input x_{in} and output x_{out} of each residual block have the same dimension, and the output of each block contains contributions from both the residual block $f(x_{\text{in}})$ and skip connection: $x_{\text{out}} = f(x_{\text{in}}) + x_{\text{in}}$.

The evolution of the Res-GLOnet architecture over the course of network training is sketched in Figure 2c. When the network is training in the early and intermediate stages of the optimization process, each residual block outputs terms that are typically larger than the skip connection contributions. As a result, the network architecture functions as a deep network, which is required during these stages of Res-GLOnets training. As network training progresses, some of the residual blocks start to output relatively small contributions and $x_{\text{out}} \approx x_{\text{in}}$ owing to the emergence of vanishing gradients. The network architecture now functions as a shallow architecture, having effectively skipped over some of the residual blocks. Note that the increasing contribution of skip connections and reduction of network complexity is not explicitly and externally controlled but evolves over the course of network training, as the loss function guides the network output distribution to a relatively simple form.

3 Optimization of an antireflection coating

We first apply our Res-GLOnet algorithm to the design of a three-layer antireflection (AR) coating for a silicon solar cell. The thin-film AR stack is designed to minimize the average reflection at an air-silicon interface over the incident angle range $[0^\circ, 60^\circ]$ and wavelength range $[400, 1100]$ nanometer for both transverse magnetic (TM) and transverse electric (TE) polarization. As a benchmark, we compare our results with those from the study by Azunre et al. [22], which provides a guaranteed global optimum solution using a parallel branch-and-bound method. The algorithm requires extensive searching through the full design space and utilizes over 19 days of CPU computation to solve for the global optimum. To be consistent with the study by Azunre et al. [22], the refractive indices of the layers in our design implementation do not take discrete categorical values from a material library but are dispersionless and continuously varying in the interval $[1.09,$




Figure 3: Optimization of a third-layer thin film antireflection (AR) coating on silicon. (a) Histogram of the average reflectivity from 100 AR coatings designed using local gradient-based optimization. The best device has an average reflectivity of 1.82%. (b) Histogram of the average reflectivity from 100 AR coatings designed using a single Res-GLOnet. The best device has an average reflectivity of 1.81%. (c) Contour plot of reflectivity from the best Res-GLOnet–designed AR coating in (b) as a function of the incidence angle and wavelength, averaged for both TE- and TM-polarized waves. GLOnet, global optimization network.

2.60]. The thicknesses of each layer are also continuous variables within the interval [5, 200] nm.

To accommodate the continuous variable nature of the refractive index values in this problem, we modify our categorical optimization scheme by setting the hyperparameter $\alpha = 1$ as a constant and specifying the material library to contain only two materials with constant refractive indices $\{m^L, m^U\}$. $m^L = 1.09$ is the lower bound of the refractive index, while $m^U = 2.60$ is the upper bound. The constraint on thickness can be satisfied by a transformation: $\mathbf{t} = t^L + \text{Sigmoid}(\tilde{\mathbf{t}}) \cdot (t^U - t^L)$. Here, the thickness directly outputted by the generator, $\tilde{\mathbf{t}}$, is normalized to [0, 1] and then linearly transformed to the interval $[t^L, t^U]$, where $t^L = 5$ and $t^U = 200$ are the lower and upper thickness bound, respectively.

As a reference, we first optimize devices using local gradient-based optimization, by replacing the ResNet generator in our Res-GLOnet algorithm with an individual device layout. The optimizations are performed with 100 different devices, initialized using random thickness and refractive index values within the limits of [1.09, 2.60] and [5, 200] nm, respectively. Each optimization is performed over 200 iterations, so that a total of 20,000 sets of calculations are performed for the entire set of optimizations. A histogram of the results (Figure 3a) shows that the optimized devices have average reflectivities that span a wide range of values, from approximately 2 to 10%, demonstrating the highly nonconvex nature of the design space. Average reflectivity is calculated as the reflectivity averaged over the wavelengths, incident angles, and polarizations covered in the design specifications. A fraction of devices are near the global optimum, and the best device has an efficiency of 1.82%.

A histogram of devices sampled from a single trained Res-GLOnet is summarized in Figure 3b. A total of 200 iterations is used together with a batch size of 20 devices, so

that a total of 4000 sets of calculations are performed. The total time that Res-GLOnet requires for training is 7 s with a single GPU. All of the devices sampled from the Res-GLOnet are near the global optimum, showing the ability for the generative network to produce a narrow distribution of devices centered at the global optimum. The best device has an efficiency of 1.81%, and its reflectivity for differing incident wavelengths and angles is plotted in Figure 3c. The design of this best device is summarized in Table 1 and is consistent with the result reported in the study by Azunre et al. [22].

4 Optimization of the incandescent light bulb filter

To explore the applicability of Res-GLOnets to more complex problems, we apply our algorithm to optimize incandescent light bulb filters that transmit visible light and reflect infrared light (Figure 4a). In this scheme, the emitter filament heats to a relatively higher temperature using recycled infrared light, thereby enhancing the emission efficiency in the visible range [29].

Table 1: Optimized structure for the AR coating of Si.

Layer #	Refractive index	Thickness (nm)
1	Air	Superstrate
	2.60	54.2
2	1.68	93.6
3	1.17	149.2
	Si	Substrate

AR, antireflection.

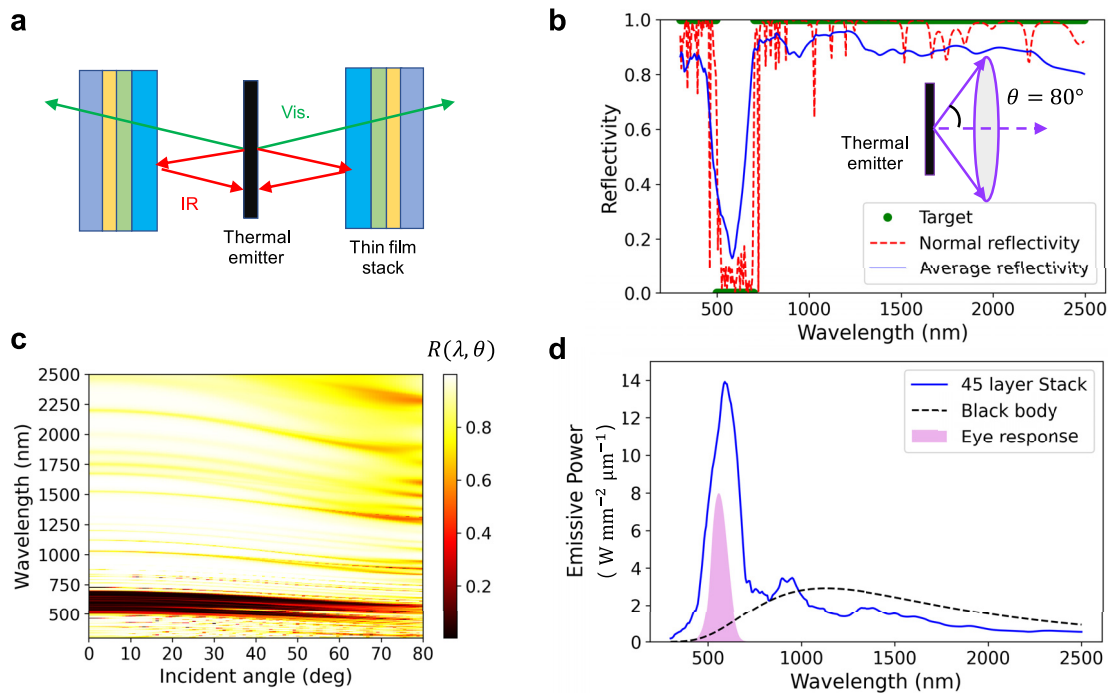


Figure 4: Thin-film stacks for incandescent light bulb filtering. (a) Schematic of an incandescent light bulb filter that transmits visible light and reflects infrared and ultraviolet light. (b) Reflection spectra of a 45-layer Res-GLOnet-optimized device, for normally incidence waves and waves averaged over a large incident solid angle, shown in the inset. (c) Reflection spectra of the device featured in (b) as a function of the incident angle, averaged for TE- and TM-polarized incident waves. (d) Emissive power of a blackbody incandescent source and an equivalent source sandwiched by the filter featured in (b). Also shown is the spectral response of the eye. GLOnet, global topology optimization network.

A range of design methods have been previously applied to this problem. In the initial demonstration of the concept, binary thin-film stacks were designed using a combination of local gradient-based optimization, used to tune the thickness of each layer, and needle optimization, which determined whether an existing layer should be removed or a new layer should be introduced [29]. A memetic algorithm was subsequently applied in which crossover, mutation, and downselecting operations were iteratively performed on a population of thin-film stacks to evolve the quality of devices [21]. Gradient-based local optimizations of device thicknesses were also periodically performed to refine the structures and accelerate algorithm convergence. In a third study, reinforcement learning (RL) was used in which an autoregressive recurrent neural network generated thin-film stacks layer by layer as a sequence [23]. Unlike the GLOnet generator, the probability distribution of the thin-film stack was explicitly outputted by the autoregressive generator. The distribution evolved by optimizing a reward function, and the gradient of the reward function with respect to the neural network weights was calculated using proximal policy optimization.

In our demonstration, we benchmark Res-GLOnets with the memetic and RL studies, which consider a material library comprising seven dielectric material types: Al_2O_3 , HfO_2 , MgF_2 , SiC , SiN , SiO_2 , and TiO_2 . The superstrate and substrate are both set to be air. The complete wavelength range under consideration is [300, 2500] nm, and the target reflection is set to be 0% for the wavelength range [500, 700] nm and 100% for all other wavelengths. The incident angles span [0, 72] degrees, and both TE and TM polarizations are considered.

We train a Res-GLOnet comprising 16 residue blocks for 1000 iterations with a batch size of 1000. The network is optimized using gradient decent with the momentum algorithm ADAM [42], and a learning rate of 1×10^{-3} is used. The broadband reflection characteristics of a 45-layer device show that the device operates with nearly ideal transmission in the [500, 700] nm interval and nearly ideal reflection at ultraviolet and near-infrared wavelengths, for both normal incidence and for incidence angles averaged over all solid angles within [0, 80] degrees (Figures 4b and 4c). The emission intensity spectrum of the light bulb with and without the thin-film filter is shown in Figure 4d. The input power is

fixed at 100 W, and the surface area of the emitter is 20 mm².

To evaluate the enhancement of visible light emission due to the filter, we compute the emissivity enhancement factor, χ , as a function of the number of thin-film layers:

$$\chi = \frac{\int_0^\infty E_{\text{emitter+stack}}(P_0, \lambda) V(\lambda) d\lambda}{\int_0^\infty E_{\text{emitter}}(P_0, \lambda) V(\lambda) d\lambda} \quad (7)$$

$E_{\text{emitter+stack}}(P_0, \lambda)$ and $E_{\text{emitter}}(P_0, \lambda)$ are the intensity emission spectrum given the input power P_0 . $V(\lambda)$ is the eye's sensitivity spectrum and is shown as the shaded region in Figure 4d. The view factor is the proportion of emitted light from the light bulb filament that can reach the light bulb filter. We use the view factor of 0.95 as was the case for the memetic study [21]. For a 45-layer device, the Res-GLOnet-optimized device achieved a χ of 17.2, and devices with as few as 30 layers still achieved a χ above 15 (Figure 5). The ability to realize high-performance devices with relatively few layers is practically important from a manufacturing and cost point of view. The 45-layer memetic algorithm and RL-optimized device have χ values of 14.8 and 16.6, respectively. We also benchmark Res-GLOnet with GLOnet based on a fixed architecture of four fully connected layers (FC-GLOnet). The benchmark, also plotted in Figure 5, shows that Res-GLOnet performs better in searching for proper devices in this nonconvex optimization problem, particularly for systems with larger numbers of thin films. The points in the plot each corresponds to the results of a single GLOnet run. In terms of computational cost, the memetic algorithm uses 600K simulations (a population size of 3000 and 200 iterations), the RL algorithm uses 30M simulations (a batch size of 3000 and 10,000 iterations), and GLOnets uses 500K simulations (a batch size of 500 and 1000 iterations). As such,

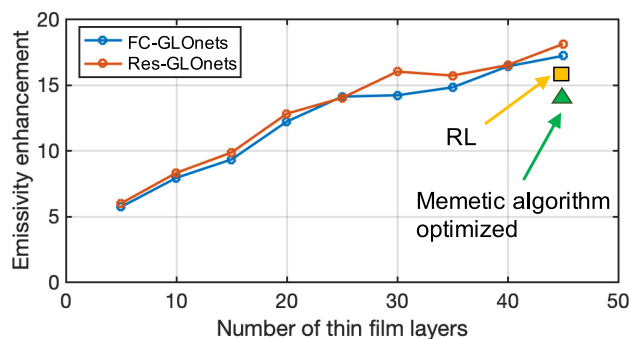


Figure 5: Plot of emissivity enhancement as a function of the number of thin-film layers, for devices optimized using Res-GLOnets and FC-GLOnets. Reference points are also plotted for devices designed using the reinforcement learning (RL) [23] and memetic [21] algorithm. GLOnet, global optimization network.

GLOnets is demonstrated to be a computationally efficient global optimization algorithm for this problem.

5 Conclusion

In summary, we show that Res-GLOnets are effective and efficient global optimizers for the multiobjective, categorical design of thin-film stacks. Categorical optimization is performed through the use of a probability matrix, which is fully differentiable and compatible with our neural network training framework. The incorporation of skip connections in our generative neural network helps it evolve from a deep to shallow architecture, which fits with our training objective and improves our search for the global optimum. Benchmarks of our algorithm with a known AR coating and incandescent light filter systems indicate that the Res-GLOnet is effective at searching for global optima, is computationally efficient, and outperforms a number of alternative design algorithms.

We anticipate that concepts developed within Res-GLOnets, particularly those in categorical optimization, can directly apply to the design of other photonics systems, such as lens design where the material type is selected from a material database. We also expect that the implementation of application-specific electromagnetic solvers, in conjunction with automatic differentiation packages, will serve as a foundational concept for many high-speed optimization algorithms beyond those for thin-film stacks. Generalizing the GLOnet algorithm to 3D photonic structures is challenging owing to the requirement of computationally expensive simulations. We envision that this roadblock can be overcome by using neural networks as fast surrogate solvers, for which much progress has been made [43, 44]. Looking ahead, we see opportunities for Res-GLOnets to apply to other fields in the physical science, ranging from material science and chemistry to mechanical engineering, where devices and systems are designed using combinations of discrete material types.

Acknowledgments: The simulations were performed in the Sherlock computing cluster at Stanford University.

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

Research funding: This work was supported by ARPA-E with Agreement Number DE-AR0001212, ONR with Agreement Number N00014-20-1-2105, and the Packard Foundation with Agreement Number 2016-65132

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

References

- [1] S. Molesky, Z. Lin, A. Y. Piggott, W. Jin, J. Vucković, and A. W. Rodriguez, “Inverse design in nanophotonics,” *Nat. Photonics*, vol. 12, no. 11, pp. 659–670, 2018.
- [2] S. D. Campbell, D. Sell, R. P. Jenkins, E. B. Whiting, J. A. Fan, and D. H. Werner, “Review of numerical optimization techniques for meta-device design,” *Opt. Mater. Express*, vol. 9, no. 4, pp. 1842–1863, 2019.
- [3] J. A. Fan, “Freeform metasurface design based on topology optimization,” *MRS Bull.*, vol. 45, no. 3, pp. 196–201, 2020.
- [4] J. Jiang, M. Chen, and J. A. Fan, “Deep neural networks for the evaluation and design of photonic devices,” arXiv preprint arXiv:2007.00084, 2020.
- [5] K. Yao, R. Unni, and Y. Zheng, “Intelligent nanophotonics: merging photonics and artificial intelligence at the nanoscale,” *Nanophotonics*, vol. 8, no. 3, pp. 339–366, 2019.
- [6] S. So, T. Badloe, J. Noh, J. Rho, and J. Bravo-Abad, “Deep learning enabled inverse design in nanophotonics,” *Nanophotonics*, vol. 9, no. 5, pp. 1041–1057, 2020.
- [7] J. Peurifoy, Y. Shen, L. Jing, et al., “Nanophotonic particle simulation and inverse design using artificial neural networks,” *Sci. Adv.*, vol. 4, no. 6, p. eaar4206, 2018.
- [8] D. Liu, Y. Tan, E. Khoram, and Z. Yu, “Training deep neural networks for the inverse design of nanophotonic structures,” *ACS Photonics*, vol. 5, no. 4, pp. 1365–1369, 2018.
- [9] 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, no. 35, p. 1901111, 2019.
- [10] J. Jiang, D. Sell, S. Hoyer, J. Hickey, J. Yang, and J. A. Fan, “Freeform diffractive metagrating design based on generative adversarial networks,” *ACS Nano*, vol. 13, no. 8, pp. 8872–8878, 2019.
- [11] Z. Liu, D. Zhu, S. P. Rodrigues, K.-T. Lee, and W. Cai, “Generative model for the inverse design of metasurfaces,” *Nano Lett.*, vol. 18, no. 10, pp. 6570–6576, 2018.
- [12] F. Wen, J. Jiang, and J. A. Fan, “Robust freeform metasurface design based on progressively growing generative networks,” *ACS Photonics*, vol. 7, no. 8, pp. 2098–2104, 2020.
- [13] T. W. Hughes, M. Minkov, I. A. D. Williamson, and S. Fan, “Adjoint method and inverse design for nonlinear nanophotonic devices,” *ACS Photonics*, vol. 5, no. 12, pp. 4781–4787, 2018.
- [14] D. Sell, J. Yang, S. Doshay, R. Yang, and J. A. Fan, “Large-angle, multifunctional metagratings based on freeform multimode geometries,” *Nano Lett.*, vol. 17, no. 6, pp. 3752–3757, 2017.
- [15] A. Y. Piggott, J. Lu, K. G. Lagoudakis, J. Petykiewicz, T. M. Babinec, and J. Vucković, “Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer,” *Nat. Photonics*, vol. 9, no. 6, pp. 374–377, 2015.
- [16] T. Phan, D. Sell, E. W. Wang, et al., “High-efficiency, large-area, topology-optimized metasurfaces,” *Light Sci. Appl.*, vol. 8, no. 1, pp. 1–9, 2019.
- [17] J. Yang, D. Sell, and J. A. Fan, “Freeform metagratings based on complex light scattering dynamics for extreme, high efficiency beam steering,” *Ann. Phys.*, vol. 530, no. 1, p. 1700302, 2018.
- [18] J. Lu and J. Vucković, “Objective-first design of high-efficiency, small-footprint couplers between arbitrary nanophotonic waveguide modes,” *Optic Express*, vol. 20, no. 7, pp. 7221–7236, 2012.
- [19] J. Jiang and J. A. Fan, “Global optimization of dielectric metasurfaces using a physics-driven neural network,” *Nano Lett.*, vol. 19, no. 8, pp. 5366–5372, 2019.
- [20] J. Jiang and J. A. Fan, “Simulator-based training of generative neural networks for the inverse design of metasurfaces,” *Nanophotonics*, vol. 9, no. 5, pp. 1059–1069, 2019.
- [21] Y. Shi, W. Li, A. Raman, and S. Fan, “Optimization of multilayer optical films with a memetic algorithm and mixed integer programming,” *ACS Photonics*, vol. 5, no. 3, pp. 684–691, 2017.
- [22] P. Azunre, J. Jean, C. Rotschild, V. Bulovic, S. G. Johnson, and M. A. Baldo, “Guaranteed global optimization of thin-film optical systems,” *New J. Phys.*, vol. 21, no. 7, p. 073050, 2019.
- [23] H. Wang, Z. Zheng, C. Ji, and L. J. Guo, “Automated optical multilayer design via deep reinforcement learning,” arXiv preprint arXiv:2006.11940, 2020.
- [24] A. P. Raman, M. Abou Anoma, L. Zhu, E. Rephaeli, and S. Fan, “Passive radiative cooling below ambient air temperature under direct sunlight,” *Nature*, vol. 515, no. 7528, pp. 540–544, 2014.
- [25] W. Li, Y. Shi, K. Chen, L. Zhu, and S. Fan, “A comprehensive photonic approach for solar cell cooling,” *ACS Photonics*, vol. 4, no. 4, pp. 774–782, 2017.
- [26] A. Lenert, D. M. Bierman, Y. Nam, et al., “A nanophotonic solar thermophotovoltaic device,” *Nat. Nanotechnol.*, vol. 9, no. 2, pp. 126–130, 2014.
- [27] Y. Shen, D. Ye, I. Celanovic, S. G. Johnson, J. D. Joannopoulos, and M. Soljačić, “Optical broadband angular selectivity,” *Science*, vol. 343, no. 6178, pp. 1499–1501, 2014.
- [28] F. Cao, Y. Huang, L. Tang, et al., “Toward a high-efficient utilization of solar radiation by quad-band solar spectral splitting,” *Adv. Mater.*, vol. 28, no. 48, pp. 10659–10663, 2016.
- [29] O. Ilic, P. Bermel, G. Chen, J. D. Joannopoulos, I. Celanovic, and M. Soljačić, “Tailoring high-temperature radiation and the resurrection of the incandescent source,” *Nat. Nanotechnol.*, vol. 11, no. 4, pp. 320–324, 2016.
- [30] M. Gerken and D. A. B. Miller, “Wavelength demultiplexer using the spatial dispersion of multilayer thin-film structures,” *IEEE Photonics Technol. Lett.*, vol. 15, no. 8, pp. 1097–1099, 2003.
- [31] W. J. Wild and H. Buhay, “Thin-film multilayer design optimization using a Monte Carlo approach,” *Opt. Lett.*, vol. 11, no. 11, pp. 745–747, 1986.
- [32] R. I. Rabady and A. Ababneh, “Global optimal design of optical multilayer thin-film filters using particle swarm optimization,” *Optik*, vol. 125, no. 1, pp. 548–553, 2014.
- [33] A. V. Tikhonravov, M. K. Trubetskov, and G. W. DeBell, “Application of the needle optimization technique to the design of optical coatings,” *Appl. Opt.*, vol. 35, no. 28, pp. 5493–5508, 1996.
- [34] V. Pervak, A. V. Tikhonravov, M. K. Trubetskov, S. Naumov, F. Krausz, and A. Apolonski, “1.5-octave chirped mirror for pulse compression down to sub-3 fs,” *Appl. Phys. B*, vol. 87, no. 1, pp. 5–12, 2007.
- [35] A. V. Tikhonravov, M. K. Trubetskov, and G. W. DeBell, “Optical coating design approaches based on the needle

- optimization technique,” *Appl. Opt.*, vol. 46, no. 5, pp. 704–710, 2007.
- [36] A. Paszke, S. Gross, F. Massa, et al., “PyTorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems*, Vancouver, Canada, Curran Associates Inc, 2019, pp. 8026–8037.
- [37] T. W. Hughes, I. A. D. Williamson, M. Minkov, and S. Fan, “Forward-mode differentiation of Maxwell’s equations,” *ACS Photonics*, vol. 6, no. 11, pp. 3010–3016, 2019.
- [38] M. Minkov, I. A. D. Williamson, L. C. Andreani, et al., “Inverse design of photonic crystals through automatic differentiation,” *ACS Photonics*, vol. 7, no. 7, pp. 1729–1741, 2020.
- [39] C. M. Bishop, *Pattern Recognition and Machine Learning*, New Delhi, India, Springer, 2006.
- [40] A. Chakrabarti, “Learning sensor multiplexing design through back-propagation,” in *Advances in Neural Information Processing Systems*, Barcelona, Spain, Curran Associates Inc, 2016, pp. 3081–3089.
- [41] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, NV, IEEE Computer Society, 2016, pp. 770–778.
- [42] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” arXiv preprint arXiv:1412.6980, 2015.
- [43] R. Pestourie, Y. Mroueh, T. V. Nguyen, P. Das, and S. G. Johnson, “Active learning of deep surrogates for PDEs: Application to metasurface design,” arXiv preprint arXiv:2008.12649, 2020.
- [44] M. V. Zhelyeznyakov, S. L. Brunton, and A. Majumdar, “Deep learning to accelerate Maxwell’s equations for inverse design of dielectric metasurfaces,” arXiv preprint arXiv:2008.10632, 2020.

Supplementary Material: The online version of this article offers supplementary material (<https://doi.org/10.1515/nanoph-2020-0407>).