

Nested deep transfer learning for modeling of multilayer thin films

Rohit Unni^{a,b}, Kan Yao^{a,b} and Yuebing Zheng^{a,b,*}^aUniversity of Texas at Austin, Walker Department of Mechanical Engineering, Austin, Texas, United States^bUniversity of Texas at Austin, Texas Materials Institute, Austin, Texas, United States

Abstract. Machine-learning techniques have gained popularity in nanophotonics research, being applied to predict optical properties, and inversely design structures. However, one limitation is the cost of acquiring training data, as complex structures require time-consuming simulations. To address this, researchers have explored using transfer learning, where pretrained networks can facilitate convergence with fewer data for related tasks, but application to more difficult tasks is still limited. In this work, a nested transfer learning approach is proposed, training models to predict structures of increasing complexity, with transfer between each model and few data used at each step. This allows modeling thin film stacks with higher optical complexity than previously reported. For the forward model, a bidirectional recurrent neural network is utilized, which excels in modeling sequential inputs. For the inverse model, a convolutional mixture density network is employed. In both cases, a relaxed choice of materials at each layer is introduced, making the approach more versatile. The final nested transfer models display high accuracy in retrieving complex arbitrary spectra and matching idealized spectra for specific application-focused cases, such as selective thermal emitters, while keeping data requirements modest. Our nested transfer learning approach represents a promising avenue for addressing data acquisition challenges.

Keywords: artificial neural networks; multilayer structures; nanophotonics; inverse design; transfer learning.

Received Mar. 31, 2024; revised manuscript received Sep. 4, 2024; accepted for publication Sep. 11, 2024; published online Oct. 8, 2024.

© The Authors. Published by SPIE and CLP under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI.

[DOI: [10.1117/1.AP.6.5.056006](https://doi.org/10.1117/1.AP.6.5.056006)]

1 Introduction

In recent years, machine-learning (ML)-based techniques have surged in popularity as tools for addressing problems in optics and photonics.¹⁻³ Deep learning (DL), employing complex many-layered neural networks, has become the predominant type of algorithm employed. DL networks transform input values to output values via a number of intermediary hidden layers.^{4,5} The weight connection values between these layers are learned through exposure to labeled training data and comparing model predictions with the true values using an objective loss function. Thus far, DL models have been applied for predicting properties of structured materials, such as their optical response in the spectral or spatial domain, faster than simulations can do.⁶⁻¹¹ DL models are also applied in the inverse direction, that is, taking a desired optical performance and

calculating the design of structures that would produce it.¹²⁻¹⁸ Successful examples of this inverse design approach include multilayer structures,¹⁹⁻²¹ metasurfaces,²²⁻²⁷ optical cloaks,²⁸⁻³⁰ among many other photonic devices.³¹⁻³⁵ In general, learning the inverse mapping is far more complex than learning the forward one. Despite this, strides have been made in developing DL models that can accurately recreate arbitrary spectra, aiming ultimately at addressing specific applications where an ideal optical response contains complex and extreme features.^{17,36}

One major limitation of DL-based inverse design, however, is the exceptionally high cost of acquiring high-quality labeled data. Most sufficiently complex structures require full-wave simulations to predict optical responses. In some circumstances, simulating even a single structure can take on the order of hours, and one may need hundreds of thousands to millions of samples to train a single model accurately, posing the largest bottleneck for building and scaling nanophotonic inverse design models. Moreover, for a given task, the model is trained with certain

*Address all correspondence to Yuebing Zheng, zheng@austin.utexas.edu

constraints and assumptions about the design being predicted, such as a fixed material for the substrate or cladding of a meta-surface and predefined geometries of the resonant elements. Once trained, the model would only be able to give useful design suggestions for that specific set of limitations. Introducing additional geometric parameters and including variables of drastically different natures (e.g., indicators of material choices) both raise the complexity of the task steeply.

One approach to addressing the above issue is the use of transfer learning. For some types of structures, although tackling a complete version with enough complexity for practical applications is exceedingly slow, simulating thousands of data points for a simplified toy version with reduced degrees of freedom can

be relatively fast and manageable. Projecting this difference to the training of DL models, rather than initializing the weight values in a network randomly, the initial layers' values are taken from another network that has already been trained to predict a similar, and in many cases, simpler task.^{37,38} In a sense, instead of learning a complex task from scratch, the model needs to learn just enough to account for the differences in the two data sets. As such, transfer learning can potentially allow for faster convergence to accurate predictions using fewer data. This relies on the assumption that the features and relations learned for the first task will also have high predictive power for the next one. Previous work has shown the ability to transfer knowledge between inverse^{39,40} and forward^{41–44} design models trained on

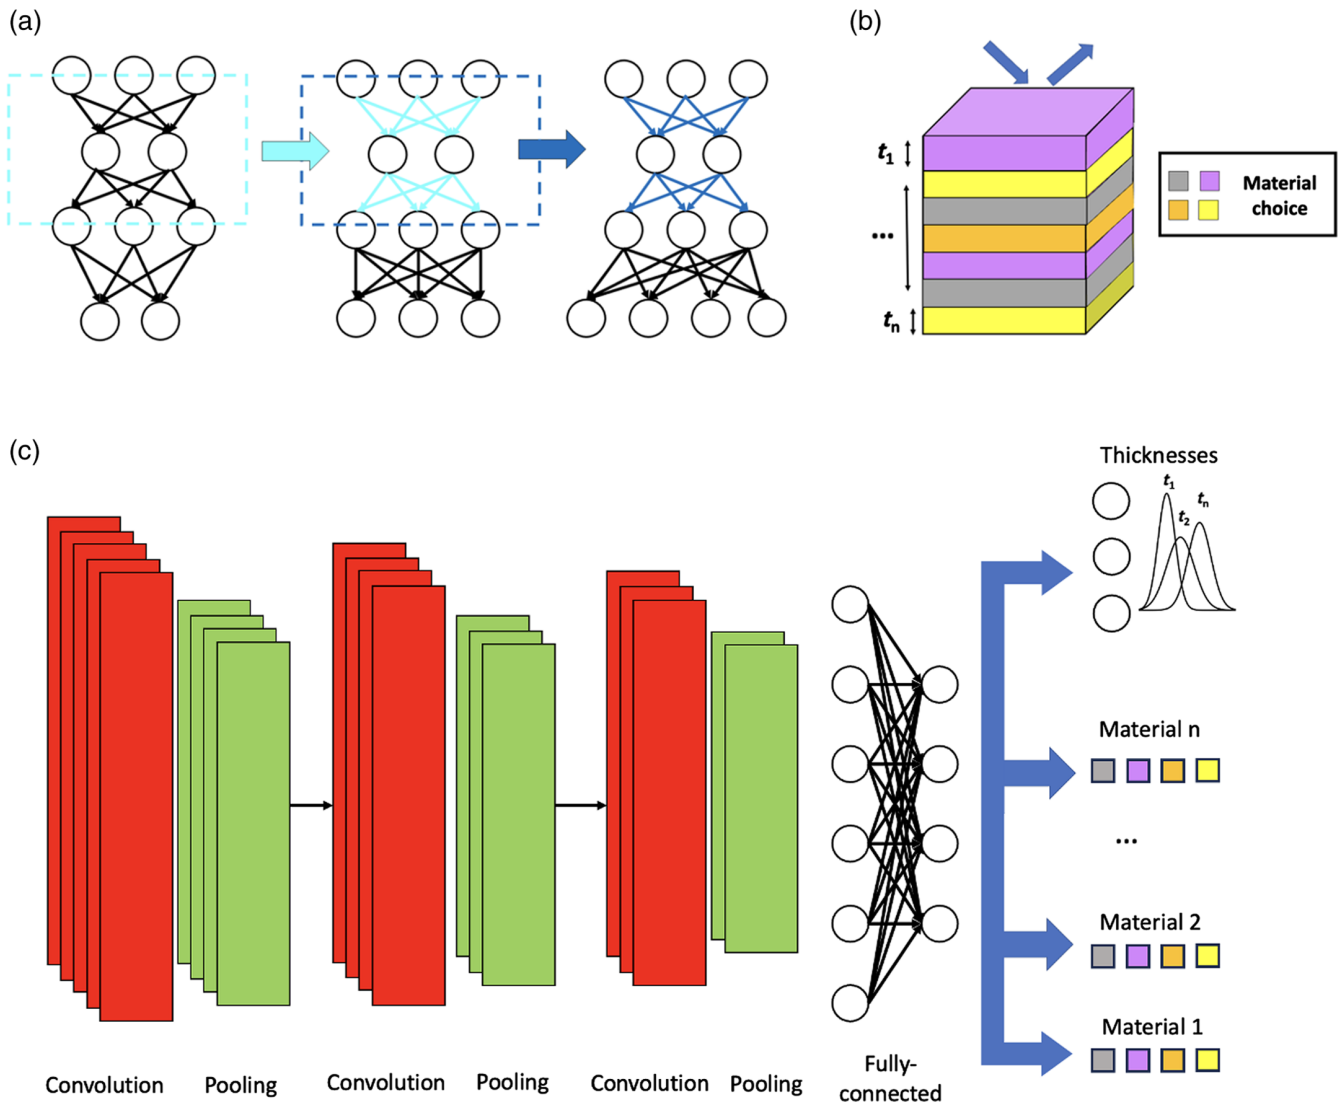


Fig. 1 (a) Schematic illustrating the principle of nested transfer. Left to right: Weights from the previous model are taken after training (box with dashed lines) and used for initialization of the weights of the next model (solid-colored lines between neurons), as the complexity of the output gradually increases. (b) Diagram of a multilayer thin-film structure. The structure has a choice of any of four materials at each layer, with the constraint that no two neighboring layers are the same material. (c) Architecture of the mixed convolutional MDN used for the inverse design. There are initial pairs of convolutional and max pooling layers leading to fully connected layers. The output is split into categorical channels predicting the material choice at each layer and a final MDN channel representing probability distributions of the layer thicknesses.

different but similar structures, as well as between structures of the same type but with different levels of complexity. While this strategy can ease data requirements, the constraint that the target and pretrained tasks need to be similar imposes some difficulties. If the original and new domains have the same level of complexity, this either means that both must be a simple task to model, for which a high amount of data is already not needed, or, if it is for a complex task, the first pretrained model will likely already need a high amount of data. Many of the existing cases of transfer learning for photonic design model tasks have more limited ranges of possible designs, such as metasurfaces with binary choices at each design variable⁴⁵ or limited geometric shapes.⁴⁶ For cases with increasing complexity, transfer learning has been more limited, such as simpler forward modeling for thin film structures with a small number of layers.^{41,43,44} A structure with higher degrees of freedom in layer numbers and materials will be able to produce a wider range of optical behaviors and therefore be more flexible for practical applications. Because of this, finding ways to utilize transfer learning for more complex inverse modeling without high data requirements is of particular interest.

In this work, we report a nested transfer learning approach wherein models are trained to predict structures of gradually increasing complexity, with transfer operations done between each model [Fig. 1(a)]. This “nesting” strategy effectively allows for a small amount of data per network and can model a significantly higher level of optical complexity than in previously reported works. We demonstrate this on both forward prediction and inverse design of multilayer thin film structures [Fig. 1(b)]. Since both neural networks and thin film structures use the term “layers” for their components, hereafter we refer them to “network layers” and “structure layers,” respectively, to distinguish them. For forward modeling, we build transfer models up to 30 structure layers deep using a bidirectional recurrent neural network (RNN) architecture. For the inverse model, we build up to 10 structure layers, using a convolutional partial mixture density network (MDN) [Fig. 1(c)]. For each data set, the material used at each structure layer is randomly selected from a prechosen list. We stress that the high degree of freedom in the design variables represents a significant challenge for modeling, surpassing those of previously reported thin-film inverse design tasks.³⁶ Because photonic devices composed of building blocks in regular shapes are typically described by a vector of discrete variables, the degrees of complexity of their design process are somewhat comparable. It is thus reasonable to assume that conclusions drawn from the study of thin films are applicable to devices showing a higher visual complexity, such as many metasurfaces and photonic/plasmonic crystals. The combination of free material choice at every layer, as well as fully continuous thickness values within a wide range, results in a design-to-response mapping that is highly sensitive to small changes. Despite this, our nested transfer approach allows us to achieve accurate retrieval without a significant increase in data requirements. We evaluate the forward and inverse models on arbitrary random spectra and can recreate closely matching designs. For the inverse model, we implement a postprocessing optimization method using the architecture to further improve results. Finally, as a proof-of-concept demonstration of the model’s ability to address realistic problems, we present a design of selective thermal emitter conceptually similar to multilayer metamaterials for thermophotovoltaics.⁴⁷

2 Materials and Methods

For the forward model, we use a bidirectional RNN [Fig. 2(a)]. For an RNN, rather than information flowing strictly to the next network layer as in a feedforward network, information can also flow from one layer back to itself.⁴⁸ A standard fully connected network does not assume that relations between any given features are more important than any others from initialization, and input variables can be arranged in any order provided they are consistent throughout the data set. In RNNs, neurons store an internal state or memory that affects how they process subsequent inputs. This allows them to handle inputs of arbitrary length, processing them differently based on recent inputs and learning context. Therefore, RNNs excel in handling sequential data such as time series forecasting and natural language processing,^{49–51} where the order of inputs contains critical information for making predictions. The reason for using an RNN for the forward modeling here is that the input data have features in common with the sequential data typically used for RNNs. For thin-film stacks, light passes through the layers in order, and therefore the final optical spectrum is heavily affected by the interactions at the interface between neighboring layers, based on the material properties and thickness of each. The data for structure layers close to each other are more important than layers further away in the structure. As such, the design variables are essentially a “sentence,” with the words made up of material and thickness data instead of letters. One type of RNN is a long-short term memory (LSTM) network. It uses a series of gates that can differentially process inputs to decide to what degree the new input should influence both the hidden state and output. This enables better capturing of information over longer sequence lengths and helps reduce the “vanishing gradient” problem that can plague traditional RNNs.^{52,53} A bidirectional LSTM extends the architecture to learn relations going in both directions. This can yield better performance for data where the inputs before and after both have high predictive power, such as in natural language.⁵⁴ Because light can be reflected into previous layers as it goes through the stack, information about layers before and after a given layer is needed, and thus a bidirectional approach is used. Though the connection and operation of LSTM layers are more complex than fully connected layers, the process of weight transfer is essentially the same.

The final network architecture for forward modeling uses a series of bidirectional LSTM layers to initially process the input data before connecting to fully connected layers to get the final output. The full architecture is shown in Sec. S1 in the [Supplementary Material](#). The transfer occurs from six structure layers all the way up to 30 structure layers, yielding accurate predictions for a higher level of complexity than that previously reported. For the loss function, we use the root mean squared error (RMSE). A different data set is generated for each structure with a different number of layers. Large-sized data sets are generated by calculating transmission and reflection using Fresnel equations.^{55,56} Most previous demonstrations of ML on thin-film optics have had fixed choices of materials at each layer based on prior physics-based intuitions guided by the researchers. This makes the modeling task much simpler, as material choice interacts with other design variables at a very fundamental level. Here, we also allow for free material choice at all structure layers to demonstrate the ability to learn a more complex mapping while using fewer overall data. Materials are chosen randomly from one of four oxides: SiO₂, TiO₂, Al₂O₃,

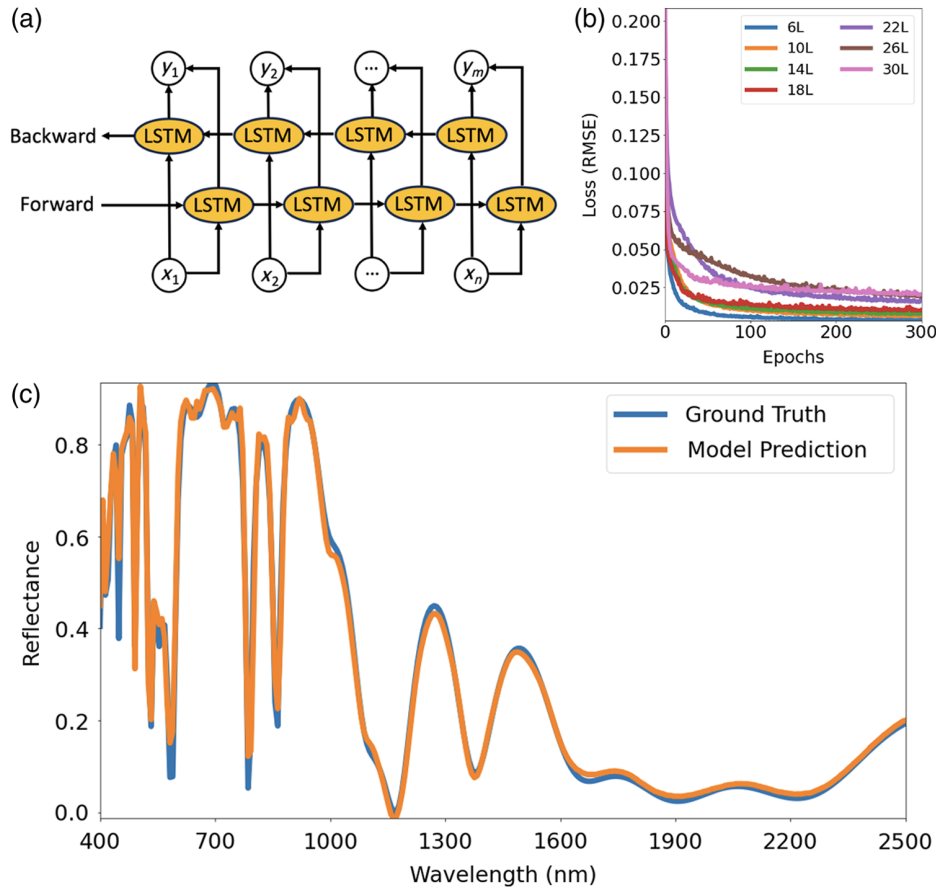


Fig. 2 (a) Diagram of a bidirectional RNN used for the forward model. (b) Training curves for nested transfer forward prediction for thin-film structures of increasing complexity (see legends), up to 30 layers. (c) Comparison of requested ground-truth spectrum (blue) and the spectrum predicted by the forward model (orange) for a randomly chosen case in the test data set.

and Ta_2O_5 . For all models demonstrated here, four possible materials are chosen during data generation, with the constraint that no two neighboring structural layers are the same material. For a 10-layer structure, this represents 4×3^9 , or over 75,000 possible combinations for material choice alone. The structure is placed on a semi-infinite glass substrate and surrounded by air cladding. We calculate the reflectance spectrum between 400 and 2500 nm, with the wavelength range discretized into 300 points.

The inverse model, unlike the forward network, needs to predict both categorical variables, the material choice at each layer, as well as continuous variables, the layer thicknesses. This complicates the modeling significantly, and to deal with this, the model branches into multiple outputs. The initial layers are three pairs of convolutional and max pooling layers that first aim to learn key spectral features such as the location and shape of peaks. This is followed by a series of fully connected layers to learn the relation and importance between these spectral features. All the initial layers use the ReLU activation function. The network then branches into $n + 1$ sets of outputs for a structure with n layers [Fig. 1(c)]. There are four output neurons for each of the first n sets, representing the relative likelihood of each of the four possible material choices for each structure layer. The last set of the output, connected to an MDN layer, comprises a series of neurons encoding parameters for several probability

distributions over the possible range of thicknesses. Each mixture is parametrized by a mean μ and a variance σ , as well as a mixing weight π . For the inverse modeling here, 32 mixtures were used. The MDN is chosen for the layer thicknesses rather than a typical fully connected layer for the final output due to its demonstrated ability to converge accurately when processing multimodal data.⁵⁷ The final outputs model two types of data, the categorical material choices and the continuous layer thicknesses, and so two different loss functions are used. The outputs representing the material choices use categorical cross-entropy as the loss function, as well as a SoftMax activation function so that the four outputs sum to one. The values then can be interpreted as an estimated probability for each of the four choices. The MDN output representing the continuous layer thicknesses uses the negative log likelihood as its loss function, which measures how well the actual probability distribution of the data matches the expected one produced by the model. No activation function is used for the final MDN layer.

3 Results

For the forward model, transfer was tested at different conditions to compare which gave the best results at 30 structure layers. The data for two comparative studies are presented in Sec. S2 in the [Supplementary Material](#). It is concluded that

the step size of four and transferring all but the final network layer of weights yielded the best results when the 30-structure layer model converged, resulting in a final nested transfer protocol. For this, 26,000 samples are first generated for a six-layer structure, of which 70% are used for training, and a model is trained from scratch for 500 epochs. After training, a new equal-sized data set is generated for a 10-layer structure and is initialized with weights from the pretrained six-layer structure model following the protocol, and this model is trained for the same number of epochs, behaving similarly in reaching convergence [Fig. 2(b)]. The transfer process is repeated, increasing the size of the thin-film structure by four layers at a time, with information accumulating with all the successive transfers. Finally, a model for 30-layer structures is trained, which can reach an RMSE on test data of 0.05, accurately reproducing arbitrary spectra [Fig. 2(c)]. Across all models used in the final nested transfer procedure, a total of 127,400 (i.e., $26,000 \times 7 \times 70\%$) samples are used for training. Previous works have reported on transfer learning for inverse design for thin-film optics; however, their use of transfer was to better learn the simpler forward model and used that forward model in conjunction with other optimization techniques.^{43,44} Alternately, the transfer has been used to more efficiently learn internal variables like mixture density parameters for a single model.⁴⁰ The results showcased here fully model both the forward and inverse directions, and on top of that, deal with a larger number of structure layers, more options of materials, and a broader band of wavelengths while achieving comparable prediction accuracy. For other types of optical structures, such as metasurfaces, transfer learning has been used with full inverse design models, albeit with simple constraints on the possible designs.³⁹ We show that the nested transfer method can model a significantly higher degree of complexity than existing benchmarks while keeping data requirements modest. To demonstrate how much the addition of the transfer protocol improves the training, we compare two models, with and without transfer, using different regression metrics. The results are shown in Sec. S2 in the [Supplementary Material](#). These improvements afforded by nested transfer enable the potential advantage of using the bidirectional RNN architecture for the accurate retrieval of optical spectra with high complexity.

As the inverse modeling is considerably more complex than the forward modeling, more data are needed. The increase in optical complexity rises exponentially with the number of possible layers. We account for this, while still using fewer data, by scaling the size of the data sets linearly with the number of layers. We also start at a lower initial layer number to allow for more consecutive transfer and learning of simple tasks. Here, an initial two-layer data set and model are generated and trained, respectively, before then transferring weights to a three-layer case. The two-layer case trains with 20,000 samples in the data set, with 70% of the data used directly for training and the remaining 30% used for validation. For the three-layer case, the weights from the initial convolutional and pooling layers are transferred, and a new data set is generated with 30,000 total samples. This process is then repeated for each layer transfer up to 10 layers, which uses 70,000 samples for training and 30,000 for validation, for a total data set size of 100,000 samples. We find that, unlike in the simpler case of forward transfer, increasing the structure layer number more than one at a time can cause overfitting of the test data. This may be due to the significantly higher complexity in the inverse modeling typically requiring larger data sets to train from scratch.

The final configuration reached is transferring from two-layers to three-layers, three-layers to four-layers, and so on, with eight weight layers in the network transferred at each step. Across all models, a total of 378,000 training samples are used. The models are trained for 300 epochs, each using an Adam optimizer with a learning rate of 0.01 and a scheduler reducing the learning rate by 70% when the average loss across all outputs does not decrease for 10 consecutive epochs [Fig. 3(a)]. The loss functions are not easily interpretable, but we can estimate the accuracy of proposed designs by simulating them and calculating the RMSE of the produced spectra compared to the original ground truths. For the MDN's output, we take the mean of each of the distributions that gives the highest likelihood value for each parameter, as all distributions calculated by the MDN for this data set tend to be unimodal or quasi-unimodal. For a 10-layer case, simply taking a single mean output from the probability distributions for the layer thicknesses and the highest probability values for the material choices of each layer, we get a response RMSE of 0.15. A selected case from the test data set is shown in Fig. 3(b), showing a decent agreement on most features. The remaining discrepancy, most likely caused by the naive sampling strategy, can be drastically reduced by implementing a postprocessing procedure. For this, a forward model needs to be trained to act as an estimator of the proposed candidate designs' viability. Using the same type of network used in the forward nested transfer protocol, we train a model for 10-layer forward prediction on the same data set for inverse modeling. This model is trained for 500 epochs without any prior transfer and can reach a test set RMSE below 0.01, yielding accurate estimation of the optical responses of arbitrary candidate designs.

The postprocessing procedure involves sampling the MDN output distributions for the design variables one at a time and fixing the best estimated value before moving on to the next variable. The full details for the procedure are given in Sec. S3 in the [Supplementary Material](#). A comparison of a random requested spectrum, initial model suggestion design, and the design after postprocessing is shown in Fig. 3(c), where even though the initial design deviates wildly from the ground truth in a relatively rare case, it is recovered through postprocessing. An unexpected phenomenon observed for this data set is that the retrieved designs do not necessarily stick to the 10-layer structure. It is not rare, as shown in Fig. 3(c), that two adjacent layers take the same material, resulting in essentially a reduced layer number. Other than the obvious cause that the distinctness constraint was only applied to data generation but not inverse design, another possible reason is the close refractive index values of the chosen oxides. Although both issues can be avoided in the implementation, the current model offers the flexibility in finding equivalent designs with fewer physical layers, partially a consequence of the weights transferred from the nine- and eight-layer models. For the task under study, the large thickness ranges and free material choice at each layer represent a significant challenge for modeling, and the complete network can still retrieve accurate solutions. In the selected case in Fig. 3(c), the postprocessing gives a 77% reduction in the RMSE between the requested and model-suggested spectra. We stress again that previous works using transfer learning on thin-film structures have primarily studied transfer between forward models and at significantly lower layer numbers and modeling complexity than we report here (see a comparison in Table S2 in the [Supplementary Material](#)). We compare these models based on

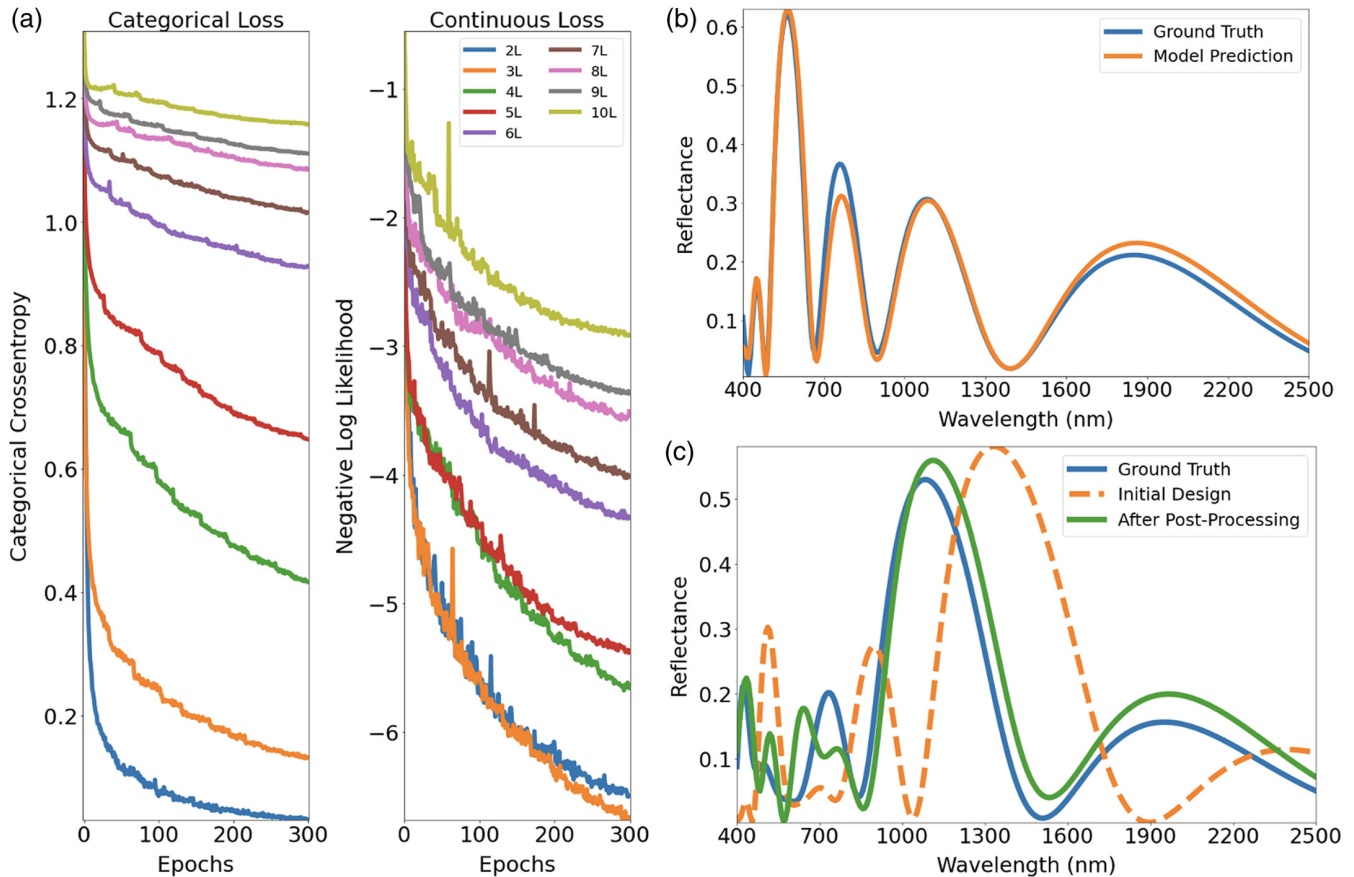


Fig. 3 (a) Training curves for nested transfer inverse design up to 10 layers. The categorical loss (left panel) refers to the outputs representing the material choices at each layer, while the continuous loss (right panel) refers to the negative log likelihood for the MDN predicting the thickness of each layer. (b) Comparison of requested ground-truth spectrum (blue) and the design suggested by the model for a randomly chosen case in the test data set without post-processing. The model-suggested design is [Al₂O₃-84 nm, SiO₂-85 nm, Al₂O₃-91 nm, SiO₂-90 nm, Al₂O₃-88 nm, SiO₂-90 nm, Al₂O₃-90 nm, SiO₂-88 nm, TiO₂-77 nm, SiO₂-92 nm]. (c) Comparison between requested ground-truth spectrum (blue) and the spectra of the original design proposed by the inverse model (orange) and of the design after postprocessing (green). The model design after postprocessing is [SiO₂-95 nm, Ta₂O₅-78 nm, SiO₂-122 nm, SiO₂-50 nm, Al₂O₃-50 nm, Ta₂O₅-98 nm, Ta₂O₅-54 nm, Al₂O₃-119 nm, TiO₂-90 nm, SiO₂-94 nm].

the maximum number of layers, number of material choices, and length of the design vector to demonstrate the increased difficulty of the modeling task. The modeling of both forward and inverse directions with free material choice at each layer gives our method more flexibility to tackle design requests for complex real-world applications.

We demonstrate this flexibility of our nested transfer procedure on specific applications. We focus on the use of thin-film structures for selective thermal emissions. Thin-film stacks can be used as optical filters to enhance the transmission, reflection, or absorption over large bandwidths and with a high contrast.^{36,56} At infrared wavelengths, these properties have well-established connections to the thermal emissivity of materials.^{58–60} In practice, materials giving spectrally selective thermal emissions are of great interest in enhancing the efficiency of photovoltaic (PV) cells.^{47,61,62} Thin-film thermal emitters are typically placed on a tungsten (W) substrate and use W as one of the materials as well. We consider a 10-layer stack and a material library of W, SiO₂, TiO₂, and Al₂O₃, placed on a semi-infinite W substrate and

illuminated at normal incidence [Fig. 4(a)]. The oxide layers have a much wider possible thickness range, from 30 to 300 nm, with the W layers range from 10 to 70 nm. Reflectance spectra from 300 to 3000 nm are generated for the samples. The same nested transfer process derived from previous testing is used, with 20,000 samples initially on a two-layer data set and model. This is successively transferred in single layer increments, with the number of training samples scaling linearly with the layer numbers until they reach 10 layers, which uses a total of 100,000 samples, with each data set split into 70% training and 30% validation. We test out the final model both in terms of reproducing arbitrary spectra from the test data set [Fig. 4(b)], as well as an unrealistic idealized spectral input for optimized performance of thermophotovoltaics [Fig. 4(c)]. The input is a steep sigmoid curve with its inflection point at the band edge of a PV cell with 0.55 eV.^{62–64} Below this threshold wavelength λ_{PV} , the thermal emitter targets near-unity emissivity to approximate the blackbody radiation, whereas above λ_{PV} , the emissivity needs to decline abruptly for better efficiency and

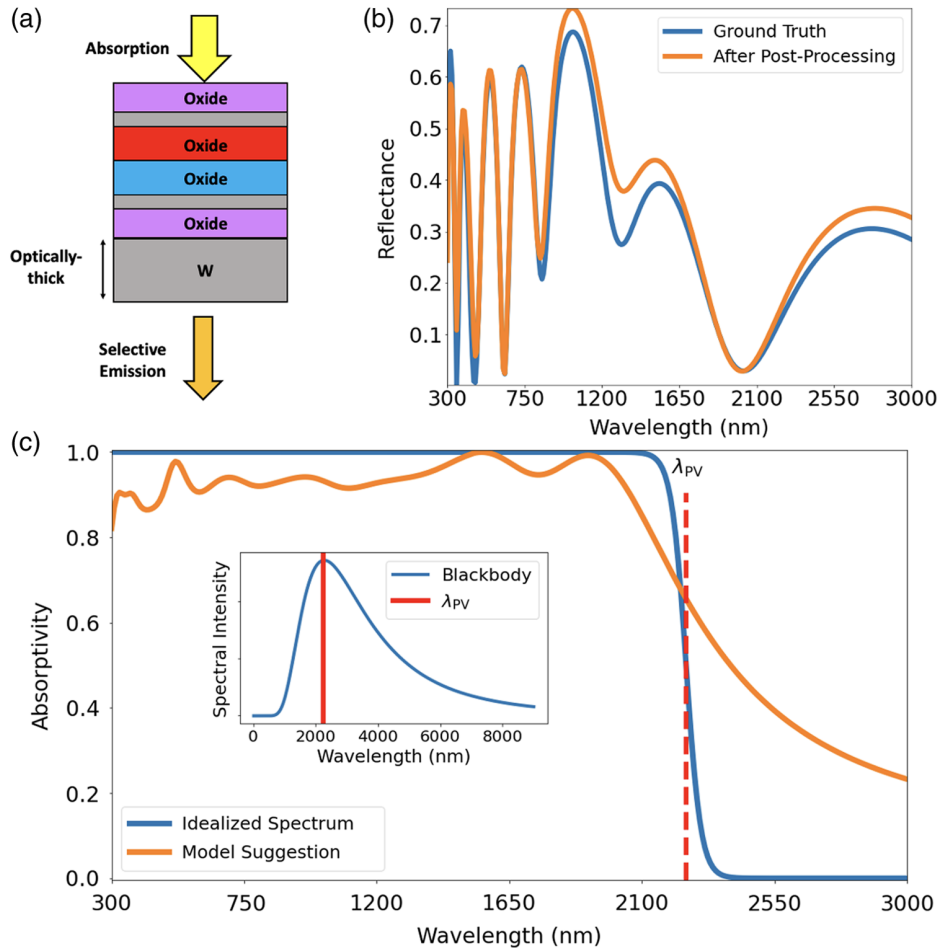


Fig. 4 (a) Diagram of thin-film structure used for selective thermal emission. (b) Comparison between requested ground-truth spectrum and spectrum produced by model suggested design for an arbitrary test data set sample. The model-suggested design is [HfO₂-124 nm, SiO₂-130 nm, W-10 nm, HfO₂-195 nm, SiO₂-119 nm, HfO₂-93 nm, W-10 nm, SiO₂-196 nm, W-20 nm, HfO₂-280 nm]. (c) Comparison of idealized absorptivity spectrum (blue) and spectrum produced by the inverse design model (orange). The red dotted line denotes the transition wavelength λ_{PV} for a PV cell with a bandgap of 0.55 eV. The inset figure shows the blackbody radiation curve at 1000°C (1273 K), with the same λ_{PV} highlighted. The model suggested design is [SiO₂-90 nm, HfO₂-54 nm, W-10 nm, SiO₂-251 nm, W-10 nm, SiO₂-97 nm, Al₂O₃-74 nm, HfO₂-105 nm, W-13 nm, SiO₂-163 nm].

thermal stability. The reflectance spectra are converted to absorptivity, which is simply the inverse of the reflectance, since there is no transmission through the structure, and the absorptivity and thermal emissivity are identical due to Kirchhoff's law of thermal radiation. The same postprocessing procedure is used here and greatly enhances the results. The extended wavelength and thickness range represent a further increase in complexity for the inverse modeling, and despite this, the nested transfer procedure and postprocessing give accurate retrieval of arbitrary spectra. Previous realizations of metamaterials for selective thermal emission have employed periodic structures of alternating W and HfO₂ to produce desired spectra.⁴⁷ Our model-suggested spectrum is better able to maximize absorptivity below the bandgap wavelength λ_{PV} , leading to an improved ultimate efficiency exceeding 40% at 1000°C (1273 K), in comparison to 19% for a blackbody emitter.^{47,65} Previously, we have demonstrated the ability to independently retrieve periodic solutions

using an MDN-based inverse design architecture with fixed materials.³⁶ When more material options are added to the library, the design suggested by the model here is aperiodic, retaining the alternating appearance of metal and dielectric layers but allowing diverse arrangements of the oxides. It is also feasible to request new designs if a different trade-off is made between the absorptions above and below the PV's band edge. By expanding the choice of materials at each position, we find alternative solutions surpassing the existing designs in fulfilling the application-specific spectral requirements, and the use of nested transfer can make these more complex design spaces more feasible to model with reasonable data set sizes.

4 Discussion

We propose a method of iterative nested transfer learning to gradually build forward prediction and inverse design models

of increasing complexity while using small data sets at each step. The forward model can accurately reproduce arbitrary spectra for 30-layer thin-film stacks. This approach is extended to inverse design models which are built up from 2 to 10 layers of thin-film stacks allowing a free material choice at each layer. A postprocessing method using a pretrained forward network is used to further reinforce the design accuracy. The forward model uses a bidirectional RNN-based architecture, and the inverse model uses a convolutional MDN architecture. The complexity arising from the broad wavelength range and free material choice represents some of the most challenging tasks to model that have been demonstrated in DL-based inverse design. The accuracy of the forward model dealing with up to 30 structure layers with the same material choice is also among the most complex modeling tasks previously shown. Despite the high degree of complexity, the nested transfer method combined with postprocessing allows for accurate recreations of arbitrary spectra while keeping data requirements modest. Finally, the same architecture and training approach are applied to a modified data set for predicting designs for selective thermal emitters, generating close approximations of unrealistic idealized spectra for thermophotovoltaic applications. While the results here are restricted to thin-film stacks, this same approach of gradually building complexity with transfer learning can be extended to a wider variety of structures, where the computational requirements for generating a suitably large data set for the desired degree of complexity may not be feasible. Even structures that require full-wave simulations can quickly generate larger data sets for simplified versions with reduced degrees of freedom and continually use small data sets as the complexity and number of design variables increase. In another vein, other than transferring information to cope with increasing layer numbers, generalization of geometry to, e.g., multilayer core shells, has proved viable.⁴¹ If the problem is formulated properly, it might be possible as well to ease the augmentation of materials, benefiting the search of all dimensions of the design space. The use of RNNs in optical modeling, especially for structures and processes that can be described by sequential inputs/outputs, is also worth further exploring. We foresee that this will enable high-performance inverse design models to be built that previously would have been computationally unfeasible, allowing for new application-specific designs to be searched for.

Code and Data Availability

The data and code used in this work can be obtained from the corresponding author upon reasonable request.

Acknowledgment

The authors acknowledge the financial support of the National Institute of General Medical Sciences of the National Institutes of Health (1R01GM146962-01).

References

1. J. Fang et al., "Decoding optical data with machine learning," *Laser Photonics Rev.* **15**(2), 2000422 (2021).
2. S. D. Campbell et al., "Review of numerical optimization techniques for meta-device design," *Opt. Mater. Express* **9**(4), 1842–1863 (2019).
3. K. Yao and Y. Zheng, *Nanophotonics and Machine Learning: Concepts, Fundamentals, and Applications*, Springer, Cham, Switzerland (2023).
4. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature* **521**, 436–444 (2015).
5. E. Alpaydin, *Introduction to Machine Learning*, MIT Press (2014).
6. P. R. Wiecha and O. L. Muskens, "Deep learning meets nanophotonics: a generalized accurate predictor for near fields and far fields of arbitrary 3D nanostructures," *Nano Lett.* **20**(1), 329–338 (2020).
7. S. Verma et al., "A comprehensive deep learning method for empirical spectral prediction and its quantitative validation of nano-structured dimers," *Sci. Rep.* **13**(1), 1129 (2023).
8. I. Sajedian, J. Kim, and J. Rho, "Finding the optical properties of plasmonic structures by image processing using a combination of convolutional neural networks and recurrent neural networks," *Microsyst. Nanoeng.* **5**(1), 27 (2019).
9. A. Hussain et al., "Deep learning approach for predicting optical properties of chalcogenide planar waveguide," in *Int. Conf. Autom. Control Mechatron. Ind. 4.0 (ACMI)*, pp. 1–6 (2021).
10. J. Jiang, M. Chen, and J. A. Fan, "Deep neural networks for the evaluation and design of photonic devices," *Nat. Rev. Mater.* **6**, 679–700 (2021).
11. P. Liu et al., "Structure-embedding network for predicting the transmission spectrum of a multilayer deep etched grating," *Opt. Lett.* **47**(23), 6185–6188 (2022).
12. R. S. Hegde, "Deep learning: a new tool for photonic nanostructure design," *Nanoscale Adv.* **2**(3), 1007–1023 (2020).
13. P. R. Wiecha et al., "Deep learning in nano-photonics: inverse design and beyond," *Photonics Res.* **9**(5), B182–B200 (2021).
14. K. Yao, R. Unni, and Y. Zheng, "Intelligent nanophotonics: merging photonics and artificial intelligence at the nanoscale," *Nanophotonics* **8**(3), 339–366 (2019).
15. P. Dai et al., "Inverse design of structural color: finding multiple solutions via conditional generative adversarial networks," *Nanophotonics* **11**(13), 3057–3069 (2022).
16. D. Zhang et al., "Inverse design of an optical film filter by a recurrent neural adjoint method: an example for a solar simulator," *J. Opt. Soc. Am. B* **38**(6), 1814–1821 (2021).
17. M. Zandehshahvar et al., "Metric learning: harnessing the power of machine learning in nanophotonics," *ACS Photonics* **10**(4), 900–909 (2023).
18. W. Ma et al., "Deep learning for the design of photonic structures," *Nat. Photonics* **15**(2), 77–90 (2021).
19. D. Liu et al., "Training deep neural networks for the inverse design of nanophotonic structures," *ACS Photonics* **5**(4), 1365–1369 (2018).
20. J. Peurifoy et al., "Nanophotonic particle simulation and inverse design using artificial neural networks," *Sci. Adv.* **4**(6), eaar4206 (2018).
21. P. Liu et al., "Deep neural networks with adaptive solution space for inverse design of multilayer deep-etched grating," *Opt. Lasers Eng.* **174**, 107933 (2024).
22. Z. A. Kudyshev et al., "Machine-learning-assisted metasurface design for high-efficiency thermal emitter optimization," *Appl. Phys. Rev.* **7**(2), 021407 (2020).
23. Z. Liu et al., "Generative model for the inverse design of metasurfaces," *Nano Lett.* **18**(10), 6560–6576 (2018).
24. Z. Liet et al., "Empowering metasurfaces with inverse design: principles and applications," *ACS Photonics* **9**(7), 2178–2192 (2022).
25. J. Zhang et al., "Experiment-based deep learning approach for power allocation with a programmable metasurface," *APL Mach. Learn.* **1**(4), 046122 (2023).
26. W. Ma et al., "Pushing the limits of functionality-multiplexing capability in metasurface design based on statistical machine learning," *Adv. Mater.* **34**(16), 2110022 (2022).
27. A. Ueno et al., "Dual-band optical collimator based on deep-learning designed, fabrication-friendly metasurfaces," *Nanophotonics* **12**(17), 3491–3499 (2023).
28. A. Sheverdin, F. Monticone, and C. Valagiannopoulos, "Photonic inverse design with neural networks: the case of invisibility in the visible," *Phys. Rev. Appl.* **14**(2), 024054 (2020).

29. Y. Chen et al., “Physics-informed neural networks for inverse problems in nano-optics and metamaterials,” *Opt. Express* **28** (8), 11618–11633 (2020).
30. A.-P. Blanchard-Dionne and O. J. F. Martin, “Successive training of a generative adversarial network for the design of an optical cloak,” *OSA Contin.* **4**(1), 87–95 (2021).
31. Y. Kiarashinejad, S. Abdollahramezani, and A. Adibi, “Deep learning approach based on dimensionality reduction for designing electromagnetic nanostructures,” *NPJ Comput. Mater.* **6**(1), 12 (2020).
32. M. Chen et al., “High speed simulation and freeform optimization of nanophotonic devices with physics-augmented deep learning,” *ACS Photonics* **9**(9), 3110–3123 (2022).
33. W. Ma, F. Cheng, and Y. Liu, “Deep-learning-enabled on-demand design of chiral metamaterials,” *ACS Nano* **12**(6), 6326–6334 (2018).
34. Z. Liu et al., “Compounding meta-atoms into metamolecules with hybrid artificial intelligence techniques,” *Adv. Mater.* **32**(6), 1904790 (2020).
35. C. Zhu et al., “Machine learning aided design and optimization of thermal metamaterials,” *Chem. Rev.* **124**(7), 4258–4331 (2024).
36. R. Unni et al., “A mixture-density-based tandem optimization network for on-demand inverse design of thin-film high reflectors,” *Nanophotonics* **10**(16), 4057–4065 (2021).
37. K. Weiss, T. M. Khoshgoftaar, and D. Wang, “A survey of transfer learning,” *J. Big Data* **3**, 9 (2016).
38. S. J. Pan and Q. Yang, “A survey on transfer learning,” *IEEE Trans. Knowl. Data Eng.* **22**(10), 1345–1359 (2010).
39. Z. Fan et al., “Transfer-learning-assisted inverse metasurface design for 30% data savings,” *Phys. Rev. Appl.* **18**(2), 024022 (2022).
40. L. Cheng, P. Singh, and F. Ferranti, “Transfer learning-assisted inverse modeling in nanophotonics based on mixture density networks,” *IEEE Access* **12**, 55218–55224 (2024).
41. Y. Qu et al., “Migrating knowledge between physical scenarios based on artificial neural networks,” *ACS Photonics* **6**(5), 1168–1174 (2019).
42. S. Yi, S. Xu, and W. Zou, “Multi-band low-noise microwave-signal-receiving system with a photonic frequency down-conversion and transfer-learning network,” *Opt. Lett.* **46**(23), 5982–5985 (2021).
43. M. Kaya and S. Hajimirza, “Using a novel transfer learning method for designing thin film solar cells with enhanced quantum efficiencies,” *Sci. Rep.* **9**(1), 5034 (2019).
44. C. Qiu et al., “Nanophotonic inverse design with deep neural networks based on knowledge transfer using imbalanced datasets,” *Opt. Express* **29**(18), 28406–28415 (2021).
45. R. Zhu et al., “Phase-to-pattern inverse design paradigm for fast realization of functional metasurfaces via transfer learning,” *Nat. Commun.* **12**(1), 2974 (2021).
46. J. Zhang et al., “Heterogeneous transfer-learning-enabled diverse metasurface design,” *Adv. Opt. Mater.* **10**(17), 2200748 (2022).
47. P. N. Dyachenko et al., “Controlling thermal emission with refractory epsilon-near-zero metamaterials via topological transitions,” *Nat. Commun.* **7**(1), 11809 (2016).
48. L. Alzubaidi et al., “Review of deep learning: concepts, CNN architectures, challenges, applications, future directions,” *J. Big Data* **8**(1), 53 (2021).
49. H. Hewamalage, C. Bergmeir, and K. Bandara, “Recurrent neural networks for time series forecasting: current status and future directions,” *Int. J. Forecast.* **37**(1), 388–427 (2021).
50. Ö. Batur Dİnler and N. Aydin, “An optimal feature parameter set based on gated recurrent unit recurrent neural networks for speech segment detection,” *Appl. Sci.* **10**, 1273 (2020).
51. A. Jagannatha and H. Yu, *Structured Prediction Models for RNN Based Sequence Labeling in Clinical Text*, pp. 856–865, Association for Computational Linguistics, Austin, Texas (2016).
52. F. Informatik et al., “Gradient flow in recurrent nets: the difficulty of learning long-term dependencies,” in *A Field Guide to Dynamical Recurrent Neural Networks*, Wiley- IEEE Press, pp. 237–243 (2001).
53. A. Sherstinsky, “Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network,” *Phys. D: Nonlinear Phenom.* **404**, 132306 (2020).
54. A. Graves and J. Schmidhuber, “Framewise phoneme classification with bidirectional LSTM networks,” in *Proc. 2005 IEEE Int. Joint Conf. Neural Networks*, Vol. 4, pp. 2047–2052 (2005).
55. L. V. Rodríguez-de Marcos et al., “Self-consistent optical constants of MgF₂, LaF₃, and CeF₃ films,” *Opt. Mater. Express* **7**(3), 989–1006 (2017).
56. M. Born and E. Wolf, *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*, Cambridge University Press (1999).
57. R. Unni, K. Yao, and Y. Zheng, “Deep convolutional mixture density network for inverse design of layered photonic structures,” *ACS Photonics* **7**(10), 2703–2712 (2020).
58. W. Li and S. Fan, “Nanophotonic control of thermal radiation for energy applications,” *Opt. Express* **26**(12), 15995–16021 (2018).
59. A. Narayanaswamy and G. Chen, “Thermal emission control with one-dimensional metallodielectric photonic crystals,” *Phys. Rev. B* **70**(12), 125101 (2004).
60. O. Ilic et al., “Tailoring high-temperature radiation and the resurrection of the incandescent source,” *Nat. Nanotechnol.* **11**(4), 320–324 (2016).
61. S. Makoto, K. Asaka, and Y. Hiroo, “High-efficiency solar-thermophotovoltaic system equipped with a monolithic planar selective absorber/emitter,” *J. Photonics Energy* **5**(1), 053099 (2015).
62. N. P. Sergeant et al., “Design of wide-angle solar-selective absorbers using aperiodic metal-dielectric stacks,” *Opt. Express* **17**(25), 22800–22812 (2009).
63. M. W. Dashiell et al., “Quaternary InGaAsSb thermophotovoltaic diodes,” *IEEE Trans. Electron Devices* **53**(12), 2879–2891 (2006).
64. M. Bosi and C. Pelosi, “The potential of III-V semiconductors as terrestrial photovoltaic devices,” *Progr. Photovolt. Res. Appl.* **15**(1), 51–68 (2007).
65. W. Shockley and H. J. Queisser, “Detailed balance limit of efficiency of p-n junction solar cells,” *J. Appl. Phys.* **32**(3), 510–519 (1961).

Rohit Unni received his PhD in materials science from the University of Texas at Austin in 2024 and his bachelor’s degree from Washington University in St. Louis in 2016. His research interests incorporate bridging nanophotonics and machine learning from multiple directions, including inverse design, computer vision, and next generation foundation models.

Kan Yao is currently a research fellow in the University of Texas at Austin. He received his PhD in Electrical Engineering from Northeastern University, Boston, USA, in 2017. His research interests span various areas of photonics, such as plasmonics, metamaterials and metasurfaces, light-matter interactions, chiroptics, quantum photonics, and device design.

Yuebing Zheng is a professor at the University of Texas at Austin. He holds the Cullen Trust for Higher Education Endowed Professorship in Engineering. He received his PhD from Pennsylvania State University in 2010 and did postdoctoral research at the University of California, Los Angeles from 2010 to 2013. His research is at the forefront of optics and photonics, where they innovate optical manipulation and measurement to transform scientific research and tackle pressing global challenges.