



# Deep neural networks for the evaluation and design of photonic devices

Jiaqi Jiang, Mingkun Chen and Jonathan A. Fan

**Abstract** | The data-science revolution is poised to transform the way photonic systems are simulated and designed. Photonic systems are, in many ways, an ideal substrate for machine learning: the objective of much of computational electromagnetics is the capture of nonlinear relationships in high-dimensional spaces, which is the core strength of neural networks. Additionally, the mainstream availability of Maxwell solvers makes the training and evaluation of neural networks broadly accessible and tailorable to specific problems. In this Review, we show how deep neural networks, configured as discriminative networks, can learn from training sets and operate as high-speed surrogate electromagnetic solvers. We also examine how deep generative networks can learn geometric features in device distributions and even be configured to serve as robust global optimizers. Fundamental data-science concepts framed within the context of photonics are also discussed, including the network-training process, delineation of different network classes and architectures, and dimensionality reduction.

Photonics has deep utility in many scientific and technological domains. Integrated photonic systems enable novel light sources<sup>1,2</sup>, classical communication platforms<sup>3,4</sup> and quantum-optical processors<sup>5</sup>. Subwavelength-scale metal and dielectric composites can scatter and localize electromagnetic waves in customizable and extreme ways, producing new modalities in sensing<sup>6</sup> and light–matter interactions<sup>7</sup>. Metamaterials<sup>8</sup> and metasurfaces<sup>9</sup>, which are structured to perform wavefront-engineering tasks, are transforming the field of optical engineering. Much of the versatility of photonics can be attributed to the strong relationship between device geometry and optical response, which allows many optical properties and functions to be tailored using structured dielectric and metallic materials.

Photonic systems are typically analysed through the framework of one of two problems. The first is the forward problem: given a structure, what is the electromagnetic response? This is the easier of the two problems and can be solved using one of a number of well-established numerical electromagnetic simulators<sup>10,11</sup>. These simulators can accurately evaluate Maxwell's equations, but it is challenging to manage computational resources when evaluating large simulation domains and large batches of simulations. The second is the inverse problem: given a desired electromagnetic response, how does one design a suitable photonic structure? The solution to the inverse problem cannot be directly evaluated and is challenging to solve because the solution space is non-convex,

meaning there exist many local optima. Approaches to this problem are often framed as an exercise in optimization<sup>12</sup> and include simulated annealing<sup>13,14</sup>, evolutionary<sup>15,16</sup>, objective-first<sup>17</sup> and adjoint-variables algorithms<sup>18,19</sup>. Although great strides have been made in solving the inverse problem, identifying the best overall device given a desired objective and constraints remains challenging.

Deep neural networks are a versatile class of machine-learning algorithms that utilize the serial stacking of nonlinear processing layers to enable the capture and modelling of highly nonlinear data relationships. They offer a fresh perspective on the forward and inverse problems, owing to the ability of neural networks to mimic nonlinear physics-based relationships, such as those between photonic-system geometries and their electromagnetic responses. The application of deep neural networks for deep learning is a fashionable area of research, which makes it difficult to separate the hype from the true utility. In spite of the hype, deep learning has the potential to strongly impact the simulation and design process of photonic technologies for a number of reasons.

First, deep learning is a proven method for the capture, interpolation and optimization of highly complex phenomena in many fields, ranging from robotic controls<sup>20</sup> and drug discovery<sup>21</sup> to image classification<sup>22</sup> and language translation<sup>23</sup>. These algorithms will only get more powerful, particularly given the recent explosive growth of the field of data science.

Department of Electrical Engineering, Stanford University, Stanford, CA, USA.  
✉e-mail: [jonfan@stanford.edu](mailto:jonfan@stanford.edu)  
<https://doi.org/10.1038/s41578-020-00260-1>

Second, deep learning is broadly accessible. Software, ranging from TensorFlow<sup>24</sup> to PyTorch<sup>25</sup>, is open source and free to use, meaning that anyone can start implementing and training neural networks. Furthermore, researchers in the machine-learning community practice a culture of openness and sharing, making many state-of-the-art algorithms openly available and easy to access. There are also many educational resources, including curricula at universities and online courses, to help researchers get up to speed with the theory and implementation of neural networks.

Third, photonic structures can be readily evaluated with a broad range of electromagnetic-simulation tools. These widely available tools enable the quantification of the near-field and far-field electromagnetic response of a structure, which facilitates the solution of inverse problems. In addition, they can be used to calculate analytic and numerical gradients, such as the impact of dielectric perturbations on a desired figure of merit. As we will discuss, the computation of such gradient terms can be combined with deep learning to yield entirely new and effective modalities of inverse design, such as global topology optimization. The implementation of electromagnetic-simulation-software tools in conjunction with deep-learning-programming packages is streamlined using application programming interfaces that come with much mainstream computational software<sup>26–28</sup>.

Fourth, there are broadly accessible computational resources that enable large numbers of electromagnetic simulations to be performed, which plays into the strengths of deep-learning approaches. Distributed computing, including cloud-based platforms, allows anyone with an internet connection to parallelize many simulations<sup>29</sup>. Additionally, the advent of new computing-hardware platforms, such as those based on graphical<sup>30</sup> and tensor<sup>31</sup> processing units, will increase the computational efficiency and capacity for both electromagnetic simulations and neural-network training.

In this Review, we first provide an overview of discriminative and generative neural networks, of how artificial neural networks are formulated and how different electromagnetic phenomena can be modelled through the processing of different data-structure types. Second, we discuss how deep discriminative networks can serve as surrogate models for electromagnetic solvers and expedite the solving of forward and inverse problems. Third, we show how generative networks are a natural framework for population-based inverse design and can be configured to perform the global optimization of nanophotonic devices. Finally, we compare and contrast deep-learning methods with classical modelling tools for electromagnetic problems, discuss pathways for future research and suggest research practices that can accelerate progress in this field. Several recent reviews on machine learning and photonics<sup>32–34</sup> cover topics complementary to those discussed here, including reinforcement learning<sup>33,35–37</sup> and the application of photonics hardware to machine-learning computation<sup>32,38</sup>. This Review provides both a tutorial introduction to basic machine-learning concepts and a comprehensive guide to current research developments. By covering

electromagnetic technologies spanning microwave to optical frequencies, it offers a broad conceptual overview of the topic.

### Principles of deep neural networks

Deep neural networks consist of multiple layers of neurons connected in series. A neuron is a mathematical function that takes one or more values as its input, performs a nonlinear operation on a weighted sum of those input values and yields a single output value (BOX 1). With layer-by-layer processing of data inputted to the network, data features with higher levels of abstraction are captured from lower-level features, and complex network input–output relations can be fitted. To train a neural network, a large training data set is first generated using electromagnetic simulations. The training data aid in the iterative adjustment of the neuron weights until the network correctly captures the data distributions in the training set. Modifications to the network weights are performed using a process called backpropagation that minimizes the network loss function, which specifies the deviation of the network output from the ground-truth training set. These terms and concepts are discussed in more detail in BOX 2.

Photonic devices modelled by neural networks are described by two types of labels (FIG. 1a). The first type comprises the physical variables describing the device, and it includes the device geometry, material and electromagnetic excitation source. These labels are delineated by the variable  $x$ . The second type describes the physical responses corresponding to a range of spectral and performance characteristics. These labels are delineated by the variable  $y$ . In electromagnetics, physical responses can be described as a single-valued function of physical variables, so that a given input  $x$  maps to a single  $y$ . For example, a thin-film stack with a fixed geometric and material configuration produces a single transmission spectrum. However, the opposite is not true: for most problems in electromagnetics, a given physical response  $y$  maps to multiple  $x$ s. For example, the same transmission spectrum can be produced using different thin-film stack configurations. As a result, different classes of neural networks need to be considered, depending on the type of device labels processed by the network. For electromagnetics, the two most commonly used network classes are discriminative and generative networks.

### Discriminative and generative networks

Discriminative networks are capable of regression and classification tasks, and can specify complex, nonlinear mapping relations between inputs and outputs<sup>39</sup>. For regression tasks, discriminative networks can interpolate relationships within training data, and the relationship between input and output mappings is that of a single-valued function that can support one-to-one or many-to-one mappings. As such, discriminative neural networks can capture the relationship  $y=f(x)$  and serve as surrogate physical models that solve the forward problem (FIG. 1b). Compared with a numerical electromagnetic solver, a trained discriminative network can evaluate the forward problem orders of magnitude faster<sup>40</sup>.

## Box 1 | Building blocks of artificial neural networks

The fundamental unit in an artificial neural network is the neuron, which receives an input vector  $\mathbf{x}$  and outputs a scalar value  $y$ , as illustrated in the figure. The neuron performs two mathematical operations, a weighted sum followed by a nonlinear mapping<sup>83</sup>. The weighted sum  $a$  is calculated as  $a = \sum_i w_i x_i + b = \mathbf{w}^T \mathbf{x} + b$ , where  $\mathbf{w}$  is a trainable weight vector that possesses the same dimension as the input  $\mathbf{x}$  and  $b$  is a trainable bias term. The nonlinear mapping applied to  $a$  is performed using a continuous, differentiable, nonlinear function  $f$ , known as the activation function, such that the output of the neuron is  $y = f(a)$ . Some of the most widely used activation functions are the sigmoid function, the rectified linear unit function and the hyperbolic tangent function.

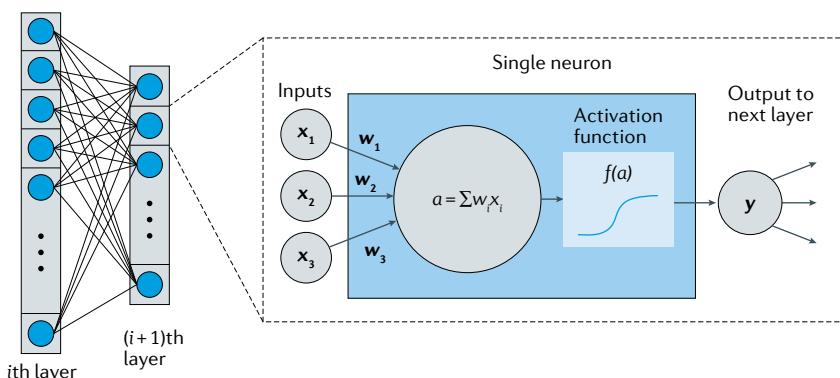
Neurons can be connected in different ways to form different modules within a deep neural network. The most commonly used modules in electromagnetics are the fully connected (FC) and convolutional layers. FC layers comprise a vector of neurons, and the inputs to each neuron in one layer are the output values from every neuron in the prior layer. The number of neurons in each layer can be arbitrary and differ between layers.

Upon stacking a large number of FC layers in sequence, the relationship between the input and output data can be specified to be increasingly


more complex. The nonlinear activation function  $f$  applied within each layer ensures that such stacking of FC layers leads to added computational complexity that cannot be captured by just a single layer.

Convolutional layers are designed to capture local spatial features within data. Convolutional layers are typically used to process image-based data structures, though they can be generally applied to vector, matrix and tensor data structures. For a convolutional layer processing a 2D image, a kernel filter is spatially displaced over the width and height of the input image with a constant displacement step. The kernel filter is a small matrix with trainable weights, and it processes small groups of neighbouring image elements. At each kernel position, a single value in the output matrix is computed using the same operations as in a single neuron: a dot product between the image elements and kernel filter, followed by a nonlinear activation. As each value in the output matrix derives from a small region of the input image, the output matrix is typically referred to as a feature map that highlights regions of the input image with kernel-like local features. For most convolutional layers, an input matrix is processed with multiple kernels, each producing a unique feature map. These maps are then stacked together to produce an output tensor.

## Fully connected layers



## Convolutional layers



Deep generative neural networks appear deceptively similar to discriminative neural networks, utilizing the same concepts in deep-network architecture and neuron-based data processing. The key difference is that one of the inputs to the network is a latent variable,  $\mathbf{z}$ , which is a random variable internal to the network. The term 'latent' refers to the fact that this variable does not have an explicit physical meaning. In generative neural networks, the latent variables are sampled from a standard probability distribution, such as a uniform or Gaussian distribution. A single instance of latent-variable sampling maps to a single network output, whereas a continuum of latent-variable samplings maps to a distribution of network outputs. The neural network can, therefore, be regarded as a function that maps a standard probability distribution to a complex output distribution.

In photonics, generative networks are typically configured to output a distribution of device layouts (FIG. 1c). The network is conditional, and its inputs include the latent variable,  $\mathbf{z}$ , and a set of labels,  $\theta$ , which is a subset of all device labels and can include physical variables and physical responses. These networks are termed 'conditional' because the outputted distributions can be considered as probability distributions conditioned on  $\theta$ . The network learns from a training set consisting of an ensemble of discrete-labelled devices, which

can be treated as samples from the distribution  $\hat{P}(\mathbf{x}|\theta)$ . A properly trained network outputs a device distribution  $P(\mathbf{x}|\theta)$  that matches the training-set distribution. Unconditional networks, for which the only input is the latent variable, can also be trained, and these networks generate devices that match an unlabelled training-set distribution. Schemes for performing inverse design with generative networks without the use of training sets also exist<sup>41,42</sup>, and are discussed later in this Review.

The stochastic nature of generative neural networks distinguishes them from discriminative networks. Whereas discriminative networks can capture the relationship between device layouts and optical response from a training set, generative networks focus on learning the properties of the device-layout distributions themselves<sup>43–48</sup>. Moreover, for a given input value of  $\theta$ , generative networks produce a distribution of outputs and, therefore, perform one-to-many mappings. We note that there are also classes of generative networks that do not utilize latent random variable inputs<sup>49</sup>, but these have limited capabilities and are not widely used to model photonic systems.

In both discriminative and generative models, the physical variable inputs and physical response outputs must be formulated as discretized data structures. Such a representation is in contrast to the continuous form of many real-world input–output types, such as freeform

Box 2 | Training of artificial neural networks

When a neural network is initialized, all neurons have randomly assigned weights. To properly specify the weights in a manner that captures a desired input–output relationship, the network weights are iteratively adjusted to push the network input–output relationship towards the values specified in the training set. This training objective can be framed as the minimization of a loss function, which quantifies the difference between the outputs of the network and the ground-truth values from the training set.

For discriminative networks performing regression, consider terms in the training set to be  $(\mathbf{x}, \hat{y})$  and the outputs of the network, given network inputs of  $\mathbf{x}$  from the training set, to be  $y$ . A common loss function for this problem is the mean squared error:

$$L(y, \hat{y}) = \frac{1}{N} \sum_{n=1}^N \left( y^{(n)} - \hat{y}^{(n)} \right)^2 \tag{B1}$$

where  $N$  is the batch size. If  $N$  is equal to the training set size, the entire training set is used in each iteration, and the training process is termed batch gradient descent. If  $N$  is equal to one, a single training set term is randomly sampled at each iteration, and the training process is termed stochastic gradient descent. If  $N$  is smaller than the training set size but greater than one, a fraction of the training set is randomly sampled at each iteration, and the training process is termed mini-batch gradient descent. This training process is the one typically used in practice, as it provides a good approximation of the gradient calculated using the entire training set while balancing the computational cost of network training.

To understand the loss function form in generative networks, the training set and generated devices must be treated in the context of probability distributions. The training set devices  $\{\mathbf{x}_i\}$  can be regarded as samples from the desired probability distribution spanning the design space,  $S$ , denoted as  $\hat{P}(\mathbf{x})$ . This distribution represents the probability that device  $\mathbf{x}$  is chosen upon a random sampling of a device from the design space. Similarly, the distribution of devices produced by the generative network can be treated as a probability distribution spanning the design space and is denoted as  $P(\mathbf{x})$ . This distribution represents the probability that device  $\mathbf{x}$  is generated by the network upon a sampling of the input latent random variable.

The goal of the training process is to match the distribution of outputted device layouts,  $P(\mathbf{x})$ , with the statistical distribution of structures within the training set,  $\hat{P}(\mathbf{x})$ . To accomplish this objective, the network loss function should quantify the dissimilarity between the two distributions. One such function is the Kullback–Leibler (KL) divergence<sup>184</sup>, also known as relative entropy, which is a metric from information theory that quantifies how different one probability distribution is from another. It is defined as:

$$D_{\text{KL}}(\hat{P}||P) = \int_S \hat{P}(\mathbf{x}) \log \frac{\hat{P}(\mathbf{x})}{P(\mathbf{x})} d\mathbf{x} \tag{B2}$$

Another related function is the Jensen–Shannon (JS) divergence<sup>185</sup>, which is a symmetric function defined as the average KL divergence of  $\hat{P}(\mathbf{x})$  and  $P(\mathbf{x})$  from their mixed distribution  $(\hat{P}(\mathbf{x}) + P(\mathbf{x}))/2$ :

$$D_{\text{JS}}(\hat{P}, P) = \frac{1}{2} D_{\text{KL}}\left(\hat{P}||\frac{\hat{P}+P}{2}\right) + \frac{1}{2} D_{\text{KL}}\left(P||\frac{\hat{P}+P}{2}\right) \tag{B3}$$

Both the KL divergence and JS divergence are minimized when  $P(\mathbf{x})$  and  $\hat{P}(\mathbf{x})$  are the same. Therefore, specifying either term as the loss function fulfils the network-training objective.

For both discriminative and generative networks, backpropagation is used to calculate adjustments to the network weights that lead to a reduction in the loss function,  $\nabla_{\mathbf{w}}L$ . To visualize backpropagation in a simple example, consider a discriminative network comprising a single neuron (BOX 1) that is being trained using stochastic gradient descent. We start from  $\nabla_y L$ , which is the gradient of the loss function with respect to  $y$  and ‘propagate’ the gradient back to  $\mathbf{w}$  using the chain rule. To compute  $\frac{\partial L}{\partial \mathbf{w}}$ , we first compute  $\frac{\partial L}{\partial a} = \frac{\partial L}{\partial y} \cdot \frac{\partial y}{\partial a}$ , which specifies how  $a$  in  $f(a)$  should be adjusted to reduce the loss function. We then compute  $\frac{\partial L}{\partial \mathbf{w}} = \frac{\partial L}{\partial a} \cdot \frac{\partial a}{\partial \mathbf{w}}$ , which specifies how  $\mathbf{w}$  in  $a(\mathbf{w})$  should be adjusted. Note that, for the chain rule to work, all of the mathematical functions involved must be differentiable. Backpropagation readily generalizes to deep networks comprising many layers of connected neurons, so that  $\nabla_{\mathbf{w}}L$  can be calculated for every neuron. Once  $\nabla_{\mathbf{w}}L$  is calculated for all neurons, all weight vectors are updated by gradient descent:  $\mathbf{w} := \mathbf{w} - \alpha \nabla_{\mathbf{w}}L$ , where  $\alpha$  is the learning rate. For mini-batch gradient descent,  $\nabla_{\mathbf{w}}L$  is calculated for each sampled training set term, and these gradients are summed up at each neuron to produce a single term for gradient descent.

device layouts and time-sequential events. Nonetheless, these forms can be readily discretized in numerical representation without loss of generality, because Maxwell’s equations can be accurately discretized.

Data structures

There are distinct data structures in electromagnetics that describe a broad range of phenomena. In this section, we discuss four types of data structures: vectors that describe discrete parameters, images that describe freeform devices, graphs that describe interacting structures and time sequences that describe time-dependent phenomena. Network architectures and layer configurations are subsequently tailored, depending on the data-structure type.

**Discrete data structures.** For relatively basic photonic-structure layouts, geometric and optical properties can be described by a vector of discrete parameters. Some of these parameters are summarized in FIG. 1a and include the height, width and period of a device geometry, the permittivity and permeability of a material, and the wavelength and angle of an electromagnetic-excitation source. Many optical properties can also be described as discrete parameters and include device efficiency, quality factor, bandgap and spectral response sampled at discrete points. Discrete data structures naturally interface with neural network layers that are fully connected (BOX 1). If the objective of the network is to relate discrete input and output data structures, a deep fully connected network architecture often suffices.

**Image data structures.** Many photonic devices have freeform geometries that cannot be parameterized by a few discrete variables but are best described as 2D or 3D images. An example is a freeform metagrating that diffracts incident light to specific orders (FIG. 2a) and is described as a pixelated image with thousands of voxels. Image data types are effectively processed using a set of convolutional layers in series (BOX 1) that can extract and process spatial features<sup>50–53</sup>. Neural networks that process image data structures using convolutional layers are termed convolutional neural networks (CNNs). If the objective of the network is to output an image data structure, such as the internal polarization or field distributions within a device, a CNN comprising all convolutional layers can be implemented<sup>54</sup>. If the objective of the network is to output a discrete data structure, such as the spectral response or efficiency of a device, high-level feature maps from a series of convolutional layers can be joined with fully connected layers for processing and conversion to the proper data structure.

**Graph data structures.** For electromagnetic systems consisting of physically interacting discrete objects, graphs are ideal data-structure representations. As an example, consider an on-chip photonics system consisting of ring resonators coupled in the near field (FIG. 2b). The physical attributes of each resonator are embedded in a node of the graph, and edges between two nodes in the graph describe the near-field interactions between two neighbouring resonators. The graph structure can be irregular,

meaning that different nodes can connect with different sets of neighbours. For the example shown here, neighbouring nodes connect only when there is a significant near-field coupling.


Graph data structures are suitably processed in graph neural networks (GNNs)<sup>55–57</sup>, which analyse and operate on aggregated information between neighbouring nodes in each layer. GNNs can learn the physical interactions between nodes through the training process and are able to generalize the nature of these interactions to different configurations of neighbouring nodes. As more layers in the network are stacked, the interactions between nodes that are more distant from each other, such as nearest-nearest neighbours, are accounted for and learned. The outputs of the GNN are abstracted representations of the node and edge properties and the graph structure, and they can be further processed using fully connected layers to output a desired discrete

physical response. Although GNNs have not yet been extensively studied in the context of photonics, they have been applied to a broad range of physical systems, including the modelling of phase transitions in glasses<sup>58</sup>, molecular fingerprint analysis<sup>59</sup> and molecular drug discovery<sup>60</sup>. The network architectures are highly specialized depending on the application, and include graph attention networks<sup>57</sup>, graph recurrent networks<sup>61</sup> and graph generative networks<sup>62</sup>.

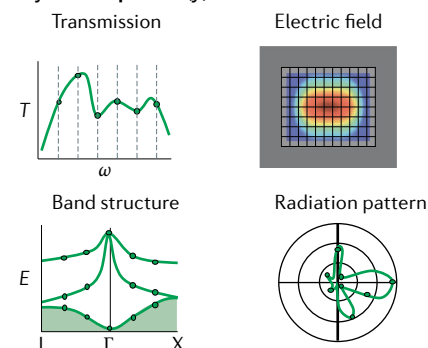
**Time-sequence data structures.** For dynamical electromagnetic systems, the physical variables and responses can be described in terms of time sequences. These continuous electromagnetic phenomena can be represented in terms of discrete time sequences without loss of generality, as long as the discrete time steps are sufficiently small. Consider electromagnetic-wave propagation in a waveguide modulated by a ring resonator (FIG. 2c).

### a Device variables ( $x, y$ )


#### Physical variables ( $x$ )




#### Physical responses ( $y$ )



### b Discriminative neural network: $f: x \rightarrow y$



### c Generative neural network: $f: (z|\theta) \rightarrow x, P_z(z) \rightarrow P(x|\theta)$



**Fig. 1 | Overview of deep learning for photonics. a** | Photonic devices are described by two types of labels, physical variables,  $x$ , and physical responses,  $y$ . **b** | Discriminative neural networks can serve as surrogate physical models. The trained network models the relationship  $y = f(x)$ , which matches with the discrete values  $(x, \hat{y})$  in the training set. **c** | Generative neural networks map latent variables,  $z$ , and conditioning labels,  $\theta$ , to a distribution of generated devices,  $P(x|\theta)$ . Upon training, the network matches  $P(x|\theta)$  to the training set distribution,  $\hat{P}(x|\theta)$ .

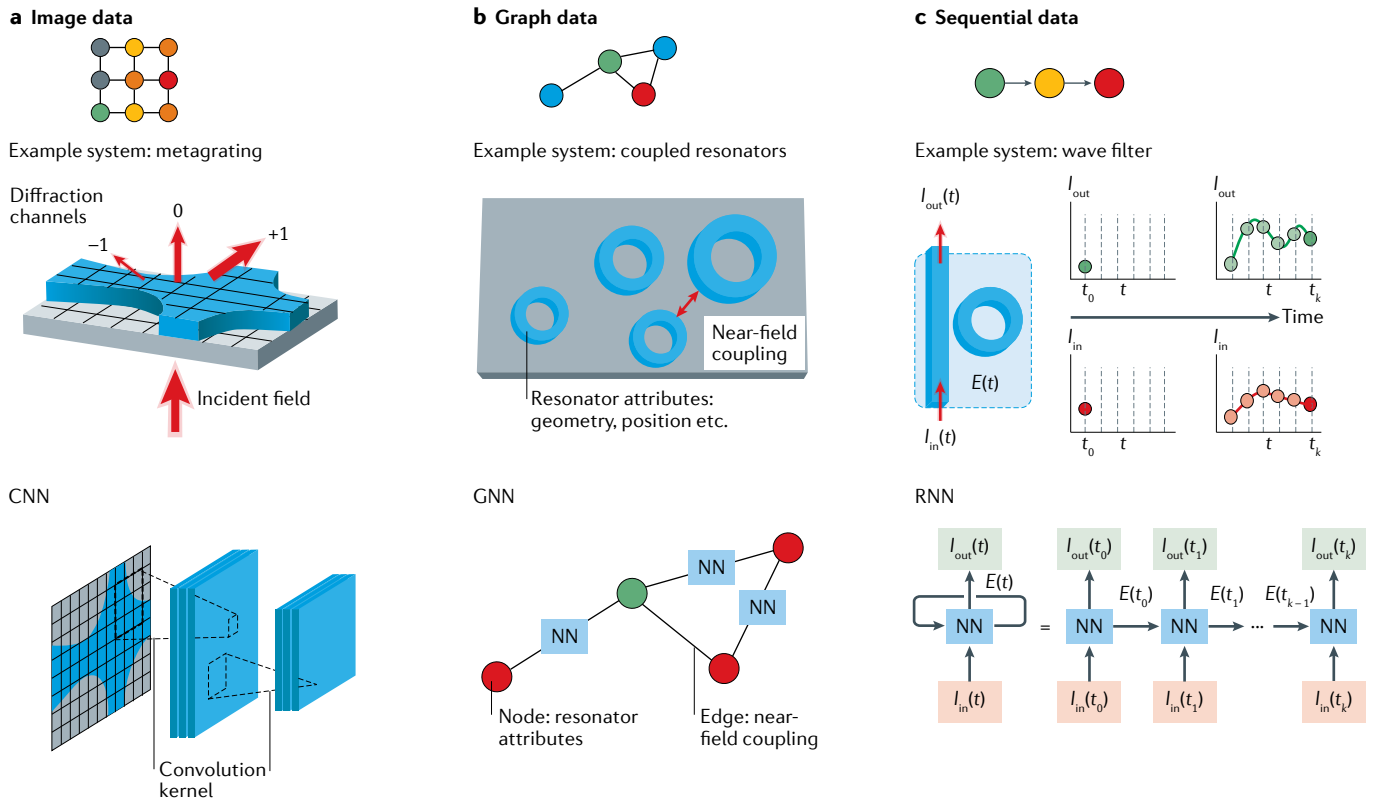


Fig. 2 | **Overview of data structures for photonics.** **a** | Image data structures are well suited to represent freeform device layouts, such as those of curvilinear nanostructures. Images are typically processed using convolutional layers in a convolutional neural network (CNN), which uses small 2D and 3D kernels to transform images into feature-map representations. **b** | Graph data structures are well suited to represent photonic systems with interacting parts, such as nanostructures coupled in the near field. Graphs are processed in graph neural networks (GNNs) that can generalize the optical response of a structure based on its node and edge attributes. **c** | Time-sequential phenomena can be captured using recurrent neural networks (RNNs). In the wave-filter example shown here, the output of the network,  $I_{out}$ , is a function of the input intensity,  $I_{in}$ , and the internal state of the network, which learns  $E(t)$ . The depiction of the RNN unrolled in time shows how the internal state of the network updates after each time step. NN, neural network.

Both the input and output signals are time sequences, and the output at a given time not only depends on the input signal at that time but also on the state of the device (its internal electric fields) at the previous time step.

Recurrent neural networks (RNNs) feed the network outputs back into the input layer, maintaining a memory that accounts for the past state of the system, which makes them ideally suited to model time-sequential systems<sup>53–66</sup>. RNNs are more general than CNNs and GNNs, and can be adapted to all the network architectures described earlier, allowing them to process discrete, image and graph data structures. For the RNN unrolled in time in the example in FIG. 2c, we see that, at the  $t_k$  time step, the current signal  $I_{in}(t_k)$  and previous electromagnetic field  $E(t_{k-1})$  are network inputs, which the RNN processes to update its state  $E(t_k)$  and output the signal  $I_{out}(t_k)$ . Whereas the output of the network varies in time, the neural network itself does not change, owing to the time-translation invariance of Maxwell’s equations. It is noted that RNNs are particularly well suited for electromagnetic-wave-phenomena modelling because the master equations describing the recurrence relations in RNNs have an exact correspondence to the equations describing wave propagation in the time domain<sup>67</sup>.

**Forward and inverse discriminative models**  
**Devices modelled by discriminative networks**

Initial demonstrations of the neural-network modelling of electromagnetic devices date back to the early 1990s in the microwave community (FIG. 3a). Microwave circuits are an analogue to nanophotonic structures, as they are based on components described by the subwavelength limits of Maxwell’s equations. Amongst the first published implementations was the use of a Hopfield neural network, which is a RNN, for microwave impedance matching<sup>68</sup>. Through an iterative process, the network could specify how changes in stub position and length could improve network matching. Network weights were determined not through a training process but by known relationships between stub position and network matching determined from simulations. Although this demonstration did not entail a classical trainable discriminative network, it showcased the early potential of neural networks in electromagnetics problems.

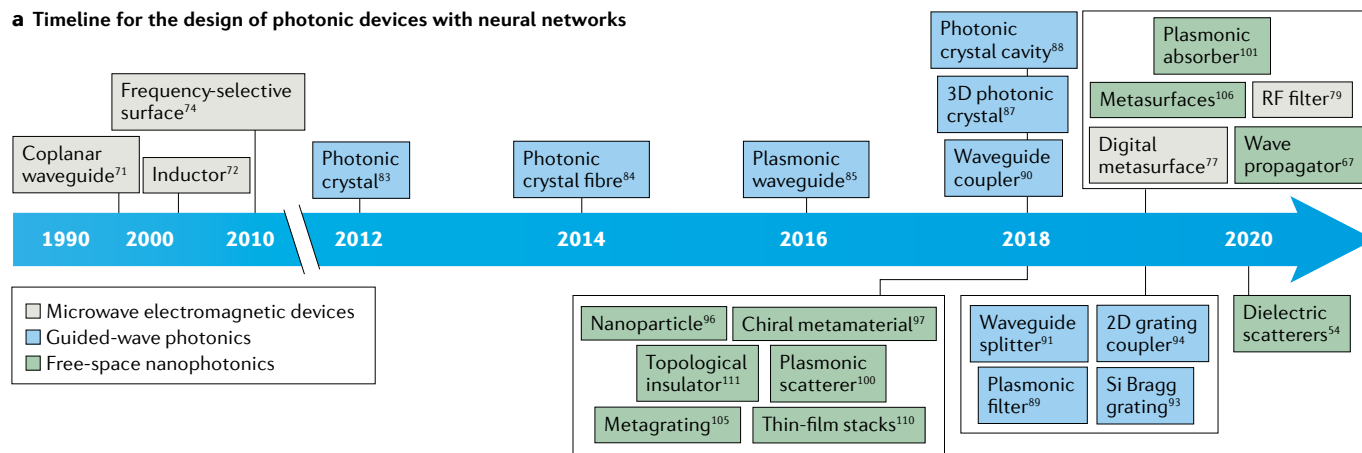
Shortly thereafter, deep fully connected discriminative networks with at least two layers were used to model more complex microwave circuit elements, including metal–semiconductor field-effect transistors<sup>69</sup>, heterojunction bipolar transistor amplifiers<sup>70</sup>, coplanar

waveguide components<sup>71</sup> and lumped 3D integrated components such as capacitors and inductors<sup>72</sup>. As the complexity of the devices increased, accurate modelling required a concept termed space mapping, in which a deep network that learns spatially coarse device features is paired with one that learns spatially fine features<sup>72,73</sup>. Over the last decade, the scope of deep discriminative networks for microwave technologies has expanded to include frequency-selective surfaces<sup>74,75</sup>, metamaterials<sup>76</sup>, metasurfaces<sup>77</sup> and filters<sup>78,79</sup>. More details on deep-learning developments in the microwave community can be found in a number of reviews<sup>80–82</sup> and are

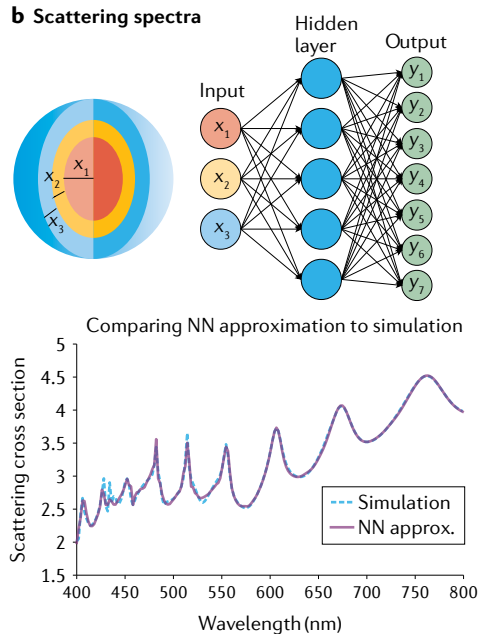
worth noting here because of their potential to translate to photonic systems.

Researchers working on silicon photonics and optical fibres started exploring the neural-network modelling of guided-wave systems in the early 2010s, around the time when deep learning as a field started undergoing tremendous growth (FIG. 3a). Initial deep discriminative networks utilized two to three total neural layers and could learn the bandgap properties of simple photonic crystals<sup>83</sup>, the dispersion properties of photonic crystal fibres<sup>84</sup> and the propagation characteristics of plasmonic transmission lines<sup>85</sup>. More recent demonstrations

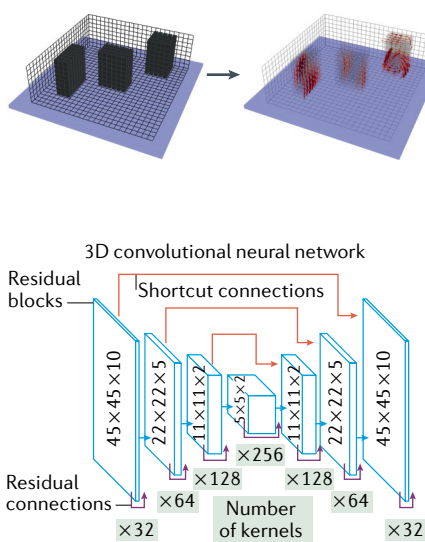
### a Timeline for the design of photonic devices with neural networks



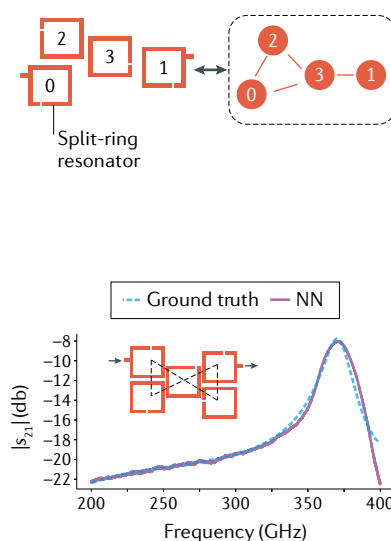
### b Scattering spectra



### c Electric polarization



### d Microwave filter



**Fig. 3 | Surrogate modelling with discriminative networks.** **a** | Initial research on the application of neural networks to electromagnetic devices started with microwave systems in the early 1990s and led to the modelling of lumped component devices, transmission lines and structured surfaces. Neural networks were then applied to model guided-wave photonic components and waveguides starting in the early 2010s. Recently, there has been immense interest to use neural networks to model various nanostructured optical media, from scatterers to metasurfaces. **b** | The scattering spectra of concentric nanoshell scatterers with shell thicknesses as inputs can be modelled using a deep, fully connected neural network (NN). **c** | The electric-polarization distribution within nanostructures under

electromagnetic excitation can be modelled using the convolutional neural network shown in the schematic. **d** | The relationship between coupled split-ring resonators constituting a microwave filter and the transmission spectrum of the filter can be modelled using a graph neural network. RF, radio frequency;  $s_{21}$ , the microwave analogue to transmission spectra. Panel **b** is reprinted with permission of AAAS from REF.<sup>96</sup>. © The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. Distributed under a Creative Commons Attribution NonCommercial License 4.0 (CC BY-NC). Panel **c** is adapted with permission from REF.<sup>54</sup>, ACS. Panel **d** is adapted with permission from REF.<sup>79</sup>, The Authors.

focused on devices with additional geometric degrees of freedom and included different classes of photonic crystal fibres<sup>86</sup>, 3D photonic crystals<sup>87</sup>, photonic crystal cavities<sup>88</sup>, plasmonic waveguide filters<sup>89</sup>, in-plane mode couplers and splitters<sup>90–92</sup>, Bragg gratings<sup>93</sup> and free-space grating couplers<sup>94,95</sup>.

Deep-learning modelling of free-space nanophotonic systems has only been researched in the last few years (FIG. 3a). The first example is from 2017 and is the modelling of the scattering spectral response of concentric metallic and dielectric nanoshells<sup>96</sup>. Discriminative networks have since modelled the spectral responses of other plasmonic systems, including chiral nanostructures<sup>97–99</sup>, planar scatters<sup>100</sup>, absorbers<sup>101</sup>, lattice structures with tailored coloration profiles<sup>102</sup>, smart windows based on phase-change materials<sup>103</sup> and nanoslit arrays supporting Fano resonances<sup>104</sup>. Discriminative networks have also been used to model photonic materials in the form of dielectric metagratings<sup>105</sup>, dielectric metasurfaces<sup>40,106</sup>, graphene-based metamaterials<sup>107,108</sup> and scatterers for colour design<sup>109</sup>, and they have been applied to thin-film dielectric stacks serving as colour filters<sup>110</sup> and topological insulators<sup>111–113</sup>.

The collection of surrogate models summarized above encompasses a wide range of deep-learning strategies that span differing network architectures, training strategies and modelling capabilities. To capture in more detail how deep discriminative networks are implemented in electromagnetic systems, we discuss a few representative examples.

**Scattering-spectra modelling.** A fully connected deep neural network was used to reconstruct the scattering properties of nanoparticles consisting of eight concentric dielectric shells of alternating silica and titania<sup>96</sup> (FIG. 3b). The input values to the network were the discrete nanoshell thicknesses, and the output was the scattering cross section between 400 and 800 nm, sampled over 200 points. The network itself contained four fully connected layers with 250 neurons each, and the spectra of 50,000 scatterers with random nanoshell thicknesses were generated for training using a transfer-matrix formalism. Scattering spectra generated by the trained neural network from a random geometry produced accurate profiles that demonstrated the ability of the network to perform high-level interpolation of the training data.

**Electric-polarization modelling.** A CNN was used to predict the vectorial polarization distribution in a nanophotonic structure, given a fixed electromagnetic excitation<sup>54</sup> (FIG. 3c). The input to the network was a 3D matrix that represented a nanophotonic structure discretized into small, subwavelength voxels. The output was matrices that represented the electric-polarization components at each voxel. A fully convolutional CNN was a sensible choice for this task because the electric-field distributions within a nanostructure are strongly spatially correlated with the detailed geometric features of the nanostructure. The network architecture utilized an ‘encoder–decoder’ scheme consisting of a series of convolutional and deconvolutional layers, which is a dimensionality-reduction scheme that enables high-level features of the input

matrix to be captured and used for data processing. The training data comprised approximately 30,000 random structures and their calculated field profiles. The trained network could predict the internal fields of a random structure with high accuracy, though approximately 5% of random-structure inputs produced predicted internal fields that strongly deviated from simulated values.

**Microwave-filter modelling.** A deep GNN was used to predict the  $s_{21}$  characteristics, which are the microwave analogue of transmission spectra, of a microwave circuit comprising three to six split-ring resonators<sup>79</sup>. The input to the network was a graphical representation of the circuit, in which the graph nodes contained information about individual ring-resonator geometries and the edges about the relative resonator positions (FIG. 3d). Each GNN layer contained two subnetworks, an edge processor that captured the near-field coupling between neighbouring resonators and a node processor that captured the electromagnetic properties of individual ring resonators based on their geometries and coupling with neighbouring resonators. The  $s_{21}$  spectra of 80,000 randomly generated circuits were generated as training data using a full-wave commercial simulator. The trained network was capable of accurately computing  $s_{21}$  for many resonator configurations four orders of magnitude faster than a commercial solver.

These representative examples capture a number of general trends common to deep discriminative networks. First, trained neural networks serve as reasonably accurate surrogate models for nanophotonic systems with limited complexity (that is, described by around ten physical parameters). An exception is the CNN mapping of nanostructure layout with the internal electric field, which is a special case due to the close correlation between nanostructure geometry and polarization profile. As the surrogate model represents a simplified approximation of the simulation space, there are always a fraction of cases for which model accuracy is poor. Second, discriminative-neural-network training is computationally expensive. Most of this expense arises from the generation of training data, which involves the simulation of tens of thousands of device examples or more using full-wave electromagnetic solvers. Third, a trained network can calculate an output response orders of magnitude faster than a full-wave solver. As such, training a discriminative neural network is advantageous only for applications in which the benefits of having a high-speed surrogate solver outweigh the substantial one-time computational cost of network training.

### Inverse design

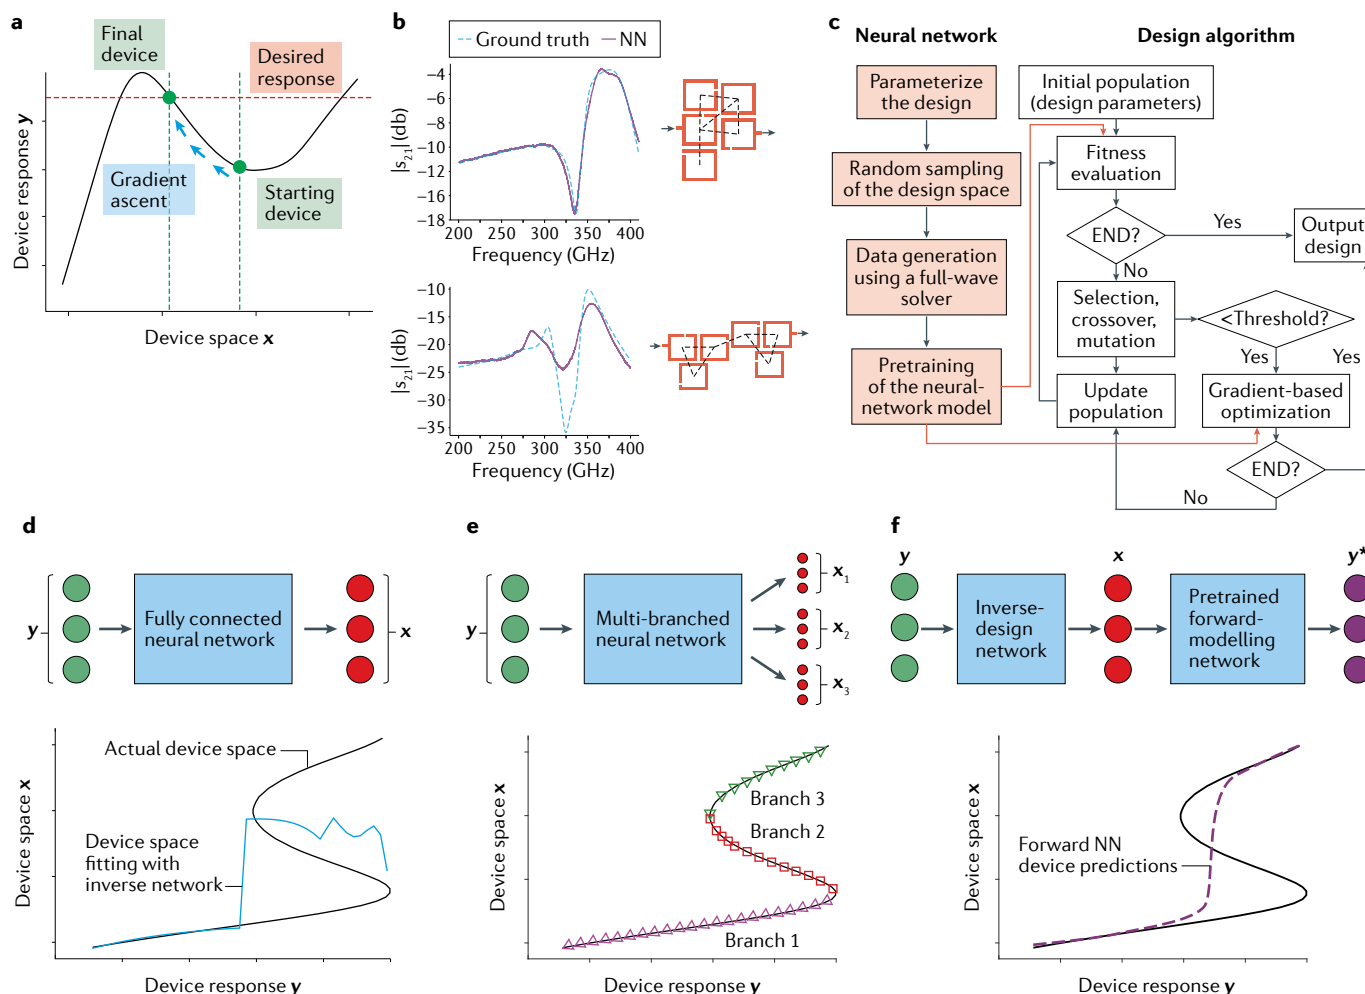
There are three general classes of inverse-design methods based on trained discriminative networks. The first class, outlined in FIG. 4a, comprises gradient-descent methods based on backpropagation. Initially, a device consisting of a random geometry or educated guess is evaluated by the trained network. The error between the outputted and desired response is then evaluated by the loss function and is iteratively reduced using backpropagation. Unlike in the network-training process, in which backpropagation reduces the loss by adjusting the



network weights, the loss here is reduced by fixing the network weights and adjusting the input-device geometry. We note that many optical-design problems do not possess unique solutions: multiple device layouts can exhibit the same desired optical response. As such, for neural networks that capture this non-uniqueness, different initial device layouts located within distinct domains of the design space will produce different final device layouts after optimization.

The backpropagation method was initially used in the 1990s for the inverse design of microwave circuits<sup>70,114</sup> and, more recently, to tailor the spectral properties of nanoparticle scatterers<sup>96</sup>, microwave filters<sup>79</sup> and photonic crystals<sup>88</sup>. As an inverse-design tool, this method can produce devices that exhibit the desired optical

response, as shown in the example of microwave filters (FIG. 4b, top). However, it does not always work well (FIG. 4b, bottom). One reason is that the neural network may not be accurately capturing the physical relationships between device geometry and response. This problem can be addressed by increasing the size of the training set. Another reason is that the device can get trapped in an undesired local optimum during the backpropagation process. This issue can be addressed by attempting optimizations with different initial device layouts, which leads to the exploration of different parts of the surrogate neural network design space. This issue can also be mitigated using alternatives to gradient descent, such as Adam optimization<sup>115</sup>, which uses momentum terms during backpropagation.



**Fig. 4 | Inverse design with discriminative networks.** **a** | Backpropagation can be used with a trained discriminative neural network to perform gradient-based inverse design. An initially random device-geometry input gets iteratively perturbed in a manner that pushes its optical response closer to that of a desired value. **b** | Examples of microwave filters designed using backpropagation with a trained graph-neural-network model. The resulting designs often match the desired spectral response (top) but not always (bottom). The corresponding device configurations are drawn on the right. **c** | Classical optimization algorithms can utilize neural networks as high-speed surrogate solvers to expedite the optimization process. In this algorithm, a surrogate solver is utilized in a genetic algorithm to coarsely search the design space, and the backpropagation method is then used to locally optimize the

device. **d** | Schematic of the inverse design space of a toy model, where the independent variable is the optical response and the dependent variable is the device layout. Due to the presence of multiple devices that exhibit the same optical response, the design space has multiple branches. Discriminative models that attempt to directly capture this design space will not get properly trained. **e** | A multi-branched neural network can be trained to fit a multi-branched inverse design space. **f** | Inverse networks can be implemented by first training a forward surrogate model, which learns to model a simplified version of the design space, and then to train an inverse network in tandem with the forward network. NN, neural network. Panel **b** is adapted with permission from REF.<sup>79</sup>, The Authors. Panel **c** is reprinted with permission from REF.<sup>78</sup>, IEEE. Panels **d–f** are adapted with permission from REF.<sup>118</sup>, IEEE.

The second class of inverse methods comprises hybrid optimization packages that utilize discriminative networks as solvers for conventional, iterative optimization algorithms. These algorithms span a wide range of well-established concepts in the optimization community and include Newton's methods<sup>93</sup>, interior-point algorithms<sup>105</sup>, evolutionary algorithms<sup>76,78,90,116</sup>, iterative multivariable approaches<sup>102</sup>, trust-region methods<sup>103</sup>, a fast-forward dictionary search<sup>40</sup> and particle swarm optimization<sup>74,77</sup>. Compared with optimization using conventional full-wave solvers, these hybrid packages are orders of magnitude faster, reducing the total optimization time from hours or days to minutes.

Unlike the backpropagation method discussed above, hybrid methods can support customized optimization strategies based on the needs of the designer. If a global search of the design space is required, global optimizers such as genetic algorithms can be used. Local optimizers such as Newton's methods are sufficient if good initial device layouts are known and the design space is sufficiently smooth. Optimizers can even be configured to combine both global and local searching of the design space. The flow chart of such a method for the optimization of microwave patch antennas and filters is shown in FIG. 4c (REF.<sup>78</sup>): a genetic algorithm coarsely searches for good regions of the design space, while gradient-based optimization locally refines the devices. The neural network accelerates the full optimization algorithm by serving both as a surrogate solver and as a gradient-based optimizer.

In the third class of methods, an inverse discriminative network can be configured such that its input is the desired optical response and its output the device geometry. While this might appear to be the most straightforward way to perform inverse design, it is difficult to execute in practice because of the non-uniqueness of optical design solutions for many problems. The issue is visualized in FIG. 4d for a toy model and shows regions of the inverse-design space in which individual inputs can take three possible outputs located along three distinct branches. During the training process, training data from a particular branch will attempt to push the surrogate model to that branch, and the net result is that the surrogate model will not converge to any branch and remain improperly trained.

One solution is to limit the design space through proper parameterization of the problem, to enforce a mostly unique mapping between device geometry and optical response. This concept was used to train an inverse network for plasmonic metasurfaces consisting of coupled metallic discs<sup>117</sup>. Another solution is to use a multi-branched neural network that can be configured to output multiple devices for a given input (FIG. 4e). In these networks, a special loss function ensures that specific network branches produce outputs that map onto unique design-space branches<sup>118</sup>. One other solution is to train a forward discriminative network serving as a surrogate model, fix its weights, and then use it to train an inverse network<sup>110</sup> (FIG. 4f). The forward surrogate model represents a simplified version of the full design space, reducing the quantity of non-unique solutions posed to the inverse network. Such tandem

networks have been effectively used for the inverse design of core-shell particles<sup>119</sup>, metasurface filters<sup>106</sup>, topological states in photonic crystals<sup>111-113</sup>, planar plasmonic scatterers<sup>100</sup> and dielectric nanostructure arrays supporting high-quality factors<sup>120</sup>.

### Dimensionality reduction

In order to train a discriminative network that adequately captures the correct mapping between the whole design and response space, a sufficient number of data points in the training set is required to adequately sample these spaces. For devices residing in a relatively low-dimensional design space, a brute-force sampling strategy is computationally costly but tractable: as we have seen, tens of thousands of training-set devices can be simulated for device geometries described by a handful of geometric parameters. However, the problem becomes increasingly intractable as the dimensionality of the design space increases, owing to the curse of dimensionality<sup>121,122</sup>, which states that the number of points in the training set required to properly sample the design space increases exponentially with the dimensionality of the design space. This scaling behaviour is plotted in FIG. 5a together with data points from ten different studies, which follow the exponential trend line. For freeform devices described as images with hundreds to thousands of voxels, a brute-force sampling strategy involving all of these voxels would require many billions of training-set devices. Even with distributed-computing resources, such an approach is not practical.

Fortunately, for many high-dimensional design spaces, candidate devices actually reside in or can be well approximated by a subspace described by a reduced number of feature parameters. For these problems, preprocessing these device representations from a high-dimensional to a low-dimensional space without critical information loss would make the discriminative modelling of these devices more tractable. In this section, we discuss three dimensionality-reduction techniques that have been applied to photonic systems.

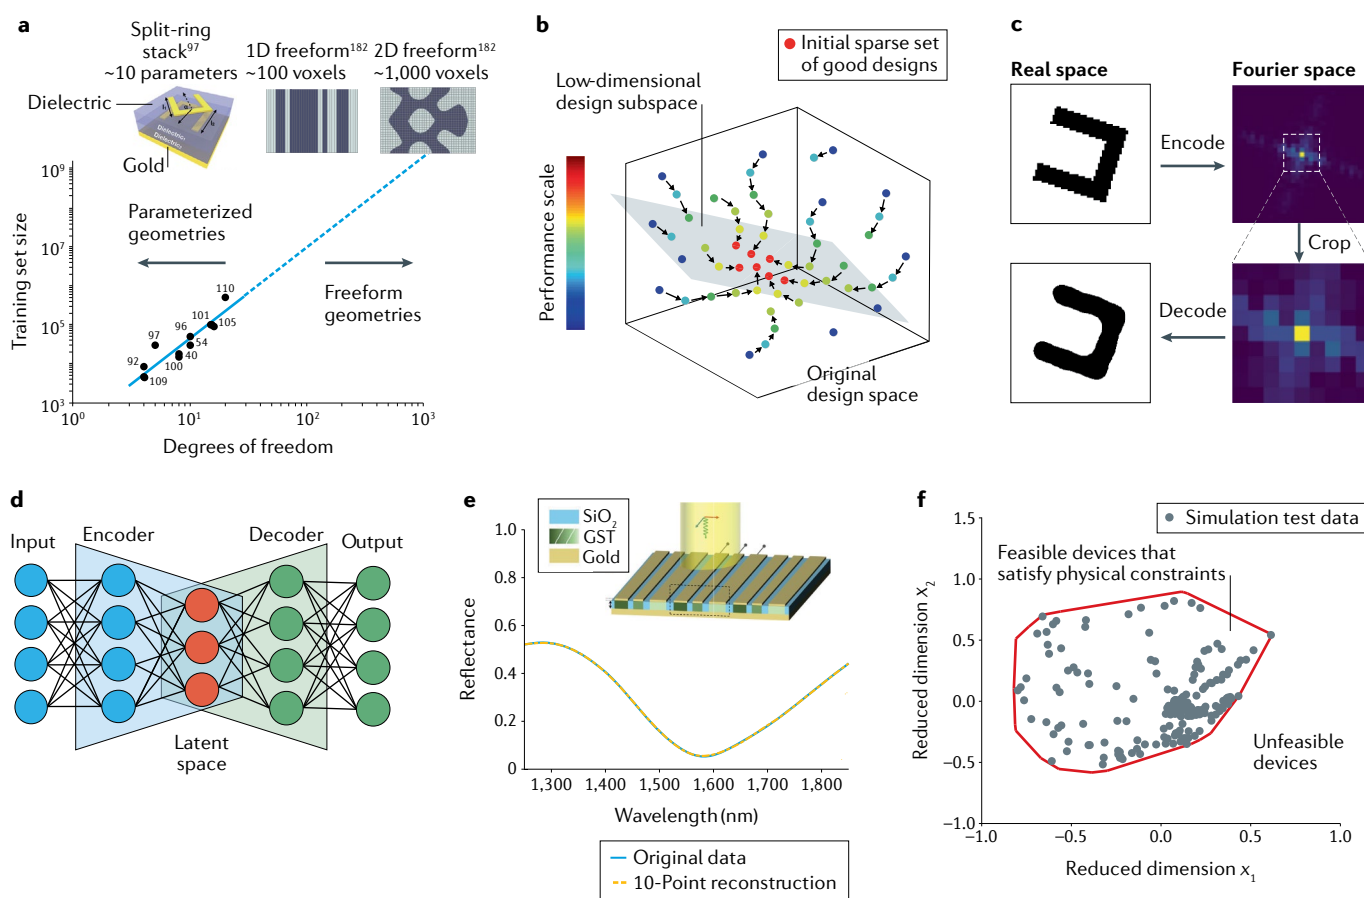
**Principal component analysis.** Principal component analysis (PCA) is a classical statistical technique in which a set of device parameters residing in a high-dimensional space is projected to a low-dimensional subspace. The first step to performing this projection is identifying a new orthogonal basis for the device parameters in the high-dimensional space. This basis is determined through the sequential specification of basis vectors with maximal component scores, which refer to the amount of information preserved when the device parameters are projected to the corresponding basis vector. The subset of these basis vectors with the largest component scores is then selected to define the low-dimensional subspace. High-dimensional device data structures projected to such a subspace retain the salient features of the original devices with minimal distortion.

PCA was used to expedite the inverse design of vertical fibre grating couplers<sup>123</sup> (FIG. 5b). Local optimization was first performed on a sparse distribution of devices within the full five-dimensional design space. PCA was then performed on this collection of locally optimized

devices, yielding a 2D subspace that captured locally optimal devices. Finally, a brute-force search within this low-dimensional subspace was performed with modest computational resources to identify even higher efficiency devices. This study did not utilize a deep neural network but the low-dimensional representation of devices based on PCA could be utilized in a discriminative network to perform device surrogate modelling in a particular design subspace.

**Fourier transformations.** Dimensionality reduction can be performed on device structures by eliminating high-spatial-frequency terms in image representations of the devices<sup>48</sup>. The concept is outlined in FIG. 5c. First, level-set functions of the device images are created, which ensures that the initial and dimensionality-reduced image-device representations are binary. Next, the

level-set functions, which represent the shape boundaries, are Fourier transformed, and high-frequency components in the Fourier space are cropped, reducing image dimensionality from  $64 \times 64$  to  $9 \times 9$ . Due to symmetry considerations, nine of the elements in this concatenated Fourier representation are unique and constitute the low-dimensional device space. To reconstruct the images in real space, the low-dimensional image representations in the Fourier domain are converted using an inverse Fourier transform and thresholding. As a demonstration, this dimensionality-reduced subspace was used to encode freeform grating elements that could diffract light to different channels. With this low-dimensional representation, a deep discriminative network that could learn the relationship between the grating subspace representation and diffractive response could be properly trained using only 12,000 devices.



**Fig. 5 | Methods for dimensionality reduction.** **a** | Plot of training-set size versus geometric degrees of freedom of the modelled device. Points from ten different studies<sup>40,54,92,96,97,100,101,105,109,110</sup> show exponential growth in training-set size as a function of degrees of freedom, which is consistent with trends from the curse of dimensionality. Insets: schematics of representative devices with different dimensionalities of the design space. **b** | Visualization of the method of principle components analysis, in which a hyperplane in the full design space of grating couplers defines a low-dimensional subspace of high-performance devices. **c** | Dimensionality reduction of images of devices is performed by encoding the image into a Fourier representation and then cropping that representation. **d** | Schematic of an autoencoder neural network. The encoder transforms input data to low-dimensional latent vectors, while the decoder attempts to reconstruct the input data from the latent vectors with minimal distortion.

**e** | An autoencoder reduces the dimensionality of the design and optical response space of a reconfigurable metasurface (shown in the inset; GST, germanium–antimony–tellurium), enabling an inverse discriminative network to be trained for the problem. The reflection spectrum of the metasurface is plotted as a function of the dimensionality of the reduced response space, showing that a dimensionality of ten or higher is sufficient for the reduced response space to delineate the original physical responses. **f** | A convex hull delineates suitable response-space regions within a data set, dimensionality reduced to two dimensions, representing the reflection spectra of differing dielectric meta-atoms. Panel **a** is adapted with permission from REF.<sup>97</sup>, ACS, and reprinted from REF.<sup>182</sup>, CC BY 4.0. Panel **b** is adapted from REF.<sup>123</sup>, CC BY 4.0. Panel **c** is adapted with permission from REF.<sup>48</sup>, © The Optical Society. Panel **e** is adapted from REF.<sup>126</sup>, CC BY 4.0. Panel **f** is adapted from REF.<sup>128</sup>, CC BY 4.0.

**Autoencoders.** Autoencoders, schematically shown in FIG. 5d, are neural networks that comprise two parts<sup>124,125</sup>. The first is an encoder network that maps input data from a high-dimensional design space to a low-dimensional latent vector. The second is a decoder network that takes the latent vector representation of the data and maps it back to the high-dimensional space, in an attempt to reconstruct the original data. Dimensionality reduction is, therefore, achieved through the encoding process. The loss function, also termed reconstruction loss, attempts to minimize the difference between the decoded and original data, and often has the form of least mean squares error. As such, the network attempts to learn the best encoding–decoding scheme to achieve a high level of dimensionality reduction, while maintaining a low level of information loss upon decoding. Compared with PCA, which is limited to linear transformations, autoencoders can learn more complex, low-dimensional representations of the data due to the highly nonlinear nature of neural networks.

Autoencoders were utilized to design reconfigurable metasurfaces<sup>126</sup> based on phase-change materials, which can perform amplitude modulation of a normally incident beam (FIG. 5e). Dimensionality reduction was performed to reduce the design space of device-geometry parameters from 10 to 5 dimensions and the response space of discretized spectra from 200 to 10 dimensions. The resulting subspace supported an approximate one-to-one mapping between the reduced design and response spaces, from which an inverse discriminative network was trained with only 4,000 devices. The trained inverse network, together with the encoder and decoder networks that mapped the data between high and low dimensions, could perform inverse design with high accuracy.

Autoencoders were also used in an inverse algorithm that could predict the layout of a digital metasurface, comprising a grid of metal or air pixels, given a desired spectral response<sup>127</sup>. To train this algorithm, an autoencoder was first used to perform a dimensionality reduction on spectra in the training set, from 1,000 to 128 data points. The resulting latent-space representation of the spectra and their corresponding device layouts were then used to train multiple support vector machines (SVMs), one for each device pixel. SVMs are classical machine-learning algorithms that serve as binary classifiers, and each SVM learned to classify latent-space inputs to be either metal or air for the corresponding pixel. The final encoder–SVM scheme produced metasurface layouts with spectral properties that closely matched the desired input spectra.

For autoencoders trained on device geometries, additional analysis of the low-dimensional latent-space representation of these devices can be performed in an attempt to further delineate the full design space of all feasible devices. These include fitting physical responses of the training-set devices to a convex geometric manifold, termed a convex hull, in the low-dimensional space (FIG. 5f), and using a SVM to classify the corresponding devices as either feasible or unfeasible. These concepts have been applied to analyse digital plasmonic nanostructures and dielectric nanopillar arrays<sup>128</sup>.

## Generative networks

### Generative networks for photonic systems

Deep generative networks have been an active topic of study in the computer-science community for the last decade and have produced impressive, eye-catching results. In one famous example, a generative network was trained to produce photorealistic images of faces, using a training set of millions of face images collected from the internet<sup>129</sup>. This example is representative of the conventional way generative networks are utilized in the computer-science community: networks are trained to generate classes of images that intrinsically exhibit a wide range of diversity, such as faces, animals and handwritten digits. Training methods exclusively rely on learning the statistical structure of the training set, as there are no analytic expressions to directly quantify the quality of generated images or to calculate gradients to directly improve the images. For example, there is no equation to score the quality of a face image or gradients to make an image more face-like<sup>130</sup>.

By contrast, in photonics inverse design, the goal is qualitatively different: it is to find just one or a handful of devices that achieve a specific design objective. In addition, electromagnetic simulators can assist in the network-training process by directly evaluating the electromagnetic fields, scattering profiles and performance gradients of generated devices. Performance gradients refer to structural perturbations that can be made to a device to improve its performance, and they can be calculated using the adjoint-variables method or auto-differentiation. With the adjoint-variables method, performance gradients that specify perturbations to the dielectric constant value at every device voxel, in a manner that improves a figure of merit, are calculated using forward and adjoint simulations<sup>18,19,131–134</sup>. Auto-differentiation is mathematically equivalent to backpropagation<sup>135–137</sup> and can directly evaluate performance gradients at every device voxel, pending the use of a differentiable electromagnetic simulator. Both methods can be performed iteratively to serve as a local freeform optimizer based on gradient descent.

These considerations have led to new implementations of generative networks for photonics inverse design; design schemes that utilize training sets for network training can be grouped into three strategies. The first is to train an unconditional network with a set of devices that samples a small, targeted subset of the design space (FIG. 6a). These training sets can be ensembles of disparate shapes that collectively exhibit a set of desired optical responses, in which case, a trained network would generate distributions of devices that more thoroughly fill out this design subspace. These training sets can also comprise variants of the same device type such that a trained network would generate even more geometric variations of that device type, some of which would perform better than those in the training set.

The second strategy is to train a conditional network with sets of high-performance devices (FIG. 6b). If the training set consists of devices operating with specific discrete values of the conditional labels, the trained network will be able to generalize and produce device layouts across the continuous spectrum of label values.

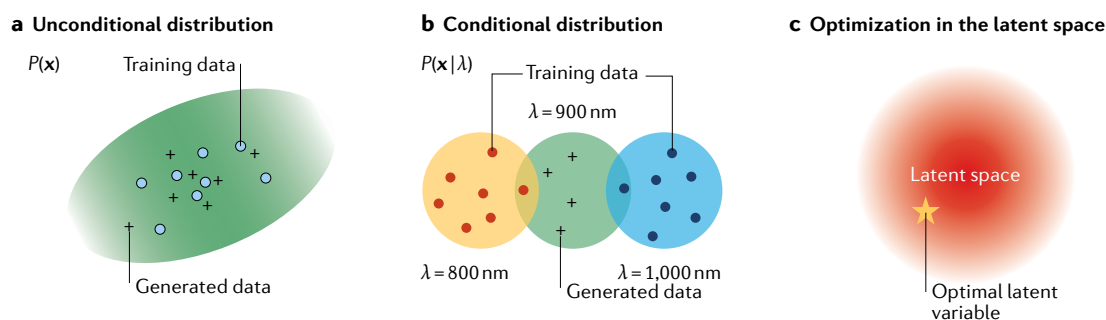


Fig. 6 | **Concepts for implementing generative neural networks for inverse design.** **a** | An unconditional network learns favourable regions of the design space from a limited training set of devices and can be sampled to find similar, better devices. **b** | A conditional network can generate distributions of devices for conditional parameter labels interpolated from the training set. In this schematic, a network learns the distributions of training-set devices operating at 800 and 1,000 nm, and can interpolate the distribution of devices operating at 900 nm. **c** | Classical optimization methods can be used to search the latent space within a trained network for optimal devices. The search space is limited to a device distribution specified by the training set.

This ability to generate devices with conditional label values interpolated from those in the training set is analogous to regression with discriminative networks.

The third strategy is to initially train either a conditional or unconditional generative network and then to use conventional optimization methods to search within the latent space for an input value that generates a structure with the desired optical properties (FIG. 6c). This method is related to that discussed previously in which a discriminative model is used as a surrogate electromagnetic solver in conjunction with conventional optimization methods. A principle difference here is that generative networks enable more control over the search space of candidate devices: the distribution of generated devices is constrained by the training set, and, furthermore, the generative network uses a proper, low-dimensional latent space to represent the training set.

### Generative model types

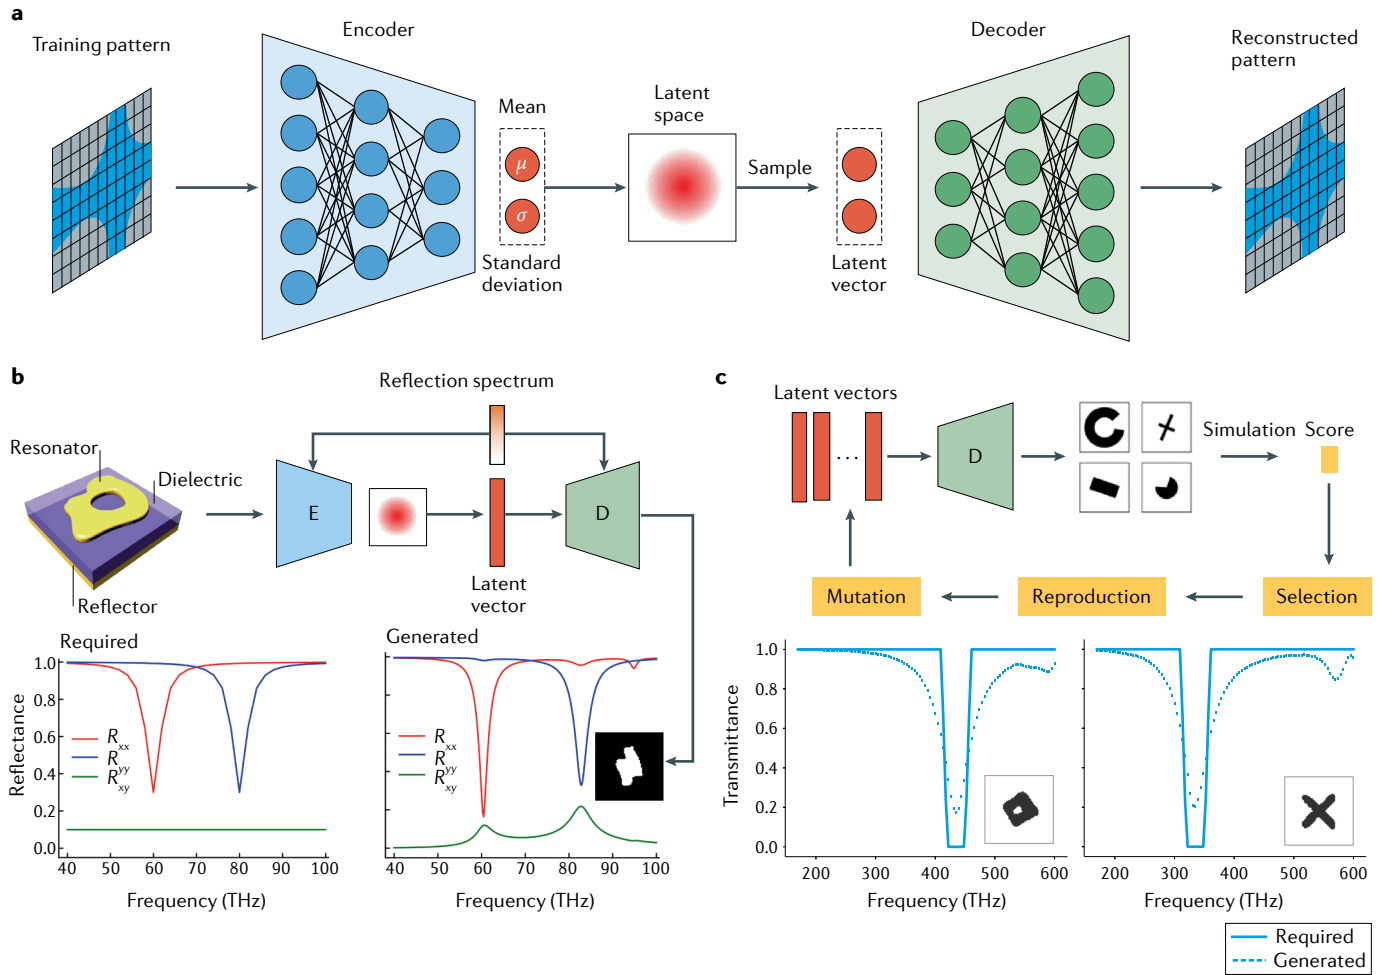
Deep generative networks are a relatively new innovation, and much of the foundation of the field has been set in the last decade. Unlike discriminative networks, which are described by relatively generic deep-network architectures, a range of generative models have been developed that assume different statistical properties about the training set and generated data distributions. Amongst the first deep generative models developed were autoregressive models<sup>138,139</sup>, which were applied to image generation in 2011 (REF.<sup>140</sup>). Images are generated pixel by pixel with the assumption that the  $i$ th generated pixel depends only on the value of all previous pixels<sup>49,141</sup>. These models do not use a latent variable but determine the value of each generated pixel by sampling an explicit conditional probability distribution. By explicit, we mean that these statistical distributions are described by analytic expressions.

Variational autoencoders (VAEs) were introduced in 2013 (REF.<sup>142</sup>) and are able to learn salient geometric features in a training set and statistically reconstruct variations of these features through sampling an explicit latent-variable distribution. Although these explicit statistical forms may not perfectly capture variations within the training set, they have explicit fitting parameters that

help simplify the training process. The most advanced and best performing generative network to date is the generative adversarial network (GAN), which was introduced in 2014 (REF.<sup>143</sup>) and learns the implicit statistical distributions of training sets. By implicit, we mean that the distributions have no predefined form. These models have the potential to better capture highly complex statistical trends within training sets but are less straightforward to train. GANs have very quickly evolved to support a high degree of sophistication, leading to demonstrations such as the photorealistic face generation described earlier. In this section, we discuss in more detail the architecture of VAEs and GANs, two of the more mainstream models used in photonics.

**Variational autoencoders.** To understand how a VAE works, we first revisit the autoencoder, which uses an encoder to map high-dimensional data to low-dimensional latent vectors and a decoder to map the latent vectors back to the high-dimensional space. The latent vectors capture principle features common to the training set. Given the objective to generate variants of structures similar to the training set, our goal is to frame the decoder as a generative neural network: by treating the latent vectors as latent variables and sampling within this latent space, the hope is that variations of different principle features get sampled and decoded into viable structures. Unfortunately, autoencoders do not properly interpolate the training set, and random sampling of the latent vectors followed by decoding produces devices with no relation to the training set.

VAEs are regularized versions of autoencoders that overcome these limitations by better managing the latent space<sup>142,144–146</sup> (FIG. 7a). We summarize two key features of VAEs with the caveat that there is a lot of probabilistic modelling required to fully understand VAEs. First, the encoder maps input data points not to discrete points in the latent space but to distributions. These latent-space distributions vary as a function of the input data and are typically set to be multivariate Gaussian distributions, each described by a mean and covariance matrix. The encoder, therefore, outputs mean and covariance matrix values that are a function of the



**Fig. 7 | Overview of variational autoencoders. a** | Schematic of a variational autoencoder (VAE). Like an autoencoder, a VAE possesses an encoder–decoder network structure and attempts to reconstruct input data with minimal distortion. However, instead of encoding input data to fixed latent vectors, the encoder maps input data to multivariate Gaussian distributions in the latent space, which is then sampled and decoded. **b** | A conditional VAE can train a decoder serving as an inverse network for meta-atoms. The inputs include a desired spectrum and latent variables, and the output is a distribution of devices. **c** | The decoder of a trained VAE can be used in the inverse design of meta-atoms by optimizing the latent space with evolutionary-optimization algorithms. D, decoder; E, encoder. Panel **b** is adapted with permission from REF.<sup>43</sup>, Wiley. Panel **c** is adapted with permission from REF.<sup>147</sup>, IEEE.

input data. The decoder now has the form of a generative network and uses these latent-space distributions as the latent-variable input to generate distributions of data. Second, the loss function includes both reconstruction loss, which is the same term used for autoencoders, and regularization loss, which is the Kullback–Leibler divergence between the multivariate Gaussian distribution returned by the encoder and a standard multivariate Gaussian distribution. The regularization loss helps to ensure that the latent space is not irregular and that samplings of the latent variable correspond to principle features learned from the training set.

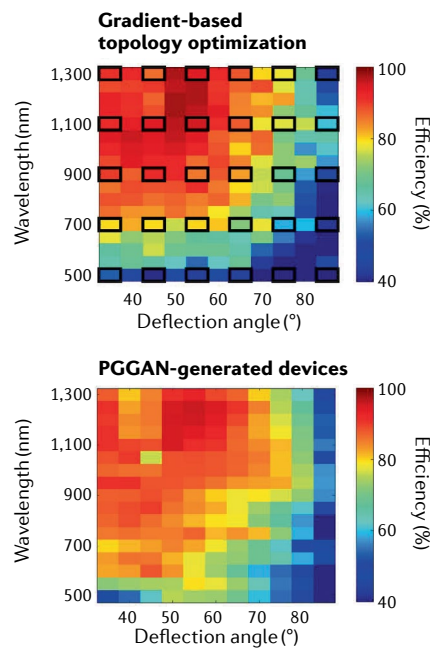
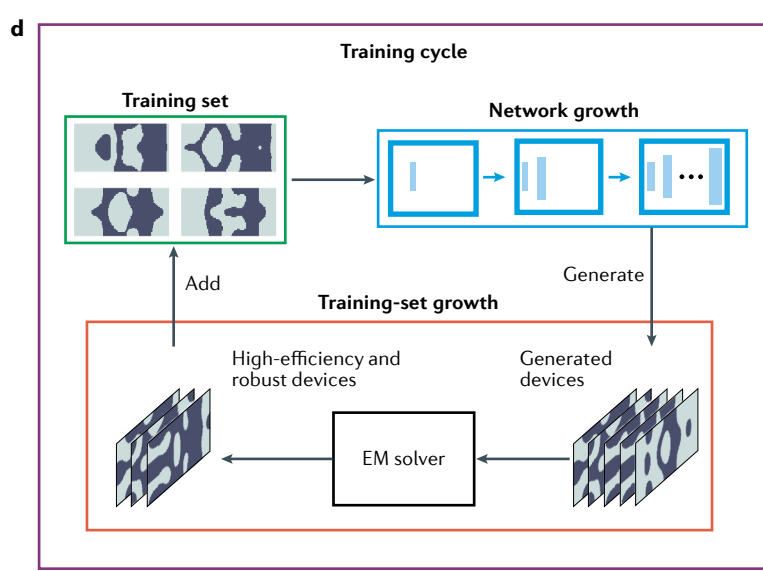
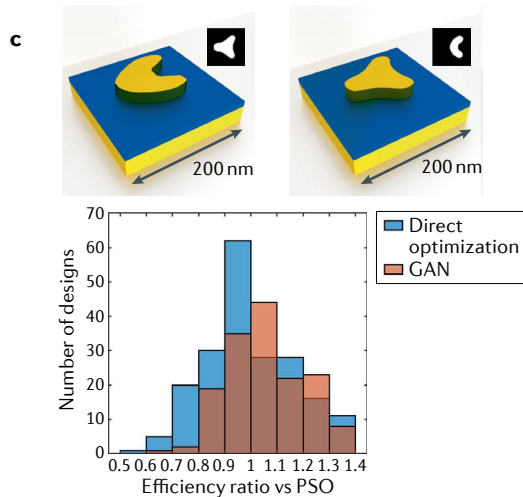
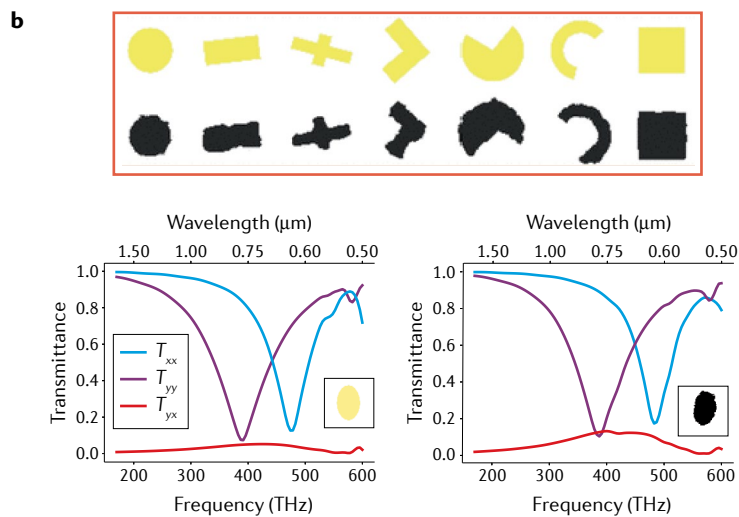
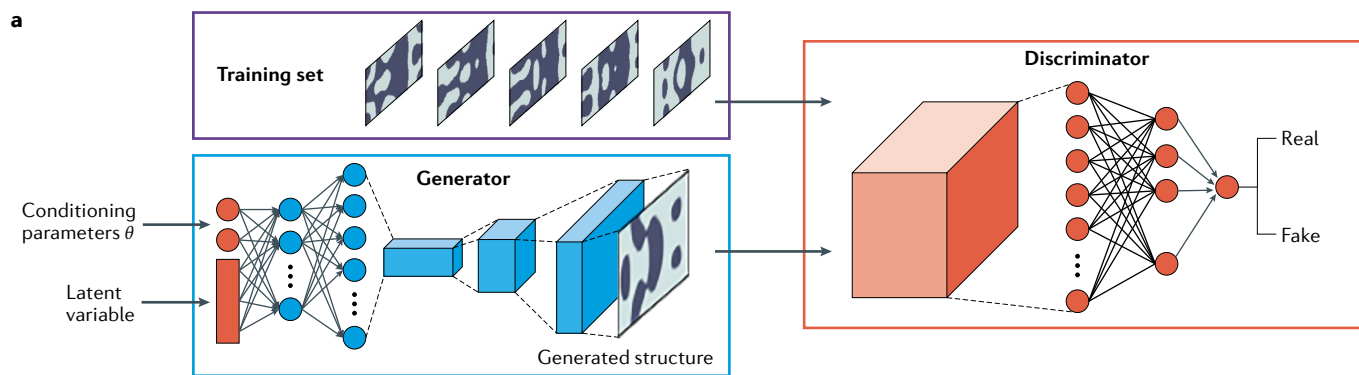
Two strategies for using VAEs in the inverse design of freeform, subwavelength-scale meta-atoms are summarized as follows. In the first strategy, the training set consists of images of meta-atom patterns and their spectra, which are both encoded into the low-dimensional latent space<sup>43</sup> (FIG. 7b). For the decoder, the inputs are the latent variable and the desired spectra, so that the

decoder is able to generate a distribution of device patterns given the spectra input. The trained decoder is able to decode the desired spectra, together with latent variables sampled from a standard Gaussian distribution, into device patterns. In the second strategy, a VAE is combined with evolutionary optimization<sup>147</sup> (FIG. 7c). The training set comprises meta-atom patterns of various shapes, such as circles, crosses and polygons, and a VAE is trained to map this subset of design patterns to the low-dimensional latent space. This latent space is then used as the basis for genetic algorithms, where a batch of latent vectors is decoded as device patterns, evaluated using electromagnetic simulations or a surrogate network, and evolved until an optimized latent vector corresponding to a suitable device pattern is identified.

**Generative adversarial networks.** GANs are a pair of neural networks that train together (FIG. 8a): a generative network that generates distributions of device images

from a latent-variable input and a discriminative network that serves as a classifier and attempts to determine if an image is from the training set or generator<sup>143,148–150</sup>. In the original GAN concept<sup>143</sup>, the training process is framed as a two-player game in which the generator attempts to fool the discriminator by generating structures that mimic the training set, while the discriminator

attempts to catch the generator by better differentiating real from fake structures. This training process is captured in the loss functions specified for each network. For the generator, the loss function minimizes Jensen–Shannon divergence between the training set and generated device distributions, in an attempt to get these device distributions to converge. For the discriminator,



the loss function maximizes Jensen–Shannon divergence, in an attempt to differentiate these two distributions as best as possible. Upon the completion of training, the fully trained discriminator will be unable to differentiate generated images from those in the training set, indicating that the generator is producing device distributions that approximate the training set. With this training method, the generator learns the implicit form of the training-set distribution without any explicit assumptions on its statistical properties. Since the inception of the GAN concept, alternative loss functions, including Wasserstein distance<sup>151</sup> and Wasserstein distance with gradient penalties<sup>149</sup>, have been implemented in so called WGANs. These loss functions help stabilize the network-training process and prevent the network from generating an overly narrow distribution of output values.

In an early example of GANs applied to photonics inverse design, a conditional generator was trained that could generate images of plasmonic nanostructures when the desired transmission spectrum was inputted<sup>144</sup> (FIG. 8b). The training set comprised a variety of freeform shapes, ranging from discs to crosses, and their corresponding transmission spectra. The training process involved the use of two discriminative networks to improve the generator. The first was the adversarial discriminator, which learned to differentiate generated from training-set patterns and ensured that the generated patterns mimicked the training set. The second was a pre-trained surrogate simulator, which evaluated the spectral response of generated nanostructures and ensured that the generated patterns exhibited the desired transmission spectrum. The trained generator produced distributions of shapes that mimicked those in the training set and served as an inverse network within this design space. A similar type of network scheme that utilized two discriminators was used to generate single-layer<sup>152</sup> and multi-layer<sup>153</sup> radio-frequency metasurfaces, where individual metasurface layers were generally represented as matrix elements stacked in a tensor. GANs were also used to generate plasmonic nanostructures, given the desired reflection spectrum as the conditional input<sup>45</sup>, and freeform dielectric meta-atom structures as

a function of amplitude and phase<sup>46</sup>. These GANs did not use a surrogate simulator during training. Instead, in REF.<sup>45</sup>, terms were added to the generator loss function to facilitate the matching of the generated shapes with the desired input spectrum, whereas in REF.<sup>46</sup>, a simulator was used to evaluate and filter for high-quality devices produced from the trained generator.

GANs have also been implemented using training data consisting of topology-optimized nanostructures. By restricting the design space of the training data to only high-performance freeform devices, the generative network exclusively focuses on learning the geometric features of freeform structures and does not expend resources exploring any other part of the design space. In one demonstration, an unconditional GAN was applied to plasmonic structures serving as thermal emitters for thermophotovoltaic systems<sup>154</sup>. The training set comprised images of different topology-optimized structures, each locally optimized from random initial dielectric distributions. Upon learning the statistical distribution of these training-set devices, the network could generate many topologically complex devices within this distribution, some of which exhibited better performance than the training-set devices (FIG. 8c). The same group also analysed thermal emitters using an adversarial autoencoder, which is a variation of the VAE model, except that an adversarial discriminative network is used instead of the Kullback–Leibler divergence to match the encoded latent space with a standard Gaussian distribution<sup>145</sup>. For their network-implementation schemes, adversarial autoencoders generated better devices than GANs<sup>155</sup>.

GANs have also been configured to learn features from topology-optimized metagratings, which are periodic metasurfaces that diffract incident light to the +1 diffraction order. Metagratings are good model systems for metasurfaces because they capture the essential light–matter interactions in diffraction processes<sup>156–161</sup>. In the first demonstration, a conventional GAN architecture, conditioned on operating wavelength and deflection angle, was trained with images of silicon-based metagratings operating for select wavelength–deflection angle pairs<sup>47</sup>. The final network could generate topologically complex devices for a continuous range of wavelength and deflection angle values, showing the ability of the network to learn and interpolate device layouts within this parameter space. However, the best generated devices were not robust or highly efficient, and they required additional optimization refinement to match the performance of training-set devices. In the following demonstration, the GAN training process was modified to incorporate progressive growth of the network architecture and training set over the course of multiple training cycles. Self-attention network layers that could capture long-range spatial correlations within images were also added to the network architecture<sup>162</sup> (FIG. 8d). When fully trained, this progressive-growing GAN could generate robust devices with efficiencies comparable to the best topology-optimized devices, showing the potential of generative networks to learn highly complex and intricate geometric trends in freeform photonic structures.

◀ Fig. 8 | **Generative adversarial networks for freeform device modelling.** **a** | Schematic of a generative adversarial network (GAN). The discriminator attempts to differentiate whether an inputted device is from the training set or generator, while the generative network attempts to fool the discriminator by generating devices mimicking the training set. The trained generator generates device distributions that match the training-set distribution. **b** | Images of generated devices and their spectral properties from a GAN conditioned on spectral response. The generator trains with two discriminators, one that ensures that the generated devices (black) match those in the training set (yellow) and one that ensures that the generated devices have spectra matching the desired conditional input. **c** | Device layouts and performance histogram of thermal-emission structures generated by an unconditional GAN. **d** | Progressive-growing GAN (PGGAN) that trains using progressive growth of the network architecture and training set over multiple cycles. The trained network can generate robust metagratings for a range of wavelengths and deflections, with efficiencies comparable to or exceeding those of the best devices designed using adjoint-based topology optimization. The black squares on the top-right plot indicate wavelength–deflection angle pairs for devices in the original training set. EM, electromagnetic; PSO, particle swarm optimization. Panel **b** is adapted with permission from REF.<sup>44</sup>, ACS. Panel **c** is adapted with permission from REF.<sup>186</sup>, Applied Computational Electromagnetics Society. Panel **d** is adapted from REF.<sup>162</sup>.



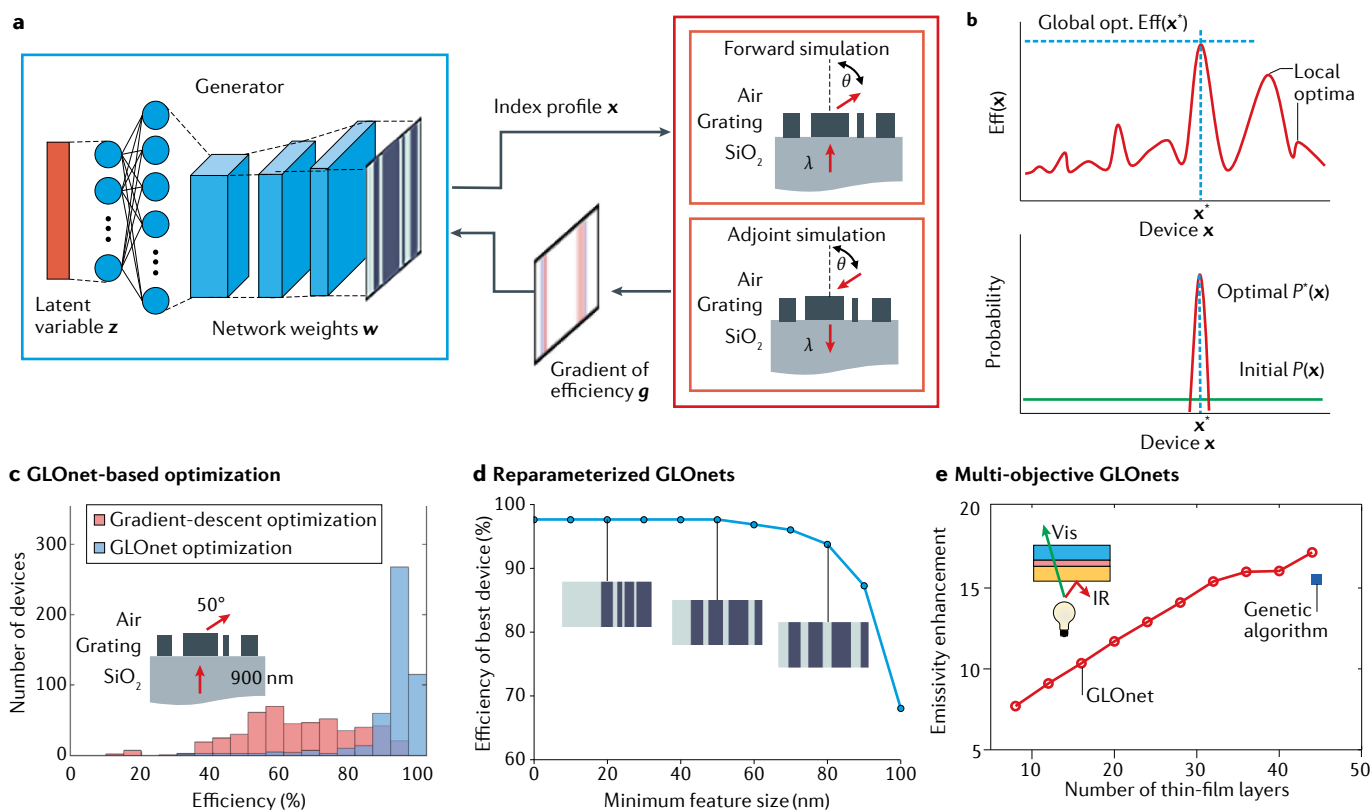
### Global topology-optimization networks

A long-standing challenge in inverse photonics design is the global optimization of freeform devices. Existing methods to perform inverse freeform design, ranging from heuristic to gradient-based topology optimization, are not able to effectively solve for the global optimum because the design space for photonic devices is vast and non-convex. In all neural-network-based inverse-design methods discussed thus far, which rely on a training set, global optimization is only possible if devices near or at the global optimum are included in the training set. Discriminative and generative neural networks can be effective at fitting training data, but they cannot perform meaningful extrapolation tasks beyond the training set.

Global topology-optimization networks (GLOnets), outlined in FIG. 9a, are a newly developed class of generative network that are capable of effectively searching the design space for the global optimum<sup>41,42</sup>. Unlike conventional implementations of generative networks, which

are trained to fit a training-set distribution, GLOnets attempt to fit a narrowly peaked function centred around the global optimum and do so without a training set (FIG. 9b). In this manner, GLOnets reframe the topology-optimization process through the dataless training of a neural network.

The basic GLOnet architecture shown in FIG. 9a is a deep generative CNN with conditional device labels and a latent variable as inputs, and it outputs a distribution of devices. During each training step, a batch of devices is generated, and the performance metric and performance gradient of each device are evaluated using a Maxwell solver. The latter can be calculated using either the adjoint-variables method or auto-differentiation. These performance metric and gradient terms are then incorporated into the loss function and backpropagated to adjust the network weights. The loss function is engineered to push the distribution of generated devices towards the global optimum and is:



**Fig. 9 | Global topology-optimization networks.** **a** | Schematic of a global topology-optimization network (GLOnet) optimizer. A generative network is trained to output a narrow distribution of devices centred around the global optimum. The training process involves evaluating generated devices with an electromagnetic solver and using those results in backpropagation to improve the mapping of the latent variable to device distribution. **b** | Schematic of the design space, where devices  $x$  have efficiencies  $\text{Eff}(x)$ . The design space has many local optima and a single globally optimal device,  $x^*$ , which has an efficiency of  $\text{Eff}(x^*)$ . Prior to training, the generative neural network is randomly initialized with weights  $w$  and generates a uniform distribution of devices spanning the entire design space,  $P_w(x)$ . Upon training completion, the final network has optimal weights  $w^*$  and produces the narrow distribution  $P_{w^*}(x)$ . **c** | Histogram of silicon metagrating efficiencies for devices designed using

the adjoint-variables method and GLOnet. The adjoint-variables-optimized devices have a broad distribution and the best device operates with 93% efficiency. The GLOnet devices have a relatively narrow distribution with high efficiencies, and the best device operates with 98% efficiency. **d** | Plot of metagrating deflection efficiency as a function of minimum feature size for globally optimized devices comprising four silicon bars, solved using a reparameterized GLOnet. **e** | Plot of emissivity enhancement as a function of the number of layers for thin-film stacks serving as an incandescent-light-bulb filter. Multi-objective GLOnets outperform the genetic-algorithm reference and can produce high-performing devices with relatively few layers. Inset: schematic of filters, based on stacks of thin films, that transmit visible light (Vis) and reflect infrared (IR) light. Panels **a** and **c** are adapted from REF.<sup>42</sup>, CC BY 4.0. Panel **d** is adapted from REF.<sup>187</sup>, © 2020 The Author(s).

$$L(\mathbf{x}, \mathbf{g}, \text{Met}) = -\frac{1}{N} \sum_{n=1}^N \frac{1}{\sigma} \exp\left(\frac{\text{Met}^{(n)}}{\sigma}\right) \mathbf{x}^{(n)} \cdot \mathbf{g}^{(n)} \quad (1)$$

where  $N$  is the batch size,  $\sigma$  a tunable hyperparameter and  $\text{Met}^{(n)}$ ,  $\mathbf{x}^{(n)}$  and  $\mathbf{g}^{(n)}$  the performance metric, device layout and performance gradient of the  $n$ th device, respectively. The biasing of the network towards the global optimum is captured by the exponential weighing of the performance metric in the loss function. Interestingly, the value of the globally optimal performance metric does not need to be known.

In an initial demonstration, GLOnets were used to globally optimize the efficiencies of metagratings consisting of silicon ridges<sup>42</sup>. Sixty-three unconditional GLOnets were trained, each searching for optimal devices with distinct combinations of operating wavelength and deflection angle, and each network was benchmarked with 500 locally optimized devices designed using the adjoint-variables method. For 57 of these networks, the best GLOnets device had the same or higher efficiency compared with the best locally optimized device. Histograms of device efficiencies from these two methods show that the distributions of generated GLOnet devices are relatively narrow and biased towards high efficiencies, which is consistent with the training goal of GLOnets (FIG. 9c). The stability of the GLOnets method was demonstrated with the training of eight different randomly initialized networks, each with the same design objective: six of the eight trained networks produced the same device possessing an efficiency of 97%.

Conditional GLOnets for metagratings, which can simultaneously optimize devices with a range of wavelengths and deflection angles, were also examined<sup>41</sup>. For this demonstration, network training with conditional labels worked well because the design space and optimal devices for different conditional labels were strongly correlated. A comparison of the best devices generated from a single GLOnet and the best devices locally optimized using the adjoint-variables method showed that 75% of the devices from the conditional GLOnet had higher efficiencies than those based on the local optimization. The computational resources required to train the GLOnet were 10 times less than those used for the local optimization of the benchmark devices. Such computational efficiency arises because the GLOnet does not expend computational resources in unpromising parts of the design space. Instead, it constantly shifts the generated-device distribution towards the globally optimal device during training.

#### **Incorporation of constraints with reparameterization**

An important consideration with all inverse-design methods is the incorporation of practical experimental constraints, such as the specification of a minimum feature size or robustness to fabrication imperfections. A typical method to incorporate these constraints in inverse design is to add terms in the figure of merit that penalize violations to these constraints<sup>163–165</sup>. Although this strategy generally pushes devices towards regions of

the design space that satisfy the desired constraints, it does not guarantee the enforcement of constraints. An alternative method that can impose hard constraints in optimization is to reparameterize the problem, in a manner where the optimizer processes devices in a latent space that can span an unconstrained range of values<sup>166</sup>. Mathematical transformations are then used to transform the latent-space representation to the real-device space, and constraints imposed within the real-device space are defined by the transformation itself. Finally, the constrained device is evaluated in the real-device space. For gradient-based optimizers, such as the adjoint-variables method or GLOnets, performance gradients are calculated for devices in the real-device space and are back-propagated to the latent-device-space representation for the optimizer to process. Backpropagation is generally possible, as long as the mathematical transformations linking the two spaces are differentiable.

Reparameterization was applied to GLOnet algorithms for silicon metagratings deflecting incident monochromatic light to a 65° angle. In this example, the topology was fixed to contain four silicon nanoridges, and the unconstrained latent-space variables mathematically transformed to ridge-width and ridge-separation values with a hard minimum-feature-size constraint. The silicon and air regions of the device were still defined to possess greyscale values with spatial profiles defined by analytic functions, which allowed gradients from the adjoint-variables method to be directly applied to this shape-optimization problem. The resulting globally optimized device efficiencies as a function of minimum feature size are summarized in FIG. 9d. The unconstrained globally optimal device has a minimum feature size of 20 nm, such that reparameterized GLOnets with minimum feature sizes equal to or smaller than 20 nm generated the same optimal device. As the minimum-feature-size constraint increased, the efficiency of the globally optimized devices decreased. Images of the device layouts show that the globally optimal devices each possess at least one feature with the minimum feature size posed by the constraint, indicating the utility of small features to enhance light-diffraction efficiency in these devices.

#### **Multi-objective GLOnets**

The metagrating topology-optimization problem above is single objective: all parameters in the problem are fixed except the refractive indices of each voxel, which are specified to be either silicon or air. Multi-objective problems are more complex and require more than just binary decisions to be made, but they can be readily handled with the GLOnets formalism without loss of generality. As an example, consider the design of multi-layer stacks in which each layer can be one of  $M$  distinct material types. With GLOnets, these multi-layer stacks are represented as matrices, where each row is a  $1 \times M$ -dimensional vector, and each term corresponds to a particular material type in a given layer. Each of these vectors is computed into a probability distribution using the softmax function, which specifies the likelihood that a particular material in a given layer is optimal. The expected refractive index of each layer given by this likelihood matrix is

calculated, evaluated with an electromagnetic solver and used to evaluate the loss function and perform backpropagation. As the training process progresses, each row of the likelihood matrix converges to have one predominant term, which is the optimal material.

Multi-objective GLOnets have been applied to a number of thin-film-stack systems<sup>167</sup>. One is anti-reflection coatings intended for broadband and broad-angle usage on silicon solar cells, where a continuum of dielectric values was selected for each layer. Existing benchmarks for a three-layer system included a brute-force search of the global optimum, which took over 19 days of CPU computation<sup>168</sup>, and a multi-start gradient optimizer<sup>168</sup>, which took 15 minutes to find the global optimum. GLOnets solved for the global optimum in seven seconds with a single GPU, demonstrating their efficiency and efficacy. GLOnets were also applied to thermal filters that could transmit visible light and reflect infrared light. The results are summarized in FIG. 9e for thin-film stacks comprising seven different dielectric material types: magnesium fluoride, silicon dioxide, silicon carbide, silicon mononitride, aluminium oxide, hafnium dioxide and titanium dioxide. The broadband reflection characteristics of a 45-layer GLOnet-optimized device showed that the device operates with nearly ideal transmission at 500–700 nm and nearly ideal reflection at near-infrared wavelengths, both for normal incidence and for incidence angles averaged over all possible solid angles. The application of GLOnets to different layer numbers and a comparison with a genetic-algorithm benchmark<sup>16</sup> showed that, for optimized devices with 45 layers, GLOnets clearly outperformed the genetic-algorithm reference. Furthermore, GLOnets could produce devices with the same performance as the genetic-algorithm reference but with approximately two-thirds the number of layers, which is important for translating these designs to experiment.

### Future research directions and practices

Deep neural networks are poised to be a disruptive force in the solving of forward and inverse-design problems in photonics. In just the last few years, discriminative networks have been shown to serve as effective surrogate models of Maxwell solvers, learning and generalizing the complex relationship between nanoscale layouts and their optical properties. Generative models have proven to serve as a new framework for the inverse design of freeform devices, through the learning of geometric features within device data sets and by dataless network training using Maxwell solvers.

Neural-network-based models are not a general replacement tool for conventional electromagnetic simulators, which will continue to be a workhorse tool for most problems, but they have complementary strengths and weaknesses. The main drawback of neural networks is that they require large training sets of thousands to millions of devices, which is a significant one-time computational cost. If conventional simulation and optimization methods can solve a problem with an equivalent or smaller computational budget, it is more judicious and straightforward to stick with conventional approaches. Another issue is that even the best trained networks cannot guarantee accuracy and should not be used in lieu of

an electromagnetic simulator when an exact physics calculation is required. It is also noted that low-dimensional electromagnetic systems described by a small number of design parameters can often be modelled and optimized using a number of classical statistical, machine learning and optimization packages<sup>169–171</sup>, many of which are available as standard numerical toolboxes in scientific computing software. Compared with the training of deep networks, these methods can work as effectively and do not require extensive hyperparameter tuning.

Neural-network-based models also have a number of strengths that make them uniquely suited for a number of problems. First, a trained neural network operates with orders-of-magnitude faster speeds than a conventional simulator and is ideal in situations where simulation time is a critical factor. Second, the regression capabilities of neural networks surpass those of classical data-fitting methods and can extend to complex, high-dimensional systems, owing to the scalability of neural networks to accommodate thousands of neurons with tunable parameters. Third, neural networks are particularly computationally efficient at simulating and designing many device variants that utilize related underlying physics. These devices range from grating couplers that require different input-mode conditions to metasurface sections that require different amplitude and phase properties, and these device variants can be readily co-designed by training a single conditional neural network. Fourth, neural-network approaches to inverse design can produce electromagnetic devices with better overall performance. Global topology optimization based on GLOnets has already been shown to supersede conventional gradient-based optimizers, and continued advancements in neural-network-based optimization promise even better and more computationally efficient design algorithms.

Looking ahead, multiple innovations will be required to push the capabilities of deep-learning algorithms towards the inverse design of complex, technologically relevant devices. First, although generic machine-learning algorithms will continue to play a role in solving photonics problems, new concepts that intimately combine the underlying physical structure of Maxwell's equations with machine learning need to be developed. GLOnets, which combine machine learning with physics-based solvers, is one such example showing how new hybrid algorithms can enhance the capabilities of neural networks. There have also been recent demonstrations that neural networks can be trained to solve differential equations<sup>172</sup>. To integrate physics with neural networks, we anticipate new innovations in network architectures, training procedures and loss-function engineering, as well as entirely new ways of using discriminative and generative networks both independently and together. We predict that dataless training, in which physics-based calculations are used to train neural networks, will serve as a particularly effective and computationally efficient means to harnessing machine learning for photonics problems.

Second, new electromagnetic simulators need to be developed that can operate at significantly faster timescales than conventional full-wave solvers. Fast solvers are

needed because, as device complexity increases, significantly larger training sets for supervised learning and larger simulation batches for dataless training methods are required. We anticipate that application-specific electromagnetic solvers will play a major role as ultra-fast solvers in deep-learning-photonics problems. One path forward is the augmentation of existing Maxwell solvers with neural-network-enhanced preconditioners<sup>173</sup>. With a neural network that can predict an approximate solution to the electromagnetic problem on hand, that solution can be used as a starting point for the solver and dramatically speed up the calculation. Specialized algorithms that can evaluate the scattering properties of structures with high computational efficiency, such as integral-equation solvers<sup>174</sup> and T-matrix approaches<sup>175</sup>, are also worth revisiting.

Third, the training and refinement of neural networks for solving photonics problems need to be better streamlined, from both data-usage and user-interface points of view. Currently, every time a new problem is proposed, a data scientist needs to train and fine-tune a neural network from scratch. One avenue that can help address data usage is transfer learning, in which a subset of network weights from a trained network solving an initial problem is applied to a network intending to solve a related problem<sup>176</sup>. The initial problem can be that of a related physical system for which a trained network already exists or it can be a simplified version of the desired problem, from which an initial neural network can be trained with computationally ‘cheap’ data. In a recent demonstration, network weights from a trained network that could predict the scattering spectra of concentric shell scatterers were transferred to a network intended to predict the spectral properties of dielectric stacks, leading to improved training accuracy of the latter<sup>177</sup>. For the user-interface problem, we anticipate that meta-learning, in which neural networks learn to learn<sup>178,179</sup>, will help automate the set-up and training process for photonics-based machine-learning algorithms. Meta-learning is currently an active area of research in the computer-science community, and, although it is a highly data-intensive proposition,

it promises to simplify the interface between algorithms and users.

Also looking ahead, it would be judicious for our community to take inspiration from the computer-science community and engage in a more open culture of sharing. In the computer-science community, extreme progress and proliferation of data science can be attributed, in part, to the willingness of computer scientists to openly share algorithms and benchmark their approaches by solving common problems. For example, it is typical in computer-vision research for groups to use established databases of labelled images, such as ImageNet<sup>180</sup> and CIFAR-10 (REF.<sup>181</sup>), as training data for algorithm benchmarking. The computer-vision community even engages in regular contests, such as the ImageNet Large Scale Visual Recognition Challenge, in which researchers attempt to solve the same image-classification task with the same training data. This community-centric design-of-experiments approach allows researchers to rapidly prototype, compare and evolve their algorithms at a rapid rate, to the benefit of the whole community.

In this spirit, an online repository for device designs and inverse-design codes for nanophotonic systems, termed MetaNet, has been developed<sup>182</sup>. As of this paper’s publication, MetaNet contains design files of over 100,000 freeform metagrating structures, as well as codes for local and global topology optimization. We hope that, with continued dialogue within the photonics community, we can agree on important design problems to tackle and to open-source training sets and basic code formulations so that we may build on each other’s algorithmic approaches. At the very least, with inverse-design strategies that produce freeform-device layouts, we need standardized methods to share device layouts so that we can benchmark and openly evaluate the capabilities of these structures, not just in terms of device performance but also other metrics, such as robustness to geometric imperfections. By working together, we can effectively push optical and photonics engineering to the next and possibly final frontier.

Published online: 17 December 2020

- Sun, J., Timurdogan, E., Yaacobi, A., Hosseini, E. S. & Watts, M. R. Large-scale nanophotonic phased array. *Nature* **493**, 195–199 (2013).
- Stern, B., Ji, X., Okawachi, Y., Gaeta, A. L. & Lipson, M. Battery-operated integrated frequency comb generator. *Nature* **562**, 401–405 (2018).
- Thomson, D. et al. Roadmap on silicon photonics. *J. Opt.* **18**, 073003 (2016).
- Cheng, Q., Bahadori, M., Glick, M., Rumley, S. & Bergman, K. Recent advances in optical technologies for data centers: a review. *Optica* **5**, 1354–1370 (2018).
- O’Brien, J. L., Furusawa, A. & Vuckovic, J. Photonic quantum technologies. *Nat. Photonics* **3**, 687–695 (2009).
- Titti, A. et al. Imaging-based molecular barcoding with pixelated dielectric metasurfaces. *Science* **360**, 1105–1109 (2018).
- Baumberg, J. J., Aizpurua, J., Mikkelsen, M. H. & Smith, D. R. Extreme nanophotonics from ultrathin metallic gaps. *Nat. Mater.* **18**, 668–678 (2019).
- Urbas, A. M. et al. Roadmap on optical metamaterials. *J. Opt.* **18**, 093005 (2016).
- Yu, N. & Capasso, F. Flat optics with designer metasurfaces. *Nat. Mater.* **13**, 139–150 (2014).
- Inan, U. S. & Marshall, R. A. *Numerical Electromagnetics* (Cambridge Univ. Press, 2014).
- Jin, J.-M. *The Finite Element Method in Electromagnetics* (Wiley, 2014).
- Campbell, S. D. et al. Review of numerical optimization techniques for meta-device design [invited]. *Opt. Mater. Express* **9**, 1842–1863 (2019).
- Kim, W. J. & O’Brien, J. D. Optimization of a two-dimensional photonic-crystal waveguide branch by simulated annealing and the finite-element method. *J. Opt. Soc. Am. B* **21**, 289–295 (2004).
- Zhao, Y. et al. Broadband diffusion metasurface based on a single anisotropic element and optimized by the simulated annealing algorithm. *Sci. Rep.* **6**, 23896 (2016).
- Preble, S., Lipson, M. & Lipson, H. Two-dimensional photonic crystals designed by evolutionary algorithms. *Appl. Phys. Lett.* **86**, 061111 (2005).
- Shi, Y., Li, W., Raman, A. & Fan, S. Optimization of multilayer optical films with a memetic algorithm and mixed integer programming. *ACS Photonics* **5**, 684–691 (2018).
- Piggott, A. Y. et al. Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer. *Nat. Photonics* **9**, 374–377 (2015).
- Lalau-Keraly, C. M., Bhargava, S., Miller, O. D. & Yablonovitch, E. Adjoint shape optimization applied to electromagnetic design. *Opt. Express* **21**, 21693–21701 (2013).
- Jensen, J. S. & Sigmund, O. Topology optimization for nano-photonics. *Laser Photonics Rev.* **5**, 308–321 (2011).
- Abbeel, P., Coates, A. & Ng, A. Y. Autonomous helicopter aerobatics through apprenticeship learning. *Int. J. Robot. Res.* **29**, 1608–1639 (2010).
- Lavecchia, A. Machine-learning approaches in drug discovery: methods and applications. *Drug. Discov. Today* **20**, 318–331 (2015).
- Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.* **25**, 1097–1105 (2012).
- Wu, Y. et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. Preprint at <https://arxiv.org/abs/1609.08144> (2016).
- Abadi, M. et al. TensorFlow: Large-scale machine learning on heterogeneous systems. (2015). Software available from tensorflow.org.
- Paszke, A. et al. in *31st Conference on Neural Information Processing Systems (NIPS 2017)* (2017).
- cmphys: A Python interface to the COMSOL Multiphysics API. <https://github.com/fellobos/cmphys> (2018).
- Python API overview (Lumerical Support, Vancouver, BC, Canada).

28. MATLAB, version 7.10.0, R2020a (MathWorks Inc.)
29. Zhang, Q., Cheng, L. & Boutaba, R. Cloud computing: state-of-the-art and research challenges. *J. Internet Serv. Appl.* **1**, 7–18 (2010).
30. Nickolls, J. & Dally, W. J. The GPU computing era. *IEEE Micro* **30**, 56–69 (2010).
31. Jouppi, N. P. et al. In-datacenter performance analysis of a tensor processing unit. *SIGARCH Comput. Archit. News* **45**, 1–12 (2017).
32. Yao, K., Unni, R. & Zheng, Y. Intelligent nanophotonics: merging photonics and artificial intelligence at the nanoscale. *Nanophotonics* **8**, 339–366 (2019).
33. So, S., Badloe, T., Noh, J., Bravo-Abad, J. & Rho, J. Deep learning enabled inverse design in nanophotonics. *Nanophotonics* **9**, 1041–1057 (2020).
34. Hegde, R. S. Deep learning: a new tool for photonic nanostructure design. *Nanoscale Adv.* **2**, 1007–1023 (2020).
35. Tunnermann, H. & Shirakawa, A. Deep reinforcement learning for coherent beam combining applications. *Opt. Express* **27**, 24223–24230 (2019).
36. Sajedian, I., Badloe, T. & Rho, J. Optimisation of colour generation from dielectric nanostructures using reinforcement learning. *Opt. Express* **27**, 5874–5883 (2019).
37. Iman Sajedian, H. L. J. R. Double-deep Q-learning to increase the efficiency of metasurface holograms. *Sci. Rep.* **9**, 10899 (2019).
38. Shen, Y. et al. Deep learning with coherent nanophotonic circuits. *Nat. Photonics* **11**, 441–446 (2017).
39. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
40. Naddell, C. C., Huang, B., Malof, J. M. & Padilla, W. J. Deep learning for accelerated all-dielectric metasurface design. *Opt. Express* **27**, 27523–27535 (2019).
41. Jiang, J. & Fan, J. A. Global optimization of dielectric metasurfaces using a physics-driven neural network. *Nano Lett.* **19**, 5366–5372 (2019).
42. Jiang, J. & Fan, J. A. Simulator-based training of generative neural networks for the inverse design of metasurfaces. *Nanophotonics* **9**, 1059–1069 (2019).
43. Ma, W., Cheng, F., Xu, Y., Wen, Q. & Li, Y. Probabilistic representation and inverse design of metamaterials based on a deep generative model with semi-supervised learning strategy. *Adv. Mater.* **31**, 1901111 (2019).
44. Liu, Z., Zhu, D., Rodrigues, S. P., Lee, K.-T. & Cai, W. Generative model for the inverse design of metasurfaces. *Nano Lett.* **18**, 6570–6576 (2018).
45. So, S. & Rho, J. Designing nanophotonic structures using conditional deep convolutional generative adversarial networks. *Nanophotonics* **8**, 1255–1261 (2019).
46. An, S. et al. Multifunctional metasurface design with a generative adversarial network. Preprint at <https://arxiv.org/abs/1908.04851> (2019).
47. Jiang, J. et al. Free-form diffractive metagrating design based on generative adversarial networks. *ACS Nano* **13**, 8872–8878 (2019).
48. Liu, Z., Zhu, Z. & Cai, W. Topological encoding method for data-driven photonics inverse design. *Opt. Express* **28**, 4825–4835 (2020).
49. Van den Oord, A. et al. Conditional image generation with pixelcnn decoders. *Adv. Neural Inf. Process. Syst.* **29**, 4790–4798 (2016).
50. Krizhevsky, A., Sutskever, I. & Hinton, G. E. in *Advances in Neural Information Processing Systems* 1097–1105 (Association for Computing Machinery, 2012).
51. Szegedy, C., Ioffe, S., Vanhoucke, V. & Alemi, A. A. in *Thirty-First AAAI Conference on Artificial Intelligence* 4278–4284 (Association for Computing Machinery, 2017).
52. Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. Preprint at [arXiv https://arxiv.org/abs/1409.1556](https://arxiv.org/abs/1409.1556) (2014).
53. Szegedy, C. et al. in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 1–9 (IEEE, 2015).
54. Wiecha, P. R. & Muskens, O. L. Deep learning meets nanophotonics: A generalized accurate predictor for near fields and far fields of arbitrary 3D nanostructures. *Nano Lett.* **20**, 329–338 (2020).
55. Henaff, M., Bruna, J. & LeCun, Y. Deep convolutional networks on graph-structured data. Preprint at [arXiv https://arxiv.org/abs/1506.05163](https://arxiv.org/abs/1506.05163) (2015).
56. Niepert, M., Ahmed, M. & Kutzkov, K. in *Proceedings of the 33rd International Conference on Machine Learning*, PMLR Vol. 48 2014–2023 (2016).
57. Veličković, P. et al. Graph attention networks. Preprint at [arXiv https://arxiv.org/abs/1710.10903](https://arxiv.org/abs/1710.10903) (2017).
58. Bapst, V. et al. Unveiling the predictive power of static structure in glassy systems. *Nat. Phys.* **16**, 448–454 (2020).
59. Duvenaud, D. K. et al. Convolutional fingerprints on graphs for learning molecular fingerprints. *Adv. Neural Inf. Process. Syst.* **28**, 2224–2232 (2015).
60. Torng, W. & Altman, R. B. Graph convolutional neural networks for predicting drug-target interactions. *J. Chem. Inf. Model.* **59**, 4131–4149 (2019).
61. Li, Y., Tarlow, D., Brockschmidt, M. & Zemel, R. Gated graph sequence neural networks. Preprint at <https://arxiv.org/abs/1511.05493> (2015).
62. Ma, T., Chen, J. & Xiao, C. Constrained generation of semantically valid graphs via regularizing variational autoencoders. *Adv. Neural Inf. Process. Syst.* **31**, 7113–7124 (2018).
63. Graves, A. Generating sequences with recurrent neural networks. Preprint at <https://arxiv.org/abs/1308.0850> (2013).
64. Sutskever, I., Vinyals, O. & Le, Q. V. Sequence to sequence learning with neural networks. *Adv. Neural Inf. Process. Syst.* **2**, 3104–3112 (2014).
65. Luong, M.-T., Pham, H. & Manning, C. D. Effective approaches to attention-based neural machine translation. Preprint at <https://arxiv.org/abs/1508.04025> (2015).
66. Weston, J., Chopra, S. & Bordes, A. Memory networks. Preprint at <https://arxiv.org/abs/1410.3916> (2014).
67. Hughes, T. W., Williamson, I. A., Minkov, M. & Fan, S. Wave physics as an analog recurrent neural network. *Sci. Adv.* **5**, eaay6946 (2019).
68. Vai, M. & Prasad, S. Automatic impedance matching with a neural network. *IEEE Microw. Guided Wave Lett.* **3**, 353–354 (1993).
69. Zaabab, A. H., Zhang, Q. J. & Nakhla, M. in *1994 IEEE MTTTS International Microwave Symposium Digest* (Cat. No. 94CH3389-4) Vol. 1 393–396 (IEEE, 1994).
70. Vai, M. M., Wu, S., Li, B. & Prasad, S. Reverse modeling of microwave circuits with bidirectional neural network models. *IEEE Trans. Microw. Theory Tech.* **46**, 1492–1494 (1998).
71. Watson, P. M. & Gupta, K. C. Design and optimization of CPW circuits using EM-ANN models for CPW components. *IEEE Trans. Microw. Theory Tech.* **45**, 2515–2523 (1997).
72. Devabhaktuni, V., Chattaraj, B., Yagoub, M. C. E. & Zhang, Q. J. in *2002 IEEE MTTTS International Microwave Symposium Digest* (Cat. No. 02CH37278) Vol. 2 1097–1100 (IEEE, 2002).
73. Bandler, J. W. et al. Space mapping: the state of the art. *IEEE Trans. Microw. Theory Tech.* **52**, 337–361 (2004).
74. da F. Silva, P. H., Cruz, R. M. S. & D'Assuncao, A. G. Blending PSO and ANN for optimal design of FSS filters with Koch island patch elements. *IEEE Trans. Magnetics* **46**, 3010–3013 (2010).
75. da Silva, M. R., de L. Nóbrega, C., da F. Silva, P. H. & D'Assunção, A. G. Optimization of FSS with Sierpinski island fractal elements using population-based search algorithms and MLP neural network. *Microw. Optical Technol. Lett.* **56**, 827–831 (2014).
76. Luna, D. R., Vasconcelos, C. F. L. & Cruz, R. M. S. in *2013 SBMO/IEEE MTTTS International Microwave Optoelectronics Conference (IMOC)* 1–4 (IEEE, 2013).
77. Zhang, Q. et al. Machine-learning designs of anisotropic digital coding metasurfaces. *Adv. Theory Simul.* **2**, 1800132 (2019).
78. Liu, Y., Lu, T., Wu, K. & Jin, J.-M. in *2018 IEEE 27th Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)* 261–263 (IEEE, 2018).
79. Zhang, G., He, H. & Katabi, D. in *Proceedings of the 36th International Conference on Machine Learning*, PMLR Vol. 97 (eds Chaudhuri, K. & Salakhutdinov, R.) 7364–7373 (2019).
80. Zhang, Q. J., Gupta, K. C. & Devabhaktuni, V. K. Artificial neural networks for RF and microwave design – from theory to practice. *IEEE Trans. Microw. Theory Tech.* **51**, 1339–1350 (2003).
81. Rayas-Sanchez, J. E. EM-based optimization of microwave circuits using artificial neural networks: The state-of-the-art. *IEEE Trans. Microw. Theory Tech.* **52**, 420–435 (2004).
82. Koziel, S., Cheng, Q. S. & Bandler, J. W. Space mapping. *IEEE Microw. Mag.* **9**, 105–122 (2008).
83. Malheiros-Silveira, G. N. & Hernandez-Figueroa, H. E. Prediction of dispersion relation and PBGs in 2-D PCs by using artificial neural networks. *IEEE Photonics Technol. Lett.* **24**, 1799–1801 (2012).
84. El-Mosalmly, D. D., Hameed, M. F. O., Areed, N. F. F. & Obayya, S. S. A. Novel neural network based optimization approach for photonic devices. *Optical Quantum Electron.* **46**, 439–453 (2014).
85. Soliman, R. R., Swillam, M. A., El-Gamal, M. A. & Adwan, E. A. Artificial neural network modeling of plasmonic transmission lines. *Appl. Opt.* **55**, 2780–2790 (2016).
86. Chugh, S., Gulistan, A., Ghosh, S. & Rahman, B. M. A. Machine learning approach for computing optical properties of a photonic crystal fiber. *Opt. Express* **27**, 36414–36425 (2019).
87. da Silva Ferreira, A., Malheiros-Silveira, G. N. & Hernández-Figueroa, H. E. Computing optical properties of photonic crystals by using multilayer perceptron and extreme learning machine. *J. Lightwave Technol.* **36**, 4066–4073 (2018).
88. Asano, T. & Noda, S. Optimization of photonic crystal nanocavities based on deep learning. *Opt. Express* **26**, 32704–32717 (2018).
89. Zhang, T. et al. Efficient spectrum prediction and inverse design for plasmonic waveguide systems based on artificial neural networks. *Photon. Res.* **7**, 368–380 (2019).
90. da Silva Ferreira, A., da Silva Santos, C. H., Gonçalves, M. S. & Figueroa, H. E. H. Towards an integrated evolutionary strategy and artificial neural network computational tool for designing photonic coupler devices. *Appl. Soft Comput.* **65**, 1–11 (2018).
91. Tahersima, M. H. et al. Deep neural network inverse design of integrated photonic power splitters. *Sci. Rep.* **9**, 1368 (2019).
92. Gabr, A. M. et al. Design and optimization of optical passive elements using artificial neural networks. *J. Opt. Soc. Am. B* **36**, 999–1007 (2019).
93. Hammond, A. M. & Camacho, R. M. Designing integrated photonic devices using artificial neural networks. *Opt. Express* **27**, 29620–29638 (2019).
94. Gostimirovic, D. & Ye, W. N. An open-source artificial neural network model for polarization-insensitive silicon-on-insulator subwavelength grating couplers. *IEEE J. Sel. Top. Quantum Electron.* **25**, 1–5 (2019).
95. Miyatake, Y., Sekine, N., Toprasertpong, K., Takagi, S. & Takenaka, M. Computational design of efficient grating couplers using artificial intelligence. *Jpn. J. Appl. Phys.* **59**, SGGE09 (2020).
96. Peurifoy, J. et al. Nanophotonic particle simulation and inverse design using artificial neural networks. *Sci. Adv.* **4**, eaar4206 (2018).
97. Ma, W., Cheng, F. & Liu, Y. Deep-learning-enabled on-demand design of chiral metamaterials. *ACS Nano* **12**, 6326–6334 (2018).
98. Li, Y. et al. Self-learning perfect optical chirality via a deep neural network. *Phys. Rev. Lett.* **123**, 213902 (2019).
99. Tao, Z. et al. Exploiting deep learning network in optical chirality tuning and manipulation of diffractive chiral metamaterials. *Nanophotonics* **9**, 2945–2956 (2020).
100. Malkiel, I. et al. Plasmonic nanostructure design and characterization via deep learning. *Light Sci. Appl.* **7**, 60 (2018).
101. Sajedian, I., Kim, J. & Rho, J. Finding the optical properties of plasmonic structures by image processing using a combination of convolutional neural networks and recurrent neural networks. *Microsyst. Nanoeng.* **5**, 27 (2019).
102. Baxter, J. et al. Plasmonic colours predicted by deep learning. *Sci. Rep.* **9**, 8074 (2019).
103. Balin, I., Garmider, V., Long, Y. & Abdulhalim, I. Training artificial neural network for optimization of nanostructured VO<sub>2</sub>-based smart window performance. *Opt. Express* **27**, A1030–A1040 (2019).
104. Blanchard-Dionne, A.-P. & Martin, O. J. F. Teaching optics to a machine learning network. *Opt. Lett.* **45**, 2922–2925 (2020).
105. Inampudi, S. & Mosallaei, H. Neural network based design of metagratings. *Appl. Phys. Lett.* **112**, 241102 (2018).
106. An, S. et al. A deep learning approach for objective-driven all-dielectric metasurface design. *ACS Photonics* **6**, 3196–3207 (2019).
107. Chen, Y., Zhu, J., Xie, Y., Feng, N. & Liu, Q. H. Smart inverse design of graphene-based photonic metamaterials by an adaptive artificial neural network. *Nanoscale* **11**, 9749–9755 (2019).
108. Phan, A. D. et al. Deep learning for the inverse design of mid-infrared graphene plasmons. *Crystals* **10**, 125 (2020).
109. Gao, L., Li, X., Liu, D., Wang, L. & Yu, Z. A bidirectional deep neural network for accurate silicon color design. *Adv. Mater.* **31**, 1905467 (2019).
110. Liu, D., Tan, Y., Khoram, E. & Yu, Z. Training deep neural networks for the inverse design of nanophotonic structures. *ACS Photonics* **5**, 1365–1369 (2018).
111. Piloizzi, L., Farrelly, F. A., Marcucci, G. & Conti, C. Machine learning inverse problem for topological photonics. *Commun. Phys.* **1**, 57 (2018).

112. Long, Y., Ren, J., Li, Y. & Chen, H. Inverse design of photonic topological state via machine learning. *Appl. Phys. Lett.* **114**, 181105 (2019).
113. Singh, R., Agarwal, A. M. & Anthony, B. W. Mapping the design space of photonic topological states via deep learning. *Opt. Express* **28**, 27893–27902 (2020).
114. Zaabab, A. H., Zhang, Q. J. & Nakhla, M. A neural network modeling approach to circuit optimization and statistical design. *IEEE Trans. Microw. Theory Tech.* **43**, 1349–1358 (1995).
115. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. Preprint at <https://arxiv.org/abs/1412.6980> (2014).
116. Hegde, R. S. Accelerating optics design optimizations with deep learning. *Opt. Eng.* **58**, 065103 (2019).
117. Lin, R., Zhai, Y., Xiong, C. & Li, X. Inverse design of plasmonic metasurfaces by convolutional neural network. *Opt. Lett.* **45**, 1362–1365 (2020).
118. Zhang, C., Jin, J., Na, W., Zhang, Q. & Yu, M. Multivalued neural network inverse modeling and applications to microwave filters. *IEEE Trans. Microw. Theory Tech.* **66**, 3781–3797 (2018).
119. So, S., Mun, J. & Rho, J. Simultaneous inverse design of materials and structures via deep learning: Demonstration of dipole resonance engineering using core-shell nanoparticles. *ACS Appl. Mater. Interfaces* **11**, 24264–24268 (2019).
120. Xu, L. et al. Enhanced light-matter interactions in dielectric nanostructures via machine-learning approach. *Adv. Photon.* **2**, 026003 (2019).
121. Zimek, A., Schubert, E. & Kriegel, H.-P. A survey on unsupervised outlier detection in high-dimensional numerical data. *Stat. Anal. Data Min.* **5**, 363–387 (2012).
122. Marimont, R. & Shapiro, M. Nearest neighbour searches and the curse of dimensionality. *IMA J. Appl. Math.* **24**, 59–70 (1979).
123. Melati, D. et al. Mapping the global design space of nanophotonic components using machine learning pattern recognition. *Nat. Commun.* **10**, 4775 (2019).
124. Kramer, M. A. Nonlinear principal component analysis using autoassociative neural networks. *AIChE J.* **37**, 235–243 (1991).
125. Baldi, P. *Proceedings of the 2011 International Conference on Unsupervised and Transfer Learning workshop* **27**, 37–49 (2012).
126. Kiarashinejad, Y., Abdollahramezani, S. & Adibi, A. Deep learning approach based on dimensionality reduction for designing electromagnetic nanostructures. *NPI Comput. Mater.* **6**, 12 (2020).
127. Shi, X., Qiu, T., Wang, J., Zhao, X. & Qu, S. Metasurface inverse design using machine learning approaches. *J. Phys. D Appl. Phys.* **53**, 275105 (2020).
128. Kiarashinejad, Y. et al. Knowledge discovery in nanophotonics using geometric deep learning. *Adv. Intell. Syst.* **2**, 1900132 (2020).
129. Karras, T., Aila, T., Laine, S. & Lehtinen, J. Progressive growing of GANs for improved quality, stability, and variation. Preprint at <https://arxiv.org/abs/1710.10196> (2018).
130. Salimans, T. et al. Improved techniques for training GANs. Preprint at <https://arxiv.org/abs/1606.03498> (2016).
131. Sell, D., Yang, J., Doshay, S., Yang, R. & Fan, J. A. Large-angle, multifunctional metagratings based on freeform multimode geometries. *Nano Lett.* **17**, 3752–3757 (2017).
132. Phan, T. et al. High-efficiency, large-area, topology-optimized metasurfaces. *Light Sci. Appl.* **8**, 48 (2019).
133. Fan, J. A. Freeform metasurface design based on topology optimization. *MRS Bull.* **45**, 196–201 (2020).
134. Hughes, T. W., Minkov, M., Williamson, I. A. & Fan, S. Adjoint method and inverse design for nonlinear nanophotonic devices. *ACS Photonics* **5**, 4781–4787 (2018).
135. Maclaurin, D., Duvenaud, D. & Adams, R. P. in *Proceedings of the 32nd International Conference on International Conference on Machine Learning* Vol. 37 2113–2122 (JMLR, 2015).
136. Minkov, M. et al. Inverse design of photonic crystals through automatic differentiation. *ACS Photonics* **7**, 1729–1741 (2020).
137. Hughes, T. W., Williamson, I. A., Minkov, M. & Fan, S. Forward-mode differentiation of Maxwell's equations. *ACS Photonics* **6**, 3010–3016 (2019).
138. Uria, B., Murray, I. & Larochelle, H. RNAE: The real-valued neural autoregressive density-estimator. Preprint at <https://arxiv.org/abs/1306.0186> (2013).
139. Oord, A. V. D. et al. WaveNet: A generative model for raw audio. Preprint at <https://arxiv.org/abs/1609.03499> (2016).
140. Larochelle, H. & Murray, I. in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics* **15**, 29–37 (2011).
141. van den Oord, A., Kalchbrenner, N. & Kavukcuoglu, K. Pixel recurrent neural networks. Preprint at <https://arxiv.org/abs/1601.06759> (2016).
142. Kingma, D. P. & Welling, M. Auto-encoding variational bayes. Preprint at <https://arxiv.org/abs/1312.6114> (2013).
143. Goodfellow, I. et al. in *Advances in Neural Information Processing Systems* Vol. 27 (eds Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D. & Weinberger, K. Q.) 2672–2680 (Curran Associates, 2014).
144. Sohn, K., Lee, H. & Yan, X. in *Advances in Neural Information Processing Systems* Vol. 28 (eds Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M. & Garnett, R.) 3483–3491 (Curran Associates, 2015).
145. Makhlani, A., Shlens, J., Jaitly, N., Goodfellow, I. & Frey, B. Adversarial autoencoders. Preprint at [arXiv https://arxiv.org/abs/1501.05644](https://arxiv.org/abs/1501.05644) (2015).
146. Kingma, D. P., Mohamed, S., Rezende, D. J. & Welling, M. Semi-supervised learning with deep generative models. *Adv. Neural Inf. Process. Syst.* **27**, 3581–3589 (2014).
147. Liu, Z., Raju, L., Zhu, D. & Cai, W. A hybrid strategy for the discovery and design of photonic structures. *IEEE J. Emerg. Sel. Top. Circuits Syst.* **10**, 126–135 (2020).
148. Radford, A., Metz, L. & Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. Preprint at <https://arxiv.org/abs/1511.06434> (2015).
149. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V. & Courville, A. C. Improved training of Wasserstein GANs. *Adv. Neural Inf. Process. Syst.* **30**, 5767–5777 (2017).
150. Brock, A., Donahue, J. & Simonyan, K. Large scale GAN training for high fidelity natural image synthesis. Preprint at <https://arxiv.org/abs/1809.11096> (2018).
151. Arjovsky, M., Chintala, S. & Bottou, L. in *Proceedings of the 34th International Conference on Machine Learning, Proceedings of Machine Learning Research* Vol. 10 (eds Precup, D. & Teh, Y. W.) 214–223 (PMLR, 2017).
152. Hodge, J. A., Mishra, K. V. & Zaghoul, A. I. in *2019 IEEE International Symposium on Phased Array System Technology (PAST)* 1–6 (IEEE, 2019).
153. Hodge, J. A., Vijay Mishra, K. & Zaghoul, A. I. in *2019 IEEE Global Conference on Signal and Information Processing (GlobalSIP)* 1–5 (IEEE, 2019).
154. Kudyshev, Z. A., Kildishev, A. V., Shalaev, V. M. & Boltasseva, A. in *CLEO: QELS\_Fundamental Science, FTh3C.2* (Optical Society of America, 2019).
155. Kudyshev, Z. A., Kildishev, A. V., Shalaev, V. M. & Boltasseva, A. Machine-learning-assisted metasurface design for high-efficiency thermal emitter optimization. *Appl. Phys. Rev.* **7**, 021407 (2020).
156. Yang, J. & Fan, J. A. Topology-optimized metasurfaces: impact of initial geometric layout. *Opt. Lett.* **42**, 3161–3164 (2017).
157. Yang, J., Sell, D. & Fan, J. A. Freeform metagratings based on complex light scattering dynamics for extreme, high efficiency beam steering. *Ann. Phys.* **530**, 1700302 (2018).
158. Yang, J. & Fan, J. A. Analysis of material selection on dielectric metasurface performance. *Opt. Express* **25**, 23899–23909 (2017).
159. Sell, D., Yang, J., Doshay, S. & Fan, J. A. Periodic dielectric metasurfaces with high-efficiency, multiwavelength functionalities. *Adv. Opt. Mater.* **5**, 1700645 (2017).
160. Sell, D. et al. Ultra-high-efficiency anomalous refraction with dielectric metasurfaces. *ACS Photonics* **5**, 2402–2407 (2018).
161. Wang, E. W., Sell, D., Phan, T. & Fan, J. A. Robust design of topology-optimized metasurfaces. *Opt. Mater. Express* **9**, 469–482 (2019).
162. Wen, F., Jiang, J. & Fan, J. A. Progressive-growing of generative adversarial networks for metasurface optimization. *ACS Photonics* **7**, 2098–2104 (2019).
163. Vercrussse, D., Sapra, N. V., Su, L., Trivedi, R. & Vucković, J. Analytical level set fabrication constraints for inverse design. *Sci. Rep.* **9**, 8999 (2019).
164. Huang, J. et al. Implementation of on-chip multi-channel focusing wavelength demultiplexer with regularized digital metamaterials. *Nanophotonics* **9**, 159–166 (2019).
165. Vercrussse, D., Sapra, N. V., Su, L. & Vucković, J. Dispersion engineering with photonic inverse design. *IEEE J. Sel. Top. Quantum Electron.* **26**, 1–6 (2019).
166. Chen, M., Jiang, J. & Fan, J. Design space reparameterization enforces hard geometric constraints in inverse-designed nanophotonic devices. *ACS Photonics* **7**, 3141–3151 (2020).
167. Jiang, J. & Fan, J. A. Multiobjective and categorical global optimization of photonic structures based on ResNet generative neural networks. *Nanophotonics* <https://doi.org/10.1515/nanoph-2020-0407> (2020).
168. Azunre, P. et al. Guaranteed global optimization of thin-film optical systems. *New J. Phys.* **21**, 073050 (2019).
169. Griebel, T. G. M. Numerical integration using sparse grids. *Numer. Algorithms* **18**, 209 (1998).
170. Schneider, P.-I. et al. Benchmarking five global optimization approaches for nano-optical shape optimization and parameter reconstruction. *ACS Photonics* **6**, 2726–2733 (2019).
171. Elsayy, M. M. et al. Global optimization of metasurface designs using statistical learning methods. *Sci. Rep.* **9**, 17918 (2019).
172. Raissi, R., Perdikaris, P. & Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019).
173. Trivedi, R., Su, L., Lu, J., Schubert, M. F. & Vucković, J. Data-driven acceleration of photonic simulations. *Sci. Rep.* **9**, 19728 (2019).
174. Chew, W. C., Tong, M. S. & Hu, B. Integral equation methods for electromagnetic and elastic waves. *Synth. Lect. Comput. Electromagnetics* **3**, 1–241 (2008).
175. Mishchenko, M. I., Travis, L. D. & Mackowski, D. W. T-matrix computations of light scattering by nonspherical particles: A review. *J. Quant. Spectrosc. Radiat. Transf.* **55**, 535–575 (1996).
176. Pan, S. J. & Yang, Q. A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* **22**, 1345–1359 (2010).
177. Qu, Y., Jing, L., Shen, Y., Qiu, M. & Soljacic, M. Migrating knowledge between physical scenarios based on artificial neural networks. *ACS Photonics* **6**, 1168–1174 (2019).
178. Andrychowicz, M. et al. Learning to learn by gradient descent by gradient descent. *Adv. Neural Inf. Process. Syst.* **29**, 3981–3989 (2016).
179. Finn, C., Abbeel, P. & Levine, S. in *Proceedings of the 34th International Conference on Machine Learning* Vol. 70 1126–1135 (JMLR, 2017).
180. Deng, J. et al. in *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition* 248–255 (2009).
181. Krizhevsky, A. *Learning Multiple Layers of Features from Tiny Images* (Univ. Toronto, 2012).
182. Jiang, J. et al. MetaNet: a new paradigm for data sharing in photonics research. *Opt. Express* **28**, 13670–13681 (2020).
183. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning* (MIT Press, 2016).
184. Kullback, S. *Information Theory and Statistics* (Courier Corporation, 1997).
185. Lin, J. Divergence measures based on the Shannon entropy. *IEEE Trans. Inf. Theory* **37**, 145–151 (1991).
186. Kudyshev, Z. A., Kildishev, A. V., Shalaev, V. & Boltasseva, A. in *2019 International Applied Computational Electromagnetics Society Symposium (ACES)* 1–2 (ACES, 2019).
187. Chen, M., Jiang, J. & Fan, J. in *2020 Conference on Lasers and Electro-Optics (CLEO): QELS\_Fundamental Science FW4B-4* (Optical Society of America, 2020).

**Acknowledgements**

The authors thank Q. Huang, L. Gan and K. Kanada for feedback on the manuscript. This work was supported by the Office of Naval Research under award number N00014-20-1-2105, ARPA-E under award number DE-AR0001212 and the David and Lucile Packard Foundation under award number 2016-65132.

**Author contributions**

All authors researched data for the article, discussed the content and contributed to the writing and revising of the manuscript.

**Competing interests**

The authors declare no competing interests.

**Publisher's note**

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.