



Machine learning and AI-based approaches for bioactive ligand discovery and GPCR-ligand recognition

Sebastian Raschka^{a,*}, Benjamin Kaufman^b

^a University of Wisconsin–Madison, Department of Statistics, United States

^b University of Wisconsin–Madison, Department of Biostatistics and Medical Informatics, United States

ARTICLE INFO

Keywords:

Molecular representations
GPCR ligands
Drug discovery
Deep learning
Machine learning
Graph convolutional neural networks

ABSTRACT

In the last decade, machine learning and artificial intelligence applications have received a significant boost in performance and attention in both academic research and industry. The success behind most of the recent state-of-the-art methods can be attributed to the latest developments in deep learning. When applied to various scientific domains that are concerned with the processing of non-tabular data, for example, image or text, deep learning has been shown to outperform not only conventional machine learning but also highly specialized tools developed by domain experts. This review aims to summarize AI-based research for GPCR bioactive ligand discovery with a particular focus on the most recent achievements and research trends. To make this article accessible to a broad audience of computational scientists, we provide instructive explanations of the underlying methodology, including overviews of the most commonly used deep learning architectures and feature representations of molecular data. We highlight the latest AI-based research that has led to the successful discovery of GPCR bioactive ligands. However, an equal focus of this review is on the discussion of machine learning-based technology that has been applied to ligand discovery in general and has the potential to pave the way for successful GPCR bioactive ligand discovery in the future. This review concludes with a brief outlook highlighting the recent research trends in deep learning, such as active learning and semi-supervised learning, which have great potential for advancing bioactive ligand discovery.

1. Introduction

G protein-coupled receptors (GPCRs) are one of the most prominent families of integral membrane proteins and are among the most widely studied targets in drug discovery and development. According to [1], 34 percent of all drugs approved by the US Food and Drug Agency target GPCRs. While the exact size of the GPCR family remains to be determined, genome analyses suggest that the GPCR family is represented by approximately 800–1000 genes in humans [2,3]. More than 150 GPCRs are characterized as orphan receptors, which means that the receptors' endogenous ligands are still unknown [4]. However, it remains unclear whether endogenous ligands exist for all orphan GPCRs [5].

DNN
ECFP
GPCR
GPU
GRU
GNN
QSAR
LSTM
MLP
NN
RNN
SLOR1
SVM
VS
VAE

Deep neural network
Extended connectivity fingerprint
G protein-coupled receptor
Graphics processing unit
Gated recurrent unit
Graph neural network
Quantitative structure–activity relationship
Long short-term memory
Multilayer perceptron
Neural network
Recurrent neural network
Sea lamprey receptor 1
Support vector machine
Virtual screening
Variational autoencoder

Since GPCRs are involved in many different cellular and biological processes, make excellent drug targets, and remain orphaned to a relatively large extent, the prediction and consequent identification of GPCR ligands is an active area of research and interest [6]. Recent

Abbreviations

AI Artificial intelligence
CNN Convolutional neural network

* Corresponding author at: Medical Science Center, 1300 University Ave, Madison, WI 53706, United States.

E-mail address: sraschka@wisc.edu (S. Raschka).

URL: <http://pages.stat.wisc.edu/sraschka/> (S. Raschka).

¹ ORCID: 0000-0001-6989-4493

<https://doi.org/10.1016/j.ymeth.2020.06.016>

Received 17 January 2020; Received in revised form 23 June 2020; Accepted 23 June 2020

1046-2023/© 2020 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

studies suggest that drugs target only approximately 10% of known GPCRs [7]. The high cost of clinical trials coupled with a relatively low success rate (approximated to be below 6.2% [8,9]), motivates the development and application of machine learning and artificial intelligence-based methods for ligand discovery, to lower costs as well as to identify candidates that would otherwise be missed using conventional methods.

1.1. GPCRs and drug discovery

Many common human diseases involve GPCR signaling, including schizophrenia, glaucoma, depression, and hypertension [2]. Despite the increasing popularity of biologics, which includes gene therapy, recombinant therapeutic proteins, antibody-drug conjugates, and tissues [10], experts see the discovery of small-molecule ligands being crucial for the future of molecular medicine and the treatment of human diseases [11,12].

GPCR ligands exhibit a wide range of characteristics and are very diverse in their physicochemical properties, shape, and size [7]. Furthermore, GPCR ligands include lipids, peptides, proteins, steroids, and other small organic molecules. Depending on their mechanism of action, GPCR ligands can be described as full or partial agonists, antagonists, or inverse agonists. This diversity among GPCR ligands poses challenges for standardized binding assays used in experimental bioactivity screening. Thus, wet-lab techniques such as high throughput screening have risen in popularity. However, high throughput screening can still be cost-prohibitive and labor-intensive. Another downside of high throughput screening is the limitation of available ligand libraries in terms of size and diversity. With drug discovery being an expensive and labor-intensive process, with estimated costs ranging between 0.5–2.6 billion US dollars, and a timeline ranging between 10–20 years, researchers have begun to favor computational methods for identifying candidate molecules for more elaborate bioactivity assays [13,14].

When it comes to the identification of protein targets of small-molecule ligands and assessing potential off-target activity, chemical proteomics are considered the gold standard [12,15–17]. However, similar to the challenges of binding assays for small molecule screening, chemical proteomics experiments are laborious and challenging to streamline [18]. Hence, many researchers are increasingly favoring computational alternatives [12] or augmenting the experimental discovery pipeline with artificial intelligence [19,20].

1.2. Computer-aided ligand discovery

Modern computational databases feature hundreds of millions of molecules and are available free of charge [21–23], which makes computer-aided ligand discovery a particularly compelling alternative to high throughput screening and other experimental approaches during the early stages of ligand discovery. The two major approaches for computer-aided ligand discovery, also known as virtual screening (VS), are ligand-based VS and structure-based VS Fig. 1.

1.2.1. Ligand-based VS

Ligand-based VS approaches focus on the structure and physicochemical properties of ligands in the absence of the receptor structure. Broadly, ligand-based VS can be described as a set of techniques for ligand-based similarity search or property prediction. Ligand-based VS approaches for identifying bioactive molecules are often based on a known bioactive molecule given the hypothesis that similar molecules are likely to bind the target receptor and exhibit a certain level of bioactivity, which has to be determined experimentally. As traditional ligand-based VS depends on pair-wise similarity measures between database molecules and a receptor agonist or antagonist, researchers have to choose a meaningful similarity or distance measure. The choice of an appropriate similarity measure, in turn, depends on the molecular representation, which we discuss in more detail in Section 2. To provide

an illustrative example of the use of ligand-based VS in GPCR bioactive ligand discovery, we consider the recent discovery of a potent inhibitor of a GPCR-mediated pheromone signaling pathway in sea lamprey [24]. Based on volumetric and physicochemical overlays between 3D conformers of database molecules and the cognate ligand, researchers were able to identify a small molecule capable of blocking a GPCR-mediated signaling response. However, when the researchers assayed the activity of the 299 top-scoring molecules, they found no correlation between molecular activity and the 3D overlay-based similarity scores. Additionally, the researchers found that changing only a single chemical group (for instance, keto to hydroxyl), which may only have a small impact on the overall similarity score, can completely suppress GPCR-mediated signaling in the pheromone pathway. This finding provides supportive evidence that general similarity measures are useful but not sufficient means for active molecule discovery. However, the virtual screening data combined with the activity measurements collected via experimental assay data offer excellent opportunities to utilize machine learning for developing custom activity classifiers and conducting quantitative structure–activity relationship (QSAR) analyses. For instance, in a follow-up study, the researchers described how machine learning techniques were used to identify the chemical features that correlated with bioactivity [25]. To facilitate this analysis, the researchers converted functional group matches between database compounds and the receptor agonist into binary feature vectors. For example, if the keto-group of the known agonist was within 1.3 Å of a query molecule's keto-group (in the in the 3D overlay), it was counted as a functional group match (1) and as a non-match (0), otherwise. After tabulating the functional group matches as binary feature vectors, the researchers fit logistic regression and random forest classification models and quantified feature importance of the functional and chemical groups. The importance of certain chemical groups was then used to inform subsequent rounds of virtual screening, to identify database molecules that were previously dismissed in early stages of the ligand-based VS pipeline.

1.2.2. Receptor structure-based VS

In contrast to ligand-based VS, receptor structure-based VS assumes knowledge of the receptor structure with molecular docking being one of the most prominent structure-based VS techniques. While molecular docking has led to many successful discoveries, one of its limitations for GPCR ligand discovery is the small number of high-resolution GPCR structures that are currently available. Recent years brought many significant improvements for protein X-ray crystallography and cryo-electron microscopy; however, high-resolution structures still only cover four out of the six different classes of GPCR: classes A, B, C, and F, of which the rhodopsin-like receptors from class A form the most significant portion [26]. Furthermore, another limitation for structure-based approaches for bioactive ligand discovery is the breadth and diversity of GPCR structures that are currently available: only 44 of the 205 GPCR structures correspond to unique GPCRs [1].

All GPCRs consist of seven transmembrane helices, which are relatively conserved across the different classes of GPCRs (A-F). However, GPCRs may differ more noticeably in the intracellular and extracellular loop domains. The extracellular loops play an essential role in many GPCR structures, since they often form the orthosteric ligand binding site or provide access to binding sites that are located within the transmembrane bundle [27]. The diversity of GPCR ligands and GPCR ligand binding sites are a direct consequence of the structural variety of the extracellular loops, which makes employing structure-based approaches extremely challenging in the absence of high-quality structural information. In those cases, ligand-based strategies may represent the only viable alternative [24]. In recent years, both traditional structure-based and ligand-based VS have been augmented or replaced by machine learning techniques. Furthermore, machine learning has also been used for predicting bioactivity from other types of data such as a bioactivity matrix (a targets-by-compounds matrix of function




Fig. 1. Conceptual overview of ligand-based and structure-based virtual screening depicting 3kPZS (a known agonist), which is a pheromone compound involved GPCR signaling, and a homology structure of its receptor, the SLOR1 (sea lamprey receptor 1) GPCR [24].

interactions) [28]. Today, machine learning is widely recognized as an essential method in the chemical biology toolbox for researching ligand binding [12]. Moreover, besides binding mode prediction and scoring, there are various other aspects of structure-based approaches that can benefit from machine learning. For instance, in a recent study, researchers combined X-ray crystal structural analysis with machine learning to identify key features distinguishing active from inactive states of class A GPCRs that were induced by bioactive ligands [29]. The researchers started with rigidity analysis [30] to obtain rigidity information for a set of ligand-bound and ligand-free GPCR crystal structures in active and inactive states. Next, the researchers partitioned the helical regions of the receptors into 29 segments, where each segment was subdivided into 3 blocks. The 87 blocks were then used to construct feature vectors as input to a machine learning classifier. In particular, each GPCR was represented by a continuous feature vector containing 87 values, each value being a rigidity index (measured on a scale ranging from 0 to 100) for that block. After assembling this tabular dataset, the researchers trained a *k*-nearest neighbor classifier that was able to predict active and inactive states with high (96%) accuracy. In addition, the researchers gained further insights, namely the identification of six key flexibility transition regions, which are involved in activating (or inactivating) the GPCR upon ligand binding. We hypothesize that if protein–ligand interactions are known or can be reliably predicted then this type of analysis can aid virtual screening of bioactive GPCR ligand as well as GPCR ligand design.

1.3. Augmenting ligand discovery with artificial intelligence and machine learning

Artificial intelligence (AI) is traditionally considered a subfield of computer science that focuses on tasks that humans are naturally good at, such as image recognition and natural language processing. Furthermore, AI can be categorized into artificial general intelligence,

i.e., human-level intelligence on a breadth of different tasks, and “narrow” AI, which focuses on solving a single task well. Examples of narrow AI include tasks such as classifying objects from images, language translation, playing chess, or predicting the bioactivity of small molecules. Currently, the most popular approach for implementing and programming narrow AI is machine learning. Machine learning is a field that focuses on the development of algorithms that allow computers to learn from representative datasets, as opposed to having domain experts developing rules manually.

The three major subcategories of machine learning are supervised learning, unsupervised learning, and reinforcement learning. In supervised learning, the goal is to predict a category label (classification) or score (regression analysis) by learning from a large collection of labeled examples. In other words, supervised learning is concerned with learning a mapping function between the so-called input features or observations (for example, fingerprint representations of small molecules) and a discrete or continuous target variable, for example, an active/inactive label or the binding affinity in the context of a specific receptor. After learning to predict the desired scores or labels from the labeled training dataset, an independent test dataset, consisting of labeled examples that were unseen during training, is used to evaluate the performance of the predictive model.

To illustrate these supervised learning concepts in the context of GPCR bioactive ligand discovery, we revisit the SLOR1 receptor signaling inhibitor projects [24,25] introduced in Section 1.2.1. Suppose the goal is to train a classification model to predict whether a (new) candidate molecule is active against SLOR1. To train such a classifier, we require a labeled dataset, which contains the target variable we want to predict. The target variable, in this case, is a binary variable with two possible values, *active* and *non-active*. Ideally, the target labels are obtained from experimental assays, for instance, by assaying a set of candidate molecules from a molecular database for bioactivity against the target receptor. (Note that it is also possible to model this prediction

problem as a regression task, where the target variable is a continuous activity value; however, in our experience, classification models are less sensitive to noise, easier to fit, and more robust overall.) Next, we have to decide about the feature representation of the molecules that are presented to the machine classifier along with the target labels. For example, we may use the functional group matching vectors based on 3D volumetric overlays with a known receptor agonist as described in Section 1.2.1 and [25]. Before we start training the classifier, we divide the dataset into a separate training and test set, where the training set is used to fit the model. After model fitting, the test set can be used to evaluate the classifier's performance using molecules that it has not seen during training. During the evaluation, the model is only presented with the feature vector representation of the molecules. The predicted activity labels are then compared to the experimentally measured activity levels to quantify predictive performance. Since model evaluation is an important aspect of machine learning, which we cannot cover in detail here, we recommend consulting the article "Model evaluation, model selection, and algorithm selection in machine learning" [31] to learn more about the best practices. If the model has achieved satisfactory performance, we may use it on a new database of molecules to predict whether they are likely to be active or not. This process is summarized in Fig. 2.

Unsupervised learning is used for tasks that do not involve target labels. A typical example of unsupervised learning is clustering, for instance, grouping small molecules based on a user-specified similarity measure.

Reinforcement learning, a third major category of machine learning, focuses on decision making in complex environments, which requires

learning an optimal series of actions in response to environmental information as opposed to a target output as in supervised learning. A classic example is the development of a chess program.

In the presence of labeled high-quality data, machine learning provides opportunities for automating the development of customized solutions that are optimized for a specific type of receptor, binding site, or small molecule – as opposed to more general scoring functions such as DrugScore [32] or AutoDock Vina [33]. Machine learning has already shown to be successful in various stages throughout typical ligand and drug discovery, including the discovery of novel drug targets and ligands [34,35], bioactivity prediction [36], predicting new binding sites in GPCRs [37,38], analyzing the association between the drug targets and diseases [39], studying binding pathways (for example, the binding pathways of opiates to μ -opioid receptors [40]), optimizing lead compounds [41], molecular *de novo* design [42], modifying molecular properties [43], and developing biomarkers to assess the efficacy of drugs [44].

1.4. Deep learning–representation learning with deep neural networks

In the last decade, deep learning research has seen a substantial increase in attention, since it has allowed researchers to develop state-of-the-art solutions for computer vision, natural language processing, and computational biology [45]. Deep learning is a subfield of machine learning that focuses on artificial neural networks with many layers that can learn multiple levels of abstractions of data that are useful for supervised and unsupervised learning tasks. Modern deep learning architectures consist of hundreds of millions of parameters [46], the so-

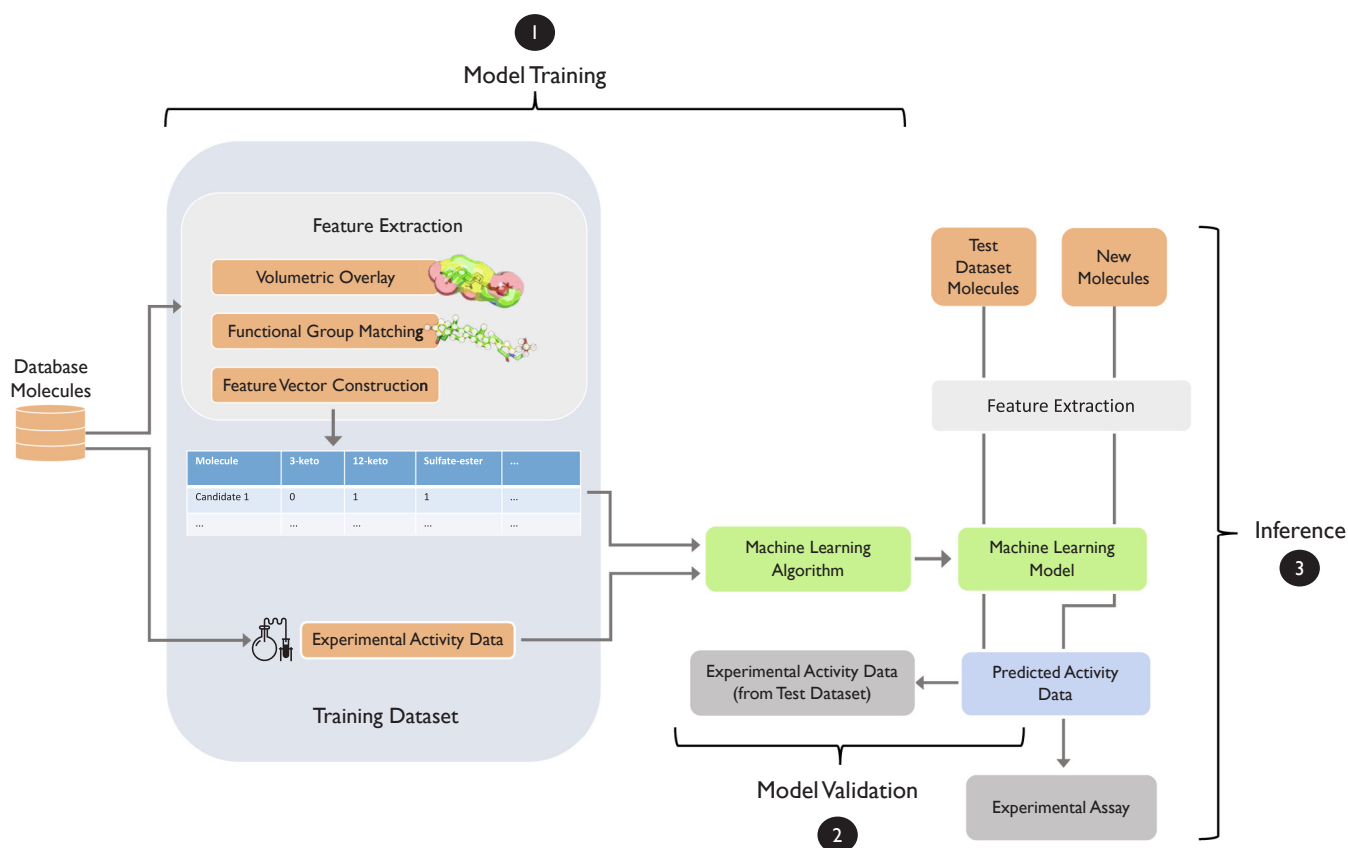


Fig. 2. Illustration of a supervised learning workflow for bioactive ligand discovery summarizing the GPCR inhibitor discovery described in [24,25]. Step 1: After extracting feature vectors for each candidate molecule, the training dataset consisting of bioactivity measurements and the feature vectors is used to train the machine learning model. Step 2: The model is evaluated on the test dataset, which contains molecules (and their experimental bioactivity measurements) to validate the model. Step 3: If the model yielded satisfactory validation performance, it could be used to predict the activity of new molecules and prioritize candidates for experimental assays.

called model weights, which are learned from a training dataset using the backpropagation algorithm [47]. An overview of the most commonly used deep learning architectures is provided in Section 1.7.

The predictive performance of traditional machine learning methods – this includes logistic regression, random forests, support vector machines (SVMs), k-nearest neighbors, and many others – heavily depends on the design of the feature engineering pipeline. For instance, the aforementioned conventional machine learning methods generally do not operate well on high dimensional datasets and are unable to extract knowledge from raw data (such as text or images) [45]. Hence, the training data has to be converted into a tabular format via manual feature engineering, where *manual* means that the feature extraction is not implicitly realized by the model but has to be carried out beforehand via an additional procedure. For example, in a virtual screening context, the two- or three-dimensional graph representations of a molecule represent *raw* data, whereas simplified molecular-input line-entry system (SMILES) string or fingerprint representations of a molecule are the outcomes of manual feature extraction. Different types of molecular representations will be covered in more detail in Section 2.

Careful feature engineering requires substantial domain expertise, and useful information contained in the data can be lost by feature engineering. Deep learning, on the other hand, relies on general-purpose algorithms that include the automatic extraction of salient information from raw data as part of the modeling architecture and optimization objective. In this respect, deep learning can be characterized as a feature or representation learning method. However, one downside of deep learning is that it is relatively data-hungry, and datasets for supervised learning require large amounts of labeled examples – dataset sizes ranging between 50 thousand and 15 million training examples are not unusual [48,49].

1.5. Open source software and datasets

Over the past decade, there has been a remarkable increase in the development of open source software for enabling data science and machine learning [50]. Many general-purpose scientific computing libraries with permissive open source licenses are now widely used in academia as well as in industry, including NumPy and SciPy [51], Matplotlib [52], and Pandas [53]. Furthermore, general-purpose libraries have been developed for the analysis of biological and structural data that lower the barrier of entry to computational biology, for example, BioPython [54] and BioPandas [55].

Under similar open source licensing terms, the Scikit-learn machine learning library [56] has been widely adopted for predictive modeling with traditional machine learning (for example, generalized linear models, SVMs, and tree-based methods such as random forests and gradient boosting). In recent years, computationally efficient and GPU-enabled libraries such as TensorFlow [57] and PyTorch [58] made deep learning accessible to the broad scientific communities similar to how Scikit-learn contributed to popularizing machine learning.

While the use of machine learning and deep learning has seen widespread adoption throughout various research areas, a bottleneck for applications to computational biology was the lack of large, annotated, high-quality datasets. However, recent years brought both improvements of experimental techniques and the “data and code sharing” culture in academia, which led to the increase of publicly available molecular activity and biomedical datasets. For instance, in drug discovery, recent efforts focused on developing an open, annotated dataset that can be utilized for therapeutic target validation [59]. MoleculeNet is a large benchmark dataset consisting of more than 700,000 molecules and their property annotations [60]. The ChEMBL large-scale bioactivity database contains more than 5.4 million bioactivity measurements for 5200 protein targets and more than 1 million ligands [61]. While the aforementioned databases are general-purpose datasets that are usually used for general algorithm development and benchmarking, they can also be used as a starting point for pre-training

deep learning models that can then be fine-tuned to smaller datasets for specific targets. This technique is called transfer learning and is discussed in greater detail in Section 6.

1.6. Interpretability and repeatability of machine learning

Machine learning models can learn to identify salient patterns in high-dimensional and complex datasets that are not obvious to a human researcher. However, machine learning and particularly deep learning models are often criticized for their black box aspect and lack of interpretability [9]. We agree that machine-learning-based methods, except for the most fundamental decision tree models and rule-based classifiers, are less interpretable than pure “if/then” rules. However, many methods exist that allow scientists to obtain insights into which features a model learns for making particular predictions. For example, in the SLOR1 context discussed in Sections 1.2.1 and 1.3, the researchers were able to identify a potent inhibitor of GPCR pheromone receptor signaling using a machine learning-aided virtual screening approach that provided structure–activity information [24]. Furthermore, by combining machine learning with feature selection algorithms, the researchers were able to identify sulfate groups on the tail end of the candidate ligands that are crucial for bioactivity in SLOR1 GPCR signaling [25]. While the methods described in [25] were applied to summarize structure–activity relationships of 3D-structural overlays of small molecules that were prioritized for experimental bioassays, the same techniques can be used for different types of data, such as hydrogen bonding patterns [62] or ligand-receptor docking poses [63].

Other commonly used methods for understanding model predictions are general-purpose methods such as LIME [64] or SHAP [65]. Neural network-specific methods include guided backpropagation [66], class activation mapping [67], gradient-weighted class activation mapping [68], and learning important features through propagating activation differences (DeepLIFT) [69]. All of these methods can be used to study how changes in the inputs affect the model’s output (predictions). Hu et al. [70] demonstrated how a simple variant of this concept could be applied to identify binding sites from 1D sequence data. Assuming that 3D structures are unavailable or too costly to obtain, the researchers trained deep neural networks on 1D representations of proteins (1-letter amino acid codes) and ligands (SMILES strings) to predict binding information. In particular, the researchers trained convolutional neural networks (see Section 1.7.2) for binary classification (binding vs. non-binding) and regression (binding affinity) on one-hot encoded versions of the 1D sequences as illustrated in Fig. 3. After model training and evaluation, parts of a protein sequence were masked to analyze how their removal impacted the prediction score. If the masking of a subsequence had a significant effect on the predicted output, one could hypothesize that the subsequence was important for binding.

Another point of criticism against the use of deep neural networks is the lack of repeatability [9]. In this context, repeatability refers to the ability to produce the exact same results upon repeating an experiment under identical conditions; in contrast, reproducibility refers to the ability to obtain similar results in different environments or conditions, for example, if a different research lab conducts similar experiments. According to [9], the issue of repeatability “arises because [machine learning] outputs are highly dependent on the initial values or weights of the network parameters or even the order in which training examples are presented to the network, as all of them are typically chosen at random.” We want to highlight that these repeatability issues can easily be circumvented by specifying the seeds of pseudo-random number generators and setting the behavior of machine learning and deep learning software libraries to deterministic. In practice, however, issues with reproducibility and repeatability arise when researchers share insufficient details about the software versions that were used to conduct the experiments. Reproducibility and repeatability require best code sharing practices, which include not only sharing the exact code but also the experimental settings and software version numbers.

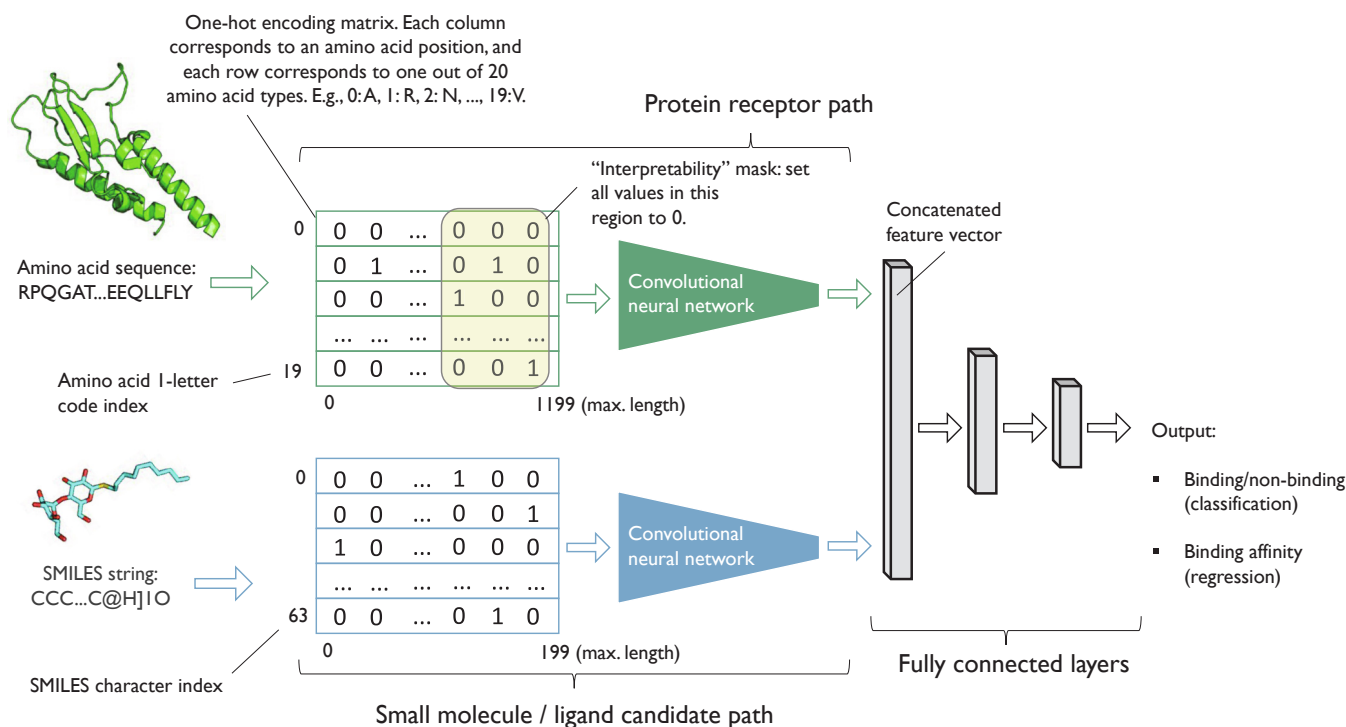


Fig. 3. Illustration of a convolutional neural network trained on 1D representations of protein receptors and ligand candidates. The column vectors in the one-hot encoded matrix are sparse, consisting of one “1” (indicating the type of amino acid or SMILES character at a given position), and the remaining values are “0”s. After training, an “interpretability” mask, which sets all values in the masked region to 0, can be used to observe how blocking out subsequences affects the binding prediction. (The crystal structure corresponds to the ligand-bound glucagon-like peptide-1 receptor extracellular domain [71], PDB code: 3C5T).

Fortunately, sharing data and code has become a recommended practice in the field of machine learning [72], and we hope that computational biology journals will adopt similar best practices such that repeatability and reproducibility issues that plagued the field in previous years can be avoided in the future.

1.7. Commonly used deep learning architectures

This section summarizes the main concepts behind machine learning with a focus on the fundamental deep learning concepts that are relevant for the remainder of this article. For an introduction and overview of general machine learning methods, we recommend [50,73]. For a more comprehensive introduction to deep learning, we recommend [45,50,74].

1.7.1. Multilayer perceptrons

Multilayer perceptrons (MLPs) are fully connected artificial neural networks (NNs) that consist of an input layer, an output layer, and at least one hidden layer between the two (Fig. 4 A). While there is no precise definition of what constitutes a *deep* neural network (DNN), an NN that has more than one hidden layer is commonly referred to as a DNN. The hidden units in the hidden layer manipulate the input information (observations or features) in a non-linear fashion so that, in the case of supervised classification, the training examples from different classes become linearly separable by the last layer. In DNNs, the early layers can be considered representation learning layers that distort the data in such a way that it can be classified by an output layer that has a similar structure to a generalized linear model. In other words, NN architectures with one or more hidden layers can learn highly complex relationships between the input features and the target label.

An MLP with only one (sufficiently large) hidden layer can already be considered a universal function approximator [75–77]. However, a

DNN with many hidden layers can achieve the same expressiveness with fewer parameters than an NN with only one hidden layer. Furthermore, constructing an architecture with multiple layers provides some form of regularization: later layers are constrained by the behavior of earlier layers. Unfortunately, a common problem with deep architectures is that increasing the number of layers exacerbates the vanishing and exploding gradient problems that arise from repeated multiplications via the chain rule in backpropagation, making it hard to parameterize very deep models. A particular focus in the deep learning field has been on developing methods that help with training very deep architectures successfully, as discussed in Section 1.7.2 and Section 1.7.3.

While the early MLP architecture was first proposed in the early 1960s [78–80], efficient ways for training such multilayer neural networks using the backpropagation procedure were formulated by several researchers independently between 1970 and 1986 [47,81,82]. Although the backpropagation procedure is still the main learning algorithm in deep learning, the training of DNNs has only become broadly feasible in recent years due to several advances towards making the general training more efficient and effective. These improvements include implementing machine learning models and training DNNs on GPUs [83–86], which are extremely well-suited for performing linear algebra operations such as matrix multiplications efficiently and on a large scale.

In addition to developing software libraries that allow researchers to develop and implement DNNs more easily (see Section 1.5), countless algorithmic advances towards training DNNs more efficiently and robustly have been made in recent years. For example, new activation functions have been proposed, such as the rectified linear unit (ReLU; Fig. 4 C) [87], which has become a popular default choice since it can help with vanishing gradient effects and generally allows for training DNNs faster and with better predictive performance. New stochastic gradient descent-based optimization algorithms such as Adam [88] can

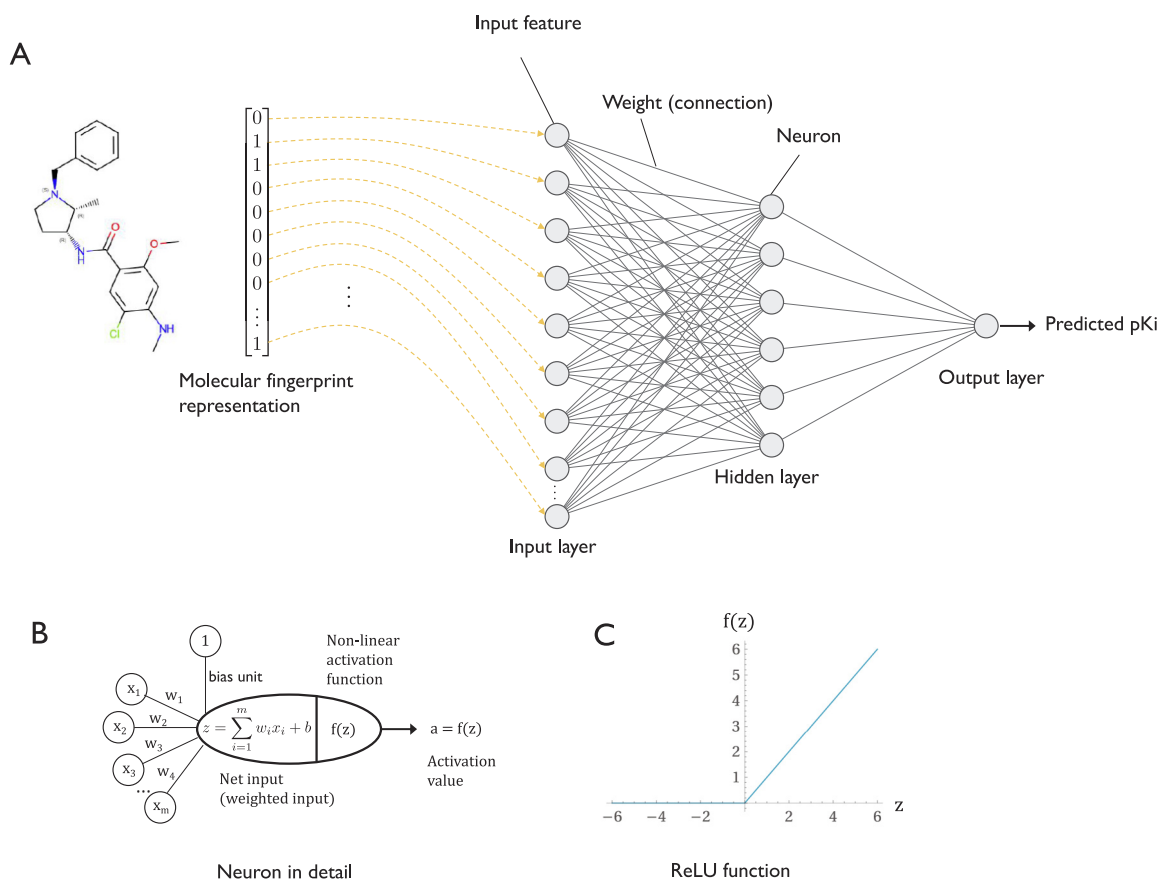


Fig. 4. (A) A fully connected neural network (multilayer perceptron; MLP) with one hidden layer. The input layer represents the input features from the dataset, for instance, a molecular fingerprint representing a small molecule (here: Dopamine D4 receptor antagonist Nemonapride). The input layer is “fully connected” to the hidden layer via weight connections. In this example, the output layer consists of a single neuron that outputs a continuous value, the predicted pK_i . (B) Depiction of a hidden layer neuron in greater detail. The neuron computes the weighted sum of all input features (plus a bias or threshold value) to calculate the so-called net input. This weighted input is then passed to a non-linear activation function, such as the rectified linear unit (ReLU) depicted in (B).

accelerate the model training and make it easier to find a proper learning rate setting compared to conventional stochastic gradient descent. Regularization methods such as Dropout [89] help reduce the effect of overfitting by dropping nodes randomly during training and thus make the model less reliant on particular inputs and connectivity patterns. Batch normalization is a method for scaling layer inputs [90], which can speed up learning further by allowing learning with larger batch sizes and converging to local minima on the loss surface in fewer iterations over the training set. While not all of these improvements exist in all modern architectures, they have been fundamental to allowing experts as well as non-experts to train DNNs successfully on a wide range of different datasets.

Koutsoukas et al. [91] applied a multilayer perceptron model for diverse bioactivity prediction (pK_i and pIC_{50}) against 7 different targets, including two GPCRs, dopamine receptor D4 and cannabinoid receptor 1 based on molecular fingerprint representation of the molecules (similar to the illustration in Fig. 4 A). The researchers also found that the MLP outperformed traditional machine learning approaches on large datasets. However, they noted that deep learning models require much more extensive hyperparameter tuning to achieve good predictive performance [91]. An approach like this can be used for ligand-based discovery, when information about the binding interface is not available.

MLPs can also be utilized for structure-based VS or QSAR studies as described [92]. Similar to [91], the researchers found that MLPs generally outperformed traditional machine learning methods like random forests and SVMs. This is a remarkable result because, in this case, the model the researchers referred to as a DNN was a simple MLP with only

one hidden layer. The descriptors used in this study are a union of atom pair and donor–acceptor pair descriptors, including information about atom types and binding interface information. Additional details about the different molecular data input representations will be provided in Section 2.

1.7.2. Convolutional neural networks

Convolutional neural networks (CNNs) are feedforward neural networks that utilize the discrete convolution operation as a filtering operation on multi-dimensional arrays (Fig. 5 A). Depending on the architectural design and discrete convolution operation, CNNs can handle various data modalities, including 1D (sequences and signals in vector form), 2D (pictures and audio spectrograms), and 3D arrays (such as videos or images with depth dimension like 3D computerized tomography scans).

Today, CNNs are the most widely used deep learning architectures. They are widely recognized for their state-of-the-art performance for image-related tasks such as object classification, detection, and segmentation. The first successful demonstration of CNNs for image classification was the design of the LeNet architecture for handwritten digit and document recognition in the 1990s [94,95]. While CNNs have shown exceptional performance on various image recognition tasks, including the segmentation of biological images [45,96], the rise of CNNs to mainstream success can be attributed to 2012 ImageNet competition, where CNNs were able to outperform traditional computer vision methods by doubling the object classification accuracy on a database consisting of millions of images [97].

Similar to MLP architectures, CNNs are feedforward neural

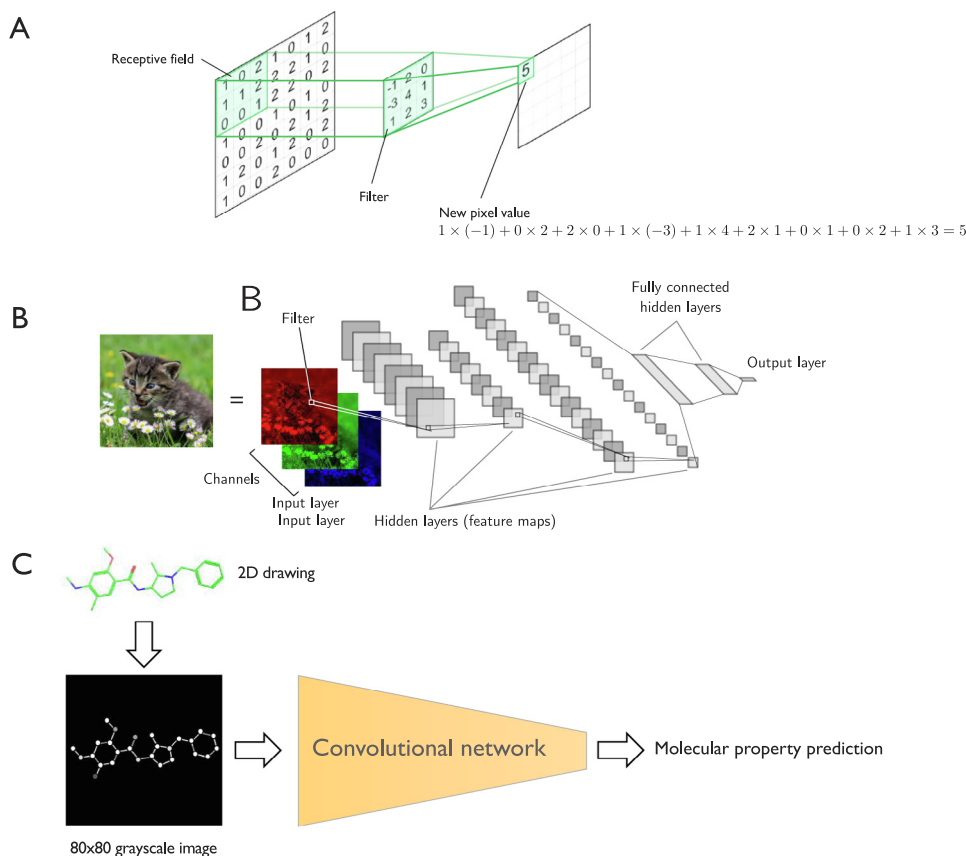


Fig. 5. (A) Illustration of a single step in a 2D convolution on an 8x8 pixel input image, using a filter of size 3 × 3. At each stage, the convolution operation performs a pairwise multiplication of the values in the receptive field and filter. The pairwise multiplication results are then summed to obtain the output pixel. For the next step, the receptive field is moved to the next position (not shown) to compute the next output pixel. This process is repeated until all output pixels are computed. (B) Depiction of a convolutional neural network (CNN) architecture for image recognition, which consists of multiple feature maps that are connected via convolutional layers. Each unit in a feature map is connected with a local patch of the previous layer's feature map through a set of filters. Traditionally, pooling layers have been inserted between every other feature map but are not necessary [66]. The fully connected hidden layers and output layer resemble an MLP architecture. (C) Application of a CNN for molecular property prediction as described in [93].

networks. However, in contrast to MLPs, CNNs rely on convolutional rather than fully connected layers at the core of the architecture. The main concepts behind CNNs that distinguish this architecture from MLPs are *sparse connectivity* and *parameter sharing*. Here, sparse connectivity refers to the fact that each unit in the hidden layer is only connected to a small number of units in the previous layer (Fig. 5 B). The connection occurs through a filter (Fig. 5 A), which is a parameter matrix that is commonly interpreted as a feature extractor. Via the filter, information between neighboring pixels is combined to compute the activation of a unit in the next hidden layer. This allows CNNs to compose hierarchies of local features to recognize complex shapes in later layers. For instance, in an image classification task, a multi-layer CNN may learn to detect simple geometric shapes like edges and curves in the earlier layers, and in later layers, it can identify more sophisticated features such as cars or buildings based on the underlying geometric shapes. These feature detectors typically only cover a small portion of the input image or feature map, and in the process of the convolution operation, they can be thought of as a sliding window over the image. Hence, the weights for the different patches of the input image or feature maps are shared within a given layer, which is inspired by the fact that a feature detector that works well in one region may also work well in another part. The advantage of weight sharing is that it reduces the number of parameters that need to be learned, which improves computational efficiency and can reduce overfitting, compared to fully connected networks.

In traditional CNN architectures, feature maps were followed by pooling layers [94,95]. These pooling layers combine similar features from neighboring regions in the previous feature map into a single feature, which is supposed to help neural networks to learn more complex relationships or geometrical shapes. However, it has been shown that pooling layers are not a requirement in modern CNN architectures [66].

Deep learning is a fast-moving field, and even revolutionary

architectures such as AlexNet, which outperformed state-of-the-art computer vision methods at the ImageNet competition in 2012 [97], have long since been succeeded by other reference architectures. For example, the VGG architecture has achieved notable performance by shrinking the receptive field and stacking many more layers than AlexNet [98]. The ResNet architecture introduced residual connections, which are shortcut connections between layers that help with vanishing gradient problems and thus improve the training of very deep neural networks [99]. Inception networks improved the extraction aspect of both local and global features by using multiple filter modules in parallel [100]. While the design of CNN architectures still remains an active area of research, recent efforts have been targeted towards improving efficiency, for example, MobileNet [101] and EfficientNet [102]. In a benchmark study, Canziani et al. compare both the predictive and computational performance of these models for further reference [46].

Goh et al. described a simple application of 2D convolutional networks for molecular property prediction. [93]. Here, the researchers trained a convolutional network, called Chemception, to predict HIV activity and toxicity (as binary classification tasks, "active" or "inactive"), and solvation (regression) from 2D molecule drawings (Fig. 5). To prepare the input data for the Chemception model, the researchers converted 2D drawings of molecules into 80 × 80 pixel grayscale images. In this representation, atoms were represented as dots with different shades of gray to encode atom type information. Similarly, bonds were encoded in two different shades of gray to distinguish between single and double bonds. The images were randomly rotated during training to make the network robust towards different ways to orient a molecule. The authors noted that Chemception slightly outperforms molecular fingerprint-based approaches in activity and solvation prediction but slightly underperforms in toxicity prediction. State-of-the-art convolutional network approaches, including 3D voxelation, are discussed in Section 2.3.

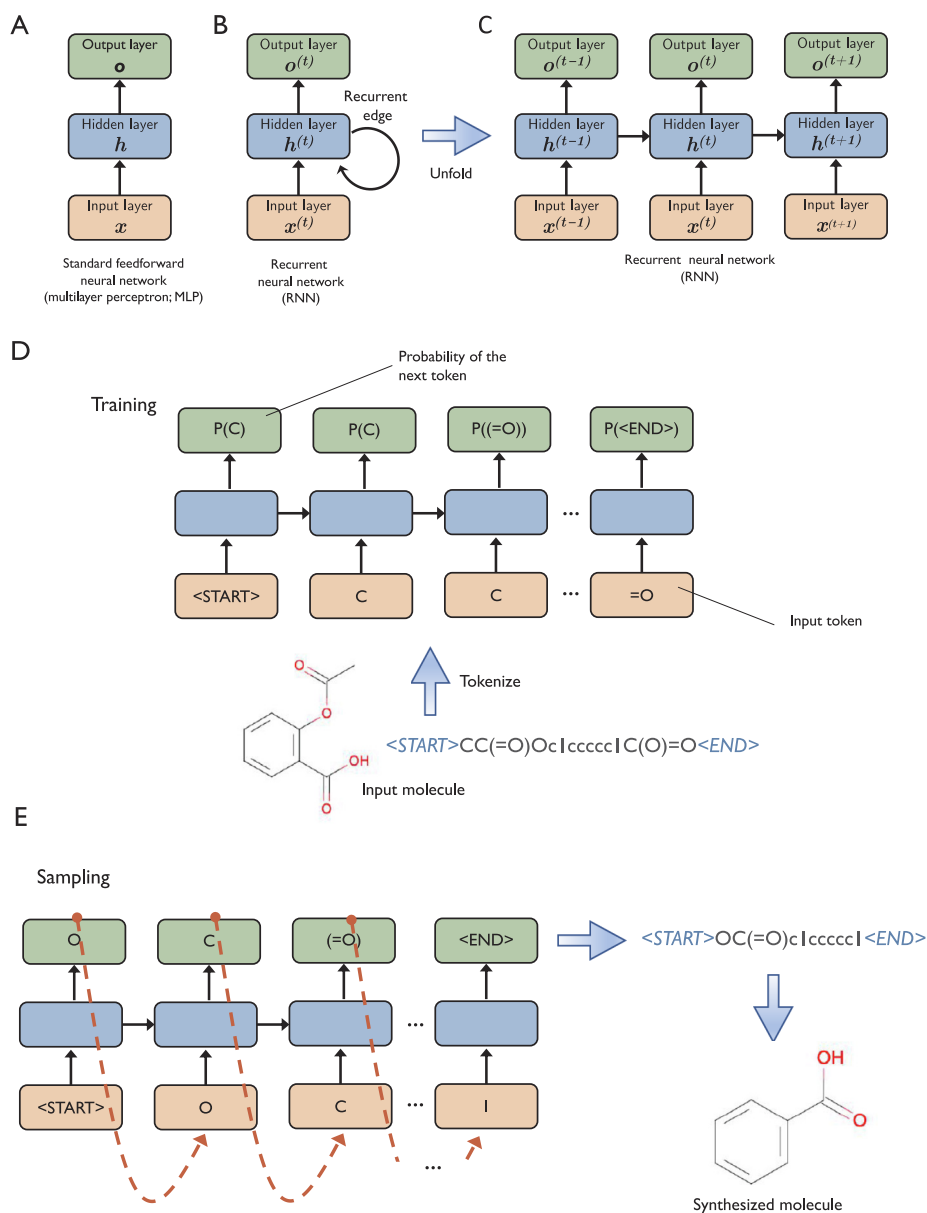


Fig. 6. Subpanels (A) and (B) depict an MLP with 1 hidden layer and an RNN with 1 hidden layer side by side for comparison. Here, the input layer x containing the input features, the hidden layer h , and the output layer o represent vectors with multiple units. Subpanel (C) depicts the RNN in its unfolded representation. In contrast to an MLP, the hidden layer of an RNN receives its inputs from both the input layer and the hidden layer of the previous time step, $t - 1$. (D) Illustration of the RNN training process. Input molecules are converted into a sequence representation (here: SMILES string), which is divided into tokens. At each time step, the RNN receives one token and returns a probability distribution over all possible token types. The RNN is trained such that the probability corresponding to the next token in the sequence is maximized. (E) To synthesize new molecules after training, a token is randomly sampled from the probability distribution at each time step (tokens with a higher probability are more likely to be drawn). The sampled token is then provided as input at the next time step. When the network outputs an $<END>$ token, the sampling is terminated. The generated SMILES string can then be used as a template for synthesizing the new molecule.

1.7.3. Recurrent neural networks

Recurrent neural networks (RNNs) were originally designed for tasks that involve sequence data, for example, textual data and audio signals such as speech. While traditional RNNs consist of fully connected layers similar to MLPs, RNNs have recurrent connections through time. This recurrence property allows RNNs to maintain an internal history of sequence elements that it processed previously (Fig. 6).

While RNNs can model arbitrary long sequences, the recurrent edges make this architecture particularly prone to vanishing and exploding gradient problems [103]. One solution to this problem is the so-called long short-term memory mechanism (LSTM) by Hochreiter and Schmidhuber [104], which is a memory cell that replaces the hidden

layer of conventional RNNs. A mechanism similar to LSTM, although slightly more computationally efficient, is the gated recurrent unit (GRU) [105].

Merk et al. [106] described an RNN model for designing new drug-like molecules with desired properties. The RNN was first trained on a large dataset consisting of SMILES strings of known bioactive molecules. Then, the researchers used transfer learning Section 6 to fine-tune the model on retinoid X and peroxisome proliferator-activated receptor agonists to generate target-specific ligands. The general workflow is summarized in Fig. 6 C–E.

The remainder of this paper is organized as follows. First, we discuss one of the common challenges with applying machine learning to bioactive ligand discovery, namely choosing an appropriate feature

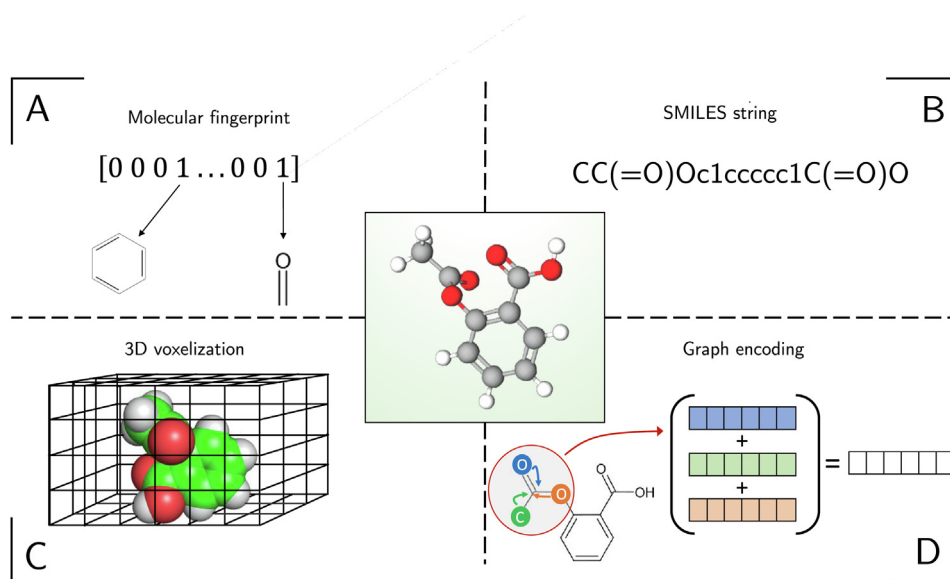


Fig. 7. Summary of commonly used molecular representation methods based on the example of Aspirin (shown in the center). (A) A molecular fingerprint encodes structural motifs into a sparse bit vector. (B) A SMILES string encoding structural information of the molecule as well as its stereochemistry. (C) A visualization of the 3D voxelization concept. Note that information about which atoms occupy which voxels would be encoded in a 4th dimension which is omitted in this visualization. (D) Illustration of how information is passed to an atom in a simple graph neural network. Note that the graph-structural information will be passed from more distant atoms when the summation is repeated (not shown).

representation for the ligand structure, and highlight recently developed deep learning methods that use these representations. Next, we discuss the latest trends for ligand-based analysis, from molecular property prediction to similarity-based virtual screening. While the ligand-based approaches assume that a high-quality receptor structure is not available, the next section reviews the recent developments for receptor structure-based bioactive ligand discovery. We then explore advances in *de novo* small-molecule design. Finally, this review concludes by motivating the use of transfer learning, which allow researchers to make better use of publicly available data in machine learning and AI-based bioactive ligand discovery.

2. Molecular feature representations

Machine learning methods excel at prediction tasks across multiple disciplines but require careful data preparation as most methods are designed to operate on tabular datasets. The standard data input format is the so-called *design matrix*, where each row represents a new training example, and the columns correspond to the different feature variables, as illustrated in the example in Fig. 2. A common challenge in conventional machine learning is how to prepare datasets as input to machine learning algorithms – in practice, machine learning practitioners have to find a sweet spot between reducing the dimensionality and retaining salient information that the model can learn from. In contrast to conventional machine learning, deep learning excels at learning from raw data, such as images and text, directly, as previously discussed in Section 1.7. However, molecular data, such as conformations of small molecules and receptors, can be challenging to represent in a standard format that most machine learning and deep learning methods have been designed for. Even if the same information can be extracted from two different data representations, an algorithm may be more effective at extracting that information from one over the other. There is no clear best representation of molecules for machine learning methods and indeed certain representations may be better for certain tasks. The following section provides a brief overview of commonly used molecular representations as well as some recent applications of them using AI-based methods.

2.1. Property-based feature vectors

A molecular descriptor is the transformation of chemical information into a numeric value [107]. Dragon [108] and Mordred descriptors [109] are examples of sets of molecular descriptors. As an alternative to

molecular descriptors, molecular fingerprints encode molecular structure in a vector format, a so-called bit vector consisting of 1's and 0's. When used as input for machine learning models, both molecular fingerprints and descriptors have historically produced state-of-the-art results on chemical machine learning tasks such as chemical odor prediction and bioactivity [41,110].

The extended connectivity fingerprint (ECFP) is among the most widely-used 2D fingerprint methods [111], and we use its generation procedure as an example of the general process for generating traditional molecular fingerprints. A fingerprint is generated by a multistep process in which each atom is associated with a series of integers. In this series the k th integer encodes information about the atom it is associated with as well as information about the atoms and bonds within k bonds of that atom – that is, the substructure of the compound that is within k bonds of the atom. Next, the integers associated with each atom are concatenated into an array format, which is then processed via a hashing algorithm to generate a bit vector of a desired length (typically 1024 or 2048 elements). This method captures information about all identified substructures in a compound, resulting in a fixed-length vector regardless of the input compound's size. ECFPs do not explicitly encode the 3D spatial information of a compound; however, specialized fingerprint methods have recently been developed that incorporate 3D-structural information [112]. Lastly, there are also fingerprints that can encode protein–ligand interactions [113].

2.2. SMILES

Simplified molecular-input line-entry system (SMILES) strings are ASCII string representations of compounds (Fig. 7 A), which are generated according to a procedure that guarantees a unique mapping from a SMILES string to a compound structure (though not the inverse) [114]. One benefit of SMILES strings over 2D molecular fingerprints like ECFP is that they encode stereochemistry explicitly. One downside for machine learning is that SMILES do not have a fixed-length; however, certain deep learning architectures designed for processing text documents, like RNNs or 1D CNNs, can handle variable-length inputs.

Recently, Hirohara et al. [115] proposed a novel molecular representation scheme by converting SMILES strings into “SMILES feature matrices,” which were used as inputs into a 1D CNN [115]. A SMILES feature matrix was constructed by mapping a SMILES string of length N to a $N \times 42$ matrix, where the k th row represents the k th character (corresponding to either atom or connectivity information) of the string, and the 42 columns correspond to properties of that character.

To address the problem of varying-length input strings, the feature matrix was padded with rows of zeros such that all feature matrices had their number of rows set to the length of the longest SMILE string. The predictive performance of the resulting CNN model was comparable with other deep learning methods on the Tox 21 dataset, which contains 8,000 compounds labeled as active or inactive for 12 proteins [116]. More interestingly though, this novel feature representation method enabled the extraction of a 64-dimensional vector from a convolutional layer referred to as the *SMILES convolution fingerprint*, which can be mapped back to the model input SMILES, providing an interpretable data-driven fingerprint.

Another notable recent development is the SMILES2vec method, which combines components from both recurrent and convolutional neural networks [117]. Here, the model input is a SMILES string converted into a one-hot encoding of characters present in the SMILES in the training dataset. The one-hot encoding then serves as an input to a 1D convolutional layer, which is followed by two GRUs before the fully-connected output layer that returns the target value. This architecture achieved equivalent performance to the state-of-the-art at the time on the ESOL solubility dataset [118].

2.3. 3D voxels

Molecular structures can be considered as 3D objects, and modeling them as such would encode their structural properties effectively. The conventional way for encoding 3D objects is by voxelization, which, in this case, takes a 3D space and discretizes it into a 3D grid. Each unit cube of the 3D grid represents a voxel (Fig. 7 C), which can be considered as the equivalent of a 2D pixel in a 3D space. Instead of merely labeling each voxel via a binary membership indicator based on what part of the compound occupies the voxel, a feature vector with discrete or continuous-valued attributes can be assigned to each voxel, which can provide further information about the atom type or charge, for example. This adds an additional dimension to the representation. Since this representation is typically associated with a large number of input features, depending on the resolution of the voxel space as well as the size of the vector representation at each voxel position, deep learning approaches used with this type of input representation typically rely on 3D convolutions. This is equivalent to using 2D convolutions, which are commonly used for image analysis (Section 1.7.2), in the 3D voxel space. Unfortunately, these convolutions can still be prohibitively slow, even at relatively low voxel resolutions, which can result in a coarse representation of the compounds properties. Since each cube represents a unit of measurement, and compounds can vary in size, this representation needs to be padded so that it can accommodate for the largest compound in the dataset, which exacerbates the computational efficiency challenges of this method. Likely owed to being computationally very intensive and inefficient, this method is not a common choice for ligand-based virtual screening. However, this representation has shown to be successful for structure-based VS where binding pocket size is consistent regardless of the interacting ligand, and capturing the 3D structure of the protein–ligand complex is important for the task.

We examine the 3D voxelization process from Ragoza et al. as a concrete example [119]. In this publication the authors create 3D voxel representations of high-scoring docking poses of protein–ligand complexes for use in a 3D CNN that computes binding affinities. This representation voxelizes a 24 Å cube centered on the docked ligand into 0.5 Å voxels. Each voxel has channels for Smina atom types in the ligand and separately for Smina atom types in the protein [120]. The value contributed to a channel by an atom of the type that is associated with that channel is based on a function of two values: First, the distance between the center of the voxel and the center of that atom, and second, the atom’s van der Waals radius. This function is a continuous piece-wise combination of a Gaussian and a quadratic based on these two values.

2.4. Graph representation

Molecules are commonly visualized as undirected graphs. In this representation, atoms represent the graph’s nodes, and the bonds are the graph’s edges. A naive approach that utilizes such graphical information is to consider pictures of molecular graphs as inputs to a DNN. This has been tried in the literature [93] using 2D CNN architectures that have been effective for image classification [93,121]. However, these models do not outperform MLPs and random forests trained on ECFPs. Two issues with this approach are that atomic properties and spatial relationships must be inferred implicitly and that the representation is sparse, with lots of white space that provides little chemically relevant information to the CNNs. Both these issues can be addressed with graph neural networks (GNNs) [122,123].

GNNs operate directly on a molecular graph (Fig. 7 D). Similar to images, graphs can encode local structure, but the local structural information is based on the graph structure rather than Euclidean distance. Like CNNs, GNNs utilize sparse connectivity and parameter sharing, but the connectivity is based on the graph’s structure. For a detailed explanation of graph convolutional operators, we refer to [123]. On a high-level, graph convolutions are based on message passing frameworks for GNNs. At each node of the graph, the following steps are performed: a *message* function is applied to each of a nodes’ neighbors individually, and the outputs are then added. An *update* function is then applied to the summed message functions and the current node with the output being the updated value for the current node. After graph convolutions are performed, the information in the graph can be aggregated by a *readout* function, which produces the networks’ prediction output.

After choosing an input representation, the training process and evaluation procedure of machine learning models tend to follow a general workflow similar to the illustration in Fig. 2. While many state-of-the-art machine learning methods are introduced in the context of general molecular benchmark datasets, most applications of these models can be mapped to this general workflow and adopted for GPCR bioactive ligand discovery. For many machine learning methods discussed in the following sections, and especially for DNNs, there are numerous training parameters to tune and many approaches for doing so. We do not focus on these approaches but refer to the respective publications that discuss them further. Similarly, there are many approaches to evaluating models that are beyond the scope of this review. Instead, the focus of this article is on the different workflows enabling bioactive ligand discovery and advances in GPCR machine learning models to help internalize the high-level approach to research in this field.

3. Ligand-based methods

Ligand-based VS methods are traditionally defined as methods that only rely on physicochemical and structural information about the ligand and sometimes measurements of ligand-receptor interactions. No other information about the target, like the receptor structure, is utilized. The use of ligand-based methods is particularly appealing when working with GPCRs for which only a small set of high-quality structures exist.

In the following subsection, we review notable advances in predicting molecular properties, which are applicable to GPCR bioactive molecule VS.

3.1. Molecular property prediction

In the following subsection, we explore recent work done in predicting molecular properties that has been applied to GPCR-related properties, and explore advances in the space that have been or could be applied to GPCRs.

3.1.1. Predicting bioactivity from conventional fingerprint representations

Researchers applied a variety of standard machine learning algorithms to the task of classifying compounds as active or inactive with two GPCRs, cannabinoid receptor 1 and cannabinoid receptor 2 [124]. In addition to the activity prediction for cannabinoid receptor 1, the researchers also predicted if compounds acted in an allosteric or orthosteric fashion. The datasets used in this study were constructed from multiple chemical databases. Orthosteric ligands were selected from ChEMBL [61] based on an inhibitory constant (K_i) cutoff. Allosteric ligands were selected from the Allosteric Database [125]. The ZINC database [21], a collection of commercially available chemical compounds, was used to add additional inactives. The authors generated three different vector representations of this ligand data: ECFP fingerprints, MACCS fingerprints, and a set of handpicked molecular descriptors. Stratified subsampling was used to create the training and test datasets; stratified subsampling is a procedure that maintains the ratio of active to inactive, or orthosteric to allosteric, compounds across sets. The data was then used to train and compare the performance of a set of machine learning classifiers on each task with each of the three data representations. The set of machine learning classifiers consisted of an SVM, MLPs with 1–5 hidden layers, a random forest, a naïve Bayes classifier, and a logistic regression classifier. There was no clear top-performing model across data representations when the models were evaluated with a variety of performance metrics. However, for individual data representations, an MLP with only one hidden layer performed best overall with the MACCS featurization, while logistic regression performed best in combination with the ECFP featurization. This all serves as an example of the standard machine learning workflow being utilized in GPCR research (Fig. 2). In addition to evaluating the predictive performance of the different models, the researchers also evaluated the importance of the features in each representation. This was done by recursively removing the features of least importance to a model and retraining models until a user-specified number of features was reached. The researchers then examined how important features differed among ligands with different properties. For example, with MACCS features, which represent molecular substructures, it was found that around 10% of the allosteric ligands had an amide group attached to an aromatic substructure, while this was only the case for 1% of orthosteric ligands. This type of analysis could be applied to other GPCRs, to elucidate further what ligand components are important for certain interactions.

3.1.2. Learning fingerprint representations

The previous example uses standard machine learning models that take a vector representation of a ligand as input. This allows feature importance to be assessed in ways that would not be readily possible with DNNs that operate on other types of input representations, such as images or graphs. However, most models that achieve state-of-the-art performance in molecular property prediction incorporate deep learning along with novel ligand representations. One such model is WDL-RF, an innovative architecture for predicting molecular properties that was trained and tested on a suite of GPCR ligand datasets to predict ligand bioactivity [126]. The novel architecture of the WDL-RF model consists of a DNN that takes a graph representation as input (Section 2.4) followed by a random forest regressor that takes the DNN output as input. The NN is trained to produce an embedding vector the authors refer to as a “data-driven molecular fingerprint.” More specifically, this DNN is a GNN (Section 2.4) that shares weights in the network based on the number of bonds attached to each atom node. Each atom’s initial feature vector includes a one-hot encoding of its element, the total number of atoms connected to it, the number of attached hydrogen atoms, its implicit valence, and an aromaticity indicator. After a series of graph convolutions, the GNN aggregates the graph embedding into a vector, the data-driven molecular fingerprint, which is used as input for the random forest that is trained to predict the molecular property of interest.

The performance of WDL-RF was evaluated on twenty-six GPCRs from families A, B, C, and F. For each GPCR, at least two hundred known ligands were collected from the Uniprot and GLASS databases [127,128]. The bioactivity data was obtained from ChEMBL [61]. In addition, the authors also included known GPCR ligands from ChEMBL that do not interact with the 26 GPCRs evaluated in this study to improve the model’s robustness. This aspect of their workflow is worth highlighting because when applying machine learning to GPCRs in general, data quantity is often a limiting factor for using more powerful methods. After splitting the dataset into training and test datasets for each GPCR, the WDL-RF was trained in two steps. First, a GNN was trained to predict ligand bioactivity. Secondly, the fingerprint embedding produced by the GNN was used to train a random forest regressor to predict bioactivity. The second step was necessary to allow fair comparisons between WDL’s data-driven fingerprint representation and other popular molecular fingerprints since random forests operate on feature vectors rather than graphs. The WDL-RF model, evaluated via the root mean squared error between predicted and actual bioactivity measures, significantly outperformed hand-crafted molecular fingerprints like ECFP and MACCS on the majority of GPCRs. Additionally, WDL-RF consistently outperformed other neural network-based fingerprints [122]. Notably, the few cases where a traditional molecular fingerprint outperformed WDL-RF corresponded to GPCRs for which only a small number of active ligands were available, close to the minimum 200 required for inclusion, which suggests that the datasets were too small to train DNNs effectively.

3.1.3. Predicting molecular properties beyond bioactivity

GNNs have also found success in predicting other molecular properties beyond bioactivity. For instance, they have achieved state-of-the-art performance in odor descriptor prediction [129]. Odor perception involves 300 to 4000 different types of olfactory receptors, which are rhodopsin-like (class A) GPCRs [130]. The odor prediction GNN used *atom-node* vectors based on atom type, charge, etc. One unique aspect of this model was that it was trained to predict 138 odor descriptors (fruity, stinky, etc.) for each molecule simultaneously, as the authors found that the correlation structure between certain descriptors helped model performance. A similar simultaneous prediction approach could be taken with other molecular properties that are relevant to GPCRs if they are believed to have correlation structure. Similar to WDL-RF, when an embedding within the GNN was used to train a random forest model it was found to outperform traditional fingerprints and molecular descriptors.

The above result suggests that different compound representations capture chemical information in ways that are utilized more or less effectively by various machine learning algorithms. If a model takes multiple representations as input, the ensemble may have better performance. The model CheMixNet followed this approach and used both a SMILES and fingerprint representation of a molecule as input [131]. The authors considered four subnetworks, three of which received a one-hot SMILES string as input and one that was trained on a MACCS fingerprint representation. Each of the three SMILES subnetworks consisted of GRU cells, 1D convolutional layers, and a combination of 1D convolutional layers and GRU cells; the fingerprint subnetwork was an MLP. The researchers trained models with combinations of these subnetworks, where the outputs of each subnetwork were concatenated and passed through additional fully connected layers before a linear or softmax layer to predict the property. The model was trained and tested multiple times to predict a variety of molecular properties. In all cases, certain subnetwork permutations outperformed or matched other state-of-the-art methods at the time. Interestingly, the best subnetwork ensemble for each dataset was not consistent, suggesting that different representations may be better at predicting different chemical properties.

3.1.4. Leveraging unsupervised data for fingerprint learning

One issue with the data-driven fingerprints discussed thus far is that they are generated with labeled training data, which restricts the amount of data that can be used to construct a robust fingerprint. The Seq2Seq model attempted to address this issue by learning a fingerprint from SMILES in an unsupervised manner [132]. This model was an autoencoder that used a GRU-based RNN to encode SMILES into a fixed-length embedding vector and then recover the original SMILES from the embedding. After the model was trained, the embedding vectors for ligands in a labeled dataset were generated to use as data-driven fingerprints in other machine learning models. Experiments showed that random forest, gradient boosting, and SVMs that were trained on the Seq2Seq fingerprints outperformed the same models trained on ECFP and the original neural fingerprint produced by Duvenaud et al. [122]. One additional benefit of fingerprints generated by Seq2Seq is that they are invertible, so it is easy to recover the original SMILES string if given the fingerprint.

Shortly after the publication of the Seq2Seq approach, researchers developed a variational autoencoder (VAE) for encoding SMILES into a continuous representation in the latent space [133]. This method also produces invertible fingerprints, and a regression model can be trained on these embeddings to predict various molecular properties. The use of a VAE is the game-changer here, as molecules with similar properties are encouraged to be closely together in latent space. As a result, slightly perturbing the embedding of a ligand and inverting it will tend to produce a different but similar ligand. This property is of interest in *de novo* synthesis which is discussed in Section 5.

3.2. Similarity-based virtual screening

Similarity-based screening methods are based on the hypothesis that active compounds share similar properties [134]. However, there is no trivial definition of compound similarity as computational methods work with approximate representations. For this reason, similarity-based methods heavily depend on and are defined by the molecular representations they operate on. Most machine learning-based methods for ligand-based VS are clustering algorithms, and recent advances in this research are focused on developing new data representation methods and similarity measures. Several examples of GPCR-specific VS in the literature [135,136]. However, recent methods were not specifically on GPCRs but rather evaluated on common benchmark datasets that include several GPCR targets. Nonetheless, their general nature could allow researchers to utilize these methods in GPCR-specific contexts in future research endeavors.

3.2.1. 2D similarity measures

Tanimoto similarity is commonly used when comparing molecular fingerprint vectors (Fig. 8 B), as it compares favorably to other similarity metrics on molecular fingerprints, which has been confirmed in analyses based on the sum of ranking differences and ANOVA analysis [137]. Fingerprint similarity-based VS is not a new methodology [138]. While there are more complex 3D approaches, these methods are still used to query large databases due to their computational efficiency. Furthermore, the efficiency of fingerprint comparison allows for more complex clustering protocols to be tractable.

A recent example of this is the Tanimoto-based fingerprint similarity method MuSSEL [139]. MuSSEL generates an ensemble of different fingerprints for a database of compounds which were labeled with a receptor as their target class and their binding affinity for the corresponding receptor. Given a query compound's fingerprint ensemble, a molecule is considered a candidate for a given target receptor class only if it meets a minimum Tanimoto similarity cutoff with at least N_1 fingerprints in the set of known agonists or antagonists. If this requirement is met, the compound must also pass the following criterion for N_2 fingerprints – where N_2 is a number that has to be specified by the user, similar to choosing N_1 . First, the compound must have at least k nearest

neighbors above a new user-specified Tanimoto cutoff, and second, the maximum difference between two of the k neighbors' binding affinities must be below a user-specified threshold. The maximum difference threshold is included to help avoid activity cliffs [140]. For each fingerprint that meets these conditions, the k neighbors' binding affinities are averaged. Finally, for each query compound, the k -neighbor averages are averaged for all the fingerprints in the ensemble that passed the second filter to predict the binding affinity of the query compound. The algorithm showed good performance on a calibration set and correctly labeled compounds in a number of case studies with GPCR targets such as adenosine A_{2A} receptor and dopamine receptor D4. This method could be applied to other GPCRs of interest to search for new actives, provided there are enough ligands that have measured activity with said GPCR.

3.2.2. Comparing 3D representations

While many similarity-based methods that utilize fingerprints like ECFP and MACCS are computationally efficient, one of the issues with such fingerprints is that they do not explicitly capture 3D information about the compounds. This is problematic in situations where a molecule's 3D structure is valuable for assessing a property, such as the steric fit of a predicted pose docked into the receptor's binding site, so methods for comparing 3D representations of compounds are of practical use.

One approach for comparing 3D representations of molecules could be to find an overlay with maximum shape overlap (Fig. 8 C). Raschka et al. recently showed that 3D overlay-based VS can be made computationally feasible on large databases if hypotheses based on prior information about the ligand-receptor interactions can be used as filtering criteria [24]. Using this hypothesis-driven approach, combined with machine learning-based molecular feature importance assessments, the researchers discovered a potent GPCR signaling inhibitor by screening more than ten million commercially available molecules [24,25].

One weakness of using just maximum shape overlap to determine molecular similarity is that it does not utilize electrostatic information. There are multiple analyses that show that electrostatics play an important role in GPCR binding sites, so 3D methods should incorporate them in some capacity [142,143]. Thus, similarity scores utilized in this space are generally based on a combination of structural overlap and the overlap of groups with like properties. Methods like ROCS [141] and WEGA [144] represent molecules as overlapping Gaussian spheres with each atom sphere containing information about its location, atom type, and charge. The maximum scoring overlay is then based on both the structural and chemical overlap. ESIM is a recent method that computes overlay similarity in a different way [145]. The model uses a set of "observer points" in 3D space that are a fixed distance from one target molecule. At each observer point spatial, angular and electrostatic values are computed that capture properties of the target and query molecule from that observer point's perspective. The closer the values for each molecule are when summed over all observer points, the higher the similarity score is. To evaluate the abovementioned eSim method, the authors used the DUD-E dataset [146]. DUD-E is a standard benchmark set for evaluating molecular conformations and docking predictions that includes 102 targets, five of which are GPCR, as well as 22,886 active compounds and their affinities against the targets. Additionally, each active has fifty decoys, molecules that have similar physio-chemical properties but different topologies than the active. An efficient virtual screening algorithm should be able to select actives while avoiding decoys effectively. The performance was evaluated by choosing an active for a target then computing eSim scores for other molecules in DUD-E in relation to that active. From there, a ROC curve could be constructed to assess how robustly the model can distinguish between actives and inactives. Statistical analysis of the ROC area under the curve (AUC-ROC) scores showed that eSim significantly outperformed other 3D similarity methods, including mainstays like ROCS and WEGA, on many of the DUD-E targets and was only rarely

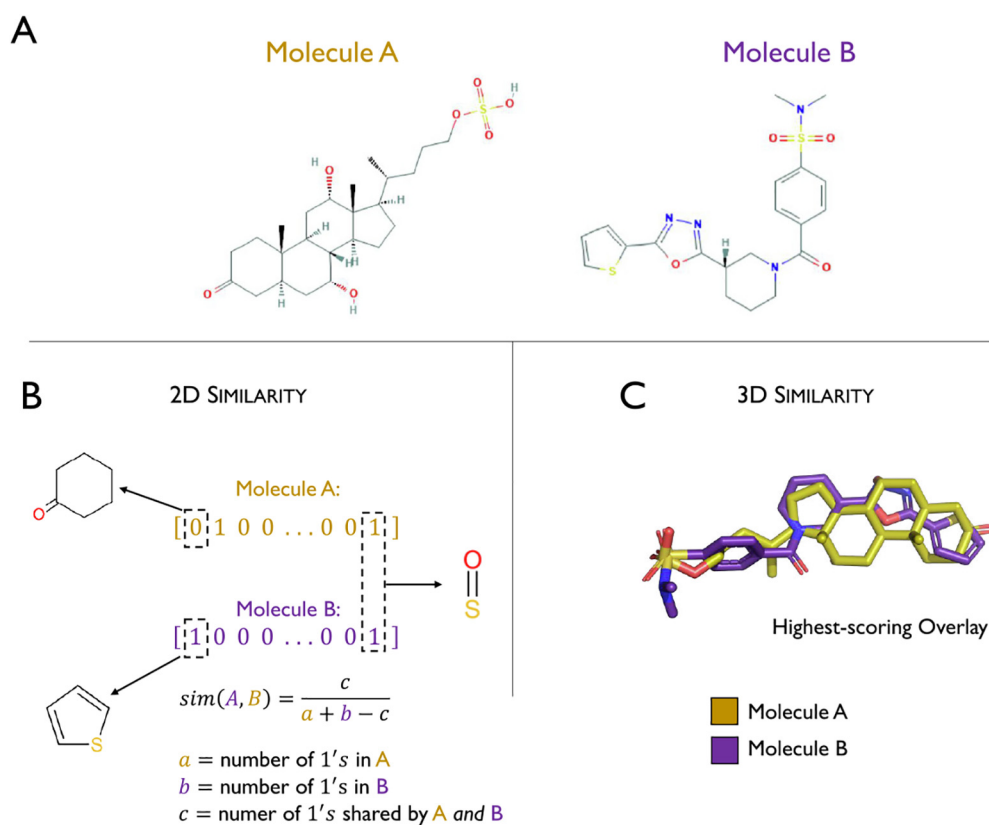


Fig. 8. Illustration of molecular similarity between two molecules (A), 3-keto-petromyzonol sulfate (molecule A) and a benzenesulfonamide (molecule B, ZINC ID ZINC38899649). (B) 2D similarity computed by comparing fingerprint representation of the molecules and computing the Tanimoto coefficient as a similarity measure between the two. (C) 3D volumetric overlay, which can be created by software such as ROCS [141], for which the Tanimoto similarity can be computed from the shapes based on positional densities.

significantly outperformed by another method. For the five GPCR in the dataset, the eSim had a more favorable AUC-ROC in all cases, with significantly better AUC-ROC performance for three of the five receptors.

3.2.3. Hybrid methods and neural network embeddings

Methods that rely on 2D fingerprints or 3D overlays for ligand-based VS utilize different information about the molecular structures – each method has individual strengths and weaknesses [147]. The HybridSimVS method combines both 2D fingerprint and 3D-structural information [148], and the proposed similarity metric is a combination of the Tanimoto similarity of the 3D shape overlay of a pair of molecules and of its 2D fingerprint. This hybrid metric outperformed MACCS, a 2D fingerprint, and WEGA, a 3D representation, when tested on the DUD-E benchmark dataset [146].

Section 3.1 summarized different neural network embeddings that can be used as inputs to machine learning models to predict various molecular properties. The same embeddings can be used for clustering-based similarity search. One example of this in practice is a clustering-based SPiDER search [149]. SPiDER uses self-organizing maps, a special type of artificial neural network approach for unsupervised learning, to create low-dimensional embeddings of the molecule space based on the CATS topological pharmacophores models and physicochemical small molecule descriptors [150]. Recently SPiDER was used in tandem with a binding affinity prediction via DEcRyPT to discover celastrol, a cannabinoid receptor 1 and 2 agonist [151,152]. DEcRyPT is based on random forest models for regression analysis, which have been trained to predict binding affinities from topological pharmacophore features. In this study, the researchers first used SPiDER to select potential cannabinoid receptor agonists via clustering. The selected molecules were then analyzed using DEcRyPT, which was trained on experimental binding affinity values for selected targets, as a predictive model that was applied to the most confident predictions from SPiDER to obtain binding affinity scores ($-\log_{10}$ values). The agonist celastrol was

identified by this approach, and the ligand was experimentally validated by biochemical assays including radiological displacement assays, followed by dynamic light scattering measurements to rule out false-positive readouts. The authors emphasized that common similarity-based searches using conventional ECFP4 Morgan-fingerprints would have disqualified celastrol as a GPCR ligand.

4. Receptor structure-based methods

As described in the previous section, many methods can be effectively utilized to identify bioactive GPCR ligands without explicitly including the receptor structure. Ligand-based methods are the only option when no receptor structure is available, as is the case for many human GPCRs [153]. However, the structure is informative if it is available and opens the door to new tasks that can be approached with machine learning. We explore applications of machine learning for two such tasks: binding site prediction and protein–ligand docking.

4.1. Binding site prediction

One of the most notable achievements in binding site prediction for identifying new targets was Nayal and Honig's random forest model, which was trained on 408 features representing structural, geometric, and physicochemical properties [154]. Their method was able to identify 1347 cavities on the surface of 99 diverse proteins, and in 100% of the cases, a drug has been experimentally found to bind at least one of these surface cavities.

A more recent method for discovering binding sites in GPCRs is implemented in the COACH-D server [37,38], which optionally allows users to provide a ligand structure that is then docked into the predicted binding site of the target protein structure to refine the ligand binding site predictions. The prediction in COACH-D is based on the COACH algorithm, which aggregates binding site predictions from both sequence and structure data based on five other methods and then uses a

linear SVM classifier for the final prediction [155].

Challenges for applying deep learning-based models to binding site identification are twofold. First, the dataset of available GPCR structures is relatively small. However, this problem can potentially be addressed by transfer learning, which is discussed in Section 6. Second, deep learning excels when the models can learn to extract patterns from the raw input data rather than hand-engineered feature descriptors as described in Section 3.1. While deep learning models based on 3D voxelization methods [156,157] and graph convolutions [122,123,158] have been successful in predicting binding affinities and molecular properties from small molecule structures, representing large molecules such as membrane receptors is more computationally prohibitive.

A recently developed deep learning method addresses the computational challenges of working with membrane receptors by splitting larger proteins into multiple parts [159]. The model BiteNet is a 3D convolutional network that takes 64 Å cubes voxelized into 1 Å units as input. Each voxel has a set of eleven channels for eleven atom types, and each channel's value is the output of an atomic density function applied to the channel's atom type. The output of BiteNet is an 8x8x8x4 tensor, where the first three dimensions correspond to the coordinates of voxel cells in the original 64 Å grid—the coordinates of 8 Å cubes within the 64 Å cube input. The last dimension's four channels correspond to a probability that the associated 8 Å cube contains a binding site and the cartesian coordinates of the most likely binding site. If a protein is larger than 64 Å in any dimension, then it is broken into overlapping 64 Å grids, which are considered separate data points. BiteNet was trained on the frames of molecular dynamic simulation trajectories of a curated set of protein–ligand complexes from PDB [160]. In a case study on the adenosine A_{2A} GPCR, the model was able to correctly identify its known binding sites.

4.2. Protein–ligand docking

Structure-based virtual screening for GPCRs is limited by the number of high-resolution GPCR structures that are available. However, advancements in techniques like Cryo-EM have increased the rate at which these structures are solved, such that many general AI-based methods for protein–ligand docking can be applied to modeling GPCR–ligand recognition in the future.

Structure-based VS methods need to be able to score protein–ligand complexes to identify favorable candidates. Scoring functions can be grouped into four classes: force field, empirical, knowledge-based, and machine learning-based [161]. Popular docking programs typically use scoring from the first three classes [33,162,163]. Functions from these three classes are sometimes referred to as classical scoring functions. Classical scoring functions are based on linear combinations of features, which restricts the ways they can utilize structural and interaction data [164]. Scoring functions based on machine learning are not restricted in this way. One of the first machine learning scoring functions to outperform classical methods was RF-score, which has been updated multiple times to further improve its performance [41,165,166]. All iterations of RF-score use the same algorithm for random forest regression. However, the researchers experimented with different ways of encoding the molecular information in the training dataset to improve the performance of RF-score, which underlines the importance of choosing a good input representation as previously discussed in Section 2. The feature representation of the protein–ligand complex used in the latest version of the RF-score (v3) [166] consists of counts of interacting protein–ligand atom pairs. Two atoms are considered interacting if they are within 10 Å in the structure. For example, a carbon in the protein that is 7 Å away from a nitrogen in the ligand in a binding pose would increment the count. These features are not symmetric; The feature for an atom of type *A* in the protein and an atom of type *B* in the ligand is a different feature than the one for an atom of type *B* in the protein and an atom of type *A* in the ligand. Additionally, RF-score v3 uses features produced by AutoDock Vina [33]. When trained and tested to model a

VS setting, RF-score outperformed classical scoring functions in various evaluation metrics [167].

Initial attempts at developing NN-based scoring functions in this space struggled with overfitting issues and were outperformed by methods that utilize conventional machine learning algorithms, like RF-score [168]. Recently, however, a set of new NN-based scoring functions have shown state-of-the-art performance [156,157,169,170]. All these networks use novel input representations to capture the structural and chemical information of the protein–ligand complex.

For models in this space, PDBbind is a relevant database that was used to train all of the NN-scoring functions we review. As of writing, it contains 17,679 protein–ligand complexes along with their experimentally-measured binding affinity data [171]. While not consisting of exclusively GPCRs, models trained on this set could easily be applied to scoring GPCR–ligand poses even if the following methods do not address them specifically.

A number of NN-based scoring functions voxelize the 3D representation of the protein–ligand complex (Section 2.3) and use functions that capture structural and chemical information to compute the channels of each voxel [156,157,170]. All of these methods are based on 3D CNNs, which can be computationally demanding due to their large number of trainable model parameters. Furthermore, the large number of parameters increases the number of training examples required to effectively train the networks without overfitting, which has been identified as an issue for DNNs in this space [119]. However, more recent methods aim to address the overfitting issue by using more modern CNN architectures. For instance, DeepAtom [157] was recently designed based on ShuffleNet v2 [172], and the recent K_{DEEP} [156] method is based on SqueezeNet [173]. These architectures were designed to achieve good computational performance while keeping the number of trainable parameters relatively small, which helps with minimizing overfitting when deeper architectures are being trained on smaller datasets. Both DeepAtom and K_{DEEP} achieved state-of-the-art binding affinity predictions on the *PDBbind v. 2016* core dataset [171]. The success of these networks suggests that advances in other subfields of deep learning could be effective in VS as well and could be a fruitful area of future research.

Other NN-based scoring functions achieved state-of-the-art performance by including more information in the input representations. OnionNet uses atom pair counts similar to RF-score, but generates them for multiple cutoffs and makes each pair exclusive to the first cutoff it appears in [169]. These values are then converted into a $n \times m$ matrix, where n is the cutoff value and m is the atom pair, and this matrix is used as input for a 2D CNN. The use of a 2D CNN reduces the parameter burden of this model so the area of the protein that can be encoded increases substantially. The previously mentioned 3D CNN models, DeepAtom and K_{DEEP} , can only capture molecular information in a 20 Å cube, while OnionNet counts at least all atom pairs within a 61 Å cube. Long-range electrostatic interactions extend beyond 20 Å and are known to be important in protein–ligand binding. Hence, OnionNet may be more effective for VS targets where these long-range interactions are important for activity prediction [174]. In practice, its performance on the *PDBbind v. 2016* core dataset was comparable to the 3D CNNs mentioned above, though there was no performance comparison conducted on individual targets where this information could be useful.

5. *De novo* small-molecule design

De novo molecule design is closely aligned with the goals of VS. While VS attempts to answer questions such as “Does this molecule activate this target?”, *de novo* design attempts to answer “Can we generate a molecule that can activate this target?”. Answering the second question in the affirmative implies that we understand what makes a molecule active against a given target. *De novo* molecule design can be tied to the idea of inverse QSAR/QSPR, deriving the structures of

AI-based GPCR bioactive ligand discovery

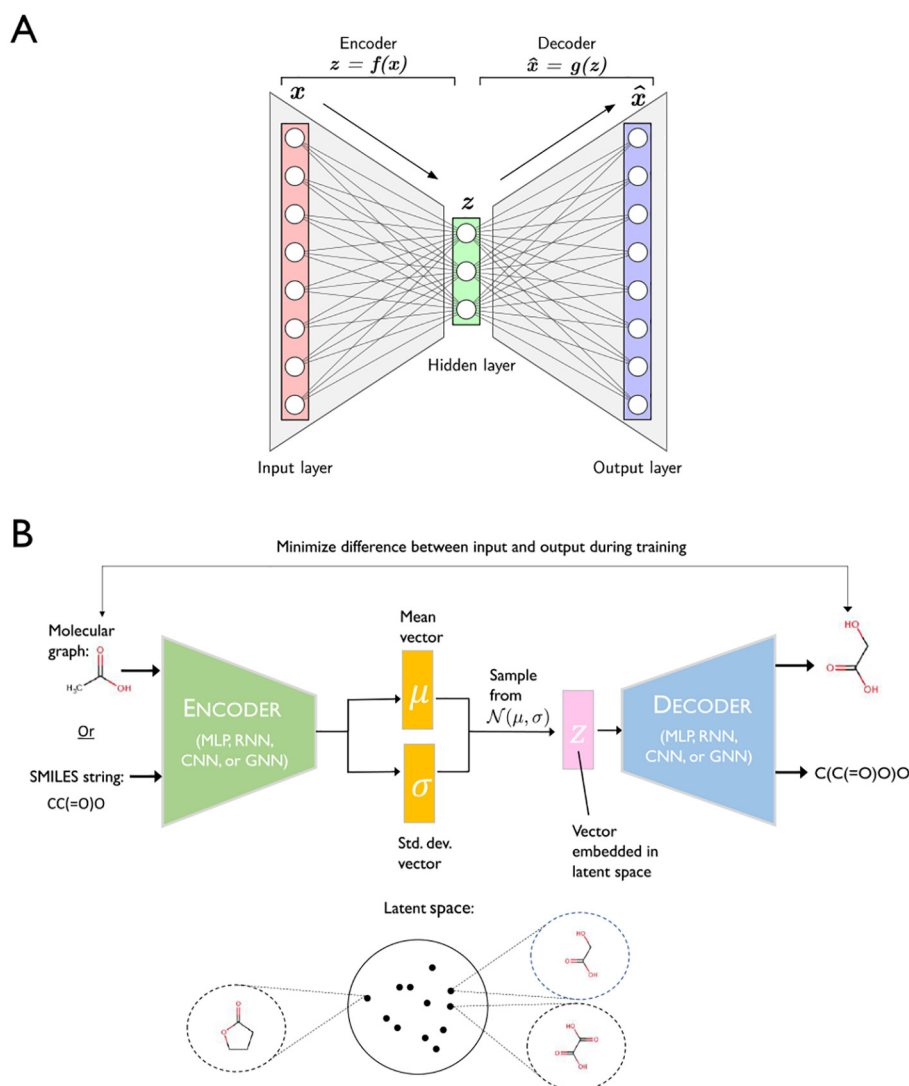


Fig. 9. (A) Fully connected autoencoder consists of two submodules: an *encoder* followed by a *decoder*. The encoder submodule compresses an input representation x into a lower-dimensional representation z . The decoder converts the lower-dimensional representation to \hat{x} , such that reconstruction resembles the original input, $\hat{x} \approx x$. (B) Illustration of a variational autoencoder architecture and how it can be applied to molecules. Given a molecular representation as input (acetic acid), an encoder network outputs both a mean vector and a standard deviation vector. These vectors are used as parameters of a normal distribution to sample a vector embedded in the latent space (z). This latent vector is used as input for a decoder network, which outputs a molecular representation (*here*: glycolic acid). The loss function for variational autoencoder training consists of two components: the first component is the reconstruction loss ensuring that the original input resembles the generated output; the second component encourages similar molecules to cluster together in the latent space. In this case the structurally similar glycolic acid and oxalic acid are close in latent space, while the less similar gamma-butyrolactone is farther away. After training, the encoder portion can be discarded, and new molecules can be sampled by sampling different z vectors from a standard normal distribution.

all molecules that are active with a target or have a certain property. This is a challenging problem, but despite its difficulty, *de novo* synthesis has recently become an active area of research, which is likely due to the recent advances in the development of deep generative models and reinforcement learning.

5.1. Generating molecules with variational autoencoders

Autoencoders refer to architectures that were traditionally used for dimensionality reduction and trained in an unsupervised manner [175]. The main idea behind autoencoders is that they learn to preserve the essential part of the data and remove the non-essential ones (Fig. 9 A).

An autoencoder consists of two neural networks: an encoder that transforms a training example into a lower-dimensional representation,

and a decoder that reconstructs the original training example from its lower-dimensional representation. For example, if both the encoder and decoder submodules of the autoencoder are fully-connected NNs with linear activation functions, the latent feature embedding produced by the encoder is equivalent to features transformed via principal component analysis except for the feature orthogonality constraint.

It shall be noted that regular autoencoders are deterministic models that can be used for data compression or to modify existing data given a user-specified objective [176–179]. A related model is the variational autoencoder, which can be considered a generative model as it is capable of synthesizing entirely new, realistic data records mimicking the data contained in the training dataset [180,181].

In essence, a VAE learns to reconstruct high-dimensional inputs from a low-dimensional continuous latent space. In contrast to a

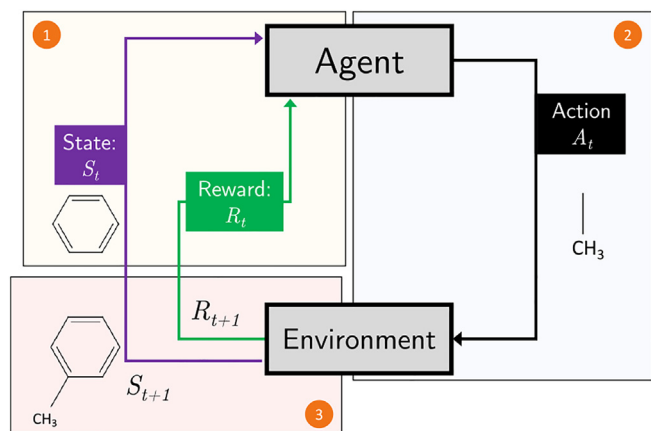


Fig. 10. Representation of the basic reinforcement learning paradigm with a simple molecular example. (1) Given a benzene ring (state S_t at iteration t) and some reward value R_t at iteration t , (2) the agent selects an action A_t that adds a methyl group to the benzene ring. (3) The environment considers this information for producing the next state (S_{t+1}) and reward (R_{t+1}). This cycle repeats until the episode is terminated.

conventional autoencoder, points closely together in the latent space are encouraged to have reconstructions with similar properties. A high-level diagram of a molecular VAE is displayed in Fig. 9 B. These models take a molecular representation, like a SMILE string [133,182] or a graph [183,184], as input. The molecular representations are then passed into an encoder network that outputs a mean vector and a standard deviation vector. These vectors are used as parameters for a multivariate normal distribution, which is sampled from to produce a low-dimensional embedding. The embedding is then passed through a decoder network to get a molecular representation as output. Thus the embedding can be viewed as a lower-dimensional representation of the molecule that is output by the decoder network.

The dimensional space that the embedding vector occupies is commonly referred to as the *latent space*. If the goal of an autoencoder is to reproduce its data input, it may seem strange that the encoder network generates parameters to sample data points in latent space instead of generating the data instances directly. However, this particular design ensures that points close in latent space result in reconstructions with similar properties. Without getting into the details, the parameter vectors, which are output by the encoder network, are used as a regularization term for reconstruction loss of the molecule input into the VAE. This term is minimized when the latent embeddings are sampled from a standard normal distribution. However, if all latent embeddings are sampled from a standard normal distribution, the reconstruction loss will be high. This encourages the model to group the latent vectors of similar molecules closely together in the embedding space. While this minimizes the regularization term, it also reduces the risk of incurring a high reconstruction loss, because the reconstruction loss for similar molecules will be less than that for different ones. This concept is also conveyed in Fig. 9 B – the molecule generated by the model differs from the input molecule, but they share a relatively large degree of structural similarity and are close in the latent space. This property of the latent space can be utilized to generate novel molecules with similar properties to selected input molecules by taking the input’s latent representation and perturbing it in a controlled manner [133,185]. One of the challenges with using VAE-based models in this space is that with current architectures, the VAE has no notion of what constitutes a valid molecule [183]. Since these models are trained on large numbers of valid compounds, certain models were still able to achieve reasonable results despite producing some invalid structures [133,182].

Jin et al. introduced a novel VAE design that only allowed valid molecules as outputs, called a junction tree VAE [183]. This model uses a GNN to encode a molecule’s graph and also a junction tree of valid

molecular fragments in the graph. The decoder network first decodes a junction tree from the latent space. Then, it decodes a molecule that satisfies that junction tree. Since all the molecular fragments the junction tree can encode are valid, a valid molecule can always be derived from the junction tree [186]. In addition to always producing valid molecules, the model outperformed other VAEs on a Bayesian optimization task, where the goal was to find a molecule that maximized octanol–water partition coefficients ($\log P$).

One shortcoming of regular VAEs is that the molecular similarity in the latent spaces does not directly translate to the similarity of a molecular property of interest. As a general example, consider the latent representation of a ligand with a thiol group that induces a positive activity response in a given target receptor. Perturbation of the latent representation of that molecule may result in similar molecules that share the thiol group feature but are not active against the target. As a consequence, the performance for VAEs in this space is usually assessed on simple chemical properties [182,183,185]. For this reason, while there is still plenty to be investigated about VAEs for *de novo* synthesis, research that has focused on the *de novo* generation of compounds for specific targets has largely utilized reinforcement learning, which will be discussed in the next section.

5.2. Reinforcement learning-based molecular design

Reinforcement learning seeks to teach an agent to perform actions that will maximize a cumulative reward over a series of iterations. In this paradigm, the definition of the reward is very flexible, which makes it possible to more directly optimize for specific traits compared to using generative models such as VAEs. As illustrated in Fig. 10, at each iteration, the agent is given a state by the environment as well as the reward from the previous iteration. The agent then selects an action based on this information. This selected action is fed into the environment, which then outputs a new state and reward. This cycle continues until the series of iterations, called an episode, is terminated, at which point the agent’s behavior is updated based on the information provided by its trajectory, the ordered list of states, actions, and rewards associated with each iteration in the episode. While this high-level overview will suffice to understand how reinforcement learning is utilized for *de novo* design, the reader is referred to [50,187], which we recommend as resources with more in-depth explanations.

Several examples exist in which reinforcement learning models have been optimized to generate GPCR bioactive molecules. One such example is the REINVENT model [42]. In REINVENT, the agent is an RNN that generates one SMILES string token at a time over the course of multiple iterations. After a batch of 128 SMILES strings are produced, the RNN agent is updated with a modified form of REINFORCE, an algorithm that updates the agent network based on the trajectories of the episodes in the current batch [188]. In one experiment, REINVENT was trained to generate compounds that are active with the GPCR Dopamine receptor D_2 (DRD2). In this experiment, the reward function was based on the output of an SVM trained to predict a compound’s DRD2 activity. After training, the agent was able to produce compounds that the SVM labeled as active against DRD2 96% of the time. The reinforcement learning model generated both novel, chemically valid compounds and, impressively, compounds that are known to be active with DRD2 even though these were not included in the training set.

DrugEx is another deep reinforcement learning model that was directly optimized for generating molecules with GPCR activity [189]. Here, the researchers focused on adenosine A_{2A} receptor, which has been targeted to treat cardiovascular and inflammatory diseases [190]. Similar to REINVENT, the DrugEx model also produces SMILES strings over the course of an episode. However, DrugEx adds a stochastic component during training. Before being used in the reinforcement learning training, the RNN agent network was individually trained with a large set of molecule SMILES from ZINC [21]. Two copies of this pre-trained network were then created, with one referred to as the *exploration*

network and the other as the *exploitation* network. Only the exploitation network was updated during the reinforcement learning training process; however, with a specified probability at each iteration, the exploration network would be queried for the next token instead. The purpose of this procedure was to explore a wider chemical space during training – afterwards, the exploration network was discarded, and only the exploitation network was used to generate new molecules. This method successfully rediscovered some known actives for adenosine A_{2A} receptor. The authors also suggested the RNN agent was able to produce molecules with a large diversity, by showing the actives generated by the model covered all clusters generated by fingerprint-based clustering on known adenosine A_{2A} receptor actives.

Other reinforcement learning-based models for *de novo* synthesis described in the literature were not specifically focused on GPCRs' bioactive molecule design but could be adopted for such tasks in the future [191,192]. One such example is Zhou et al.'s Molecule Deep Q-Networks (MOLDQN) approach, which modifies deep Q-networks for molecule generation [191,193]. This model's agent network takes the current molecular graph's Morgan fingerprint as input and selects an action to modify the molecular graph. The actions include adding an atom, adding a bond, or increasing a bond's order. Additionally, the actions that are allowed at a given iteration are restricted if they are invalid so the system will always produce valid molecules. Similar to DrugEx, the model encourages exploration by selecting a random valid action at a given iteration with some probability ϵ . The model achieved state-of-the-art performance when producing molecules that maximized for logP and quantitative estimates of drug likeness separately [194].

While many publications in this area demonstrate the ability to optimize for molecular properties, they unfortunately lack experimental follow-up procedures for further model evaluation. As a notable counter-example, we want to highlight GENTRL, which was used to discover novel inhibitors of discodin domain receptor 1, a tyrosine kinase [14]. The inhibitors were generated computationally and then synthesized and experimentally validated. The experiments *in silico* and *in vitro* were completed in approximately 46 days at a fraction of the cost of a high throughput screening approach.

6. Transfer learning

Machine learning, and deep learning in particular, requires large training datasets. It's not atypical for modern deep neural networks to have millions of trainable parameters, which require sufficient data for successful parameterization. Transfer learning refers to the process of adapting models that have been trained on one task to another, typically similar task [195]. The most common form of transfer learning is to take a DNN that was trained on a large general-purpose dataset and fine-tune it to a smaller dataset of interest. As an example, suppose there is only a limited number of ligands with bioactivity data for a certain class A GPCR of interest. However, a larger bioactivity dataset exists for a second GPCR from the same subfamily that has a similar binding site. In this scenario, one could apply transfer learning by training a DNN to predict bioactive ligands for the second GPCR and then fine-tune the model weights on the smaller dataset of the GPCR of interest. Typically, the more similar the initial dataset is to the dataset for the target task, the more successful transfer learning can be. In our example, if the ligands in the larger dataset are physio-chemically similar to those in the smaller one, one could expect transfer learning to be more successful.

One example of transfer learning being applied to improve the performance of models with limited available data is in a recent structure-based method that sought to train a 3D CNN to classify molecule activity from high-scoring docking poses [196]. When the network was trained on targets in a DUD-E based training set and then fine-tuned on a smaller, protein family-specific dataset, the model performed better on test set targets in that protein family than the same network that was trained only on the family-specific set [146]. Family-

specific models typically outperform general models, but data availability may limit their performance [197]. This work shows that transfer learning can help to address this issue. While it has not yet been applied further in bioactivity predictions, transfer learning has been successful in improving the performance of models that predict standard molecular properties like solvency as well as in models that predict quantum mechanical approximations [198,199].

7. Future directions

Currently, one of the most neglected areas of machine learning in bioactive ligand discovery is active learning, which aims to combine artificial and human intelligence. Active learning is a branch of machine learning that focuses on selecting labeled data for supervised learning to improve non-confident predictions and fill the knowledge gaps of the model [200]. Knowledge gaps can be filled through human interaction, by sampling optimal unlabeled training examples for annotations by humans. We think that utilizing domain knowledge from experts is crucial in a highly specialized field such as GPCR bioactive ligand discovery, and we believe that it is highly advantageous to develop machine learning systems that include human feedback loops.

Another avenue that remains largely unexplored for computational ligand discovery is a combination of active learning and transfer learning (Section 6), called active transfer learning. Here, a separate model is trained on an existing validation set to predict whether a given model is able to make correct predictions or not on a given unlabeled dataset, which then can highlight problematic cases for human review [200].

Another recent development and promising research trend in deep learning is semi-supervised learning, which is particularly useful if pre-trained models for transfer learning are not available for the target domain or are infeasible to obtain. Semi-supervised learning is the process of deriving and utilizing label information directly from the data itself rather than having humans annotating it. In a language model, for example, semi-supervised learning can be utilized by training a model to predict the next word in a sequence [201] or, in the context of computer vision, this could be the composition of an image into a jigsaw puzzle that the DNN learns to assemble [202]. The main idea of this approach is to choose a task that requires an understanding of the underlying data in order to be solved. This stage of self-supervised training can be regarded as pre-training, and the model can then be adopted and fine-tuned to solve the target task downstream similar to conventional transfer learning.

In addition to applications of new machine learning paradigms to VS for GPCRs it is also worth considering how machine learning could be used to assist other areas of GPCR research. One potential application for machine learning that could push GPCR research forward is to assist in deorphanization. While this intersection is largely unexplored there is one recent method where researchers used an SVM trained on peptide-descriptor vectors for GPCR neuropeptide ligands to predict if a set of orphaned neuropeptides from *C. elegans* interacted with any putative GPCRs in the organism [203]. After predicting 22 putative neuropeptide-GPCR pairs, the authors experimentally validated 11 of them with cell-based signaling assays to provide evidence for nine new receptors. A similar approach could be applied to general chemical ligands by modifying the data representation. Such an approach could also be effective in other organisms or with more advanced machine learning models.

Binding site modeling is another area that modern machine learning methods could drive forward. As mentioned throughout the paper, one challenge when working with GPCRs is that no high-resolution structures are available. Accurate computational approaches for predicting the structure of a binding site could have an immensely beneficial impact on guiding chemical screening. It is worth mentioning the recently published deep learning model AlphaFold obtained state-of-the-art performance on the protein folding problem, far surpassing traditional

approaches to it [204]. While a prediction of binding site structure would ideally require higher resolutions than this model provides, it is a step in the right direction.

8. Conclusions

The confluence of the meteoric rise of deep learning methods, the increase in large chemical datasets, and the further ease of access to high powered computing resources have spurred the field of chemical machine learning forward at a rapid pace. The many representations of chemical data allow a variety of methods used in other fields to be put to task on chemical quandaries and have encouraged the development of new machine learning models. As we have discussed, many of these models have been or could be applied to VS on GPCRs, though the types of GPCR data currently available may limit the general applicability of certain models. However, as the amount of bioactivity data available continues to increase, so will the performance of these models. One area in which the space is currently lacking is the experimental validation of the computational method development.

Hopefully, by providing an overview of deep learning methods and detailed explanations of new approaches in the space, we can encourage the more experimentally-minded to utilize these methods, or at least to collaborate more confidently. In the other direction, it is important to convey to the methodologically-inclined that while it is great to develop a new method that performs well on a benchmark, experimental validation is the gold standard, so it is beneficial to seek out experimental collaboration when possible. This will lead to more actionable results and push the development of machine learning VS methods even faster in the future.

Acknowledgements

Support for this work was provided by the Office of the Vice Chancellor for Research and Graduate Education at the University of Wisconsin-Madison with funding from the Wisconsin Alumni Research Foundation. In addition, we are thankful for the support provided by the NLM Biomedical Informatics Training Program (Grant No. 5T15LM007359). Also, we would like to thank Kylie Moynihan for helpful feedback on the manuscript.

References

- [1] A.S. Hauser, M.M. Attwood, M. Rask-Andersen, H.B. Schiöth, D.E. Gloriam, Trends in GPCR drug discovery: new agents, targets and indications, *Nature Reviews Drug Discovery* 16 (12) (2017) 829.
- [2] S.L. Garland, Are GPCRs still a source of new targets? *Journal of Biomolecular Screening* 18 (9) (2013) 947–966.
- [3] W. Thomsen, J. Frazer, D. Unett, Functional assays for screening GPCR targets, *Current Opinion in Biotechnology* 16 (6) (2005) 655–665.
- [4] T.K. Bjarnadóttir, D.E. Gloriam, S.H. Hellstrand, H. Kristiansson, R. Fredriksson, H.B. Schiöth, Comprehensive repertoire and phylogenetic analysis of the G protein-coupled receptors in human and mouse, *Genomics* 88 (3) (2006) 263–273.
- [5] A.P. Davenport, S.P. Alexander, J.L. Sharman, A.J. Pawson, H.E. Benson, A.E. Monaghan, et al., International union of basic and clinical pharmacology. LXXXVIII.G protein-coupled receptor list: recommendations for new pairings with cognate ligands, *Pharmacological Reviews* 65 (3) (2013) 967–986.
- [6] S. Raschka, Automated discovery of GPCR bioactive ligands, *Current Opinion in Structural Biology* 55 (2019) 17–24.
- [7] C. Southan, J.L. Sharman, H.E. Benson, E. Faccenda, A.J. Pawson, S.P. Alexander, et al., The IUPHAR/BPS Guide to PHARMACOLOGY in 2016: towards curated quantitative interactions between 1300 protein targets and 6000 ligands, *Nucleic Acids Research* 44 (D1) (2015) D1054–D1068.
- [8] C.H. Wong, K.W. Siah, A.W. Lo, Estimation of clinical trial success rates and related parameters, *Biostatistics* 20 (2) (2019) 273–286.
- [9] J. Vamathevan, D. Clark, P. Czodrowski, I. Dunham, E. Ferran, G. Lee, et al., Applications of machine learning in drug discovery and development, *Nature Reviews Drug Discovery* 1 (2019):.
- [10] U. Food, Administration D. FDA (Eds.). What Are Biologics; 2020. Accessed: 2020-01-02. URL:<https://www.fda.gov/about-fda/center-biologics-evaluation-and-research-cber/what-are-biologics-questions-and-answers>.
- [11] A. Mullard, 2018 FDA drug approvals. NLM (Medline), 2019.
- [12] T. Rodrigues, G.J. Bernardes, Machine learning for target discovery in drug development, *Current Opinion in Chemical Biology* 56 (2020) 16–22.
- [13] S.M. Paul, D.S. Mytelka, C.T. Dunwiddie, C.C. Persinger, B.H. Munos, S.R. Lindborg, et al., How to improve R&D productivity: the pharmaceutical industry's grand challenge, *Nature Reviews Drug Discovery* 9 (3) (2010) 203.
- [14] A. Zhavoronkov, Y.A. Ivanenkov, A. Aliper, M.S. Veselov, V.A. Aladinskiy, A.V. Aladinskaya, et al., Deep learning enables rapid identification of potent DDR1 kinase inhibitors, *Nature Biotechnology* 37 (9) (2019) 1038–1040.
- [15] C.G. Parker, A. Galmozzi, Y. Wang, B.E. Correia, K. Sasaki, C.M. Joslyn, et al., Ligand and target discovery by fragment-based screening in human cells, *Cell* 168 (3) (2017) 527–541.
- [16] L. Bar-Peled, E.K. Kemper, R.M. Suci, E.V. Vinogradova, K.M. Backus, B.D. Horning, et al., Chemical proteomics identifies druggable vulnerabilities in a genetically defined cancer, *Cell* 171 (3) (2017) 696–709.
- [17] R.E. Moellering, B.F. Cravatt, How chemoproteomics can enable drug discovery and development, *Chemistry & Biology* 19 (1) (2012) 11–22.
- [18] L. Laraia, H. Waldmann, Natural product inspired compound collections: evolutionary principle, chemical synthesis, phenotypic screening, and target identification, *Drug Discovery Today: Technologies* 23 (2017) 75–82.
- [19] V. Duros, J. Grizou, W. Xuan, Z. Hosni, D.L. Long, H.N. Miras, et al., Human versus robots in the discovery and crystallization of gigantic polyoxometalates, *Angewandte Chemie* 129 (36) (2017) 10955–10960.
- [20] F. Häse, L.M. Roch, A. Aspuru-Guzik, Next-generation experimentation with self-driving laboratories, *Trends in Chemistry* 1 (3) (2019) 282–291.
- [21] T. Sterling, J.J. Irwin, ZINC 15-ligand discovery for everyone, *Journal of Chemical Information and Modeling* 55 (11) (2015) 2324–2337.
- [22] J. Sunseri, D.R. Koes, Pharmit: interactive exploration of chemical space, *Nucleic Acids Research* 44 (W1) (2016) W442–W448.
- [23] A.P. Bento, A. Gaulton, A. Hersey, L.J. Bellis, J. Chambers, M. Davies, et al., The ChEMBL bioactivity database: an update, *Nucleic Acids Research* 42 (D1) (2014) D1083–D1090.
- [24] S. Raschka, A.M. Scott, N. Liu, S. Gunturu, M. Huertas, W. Li, et al., Enabling the hypothesis-driven prioritization of ligand candidates in big databases: Screenlamp and its application to GPCR inhibitor discovery for invasive species control, *Journal of Computer-aided Molecular Design* 32 (3) (2018) 415–433.
- [25] S. Raschka, A.M. Scott, M. Huertas, W. Li, L.A. Kuhn, Automated inference of chemical discriminants of biological activity, in: *Computational Drug Discovery and Design*, Springer, 2018, pp. 307–338.
- [26] S. Basith, M. Cui, S.J. Macalino, J. Park, N.A. Clavio, S. Kang, et al., Exploring G protein-coupled receptors (GPCRs) ligand space via cheminformatics approaches: impact on rational drug design, *Frontiers in Pharmacology* 9 (2018) 128.
- [27] M. Wheatley, D. Wootten, M.T. Conner, J. Simms, R. Kendrick, R.T. Logan, et al., Lifting the lid on GPCRs: the role of extracellular loops, *British Journal of Pharmacology* 165 (6) (2012) 1688–1703.
- [28] H. Zhang, S.S. Ericksen, Lee Cp, G.E. Ananiev, N. Wlodarchak, P. Yu, et al., Predicting kinase inhibitors using bioactivity matrix derived informer sets, *PLOS Computational Biology* 15 (8) (2019): e1006813 .
- [29] J. Bemister-Buffington, A.J. Wolf, S. Raschka, L.A. Kuhn, Machine Learning to Identify Flexibility Signatures of Class A GPCR Inhibition, *Biomolecules* 10 (3) (2020) 454.
- [30] D.J. Jacobs, A.J. Rader, L.A. Kuhn, M.F. Thorpe, Protein flexibility predictions using graph theory, *Proteins: Structure, Function, and Bioinformatics* 44 (2) (2001) 150–165.
- [31] S. Raschka, Model evaluation, model selection, and algorithm selection in machine learning. arXiv preprint arXiv:1811.2808, 2018.
- [32] G. Neudert, G. Klebe, DSX: a knowledge-based scoring function for the assessment of protein–ligand complexes, *Journal of Chemical Information and Modeling* 51 (10) (2011) 2731–2745.
- [33] O. Trott, A.J. Olson, AutoDock Vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading, *Journal of Computational Chemistry* 31 (2) (2010) 455–461.
- [34] J. Jeon, S. Nim, J. Teyra, A. Datti, J.L. Wrana, S.S. Sidhu, et al., A systematic approach to identify novel cancer drug targets using machine learning, inhibitor design and high-throughput screening, *Genome Medicine* 6 (7) (2014) 57.
- [35] S. Riniker, Y. Wang, J.L. Jenkins, G.A. Landrum, Using information from historical high-throughput screens to predict active compounds, *Journal of Chemical Information and Modeling* 54 (7) (2014) 1880–1891.
- [36] A.S. Rifaioğlu, H. Atas, M.J. Martin, R. Cetin-Atalay, V. Atalay, T. Dogan, Recent applications of deep learning and machine intelligence on in silico drug discovery: methods, tools and databases, *Briefings in Bioinformatics* 10 (2018).
- [37] H.S. Chan, Y. Li, T. Dahoun, H. Vogel, S. Yuan, New binding sites, new opportunities for GPCR drug discovery, *Trends in Biochemical Sciences* (2019).
- [38] Q. Wu, Z. Peng, Y. Zhang, J. Yang, COACH-D: improved protein–ligand binding sites prediction with refined ligand-binding poses through molecular docking, *Nucleic Acids Research* 46 (W1) (2018) W438–W442.
- [39] E. Ferrero, I. Dunham, P. Sanseau, In silico prediction of novel therapeutic targets using gene–disease association data, *Journal of Translational Medicine* 15 (1) (2017) 182.
- [40] A.B. Farimani, E. Feinberg, V. Pande, Binding pathway of opiates to μ -opioid receptors revealed by machine learning, *Biophysical Journal* 114 (3) (2018) 62a–63a.
- [41] P.J. Ballester, Mitchell J.B.A machine learning approach to predicting protein–ligand binding affinity with applications to molecular docking, *Bioinformatics* 26 (9) (2010) 1169–1175.
- [42] M. Olivecrona, T. Blaschke, O. Engkvist, H. Chen, Molecular de-novo design through deep reinforcement learning, *Journal of Cheminformatics* 9 (1) (2017) 48.
- [43] A. Kadurin, S. Nikolenko, K. Khrabrov, A. Aliper, A. Zhavoronkov, druGAN: an advanced generative adversarial autoencoder model for de novo generation of new

- molecules with desired molecular properties in silico, *Molecular Pharmaceutics* 14 (9) (2017) 3098–3104.
- [44] P. Mamoshina, M. Volosnikova, I.V. Ozerov, E. Putin, E. Skibina, F. Cortese, et al., Machine learning on human muscle transcriptomic data for biomarker discovery and tissue-specific drug target identification, *Frontiers in Genetics* 9 (2018) 242.
- [45] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (7553) (2015) 436.
- [46] A. Canziani, A. Paszke, E. Cukurciello, An analysis of deep neural network models for practical applications, ArXiv preprint arXiv:160507678, 2016.
- [47] D.E. Rumelhart, G.E. Hinton, R.J. Williams, Learning representations by back-propagating errors, *Nature* 323 (6088) (1986) 533–536.
- [48] W. Cao, V. Mirjalili, S. Raschka, Rank-consistent ordinal regression for neural networks, ArXiv preprint arXiv:190107884, 2019.
- [49] J. Deng, W. Dong, R. Socher, L.J. Li, K. Li, L. Fei-Fei, ImageNet: A large-scale hierarchical image database, 2009 IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2009, pp. 248–255.
- [50] S. Raschka, V. Mirjalili, Python Machine Learning: Machine Learning and Deep Learning with Python, scikit-learn, and TensorFlow 2, Packt Publishing Ltd, 2019.
- [51] S. Van Der Walt, S.C. Colbert, G. Varoquaux, The NumPy array: a structure for efficient numerical computation, *Computing in Science & Engineering* 13 (2) (2011) 22.
- [52] J.D. Hunter, Matplotlib: A 2D graphics environment, *Computing in Science & Engineering* 9 (3) (2007) 90.
- [53] W. McKinney, et al., Data structures for statistical computing in python, in: Proceedings of the 9th Python in Science Conference. vol. 445. Austin, TX, 2010. pp. 51–56.
- [54] P.J. Cock, T. Antao, J.T. Chang, B.A. Chapman, C.J. Cox, A. Dalke, et al., Biopython: freely available Python tools for computational molecular biology and bioinformatics, *Bioinformatics* 25 (11) (2009) 1422–1423.
- [55] S. Raschka, BioPandas: Working with molecular structures in pandas DataFrames, *Journal of Open Source Software* 2 (14) (2017) 279.
- [56] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, et al., Scikit-learn: Machine learning in Python, *Journal of Machine Learning Research* 12 (Oct) (2011) 2825–2830.
- [57] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, et al., Tensorflow: A system for large-scale machine learning, in: 12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16), 2016, pp. 265–283.
- [58] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, et al., PyTorch: An imperative style, high-performance deep learning library, in: Advances in Neural Information Processing Systems, 2019, pp. 8024–8035.
- [59] G. Kocielnny, P. An, D. Carvalho-Silva, J.A. Cham, L. Fumis, R. Gasparyan, et al., Open Targets: a platform for therapeutic target identification and validation, *Nucleic Acids Research* 45 (D1) (2016) D985–D994.
- [60] Z. Wu, B. Ramsundar, E.N. Feinberg, J. Gomes, C. Geniesse, A.S. Pappu, et al., MoleculeNet: a benchmark for molecular machine learning, *Chemical Science* 9 (2) (2018) 513–530.
- [61] A. Gaulton, A. Hersey, M. Nowotka, A.P. Bento, J. Chambers, D. Mendez, et al., The ChEMBL database in 2017, *Nucleic Acids Research* 45 (D1) (2016) D945–D954.
- [62] S. Raschka, A.J. Wolf, J. Bemister-Buffington, L.A. Kuhn, Protein–ligand interfaces are polarized: discovery of a strong trend for intermolecular hydrogen bonds to favor donors on the protein side with implications for predicting and designing ligand complexes, *Journal of Computer-aided Molecular Design* 32 (4) (2018) 511–528.
- [63] S. Raschka, J. Bemister-Buffington, L.A. Kuhn, Detecting the native ligand orientation by interfacial rigidity, *SiteInterlock. Proteins: Structure, Function, and Bioinformatics* 84 (12) (2016) 1888–1901.
- [64] M.T. Ribeiro, S. Singh, C. Guestrin, Why should I trust you?: Explaining the predictions of any classifier, Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2016, pp. 1135–1144.
- [65] S.M. Lundberg, S.I. Lee, A unified approach to interpreting model predictions, in: Advances in Neural Information Processing Systems, 2017, pp. 4765–4774.
- [66] J. Springenberg, A. Dosovitskiy, T. Brox, M. Riedmiller, Striving for simplicity: the all convolutional net, in: ICLR (workshop track), 2014, pp. 1–14.
- [67] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, A. Torralba, Learning deep features for discriminative localization, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 2921–2929.
- [68] R.R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, Grad-cam: Visual explanations from deep networks via gradient-based localization, Proceedings of the IEEE International Conference on Computer Vision, 2017, pp. 618–626.
- [69] A. Shrikumar, P. Greenside, A. Kundaje, Learning important features through propagating activation differences, Proceedings of the 34th International Conference on Machine Learning–Volume 70. JMLR. org, 2017, pp. 3145–3153.
- [70] F. Hu, J. Jiang, P. Yin, Interpretable prediction of protein–ligand interaction by convolutional neural network, in: 2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), IEEE, 2019, pp. 656–659.
- [71] S. Runge, H. Thøgersen, K. Madsen, J. Lau, R. Rudolph, Crystal structure of the ligand-bound glucagon-like peptide-1 receptor extracellular domain, *Journal of Biological Chemistry* 283 (17) (2008) 11340–11347.
- [72] committee N. Reproducibility Checklist, Accessed: 2020–01-02. URL:https://nips.cc/Conferences/2019.
- [73] A. Burkov, The hundred-page machine learning book, Andriy Burkov (2019).
- [74] I. Goodfellow, Y. Bengio, A. Courville, Deep Learning, MIT press, 2016.
- [75] B.C. Csáji, Approximation with Artificial Neural Networks, Faculty of Sciences, Etsv Lornd University, Hungary, 2001, pp. 24–48.
- [76] G. Cybenko, Approximations by superpositions of a sigmoidal function, *Mathematics of Control, Signals and Systems* 2 (1989) 183–192.
- [77] K. Hornik, M. Stinchcombe, H. White, Multilayer feedforward networks are universal approximators, *Neural Networks* 2 (5) (1989) 359–366.
- [78] J. Schmidhuber, Deep learning in neural networks: An overview, *Neural Networks* 61 (2015) 85–117.
- [79] A.G. Ivakhnenko, V.G. Lapa, Cybernetic Predicting Devices, Purdue University Lafayette and Ind. School of Electrical Engineering, 1966.
- [80] A.G. Ivakhnenko, V.G. Lapa, Cybernetics and Forecasting Techniques, North-Holland, 1967.
- [81] S. Linnainmaa, The representation of the cumulative rounding error of an algorithm as a Taylor expansion of the local rounding errors, Master’s Thesis (in Finnish), Univ Helsinki, 1970, pp. 6–7.
- [82] P. Werbos, Beyond regression: new tools for prediction and analysis in the behavioral sciences, Ph D dissertation, Harvard University, 1974.
- [83] D. Steinkraus, I. Buck, P. Simard, Using GPUs for machine learning algorithms, 8th International Conference on Document Analysis and Recognition, IEEE, 2005, pp. 1115–1120.
- [84] K. Chellapilla, S. Puri, P. Simard, High performance convolutional neural networks for document processing, in: Tenth International Workshop on Frontiers in Handwriting Recognition, 2006, pp. 1–7.
- [85] R. Raina, A. Madhavan, A.Y. Ng, Large-scale deep unsupervised learning using graphics processors, Proceedings of the 26th International Conference on Machine Learning, ACM, 2009, pp. 873–880.
- [86] D. Cireşan, U. Meier, J. Schmidhuber, Multi-column deep neural networks for image classification. ArXiv preprint arXiv:12022745, 2012.
- [87] V. Nair, G.E. Hinton, Rectified linear units improve restricted boltzmann machines, Proceedings of the 27th International Conference on Machine Learning, 2010, pp. 807–814.
- [88] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization. ArXiv preprint arXiv:1412.6980, 2014.
- [89] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: a simple way to prevent neural networks from overfitting, *Journal of Machine Learning Research* 15 (1) (2014) 1929–1958.
- [90] S. Ioffe, Szegedy C. Batch, Normalization: accelerating deep network training by reducing internal covariate shift, International Conference on Machine Learning, 2015, pp. 448–456.
- [91] A. Koutsoukas, K.J. Monaghan, X. Li, J. Huan, Deep-learning: investigating deep neural networks hyper-parameters and comparison of performance to shallow methods for modeling bioactivity data, *Journal of Cheminformatics* 9 (1) (2017) 42.
- [92] J. Ma, R.P. Sheridan, A. Liaw, G.E. Dahl, V. Svetnik, Deep neural nets as a method for quantitative structure–activity relationships, *Journal of Chemical Information and Modeling* 55 (2) (2015) 263–274.
- [93] G.B. Goh, C. Siegel, A. Vishnu, N.O. Hodas, N. Baker, Chemception: a deep neural network with minimal chemistry knowledge matches the performance of expert-developed QSAR/QSPR models, ArXiv preprint arXiv:170606689, 2017.
- [94] Y. LeCun, B.E. Boser, J.S. Denker, D. Henderson, R.E. Howard, W.E. Hubbard, et al., Handwritten digit recognition with a back-propagation network, in: Advances in Neural Information Processing Systems, 1990, pp. 396–404.
- [95] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, et al., Gradient-based learning applied to document recognition, Proceedings of the IEEE 86 (11) (1998) 2278–2324.
- [96] F. Ning, D. Delhomme, Y. LeCun, F. Piano, L. Bottou, P.E. Barbano, Toward automatic phenotyping of developing embryos from videos, *IEEE Transactions on Image Processing* 14 (2005) 1360–1371.
- [97] A. Krizhevsky, I. Sutskever, G.E. Hinton, ImageNet classification with deep convolutional neural networks, in: Advances in Neural Information Processing Systems, 2012, pp. 1097–1105.
- [98] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, ArXiv preprint arXiv:14091556, 2014.
- [99] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 770–778.
- [100] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, Z. Wojna, Rethinking the inception architecture for computer vision, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 2818–2826.
- [101] A.G. Howard M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, et al., MobileNets: Efficient convolutional neural networks for mobile vision applications, ArXiv preprint arXiv:170404861, 2017.
- [102] M. Tan, Q.V. Le, EfficientNet: rethinking model scaling for convolutional neural networks, ArXiv preprint arXiv:190511946, 2019.
- [103] R. Pascanu, T. Mikolov, Y. Bengio, On the difficulty of training recurrent neural networks, International Conference on Machine Learning, 2013, pp. 1310–1318.
- [104] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Computation* 9 (8) (1997) 1735–1780.
- [105] R. Jozefowicz, W. Zaremba, I. Sutskever, An empirical exploration of recurrent network architectures, International Conference on Machine Learning, 2015, pp. 2342–2350.
- [106] D. Merk, L. Friedrich, F. Frisoni, G. Schneider, De novo design of bioactive small molecules by artificial intelligence, *Molecular Informatics* 37 (1–2) (2018) 1700153.
- [107] R. Todeschini, V. Consonni, Molecular descriptors for chemoinformatics: volume I: alphabetical listing/volume II: appendices, references, vol. 41, John Wiley & Sons, 2009.
- [108] A. Mauri, V. Consonni, M. Pavan, R. Todeschini, Dragon software: An easy approach to molecular descriptor calculations, *Match* 56 (2) (2006) 237–248.

- [109] H. Moriwaki, Y.S. Tian, N. Kawashita, T. Takagi, Mordred: a molecular descriptor calculator, *Journal of Cheminformatics* 10 (1) (2018) 4.
- [110] A. Keller, R.C. Gerkin, Y. Guan, A. Dhurandhar, G. Turu, B. Szalai, et al., Predicting human olfactory perception from chemical features of odor molecules, *Science* 355 (6327) (2017) 820–826.
- [111] D. Rogers, M. Hahn, Extended-connectivity fingerprints, *Journal of Chemical Information and Modeling* 50 (5) (2010) 742–754.
- [112] S.D. Axen, X.P. Huang, E.L. Caceres, L. Gendele, B.L. Roth, M.J. Keiser, A simple representation of three-dimensional molecular structure, *Journal of Medicinal Chemistry* 60 (17) (2017) 7393–7409.
- [113] C. Da, D. Kireev, Structural protein–ligand interaction fingerprints (SPLIF) for structure-based virtual screening: method and benchmark study, *Journal of Chemical Information and Modeling* 54 (9) (2014) 2555–2561.
- [114] D. Smiles Weininger, A chemical language and information system. 1. Introduction to methodology and encoding rules, *Journal of Chemical Information and Computer Sciences* 28 (1) (1988) 31–36.
- [115] M. Hirohara, Y. Saito, Y. Koda, K. Sato, Y. Sakakibara, Convolutional neural network based on SMILES representation of compounds for detecting chemical motif, *BMC Bioinformatics* 19 (19) (2018) 526.
- [116] R. Huang, M. Xia, D.T. Nguyen, T. Zhao, S. Sakamuru, J. Zhao, et al., Tox21Challenge to build predictive models of nuclear receptor and stress response pathways as mediated by exposure to environmental chemicals and drugs, *Frontiers in Environmental Science* 3 (2016) 85.
- [117] G.B. Goh, N.O. Hodas, C. Siegel, A. Vishnu, SMILES2vec: An interpretable general-purpose deep neural network for predicting chemical properties, *ArXiv preprint arXiv:1712.02034*, 2017.
- [118] J.S. Delaney, ESOL: estimating aqueous solubility directly from molecular structure, *Journal of Chemical Information and Computer Sciences* 44 (3) (2004) 1000–1005.
- [119] M. Ragoza, J. Hochuli, E. Idrobo, J. Sunseri, D.R. Koes, Protein–ligand scoring with convolutional neural networks, *Journal of Chemical Information and Modeling* 57 (4) (2017) 942–957.
- [120] D.R. Koes, M.P. Baumgartner, C.J. Camacho, Lessons learned in empirical scoring with smina from the CSAR 2011 benchmarking exercise, *Journal of Chemical Information and Modeling* 53 (8) (2013) 1893–1904.
- [121] J.G. Meyer, S. Liu, L.J. Miller, J.J. Coon, A. Gitter, Learning drug function from chemical structure with convolutional neural networks and random forests, *BioRxiv* (2019;), 482877 .
- [122] D.K. Duvenaud, D. Maclaurin, J. Iparraguirre, R. Bombarell, T. Hirzel, A. Aspuru-Guzik, et al., Convolutional networks on graphs for learning molecular fingerprints, *Advances in Neural Information Processing Systems* (2015) 2224–2232.
- [123] J. Gilmer, S.S. Schoenholz, P.F. Riley, O. Vinyals, G.E. Dahl, Neural message passing for quantum chemistry, *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 1263–1272.
- [124] Y. Bian, Y. Jing, L. Wang, S. Ma, J.J. Jun, X.Q. Xie, Prediction of orthosteric and allosteric regulations on cannabinoid receptors using supervised machine learning classifiers, *Molecular Pharmacology* (2019).
- [125] Q. Shen, G. Wang, S. Li, X. Liu, S. Lu, Z. Chen, et al., ASD v3. 0: unraveling allosteric regulation with structural mechanisms and biological networks, *Nucleic Acids Research* 44 (D1) (2015) D527–D535.
- [126] J. Wu, Q. Zhang, W. Wu, T. Pang, H. Hu, W.K. Chan, et al., WDL-RF: predicting bioactivities of ligand molecules acting with G protein-coupled receptors by combining weighted deep learning and random forest, *Bioinformatics* 34 (13) (2018) 2271–2282.
- [127] U. Consortium, UniProt: a worldwide hub of protein knowledge, *Nucleic Acids Research* 47 (D1) (2019) D506–D515.
- [128] W.K. Chan, H. Zhang, J. Yang, J.R. Brender, J. Hur, A. Özgür, et al., GLASS: a comprehensive database for experimentally validated GPCR–ligand associations, *Bioinformatics* 31 (18) (2015) 3035–3042.
- [129] B. Sanchez-Lengeling, J.N. Wei, B.K. Lee, R.C. Gerkin, A. Aspuru-Guzik, A.B. Wiltschko, Machine learning for scent: learning generalizable perceptual representations of small molecules, *ArXiv preprint arXiv:1910.06855*, 2019.
- [130] C.Y. Su, K. Menuz, J.R. Carlson, Olfactory perception: receptors, cells, and circuits, *Cell* 139 (1) (2009) 45–59.
- [131] A. Paul, D. Jha, R. Al-Bahrani, W.K. Liao, A. Choudhary, A. Agrawal, CheMixNet: Mixed DNN architectures for predicting chemical properties using multiple molecular representations, *ArXiv preprint arXiv:1811.08283*, 2018.
- [132] Z. Xu, S. Wang, F. Zhu, J. Huang, Seq2Seq fingerprint: An unsupervised deep molecular embedding for drug discovery, *Proceedings of the 8th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, ACM, 2017, pp. 285–294.
- [133] R. Gómez-Bombarelli, J.N. Wei, D. Duvenaud, J.M. Hernández-Lobato, B. Sánchez-Lengeling, D. Sheberla, et al., Automatic chemical design using a data-driven continuous representation of molecules, *ACS Central Science* 4 (2) (2018) 268–276.
- [134] M.A. Johnson, G.M. Maggiora, *Concepts and Applications of Molecular Similarity*, Wiley, 1990.
- [135] I. Vogt, H.E. Ahmed, J. Auer, J. Bajorath, Exploring structure–selectivity relationships of biogenic amine GPCR antagonists using similarity searching and dynamic compound mapping, *Molecular Diversity* 12 (1) (2008) 25–40.
- [136] M. Luo, X.S. Wang, A. Tropsha, Comparative Analysis of QSAR-based vs. Chemical Similarity Based Predictors of GPCRs Binding Affinity, *Molecular Informatics* 35 (1) (2016) 36–41.
- [137] D. Bajusz, A. Racz, K. Heberger, Why is Tanimoto index an appropriate choice for fingerprint-based similarity calculations? *Journal of Cheminformatics* (2015;7(20)).
- [138] P. Willett, Similarity-based virtual screening using 2D fingerprints, *Drug Discovery Today* 11 (23–24) (2006) 1046–1053.
- [139] D. Alberga, D. Triscuzzi, M. Montaruli, F. Leonetti, G.F. Mangiatordi, O. Nicolotti, A new approach for drug target and bioactivity prediction: The multifingerprint similarity search algorithm (MuSSEL), *Journal of Chemical Information and Modeling* 59 (1) (2018) 586–596.
- [140] D. Stumpfe, J. Bajorath, Exploring activity cliffs in medicinal chemistry: mini-pherchem, *Journal of Medicinal Chemistry* 55 (7) (2012) 2932–2942.
- [141] P.C. Hawkins, A.G. Skillman, A. Nicholls, Comparison of shape-matching and docking as virtual screening tools, *Journal of Medicinal Chemistry* 50 (1) (2007) 74–82.
- [142] F.A. Baltoumas, M.C. Theodoropoulou, S.J. Hamodrakas, Interactions of the α -subunits of heterotrimeric G-proteins with GPCRs, effectors and RGS proteins: a critical review and analysis of interacting surfaces, conformational shifts, structural diversity and electrostatic potentials, *Journal of Structural Biology* 182 (3) (2013) 209–218.
- [143] J.A. Javitch, J.A. Ballesteros, J. Chen, V. Chiappa, M.M. Simpson, Electrostatic and aromatic microdomains within the binding-site crevice of the D2 receptor: contributions of the second membrane-spanning segment, *Biochemistry* 38 (25) (1999) 7961–7968.
- [144] X. Yan, J. Li, Z. Liu, M. Zheng, H. Ge, J. Xu, Enhancing molecular shape comparison by weighted Gaussian functions, *Journal of Chemical Information and Modeling* 53 (8) (2013) 1967–1978.
- [145] A.E. Cleves, S.R. Johnson, A.N. Jain, Electrostatic-field and surface-shape similarity for virtual screening and pose prediction, *Journal of Computer-aided Molecular Design* 33 (10) (2019) 865–886.
- [146] M.M. Mysinger, M. Carchia, J.J. Irwin, B.K. Shoichet, Directory of useful decoys, enhanced (DUD-E): better ligands and decoys for better benchmarking, *Journal of Medicinal Chemistry* 55 (14) (2012) 6582–6594.
- [147] G. Hu, G. Kuang, W. Xiao, W. Li, G. Liu, Y. Tang, Performance evaluation of 2D fingerprint and 3D shape similarity methods in virtual screening, *Journal of Chemical Information and Modeling* 52 (5) (2012) 1103–1113.
- [148] J. Shang, X. Dai, Y. Li, M. Pistolozzi, L. Wang, HybridSim-VS: a web server for large-scale ligand-based virtual screening using hybrid similarity recognition techniques, *Bioinformatics* 33 (21) (2017) 3480–3481.
- [149] D. Reker, T. Rodrigues, P. Schneider, G. Schneider, Identifying the macro-molecular targets of de novo-designed chemical entities through self-organizing map consensus, *Proceedings of the National Academy of Sciences* 111 (11) (2014) 4067–4072.
- [150] M. Reutlinger, C.P. Koch, D. Reker, N. Todoroff, P. Schneider, T. Rodrigues, et al., Chemically advanced template search (CATS) for scaffold-hopping and prospective target prediction for ‘orphan’ molecules, *Molecular Informatics* 32 (2) (2013) 133–138.
- [151] T. Rodrigues, M. Werner, J. Roth, E.H. da Cruz, M.C. Marques, P. Akkapeddi, et al., Machine intelligence decrypts β -lapachone as an allosteric 5-lipoxygenase inhibitor, *Chemical Science* 9 (34) (2018) 6899–6903.
- [152] T. Rodrigues, B.P. de Almeida, N.L. Barbosa-Morais, G.J. Bernardes, Dissecting celastrol with machine learning to unveil dark pharmacology, *Chemical Communications* 55 (45) (2019) 6369–6372.
- [153] M.C. Lagerström, H.B. Schiöth, Structural diversity of G protein-coupled receptors and significance for drug discovery, *Nature Reviews Drug Discovery* 7 (4) (2008) 339–357.
- [154] M. Nayal, B. Honig, On the nature of cavities on protein surfaces: application to the identification of drug-binding sites, *Proteins: Structure, Function, and Bioinformatics* 63 (4) (2006) 892–906.
- [155] J. Yang, A. Roy, Y. Zhang, Protein–ligand binding site recognition using complementary binding-specific substructure comparison and sequence profile alignment, *Bioinformatics* 29 (20) (2013) 2588–2595.
- [156] J. Jiménez, M. Skalic, G. Martínez-Rosell, G.K. De Fabritiis, DEEP: protein–ligand absolute binding affinity prediction via 3D-convolutional neural networks, *Journal of Chemical Information and Modeling* 58 (2) (2018) 287–296.
- [157] Y. Li, M.A. Rezaei, C. Li, X. Li, D. Wu, DeepAtom: A Framework for Protein–Ligand Binding Affinity Prediction, *ArXiv preprint arXiv:1912.00318*, 2019.
- [158] K.T. Schütt, F. Arbabzadah, S. Chmiela, K.R. Müller, A. Tkatchenko, Quantum-chemical insights from deep tensor neural networks, *Nature Communications* 8 (2017) 13890.
- [159] I. Kozlovskii, P. Popov, Spatiotemporal identification of druggable binding sites using deep learning, *bioRxiv*, 2020.
- [160] H.M. Berman, J. Westbrook, Z. Feng, G. Gilliland, T.N. Bhat, H. Weissig, et al., The protein data bank, *Nucleic Acids Research* 28 (1) (2000) 235–242.
- [161] J. Li, A. Fu, L. Zhang, An overview of scoring functions used for protein–ligand interactions in molecular docking, *Interdisciplinary Sciences: Computational Life Sciences* (2019) 1–9.
- [162] M.L. Verdonk, J.C. Cole, M.J. Hartshorn, C.W. Murray, R.D. Taylor, Improved protein–ligand docking using GOLD, *Proteins: Structure, Function, and Bioinformatics* 52 (4) (2003) 609–623.
- [163] S. Vilar, G. Cozza, S. Moro, Medicinal chemistry and the molecular operating environment (MOE): application of QSAR and molecular docking to drug discovery, *Current Topics in Medicinal Chemistry* 8 (18) (2008) 1555–1572.
- [164] H. Li, J. Peng, P. Sidorov, Y. Leung, K.S. Leung, M.H. Wong, et al., Classical scoring functions for docking are unable to exploit large volumes of structural and interaction data, *Bioinformatics* (2019).
- [165] P.J. Ballester, A. Schreyer, T.L. Blundell, Does a more precise chemical description of protein–ligand complexes lead to more accurate prediction of binding affinity? *Journal of Chemical Information and Modeling* 54 (3) (2014) 944–955.
- [166] H. Li, K.S. Leung, M.H. Wong, P.J. Ballester, Improving AutoDock Vina using

- random forest: the growing accuracy of binding affinity prediction by the effective exploitation of larger data sets, *Molecular Informatics* 34 (2–3) (2015) 115–126.
- [167] M. Wójcikowski, P.J. Ballester, P. Siedlecki, Performance of machine-learning scoring functions in structure-based virtual screening, *Scientific Reports* 7 (2017) 46710.
- [168] J. Sunseri, M. Ragoza, J. Collins, D.R. Koes, A D3R prospective evaluation of machine learning for protein–ligand scoring, *Journal of Computer-aided Molecular Design* 30 (9) (2016) 761–771.
- [169] L. Zheng, J. Fan, Y. Mu, OnionNet: a multiple-layer inter-molecular contact based convolutional neural network for protein–ligand binding affinity prediction, *ACS Omega* 4 (14) (2019) 15956–15965.
- [170] M.M. Stepniewska-Dziubinska, P. Zielenkiewicz, P. Siedlecki, Development and evaluation of a deep learning model for protein–ligand binding affinity prediction, *Bioinformatics* 34 (21) (2018) 3666–3674.
- [171] R. Wang, X. Fang, Y. Lu, C.Y. Yang, S. Wang, The PDBbind database: methodologies and updates, *Journal of Medicinal Chemistry* 48 (12) (2005) 4111–4119.
- [172] N. Ma, X. Zhang, H.T. Zheng, J. Sun, Shufflenet v2: Practical guidelines for efficient cnn architecture design, *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 116–131.
- [173] F.N. Iandola, S. Han, M.W. Moskewicz, K. Ashraf, W.J. Dally, K. Keutzer, SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and < 0.5 MB model size, *ArXiv preprint arXiv:160207360*, 2016.
- [174] O. Dagliyan, E.A. Proctor, K.M. D’Auria, F. Ding, N.V. Dokholyan, Structural and dynamic determinants of protein–peptide recognition, *Structure* 19 (12) (2011) 1837–1845.
- [175] G.E. Hinton, R.R. Salakhutdinov, Reducing the dimensionality of data with neural networks, *Science* 313 (5786) (2006) 504–507.
- [176] V. Mirjalili, S. Raschka, A. Nambodiri, A. Ross, Semi-Adversarial Networks: Convolutional autoencoders for imparting privacy to face images, *Proceedings of 11th IAPR International Conference on Biometrics, IEEE, Gold Coast, Australia, 2018*.
- [177] V. Mirjalili, S. Raschka, A. Ross, FlowSAN: privacy-enhancing semi-adversarial networks to confound arbitrary face-based gender classifiers, *IEEE Access* 7 (2019) 99735–99745.
- [178] V. Mirjalili, S. Raschka, A. Ross, Gender privacy: An ensemble of semi adversarial networks for confounding arbitrary gender classifiers, *Proceedings of the 9th International Conference on Biometrics: Theory, Applications and Systems, IEEE, Los Angeles, USA, 2018*.
- [179] V. Mirjalili, S. Raschka, A. Ross, PrivacyNet: semi-adversarial networks for multi-attribute face privacy, *ArXiv preprint arXiv:200100561*, 2020.
- [180] D.P. Kingma, M. Welling, Auto-encoding variational bayes, *ArXiv preprint arXiv:1312.6114*, 2013.
- [181] Doersch C. Tutorial on variational autoencoders. *ArXiv preprint arXiv:160605908*. 2016;.
- [182] M.J. Kusner, B. Paige, J.M. Hernández-Lobato, Grammar variational autoencoder, *Proceedings of the 34th International Conference on Machine Learning. JMLR. org*, 2017, pp. 1945–1954.
- [183] W. Jin, R. Barzilay, T. Jaakkola, Junction tree variational autoencoder for molecular graph generation, *ArXiv preprint arXiv:180204364*, 2018.
- [184] M. Simonovsky, N. Komodakis, Graphvae: Towards generation of small graphs using variational autoencoders, *International Conference on Artificial Neural Networks, Springer, 2018*, pp. 412–422.
- [185] H. Dai, Y. Tian, B. Dai, S. Skiena, L. Song, Syntax-directed variational autoencoder for structured data, *ArXiv preprint arXiv:180208786*, 2018.
- [186] W. Wiegand, Variational approximations between mean field theory and the junction tree algorithm, *Proceedings of the 16 Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann Publishers Inc., 2000*, pp. 626–633.
- [187] R.S. Sutton, A.G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, 2018.
- [188] R.J. Williams, Simple statistical gradient-following algorithms for connectionist reinforcement learning, *Machine Learning* 8 (3–4) (1992) 229–256.
- [189] X. Liu, K. Ye, H.W. van Vlijmen, A.P. IJzerman, G.J. van Westen, An exploration strategy improves the diversity of de novo ligands using deep reinforcement learning: a case for the adenosine A2A receptor, *Journal of Cheminformatics* 11 (1) (2019) 35.
- [190] J.F. Chen, H.K. Eltzschig, B.B. Fredholm, Adenosine receptors as drug targets—what are the challenges? *Nature Reviews Drug Discovery* 12 (4) (2013) 265.
- [191] Z. Zhou, S. Kearnes, L. Li, R.N. Zare, P. Riley, Optimization of molecules via deep reinforcement learning, *Scientific Reports* 9 (1) (2019) 1–10.
- [192] J. You, B. Liu, Z. Ying, V. Pande, J. Leskovec, Graph convolutional policy network for goal-directed molecular graph generation, in: *Advances in Neural Information Processing Systems, 2018*, pp. 6410–6421.
- [193] V. Mnih, K. Kavukcuoglu, D. Silver, A.A. Rusu, J. Veness, M.G. Bellemare, et al., Human-level control through deep reinforcement learning, *Nature* 518 (7540) (2015) 529.
- [194] P. Leeson, Drug discovery: Chemical beauty contest, *Nature* 481 (7382) (2012) 455.
- [195] S.J. Pan, Q. Yang, A survey on transfer learning, *IEEE Transactions on Knowledge and Data Engineering* 22 (10) (2009) 1345–1359.
- [196] F. Imrie, A.R. Bradley, M. van der Schaar, C.M. Deane, Protein family-specific models using deep neural networks and transfer learning improve virtual screening and highlight the need for more data, *Journal of Chemical Information and Modeling* 58 (11) (2018) 2319–2330.
- [197] G.A. Ross, G.M. Morris, P.C. Biggin, One size does not fit all: the limits of structure-based models in drug discovery, *Journal of Chemical Theory and Computation* 9 (9) (2013) 4266–4274.
- [198] G.B. Goh, C. Siegel, A. Vishnu, N. Hodas, Using rule-based labels for weak supervised learning: a ChemNet for transferable chemical property prediction, *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, ACM, 2018*, pp. 302–310.
- [199] J.S. Smith, B.T. Nebgen, R. Zubatyuk, N. Lubbers, C. Devereux, K. Barros, et al., Approaching coupled cluster accuracy with a general-purpose neural network potential through transfer learning, *Nature Communications* 10 (1) (2019) 1–8.
- [200] R. Munro, Human-in-the-loop machine learning, *Manning* (2020).
- [201] Howard J, Ruder S. Universal language model fine-tuning for text classification. *ArXiv preprint arXiv:180106146*. 2018;.
- [202] M. Noroozi, P. Favaro, Unsupervised learning of visual representations by solving jigsaw puzzles, *European Conference on Computer Vision, Springer, 2016*, pp. 69–84.
- [203] A. Shiraishi, T. Okuda, N. Miyasaka, T. Osugi, Y. Okuno, J. Inoue, et al., Repertoires of G protein-coupled receptors for Ciona-specific neuropeptides, *Proceedings of the National Academy of Sciences* 116 (16) (2019) 7847–7856.
- [204] A.W. Senior, R. Evans, J. Jumper, J. Kirkpatrick, L. Sifre, T. Green, et al., Improved protein structure prediction using potentials from deep learning, *Nature* (2020) 1–5.