

Artificial Intelligence for Photonics and Photonic Materials

Davide Piccinotti¹, Kevin F. MacDonald¹, Simon Gregory², Ian Youngs² and Nikolay I. Zheludev^{1,3}

¹ Optoelectronics Research Centre and Centre for Photonic Metamaterials, University of Southampton, Southampton, SO17 1BJ, UK
² Defence Science and Technology Laboratory, Salisbury, SP4 0JQ, UK
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

Abstract

Artificial intelligence (AI) is the most important new methodology in scientific research since adoption of quantum mechanics and it is providing exciting results in multiple 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, metamaterials, materials discovery, photonic materials, plasmonics

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 going 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 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.

In the domain of nanophotonics and metamaterials, the design and optimization of devices to control light-matter interactions and manipulate electromagnetic fields at the subwavelength scale is computationally expensive (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 (c.f. all-optical computing), with reported approaches including photonic reservoir and neuromorphic 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” 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 metastructures to perform analogue computing and solve integral equations using wave interactions.

In summary, 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, on the other hand, photonic systems offer an extremely interesting platform for the implementation of AI due to the inherent parallelism of photonics technology.

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 “Artificial Intelligence in Nanophotonics” was held recently (June 2019). 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 is the study of agents that perceive the world around them, form plans and make decisions to accomplish their goals [8]. AI currently comprises a huge number of subfields, going from the general learning and perception to specific tasks, such as playing chess, proving mathematical theorems, writing poetry, driving cars and medical diagnosis. AI is an universal methodology and so it can find application in any intellectual task [9]. A particular subfield of AI named machine learning (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]. Within this approach, 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 (namely a word), the problem is known as classification or pattern recognition, instead when y is real-valued (namely a 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 apply that learned relationship to classify completely new images (without labels) that the machine has not seen before [8].

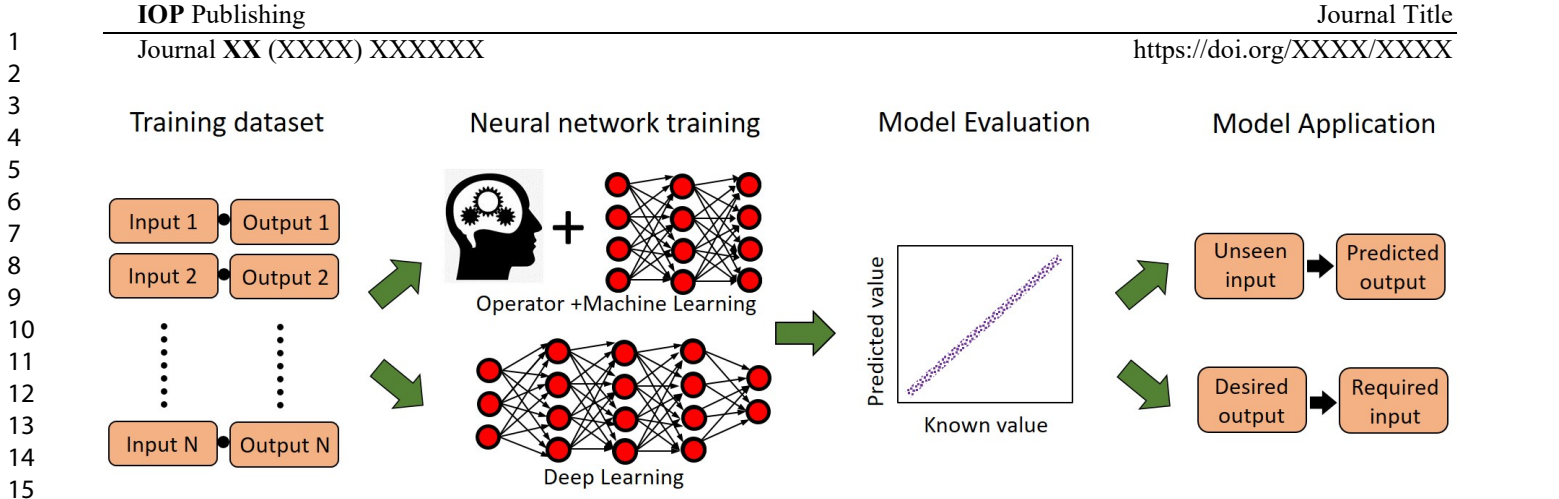



Figure 1: Supervised learning and deep learning. In supervised learning, the training dataset consists of a large number of input values (e.g. material compositions) linked to the corresponding output values (e.g. optical properties). At the training phase a model for a relationship between material 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.

Instead, the unsupervised learning approach (Figure 2) is based on unlabelled datasets, namely 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] that consists on learning useful properties of the structure of the given dataset without the guide 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 2a) and reducing dimensionality (Figure 2b) 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 more 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, namely there is a feedback loop between the learning system and its experiences [1, 13]. In this framework, the 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 promises, for instance Melnikov *et al.* [14, 15] investigated the potential of intelligent machines in the context of scientific research using a reinforcement learning methodology.

(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).

Conventional ML models have a simple two-layer structure meaning that inputs and outputs are directly connected via a certain function [12]. However, the study of the 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 a sort of “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), one of the names given to deep learning during its “history” [1]. Deep networks or ANNs comprise an input and output layer linked via a certain number of so-called hidden layers. This type of structure allows to directly insert raw data into the deep learning algorithms without the need of human-designed feature extractor that instead is needed in conventional ML. Indeed, these multiple hidden layers are able to learn directly from the data the suitable feature extractors and 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, slightly more abstract level [1]. The composition of these transformations enables the learning of very complex function. 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 refs. [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 ref. [17]. Many different architectures have been developed for ANNs; recently the Asimov Institute has implemented 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 very promising for many tasks in natural language understanding, such as topic classification, sentiment analysis, question answering, and language translation [11].

2.1 Interpretation of deep neural networks

Despite the astonishing results achieved with deep neural networks in many different fields, there is still a fundamental issue regarding how these models make their predictions, namely they are considered 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 ref. [20].

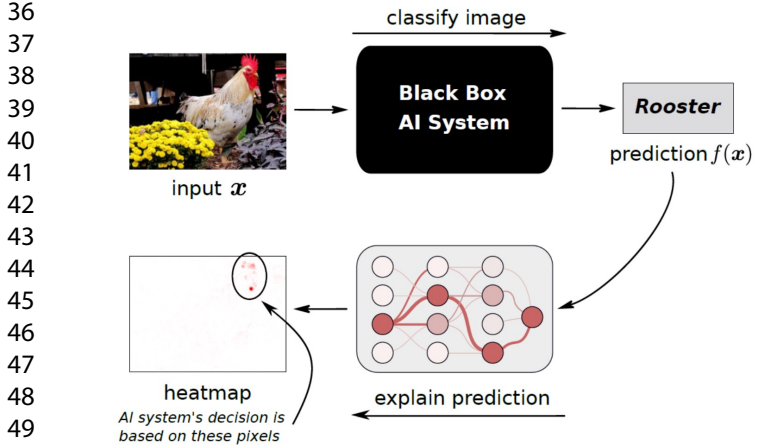


Figure 3: Explaining the predictions of an AI system. The input image is correctly classified as “rooster”. A heatmap derived, for example, via sensitivity analysis (SA) or layer-wise relevance propagation (LRP) reveals the importance of each input pixel for the prediction, helping to verify that the system works as intended [20]. In this example, the rooster’s red comb and wattle are identified as defining characteristics. Figure adapted from Samek, *et al.* [20] (International Telecommunication Union, 2017).

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. Sensitivity analysis 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 refs. [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 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 (Figure 3) have developed with the aim of accelerating materials transfer from computer simulations to lab experimentation and consumer products [23]. Acceleration of materials by design research is highlight by the almost exponential growth in the number of publications on materials design [23]. This was stimulated by the 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 the rational design of materials through the effective use of data-driven methods is the fundamental principle of the U. S. 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 in less common ML tools that allow to grow small experimentally generated databases into larger ones and along the way to predict new materials [2].

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

- 1) Data collection: it consists on the collections of scientific raw data that can be done experimentally via the measurements of certain properties of interest or extracted from first-principles calculations like DFT.
- 2) 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.
- 3) 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:
 - Naive Bayes classifiers [12]
 - k-nearest-neighbour methods [12]
 - Decision trees [12]
 - Kernel methods [12]
 - ANNs and deep neural networks [1]
- 4) 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 on refs. [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 datasets examples, the second one is on small datasets examples (Figure 3) and multiobjective optimization, and the last one 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 on large datasets

Today, the biggest databases of structural and other material information are those recently created by HT-DFT calculations [2]. These large databases (which have entries in the order of millions) 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 and due to its potential for discovering new materials and efficiently searching materials databases has become very interesting in materials science [28-30]. Nevertheless, materials informatics [31] have still some key aspects that need to be improved such as [23, 30]:

- 1) Quality of data: the management and inspection of materials databases for the presence of errors is difficult due 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].
- 2) 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 these problems prevent the creation of a global database [30].
- 3) 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 and a critical step in materials informatics [30].
- 4) Platform design: so far, these processes are realized with different software using different programming languages. Uniting all these processes under one software platform available in the Internet will benefit researchers that could add new data, extract existing data or predict unknown materials in a more consistent and systematic way [30].

Despite these critical aspects, 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 synchrotron beamline from combinatorial libraries as they are measured and rapidly cross-referenced with pre-selected entries from the Inorganic Crystal Structure Database (ICSD). They have used this approach to identify a novel magnetic phase with enhanced magnetic anisotropy which is a candidate for rare-earth free permanent magnet. More examples related to big data science generated by synchrotrons and strategies adopted nowadays 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 the full UV (ultraviolet) – visible absorption spectra from just sample images (Figure 4d). 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 very recent piece of research work done 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 find manually the descriptors (or features) that are necessary on 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 allows also to unveil 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 machine and deep learning main capability is to deal with big data. Therefore, these techniques can be employed for the automatic extraction of information in a process called text mining. An astonishing 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 on 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 on 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 return directly 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 multicomponents, solid solutions, defects and metastable compositions; but in general,

experiments are expensive and time consuming, thus 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 been recently 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

reported in Figure 3; the first step is the establishment of a training dataset that consists on 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 need to be optimized.

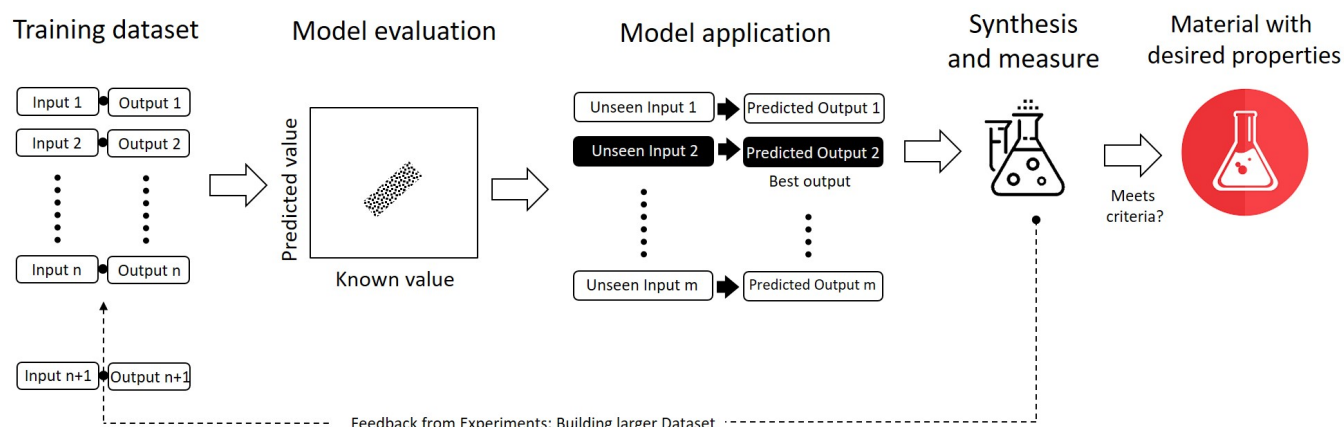


Figure 4: Material discovery based on small datasets (after ref. [43]). In search for materials with desired properties, a small set of experimental data on structure and corresponding properties is initially used for training. The model is applied to a set of unexplored structures to identify the best 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.

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 the successive experiments to find a material with an enhanced targeted property. This method is also called 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 requires the creation of big datasets generated in a “random fashion”. More details about this process can be found in ref. [2].

After a brief description of the general methodology used for working with small datasets in materials science, we present 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 alloys application). Overall, they synthesised 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.

In another example, 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 been used also for guiding DFT calculations towards targeted regions, one 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; hence ML substantially improved the efficiency of discovery.

The active learning approach represented in Figure 3 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 that is called multiobjective optimization can be found in ref. [46]. Here we illustrate the basic concepts of multiobjective optimization using an example from ref. [43]. This research 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, namely 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 the research work done by 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, namely they started by considering a specific electronic structure and find 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 be applied also 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; however, the use of these methods in the search of new optical and photonic materials is not well established even if 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, since the continuous growing in the demand performance and integration level has become computationally expensive and time inefficient. Hence, the application of ML techniques enables to overcome physical intuitions 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 nanoscale. Conventionally, the guiding principles for the design of nanophotonics devices are provided by the physical intuitions revealed by the study of simple systems, the experience obtained from previous practice, and the 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 costly and time inefficient as the complexity of nanophotonic devices rises. Another approach is based on the so-called inverse design where targeted photonic functionalities are obtained by adjusting the parameters of the structure until the desired response is attained [4]. There are many different optimization techniques reported in literature for the inverse design problem such as genetic algorithm, particle swarm optimization, topology optimization and many others [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 super-oscillations 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 D. Miller [54, 55] in the mathematic 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 the 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]. With the recent rise of AI techniques, machine and deep learning algorithms have become easily available in many research areas outside the typical computer science topics like computer vision, speech recognition and strategy making, etc. In particular, the recent application of deep learning (or ANNs) to nanophotonics design problems has provided a significant design flexibility compared with conventional optimization methods. Indeed, the advantages of ANNs over traditional optimization approaches have been highlighted in many recent publications [58-60] showing that ANNs enable 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 to tackle the direct and inverse problem at the same time [62, 63] and can help to find complex, nonintuitive relationship between the structure and its optical response as shown in Figure 4a [5].


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



(b) Neural network modelling of laser machining



(c) Deeply subwavelength superoscillatory imaging (DSSI)



(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 the design of (a) multilayer nanoparticles with prescribed light scattering spectra; (b) for the prediction of the 3D surface profile obtained by a laser machining process; (c) for a new kind of imaging technique named deeply subwavelength superoscillatory imaging; and (d) for predicting the full UV (ultraviolet) – visible absorption spectra from just sample images. Figures adapted from (a) Peurifoy *et al.* [5] (AAAS, 2018); (b) Heath, *et al.* [64] (OSA Publishing, 2018); (c) Pu, *et al.* [65]; (d) Stein, *et al.* [35] (RSC Publishing, 2019).

In this work, 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 also the 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 design and optimize three-dimensional chiral metamaterials with strong chiroptical responses at predesignated wavelengths (Figure 5b). The investigated chiral metamaterial has a unit cell that consists of two stacked gold split ring resonators (SRRs) 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 25000 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, 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 5d) 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.

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, with a mean average error below 5% for the validation set. Subsequently, they have proposed a scalable workflow design process connecting the unique performance of auxetics to machine learning design.

Most of these examples are based on supervised learning while Liu *et al.* (Figure 5c) [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 5c).

Another example of the application of unsupervised learning is the work of Kudyshev *et al.* [78] where they used this method for the design of thermal emitters with a non-trivial topology for thermophotovoltaics applications (Figure 5a). They coupled a GAN with topology optimization showing 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.

Despite designing issues, machine and deep learning can help optics and nanophotonics also in different tasks, for example Luo *et al.* [79] used a ML method to estimate the optical properties of black carbon fractal aggregates whereas the complex morphology of black carbon aggregates makes this problem computationally very expensive 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 field energies and specific spatial properties. Additionally, a neural network modelling approach has been used for the simulations of laser machining processes [64], namely they have demonstrated that a neural network can predict the 3D surface profile of the laser machined surface after exposure to different spatial intensity profiles (Figure 4b). 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.


Furthermore, Pu *et al.* [65, 81] (Figure 4c) have shown that the combination of machine learning with photonics can revolutionize one of the most important field in optics, imaging. In a theoretical article [65], they first introduced a new imaging technique termed Deeply Subwavelength Superoscillatory Imaging (DSSI) that has the potential to reveal 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" exceeding $\lambda/200$ for a dimer comprising two subwavelength opaque particles. Then, they have demonstrated this new imaging technique experimentally [81] by retrieving the parameters of a physical object from its scattering pattern with resolution exceeding $\lambda/20$. And in an 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 – approaching that of the electron

and ion beam lithographic techniques by which such structures are typically fabricated.


(a) Topology optimization of metasurfaces for efficient thermal emitter design



(b) Optimizing chiro-optical response of chiral metamaterials



(c) Metamolecular engineering of metasurfaces



(d) Deep learning to increase the efficiency of metasurface holograms




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 like (a) efficient metasurface absorbers for photovoltaic systems; (b) metamaterials with strong chiro-optical responses; (c) metasurfaces with desired reflection and transmission spectra; and (d) metasurface holograms with high efficiency. Figures adapted from (a) Kudyshev, *et al.* [78] (AIP Publishing, 2020); (b) Ma, *et al.* [67] (ACS, 2018); (c) Liu, *et al.* [76] (ACS, 2018); (d) Sajedian, *et al.* [73] (NPG, 2019).

5. Optics and Nanophotonics for AI

With the advent of big data, the hardware development has become a fundamental issue that requires the improvement of key factors such as the speed, energy consumption, and information density of computing. Traditional electronic components like the central processing units (CPUs) are not appropriate for the implementation of the emerging techniques in AI like artificial neural networks (ANNs), because traditional computers must emulate millions of artificial neurons and calculate each of them in turn. Therefore, new hardware architectures that are physically structured like artificial neural networks are more advantageous and they are urgently needed for accelerating AI and deep learning [83].

Many different all-optical implementations are 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 the many different all-optical architectures used for the realization of AI methods.

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 neural network composed of connected artificial neurons (Figure 6a). 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 that is named 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, namely 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 lasers grown directly on silicon substrates [85].

An artificial neural network 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 quantum dots and these devices include multisection lasers, ring lasers, photonic crystal nanocavities, tunnelling diode attached to laser diodes, and semiconductor lasers with feedback [86]. An example of an integrated laser neural network is shown in Figure 6 (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] has 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, namely during the 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 itself 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 6, 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].

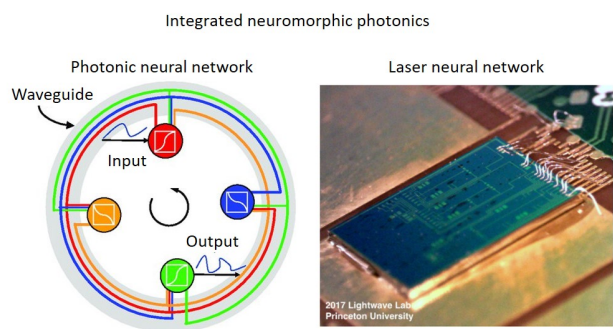


Figure 7: Neuromorphic optical computing. Schematic of a photonic neural network (left) highlighting that one waveguide can carry signals from multiple connections at the same time solving the interconnect bottleneck; (right) picture of a laser neural network on silicon photonics platform capable of operating roughly 100 million times the speed of biological neurons. Figures reproduced from Shastri, *et al.* [87] (SPIE, 2018).

5.2 Photonic reservoir computing

Photonic reservoir computing is a novel concept that has emerged from neuromorphic optical computing [89]. Fundamentally, it is a bio-inspired approach to realise a reservoir (a fixed complex system) computer in optics, where information is encoded in the intensity and phase of the optical field. Reservoir computing is a computational approach for the implementation of neural networks. The architecture of a standard reservoir consists of three basic components: the input layer, the reservoir and the output layer. Typically, an input signal is inserted into the reservoir that maps the input to a higher dimension. 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 the peculiarity of 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 mass production [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 topology in which each neuron output can reach every other neuron, including itself) or a wide variety of other systems, such as time-delayed feedback [86].

In general, there are two different class of photonic reservoirs: spatially extended reservoirs which consist of a network of spatially distributed nonlinear nodes (Figure 7a) and delay-based reservoirs which consist of a single nonlinear node multiplexed in time (Figure 7b) [90, 91].

The platforms used for the physical realization of spatially extended reservoirs are silicon photonics and diffractive imaging using diffractive optical elements. An example of spatially distributed reservoir is reported by Vandoorne *et al.* [92] where 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.

(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. Figures adapted from Tanaka, *et al.* [91] (Elsevier, 2019).

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 performances 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, 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 so-called hidden layers. Photonic deep neural networks differ from neuromorphic optical computing and photonic reservoir computing, because they are not “spiking systems” (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 one was suggested by Shen *et al.* [95] that rely on nanophotonic circuits and 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 8a). 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 (MZIs) 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 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 the one of a conventional digital computer [95].

Instead, Lin *et al.* [96] introduced an all-optical deep learning framework named diffractive deep neural network where the network is physically formed by multiple layers of diffractive surfaces (Figure 8b). 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.

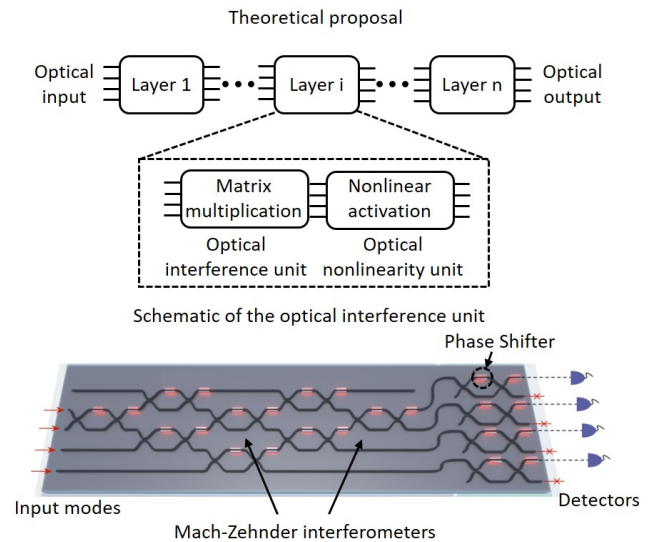
Both approaches are promising for the implementation of optical neural networks, nevertheless they both suffer from two important limitations: the training process and the nonlinear activation function are not yet implemented at the hardware level. Therefore, there is a lot of interest for the implementation of these functionality. For instance, Hughes *et al.* [97] 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 artificial neural network.

Miscuglio *et al.* [98] 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 quantum dot (QD) between a pair of gold nanoparticles integrated in a waveguide platform and a film of C₆₀. The first structure exhibited induced transparency reaching a fully nonlinearity optical modulation of the transmitted signal up to 3 dB; while the second one displayed a nonlinear optical response as 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% that are comparable with software based trained neural networks.

Another related work has been reported by Williamson *et al.* [99] where they have introduced an electro-optic hardware platform for nonlinear activation functions in optical neural networks.

(a) Photonic deep neural network for vowel recognition



(b) Diffractive deep neural network for handwritten digits classification

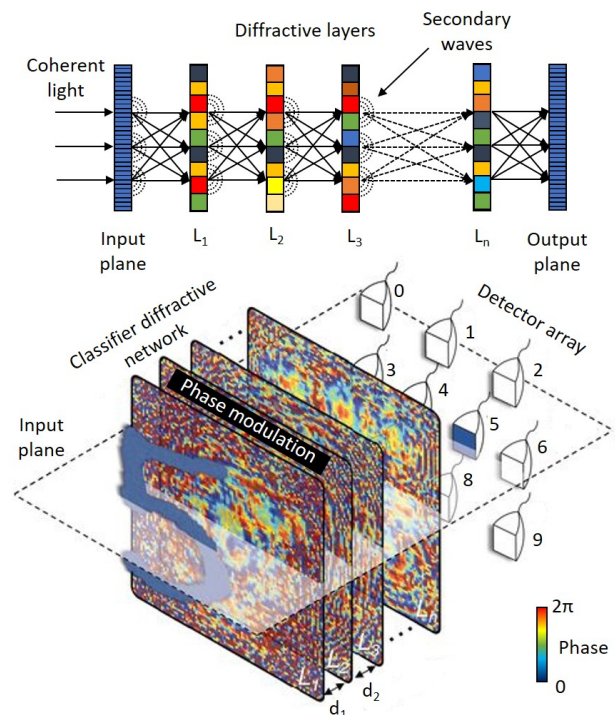


Figure 9: Photonic deep neural networks. (a) Simplified schematic, after ref. [95], of the proposed fully optical neural network architecture where each layer includes an optical interference unit (OIU) to perform the linear matrix multiplication and an optical

nonlinear unit (ONU) that acts as the nonlinear activation. The OIU is realized with a programmable nanophotonic processor (PNP) that is composed of a cascaded array of Mach-Zehnder interferometers in a silicon photonic integrated circuit capable of vowel recognition. (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 wave; after the training process performed in a conventional computer, this network has been tested on classification of handwritten digits, a benchmark machine learning task. Figure adapted from Lin, *et al.* [96] (AAAS, 2018).

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 machine learning 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.* [100] 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.* [101] 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 has been introduced by Ovshinsky [102] 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 8a).

Another example of the solution of the Hamiltonian path problem through optical networks is the one by Vazquez *et al.* [103] where they used a femtosecond laser-written optical oracle based on cascaded directional couplers in glass.


Moreover, Hu *et al.* [104] employed an optical network with nonlinear waveguides for implementing the optimization algorithm of the famous "ant colony" problem. Ant colonies progressively optimize pathway to food discovered by one of the ants through identifying the discovered route with volatile chemicals (pheromones) secreted on the way back from the food deposit. They have experimentally shown that photons traveling through the network behave like ants that dynamically modify the environment to find the shortest pathway to any chosen point in the graph. This demonstration illustrates how transient nonlinearity 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 very recent example on optical computing has been proposed by Estakhri *et al.* [7] where they explored a novel aspect of metamaterial-based optical processing. In particular, they used a metamaterial-waveguide network as analog computing system for solving integral equations of general format (Figure 8b). 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

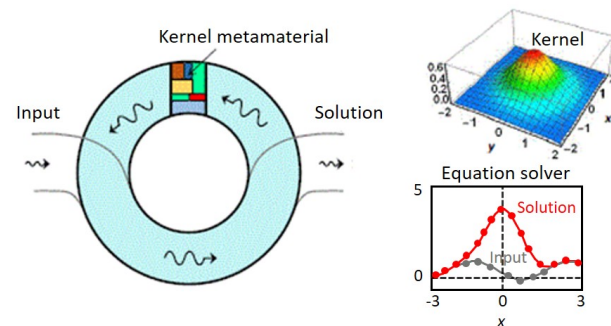
through solving a generic integral equation and using a set of waveguides as the input and output to the designed metastructures. By exploiting subwavelength-scale light-matter interactions in a metamaterial platform, their wave-based, material-based analog computer may provide a route to achieve chip-scale, fast, and integrable computing elements.

An alternative optical computing approach has been proposed by Pierangeli *et al.* [105] where they have designed and experimentally demonstrated the use of spatial light modulation for calculating the ground state of an Ising Hamiltonian (Figure 8c). 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 machine learning. Moreover, they have recently demonstrated that an optimal noise level enhances the performance of spatial-photonic Ising machines on frustrated spin problems [106]. Their experimental results identify noise as a potentially valuable resource for optical computing, opening important possibilities for realizing classical and quantum annealing.

(a) Optical oracle for NP-complete problems



(b) Solving integral equations using inverse-designed metastructures



(c) Photonic Ising machine by spatial light modulation

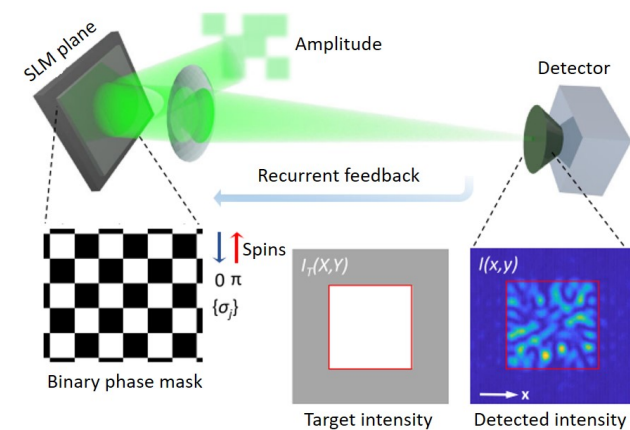


Figure 10: All-optical photonic cognitive networks. Illustrations of all-optical photonic cognitive networks include: (a) an optical oracle where an optical pulse interrogates a fibre network to solve the Hamiltonian path problem, the mathematical complexity problem of finding whether a set of towns can be travelled via a path in which each town is visited only once; optical analog computing is represented in panel b where Estakhri *et al.* [7] have introduced a metamaterial platform capable of solving integral equations using monochromatic electromagnetic fields; (c) a photonic Ising machine based on a spatial light modulator (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. Figures adapted from (a) Wu, *et al.* [6] (NPG, 2014); (b) Estakhri, *et al.* [7] (AAAS, 2019); (c) Pierangeli, *et al.* [105] (APS, 2019).

Furthermore, Miller [107, 108] 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.

In this section we have presented an overview of the many different all-optical architectures used for the realization of AI methods covering neuromorphic optical computing, photonic reservoir computing, photonic deep neural networks, and all-optical photonic cognitive networks.

So far, neuromorphic optical computing seems the most mature approach since it can be implemented on silicon photonic platforms that is 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 seems a very 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 that 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. Hence, 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 machine. 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 could 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 strength of the AI paradigm lies in its ability to address what would otherwise be intractable computational problems.

In doing so it can provide not only new or optimized solutions and predictions, but also new physical insight to the system under consideration.

This review focuses on the potential applications of AI at the intersection between optics/(nano)photonics and materials science. So far in this space, ML algorithms have typically been applied to large computationally-generated databases for the purpose of predicting new materials and unveiling trends that would not readily be revealed by conventional data mining techniques. They have also been used to grow small experimental datasets in iterative processes to predict new materials [2] and perform multiobjective optimization of properties. For example, in a search for NiTi-based shape-memory alloys with very small thermal hysteresis, Xue *et al.* [3] started with only 22 experimental samples and found the optimal compound (with a hysteresis of only 1.84 K) on the sixth iteration. In another example, Yuan *et al.* [44] has 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.

In the domain of optics and photonics, more work has so far been presented on applications of photonics to the implementation of AI (or all-optical computing in its various forms) than has been published on applications of AI to the design, development and optimization of photonic (meta)materials and devices. In regard to the latter, the design and optimization of nanostructured materials intended to harness light-matter interactions and manipulate electromagnetic fields at the subwavelength scale is computationally expensive. 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. In another example, Liu *et al.* [76] proposed a generative, deep, network model which can discover and optimize efficiently unit cell patterns of metasurfaces in response to user-defined, on-demand spectra at the input. Interestingly, this model is mostly based on unsupervised learning, which guarantees efficient generation of structural patterns independent of human experience.

In the former, reported approaches include photonic reservoir and neuromorphic 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 treated as an optical oracle for the “Hamiltonian path problem” 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.

Many of the world's leading (nano)photonics research groups are now exploring applications of AI to the field and common themes are beginning to emerge, at least in regard to broadly categorized problem types amenable to the application of AI methodologies:

- 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).

Acknowledgements

The authors would like to thank Brian Hayden and Ilya Kuprov for stimulating discussion. This work is supported by the UK's Engineering and Physical Sciences Research Council (Grant Nos. EP/M009122/1) and the MOE Singapore (Grant No. MOE2016-T3-1-006).

Data availability

No new data are presented in this review.

References

[1] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.

[2] J. E. Gubernatis and T. Lookman, "Machine learning in materials design and discovery: Examples from the present and suggestions for the future," *Phys. Rev. Mater.*, vol. 2, 12, 2018

[3] D. Z. Xue, P. V. Balachandran, J. Hogden, J. Theiler, D. Q. Xue, and T. Lookman, "Accelerated search for materials with targeted properties by adaptive design," *Nat. Commun.*, 11241, vol. 7, p. 9, 2016

[4] K. Yao, R. Unni, and Y. B. Zheng, "Intelligent nanophotonics: merging photonics and artificial intelligence at the nanoscale," *Nanophotonics*, vol. 8, 3, pp. 339-366, 2019

[5] J. Peurifoy *et al.*, "Nanophotonic particle simulation and inverse design using artificial neural networks," *Sci. Adv.*, vol. 4, 6, 2018

[6] K. Wu, J. G. de Abajo, C. Soci, P. P. Shum, and N. I. Zheludev, "An optical fiber network oracle for NP-complete problems," *Light-Sci. Appl.*, e147, vol. 3, p. 5, 2014

[7] N. M. Estakhri, B. Edwards, and N. Engheta, "Inverse-designed metastructures that solve equations," *Science*, vol. 363, 6433, p. 1333, 2019

[8] V. Maini and S. Sabri, "Machine Learning for Humans," 2019.

[9] S. Russell and P. Norving, *Artificial Intelligence: a Modern Approach*. Pearson; 3rd Ed., 2010.

[10] S. Dutta, "An overview on the evolution and adoption of deep learning applications used in the industry," *Wiley Interdisciplinary Reviews-Data Mining and Knowledge Discovery*, e1257, vol. 8, 4, 2018

[11] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, 7553, pp. 436-44, 2015

[12] K. P. Murphy, *Machine learning: a probabilistic perspective*. MIT press, 2012.

[13] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT press, 1998.

[14] A. A. Melnikov *et al.*, "Active learning machine learns to create new quantum experiments," *Proc. Natl. Acad. Sci. U.S.A.*, vol. 115, 6, pp. 1221-1226, 2018

[15] F. Flamini, A. Hamann, S. Jerbi, L. M. Trenkwalder, H. P. Nautrup, and H. J. Briegel, "Photonic architecture for reinforcement learning," *New Journal of Physics*, 045002, vol. 22, 4, 2020

[16] L. Zajmi, F. Y. H. Ahmed, and A. A. Jaharadak, "Concepts, Methods, and Performances of Particle Swarm Optimization, Backpropagation, and Neural Networks," *Appl. Comput. Intell. Soft Comput.*, 9547212, 2018

[17] J. Schmidhuber, "Deep learning in neural networks: An overview," *Neural Netw.*, vol. 61, pp. 85-117, 2015

[18] <http://www.asimovinstitute.org/neural-network-zoo/>.

[19] G. Montavon, W. Samek, and K. R. Muller, "Methods for interpreting and understanding deep neural networks," *Digit. Signal Prog.*, vol. 73, pp. 1-15, 2018

[20] W. Samek, T. Wiegand, and K.-R. Müller, "Explainable Artificial Intelligence: Understanding, Visualizing and Interpreting Deep Learning Models," *ITU Journal: ICT Discoveries - Special Issue 1 - The Impact of Artificial Intelligence (AI) on Communication Networks and Services*, vol. 1, pp. 1-10, 2017

[21] J. L. Amey, J. Keeley, T. Choudhury, and I. Kuprov, "Neural network interpretation using descrambler groups," arXiv:1912.01498, 2019

[22] A. Aspuru-Guzik, K. Persson, and H. Tribukait-Vasconcelos, "Materials Acceleration Platform—Accelerating Advanced Energy Materials Discovery by Integrating High-Throughput Methods with Artificial Intelligence," 2018.

[23] K. Alberi *et al.*, "The 2019 materials by design roadmap," *J. Phys. D-Appl. Phys.*, 013001, vol. 52, 1, p. 48, 2019

[24] B. D. Conduit *et al.*, "Probabilistic neural network identification of an alloy for direct laser deposition," *Mater. Des.*, 107644, vol. 168, p. 9, 2019

[25] M. L. Green *et al.*, "Fulfilling the promise of the materials genome initiative with high-throughput experimental methodologies," *Appl. Phys. Rev.*, vol. 4, 1, p. 011105, 2017

[26] K. T. Butler, D. W. Davies, H. Cartwright, O. Isayev, and A. Walsh, "Machine learning for molecular and materials science," *Nature*, vol. 559, 7715, pp. 547-555, 2018

[27] T. Mueller, A. G. Kusne, and R. Ramprasad, "Machine learning in Materials Science: recent progress and emerging applications," vol. 29, A. L. Parrill and K. B. Lipkowitz, Eds. (Reviews in Computational Chemistry: Wiley-Blackwell, 2016, pp. 186-273.

- [28] F. Oba and Y. Kumagai, "Design and exploration of semiconductors from first principles: A review of recent advances," *Appl. Phys. Express*, 060101, vol. 11, 6, 2018
- [29] A. Jain, G. Hautier, S. P. Ong, and K. Persson, "New opportunities for materials informatics: Resources and data mining techniques for uncovering hidden relationships," *J. Mater. Res.*, vol. 31, 8, pp. 977-994, 2016
- [30] K. Takahashi and Y. Tanaka, "Materials informatics: a journey towards material design and synthesis," *Dalton Transactions*, vol. 45, 26, pp. 10497-10499, 2016
- [31] K. Takahashi and Y. Tanaka, "Material synthesis and design from first principle calculations and machine learning," *Comput. Mater. Sci.*, vol. 112, pp. 364-367, 2016
- [32] S. Curtarolo, G. L. W. Hart, M. B. Nardelli, N. Mingo, S. Sanvito, and O. Levy, "The high-throughput highway to computational materials design," *Nat. Mater.*, vol. 12, p. 191, 2013
- [33] A. G. Kusne *et al.*, "On-the-fly machine-learning for high-throughput experiments: search for rare-earth-free permanent magnets," *Sci. Rep.*, 6367, vol. 4, p. 7, 2014
- [34] C. Wang, U. Steiner, and A. Sepe, "Synchrotron Big Data Science," *Small*, 1802291, vol. 14, 46, 2018
- [35] H. S. Stein, D. Guevarra, P. F. Newhouse, E. Soedarmadji, and J. M. Gregoire, "Machine learning of optical properties of materials - predicting spectra from images and images from spectra," *Chem. Sci.*, vol. 10, 1, pp. 47-55, 2019
- [36] D. Jha *et al.*, "ElemNet: Deep Learning the Chemistry of Materials From Only Elemental Composition," *Sci. Rep.*, 17593, vol. 8, p. 13, 2018
- [37] W. Lu, R. Xiao, J. Yang, H. Li, and W. Zhang, "Data mining-aided materials discovery and optimization," *J. Materiomics*, vol. 3, 3, pp. 191-201, 2017
- [38] V. Tshitoyan *et al.*, "Unsupervised word embeddings capture latent knowledge from materials science literature," *Nature*, vol. 571, 7763, p. 95, 2019
- [39] M. A. Mosquera, B. Fu, K. L. Kohlstedt, G. C. Schatz, and M. A. Ratner, "Wave Functions, Density Functionals, and Artificial Intelligence for Materials and Energy Research: Future Prospects and Challenges," *ACS Energy Lett.*, vol. 3, 1, pp. 155-162, 2018
- [40] A. Keane, A. Forrester, and A. Sobester, *Engineering Design via Surrogate Modelling: A Practical Guide*. American Institute of Aeronautics and Astronautics, Inc., 2008.
- [41] E. Brochu, V. M. Cora, and N. De Freitas, *A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning*. 2010.
- [42] E. G. Ryan, C. C. Drovandi, J. M. McGree, and A. N. Pettitt, "A Review of Modern Computational Algorithms for Bayesian Optimal Design," *International Statistical Review*, vol. 84, 1, pp. 128-154, 2015
- [43] D. Z. Xue *et al.*, "An informatics approach to transformation temperatures of NiTi-based shape memory alloys," *Acta Mater.*, vol. 125, pp. 532-541, 2017
- [44] R. Yuan *et al.*, "Accelerated Discovery of Large Electrostrains in BaTiO₃-Based Piezoelectrics Using Active Learning," *Adv. Mater.*, vol. 30, 7, p. 1702884, 2018
- [45] A. Seko, T. Maekawa, K. Tsuda, and I. Tanaka, "Machine learning with systematic density-functional theory calculations: Application to melting temperatures of single- and binary-component solids," *Phys. Rev. B*, vol. 89, 5, 2014
- [46] A. M. Gopakumar, P. V. Balachandran, D. Z. Xue, J. E. Gubernatis, and T. Lookman, "Multi-objective Optimization for Materials Discovery via Adaptive Design," *Sci. Rep.*, 3738, vol. 8, p. 12, 2018
- [47] B. Sanchez-Lengeling and A. Aspuru-Guzik, "Inverse molecular design using machine learning: Generative models for matter engineering," *Science*, vol. 361, 6400, pp. 360-365, 2018
- [48] A. Franceschetti and A. Zunger, "The inverse hand-structure problem of finding an atomic configuration with given electronic properties," *Nature*, vol. 402, 6757, pp. 60-63, 1999
- [49] D. Piccinotti, B. Gholipour, J. Yao, K. F. MacDonald, B. E. Hayden, and N. I. Zheludev, "Stoichiometric Engineering of Chalcogenide Semiconductor Alloys for Nanophotonic Applications," *Adv. Mater.*, vol. 31, 14, p. e1807083, 2019
- [50] S. M. Choudhury *et al.*, "Material platforms for optical metasurfaces," *Nanophotonics*, vol. 7, 6, pp. 959-987, 2018
- [51] J. Zhou, B. Huang, Z. Yan, and J.-C. G. Bünzli, "Emerging role of machine learning in light-matter interaction," *Light Sci. Appl.*, vol. 8, 1, p. 84, 2019
- [52] E. T. F. Rogers and N. I. Zheludev, "Optical super-oscillations: sub-wavelength light focusing and super-resolution imaging," *J. Opt.*, 094008, vol. 15, 9, p. 23, 2013
- [53] C. Forestiere, M. Donelli, G. F. Walsh, E. Zeni, G. Miano, and L. Dal Negro, "Particle-swarm optimization of broadband nanoplasmonic arrays," *Optics Letters*, vol. 35, 2, pp. 133-135, 2010
- [54] D. A. B. Miller, "All linear optical devices are mode converters," *Opt. Express*, vol. 20, 21, p. 23985, 2012
- [55] D. A. B. Miller, "How complicated must an optical component be?," *J. Opt. Soc. Am. A*, vol. 30, 2, p. 238, 2013
- [56] V. Liu, D. A. B. Miller, and S. H. Fan, "Ultra-compact photonic crystal waveguide spatial mode converter and its connection to the optical diode effect," *Opt. Express*, vol. 20, 27, pp. 28388-28397, 2012
- [57] D. A. B. Miller, "Self-configuring universal linear optical component," *Photonics Res.*, vol. 1, 1, pp. 1-15, 2013
- [58] S. An *et al.*, "A Freeform Dielectric Metasurface Modeling Approach Based on Deep Neural Networks," arXiv:2001.00121, 2020
- [59] Z. C. Liu, Z. M. Zhu, and W. S. Cai, "Topological encoding method for data-driven photonics inverse design," *Opt. Express*, vol. 28, 4, pp. 4825-4835, 2020
- [60] S. So, T. Badloe, J. Noh, J. Rho, and J. Bravo-Abad, "Deep learning enabled inverse design in nanophotonics," *Nanophotonics*, vol. 9, 5, p. 1041, 2020
- [61] D. Gostimirovic and W. N. Ye, "Automating Photonic Design with Machine Learning," presented at the 2018 IEEE 15th International Conference on Group IV Photonics, 2018.
- [62] I. Malkiel, M. Mrejen, A. Nagler, U. Arieli, L. Wolf, and H. Suchowski, "Plasmonic nanostructure design and characterization via Deep Learning," *Light Sci. Appl.*, vol. 7, 1, 2018
- [63] A. Mall, A. Patil, A. Sethi, and A. Kumar, "A Cyclical Deep Learning Based Framework For Simultaneous Inverse and Forward design of Nanophotonic Metasurfaces," arXiv:2005.12796, 2020
- [64] D. J. Heath *et al.*, "Machine learning for 3D simulated visualization of laser machining," *Opt. Express*, vol. 26, 17, pp. 21574-21584, 2018
- [65] T. Pu, V. Savinov, G. Yuan, N. Papasimakis, and N. I. Zheludev, "Unlabelled Far-field Deeply Subwavelength Superoscillatory Imaging (DSSI)," arXiv, 2019
- [66] S. So, J. Mun, and J. Rho, "Simultaneous Inverse Design of Materials and Structures via Deep Learning: Demonstration of Dipole Resonance Engineering Using Core-Shell

- Nanoparticles," *ACS Appl. Mater. Interfaces*, vol. 11, 27, pp. 24264-24268, 2019
- [67] W. Ma, F. Cheng, and Y. Liu, "Deep-Learning-Enabled On-Demand Design of Chiral Metamaterials," *ACS Nano*, 2018
- [68] L. Piloizzi, F. A. Farrelly, G. Marcucci, and C. Conti, "Machine learning inverse problem for topological photonics," *Commun. Phys.*, vol. 1, 1, p. 57, 2018
- [69] A. D. Silva Ferreira, G. N. Malheiros-Silveira, and H. E. Hernandez-Figueroa, "Computing Optical Properties of Photonic Crystals by Using Multilayer Perceptron and Extreme Learning Machine," *J. Light. Technol.*, vol. 36, 18, pp. 4066-4073, 2018
- [70] M. Turdudiev, E. Bor, C. Latifoglu, I. H. Giden, Y. S. Hanay, and H. Kurt, "Ultracompact Photonic Structure Design for Strong Light Confinement and Coupling Into Nanowaveguide," *J. Light. Technol.*, vol. 36, 14, pp. 2812-2819, 2018
- [71] Y. S. Chen, J. F. Zhu, Y. N. Xie, N. X. Feng, and Q. H. Liu, "Smart inverse design of graphene-based photonic metamaterials by an adaptive artificial neural network," *Nanoscale*, vol. 11, 19, pp. 9749-9755, 2019
- [72] Y. Long, J. Ren, Y. H. Li, and H. Chen, "Inverse design of photonic topological state via machine learning," *Appl. Phys. Lett.*, 181105, vol. 114, 18, p. 5, 2019
- [73] I. Sajedian, H. Lee, and J. Rho, "Double-deep Q-learning to increase the efficiency of metasurface holograms," *Sci. Rep.*, 10899, vol. 9, 2019
- [74] V. Mnih *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, 7540, pp. 529-533, 2015
- [75] J. K. Wilt, C. Yang, and G. X. Gu, "Accelerating Auxetic Metamaterial Design with Deep Learning," *Advanced Engineering Materials*, vol. 22, 5, p. 1901266, 2020
- [76] Z. C. Liu, D. Y. Zhu, S. P. Rodrigues, K. T. Lee, and W. S. Cai, "Generative Model for the Inverse Design of Metasurfaces," *Nano Lett.*, vol. 18, 10, pp. 6570-6576, 2018
- [77] I. J. Goodfellow *et al.*, "Generative Adversarial Nets," vol. 27, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds. La Jolla: Neural Information Processing Systems (Nips), 2014.
- [78] Z. A. Kudyshev, A. V. Kildishev, V. M. Shalaev, and A. Boltasseva, "Machine-learning-assisted metasurface design for high-efficiency thermal emitter optimization," *Appl. Phys. Rev.*, 021407, vol. 7, 2, p. 10, 2020
- [79] J. Luo, Y. Zhang, F. Wang, J. Wang, and Q. Zhang, "Applying machine learning to estimate the optical properties of black carbon fractal aggregates," *J. Quant. Spectrosc. Radiat. Transf.*, vol. 215, pp. 1-8, 2018
- [80] C. Barth and C. Becker, "Machine learning classification for field distributions of photonic modes," *Commun. Phys.*, 58, vol. 1, 2018
- [81] T. Pu, J. Y. Ou, N. Papasimakis, and N. I. Zheludev, "Label-free deeply subwavelength optical microscopy," *Appl. Phys. Lett.*, 131105, vol. 116, 13, p. 4, 2020
- [82] C. Rendón-Barraza, E. Aik Chan, G. Yuan, G. Adamo, T. Pu, and N. I. Zheludev, "Optical Metrology of Sub-Wavelength Objects Enabled by Artificial Intelligence," arXiv:2005.04905, 2020
- [83] Q. Zhang, H. Yu, M. Barbiero, B. Wang, and M. Gu, "Artificial neural networks enabled by nanophotonics," *Light Sci. Appl.*, vol. 8, 1, p. 42, 2019
- [84] T. Ferreira de Lima, B. J. Shastri, A. N. Tait, M. A. Nahmias, and P. R. Prucnal, "Progress in neuromorphic photonics," *Nanophotonics*, vol. 6, 3, 2017
- [85] T. F. de Lima *et al.*, "Machine Learning With Neuromorphic Photonics," *J. Light. Technol.*, vol. 37, 5, pp. 1515-1534, 2019
- [86] B. J. Shastri, A. N. Tait, T. Ferreira de Lima, M. A. Nahmias, H.-T. Peng, and P. R. Prucnal, "Principles of Neuromorphic Photonics," *arXiv*, 2017
- [87] B. J. Shastri, M. A. Nahmias, A. N. Tait, T. F. de Lima, H. T. Peng, and P. R. Prucnal, "Integrated Neuromorphic Photonics," in *Active Photonic Platforms X (SPIE)*, San Diego, 2018, vol. 10721, p. 107211M, 2018.
- [88] A. N. Tait *et al.*, "Silicon Photonic Modulator Neuron," *Phys. Rev. Appl.*, 064043, vol. 11, 6, 2019
- [89] G. Van der Sande, D. Brunner, and M. C. Soriano, "Advances in photonic reservoir computing," *Nanophotonics*, vol. 6, 3, 2017
- [90] L. Appeltant *et al.*, "Information processing using a single dynamical node as complex system," *Nat. Commun.*, 468, vol. 2, p. 6, 2011
- [91] G. Tanaka *et al.*, "Recent advances in physical reservoir computing: A review," *Neural Netw.*, vol. 115, pp. 100-123, 2019
- [92] K. Vandoorne *et al.*, "Experimental demonstration of reservoir computing on a silicon photonics chip," *Nat. Commun.*, 3541, vol. 5, p. 6, 2014
- [93] I. Estébanez, J. Schwind, I. Fischer, and A. Argyris, "Accelerating photonic computing by bandwidth enhancement of a time-delay reservoir," *Nanophotonics*, 20200184, (published online ahead of print), 2020
- [94] A. Dejonckheere *et al.*, "All-optical reservoir computer based on saturation of absorption," *Opt. Express*, vol. 22, 9, pp. 10868-10881, 2014
- [95] Y. Shen *et al.*, "Deep learning with coherent nanophotonic circuits," *Nat. Photonics*, vol. 11, p. 441, 2017
- [96] X. Lin *et al.*, "All-optical machine learning using diffractive deep neural networks," *Science*, vol. 361, 6406, pp. 1004-1008, 2018
- [97] T. W. Hughes, M. Minkov, Y. Shi, and S. H. Fan, "Training of photonic neural networks through in situ backpropagation and gradient measurement," *Optica*, vol. 5, 7, pp. 864-871, 2018
- [98] M. Miscuglio *et al.*, "All-optical nonlinear activation function for photonic neural networks," *Opt. Mater. Express*, vol. 8, 12, pp. 3851-3863, 2018
- [99] I. A. D. Williamson, T. W. Hughes, M. Minkov, B. Bartlett, S. Pai, and S. H. Fan, "Reprogrammable Electro-Optic Nonlinear Activation Functions for Optical Neural Networks," *IEEE J. Sel. Top. Quantum Electron.*, 7700412, vol. 26, 1, p. 12, 2020
- [100] J. Moughames *et al.*, "Three-dimensional waveguide interconnects for scalable integration of photonic neural networks," *Optica*, vol. 7, 6, pp. 640-646, 2020
- [101] E. Khoram *et al.*, "Nanophotonic media for artificial neural inference," *Photonics Res.*, vol. 7, 8, pp. 823-827, 2019
- [102] S. R. Ovshinsky, "Optical cognitive information processing - A new field," *Jpn. J. Appl. Phys. Part 1 - Regul. Pap. Short Notes Rev. Pap.*, vol. 43, 7B, pp. 4695-4699, 2004
- [103] M. R. Vazquez *et al.*, "Optical NP problem solver on laser-written waveguide platform," *Opt. Express*, vol. 26, 2, pp. 702-710, 2018
- [104] W. C. Hu, K. Wu, P. P. Shum, N. I. Zheludev, and C. Soci, "All-Optical Implementation of the Ant Colony Optimization Algorithm," *Sci. Rep.*, 26283, vol. 6, p. 7, 2016
- [105] D. Pierangeli, G. Marcucci, and C. Conti, "Large-Scale Photonic Ising Machine by Spatial Light Modulation," *Phys. Rev. Lett.*, 213902, vol. 122, 21, p. 6, 2019

- [106] D. Pierangeli, G. Marcucci, D. Brunner, and C. Conti, "Noise-enhanced spatial-photonic Ising machine," *Nanophotonics*, 20200119, (published online ahead of print), 2020
- [107] D. A. B. Miller, "Perfect optics with imperfect components," *Optica*, vol. 2, 8, p. 747, 2015
- [108] D. A. B. Miller, "Setting up meshes of interferometers - reversed local light interference method," *Opt. Express*, vol. 25, 23, pp. 29233-29248, 2017