



Advanced machine-learning techniques in drug discovery

Moe Elbadawi¹, Simon Gaisford^{1,2} and Abdul W. Basit^{1,2}

¹ Department of Pharmaceutics, UCL School of Pharmacy, University College London, 29-39 Brunswick Square, London, WC1N 1AX, UK

² FabRx Ltd, 3 Romney Road, Ashford, TN24 0RW, UK



The popularity of machine learning (ML) across drug discovery continues to grow, yielding impressive results. As their use increases, so do their limitations become apparent. Such limitations include their need for big data, sparsity in data, and their lack of interpretability. It has also become apparent that the techniques are not truly autonomous, requiring retraining even post deployment. In this review, we detail the use of advanced techniques to circumvent these challenges, with examples drawn from drug discovery and allied disciplines. In addition, we present emerging techniques and their potential role in drug discovery. The techniques presented herein are anticipated to expand the applicability of ML in drug discovery.

Introduction

The application of ML applied in the field of drug discovery continues to grow, facilitating research in numerous avenues. The success of ML is demonstrated by the increasing number of pharmaceutical companies in which ML is central to their business model (Table 1). In addition, ML has also been explored by large pharmaceutical companies for drug discovery [1–6]. Such success is a testament to the necessity and utility of ML for drug discovery, and an unambiguous indication that drug discovery will be intrinsically tied with ML. The goal is to reduce the resource- and labour-intensiveness of drug discovery, primarily the high-throughput screening (HTS) technique. Another aim of ML is to obviate the need for animal testing, which has received negative publicity of late.

The success of ML lies in its ability to discern patterns in complex and large volume data sets [7]. In addition, ML techniques (MLT) can be developed using common programming languages, including Python and R, which are accessible to most researchers. Furthermore, there are third-party software that provide access to ML techniques for researchers unfamiliar with coding, such as Apple's Create ML. Despite their simplicity,

third-party software are limited in their capacity to perform ML techniques, as well as other aspects of the ML pipeline.

Conventional MLTs have been thoroughly explored in drug discovery [8–10]. Such techniques include both supervised and unsupervised MLTs, including k-Nearest Neighbour (kNN), decision tree, random forest, support vector machines (SVM), artificial neural networks (ANN), principal component analysis (PCA), and k-means. Their appeal stems from their simplicity, computationally undemanding, yet improved prediction accuracy compared with traditional predictive algorithms [11]. Equally, the underlying mechanisms for conventional techniques can be cognitively comprehended by noncomputer scientist researchers. For example, for kNN, the user has one parameter to control, the *k* value, which in turn determines the classification search space based on a plurality vote. Another example is SVM, which delineates categories using a hyperplane in conjunction with support vectors to maximise the distance between the different categories. SVM benefits from using the kernel trick, which allows for nonlinear mapping of the data, which has been widely used for nonlinear data sets [12]. The technique is also available for PCA (kernel PCA; kPCA) [12]. A recent study found that kPCA can be used to improve the classification of linear models, with comparable performance to nonlinear models, although at a significantly faster rate [13].

Corresponding author: Basit, A.W. (a.basit@ucl.ac.uk)

TABLE 1

Examples of pharmaceutical companies in which ML is central to their business model

Company	Application	URL
AIQ Solutions	Patient heterogeneity	www.aiq-solutions.com/
Atomwise	<i>De novo</i> drug	www.atomwise.com/
Benevolent AI	<i>De novo</i> drug	www.benevolent.com/
Bioxcel Therapeutics	Clinical trials	www.bioxceltherapeutics.com/
BullFrog AI	Identifying niche patient population	www.bullfrogai.com/
CytoReason	Target discovery	www.cytoreason.com/
DeepCure	<i>De novo</i> small molecules	www.deepcure.ai/
EVQLV	Biologics discovery	https://evqlv.com/
Genesis Therapeutics	Therapeutic target	www.genesistherapeutics.ai/
Genome Biologics	Drug repurposing	https://genomebiologics.com/
Genomenon	Drug discovery and clinical trials	www.genomenon.com/
HealX	Therapeutic target	https://healx.io/
InSilico Medicine	Therapeutic target	https://insilico.com/
InveniAI	Disease indication	www.inveniai.com/
Kintai Therapeutics	<i>De novo</i> compounds	www.kintaitx.com/
MabSilico	Biomarker discovery	www.mabsilico.com/
MIMS	General purpose	www.mims.ai/
Pepticom	Therapeutic targets	https://pepticom.com/
Recursion	<i>De novo</i> drug	www.recursionpharma.com/
Reveal Biosciences	Pathology	www.revealbio.com/
Socium Inc	Biomarker identification	www.socium.co.jp/en/about.html
Standigm	<i>De novo</i> drug	www.standigm.com/
Thinkcyte	General purpose	https://thinkcyte.com/
Trials.ai	Designing clinical trials	www.trials.ai/
TWOXAR	Small-molecule discovery and development	www.twoxar.com/

^aSource accession date: August 29, 2020.

Despite their simplicity, conventional MLTs have their drawbacks. kNN suffers from the curse of dimensionality, wherein, at high dimensional space, the predictive performance begins to weaken [14]. Similarly, the performance of SVM begins to degrade when the number of dimensions is greater than the sample size [15]. Increasing the number of trees in random forest improves the predictive accuracy, although a large number of tree results produces an algorithm inefficient for real-time monitoring [16,17]. However, there are two chief criticisms of MLT, which are their demand for big data and lack of transparency. Addressing these limitations is required given that the collection of data can be challenging, costly, and time-consuming. In addition, transparency might facilitate the user's understanding of the discovery process and minimise their reliance on ML to understand the process. Another limitation with conventional MLTs is their lack of autonomy. For example, supervised learning requires labelling of the target variable (i.e., the variable to be predicted). In addition, once deployed, for example as a web-based software, it will require post-production maintenance, particularly as the data set evolves. To address these limitations, new techniques have been adopted by research communities and with promising results. It is anticipated that these advanced techniques will further expand the application of ML. Ultimately, the goal is to achieve artificial intelligence (AI) in the drug discovery pipeline [18]. AI is a broad branch in computer science that seeks to create human intelligence using machines, of which ML is central to achieving this goal. In recent years, a subset of ML, deep learning, as emerged as a technique capable of achieving high accuracies from big data, while handling both structured and unstructured data.

As mentioned earlier, ML in drug discovery continues to grow. This growth is accompanied by suitable reviews discussing the fundamentals and the application of conventional MLTs [8], and deep learning [19]. There is also a recent review of natural language processing, a field that is gaining attention in drug discovery [20]. Here, we focus on advanced techniques that have not received sufficient attention, albeit that have strong potential to advance the field. We prioritise examples used in drug discovery, although, if not available, we draw examples from allied fields. The reviewed techniques include reinforcement learning (RL), transfer learning, and multitask learning. In their well-received review centred on ML for drug discovery, Lo *et al.* remarked that techniques with increased visibility, as well methods for preventing overfitting, warrant further development [8]. We address their remark by describing Bayesian neural networks (BNNs) and explainable algorithms. We also detail the emergence of hybrid quantum-ML and recommender systems.

Advanced machine-learning techniques

Some of the criticisms of MLTs include the need for large data sets and for human intervention. From these remarks, advanced techniques were investigated to address the shortcomings of conventional MLT, and thereby further widen their applicability. These advanced techniques include RL, which bridges the gap toward self-autonomous learning techniques; transfer learning, and multi-task learning for developing predictive models where big data are lacking. Here, we provide an overview of these advanced techniques and illustrate examples of their application in drug discovery where possible. A summary of the techniques are tabulated in Table 2.

TABLE 2

Overview of advanced techniques described in this review

Advanced MLT	Summary	Refs
RL	Continuous learning through a reward and penalty system	[32–35,37,39]
Transfer learning	Making predictions on a data set using knowledge developed from another, larger data set	[46,54–58]
Multitask learning	Simultaneous learning of multiple tasks	[67–71]
Active learning	Semiautomated learning using human feedback	[72,75–77]
Generative models	ML for generating new data; can be used for <i>de novo</i> , data augmentation, and dimensionality reduction	[59,60,79–81]
BNNs	Using Bayesian probability distribution for neural network weights and biases; can handle small data sets	[84–86]

Reinforcement learning

RL is an exhilarating subcategory of ML that is sparking interest across both academia and industry. It has been around since the 1950s and its recent rise in popularity was sparked when RL models were victorious in a game of *Go* against professional human opponents, where no algorithm before was able to achieve this remarkable feat. The game *Go* is one of the world's oldest continuously played games [21], and is used as a benchmark for AI because the number of possible configurations in the game is thought to be 250^{150} [22]. This far exceeds both the number of proteins in the human body and the number of protons in the universe [23].

RL distinguishes itself from supervised and unsupervised learning in that it is a form of continuous learning while being autonomous. This is because RL algorithms produce judgements, whereas most supervised and unsupervised algorithms make predictions. This ability of RL to rapidly respond to dynamic environments is why it is being used for gaming, robotics, and trading in the finance sector [24]. Indeed, there are applications where RL outperformed classification tasks compared with supervised learning [25], but it is the ability of RL to continuously learn with minimal human interference that is desirable [26].

The concept of RL draws inspiration from the reward mechanism found in animals [27]. In RL, the system is not presented with examples of desired strategies. Rather, RL empirically learns the optimal decision to take through receiving reinforcement signals from its environment. The main components of RL are an agent, environment, state, policy, and reward function [28]. An agent is trained by interacting with the environment, which can have multiple states (i.e., scenarios). The agent will select an action for a given state and will receive either a positive or a negative (i.e., penalise) reward. The agent will continue taking actions for each of the different states while looking to increase the cumulative reward it receives. The reward is a mathematical formula and is defined by the user with a specific goal in mind [29]. Using gaming as an example, the agent's goal, or policy, is to win the game and it will receive +1 for when it does, and –1 for when it loses. In the case of financial trading, the policy can be to maximise profits and, hence, the agent will be rewarded for taking the series of actions that result in maximising the profit [30]. There are multiple versions of the reward function [31].

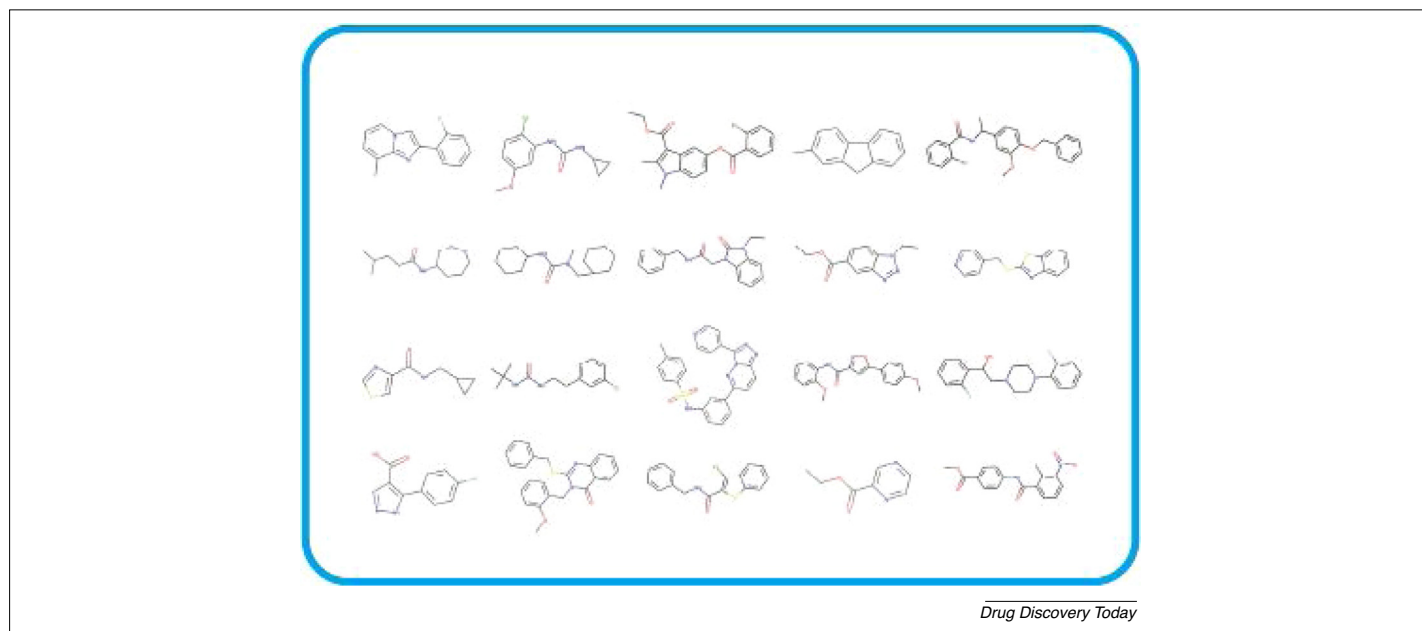
Contemporary RL has centred on *de novo* molecule designs [32–35] or molecule optimisation [36]. A noteworthy study that combined both aspects was conducted by Popova *et al.* for the *de novo* design of drugs (Fig. 1a) [37]. With this approach, RL was combined with two deep-learning techniques. One technique, the generative model, acted as the agent and generated ostensibly

chemically feasible molecules. The other technique, the predictive model, acted as the critic, whereby it rewarded or penalised the generative model for every generated molecule. Using this approach, the researchers used ~1.5 million structures from the ChEMBL21 database to train the generative model based on their SMILES strings. The results were that 1 million compounds were generated, from which 95% were confirmed to be feasible using the structure checker from ChemAxon. Moreover, they discovered that ~32 000 molecules of *de novo*-generated structures existed in a separate database (ZINC). The study went further and demonstrated that novel compounds optimised for desirable physical properties, chemical complexity, or biological activity were attainable via deep RL. Although the study demonstrated that RL can be exploited to generate new compounds, further work is needed to refine the model. For example, the strategy adopted might not guarantee drug-specific compounds [38]. Moreover, the study used SMILES, which, despite being a simple and elegant representation of compounds, issues have been raised with its use in generative models [32].

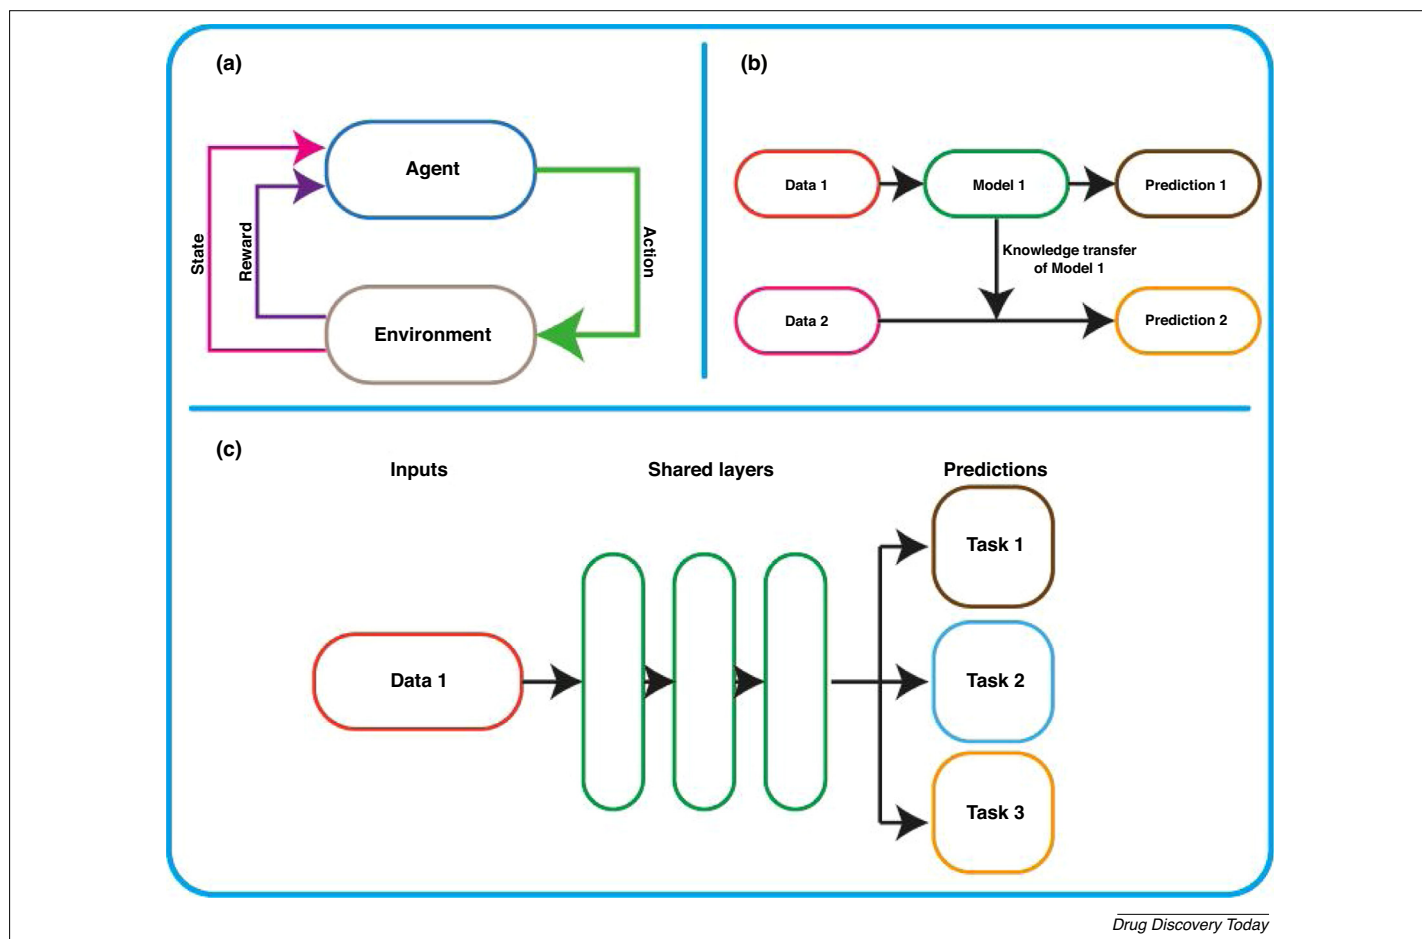
In a separate study, Zhavoronkov *et al.* developed a model for *de novo* for specific compounds: DDR1 kinase inhibitors (Fig. 1b) [39]. Their aim was to demonstrate the effectiveness of RL for rapidly identifying potent compounds, thereby demonstrating that RL can address important drawbacks of drug development, namely the slow development phase and drug selectivity. In just 46 days, the authors were able to design, synthesise, and perform both *in vitro* and *in vivo* tests. However, one of the generated compounds was similar to both a compound that was used to train the model, as well as an existing marketed drug [40]. Hence, despite the success of demonstrating how RL can expedite the drug discovery pipeline, future models will need to be coded such that newly generated compounds are dissimilar from both the input data and existing marketed compounds. Although in the pharmaceutical discipline, the use of RL has been limited to drug design, the wider medical community has explored other potentials for the algorithm. In a step towards personalised dosage, several simulation-based studies explored using RL to provide dynamic decision-making for sepsis treatment [41], anaesthetic drug delivery control [42], and detection of diabetic retinopathy [43]. The use of RL has also been extended to 'omics, bioimaging, and medical studies [28]. A schematic representation of RL is illustrated in Fig. 2a.

Transfer learning

If data are in short supply, then there are techniques that can be used to circumvent this problem. One such technique is transfer learning, which is the process of transferring knowledge acquired from solving one task to another related task. Transfer learning is

**FIGURE 1**

Examples of generated compounds using reinforcement learning (RL) reported by (a) Popova *et al.* [37] and (b) Zhavoronkov *et al.* [39].

**FIGURE 2**

Schematic representations of (a) reinforcement learning (RL), (b) transfer learning and (c) multitask learning.

an increasingly popular ML framework, particularly in medical image classification [44,45], that encompasses a range of techniques. Transfer learning is the improvement of learning a new task through the transfer of knowledge from a related task that has already been learned. The technique leverages the features generated from a large data set, **A** that is used to predict its target variable \mathbf{Y}_a , and sequentially transfer the knowledge to predict a different target, \mathbf{Y}_b , from a data set, **B**, which has insufficient data. In the context of deep learning, the learned weights of the models are trained using the larger data set and then transferred to perform models for new similar tasks (Fig. 2b). The approach has been found to outperform conventional MLTs that were trained on the smaller data set. Furthermore, transfer learning can be rapidly deployed for new models because the optimisation process has already been performed. It makes the assumption that the predictive features in the larger data set can in principle be applied to a different yet related task. In addition, if the features are physically related, the features learned can be transferred partially as input features for the target domain [46]. Transfer learning frameworks can comprise supervised and unsupervised learning techniques, where the latter is lacking labelled output variables for the target domain [47]. Transfer learning has been implemented using spectral [48,49], images [50,51], audio, text [52], and numeric [53] data types.

Turki *et al.* illustrated the potency of transfer learning in predicting the drug sensitivity of patients with multiple myeloma, where there was a lack of gene expression data, and acquiring new data was costly [54]. Using SVM and ridge regressions, the researchers trained the model on data from patients with lung and breast cancer, which were in abundance, and subsequently applied it to the multiple myeloma data set. The authors recorded a higher accuracy compared with their baseline. Most gene data sets generated by individual researchers are too small for MLTs. Taroni *et al.* leveraged the large, public expression compendia for transfer learning [55] and demonstrated that it was possible to describe biological processes more effectively than by using models trained only on their original features when using transfer learning. kNN regression-based transfer learning was combined with latent regression prediction to predict the sensitivity of different anticancer compounds [56]. Transfer learning was recently used to identify adverse drug reactions based on a model developed for automatic text classification of sentences to detect mentions of adverse drug reaction [57]. A large corpora source was used to train the model, and the knowledge gained was sequentially applied to a small-scale corpora. Other applications of transfer learning include incorporating the technique in de novo drug design [58–60].

ML has also been applied in material science, although its use is not as developed as in drug discovery and development. Material science is of interest to pharmaceutical formulation, and indeed is an allied field, sharing similar research concepts and approaches. Recently, transfer learning was applied to various materials, including small molecules, polymers, and inorganic crystalline materials [46]. The study was able to successfully apply transfer learning to a data set with a small number of observations. In addition, underlying links between small molecules and polymers, and between inorganic and organic chemistry, were revealed. For example, a mean absolute error and correlation values of 0.063 and 0.832, respectively, for predicting the refractive index were

obtained using the transferred features. By contrast, a notably poor error and correlation of 0.833 and 0.541, respectively, were obtained without transfer learning.

Multitask learning

Whereas transfer learning is the sequential learning and subsequent transfer of knowledge to another task, multitask learning is the simultaneous learning of different tasks in one model. It was observed that learning related tasks simultaneously led to an improved predictive performance than when learning the tasks individually (i.e., single task learning). The benefits of multitask learning are particularly useful in low-volume data sets and/or when noise is significant [61]. In addition, multitask learning was found to outperform traditional MLT, particularly when data were relatively sparse. Using the example of a neural network, a traditional architecture learns a single task at a time that outputs a single layer for the predictive task. By contrast, multitask learning outputs multiple hidden layers corresponding to the number of tasks predicted. The related tasks could be uncorrelated at the output layer, but they should be correlated at the internal representation level. Multitask learning allows for the inductive transfer of knowledge between tasks. This optimises multiple loss functions that can enable models to better generalise across multiple tasks. The improved predictability of multitask learning can be attributed to different factors [62]. With multitask learning, the data are amplified because of the extra information shared between the related tasks (Fig. 2c). The multiple tasks are able to learn from one another and are able to filter between relevant and irrelevant features, particularly where data are few and/or significant noise is present. Furthermore, bias and overfitting are mitigated, because the multiple tasks learn cooperatively. In the case of overfitting, multitask learning affords the multiple tasks to help each other to create a smoother dependence on common features. Multitask learning can be used for both supervised and unsupervised learning [63,64], and can be realised with different MLTs, such as neural networks, kNN [63], Bayesian multiple linear regression [65], and SVM [66].

In drug discovery, multitask learning has found application in addressing the effect of multitarget drugs. Such candidates were studied because their severe adverse effects, which is a negative consequence of acting on multiple targets. Of equal importance, it was recently demonstrated that multitarget drugs have been found to be more effective than single-target drugs for several complex diseases, such as cancer and metabolic diseases. This rationale was leveraged by Li *et al.*, who showed that multitask learning could discover useful multiple targets that are affected by the same drug [67]. The researchers used unsupervised ML for their approach and both expression data and compound structure information. Yang *et al.* developed a multitask framework, called Macau, for large-scale drug screening, while simultaneously deriving interpretable insights about the interactions between the characteristics of the drugs and the cell lines [68]. Their algorithm used Bayesian multitask multi-relation to explore the interaction between the drug targets and signalling pathway activation using drug and gene data. Gene expressions were used as molecular inputs to predict signalling pathways; whereas, for the drug, their nominal targets were used as inputs. The rationale for their work was that the interaction between drug targets and signalling pathways can

provide novel in-depth views of cellular mechanisms and drug mode of action.

In addition to sequential learning, multitask learning can be combined with gradient-boosting decision trees for small data sets [69]. Four data sets were investigated using this approach, with test sizes of 7413, 1792, 823, and 353 compounds. For the smallest set of 353 compounds, the R^2 values when gradient boosting and multitask learning were used were 0.472 and 0.721, respectively. Combining the two techniques resulted in a R^2 value of 0.733, which is an improvement on both individual techniques.

Multitask learning was also revealed by Weng *et al.* to simultaneously learn both classification and regression task analyses for drug–target interactions [70]. Classification tasks are prone to higher bias, whereas regression models are susceptible to overfitting because of the large variance encountered. Thus, to address the trade-off between bias and variance, a convolutional neural network model was developed to simultaneously optimise the regression and classification loss, using shared features. In another application, Han *et al.* used multitask learning for sentiment analysis of drug reviews [71]. The main objective was to identify people's sentiment, opinions, and attitudes from a collection of 4200 drug reviews. In addition, Zubatyuk *et al.* combined multitask and multimodal learning to overcome sparsity in training data. Another key benefit of their approach is that the results were comparable to the density functional theory (DFT) method, which is a considerably more expensive modelling method.

Active learning

Active learning is a unique semiautomated ML approach that also seeks to address the issue of low-labelled data sets using user feedback. In contrast to passive learning, active learning is ideal where there is an abundance of unsupervised training data that require costly and resource-intensive experiments to label. Consequently, the user can conduct experiments and subsequently label the data for a subset of the data set and use active learning to obtain the predictions for the remaining unlabelled data. Using this approach, active learning makes queries of samples that it is unsure of. For example, in using ML to predict the penetration of drugs through the blood–brain barrier, one can perform the experiment on 10% of the molecules, and train the model using said

10% to make predictions for the other 90%. Where the model is uncertain, it will make a query and the researcher can then perform the experiments on those samples. Hence, compared with passive learning, it has the potential to require considerably fewer labelled data [72], and thereby accelerate the drug discovery process while minimising costs. Further information regarding active learning, including sampling method and query strategies, can be found in [73].

Active learning models can be built using conventional MLTs, such as SVM, and also deep learning [73,74]. Recent work demonstrated that active learning can be used for predicting small-molecule bioactivity, ligand–target interactions, and toxicity [75–78].

Generative models

As described earlier, generative models are MLTs capable of generating new samples. This was leveraged for RL *de novo* applications, but generative models can also be used as standalone techniques. Generative models distinguish themselves from discriminative models by directly learning from the input data and do not necessarily require explicit rules to be coded by users. Generative models can generate new data instances through implementing a probabilistic estimator of data distribution, where the new data lie within the distribution. In other words, generative models are able to generate new samples for a given distribution. This contrasts with discriminative models, which reveal the probability of the labelled data given the data instance, regardless of whether the data instance is valid (Fig. 3). Recent studies used deep-learning generative models, which, in addition to generating new compounds, can be used for data augmentation when working with small data sets, and dimensionality reduction [79–81]. As mentioned earlier, newly generated molecules will need to be thoroughly assessed to ensure that they are distinct from compounds that already exist in the market and/or different to compounds fed into the model.

Bayesian neural networks

BNNs are ensemble models that combine multiple neural network models using Bayesian inference [82]. Unlike conventional neural networks, which require large amount of data for training, BNN

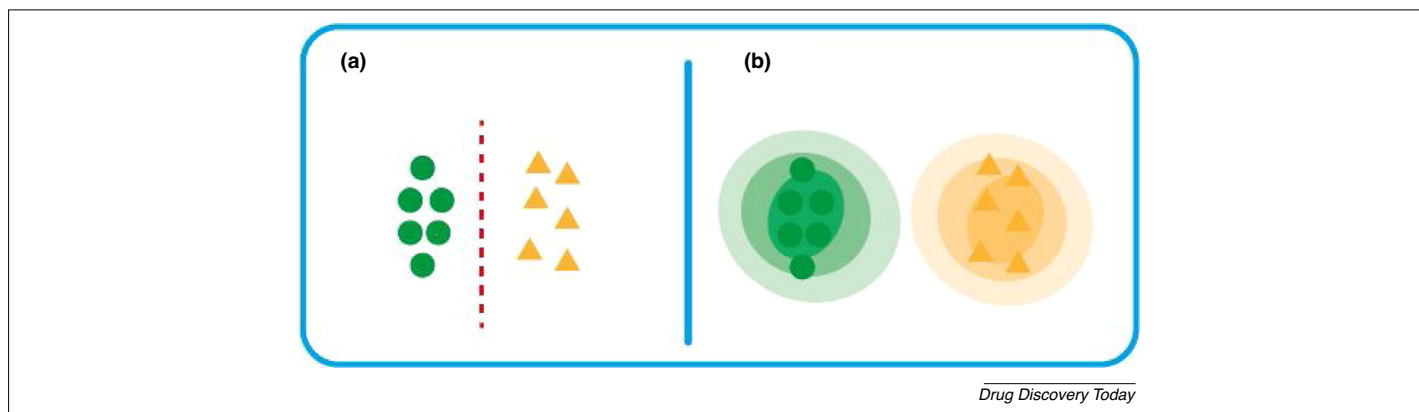


FIGURE 3

Differences between (a) discriminative and (b) generative modelling. Discriminative modelling seeks to classify through establishing, for example, decision boundaries. By contrast, generative models look at the probability distribution of the classes.

can handle small data sets because of their ability to avoid overfitting. Overfitting is a problem associated with most conventional MLTs, which BNN avoids through prior probability distribution to compute the average across numerous models during training, which yields a regularisation effect to the network [83]. In other words, the weights and biases for neurons are not a single value but rather sampled from a distribution, which is regularly updated to train the BNN. The use of BNN has not been thoroughly explored for drug discovery. A recent study revealed that Bayesian graph networks outperformed conventional graph networks in predicting the inhibitory activity of molecules, using the ChEMBL data set [84]. BNNs were also used to identify genes associated with anticancer drug sensitivities using data gathered from the cancer cell line encyclopaedia study [85]. More recently, BNN were applied for identifying drug-likeness, where the Bayesian error distribution of individual classifiers can yield an accuracy of 93% for distinguishing drug-like from nondrug-like molecules [86]. Although BNNs are able to address some of the shortcomings of neural networks, they require a comparatively large effort to design the neural net, which can lead to establishing casual influences that are recognised by the individual programming it.

Explainable algorithms

The use of ML is indeed to facilitate and expedite decision-making, particularly for routine tasks. Thus, it might not be necessary to understand the decision-making process achieved by the model. However, understanding the decision process made by ML will instil confidence in researchers. Interpreting the model can help researchers troubleshoot when the model appears erroneous. In addition, the insight from the decision process could lead to plausible research questions. In addition, it can facilitate research understanding by providing insight into the decision making. Equally, transparency might also instil trust in regulatory bodies if the technology is to be commercialised.

A recent example of explainable ML was applied to quality structure–activity relationship modelling, wherein semisupervised regression trees were found to outperform supervised regression trees [87]. Using a different strategy for predicting activity, Rodriguez-Perez and Bajorath developed a method that elucidates the prediction process of conventional techniques, as well as ensemble and deep-learning models [88]. The focus of their work was to eliminate the ‘black-box’ nature of ML models. The approach was based on Shapley values that was initially developed for game theory, but were demonstrated by the authors to be applicable to ML. In their approach, each feature was assigned an importance value for a given prediction and, in turn, it provided an overview of which features have the most contribution to a model. Moreover, their approach uncovered model errors and consequently provided rationales for inaccurate predictions, which otherwise could not have been readily rationalised.

Emerging machine-learning techniques

Hybrid quantum-machine learning

The hybridisation of ML with quantum computing has emerged as a powerful technology in predictive analysis [89]. The main promise of quantum computing is the efficiency to solve complex problems that are prohibitively expensive for classical computers [90]. In classical models, the processing units compute bits that are

either 0 or 1, whereas for quantum computing, the quantum bits, qubits, are in a superimposed state of both 0 and 1 [91]. The qubits are processed by quantum logic gates, which, in contrast to classical logic gates, are reversible. This yields computing prowess that prevents loss of information [92], faster analysis, and low power consumption [93]. The qubits and quantum gates are components of the quantum circuit that has been demonstrated to perform tasks that were quadratic, polynomial, or exponentially faster than their classical counterparts [94–97]. The definition of hybrid quantum ML is yet to be decided upon. To date, it encompasses the use of quantum computers to execute ML algorithms or adopting quantum information processing into ML [94,98]. The former approach can be regarded of as quantum-enhanced ML, whereas the latter can be regarded as quantum-inspired ML. Examples of hybrid quantum ML include supervised [99], unsupervised [100], and RL [101].

The advantages of H-QML can indeed be leveraged in pharmaceutical sciences, however, at the time of writing, the technology has not yet been applied. In 2018, International Business Machines Corporation (IBM) published an article on *The potential of quantum computing for drug discovery*, wherein the authors included the potential of quantum ML in the scope of their review [102]. More recently, Google LLC released an open-access quantum ML framework for python that will enable researchers to use hybrid quantum ML [103]. Therefore, the promise of hybrid-quantum ML in pharmaceutical sciences is likely to be realised soon.

Recommendation systems

Recommendation systems gained fame in 2006 with the announcement of a Netflix competition seeking to create accurate user preference content for its users. A recommendation system is a ML framework that is based on data establishing links between a set of users (e.g., customers) to a set of items (e.g., products) [104]. Recommendation systems are heavily used in e-commerce, for example by Amazon and YouTube, to drive their sales [105]. The advantageous of such techniques are their ability to handle sparsity in data, to make predictions if prior information is unavailable, and to provide transparency by explaining how the recommender system makes the decision [106].

Recommender systems have been investigated for medical applications, where the right treatment is proposed based on the patient’s medical history [107,108]. However, applications in drug discovery and development are yet to be established. Sosnina *et al.* developed a recommender system for compound–target interaction prediction for antiviral drug discovery [109]. The authors used a content-based filtering recommender system, which is suitable for sparse data and interpretability. In addition, their model made it possible to perform cold-start prediction, in which predictions can be made where there is no experimental data. Given that data in drug discovery and development are afflicted by all three issues, it is anticipated that the use of recommender system will increase.

Concluding remarks

Here, we have presented examples of MLT used to circumvent the issues surrounding conventional techniques. We have detailed the use of ML for automating processes without human involvement; the use of transfer learning and multitask learning for when big

data are lacking; BNNs for avoiding overfitting; and explainable algorithms that can shed light the decision-making process of a model. In addition, emerging techniques and their potential involvement in drug discovery were also discussed. Hybrid quantum-ML has the potential to further improve prediction performance, whereas recommendation systems can address data sparsity. It is anticipated that the use of the techniques discussed herein will be adopted in the near future, and that their application will further

progress research in drug discovery. Ultimately, the quality of the predictions made by the models will depend on the quality of the data. Thus, the application of ML in drug discovery will benefit from a strategic and unified database.

Acknowledgements

The authors thank the Engineering and Physical Sciences Research Council (EPSRC), UK for its financial support (EP/S009000/1).

References

- Rohall, S.L. *et al.* (2020) An artificial intelligence approach to proactively inspire drug discovery with recommendations. *J. Med. Chem.* 63, 8824–8834
- Montanari, F. *et al.* (2020) Modeling physico-chemical ADMET endpoints with multitask graph convolutional networks. *Molecules* 25, 44
- Kotsias, P.-C. *et al.* (2020) Direct steering of de novo molecular generation with descriptor conditional recurrent neural networks. *Nat. Mach. Intell.* 2, 254–265
- Davies, M. *et al.* (2020) Improving the accuracy of predicted human pharmacokinetics: lessons learned from the AstraZeneca drug pipeline over two decades. *Trends Pharm. Sci.* 41, 390–408
- Zarringhalam, K. *et al.* (2018) Robust phenotype prediction from gene expression data using differential shrinkage of co-regulated genes. *Sci. Rep.* 8, 1237
- Zhou, H. *et al.* (2019) Molecular profiling of ulcerative colitis subjects from the TURANDOT trial reveals novel pharmacodynamic/efficacy biomarkers. *J. Crohn's Colitis* 13, 702–713
- Elbadawi, M. *et al.* (2020) M3DISEEN: a novel machine learning approach for predicting the 3D printability of medicines. *Int. J. Pharm.* 590, 119837
- Lo, Y.-C. *et al.* (2018) Machine learning in chemoinformatics and drug discovery. *Drug Discov. Today* 23, 1538–1546
- Vamathevan, J. *et al.* (2019) Applications of machine learning in drug discovery and development. *Nat. Rev. Drug Discov.* 18, 463–477
- Réda, C. *et al.* (2020) Machine learning applications in drug development. *Comput. Struct. Biotechnol. J.* 18, 241–252
- Elbadawi, M. *et al.* (2020) 3D printing tablets: predicting printability and drug dissolution from rheological data. *Int. J. Pharm.* 590, 119868
- Hoffmann, H. (2007) Kernel PCA for novelty detection. *Pattern Recognit.* 40, 863–874
- Rensi, S.E. and Altman, R.B. (2017) Shallow representation learning via kernel PCA improves QSAR modelability. *J. Chem. Inform. Model.* 57, 1859–1867
- Bzdok, D. *et al.* (2018) Machine learning: supervised methods. *Nat. Methods* 15, 5–6
- Mishra, P.K. and Saroha, G. (2016) A study on classification for static and moving object in video surveillance system. *Int. J. Image Graphics Signal Process.* 8, 76
- Sathya, K. and Karthiban, R. (2020) Performance analysis of heart disease classification for computer diagnosis system. *2020 International Conference on Computer Communication and Informatics (ICCCI)*. <http://dx.doi.org/10.1109/ICCCI48352.2020.9104089>
- Malik, M.N. *et al.* (2019) ADLAuth: passive authentication based on activity of daily living using heterogeneous sensing in smart cities. *Sensors* 19, 2466
- Chan, H.C.S. *et al.* (2019) Advancing drug discovery via artificial intelligence. *Trends Pharm. Sci.* 40, 592–604
- Chen, H. *et al.* (2018) The rise of deep learning in drug discovery. *Drug Discov. Today* 23, 1241–1250
- Öztürk, H. *et al.* (2020) Exploring chemical space using natural language processing methodologies for drug discovery. *Drug Discov. Today* 25, 689–705
- Richards, N. *et al.* (1998) Evolving neural networks to play Go. *Appl. Intell.* 8, 85–96
- Silver, D. *et al.* (2016) Mastering the game of Go with deep neural networks and tree search. *Nature* 529, 484–489
- Gibb, B. (2017) Hedgehogs and foxes (and a bear). *Nat. Chem.* 9, 103–104
- Polydoros, A.S. and Nalpantidis, L. (2017) Survey of model-based reinforcement learning: applications on robotics. *J. Intell. Robot. Syst.* 86, 153–173
- Quah, K.H. *et al.* (2005) Reinforcement learning combined with a fuzzy adaptive learning control network (FALCON-R) for pattern classification. *Pattern Recognit.* 38, 513–526
- Abdulhai, B. *et al.* (2003) Reinforcement learning for true adaptive traffic signal control. *J. Transp. Eng.* 129, 278–285
- Niv, Y. (2009) Reinforcement learning in the brain. *J. Math. Psychol.* 53, 139–154
- Mahmud, M. *et al.* (2018) Applications of deep learning and reinforcement learning to biological data. *IEEE Trans. Neural Netw. Learn. Syst.* 29, 2063–2079
- Valasek, J. *et al.* (2012) Morphing unmanned air vehicle intelligent shape and flight control. In *Morphing Aerospace Vehicles and Structures*. pp. 55–86, AIAA
- Moody, J. *et al.* (1998) Performance functions and reinforcement learning for trading systems and portfolios. *J. Forecasting* 17, 441–470
- Arulkumaran, K. *et al.* (2017) A brief survey of deep reinforcement learning. *arXiv* 2017. arXiv:1708.05866
- Ståhl, N. *et al.* (2019) Deep reinforcement learning for multiparameter optimization in de novo drug design. *J. Chem. Inf. Model.* 59, 3166–3176
- Putin, E. *et al.* (2018) Reinforced adversarial neural computer for de novo molecular design. *J. Chem. Inf. Model.* 58, 1194–1204
- Sanchez-Lengeling, B. *et al.* (2017) Optimizing distributions over molecular space. An objective-reinforced generative adversarial network for inverse-design chemistry (ORGANIC). *ChemRxiv* 2017. chemrxiv.5309668.v3
- Olivecrona, M. *et al.* (2017) Molecular de-novo design through deep reinforcement learning. *J. Cheminf.* 9, 48
- Zhou, Z. *et al.* (2019) Optimization of molecules via deep reinforcement learning. *Sci. Rep.* 9, 10752
- Popova, M. *et al.* (2018) Deep reinforcement learning for de novo drug design. *Sci. Adv.* 4, eaap7885
- Gao, K. *et al.* (2020) Generative network complex for the automated generation of druglike molecules. *arXiv* 2020. arXiv:2005.14286
- Zhavoronkov, A. *et al.* (2019) Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nat. Biotechnol.* 37, 1038–1040
- Walters, W.P. and Murcko, M. (2020) Assessing the impact of generative AI on medicinal chemistry. *Nat. Biotechnol.* 38, 143–145
- Bennett, C.C. and Hauser, K. (2013) Artificial intelligence framework for simulating clinical decision-making: a Markov decision process approach. *Artif. Intel. Med.* 57, 9–19
- Lowery, C. and Faisal, A.A. (2013) Towards efficient, personalized anesthesia using continuous reinforcement learning for propofol infusion control. *2013 6th International IEEE/EMBS Conference on Neural Engineering (NER)* 1414–1417 IEEE
- Gulshan, V. *et al.* (2016) Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA* 316, 2402–2410
- Cheplygina, V. *et al.* (2019) Not-so-supervised: a survey of semi-supervised, multi-instance, and transfer learning in medical image analysis. *Med. Image Anal.* 54, 280–296
- Huynh, B. *et al.* (2016) Digital mammographic tumor classification using transfer learning from deep convolutional neural networks. *J. Med. Imaging* 3, 034501
- Yamada, H. *et al.* (2019) Predicting materials properties with little data using shotgun transfer learning. *ACS Cent. Sci.* 5, 1717–1730
- Arnold, A. *et al.* (2007) A comparative study of methods for transductive transfer learning. *Seventh IEEE International Conference on Data Mining Workshops (ICDMW 2007)* 77–82
- Jayaram, V. *et al.* (2016) Transfer learning in brain-computer interfaces. *IEEE Comput. Intell. Mag.* 11, 20–31
- Makin, J.G. *et al.* (2020) Machine translation of cortical activity to text with an encoder-decoder framework. *Nat. Neurosci.* 23, 575–582
- Yildirim, O. *et al.* (2019) Automated detection of diabetic subject using pre-trained 2D-CNN models with frequency spectrum images extracted from heart rate signals. *Comput. Biol. Med.* 113, 103387
- Celik, Y. *et al.* (2020) Automated invasive ductal carcinoma detection based using deep transfer learning with whole-slide images. *Pattern Recognit. Lett.* 133, 232–239
- Pan, W. *et al.* (2012) Transfer learning for text mining. In *Mining Text Data* (Aggarwal, C.C. and Zhai, C., eds), pp. 223–257, Springer US
- Qureshi, A.S. *et al.* (2017) Wind power prediction using deep neural network based meta regression and transfer learning. *Appl. Soft Comput.* 58, 742–755
- Turki, T. *et al.* (2017) Transfer learning approaches to improve drug sensitivity prediction in multiple myeloma patients. *IEEE Access* 5, 7381–7393
- Taroni, J.N. *et al.* (2019) MultiPLIER: a transfer learning framework for transcriptomics reveals systemic features of rare disease. *Cell Syst.* 8, 380–394

- 56 Dhruva, S.R. *et al.* (2018) Application of transfer learning for cancer drug sensitivity prediction. *BMC Bioinf.* 19, 497
- 57 Li, Z. *et al.* (2020) Exploiting adversarial transfer learning for adverse drug reaction detection from texts. *J. Biomed. Inf.* 106, 103431
- 58 Yasonik, J. (2020) Multiobjective de novo drug design with recurrent neural networks and nondominated sorting. *J. Cheminf.* 12, 14
- 59 Gupta, A. *et al.* (2018) Generative recurrent networks for de novo drug design. *Mol. Inf.* 37, 1700111
- 60 Segler, M.H.S. *et al.* (2018) Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS Cent. Sci.* 4, 120–131
- 61 Wu, K. and Wei, G.-W. (2018) Quantitative toxicity prediction using topology based multitask deep neural networks. *J. Chem. Inf. Model.* 58, 520–531
- 62 Varnek, A. *et al.* (2009) Inductive transfer of knowledge: application of multi-task learning and feature net approaches to model tissue-air partition coefficients. *J. Chem. Inf. Model.* 49, 133–144
- 63 Zhang, Y. and Yang, Q. (2017) A survey on multi-task learning. *arXiv* 2017 arXiv:1707.08114
- 64 Caruana, R. (1997) Multitask learning. *Mach. Learn.* 28, 41–75
- 65 Ammad-ud-din, M. *et al.* (2017) Systematic identification of feature combinations for predicting drug response with Bayesian multi-view multi-task linear regression. *Bioinformatics* 33, i359–i368
- 66 He, X. *et al.* (2014) Multi-task learning with one-class SVM. *Neurocomputing* 133, 416–426
- 67 Li, L. *et al.* (2018) Multi-target drug repositioning by bipartite block-wise sparse multi-task learning. *BMC Syst. Biol.* 12, 55
- 68 Yang, M. *et al.* (2018) Linking drug target and pathway activation for effective therapy using multi-task learning. *Sci. Rep.* 8, 8322
- 69 Jiang, J. *et al.* (2020) Boosting tree-assisted multitask deep learning for small scientific datasets. *J. Chem. Inf. Model.* 60, 1235–1244
- 70 Weng, Y. *et al.* (2019) Drug target interaction prediction using multi-task learning and co-attention. *2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)* 528–533 IEEE
- 71 Han, Y. *et al.* (2020) Aspect-level drug reviews sentiment analysis based on double BiGRU and knowledge transfer. *IEEE Access* 8, 21314–21325
- 72 Chen, C.-T. and Gu, G.X. (2020) Generative deep neural networks for inverse materials design using backpropagation and active learning. *Adv. Sci.* 7, 1902607
- 73 Reker, D. (2020) Practical considerations for active machine learning in drug discovery. *Drug Discov. Today Technol.* . <http://dx.doi.org/10.1016/j.ddtec.2020.06.001>
- 74 Warmuth, M.K. *et al.* (2003) Active learning with support vector machines in the drug discovery process. *J. Chem. Inf. Comput. Sci.* 43, 667–673
- 75 Reker, D. *et al.* (2017) Active learning for computational chemogenomics. *Future Med. Chem.* 9, 381–402
- 76 Zhang, Y. and Lee, A.A. (2019) Bayesian semi-supervised learning for uncertainty-calibrated prediction of molecular properties and active learning. *Chem. Sci.* 10, 8154–8163
- 77 Borkowski, O. *et al.* (2020) Large scale active-learning-guided exploration for *in vitro* protein production optimization. *Nat. Commun.* 11, 1872
- 78 Habib Polash, A. *et al.* (2020) Active learning efficiently converges on rational limits of toxicity prediction and identifies patterns for molecule design. *Comput. Toxicol.* 15, 100129
- 79 Lin, E. *et al.* (2020) A deep adversarial variational autoencoder model for dimensionality reduction in single-cell RNA sequencing analysis. *BMC Bioinf.* 21, 64
- 80 Ding, J. *et al.* (2018) Interpretable dimensionality reduction of single cell transcriptome data with deep generative models. *Nat. Commun.* 9, 2002
- 81 Arús-Pous, J. *et al.* (2020) SMILES-based deep generative scaffold decorator for de-novo drug design. *J. Cheminf.* 12, 38
- 82 Cai, R. *et al.* (2018) VIBNN: Hardware acceleration of Bayesian neural networks. *ACM SIGPLAN Notices* 53, 476–488
- 83 Shridhar, K. *et al.* (2019) A comprehensive guide to bayesian convolutional neural network with variational inference. *arXiv* 2019 arXiv:1901.02731
- 84 Ryu, S. *et al.* (2019) A Bayesian graph convolutional network for reliable prediction of molecular properties with uncertainty quantification. *Chem. Sci.* 10, 8438–8446
- 85 Liang, F. *et al.* (2018) Bayesian neural networks for selection of drug sensitive genes. *J. Am. Stat. Assoc.* 113, 955–972
- 86 Beker, W. *et al.* (2020) Minimal-uncertainty prediction of general drug-likeness based on Bayesian neural networks. *Nat. Mach. Intell.* 2, 457–465
- 87 Levatić, J. *et al.* (2020) Semi-supervised regression trees with application to QSAR modelling. *Expert Syst. Appl.* 158, 113569
- 88 Rodríguez-Pérez, R. and Bajorath, J. (2019) Interpretation of compound activity predictions from complex machine learning models using local approximations and Shapley values. *J. Med. Chem.* 63, 8761–8777
- 89 Xia, R. and Kais, S. (2018) Quantum machine learning for electronic structure calculations. *Nat. Commun.* 9, 4195
- 90 Risté, D. *et al.* (2017) Demonstration of quantum advantage in machine learning. *NPJ Quantum Inf.* 3, 16
- 91 Steane, A. (1998) Quantum computing. *Rep. Prog. Phys.* 61, 117–173
- 92 Wittek, P. (2014) *Quantum Machine Learning: What Quantum Computing Means to Data Mining*. Academic Press
- 93 Al-Rabadi, A.N. (2012) *Reversible Logic Synthesis: From Fundamentals to Quantum Computing*. Springer Science & Business Media
- 94 Biamonte, J. *et al.* (2017) Quantum machine learning. *Nature* 549, 195–202
- 95 Li, J.-A. *et al.* (2020) Quantum reinforcement learning during human decision-making. *Nat. Hum. Behav.* 4, 294–307
- 96 Aïmeur, E. *et al.* (2013) Quantum speed-up for unsupervised learning. *Mach. Learn.* 90, 261–287
- 97 Li, Z. *et al.* (2015) Experimental realization of a quantum support vector machine. *Phys. Rev. Lett.* 114, 140504
- 98 Wan, K.H. *et al.* (2017) Quantum generalisation of feedforward neural networks. *NPJ Quantum Inf.* 3, 36
- 99 Havlíček, V. *et al.* (2019) Supervised learning with quantum-enhanced feature spaces. *Nature* 567, 209–212
- 100 Zhang, Y. and Ni, Q. (2020) Recent advances in quantum machine learning. *Quantum Eng.* 2, e34
- 101 Albarrán-Arriagada, F. *et al.* (2018) Measurement-based adaptation protocol with quantum reinforcement learning. *Phys. Rev. A* 98, 042315
- 102 Cao, Y. *et al.* (2018) Potential of quantum computing for drug discovery. *IBM J. Res. Dev.* 62, 6:1–6:20
- 103 Broughton, M. *et al.* (2020) Tensorflow quantum: a software framework for quantum machine learning. *arXiv* 2020 arXiv:2003.02989
- 104 Jordan, M.I. and Mitchell, T.M. (2015) Machine learning: trends, perspectives, and prospects. *Science* 349, 255–260
- 105 Smith, B. and Linden, G. (2017) Two decades of recommender systems at Amazon. *IEEE Internet Comput.* 21, 12–18
- 106 Chadha, A. and Kaur, P. (2015) Comparative analysis of recommendation system. *2015 4th International Symposium on Emerging Trends and Technologies in Libraries and Information Services* 313–318 IEEE
- 107 Chen, R.-C. *et al.* (2012) A recommendation system based on domain ontology and SWRL for anti-diabetic drugs selection. *Expert Syst. Appl.* 39, 3995–4006
- 108 Van-Dai, T. *et al.* (2016) Big data stream computing in healthcare real-time analytics. *2016 IEEE International Conference on Cloud Computing and Big Data Analysis (ICCCBDA)* 37–42 IEEE
- 109 Sosnina, E.A. *et al.* (2020) Recommender systems in antiviral drug discovery. *ACS Omega* 5, 15039–15051