

KEY ISSUES REVIEW

Artificial intelligence for photonics and photonic materials

To cite this article: Davide Piccinotti *et al* 2021 *Rep. Prog. Phys.* **84** 012401

View the [article online](#) for updates and enhancements.

You may also like

- [Effective field theories of post-Newtonian gravity: a comprehensive review](#)
Michèle Levi
- [A review on partial-wave dynamics with chiral effective field theory and dispersion relation](#)
De-Liang Yao, Ling-Yun Dai, Han-Qing Zheng *et al.*
- [Observational diversity of magnetized neutron stars](#)
Teruaki Enoto, Shota Kisaka and Shinpei Shibata

Recent citations

- [Sunil Sharma and Lokesh Tharani](#)
- [Deep Learning the Electromagnetic Properties of Metamaterials—A Comprehensive Review](#)
Omar Khatib *et al*
- [Deep learning in nano-photonics: inverse design and beyond](#)
Peter R. Wiecha *et al*




IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

Key Issues Review

Artificial intelligence for photonics and photonic materials

Davide Piccinotti^{1,*} , Kevin F MacDonald^{1,*} , Simon A Gregory², Ian Youngs² and Nikolay I Zheludev^{1,3} 

¹ Optoelectronics Research Centre and Centre for Photonic Metamaterials, University of Southampton, Southampton, SO17 1BJ, United Kingdom

² Defence Science and Technology Laboratory, Salisbury, SP4 0JQ, United Kingdom

³ Centre for Disruptive Photonic Technologies, The Photonics Institute, School of Physical and Mathematical Sciences, Nanyang Technological University, 637371 Singapore

E-mail: D.Piccinotti@soton.ac.uk and kfm1@soton.ac.uk

Received 25 February 2020, revised 9 July 2020

Accepted for publication 2 September 2020

Published 18 December 2020



Abstract

Artificial intelligence (AI) is the most important new methodology in scientific research since the adoption of quantum mechanics and it is providing exciting results in numerous fields of science and technology. In this review we summarize research and discuss future opportunities for AI in the domains of photonics, nanophotonics, plasmonics and photonic materials discovery, including metamaterials.

Keywords: machine learning, deep learning, photonics, nanophotonics, photonic devices, materials discovery, photonic materials, metamaterials, plasmonics

(Some figures may appear in colour only in the online journal)

1. Introduction

Artificial intelligence (AI) is currently one of the most active fields in research, mainly due to the impressive results obtained via machine learning (ML) in fields ranging from genetics and synthetic chemistry to speech recognition and image processing. The technology enables computers to learn on their own, without being explicitly programmed, to identify patterns in data, build models that explain the world, and make predictions that do not explicitly follow from predefined rules and models [1].

AI represents a potent addition to the computational toolkit for analysis and interpretation of data, predictive modelling, and the automated, (self-)adaptive design and control of systems and processes. Its greatest power lies in its ability to address big data, multi-dimensional and high-complexity

problems that are intractable with conventional numerical and analytical computational techniques. It will become a major instrument/methodology for and driver of discovery and development in all areas of science and technology, including photonics/electromagnetism—the focus of this review.

ML methods hold great potential for use in the development of new materials. So far in this field, ML algorithms have typically been applied to large databases (themselves often generated computationally) for the purpose of predicting new materials and revealing trends that would not easily be seen by conventional data mining techniques. They have also been used to grow small experimental datasets in an iterative process that predicts new materials [2]. ML can furthermore facilitate multiobjective optimization of material properties; for example, in the search for NiTi-based shape-memory alloys with very small thermal hysteresis, Xue *et al* [3] began with only 22 experimental samples and found the optimal compound (with a hysteresis of only 1.84 K) on the sixth iteration.

* Author to whom any correspondence should be addressed.

†Corresponding editor: Professor Masud Mansuripur†.

In the domain of nanophotonics and metamaterials, the design and optimization of devices to control light–matter interactions and manipulate electromagnetic fields at the sub-wavelength scale is computationally demanding (indeed, often prohibitively so). Here, ML techniques present opportunities both to extend physical insight and to search parameter spaces in a more efficient way, leading to data-driven, on-demand design of novel devices [4]. For instance, Peurifoy *et al* [5] trained a neural network to approximate light scattering by multilayer nanoparticles and, by taking advantage of the fact that it could do so orders of magnitude faster than conventional simulations, employed it to solve nanophotonic inverse design problems via back propagation.

Recent years have also seen growing interest in photonic implementations of AI (cf. all-optical computing), with reported approaches including photonic reservoir and neuro-morphic optical computing, photonic deep neural networks, and all-optical photonic cognitive networks (e.g. optical oracles). For example, Wu *et al* [6] demonstrated that a fibre network can be considered as an optical oracle for the ‘Hamiltonian path’ problem. It is able to address the famous mathematical complexity problem (of determining whether a set of locations can be visited via a path that passes each location only once) hundreds of times faster than brute-force (electronic) computing. Estakhri *et al* [7] have designed meta-structures to perform analogue computing and solve integral equations using wave interactions.

It is now evident that AI and photonics provide huge opportunities for both fields, it is essentially a two-way road: photonics is an ideal platform for ML methodology since the Maxwell equations can be used to generate extremely reliable source of training data and, at the same time photonic systems offer an extremely interesting platform for the implementation of AI due to the inherent parallelism of photonics.

The use of ML in materials science and (nano)photonics is still in its infancy and the landscape of existing work at this intersection is relatively sparse. Nonetheless, key themes and activity areas are emerging, and the first international workshop on ‘AI in Nanophotonics’ was held in 2019 (as part of the 15th International Work-Conference on Artificial Neural Networks, IWANN2019, June 12-14, 2019, Gran Canaria, Spain). Broadly, we consider that there are four areas in which the AI/ML toolkit may be applied:

- Engineering of artificial electromagnetic materials (optimization of metamaterial and other nanostructural geometries to achieve specific, targeted optical properties);
- Optimization of (nano)photonic device, system and network architectures (e.g. adaptively reconfigurable waveguide structures to perform arbitrary input/output operations);
- Discovery and optimization of photonic materials, their synthesis and processing (identification of materials providing selected desirable properties for photonic applications);
- Development of sources of complex electromagnetic fields, e.g. phase singularities, super-oscillatory foci, etc. (design of field structures themselves and evolution of

strategies for generating, characterizing and utilizing such fields).

2. AI methods

AI currently comprises a huge number of subfields, [8] ranging from general learning and performance of to specific tasks such as playing chess, to proving mathematical theorems, writing poetry, driving cars and medical diagnosis [9]. A particular subfield of AI named ML has shown impressive results on classification and prediction tasks, becoming a vital part of many industries [10]; the ML approach is based on its intrinsic ability to solve tasks that are too challenging for conventional programmes [1]. The word ‘task’ is used in ML to define a piece of work allocated to an algorithm and the process of learning itself is not a task. Learning is the procedure through which the algorithm reaches the ability to perform the task [1]. Most ML algorithms can be classified as supervised learning, unsupervised learning and reinforcement learning based on how the learning process is designed [1]. The most common form of ML is supervised learning (figure 1) [11]. The starting point is to learn a mapping function from inputs x to outputs y , given a labelled set of input-output pairs $D = \{(x_i, y_i)\}_{i=1}^N$ where D is called the training set and N is the number of training examples [12]. After this first step, named training, the algorithm has to make predictions on novel inputs (unseen data) and this process is called generalization or cross-validation. When the output y is a categorical or nominal variable (e.g. a word), the problem is known as classification or pattern recognition; when y is real-valued number, the problem is known as regression [12]. Learning to classify handwritten digits is a standard example of a classification task and to attain this result, a supervised learning algorithm takes thousands of pictures of handwritten digits labelled with the number each image represents. Then, the algorithm learns the relationship between the images and their associated numbers, and applies that learned relationship to classify completely new images (without labels) that the machine has not seen before [8].

The unsupervised learning approach (figure 2) is based on unlabelled datasets: an unsupervised algorithm receives only given inputs $D = \{(x_i)\}_{i=1}^N$ and the goal is to find ‘interesting patterns’ in the data [12]. Useful properties of the structure of the given dataset are derived without the guidance of a ‘teacher’ or ‘instructor’ that is present on supervised learning [1]. Two typical examples of unsupervised tasks are clustering the data into groups by similarity (figure 2(a)) and reducing dimensionality (figure 2(b)) to compress the data while maintaining its structure and usefulness [8, 12]. Unsupervised learning is more characteristic of human and animal learning since we acquire our knowledge mainly by observing the things around us and not by being told the name of every object [11, 12]. An advantage of unsupervised learning is that it does not need manually labelled data, thus its application is broader than supervised learning [12].

The third type of ML, known as reinforcement learning, is a computational approach based on learning from interaction with an environment, whereby there is a feedback loop between the learning system and its experiences [1, 13]. The

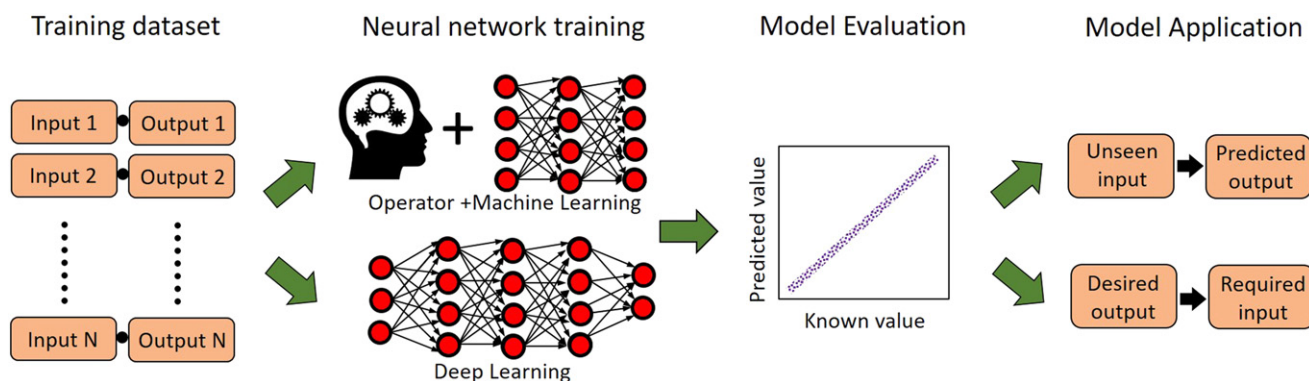



Figure 1. Supervised learning and deep learning. In supervised learning, the training dataset consists of a large number of input values (e.g. material compositions) linked to the corresponding output values (e.g. optical properties). At the training phase a model for a relationship between material composition and its properties is established. This can be achieved by a ‘deep learning’ algorithm or by a combination of a simpler machine algorithm and human analysis. The algorithm often takes the form of a stacked neural network connecting input and output data and made of nodes. A node combines input from the data with a set of weights, thereby assigning significance to inputs with regard to the task the algorithm is trying to develop. The weighted inputs are summed and passed through a node’s activation function to the next layer of the network, to determine to what extent that signal should progress further through the network to affect the ultimate outcome. Once the model is established and evaluated, it can be applied for predictive modelling, for instance for predicting of materials optical properties for previously unseen material compositions. Moreover, predictive modelling can be reversed to identify the required composition of material to achieve desired properties.

(a) Unsupervised learning for clustering



(b) Unsupervised learning for dimensionality reduction




Figure 2. Unsupervised learning. Unsupervised learning is based on unlabelled datasets and the goal of unsupervised algorithms is to find ‘interesting patterns’ in the data (clustering) (a), or to reduce dimensionality to compress the data while maintaining its structure and usefulness (b).

learning algorithm has to accomplish a specific goal for which it is rewarded. So far, this approach is less commonly used but it holds great promise. For instance, Melnikov *et al* [14, 15] investigated the potential of intelligent machines in the context of scientific research using a reinforcement learning methodology.

Conventional ML models have a simple two-layer structure where inputs and outputs are directly connected via a certain function [12]. However, the study of the human brain reveals the existence of many levels of processing and it is believed that each level is learning features or representations at increasing levels of abstraction [12]. This observation has stimulated the development of deep learning, a ML method that attempts to replicate this kind of architecture in a computer [12]. Therefore, deep learning can be considered as an ‘evolution’ of ML and deep learning algorithms can be classified in the same three categories used for ML algorithms.

A deep learning architecture is composed of a multilayer structure of nodes that can learn and compute non-linear input-output mappings [11]. This type of architecture is also called artificial neural network (ANN) [1]. ANNs comprise an input and output layer linked via a certain number of so-called hidden layers. This type of structure allows for direct insertion of raw data into the deep learning algorithm without the need of a human-designed feature extractor that is needed in conventional ML. Indeed, these multiple hidden layers are able to learn the suitable feature extractors directly from the data. So starting from the raw data (for instance a picture) each module or layer transform the representation at one level into a representation at a higher, more abstract level [1]. The composition of these transformations enables the learning of very complex functions. In the context of a classification task, the hidden layers can be considered to distort the input nonlinearly in such a way that the different groups become linearly separable by the last layer [11]. In summary, deep learning is a particular type of ML where the world is seen as a hierarchical structure of concepts and each concept can be described in terms of simpler concepts with more abstract representations computed based on less abstract ones [1]. A more detailed description of ANNs can be found in [11, 16]. In particular, a detailed review of deep supervised learning, unsupervised learning, reinforcement learning and evolutionary computation with an historical perspective summarizing the relevant work (much of it from the previous millennium) is the subject of [17]. Many different architectures have been developed for ANNs; recently the Asimov Institute has introduced a chart containing mostly of these architectures; this chart can be found at www.asimovinstitute.org/neural-network-zoo/ [18].

In recent years, deep learning has become more popular and useful mainly due to the availability of larger computational power, larger datasets and techniques to train deeper networks [1]. It has found many applications in several areas of science, business and government thanks to its capability for discovering complex structures in high-dimensional data. Moreover, deep learning has set a new standard in many different fields like image and speech recognition, drugs discovery and medical diagnosis. Furthermore, deep learning is powerful

in natural language tasks, such as topic classification, sentiment analysis, question answering, and language translation [11].

2.1. Interpretation of deep neural networks

Despite the impressive results achieved with deep neural networks in many different fields, there is still a fundamental issue regarding how these models make their predictions in that they function as ‘black boxes’ that provide results without explaining them [19]. This aspect is particularly important in areas like medical applications or self-driving cars, and so there is an increasing interest in methods for visualizing, explaining and interpreting deep learning models [19–21]. This includes the development of techniques that enable a better comprehension of what the model has learned (i.e. its representation) and also methods for explaining individual predictions, to verify predictions, to reveal hidden flaws and biases that may lead to failure, to ensure legislative compliance, and to learn more about the problem posed in the first place.

Two common techniques for explaining the individual predictions of a deep learning model in terms of input variables are sensitivity analysis (SA) and layer-wise relevance propagation (LRP). These are primarily used for interpretation of deep learning models applied to classification tasks, such as image and speech recognition. We consider an example of the former as presented in [20].

The interpretation procedure is summarized in figure 3: after an input image is correctly classified, the explanation process generates a ‘heatmap’ visualizing the importance of each pixel for the prediction. SA assumes that the most relevant input features are those to which the output is most sensitive—it shows how much changes in each pixel affect the prediction. LRP explains the classifier’s decisions by decomposition; essentially, by propagating the output backwards to arrive at a relevance score for each input variable (pixel)—it shows how much each pixel contributes to the prediction. A detailed description of these techniques is given in [19, 20]. These techniques for interpreting and understanding what the model has learned have become an essential element of a robust validation procedure. Furthermore, the application of these methods together with deep nonlinear ML models has offered new understandings from complex physical, chemical or biological systems [19].

Another method for neural network interpretation has recently been reported by Amey *et al* [21], wherein a group-theoretical procedure is employed to bring intermediate signalling (hidden layers) into a human-readable form that may allow neural network functionality to be replicated using deterministic digital signal processing. The technique is based upon assumptions that fully connected deep neural networks are interpretable and that the descrambled signals have identifiable and quantifiable features—for example, smoothness or locality. It opens a path to deeper study of neural networks and potentially the derivation of new knowledge: so far, networks have only been seen to reinvent mathematics already

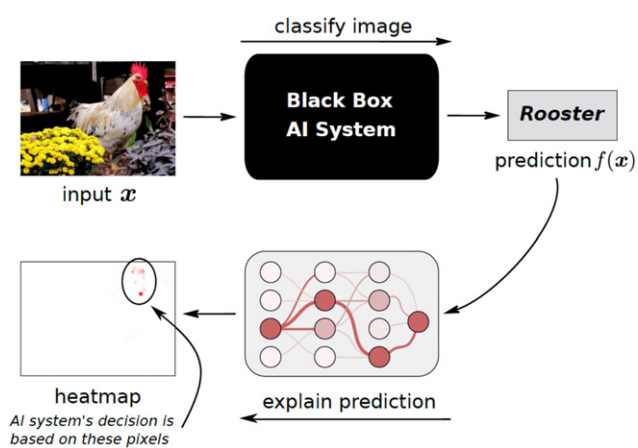


Figure 3. Explaining the predictions of an AI system. The input image is correctly classified as ‘rooster’. A heatmap derived, for example, via SA or LRP reveals the importance of each input pixel for the prediction, helping to verify that the system works as intended [20]. In this example, the rooster’s red comb and wattle are identified as defining characteristics. Adapted from [20]. CC BY-NC-ND 3.0.

known to humans; at some point though, previously unknown mathematics may emerge.

3. AI for materials discovery

The discovery and development of new materials plays an important role in the technological progress, as shown by the way silicon has revolutionized the microelectronics industry. Indeed advances in technological areas such as renewable and sustainable energy depend on the ability to design and realize materials with optimal properties [22]. Materials discovery and design efforts require an interplay between materials prediction, synthesis and characterization. The increased use of computational tools, the generation of materials databases, and advances in experimental methods have significantly accelerated these activities [23]. In this context, ‘materials by design’ approaches have developed with the aim of accelerating materials transfer from computer simulations to lab experimentation and consumer products [23]. The acceleration of materials by design research is highlighted by the almost exponential growth in the number of publications on materials design [23]. This was stimulated by progress on computational techniques like the development of user friendly *ab initio* codes and automation of these codes to run high-throughput computations [23], and more recently by the use of ML methods [24]. Indeed, the concept of rational design of materials through the effective use of data-driven methods is the fundamental principle of the US Materials Genome Initiative [25]. The introduction of ML in materials science was firstly motivated by the generation of large databases mainly by computational methods such as high-throughput density functional theory (HT-DFT) calculations and in some cases by high-throughput experiments [2]. However, in the last few years another approach has been established based on less

common ML tools, which allows one to grow small experimentally generated databases into larger sets and along the way to predict new materials [2].

The application of ML to materials discovery involves the following steps [26]:

- (a) Data collection: the collections of scientific raw data. Can be performed experimentally via the measurements of certain properties of interest or extracted from first-principles calculations like DFT.
- (b) Data representation: the form in which raw data are presented often affects the learning algorithm; the process of converting raw data into something more appropriate for an algorithm is called featurization or feature engineering.
- (c) Choice of model types (or learners): several different learning algorithms can be used based on the type of data and the specific application. Some of the common algorithms are the following:
 1. Naive Bayes classifiers [12]
 2. k-nearest-neighbour methods [12]
 3. Decision trees [12]
 4. Kernel methods [12]
 5. ANNs and deep neural networks [1]
- (d) Model optimization: the selected algorithm is applied to unseen data and its accuracy is evaluated. The three main sources of error are: model bias, model variance and irreducible errors; the total error of the model is the sum of these.

A detailed review of ML methods used in materials science can be found in [26, 27]. This section on ‘AI for materials discovery’ covers the use of ML methods for the development of new materials with any desired property and is broadly divided in three subsections: the first one is on large dataset examples, the second is on small dataset examples (figure 4) and multiobjective optimization, and the last is about methods for achieving inverse design, which aim to discover tailored materials from the starting point of a particular desired functionality.

3.1. Application of AI methods to large datasets

Today, the biggest databases of structural and other material information are those recently created by high-throughput density functional theory (HT-DFT) calculations [2]. These large databases (which have millions of entries) are used as training data for ML algorithms in the context of supervised learning. The predictive ability of these algorithms is tested on unseen data and verified experimentally. This blend of computation and experimentation with data science for the development of materials is named materials informatics. Due to its potential for discovering new materials and efficiently searching materials databases, materials informatics has become important tool in materials science [28–30]. Nevertheless, materials informatics techniques [31] can still be improved in some key regards, such as [23, 30]:

- (a) Quality of data: the management and inspection of materials databases for the presence of errors is difficult due

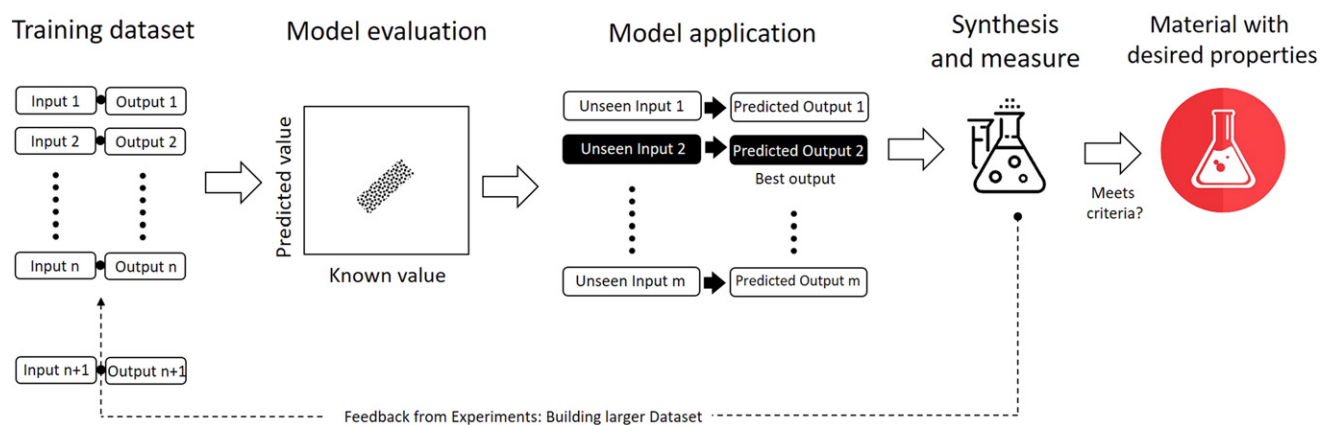


Figure 4. Material discovery based on small datasets (after [43]). In the search for materials with desired properties, a small set of experimental data on structure and corresponding properties is initially used for training. The model is applied to a set of unexplored structures to identify the best material with properties closest to the target. This material is synthesised and results of its characterization are added to the dataset to improve the model in iterative fashion.

to the large number of sample points that they include; however, the identification and removal of such errors is fundamental for obtaining accurate predictions from ML algorithms [30].

- (b) Database management: the collection and organization of big data is quite challenging, especially for experimental databases since they depend on the experimental methodology used. Computational databases are in general easier to organize because data are generated using supercomputers, nevertheless they are limited on the material properties that can be calculated if first-principle calculations are used. Furthermore, the collection of data for materials informatics requires both positive and negative results for the efficient application of ML algorithms while most of the time negative results are not considered. Other issues include database accessibility, since some databases are not open source, and the lack of a global standard method of organizing data; both of these problems prevent the creation of a global database [30].
- (c) Data representation: the parametrization of raw data is an essential step in ML and this process defines the variables that are used by the learning algorithm. These variables are also called features or descriptors and connect the microscopic parameters (for example, formation and defect energies, atomic environments, band structure, density of states or magnetic moments) to macroscopic properties of the materials (for example, mobility, susceptibility or critical temperatures) [32]. Discovering descriptors corresponding to material properties is usually not a trivial process but is a critical step in materials informatics [30].
- (d) Platform design: so far, these processes are realized with different software using different programming languages. Uniting all these processes under one software platform available via the Internet will benefit researchers, who could add new data, extract existing data or predict

unknown materials in a more consistent and systematic way [30].

Despite these constraints, the application of ML methods to big data has provided many successful results, for example, Jain *et al* [29] have developed a probabilistic model for predicting the dominant character (element, charge state, and orbital type) of the electronic states near the valence band maximum and conduction band minimum over a broad range of compounds. Many materials properties, including the band gap (important for light capture) and Seebeck coefficient (important for thermoelectrics) depend critically on the details of these band edges. Understanding the character of these states allows one to determine the type of modifications needed to achieve desired properties. Another interesting example is the research work conducted by Kusne *et al* [33], where they have developed a learning algorithm which allows on-the-fly analysis of diffraction data collected at a synchrotron beamline from combinatorial libraries as they are measured and rapidly cross-referenced with pre-selected entries from the inorganic crystal structure database. They have used this approach to identify a novel magnetic phase with enhanced magnetic anisotropy, which is a candidate for rare-earth-free permanent magnets. More examples related to big data science generated by synchrotrons and strategies adopted to deal with these big data can be found in a recent review article [34]. In addition to these examples, Stein *et al* [35] have built and trained three distinct ML models that are able to predict full UV (ultraviolet)—visible absorption spectra from just sample images (figure 5(d)). Furthermore, they have shown that band gap energies extracted from the predicted spectra are very accurate and have claimed that this algorithm is better than common *ab initio* methods for phase-pure materials. These results have been achieved through the training of an ANN using an experimental dataset composed of pairs of RGB (red-green-blue) images of 178 994 distinct materials samples acquired using a commercial scanner and

the corresponding optical absorption spectra recorded using a UV-visible spectrometer.

Another recent piece of research work by Jha *et al* [36] has demonstrated that by using a deep learning approach it is possible to bypass the problematic step of data representation, since deep learning models can work directly on raw data without the need to manually find the descriptors (or features) that are necessary in ML algorithms. They implemented a deep neural network model, called Elemnet, that can predict material properties starting from only the elemental compositions as inputs. By using Elemnet, they scanned around 450 million candidate compositions for novel ternary and quaternary compounds and predicted that new stable compounds could be found in about 368 000 different chemical systems.

The application of ML methods to large datasets also allows the unveiling of information that is already included with the data but not apparent, this process is called data mining. Data mining methods are very helpful in materials discovery and optimization due to their good performance, speed, and simplicity, both in obtaining classification diagrams and in construction of quantitative structure-property relationship models. These methods have been used in the process of discovery of lithium battery materials and also in the development of novel thermoelectric materials [37].

The increasing amount of scientific literature is overwhelming researchers while a key strength of machine and deep learning lies in dealing with big data. Therefore, these techniques can be employed for the automatic extraction of information via a process called text mining. An remarkable application of this method has been reported by Tshitoyan *et al* [38], where they have shown that materials science knowledge present in the published literature can be efficiently encoded as information-dense word embeddings (vector representations of words) without human labelling or supervision. Without any explicit insertion of chemical knowledge, these embeddings capture complex materials science concepts such as the underlying structure of the periodic table and structure-property relationships in materials. Additionally, they have demonstrated that this unsupervised method can recommend materials for functional applications several years before their discovery. Their findings highlight the possibility of extracting knowledge and relationships from the massive body of scientific literature in a collective manner and point towards a generalized approach to the mining of scientific literature.

In the context of theoretical methods in materials science, approaches based on density functional theory (DFT) have been successful in predicting the properties of many classes of compounds, offering high accuracy at reasonable cost. However, DFT and related electronic structure techniques are limited by the exchange-correlation functional that describes non-classical interactions between electrons. Accurate universal density functionals can be learned from data extracted from structure-property databases and early examples include the Bayesian error-estimation functional and combinatorially optimized DFT functionals. Moreover, the use of ML algorithms could provide a new method beyond the standard approach to

DFT that consists of learning density-to-energy and density-to-potential maps directly from training systems without the need to solve the Kohn–Sham equations [26, 39].

3.2. Application of AI methods to small datasets

The introduction of ML methods in materials science was firstly motivated by the generation of large databases mainly by computational methods; however, methods like DFT cannot directly return information on significant functionalities such as superconductivity or magnetic properties and datasets generated computationally are usually calculated on ideal systems at thermodynamic equilibrium. Hence, experimental datasets are fundamental for real material problems that involve multi-components, solid solutions, defects and metastable compositions; but in general, experiments are expensive and time consuming, thus it is very difficult to implement experimental datasets with thousands or millions of entries. For these reasons, ML methods that can be applied to small datasets (between 10 and 100 samples) have recently been employed in materials science. These are known as adaptive (or active) learning methods [40], part of the so-called Bayesian global optimization process [41, 42]. The workflow of the methodology used for small datasets is represented in figure 4; the first step is the establishment of a training dataset that consists of a small set of experimental data on structure and corresponding properties. The second step is the development of a model based on the learning of the relationship between the material descriptors (or features that are essentially the variables used in ML) and the material property that needs to be optimized.

Then, this model is applied to a set of unexplored structures to identify the best material with properties closest to the target. Finally, this material is synthesised and results of its characterization are added to the dataset to improve the model in iterative fashion [43].

The main advantage of this methodology compared with conventional ML methods is that it enables one to start with a relatively low number of data points (between 10 and 100) and then guide successive experiments to find a material with an enhanced targeted property. This method is also called an adaptive or active learning approach since it is a guided process that reduces the number of experiments or calculations needed compared to other techniques which require the creation of big datasets generated in a random fashion. More details about this process can be found in [2].

To give some recent examples where this approach accelerated the discovery of new materials with targeted properties:

In the search for NiTi-based shape memory alloys with very small thermal hysteresis, Xue *et al* [3] used 22 initial experimental samples and found the best compound with a thermal hysteresis as small as 1.84 K on the sixth iteration (a small hysteresis is crucial for realizing shape memory alloy application). Overall, they synthesized and characterized 36 predicted compositions from a potential space of 800 000 compositions. Of these, 14 had smaller thermal hysteresis than any of the 22 in the original data set.

Yuan *et al* [44] have found and synthesized the piezoelectric composition with the largest electrostrain of 0.23% in the BaTiO₃ family. Their training dataset consisted of only 61 compounds while there are potentially 605 000 possible compositions.

This active learning approach has also been used for guiding DFT calculations towards targeted regions, one such example is from Seko *et al* [45] where they combined systematic DFT calculations and ML techniques for predicting the melting temperature of single and binary compounds. They have shown that the average number of data points required for finding the compounds with the highest melting temperature over thirty trials using the Kriging method (a specific ML algorithm) compared to using random compound selections, are 16.1 and 133.4, respectively; i.e. that ML substantially improved the efficiency of discovery.

The active learning approach represented in figure 4 can be generalized to include more than one targeted property, since in many cases more than one property is of interest. A detailed description of this method, called multiobjective optimization, can be found in [46]. Here we illustrate the basic concepts of multiobjective optimization using an example from [43]. This work is essentially an extension of the previous example on NiTi-based shape memory alloys [3] where in this case the objective is to search for a composition that at the same time minimize the thermal hysteresis and the transition temperature. They used a training set of 53 different material compositions and this optimization strategy decreases the number of measurements required to find the optimal composition by nearly 20% compared to random selection. The same strategy can be used to discover new materials with other desired properties.

3.3. Inverse design using machine learning


This last subsection describes another approach to materials discovery that is based on inverse design, whereby the aim is to discover tailored materials from the starting point of a desired functionality. So far, this approach has been applied mainly in molecular chemistry for the design of prospective drugs, synthetic routes to organic compounds, and optimization of photovoltaics and redox flow batteries [47].

In the context of solid-state materials this approach is less common, but an interesting example is found in the work of Franceschetti and Zunger [48] where they describe a theoretical method that addresses the problem of finding the atomic configuration of a complex, multi-component system having a target electronic-structure property, starting by considering a specific electronic structure and finding the corresponding material composition with this desired property. They have predicted the composition with the largest optical bandgap for the AlGaAs alloy family. In this example ML methods are not used, but considering the successful results achieved in molecular chemistry, it seems that ML could also be applied in inverse design methods for solid-state materials to accelerate the discovery of new materials with specific targeted functionalities.


4. AI in photonics and nanophotonics

So far, the application of ML methods described in the previous section has been used to develop new materials with many different types of functionality but the use of these methods in the search of new optical and photonic materials is not well established, even though there is a considerable ongoing interest in the identification of new material platforms for plasmonics and photonic metamaterial applications [49, 50]. Nevertheless, ML methods have been introduced recently in optics and nanophotonics mainly for the design and optimization of nanophotonic devices, which can be computationally demanding and time-inefficient tasks. The application of ML techniques here enables one to overcome the constraint of physical intuition and to explore the parameter space more effectively, leading to data-driven, on-demand design of novel devices [4, 51]. In this section, we review the main advances in the application of ML methods for improving the design of novel devices. Nanophotonics or nano-optics is the study of light and its interaction with matter at the (sub-wavelength) nanoscale. Conventionally, the guiding principles for the design of nanophotonic devices are provided by physical intuition based upon the study of simple systems, experience obtained from previous practice, and intuitive reasoning [4]. The optical properties of the initial design are usually computed by simulations solving the Maxwell's equations and in order to obtain the desired optical response, the initial design is usually adjusted by performing multiple simulations until the target response is reached. This is a direct design process and it becomes computationally demanding and slow as the complexity of devices increases. Another approach is based on so-called inverse design, where targeted photonic functionalities are obtained by working backwards from the desired response to the required parameters of the structure [4]. There are many different optimization techniques reported in literature for structural design problems, such as genetic algorithm, particle swarm optimization, and topology optimization [4]. The use of these optimization methods has been quite successful, for example a genetic algorithm has been used for designing binary masks that create optical superoscillations for imaging with the potential for manufacturing with light and data-storage applications [52] and a particle swarm optimization algorithm has been employed to design metal nanoparticle arrays that produce broadband plasmonic field enhancement over the entire visible spectral range [53]. Another approach for designing optical devices has been implemented by Miller [54, 55] in the mathematical framework of singular value decomposition, and by using this approach several different optical devices have been realised such as an extremely compact photonic crystal waveguide spatial mode converter which converts the fundamental even mode to a higher order odd mode with nearly 100% efficiency [56] and an optical device that can perform any linear function or coupling between inputs and outputs [57]. The recent application of deep learning (or ANNs) to nanophotonics design problems has provided significant flexibility compared with conventional optimization methods. Indeed, the advantages of ANNs over traditional optimization approaches have been

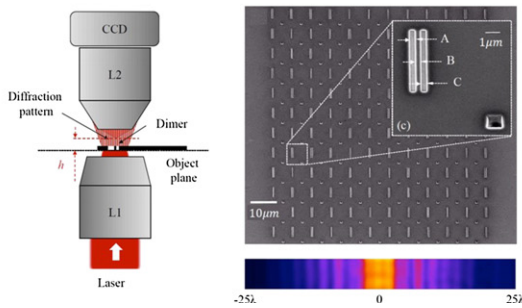
(a) Inverse design of dielectric nanoparticles from the scattering spectra



(b) Neural network modelling of laser machining



(c) AI-assisted imaging with scattered light



(d) Prediction of optical properties from a photograph




Figure 5. Artificial intelligence in photonic applications. Supervised deep learning has been applied to problems such as (a) the design of multilayer nanoparticles with prescribed light scattering spectra; (b) the prediction of the 3D surface profile obtained by a laser machining process; (c) AI-assisted imaging with scattered light; and (d) predicting the full UV (ultraviolet)—visible absorption spectra from just sample images. (a) From [5]. Adapted with permission from AAAS. (b) Adapted with permission from [64]. © The Optical Society. (c) Adapted with permission from [81]. (d) Adapted from [35] with permission of The Royal Society of Chemistry.

highlighted in a number of recent publications [58–60] showing that ANNs enable one to automate and solve design problems in a much faster way than conventional optimization methods (once the neural network is trained) [61] and, additionally, deep learning allows one to tackle the direct and inverse problem at the same time [62, 63] and can help to find complex, nonintuitive relationships between the structure and its optical response, as shown in figure 5(a): Peurifoy *et al* [5] have used an ANN to approximate light scattering by multilayer nanoparticles. Once the neural network is trained, it can simulate the optical properties of nanoparticles orders of magnitude faster than conventional simulations. Furthermore, the trained neural network can be used to solve nanophotonic inverse design problems by using back propagation, where the gradient is analytical, not numerical. A further extension of this work has been recently reported by So *et al* [66] where they include material information in the inverse design of core-shell nanoparticles.

Another recent example of the use of ANNs in nanophotonics has been reported by Ma *et al* [67] where they developed a deep learning-based model, comprising two bidirectional neural networks, capable of automatically designing and optimizing three-dimensional chiral metamaterials with strong chiroptical responses at predesignated wavelengths (figure 6(b)). The investigated chiral metamaterial has a unit cell that consists of two stacked gold split ring resonators twisted at a certain angle and separated by two spacing dielectric layers with a continuous gold reflector at the bottom. Firstly, they have solved the forward problem, namely they have trained the deep learning model in a supervised way using 25 000 data points in order

to find the relationship between the five design parameters of the metamaterial structure with its corresponding chiroptical properties. Then, they have shown that this model can also solve the inverse design problem, namely the model allows the retrieval of the geometric parameters of the metamaterial structure starting from specific targeted optical responses.

Further application of ANNs included the identification of the parameters of a complex topological insulator in order to obtain protected edge states at target frequencies [68], the computation of the dispersion relations of photonic crystals [69], and the design of the smallest photonic lens for subwavelength focusing of light compared to the current state of the art [70]. Additionally, the deep learning methodology has been used recently for the inverse design of graphene-based metamaterials with on-demand optical responses [71], it has been applied for designing photonics crystals with targeted photonics topological properties [72] and also for finding the material and the optimal geometrical properties of metasurface holograms (figure 6(d)) with high efficiency [73]. In order to achieve this, Sajedian *et al* [73] have used a double deep Q-learning network (a deep reinforcement learning method [74]) that acts like an intelligent sweep and could identify the optimal results in ~ 5.7 billion states after only 2169 steps. The optimal results were found between 23 different material types and various geometrical properties for a three-layer structure. The computed transmission efficiency was 32% for high-quality metasurface holograms; this is two times bigger than the previously reported results under the same conditions. The found structure is transmission-type, polarization-independent and works in the visible region.

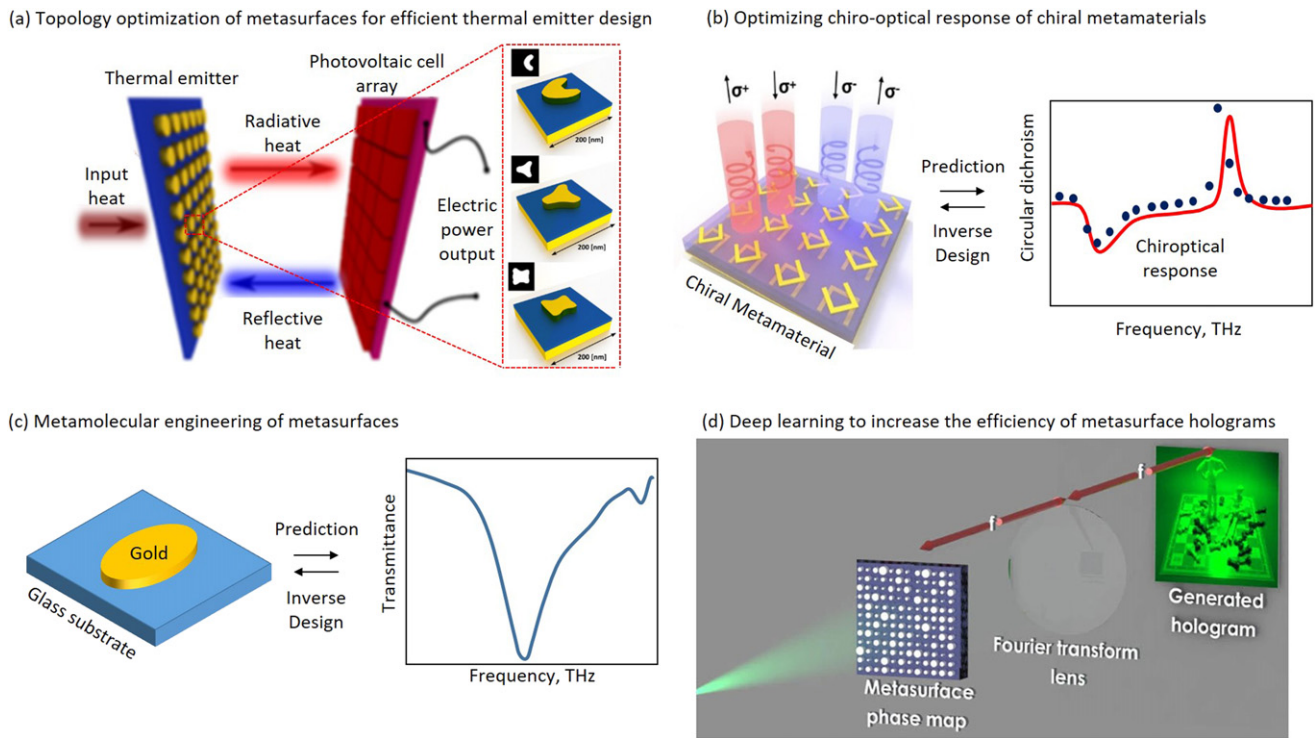


Figure 6. Artificial intelligence for optimization of metamaterials. The deep learning approach has been used for the design and optimization of metamaterials with different functionalities: (a) efficient metasurface absorbers for photovoltaic systems; (b) metamaterials with strong chiroptical responses; (c) metasurfaces with desired reflection and transmission spectra; and (d) metasurface holograms with high efficiency. (a) Adapted from [78], with the permission of AIP Publishing. (b) Adapted with permission from [67]. Copyright (2018) American Chemical Society. (c) Adapted with permission from [76]. Copyright (2018) American Chemical Society. (d) Adapted from [73]. **CC BY 3.0.**

Deep learning has also been used recently by Wilt *et al* [75] for the design of auxetic metamaterials (materials with negative Poisson's ratio): pseudorandomized images of experimental specimens and their respective computational deformation results were used to train a regressive model and predict the deviation from optimal behaviour. Subsequently, they have proposed a scalable workflow design process connecting the unique performance of auxetics to ML design.

Most of these examples are based on supervised learning, while Liu *et al* (figure 6(c)) [76] proposed a generative, deep, network model mostly based on unsupervised learning, which ensures the generation of structural patterns independent of human experience that can discover and optimize unit cell patterns of metasurfaces in response to user-defined, on-demand spectra at the input. The network architecture consists of three neural networks: the simulator, the generator and the critic. The simulator is used for solving the direct problem, namely this ANN takes the metasurface patterns as input and calculate the corresponding transmission spectra after being trained with 6500 full wave finite element simulations. The generator and the critic work together to solve the inverse problem, essentially this system can generate metasurface patterns in response to an arbitrarily input optical spectra. These two ANNs constitute a generative adversarial network (GAN) [77], a system where the two networks contest with each other to generate new data with the same statistics as the training set. In this case, the training set comprised a certain number of geometric shapes and the GAN learns to generate geometric

pattern similar to the ones of the training set. In this way, specific geometrical constraints are applied at the input that help to narrow down the potential candidates (for instance unrealistic patterns for actual nanofabrication are excluded) and thereby accelerate the convergence to a solution. As a representative case study, they considered a unit cell of a metasurface which has a single layered gold pattern in a square lattice situated on a glass substrate (figure 6(c)).

Another example of the application of unsupervised learning is the work of Kudyshev *et al* [78] in which they used this method for the design of thermal emitters with a non-trivial shape for thermophotovoltaics applications (figure 6(a)). They coupled a GAN with topology optimization showing that that the topology-optimized design can be up to 45% more efficient in comparison with a cylindrical emitter. The thermal emitters are realized with metasurfaces that consist of a three-layered structure composed of 300 nm thick TiN back reflector, 40 nm thick SiN dielectric spacer and 100 nm thick TiN top layer.

Beyond design issues, machine and deep learning can help optics and nanophotonics in other tasks, for example Luo *et al* [79] used a ML method to estimate the optical properties of black carbon fractal aggregates, where the complex morphology of black carbon aggregates makes this problem computationally very demanding for broadband applications. Another example is reported by Barth and Becker [80] where they presented a method that combines finite element simulations and clustering for the identification of photonic modes with large

local fields with specific spatial properties. Furthermore, a neural network modelling approach has been used for simulations of laser machining processes [64], wherein it has been demonstrated that a neural network can predict the 3D surface profile of the laser machined surface after exposure to different spatial intensity profiles (figure 5(b)). This has been achieved through an ANN trained on an experimental dataset composed of pairs of bitmap patterns used to spatially shape laser intensity profiles and measured height profiles of the corresponding laser machined surface. During the training, the predicted output from a particular input bitmap was compared to the experimentally measured output for the corresponding input bitmap from the training data set, and the difference (i.e. the error) was determined. This error was propagated backwards through the ANN, and the weightings of the neurons were changed accordingly until this error was minimized. The trained ANN has the predictive capability to simulate the depth profile resulting from laser machining using previously untested spatial intensity profiles.

Pu *et al* [65, 81] (figure 5(c)) have shown that the combination of ML with topologically structured light illumination of the object can revolutionize one of the most important fields in optics, superresolution microscopy. They introduced a new imaging technique termed DSSI (deeply subwavelength superoscillatory imaging) that reveals the fine structure of a physical object through its far-field scattering pattern under superoscillatory illumination, with a resolution far beyond the conventional ‘diffraction limit’ [65]. Even with plane wave illumination ML can retrieve parameters of a physical object from its scattering pattern with resolution exceeding $\lambda/20$ [81]. In a related application of AI to optical metrology, Rendón-Barraza *et al* [82] have demonstrated how the physical size of sub-wavelength objects can be determined with accuracy exceeding $\lambda/800$ via a deep learning-enabled analysis of diffraction patterns. Using a 633 nm laser source, they measure the width of sub-wavelength slits in an opaque screen with accuracy of 0.77 nm—challenging that of the electron and ion beam lithographic techniques by which such structures are typically fabricated.

5. Optics and nanophotonics for AI

With the advent of big data, hardware development has become a fundamental issue relevant to the improvement of key performance metrics such as the speed, energy consumption, and information density of computing. Traditional electronic components like the central processing units (CPUs) are not best-suited for the implementation of emerging techniques in AI like artificial neural networks, because to emulate millions of artificial neurons traditional computers must calculate for each of them in turn—a methodology that has been driving the use of graphics processing units for implementing neural networks. This has also promoted the search for new hardware architectures that are physically structured like ANNs [83].

Parallelism is easy to archive in photonics and a range of all-

optical artificial intelligence schemes and cognitive solutions have been reported in literature, such as neuromorphic optical computing, photonic reservoir computing, photonic deep neural networks, and all-optical photonic cognitive networks (e.g. optical oracles). In this section we provide a summary of these.

5.1. Neuromorphic optical computing

The concept of neuromorphic computing was firstly introduced in electronics with the aim of developing electronic analog circuits to mimic the architecture and processes present in the nervous system of animals [84, 85]. A neuromorphic optical computing (or neuromorphic photonic) system is fundamentally a network composed of connected artificial neurons (figure 7(a)). Each neuron corresponds to a node of the ANN and is capable of only three fundamental mathematical operations: vector multiplication (weighting), spatial summation (addition) and a nonlinear transformation (activation function). The input of a neuron is a linear combination (or weighted addition) of the outputs of the neurons connected to it. Then, this neuron integrates the combined signal and delivers a nonlinear response known as the activation function [84]. The peculiar characteristic of neural networks used in neuromorphic optical computing is that they are based on ‘spike processes’ mimicking physiological neurons which communicate with each other using trains of electrical pulses called action potentials or spikes.

The main difference in the working principle between neuromorphic computing and conventional computers (which are based on the so-called von Neumann architecture) is parallelization: von Neumann processors depend on a single point-to-point link between memory and CPU while a neuromorphic processor typically necessitates a large number of interconnects [84]. The implementation of neuromorphic computing using a photonic platform can provide several advantages compared to electronics in terms of energy efficiency, bandwidth and latency.

So far, the platforms used for the physical implementation of neuromorphic photonic systems are silicon photonics (the most mature technology [85]), photonic hybrids composed of III–V semiconductors (like GaAs and InP) on silicon-on-insulator substrates and quantum dot (QD) lasers grown directly on silicon substrates [85].

An ANN is characterized by two fundamental elements: nonlinear nodes (artificial neurons) and interconnections (network). In neuromorphic photonic systems the artificial neurons can be realized with optoelectronics devices such as excitable lasers or integrated modulators [85]. Excitable lasers emulate the spiking behaviour of biological neurons and are characterized by three fundamental features: the system has only one stable state at which it can indefinitely stay at rest; when excited above a certain threshold, the system emits a ‘spike’; and after that, the system decays back to rest in the course of a ‘refractory period’ during which it is temporarily less likely to emit another spike [84]. In general, excitable lasers use III–V quantum wells or QDs and these

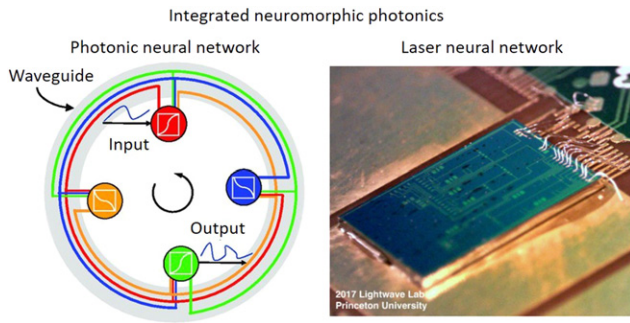


Figure 7. Neuromorphic optical computing. Schematic of a photonic neural network (left) highlighting that one waveguide can carry signals from multiple connections at the same time, solving the interconnect bottleneck; (right) picture of a laser neural network on a silicon photonic platform, capable of operating at roughly 100 million times the speed of biological neurons. Reproduced with permission from [87]. © The Optical Society.

devices include multisection lasers, ring lasers, photonic crystal nanocavities, tunnelling diodes attached to laser diodes, and semiconductor lasers with feedback [86]. An example of an integrated laser neural network is shown in figure 7 (right); this picture represents an artificial neuron capable of operating roughly 100 million times the speed of biological neurons, due to the speed of optoelectronic systems compared to biochemical interactions [87].


Tait *et al* [88] have recently demonstrated a silicon photonic modulator neuron that consists of a balanced photodetector directly connected to a microring (MRR) modulator. This device takes two optical inputs, subtracts their photocurrents electronically, and remodulates a signal onto a new wavelength. The output signal is a nonlinear function of inputs determined by the electro-optic transfer function of the modulator.

Photonic neurons are interconnected with optical waveguides forming a network; these connections must be reconfigurable and provide the weighting functionality, such that during training the weights that connect each neuron are adjusted until the neural network achieves the targeted performance. The physical implementation of these interconnections can be realized using MRR weight banks that act as reconfigurable filters and are essentially waveguides bent back on themselves to create an interference condition [84]. The resonance of these devices can be tuned thermally or electronically for the adjustment of the weights that connect each neuron. An efficient approach to use the full capacity of a waveguide is wavelength division multiplexing (WDM). WDM is particularly important for solving the interconnect bottleneck (figure 7, left), a typical problem encountered in the implementation of neural networks, since a waveguide can carry signals from multiple connections at the same time [87].

5.2. Photonic reservoir computing

Photonic reservoir computing is a concept that has emerged from neuromorphic optical computing [89]. The architecture

(a) Spatially extended reservoir



(b) Delay-based reservoir




Figure 8. Photonic reservoir computing. (a) Schematic of a spatially distributed reservoir: the input is coupled into the reservoir via a randomly connected input layer to the N nodes in the reservoir. The connections between reservoir nodes are randomly chosen and kept fixed. (b) Schematic of a delay-based reservoir: the reservoir is obtained by dividing the delay loop into N intervals and using time multiplexing. The input states are sampled and held for a duration τ , where τ is the delay in the feedback loop. Adapted from [91]. CC BY 3.0.

of a reservoir computer consists of three basic components: the input layer, the reservoir and the output layer. An input signal is inserted into the reservoir, which maps the input to a higher dimension by a reservoir. A simple readout mechanism is trained to read the state of the reservoir and map it to the desired output. The reservoir is fundamentally a neural network with randomized but fixed connections. Indeed, in photonic reservoir computing only the output connections are trained while the reservoir itself and input layer connections remain unaltered and do not need to be reconfigured individually. This is the main advantage of reservoir computing, since the limited number of connections which must be modified individually strongly aids implementations in hardware and simplifies its fabrication [89]. Furthermore, the computational concept is fully parallel, starting with multi-valued input data, continuing with the creation of high-dimensional reservoir responses until final computation of multi-valued output data. The reservoir can be implemented by a recurrent neural network (a network in which each neuron output can reach every other neuron, including itself) or a wide variety of other approaches, such as using time-delayed feedback [86].

Two different classes of photonic reservoir have been investigated: spatially extended reservoirs which consist of a network of spatially distributed nonlinear nodes (figure 8(a)) and delay-based reservoirs which consist of a single nonlinear node multiplexed in time (figure 8(b)) [90, 91].

Platforms used for the physical realization of spatially extended reservoirs include silicon photonics and diffractive optical elements. An example of spatially distributed reservoir is reported by Vandoorne *et al* [92] wherein a linear photonic network consisting of optical waveguides, splitters and combiners has been implemented using a SOI system. This generic chip can perform arbitrary Boolean logic operations with memory as well as 5-bit header recognition up to 12.5 Gbit s⁻¹, without power consumption in the reservoir, and it can also perform spoken digit recognition.

Delay-based reservoirs have been implemented using two different approaches: optoelectronic reservoirs that use laser sources, optical fibres [93] and Mach–Zehnder modulators, and all-optical reservoirs that are based on semiconductor lasers, semiconductor optical amplifiers or passive optical cavities [89]. An example of a delay-based reservoir is described by Dejonckheere *et al* [94] and it is based on a fully passive nonlinearity, namely the saturable absorption of a semiconductor mirror placed in a ring-like optical cavity. This system has been tested on benchmark tasks such as speech recognition, showing performance comparable to other photonic reservoir computers.


5.3. Photonic deep neural networks

Integrated optics is considered a promising hardware platform for implementing ML algorithms. A special interest is reserved for ANNs, since matrix-vector multiplications, which are used heavily in ANNs, can be done efficiently in photonic circuits. Moreover, the speed and energy efficiency of nanophotonic circuits can be much higher than those of their electronic counterparts [4]. Deep neural networks comprise an input and output layer linked via a certain number of hidden layers. Photonic deep neural networks differ from neuromorphic optical computing and photonic reservoir computing, because they are not ‘spiking systems’ (the main distinctive feature of neuromorphic computing) and their connections are reconfigurable while in reservoir computing they are fixed.

So far, two different approaches have been used for the physical realization of photonic networks: the first was suggested by Shen *et al* [95] and relied on nanophotonic circuits; the other proposed by Lin *et al* [96] is based on diffractive optical elements.

Shen *et al* proposed a theoretical fully optical neural network architecture where each layer of the network is composed of an optical interference unit (OIU) to perform the linear matrix multiplication and an optical nonlinear unit (ONU) that acts as the nonlinear activation (figure 9(a)) [97]. In principle, this architecture enables the realization of an ANN with an arbitrary number of layers fully in the optical domain. For the physical implementation, the OIU is realized with a programmable nanophotonic processor (PNP) and the ONU is simulated on a conventional computer. The PNP is a silicon photonic integrated circuit composed of 56 Mach–Zehnder interferometers and 213 phase shifting elements. Each interferometer is composed of two evanescent-mode waveguide couplers sandwiching an internal thermo-optic phase shifter to control the splitting ratio of the output modes, followed by

(a) Photonic deep neural network for vowel recognition



(b) Diffractive deep neural network for handwritten digits classification

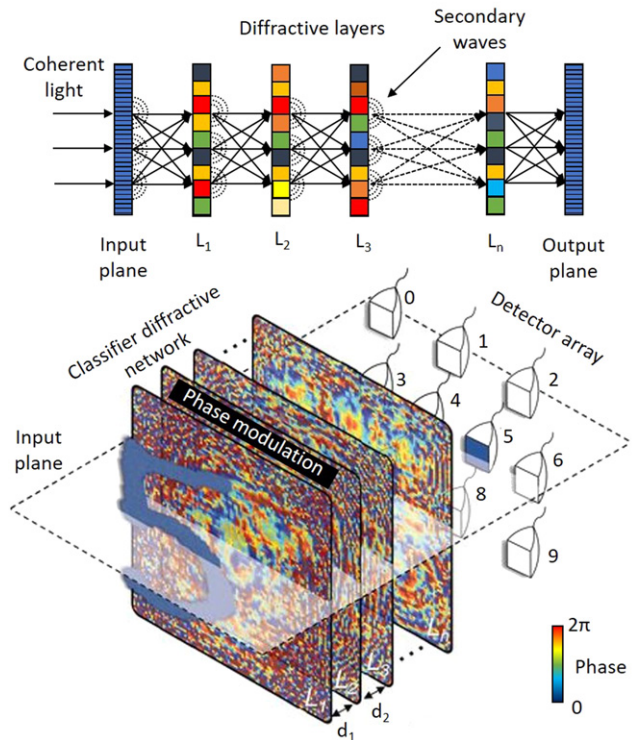


Figure 9. Photonic deep neural networks. (a) Simplified schematic, after [97], of a proposed fully optical neural network architecture where each layer includes an OIU to perform the linear matrix multiplication and an ONU that acts as the nonlinear activation. (b) A diffractive deep neural network formed by multiple layers of diffractive surfaces where each point on a given layer acts as a secondary source of waves; after the training process performed in a conventional computer, this network has been tested on classification of handwritten digits, a benchmark ML task. (a) Adapted from [97]. CC BY 3.0. (b) From [96]. Adapted with permission from AAAS.

a second modulator to control the relative phase of the output modes. By adjusting the phase shifters, it is possible to program the PNP to behave as an OIU. They experimentally demonstrated that this system is capable of vowel recognition with an accuracy comparable of that of a conventional digital computer [95].

Lin *et al* [96] introduced an all-optical deep learning framework called a diffractive deep neural network where the network is physically formed by multiple layers of diffractive surfaces (figure 9(b)). Each point on a given layer acts as a secondary source of a wave (Huygens’ principle) which has an amplitude and phase determined by the product of the input

wave and the complex valued transmission or reflection coefficient at that point. Hence, an artificial neuron in this network is connected to other neurons of the following layer through a secondary wave that is modulated in amplitude and phase by both the input interference pattern created by the previous layers and the local transmission/reflection coefficient at that point. This network is numerically simulated and trained on a conventional computer; once the optimal structure is achieved this is physically realized via 3D printing. They experimentally demonstrated that this network can perform classification of handwritten digits and fashion products with good accuracy.

Both approaches are promising for the implementation of optical neural networks, nevertheless they both suffer from two important limitations: the training process and the nonlinear activation function are not yet implemented at the hardware level. Therefore, there is considerable interest in the implementation of these functionalities. For instance, Hughes *et al* [98] introduced a method that enables highly efficient, *in situ* training of a photonic neural network. They used adjoint variable methods to derive the photonic analogue of the backpropagation algorithm, which is the standard method for computing gradients of conventional neural networks. This method works by physically propagating the adjoint field and interfering its time-reversed copy with the original field. The gradient information can then be directly measured out as an in-situ intensity measurement. As an application, they demonstrated the training of a numerically simulated photonic ANN.

Miscuglio *et al* [99] have discussed two independent approaches for implementing a nonlinear activation function for optical neural networks. Their method is based on two different nanophotonic structures: a system consisting of a single QD between a pair of gold nanoparticles integrated in a waveguide platform and a film of C₆₀. The first structure exhibited induced transparency, with the optical nonlinearity providing an extinction ratio of ~ 3 dB, while the second displayed a nonlinear optical response as a function of the impinging power density with a modulation range of approximately 7 dB. The proposed nonlinear optical responses were used as activation functions for the simulation of fully-connected neural networks. They tested these nonlinear activation functions on a standardized neural network training set, MNIST (Modified National Institute of Standards and Technology database) classifiers of handwritten digits, obtaining classification accuracies of 97% and near 100% comparable with software based trained neural networks.

Another related work has been reported by Williamson *et al* [100] where they have introduced an electro-optic hardware platform for nonlinear activation functions in optical neural networks. The optical-to-optical nonlinearity operates by converting a small portion of the input optical signal into an analog electric signal, which is used to intensity-modulate the original optical signal with no reduction in processing speed.

Their scheme allows for complete nonlinear ON–OFF contrast in transmission at relatively low optical power thresholds and eliminates the requirement of having additional optical sources between each of the layers of the network. Moreover, the activation function is reconfigurable via electrical bias, allowing it to be programmed or trained to synthesize

a variety of nonlinear responses. Using numerical simulations, they have demonstrated that this activation function significantly improves the accuracy of optical neural networks, allowing them to perform well on two benchmark ML tasks: learning a multi-input exclusive-OR (XOR) logic function and classification of images of handwritten numbers from the MNIST dataset. The addition of the nonlinear activation function improves test accuracy on the MNIST task from 85% to 94%.

The implementation of photonic deep neural networks requires integrated and parallel photonic interconnects, which correspond to the large-scale vector matrix products that are at the basis of neural network computation. However, parallel photonic waveguide circuits realized in two dimensions are strongly limited in size due to scaling constraints. To overcome this limitation, Moughames *et al* [101] have used three-dimensional (3D) printed photonic waveguides. 3D optical couplers with fractal topology efficiently connect large numbers of input and output channels and they have shown that the required substrate area and structure height scale linearly. Going beyond simple couplers, they have introduced functional circuits for discrete spatial filters identical to those used in deep convolutional neural networks.


There are also many suggested architectures for optical neural networks in literature that are not yet physically implemented, but only simulated on conventional computers. In this context, Khoram *et al* [102] have recently proposed a novel approach to realize artificial neural computing in a continuous and layer-free fashion that goes beyond the paradigm of layered feed-forward networks. In their system, computation is performed by a host material with numerous subwavelength scatterers that transforms the wavefront of an input light to realize sophisticated computing tasks such as image recognition. These computing media can be as small as tens of wavelengths and offer ultra-high computing density.

5.4. All-optical photonic cognitive networks

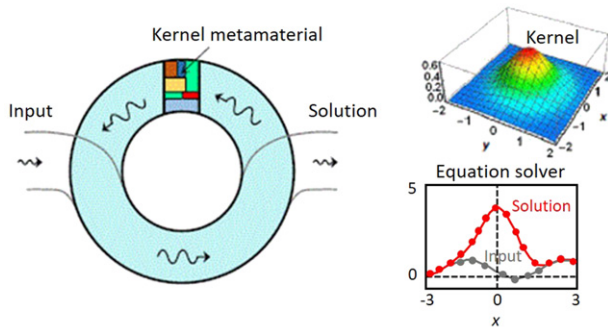
The idea of optical cognitive information processing was introduced by Ovshinsky [103], where he discussed the fact that ovonic phase change (non-volatile change of electrical and/or optical properties due to phase transition, e.g. from a crystalline phase to an amorphous phase) memories can behave in a similar way to nerve cells, which due to their plasticity have the ability to change behaviours as a result of experience. This adaptive behaviour results in learning. Indeed, a cognitive process is essentially an acquisition of knowledge and understanding of the surrounding environment through experience.

It has been demonstrated that optical networks are capable of performing cognitive processes and also solving computationally hard problems in an efficient way; for example Wu *et al* [6], using a simple fibre network, provided a proof-of-principle demonstration that this network can be treated as an optical oracle for the Hamiltonian path problem, the famous mathematical complexity problem of finding whether a set of towns can be travelled via a path in which each town is visited only once (figure 10(a)).

(a) Optical oracle for NP-complete problems



(b) Solving integral equations using inverse-designed metastructures



(c) Photonic Ising machine by spatial light modulation




Figure 10. All-optical photonic cognitive networks. Illustrations of all-optical photonic cognitive networks include: (a) an optical oracle where an optical pulse interrogates a fibre network to solve the Hamiltonian path problem, the mathematical complexity problem of finding whether a set of towns can be travelled via a path in which each town is visited only once; (b) a metamaterial platform capable of solving integral equations using monochromatic electromagnetic fields; (c) a photonic Ising machine based on a SLM used for the calculation of the ground state of an Ising Hamiltonian with thousands of spins. The spin variables are encoded in a binary phase modulation of the field and light propagation can be tailored to minimize an Ising Hamiltonian with spin couplings set by input amplitude modulation and a feedback scheme. (a) Adapted from [6]. CC BY 3.0. (b) From [7]. Adapted with permission from AAAS. (c) Adapted figure with permission from [106], Copyright (2019) by the American Physical Society.

Another example of the solution of the Hamiltonian path problem through optical networks has been reported by

Vazquez *et al* [104] using they used a femtosecond laser-written optical oracle based on cascaded directional couplers in glass.

Hu *et al* [105] employed an optical network with nonlinear waveguides for implementing the optimization algorithm of the famous ‘ant colony’ problem. Ant colonies progressively optimize a pathway to food discovered by one of the ants by marking the discovered route with volatile chemicals (pheromones) secreted on the way back from the food deposit. They have experimentally shown that photons travelling through the network behave like ants that dynamically modify the environment to find the shortest pathway to any chosen point. This demonstration illustrates how transient non-linearity in the optical system can be exploited to tackle complex optimization problems directly, on the hardware level, which may be used for self-routing of optical signals in transparent communication networks and energy flow in photonic systems.

Another recent example of optical computing has been proposed by Estakhri *et al* [7], exploring a novel aspect of metamaterial-based optical processing. Specifically, they used a metamaterial-waveguide network as an analog computing system for solving integral equations of general format (figure 10(b)). For an arbitrary wave as the input function to an equation associated with a prescribed integral operator, the solution of such an equation is generated as a complex-valued output electromagnetic field. Their approach is experimentally demonstrated at microwave frequencies through solving a generic integral equation and using a set of waveguides as the input and output to the designed metastructures. By exploiting subwavelength-scale light–matter interactions in a metamaterial platform, their wave-based, material-based analog computer may provide a route to achieve chip-scale, fast, and integrable computing elements.

An alternative optical computing approach has been proposed by Pierangeli *et al* [106] where by they experimentally demonstrated the use of spatial light modulation for calculating the ground state of an Ising Hamiltonian (figure 10(c)). The phase matrix on a spatial light modulator (SLM) acts as a lattice of spins for which the interaction is ruled by the constrained optical intensity in the far field and can be programmed by input amplitude modulation. Feedback from the detection plane allows the spatial phase distribution to evolve towards the minimum of the selected spin model. They have realized configurations with thousands of spins that settle in the ground state in a low-temperature ferromagnetic-like phase with all-to-all and tunable pairwise interactions. Their results open the route to classical and quantum photonic Ising machines that exploit light spatial degrees of freedom for parallel processing of a vast number of spins with programmable couplings and they have claimed that this computational system could provide new ultrafast hardware for ML. Moreover, they have recently demonstrated that an optimal noise level enhances the performance of spatial-photonic Ising machines on frustrated spin problems [107]. Their experimental results identify noise as a potentially valuable resource for optical computing, opening important possibilities for realizing classical and quantum annealing.

Miller [108, 109] proposed an approach that allows a broad category of optical networks to be set up progressively and automatically, including the correction of fabrication imprecision. By exploiting a new configuration algorithm, with embedded monitoring detectors, broad categories of linear optical networks can be refined, set up and/or calibrated progressively, based only on calibrated and controlled inputs. Specifically, this method can be applied generally to ‘forward-only’ networks in which the light only flows in one direction. It will work for lossless networks and those in which the loss is substantially equal on all interfering paths, as is the case for a wide range of Mach–Zehnder networks.

Among the many all-optical architectures used for the realization of AI methods, neuromorphic optical computing so far seems the most mature approach since it can be implemented on silicon photonic platforms, which are a well-established technology, and integrated systems have been already realized using optoelectronics devices such as excitable lasers and integrated modulators [85]. Photonic reservoir computing has been introduced more recently, but it appears to be a promising approach for the implementation of AI hardware, since the reservoir does not need to be trained and so hardware requirements are minimal compared with other methods such as neuromorphic optical computing which requires reconfigurable interconnections [89].

Photonic deep neural networks have been successfully implemented using nanophotonic circuits and diffractive optical elements. However, the training process and the nonlinear activation function are not yet implemented at the hardware level. Finding solutions for these challenges will be crucial for the realization of all-optical neural networks [4].

All-optical photonic cognitive networks comprise several different approaches such as optical oracles, optical analog computing and photonic Ising machines. These methods have been used to successfully tackle computationally challenging tasks like NP (nondeterministic polynomial time) problems, solving integral equations and calculation of the ground state of an Ising Hamiltonian with thousands of spins. In particular, the use of metamaterials here can be useful for the miniaturization of the hardware since metamaterials offer a highly compact and novel technological solution for controlling light at the nanoscale.

6. Conclusions

The intriguing strength of AI lies in its ability to address some problems that are intractable by other known methods. In doing so it can provide not only new or optimized solutions and predictions, but also new physical insight to the system under consideration.


Many of the world’s leading (nano)photonics research groups are now exploring applications of AI to the field and developing photonic implementations of AI. The key challenges being addressed include:

- Engineering of artificial electromagnetic materials (optimization of metamaterial and other nanostructural geometries to achieve specific, targeted optical properties);
- Optimization of (nano)photonic device, system and network architectures (e.g. adaptively reconfigurable waveguide structures to perform arbitrary input/output operations);
- Discovery and optimization of photonic materials, their synthesis and processing (identification of materials providing selected desirable properties for photonic applications);
- Development of sources of complex electromagnetic fields, e.g. phase singularities, super-oscillatory foci, etc. (design of field structures themselves and evolution of strategies for generating, characterizing and utilizing such fields).
- Development of photonic reservoir and neuromorphic optical computing, photonic deep neural networks, all-optical photonic cognitive networks and optical oracles.

Acknowledgments

The authors would like to thank Brian Hayden and Ilya Kuprov for stimulating discussions. This work is supported by The Defence Science and Technology Laboratory (Dstl), the UK’s Engineering and Physical Sciences Research Council (EPSRC Grant No. EP/M009122/1) and the MOE Singapore (Grant No. MOE2016-T3-1-006).

ORCID iDs

Davide Piccinotti  <https://orcid.org/0000-0002-2358-6417>
 Kevin F MacDonald  <https://orcid.org/0000-0002-3877-2976>
 Nikolay I Zheludev  <https://orcid.org/0000-0002-1013-6636>

References

- [1] Goodfellow I, Bengio Y and Courville A 2016 *Deep Learning* (Cambridge, MA: MIT Press)
- [2] Gubernatis JE and Lookman T 2018 Machine learning in materials design and discovery: examples from the present and suggestions for the future *Phys. Rev. Mater.* **2** 120301
- [3] Xue D Z, Balachandran P V, Hogden J, Theiler J, Xue D Q and Lookman T 2016 Accelerated search for materials with targeted properties by adaptive design *Nat. Commun.* **7** 11241
- [4] Yao K, Unni R and Zheng Y 2019 Intelligent nanophotonics: merging photonics and artificial intelligence at the nanoscale *Nanophotonics* **8** 339–66
- [5] Peurifoy J, Shen Y, Jing L, Yang Y, Cano-Renteria F, DeLacy B G, Joannopoulos J D, Tegmark M and Soljačić M 2018 Nanophotonic particle simulation and inverse design using artificial neural networks *Sci. Adv.* **4** eaar4206
- [6] Wu K, García de Abajo J, Soci C, Ping Shum P and Zheludev N I 2014 An optical fiber network oracle for NP-complete problems *Light: Sci. Appl.* **3** e147
- [7] Estakhri N M, Edwards B and Engheta N 2019 Inverse-designed metastructures that solve equations *Science* **363** 1333
- [8] Maini V and Sabri S 2019 Machine learning for humans <https://medium.com/machine-learning-for-humans>

- [9] Russell S and Norving P 2010 *Artificial Intelligence: A Modern Approach* 3rd edn (New York: Pearson)
- [10] Dutta S 2018 An overview on the evolution and adoption of deep learning applications used in the industry *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **8** e1257
- [11] LeCun Y, Bengio Y and Hinton G 2015 Deep learning *Nature* **521** 436–44
- [12] Murphy K P 2012 *Machine Learning: A Probabilistic Perspective* (Cambridge, MA: MIT Press)
- [13] Sutton R S and Barto A G 1998 *Reinforcement Learning: An Introduction* (Cambridge, MA: MIT Press)
- [14] Melnikov A A, Poulsen Nautrup H, Krenn M, Dunjko V, Tiersch M, Zeilinger A and Briegel H J 2018 Active learning machine learns to create new quantum experiments *Proc. Natl Acad. Sci. USA* **115** 1221–6
- [15] Flamini F, Hamann A, Jerbi S, Trenkwalder L M, Nautrup H P and Briegel H J 2020 Photonic architecture for reinforcement learning *New J. Phys.* **22** 045002
- [16] Zajmi L, Ahmed F Y H and Jaharadak A A 2018 Concepts, methods, and performances of particle swarm optimization, backpropagation, and neural networks *Appl. Comput. Intell. Soft Comput.* **2018** 9547212
- [17] Schmidhuber J 2015 Deep learning in neural networks: an overview *Neural Netw.* **61** 85–117
- [18] The neural network zoo posted on September 4, 2016 by Fjodor Van Veen www.asimovinstitute.org/neural-network-zoo/
- [19] Montavon G, Samek W and Müller K-R 2018 Methods for interpreting and understanding deep neural networks *Digit. Signal Process.* **73** 1–15
- [20] Samek W, Wiegand T and Müller K-R 2017 Explainable artificial intelligence: understanding, visualizing and interpreting deep learning models *ITU Journal: ICT Discoveries* **1** 1–10
- [21] Amey J L, Keeley J, Choudhury T and Kuprov I 2019 Neural network interpretation using descrambler groups (arXiv:1912.01498)
- [22] Aspuru-Guzik A, Persson K and Tribukait-Vasconcelos H 2018 *Materials Acceleration Platform—Accelerating Advanced Energy Materials Discovery by Integrating High-Throughput Methods with Artificial Intelligence* Mission Innovation: Innovation Challenge 6
- [23] Alberi K *et al* 2019 The 2019 materials by design roadmap *J. Phys. D: Appl. Phys.* **52** 013001
- [24] Conduit B D, Illston T, Baker S, Duggappa D V, Harding S, Stone H J and Conduit G J 2019 Probabilistic neural network identification of an alloy for direct laser deposition *Mater. Des.* **168** 107644
- [25] Green M L *et al* 2017 Fulfilling the promise of the materials genome initiative with high-throughput experimental methodologies *Appl. Phys. Rev.* **4** 011105
- [26] Butler K T, Davies D W, Cartwright H, Isayev O and Walsh A 2018 Machine learning for molecular and materials science *Nature* **559** 547–55
- [27] Mueller T, Kusne A G and Ramprasad R 2016 *Machine Learning in Materials Science: Recent Progress and Emerging Applications (Reviews in Computational Chemistry)* vol 29 ed A L Parrill and K B Lipkowitz (New York: Wiley) pp 186–273
- [28] Oba F and Kumagai Y 2018 Design and exploration of semiconductors from first principles: a review of recent advances *Appl. Phys. Express* **11** 060101
- [29] Jain A, Hautier G, Ong S P and Persson K 2016 New opportunities for materials informatics: resources and data mining techniques for uncovering hidden relationships *J. Mater. Res.* **31** 977–94
- [30] Takahashi K and Tanaka Y 2016 Materials informatics: a journey towards material design and synthesis *Dalton Trans.* **45** 10497–9
- [31] Takahashi K and Tanaka Y 2016 Material synthesis and design from first principle calculations and machine learning *Comput. Mater. Sci.* **112** 364–7
- [32] Curtarolo S, Hart G L W, Nardelli M B, Mingo N, Sanvito S and Levy O 2013 The high-throughput highway to computational materials design *Nat. Mater.* **12** 191
- [33] Kusne A G *et al* 2014 On-the-fly machine-learning for high-throughput experiments: search for rare-earth-free permanent magnets *Sci. Rep.* **4** 6367
- [34] Wang C, Steiner U and Sepe A 2018 Synchrotron big data science *Small* **14** 1802291
- [35] Stein H S, Guevarra D, Newhouse P F, Soedarmadji E and Gregoire J M 2019 Machine learning of optical properties of materials—predicting spectra from images and images from spectra *Chem. Sci.* **10** 47–55
- [36] Jha D *et al* 2018 ElemNet: deep learning the chemistry of materials from only elemental composition *Sci. Rep.* **8** 17593
- [37] Lu W, Xiao R, Yang J, Li H and Zhang W 2017 Data mining-aided materials discovery and optimization *J. Materomics* **3** 191–201
- [38] Tshitoyan V, Dagdelen J, Weston L, Dunn A, Rong Z, Kononova O, Persson K A, Ceder G and Jain A 2019 Unsupervised word embeddings capture latent knowledge from materials science literature *Nature* **571** 95
- [39] Mosquera M A, Fu B, Kohlstedt K L, Schatz G C and Ratner M A 2018 Wave functions, density functionals, and artificial intelligence for materials and energy research: future prospects and challenges *ACS Energy Lett.* **3** 155–62
- [40] Keane A, Forrester A and Sobester A 2008 *Engineering Design via Surrogate Modelling: A Practical Guide* (New York: Wiley)
- [41] Brochu E, Cora V M and De Freitas N 2010 A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning (arXiv:1012.2599)
- [42] Ryan E G, Drovandi C C, McGree J M and Pettitt A N 2015 A review of modern computational algorithms for Bayesian optimal design *Int. Stat. Rev.* **84** 128–54
- [43] Xue D, Xue D, Yuan R, Zhou Y, Balachandran P V, Ding X, Sun J and Lookman T 2017 An informatics approach to transformation temperatures of NiTi-based shape memory alloys *Acta Mater.* **125** 532–41
- [44] Yuan R, Liu Z, Balachandran P V, Xue D, Zhou Y, Ding X, Sun J, Xue D and Lookman T 2018 Accelerated discovery of large electrostrains in BaTiO₃-based piezoelectrics using active learning *Adv. Mater.* **30** 1702884
- [45] Seko A, Maekawa T, Tsuda K and Tanaka I 2014 Machine learning with systematic density-functional theory calculations: application to melting temperatures of single- and binary-component solids *Phys. Rev. B* **89** 054303
- [46] Gopakumar A M, Balachandran P V, Xue D Z, Gubernatis J E and Lookman T 2018 Multi-objective Optimization for materials discovery via adaptive design *Sci. Rep.* **8** 3738
- [47] Sanchez-Lengeling B and Aspuru-Guzik A 2018 Inverse molecular design using machine learning: generative models for matter engineering *Science* **361** 360–5
- [48] Franceschetti A and Zunger A 1999 The inverse band-structure problem of finding an atomic configuration with given electronic properties *Nature* **402** 60–3
- [49] Piccinotti D, Gholipour B, Yao J, MacDonald K F, Hayden B E and Zheludev N I 2019 Stoichiometric engineering of chalcogenide semiconductor alloys for nanophotonic applications *Adv. Mater.* **31** e1807083
- [50] Choudhury S M, Wang D, Chaudhuri K, DeVault C, Kildishev A V, Boltasseva A and Shalaev V M 2018 Material platforms for optical metasurfaces *Nanophotonics* **7** 959–87

- [51] Zhou J, Huang B, Yan Z and Bünzli J-C G 2019 Emerging role of machine learning in light-matter interaction *Light: Sci. Appl.* **8** 84
- [52] Rogers E T F and Zheludev N I 2013 Optical superoscillations: sub-wavelength light focusing and super-resolution imaging *J. Opt.* **15** 094008
- [53] Forestiere C, Donelli M, Walsh G F, Zeni E, Miano G and Dal Negro L 2010 Particle-swarm optimization of broadband nanoplasmonic arrays *Opt. Lett.* **35** 133–5
- [54] Miller D A B 2012 All linear optical devices are mode converters *Opt. Express* **20** 23985
- [55] Miller D A B 2013 How complicated must an optical component be? *J. Opt. Soc. Am. A* **30** 238
- [56] Liu V, Miller D A B and Fan S 2012 Ultra-compact photonic crystal waveguide spatial mode converter and its connection to the optical diode effect *Opt. Express* **20** 28388–97
- [57] Miller D A B 2013 Self-configuring universal linear optical component [Invited] *Photon. Res.* **1** 1–15
- [58] An S *et al* 2020 A freeform dielectric metasurface modeling approach based on deep neural networks (arXiv:2001.00121)
- [59] Liu Z, Zhu Z and Cai W 2020 Topological encoding method for data-driven photonics inverse design *Opt. Express* **28** 4825–35
- [60] So S, Badloe T, Noh J, Bravo-Abad J and Rho J 2020 Deep learning enabled inverse design in nanophotonics *Nanophotonics* **9** 1041
- [61] Gostimirovic D and Ye W N 2018 Automating photonic design with machine learning 2018 *IEEE 15th Int. Conf. on Group IV Photonics*
- [62] Malkiel I, Mrejen M, Nagler A, Arieli U, Wolf L and Suchowski H 2018 Plasmonic nanostructure design and characterization via deep learning *Light Sci. Appl.* **7**
- [63] Mall A, Patil A, Sethi A and Kumar A 2020 A cyclical deep learning based framework for simultaneous inverse and forward design of nanophotonic metasurfaces (arXiv:2005.12796)
- [64] Heath D J, Grant-Jacob J A, Xie Y, Mackay B S, Baker J A G, Eason R W and Mills B 2018 Machine learning for 3D simulated visualization of laser machining *Opt. Express* **26** 21574–84
- [65] Pu T, Savinov V, Yuan G, Papisimakis N and Zheludev N I 2019 Unlabelled far-field deeply subwavelength imaging (DSSI) (arXiv:1908.00946)
- [66] So S, Mun J and Rho J 2019 Simultaneous inverse design of materials and structures via deep learning: demonstration of dipole resonance engineering using core-shell nanoparticles *ACS Appl. Mater. Interfaces* **11** 24264–8
- [67] Ma W, Cheng F and Liu Y 2018 Deep-learning-enabled on-demand design of chiral metamaterials *ACS Nano* **12** 6326–34
- [68] Pillozzi L, Farrelly F A, Marcucci G and Conti C 2018 Machine learning inverse problem for topological photonics *Commun. Phys.* **1** 57
- [69] Ferreira A D S, Malheiros-Silveira G N and Hernandez-Figueroa H E 2018 Computing optical properties of photonic crystals by using multilayer perceptron and extreme learning machine *J. Lightwave Technol.* **36** 4066–73
- [70] Turduev M, Bor E, Latifoglu C, Giden I H, Hanay Y S and Kurt H 2018 Ultracompact photonic structure design for strong light confinement and coupling into nanowaveguide *J. Lightwave Technol.* **36** 2812–9
- [71] Chen Y, Zhu J, Xie Y, Feng N and Liu Q H 2019 Smart inverse design of graphene-based photonic metamaterials by an adaptive artificial neural network *Nanoscale* **11** 9749–55
- [72] Long Y, Ren J, Li Y and Chen H 2019 Inverse design of photonic topological state via machine learning *Appl. Phys. Lett.* **114** 181105
- [73] Sajedian I, Lee H and Rho J 2019 Double-deep Q-learning to increase the efficiency of metasurface holograms *Sci. Rep.* **9** 10899
- [74] Mnih V *et al* 2015 Human-level control through deep reinforcement learning *Nature* **518** 529–33
- [75] Wilt J K, Yang C and Gu G X 2020 Accelerating auxetic metamaterial design with deep learning *Adv. Eng. Mater.* **22** 1901266
- [76] Liu Z, Zhu D, Rodrigues S P, Lee K-T and Cai W 2018 Generative model for the inverse design of metasurfaces *Nano Lett.* **18** 6570–6
- [77] Goodfellow I J *et al* 2014 *Generative Adversarial Nets (La Jolla: Neural Information Processing Systems)* vol 27 ed Z Ghahramani, M Welling, C Cortes, N D Lawrence and K Q Weinberger (La Jolla, CA: Neural Information Processing Systems)
- [78] Kudyshev Z A, Kildishev A V, Shalaev V M and Boltas-seva A 2020 Machine-learning-assisted metasurface design for high-efficiency thermal emitter optimization *Appl. Phys. Rev.* **7** 021407
- [79] Luo J, Zhang Y, Wang F, Wang J and Zhang Q 2018 Applying machine learning to estimate the optical properties of black carbon fractal aggregates *J. Quant. Spectrosc. Radiat. Transfer* **215** 1–8
- [80] Barth C and Becker C 2018 Machine learning classification for field distributions of photonic modes *Commun. Phys.* **1** 58
- [81] Pu T, Ou J Y, Papisimakis N and Zheludev N I 2020 Label-free deeply subwavelength optical microscopy *Appl. Phys. Lett.* **116** 131105
- [82] Rendón-Barraza C, Aik Chan E, Yuan G, Adamo G, Pu T and Zheludev N I 2020 Optical metrology of sub-wavelength objects enabled by artificial intelligence (arXiv:2005.04905)
- [83] Zhang Q, Yu H, Barbiero M, Wang B and Gu M 2019 Artificial neural networks enabled by nanophotonics *Light Sci. Appl.* **8** 42
- [84] de Lima T F, Shastri B J, Tait A N, Nahmias M A and Prucnal P R 2017 Progress in neuromorphic photonics *Nanophotonics* **6** 577–99
- [85] de Lima T F, Peng H-T, Tait A N, Nahmias M A, Miller H B, Shastri B J and Prucnal P R 2019 Machine learning with neuromorphic photonics *J. Lightwave Technol.* **37** 1515–34
- [86] Shastri B J, Tait A N, Ferreira de Lima T, Nahmias M A, Peng H-T and Prucnal P R 2017 Principles of neuromorphic photonics (arXiv:1801.00016)
- [87] Nahmias M A, Shastri B J, Tait A N, Ferreira de Lima T and Prucnal P R 2018 Neuromorphic photonics *Opt. Photonics News* **29** 34–41
- [88] Tait A N *et al* 2019 Silicon photonic modulator neuron *Phys. Rev. Appl.* **11** 064043
- [89] Van der Sande G, Brunner D and Soriano M C 2017 Advances in photonic reservoir computing *Nanophotonics* **6** 561–576
- [90] Appeltant L *et al* 2011 Information processing using a single dynamical node as complex system *Nat. Commun.* **2** 468
- [91] Tanaka G, Yamane T, Héroux J B, Nakane R, Kanazawa N, Takeda S, Numata H, Nakano D and Hirose A 2019 Recent advances in physical reservoir computing: a review *Neural Netw.* **115** 100–23
- [92] Vandoorne K *et al* 2014 Experimental demonstration of reservoir computing on a silicon photonics chip *Nat. Commun.* **5** 3541
- [93] Estébanez I, Schwind J, Fischer I and Argyris A 2020 Accelerating photonic computing by bandwidth enhancement of a time-delay reservoir *Nanophotonics* **20200184**

- [94] Dejonckheere A, Duport F, Smerieri A, Fang L, Oudar J-L, Haelterman M and Massar S 2014 All-optical reservoir computer based on saturation of absorption *Opt. Express* **22** 10868–81
- [95] Shen Y *et al* 2017 Deep learning with coherent nanophotonic circuits *Nat. Photon.* **11** 441
- [96] Lin X, Rivenson Y, Yardimci N T, Veli M, Luo Y, Jarrahi M and Ozcan A 2018 All-optical machine learning using diffractive deep neural networks *Science* **361** 1004–8
- [97] Harris N C *et al* 2018 Linear programmable nanophotonic processors *Optica* **5** 1623–31
- [98] Hughes T W, Minkov M, Shi Y and Fan S 2018 Training of photonic neural networks through *in situ* backpropagation and gradient measurement *Optica* **5** 864–71
- [99] Miscuglio M, Mehrabian A, Hu Z, Azzam S I, George J, Kildishev A V, Pelton M and Sorger V J 2018 All-optical nonlinear activation function for photonic neural networks [Invited] *Opt. Mater. Express* **8** 3851–63
- [100] Williamson I A D, Hughes T W, Minkov M, Bartlett B, Pai S and Fan S H 2020 Reprogrammable electro-optic nonlinear activation functions for optical neural networks *IEEE J. Sel. Top. Quantum Electron.* **26** 7700412
- [101] Moughames J, Porte X, Thiel M, Ulliac G, Larger L, Jacquot M, Kadic M and Brunner D 2020 Three-dimensional waveguide interconnects for scalable integration of photonic neural networks *Optica* **7** 640–6
- [102] Khoram E, Chen A, Liu D, Ying L, Wang Q, Yuan M and Yu Z 2019 Nanophotonic media for artificial neural inference *Photon. Res.* **7** 823–7
- [103] Ovshinsky S R 2004 Optical cognitive information processing—a new field *Japan J. Appl. Phys.* **43** 4695–9
- [104] Vázquez M R, Bharadwaj V, Sotillo B, Lo S-Z A, Ramponi R, Zheludev N I, Lanzani G, Eaton S M and Soci C 2018 Optical NP problem solver on laser-written waveguide platform *Opt. Express* **26** 702–10
- [105] Hu W C, Wu K, Shum P P, Zheludev N I and Soci C 2016 All-optical implementation of the ant colony optimization algorithm *Sci. Rep.* **6** 26283
- [106] Pierangeli D, Marcucci G and Conti C 2019 Large-scale photonic ising machine by spatial light modulation *Phys. Rev. Lett.* **122** 213902
- [107] Pierangeli D, Marcucci G, Brunner D and Conti C 2020 Noise-enhanced spatial-photonic ising machine *Nanophotonics* 20200119
- [108] Miller D A B 2015 Perfect optics with imperfect components *Optica* **2** 747
- [109] Miller D A B 2017 Setting up meshes of interferometers - reversed local light interference method *Opt. Express* **25** 29233–48