



Generative Network Complex for the Automated Generation of Drug-like Molecules

Kaifu Gao, Duc Duy Nguyen, Meihua Tu, and Guo-Wei Wei*



properties but also are cheaply available to low-income people. In this work, we develop a generative network complex (GNC) to generate new drug-like molecules based on the multiproperty optimization via the gradient descent in the latent space of an autoencoder. In our GNC, both multiple chemical properties and similarity scores are optimized to generate drug-like molecules with desired chemical properties. To further validate the reliability of the predictions, these molecules are reevaluated and screened by independent 2D fingerprint-based predictors to come up with a few hundreds of new drug candidates. As a demonstration, we



apply our GNC to generate a large number of new BACE1 inhibitors, as well as thousands of novel alternative drug candidates for eight existing market drugs, including Ceritinib, Ribociclib, Acalabrutinib, Idelalisib, Dabrafenib, Macimorelin, Enzalutamide, and Panobinostat.

I. INTRODUCTION

Drug discovery ultimately tests our understanding of molecular biology, medicinal chemistry, genetics, physiology, and pharmacology, the status of biotechnology, the utility of computational sciences, and the maturity of mathematical biology. Technically, drug discovery involves target discovery, lead discovery, lead optimization, preclinical development, three phases of clinical trials, and, finally, launching to the market only if a drug can be demonstrated to be safe and effective in every stage. Among them, lead discovery, lead optimization, and preclinical development disqualify tens of thousands of compounds based on their binding affinity, solubility, partition coefficient, clearance, permeability, toxicity, pharmacokinetics, etc., leaving only about tens of them to clinical trials. Therefore, drug discovery is expensive and time-consuming currently: it takes about \$2.6 billion dollars and more than ten years, on average, to bring a new drug to the market.¹ Reducing the expense and speeding up the process is one of the top priorities of the pharmaceutical industry.

One of the key challenges in small molecule drug discovery is to find novel chemical compounds with desirable properties. Much effort has been taken to optimize this critical step in the drug discovery pipeline. For example, the development of highthroughput screening (HTS) has led to an unprecedented increase in the number of potential targets and leads.² HTS can quickly conduct millions of tests to identify active compounds of interest from compound libraries.³ While there has been an increase in the number of potential targets and leads, the number of newly generated molecular entities has remained stable because of a high attrition rate by the elimination of leads with inappropriate physicochemical or pharmacological properties during preclinical development and clinical phases.^{4,5} Rational drug design (RDD) approaches are proposed to better identify candidates with the highest probability of success.⁶ These methods aim at finding new medications based on the knowledge of biologically druggable targets.^{1,7}

More recently, computer-aided drug design (CADD) has emerged as a useful approach in reducing the expense and period of drug discovery.⁸ Computational techniques have been developed for both virtual screening (VS) and optimizing ADME properties of lead compounds. Primarily, these methods are designed as *in silico* filters to eliminate compounds with undesirable properties. These filters are widely applied for the assembly of compound libraries using combinatorial chemistry.⁹ The integration of early ADME profiling of lead chemicals has contributed to speeding up lead selection for phase-I trials without large amounts of revenue loss.¹⁰ Currently, compounds are added in libraries based on target-focused design or diversity

Special Issue:Generative Models for Molecular DesignReceived:May 30, 2020Published:July 20, 2020





considerations.¹¹ VS and HTS can screen compound libraries to a subset of compounds whose properties are in agreement with various criteria.¹²

Despite these efforts, current databases of chemical compounds remain small when compared with the chemical space spanned by all possible energetically stable stoichiometric combinations of atoms and topologies in molecules. It is estimated that there are about 10⁶⁰ distinct molecules; among them, roughly 10³⁰ are drug-like.³ As a result, computational techniques are also being developed for *de novo* design of drug-like molecules¹³ and generating large virtual chemical libraries, which can be screened more efficiently for *in silico* drug discovery.

Among the available computational techniques, deep neural networks (DNNs) have gathered much interest for their ability to extract features and learn physical principles from training data. Currently, DNN-based architectures have been successfully applied in a wide variety of fields in the biological and biomedical sciences.^{14,15}

More interestingly, many deep generative models based on sequence-to-sequence autoencoders (Seq2seq AEs),¹⁶ variational autoencoders (VAEs),¹⁷ adversarial autoencoders (AAEs),¹⁸ generative adversarial networks (GANs),¹⁹ reinforcement learning,²⁰ etc. have been proposed for exploring the vast drug-like chemical space and generating new drug-like molecules. Winter et al.^{21,22} performed the optimization based on particle swarm optimization on the continuous latent space of a Seq2seq AE to generate new molecules with desired properties. Gomez-Bombarelli et al.²³ used a VAE to encode a molecule in the continuous latent space for exploring associated properties. Skalic et al.²⁴ combined a conditional VAE and a captioning network to generate previously unseen compounds from input voxelized molecular representations. Kadurin et al.²⁵ built an AAE to create new compounds. Sattarov et al.²⁶ combined deep autoencoder RNNs with generative topographic mapping to carry out de novo molecule design. A policy-based reinforcement learning approach was proposed to tune RNNs for episodic tasks,^{27,28} and extended to design desired molecules.²⁹ Zhou et al.³⁰ also proposed a strategy to optimize molecules by combining domain knowledge of chemistry and state-of-theart reinforcement learning techniques.

However, the generative strategies mentioned above are not drug-specified. What is vital to drug discovery is to design potential drug candidates for specific drug targets. In a regular drug discovery procedure, the starting point is target identification, followed by lead generation. Then, lead optimization is performed to make lead compounds more drug-like.²

It is useful to find new lead compounds to replace existing market drugs for several reasons. First, existing market drugs might not be optimal. For example, they might not be potent enough, be too toxic and harmful to human health, or be too hard to synthesize. Additionally, they might have side effects, insufficient aqueous solubility (log S) that reduces the absorption of drugs,³¹ or higher partition coefficients (log P) that lead to improper drug distributions within the body.³² Moreover, for very expensive drugs, it is desirable to find cheap alternatives.

With the generation of new alternative lead compounds in mind, one can make use of existing drug data sets to develop drug-specified generative models. In this process, it is critical to apply a similarity restraint to generate hundreds or even thousands of new drug-like molecules inside the chemical space close to the reference molecule. This similarity restraint enables us to generate new molecules that remain effective to the target. Moreover, from the viewpoint of machine learning, higher similarities to existing data always lead to more reliable predictions in generating molecules. The generative model can also realize lead optimization: by incorporating optimizers, generated molecules are designated to have one or more chemical properties better than the reference molecule. As a result, a large number of alternative drug candidates are created by drug-specified generative models. These candidates could be an effective and specified library to further screen for better or cheaper drug alternatives.

Therefore, in this work, we develop a generative network complex (GNC) based on the multiple-property optimization via gradient descent in the latent space to automatically generate new drug-like molecules. One workflow of our GNC consists of three following stages

- 1. The SMILES string of a seed molecule are encoded into a vector in the latent space by a pretrained encoder.
- Starting with the latent vector of the seed molecule, a DNN-based drug-like molecule generator modifies the vector via gradient descent, so that new latent vectors that satisfy multiple property restraints including chemical properties and similarity scores to the desired objectives are created.
- 3. A pretrained decoder decodes these new latent vectors into the SMILES strings of newly generated molecules.

The remainder of the paper is organized as follows. Section II introduces our new GNC framework formed by the seq2seq AE and drug-like molecule generator. Section III discusses its reliability test on the BACE1 target and, more importantly, presents the performance of our GNC on a variety of existing drug targets. Insight into the roles of the multiple property restraints is offered in Section IV. The conclusion is given in Section V.

II. METHODS

II.A. Sequence to Sequence Autoencoder (seq2seq AE). The seq2seq model is an autoencoder architecture originated from natural language processing.¹⁶ It has been demonstrated as a breakthrough in language translation. The basic strategy of the seq2seq AE is to map an input sequence to a fixed-sized vector in the latent space using a gated recurrent unit (GRU)³³ or a long short-term memory (LSTM) network,³⁴ and then map the vector to a target sequence with another GRU or LSTM network. Thus, the latent vector is an intermediate representation containing the "meaning" of the input sequence.

In our case, input and output sequences are both SMILES strings—a one-dimensional "language" of chemical structures.³⁵ The seq2seq AE is trained to have a high reconstruction ratio between inputs and outputs so that the latent vectors contain faithful information on the chemical structures (see the upper part of Figure 1). Here we utilize a pretrained seq2seq model from a recent work.²²

II.B. Drug-like Molecule Generator Based on the Multiproperty Optimization. In our new GNC, we design a drug-like molecule generator elaborately, so that generated molecules not only satisfy desired properties but also share common pharmacophores with reference compounds. From a seed molecule, one generative workflow of the GNC is depicted in Figure 1 and described as below.



Figure 1. Schematic illustration of our generative network complex.

- 1. Randomly pick a low-binding-affinity molecule from a target-specified training set as the seed, and then the SMILES string of the seed molecule is encoded by a pretrained encoder (in our case a GRU encoder) into a latent vector.
- 2. The latent vector of the seed is fed into our DNN molecule generator. In every epoch, the generator comes up with a new vector $X \in \mathbb{R}^n$, and the deep learning network is instructed to evaluate X via the following loss function

$$\mathcal{L}(X) = \frac{1}{n} \sum_{i=1}^{n} k_i |\hat{y}_i(X) - y_{i0}|$$
(1)

where k_i is the *i*th predefined weight serving the purpose of emphasizing or deemphasizing different property restraints and $\hat{y}_i(X)$ is the predicted *i*th property value by a pretrained predictor \mathcal{M}_i . Additionally, y_{i0} is the objective value of the *i*th property. The restrained properties can be binding affinity (BA), the similarity score (Sim) to a reference molecule or others such as partition coefficient (Log P), Lipinski's rule of five, ³⁶ etc. A guideline for y_{i0} includes, in the BA restraint, often setting $y_{\Delta G} < -9.6$ kcal/mol; in the Log P condition, it is common to set $y_{\text{LogP}} < 5$, etc.

- 3. Gradient descent is used to minimize the loss function in eq 1 until the maximum number of epochs is reached.
- 4. The generated latent vectors satisfying the desired restraints are decoded into SMILES strings through a pretrained decoder, as shown in Figure 1.

To create a variety of novel drug-like molecules originated from leads or existing drugs (reference molecules), one can adopt different seed molecules as well as different objective values to achieve desired properties and similarity scores. The ultimate purpose of our molecule generator is to keep modifying the latent vector to satisfy the multiple druggable property restraints. Figure 2 illustrates the general mechanism of our generator. Notably, the pretrained predictor weights inside our model stay intact throughout the backpropagation of the training process to the generator. The loss function converges when all conditions are met, and then, resulting latent vectors are decoded into SMILES strings.

II.C. Parameters of the Molecule Generator. In our model, the dimension of the latent space is 512, so the input and output dimensions of the DNN molecule generator are also 512. The DNN generator has two hidden layers, with 1024 neurons



Figure 2. Illustration of the latent-space molecule generator. Here, *X* is a latent vector representing a molecule and ΔX is the negative gradient of the loss function.

in each layer. The activation function is tanh, the learning rate is 0.1, and the momentum is also 0.1. In this work, we are interested in binding affinity and similarity score restraints. The regularization coefficients of these two restraints ($k_{\Delta G}$ and $k_{\rm Sim}$) are set to 1 and 10, respectively. The similarity score restraint is determined via the Tanimoto coefficient between a generated latent vector and the latent vector of a reference molecule.

The binding affinity restraint relies on pretrained binding affinity predictors. A pretrained binding affinity predictor (LV-BP) takes latent vectors as its inputs and return predicted binding affinities. Therefore, typically, the input dimension of the predictor is 512, and the output dimension is 1. The DNN predictor has three hidden layers with 1024, 1536, and 1024 neurons, respectively. The ReLU activation function is applied. The learning rate is 0.001, the number of training epochs is 4000, and the batch size is 4. The predictor network is trained on target-specified data sets carefully selected from public databases such as ChEMBL.³⁷ The generator and predictor are both programmed in the framework of PyTorch (Version 1.0.0).³⁸

In the current work, for each generation task, we randomly pick 50 low-binding-affinity molecules from the preselected data set as seeds. For each seed, the generative network is run in a total of 2000 epochs, which takes less than 10 min under the supercomputer equipped with one Nvidia K80 GPU card. In practice, to quickly fill up the potential chemical search space, one can use more seeds and run more epochs for each seed.

II.D. Binding Affinity Reevaluation by the 2D-Fingerprint Predictor. Besides generating new molecules, the LV-BP in our GNC also predicts their binding affinities. However, no experimental values are available to validate these predicted affinities. Therefore, we cross-validate them using alternative binding affinity predictors. In the present work, we construct machine learning predictors based on 2D fingerprints (2DFP-BPs) to reevaluate the affinities of generated compounds. The 2D fingerprints computed from their SMILES strings are inputs to these 2DFP-BPs. If the predictions from the LV-BP and 2DFP-BPs are consistent, we regard the predictions as reliable.

According to our previous tests,³⁹ the consensus of ECFP4,⁴⁰ Estate1⁴¹ and Estate2⁴¹ fingerprints performs best on bindingaffinity prediction tasks. Therefore, this work also makes use of this consensus. We employ the RDKit software (version 2018.09.3)⁴² to generate 2D fingerprints from SMILES strings. Since the training sets in our current cases are not so large, we apply gradient boosting decision tree (GBDT)⁴³ model due to its accuracy and speed when handling small data sets. This

pubs.acs.org/jcim

Table 1. Information a	bout the Eight Marl	cet Drugs Used in t	he Present Study
------------------------	---------------------	---------------------	------------------

Drug name	ChEMBL ID	Treatment	FDA approval year	ΔG (kcal/mol)	Filtered training set size	ΔG range of training set (kcal/mol)
Ceritinib	2403108	Nonsmall cell lung cancer	2014	-10.77	1203	-5.26 to -13.93
Ribociclib	3545110	Breast cancer	2017	-10.98 and -10.16	918 and 289	-6.31 to -12.98 and -4.11 to -14.13
Acalabrutinib	3707348	Mantle cell lymphoma	2017	-11.67	1451	-4.21 to -14.05
Idelalisib	2216870	Blood cancers	2014	-10.43	1959	-1.92 to -14.16
Macimorelin	278623	Adult growth hormone deficiency	2017	-10.48	608	-4.18 to -14.68
Dabrafenib	2028663	Cancers with a mutated gene BRAF	2013	-12.35	2254	-5.49 to -14.68
Enzalutamide	1082407	Prostate cancer	2012	-10.53	1386	-3.83 to -13.72
Panobinostat	483254	Cancers	2015	-12.46	1645	-2.91 to -12.46

GBDT predictor is constructed using the gradient boosting regressor module in scikit-learn (version 0.20.1)⁴⁴ and the following parameters: n_estimators = 10000, max_depth = 7, min_samples_split = 3, learning_rate = 0.01, subsample = 0.3, and max_features = sqrt.

The criteria used in our reevaluation are the root-mean-square error (RMSE), Pearson correlation coefficient (R_P), and active ratio. Here, the active ratio means the ratio of the number of the active molecules indicated by both the 2DFP-BP and LV-BP to the number of the active ones indicated by the LV-BP.

II.E. Multitask DNN Predictors. Multitask DNN predictors⁴⁵ for both latent vectors and 2D fingerprints are built for the drug Ribociclib with two different targets.

The latent-vector based model has three hidden layers with 1024, 1536, and 1024 neurons. For the 2D-fingerprint based models, because the three different 2D fingerprints ECFP4, Estate1, and Estate2 have 2048, 79, and 79 features, respectively, two different network architectures are used. For ECFP4, the numbers of neurons in the three hidden layers are 2500, 1500, and 500, respectively. For Estate1 and Estate2, their numbers of neurons are 500, 1000, and 500. Other parameters are the same as those of our single task predictors. These multitask models are also programmed in the framework of PyTorch (Version 1.0.0).

II.F. Drug–Target Interaction and Common Pharmacophore Analysis. The interactions between drugs and their targets, as well as the pharmacophores of the drugs, are investigated. The purpose is to explore whether our generated molecules can still bind to the drug targets.

Drug-target interactions are analyzed via the protein-ligand interaction profiles.⁴⁶ It can identify drug-target interaction types such as hydrogen bonds, hydrophobic interactions, etc.

However, the interaction analysis itself could not determine whether interactions are critical or not to the drug-target binding. By using the Phase module in Schrödinger (version 2018-4),⁴⁷ we build pharmacophore models via searching common pharmacophores in all the active compounds to the target. Since these pharmacophores are widespread to all the active compounds, they are critical to the binding. Thus, if generated molecules still contain such pharmacophores, they are potential binders.

It could be time-consuming to recognize common pharmacophores of hundreds of compounds. To avoid this obstacle, we group compounds into 50 clusters via the k-means algorithm implemented by scikit-learn.⁴⁴ We, then, collect the centroids of these 50 clusters for the common pharmacophore search.

II.G. Data Sets. In this work, first, we explore the effect of different objective values in our generator on the binding-affinity

prediction reliability to generated molecules. We carry out this reliability test on the Beta-Secretase 1 (BACE1) data set.

BACE1 is a transmembrane aspartic acid protease human protein encoded by the BACE1 gene. It is essential for the generation of β -amyloid peptide in neural tissue,⁴⁸ a component of amyloid plaques widely believed to be critical in the development of Alzheimer's, rendering BACE1 an attractive therapeutic target for this devastating disease.⁴⁹ We download 3916 BACE1 compounds from the ChEMBL database. In the seq2seq autoencoder we utilized, there is a molecule filter that only selects organic molecules with more than 3 heavy atoms, their weights between 12 and 600, and their Log P values between -5 and 7.²² As a result, a total of 3151 molecules in the BACE1 data set pass this filter and are used as the training set.

More importantly, we employ our GNC to design alternative promising drug candidates for the eight drugs on the market. For each drug, we construct a data set of the compounds that bind to the same drug target from the ChEMBL database. The collected compounds are also filtered by the filter in the seq2seq autoencoder. Table 1 lists information regarding these eight drugs, namely drug name, ChEMBL ID, FDA approval year, experimental drug affinity (ΔG), filtered training set size, and affinity range of the training set.

SciFinder is one of the most comprehensive databases for chemical literature and substances (https://scifinder-n.cas.org). It contains more than 143 million organic and inorganic substances. Here, we search our generated molecules in this database to confirm they are new and have never existed before.

III. EXPERIMENTS

III.A. Designing BACE1 Inhibitors. *III.A.1. Accuracy of the seq2seq AE and Predictors.* We first test the accuracy of the seq2seq autoencoder and the LV-BP and 2DFP-BP predictors.

When performing the seq2seq model on the filtered BACE1 data set with 3151 molecules, the reconstruction ratio is 96.2%. This high ratio guarantees the essential information on these input molecules is encoded into the corresponding latent vectors.

Subsequently, these latent vectors are used as the features to train our latent-vector DNN binding affinity predictor (LV-BP); the labels are their corresponding experimental binding affinities. In a 5-fold cross-validation test on the BACE1 data set, the LV-BP achieves an average Pearson correlation coefficient (R_P) of 0.871 and an average RMSE of 0.704 kcal/mol.

The 2D-fingerprint GBDT binding affinity predictor (2DFP-BP) is used to reevaluate the predictions from the LV-BP. We also test this 2DFP-BP by the 5-fold cross-validation. The average $R_{\rm P}$ and RMSE are 0.874 and 0.692 kcal/mol, respectively, quite comparable to the LV-BP.

III.A.2. Convergence Analysis. Here we conduct our GNC to generate new molecules and analyze how these new molecules evolve over a generation course. We start this experiment with a seed molecule picked from the BACE1 data set; this molecule is far from active with binding free energy (ΔG) = -6.81 kcal/mol. The reference molecule is also from the BACE1 data set, it is highly active with ΔG = -12.02 kcal/mol. The binding affinity objective $y_{\Delta G}$ is set to be -12.02 kcal/mol, and the similarity score to the reference molecule is targeted to y_{sim} = 1.0.

Figure 3a depicts the loss function values computed at every epoch; Figure 3b and c illustrate the LV-BP predicted binding



Figure 3. Convergence of the loss function, the LV-BP predicted ΔG , and the similarity score to the reference molecule during a molecule generation course. In this example, the ΔG s of the seed molecule and the reference molecule are -6.81 kcal/mol and -12.02 kcal/mol, respectively. To force generated molecules evolving toward the reference molecule, we set the similarity score objective and the ΔG objective to be 1.00 and -12.02 kcal/mol, respectively.

affinities and similarity scores, respectively. From these figures, one can observe that our GNC produces new potential BACE1 inhibitors with desirable binding affinities in less than 3000 epochs.

Figure 4 shows a series of generated molecules over the evolution from the seed to the reference molecule. The starting point is the seed molecule, and its binding affinity and similarity score to the reference molecule are as low as -6.81 kcal/mol and 0.01, respectively. By receiving the feedback from the gradient descent in the generator, the similarity score gradually rises to 1.0. The improvement to the binding affinity is even faster: while a created molecule has a similarity score of 0.28, its LV-BP predicted ΔG already reaches -11.30 kcal/mol; while the similarity score is 0.90, the LV-BP predicted ΔG is -12.00 kcal/mol.

III.A.3. Reliability Test on the Designed BACE1 Inhibitors. Using our GNC, we also generate millions of compounds targeting a wide range of binding affinities and similarity scores. Then the prediction reliability with these different ranges of binding affinities and similarity scores is tested. Individually, the similarity score objectives, y_{sim} , vary from 0.50 to 0.95 with an increment of 0.025; the binding affinity objectives, $y_{\Delta G}$, receive values from -9.6 kcal/mol to -13.1 kcal/mol with an increment of -0.25 kcal/mol. Here we select -9.6 kcal/mol as the starting point since this value is a widely accepted threshold to identify active compounds; the end point of $\Delta G = -13.1$ kcal/mol is the highest binding affinity value in the BACE1 data set (see Figure 5).

The reliability test is based on reevaluating the LV-BP predicted binding affinities by the 2DFP-BP. The reliability criteria are the RMSE, active ratio, and $R_{\rm P}$ between the LV-BP and 2DFP-BP prediction. The heatmaps in Figure 6 show these evaluation metric values corresponding to different similarity score and binding affinity restraints. Few blank points are present in each heatmap due to no available generation meeting these specific restraints.

Figure 6a plots the RMSE metrics. It reveals the most reliable region, i.e., having low RMSE, is y_{sim} above 0.925, and the $y_{\Delta G}$ is between -10.1 and -12.1 kcal/mol. This is expected since machine learning models can render accurate predictions if generated structures are highly similar to the training data (see Figure 5). Besides, a large population of training samples have ΔG s between -10.1 and -12.1 kcal/mol, leading to more reliable predictions to molecules generated inside this range. Outside this range, as both $y_{\Delta G}$ and y_{sim} decrease, the RMSEs between the LV-BP and 2DFP-BP predictions increase. Specifically, if $y_{\Delta G} < -12.1$ kcal/mol and $y_{sim} < 0.675$, the RMSEs are always over 3.2 kcal/mol.

Figure 6c depicts the $R_{\rm p}$'s between the LV-BP and 2DFP-BP predictions with respect to $y_{\Delta G}$ and $y_{\rm sim}$. Similar to the manner of the RMSE distribution, -12.1 kcal/mol $\leq y_{\Delta G} \leq 10.1$ kcal/mol and $0.9 \leq y_{\rm sim} \leq 0.95$ lead to the $R_{\rm p}$ values consistently higher than 0.8.

The last component of our reliable analysis regards the loss function magnitudes of our GNC generator plotted in Figure 6d. In most cases, the loss function values are less than 0.05. However, in some special situations, our network cannot maintain the loss values lower than 0.15. At these points, we cannot find any generated molecules subject to the multiproperty restraints simultaneously, which renders the blank spots in the criteria plots in Figures 6a, b, and c.

In summary, to generate molecules with reliable predictions, one should set the binding affinity objective $y_{\Delta G}$ in a region filled with a large population of training data. Besides, the similarity score restraint y_{sim} should be high. However, in some circumstances, the generated molecules that have high predicted affinities should also be included in further consideration.

III.B. Designing Alternative Drug Candidates. In this section, we utilize our GNC to produce alternative drug-like molecules with high binding affinities to the existing drugs' targets, which provides effective libraries for further improvement or searching for cheaper drug alternatives. This work discusses eight drugs and their targets with the information regarding the names, ChEMBL IDs, energies, etc. summarized in Table 1. All of these drugs were approved by the FDA in the recent decade to treat critical diseases, especially a variety of cancers. Notably, the drug Ribociclib has two different targets (Cyclin-dependent kinase 4 and Cyclin-dependent kinase 6), so Ribociclib has two sets of ΔGs and two sets of training compounds.



Figure 4. A series of generated molecules over the evolution from the seed to the reference molecule.

III.B.1. Single-Target Drug: Ceritinib. Statistics of the Drug Ceritinib. The brand name of Ceritinib (ChEMBL ID: CHEMBL2403108) is Zykadia. It was developed by Novartis and approved by the FDA in April 2014 to treat different types of nonsmall cell lung cancers. Ceritinib is extremely expensive, with the monthly cost of Ceritinib-based treatment in the US being approximately \$11,428.

Ceritinib inhibits the ALK tyrosine kinase receptor (ALK, CHEMBL ID: CHEMBL4247). The ChEMBL database provides 1407 molecules with experimental binding affinity labels to this target available. After going through the filter in our model, 1203 molecules are left to train our generator and predictor. Figure 7a depicts the ΔG distribution of this training

set; it unveils that hundreds of training samples have their binding affinities close to or even higher than Ceritinib. The similarity score distribution in Figure 7b indicates the training set includes 355 samples with their similarity scores to Ceritinib being over 0.3, with 56 samples with such scores over 0.6. These promising analytical assessments enable our GNC to design potential inhibitors to the target ALK.

Designing New Drug-like Molecules. Here we use Ceritinib as the reference molecule to design alternative Ceritinib drugs. Section III.A.3 suggests, to generate new molecules with desirable properties, the binding affinity objectives should be inside a region with plenty of training data, and the similarity scores to the reference molecule y_{sim} should be restrained to be



Figure 5. Experimental ΔG distribution of the filtered BACE1 data set.

high. However, high similarity scores could lead to quite limited chemical space. Our solution to this drawback is, first, extend similarity score restraint to a broader range, then reevaluate generated compounds using the 2DFP-BP, and only pick the ones with low discrepancies between the LV-BP and 2DFP-BP predicted affinities.



Figure 7. (a) ΔG distribution of the filtered training set to the target ALK. The red bar indicates the interval containing the ΔG of Ceritinib. (b) Similarity score distribution of the other molecules in this set to Ceritinib.

Following this strategy, we set the similarity score restraints varying from 0.3 to 0.9 with an increment of 0.025; this is because more than 300 training samples have their similarity scores over 0.3, which is supportive to generate compounds in this similarity range. The binding affinities are aimed at the interval from -10.5 kcal/mol to -12.25 kcal/mol with an



Figure 6. Reliability test to our GNC generator on the BACE1 data set. The prediction reliability with different binding affinity objectives $(y_{\Delta G})$ and similarity score objectives (y_{sim}) is tested. The discrepancies between the LV-BP and the 2DFP-BP predictions are evaluated by different criteria: (a) RMSE; (b) activate ratio; (c) Pearson correlation; (d) loss function value.

increment of -0.25 kcal/mol; this binding affinity region covers hundreds of training samples as well as the drug Ceritinib itself. After reevaluating by the 2DFP-BP, any generated molecules with the relative errors between the LV-BP and 2DFP-BP predicted affinities over 5% are thrown away.

2DFP-BP Reevaluation. The details of the 2DFP-BP are offered in section II.D. With Ceritinib as the reference, our GNC model creates 1095 novel active drug-like molecules. Upon eliminating the ones with high discrepancies in their LV-BP and 2DFP-BP predictions, 629 molecules are left. Figure 8 indicates,



Figure 8. Correlation plot between the LV-BP and 2DFP-BP predicted ΔGs of the generated molecules to the target ALK, the ones having high relative errors between the two predictions (>5%) are already eliminated.

for these 629 molecules, the correlation between the two predictions is quite promising, with the RMSE = 0.28 kcal/mol, $R_p = 0.80$, and active ratio = 0.95, respectively. This statistical information endorses the drug-likeness potential of our AI-generated molecules.

The LV-BP and 2DFP-BP predicted binding affinities of these 629 molecules are also averaged, and their distribution is shown in Figure 9a. This figure reveals the preferred affinities of the generated compounds is from -9.8 kcal/mol to -10.8 kcal/mol, which is also the most popular affinity region of the training



Figure 9. Average predicted binding affinities to the target ALK of the generated molecules and their similarity scores to the reference drug Ceritinib. (a) Distribution of the averages of the LV-BP and 2DFP-BP predicted ΔG s to the target ALK. (b) Similarity score distribution to the reference drug Ceritinib.

samples. Figure 9b illustrates their similarity score distribution to the reference drug Ceritinib.

Top 6 Drug Candidates. Ranked by the average predicted binding affinities of these 629 molecules, we select the top six drug candidates. Their 2D draws are plotted in Figure 10. The relative errors between their LV-BP and 2DFP-BP predictions are 1.3%, 3.1%, 0.1%, 4.0%, 3.5%, and 3.9%, respectively. It is delighted to see that their average predicted affinities are much higher than that of the reference drug Ceritinib. Moreover, these candidates have similarity scores to the reference drug from 0.54 to 0.69. They are brand new and do not exist in the SciFinder database. We also predict their toxicity in terms of LD50 using our own GBDT model,³⁹ log P values and synthetic accessibility scores (SAS) using RDKit, and log S values by Alog PS 2.1,⁵⁰ and report them in Figure 10; these values are comparable to those of the reference drug and in reasonable ranges.

Interaction and Pharmacophore Analysis. One can observe from Figure 10, even though the similarity scores are only between 0.56 and 0.69, these generated compounds share some moieties with the reference drug Ceritinib. This designates these common moieties to possibly involve critical interactions with the binding site of the target.

To verify our generated compounds still contain critical moieties to the binding, from experimental structures, we investigate the drug-target interactions, as well as the common pharmacophores among all the active compounds to the target. Figure 11a shows the interaction details between the drug Ceritinib and its target in the 3D crystal structure of their complex (PDB ID 4MKC⁵¹). It reveals their interactions include one hydrogen bond with one N atom in the pyrimidine of the drug (marked by 1 in the Figure 11a) as the acceptor, one hydrogen bond with one N atom in the chain of the drug (marked by 2 in the Figure 11a) as the donor, another hydrogen bond with one O atom in the drug as the acceptor, one hydrophobic interaction between one benzene ring (marked by 3 in the Figure 11a), and the target, as well as other hydrophobic interactions.

The common pharmacophore analysis in Figure 11b is consistent with the interaction analysis. The N atom in the pyrimidine of the drug (marked by 1 in Figure 11a and 11b) and the N atom in the chain of the drug (marked by 2 in Figure 11a and 11b) are critical pharmacophores; they play the roles of an acceptor and a donor of hydrogen bonds, respectively. Another pharmacophore is the benzene ring forming a hydrophobic interaction with the target. The pharmacophore analysis also reveals more potential interaction modes, such as the hydrophobic interaction between the Cl atom and target, and the hydrogen bond with the other N atom in the chain of the drug as its donor.

As illustrated in Figure 11b, all these critical pharmacophores are retained in our top six generated compounds, which strongly supports that these compounds are potential inhibitors to the target.

III.B.2. Double-Target Drug: Ribociclib. Statistics of the Drug Ribociclib. Here, we test the generative power of our GNC on multitarget drugs. In this case study, the multiproperty restraints consist of multiple binding affinity criteria and the similarity score to a reference molecule. The drug we test here is Ribociclib (ChEMBL ID: CHEMBL3545110, brand name: Kisqali). It was developed by Novartis and Astex Pharmaceuticals and approved by the FDA in 2017 to treat certain kinds of breast cancers. The monthly cost of Ribociclib treatment is \$10,950 in the US.



Tox=2.68 log P=5.95 log S=-6.12 SAS=3.47

(f) ΔG=-11.36 kcal/mol Sim=0.63 Tox=2.91 log P=6.67 log S=-5.59 SAS=3.32

Figure 10. Drug Ceritinib and its top six generated molecules. The predicted ΔG_S to the target ALK, similarity scores (Sim) to the drug, calculated toxicity (Tox), log P, log S values, and synthetic accessibility scores (SASs) are also present.

Ribociclib inhibits two different targets, namely the Cyclindependent kinase 4 (CDK4, CHEMBL ID: CHEMBL331) and Cyclin-dependent kinase 6 (CDK6, CHEMBL ID: CHEMBL2508). In the ChEMBL database, 919 molecules have CDK4 binding data and 289 molecules have CDK6 binding data. After filtering, 918 and 289 molecules are retained, respectively, providing a small training set to the CDK6. Figure 12a and b plot the ΔG distributions of these two training sets. It reveals, in both the CDK4 and CDK6 sets, hundreds of samples have binding affinities close to or even higher than Ribociclib. Figure 12c and d show the similarity score distributions to Ribociclib of the two training sets. Both these sets include more than 200 samples with similarity scores to the drug over 0.3, with more than 60 samples with such scores over 0.6.



Figure 11. Pharmacophore analysis to the ALK tyrosine kinase's binding site. (a) Interaction plot between the drug Ceritinib and ALK tyrosine kinase from the 3D experimental structure with the PDB ID 4MKC. (b) 3D alignments of the common pharmacophores obtained from all the active compounds to the target ALK with the 3D experimental structure of the drug Ceritinib; the top six generated molecules are also aligned to it.



Figure 12. (a) Experimental ΔG distribution of the training set to the target CDK4 with the interval containing the ΔG of the drug Ribociclib to CDK4 marked in red. (b) Experimental ΔG distribution of the CDK6 training set with the interval containing the ΔG of Ribociclib to CDK4 in red. (c) Similarity score distribution of the other samples in the CDK4 training set to Ribociclib. (d) Similarity score distribution of the other samples in the CDK6 training set to Ribociclib.

Multitask Predictor. Since the CDK6 training set is small, it is challenging to train an accurate predictor for this target. However, the two targets, CDK4 and CDK6, are similar due to the calculation via SWISS-MODEL⁵² with the sequence identity being 71.1%. Therefore, multitask deep learning can enhance reliability.

In our multitask architecture, the binding affinity predictor in our generator offers two outputs, each for one of the two targets, so the predictor is trained by the two training sets simultaneously. As a result, in a 5-fold crossing-validation test, the multitask model significantly improves the performance on the small data set.

As illustrated in Table 2, the target CDK4 with a 918molecule training set does not benefit so much from the multitask. However, in the case of the target CDK6 only with a 289-molecule training set, the improvement is dramatic: the R_p 's rise from 0.524 to 0.811 by the LV-BP and from 0.485 to 0.779 by the 2FP-BP. These results demonstrate the efficiency of the multitask architecture.

Generation of New Drug-like Molecules. To design alternative Ribociclib drugs in broader chemical space, we set

Table 2. R_p s of ΔG Predictions from the 5-Fold Cross-Validation Tests on the Two Targets of the Drug Ribociclib by the Single Task and Multitask Predictors

	Target CDK4		Target CDK6	
	Single task	Multitask	Single task	Multitask
LV-BP	0.791	0.804	0.524	0.811
2FP-BP	0.824	0.836	0.485	0.779

the similarity score restraints to Ribociclib from 0.30 to 0.90, with an increment of 0.025. The binding affinities to the target CDK4 are aimed at the interval between -10.80 kcal/mol and -12.00 kcal/mol with an increment of -0.2 kcal/mol; this interval covers the ΔGs of Ribociclib as well as many other training samples. For the same reason, the objectives of the binding affinities to the target CDK6 are set from -10.2 kcal/mol. Totally, following this scheme, we create 1080 novel molecules.

Reevaluation by 2DFP-BP. The predicted bind affinities of these 1080 generated compounds are reevaluated by the 2DFP-BP model incorporated with the multitask DNN. The parameters of this architecture are introduced in section II.E. After excluding the ones with high discrepancies between the LV-BP and 2DFP-BP predictions, 271 molecules remain.

Figure 13 depicts the correlation plots between their LV-BP and 2DFP-BP predicted ΔGs to the two targets. The correlations of the ΔGs to both the two targets are promising. Specifically, the ΔGs to the target CDK4 can achieve an RMSE of 0.29 kcal/mol, an $R_{\rm P}$ of 0.69, and an active ratio as high as 0.98; the ΔGs to the target CDK6 also have an RMSE of 0.27 kcal/mol, a $R_{\rm P}$ of 0.65, and an active ratio of 0.78.

Their LV-BP and 2DFP-BP predicted binding affinities to the two targets are also averaged, and the distributions are shown in Figure 14a,b. Figure 14c illustrates their similarity score distribution to the reference drug Ribociclib.

Top Six Drug Candidates. According to the average predicted binding affinities of these 271 molecules, we select the top six drug candidates. Figure 15 represents their 2D plots. To the target CDK4, the relative errors between their LV-BP and 2DFP-BP predictions are 4.6%, 3.0%, 3.2%, 0.2%, 1.6%, and 0.7%, respectively; to CDK6, the relative errors between the two



Figure 14. Average predicted ΔG to the targets CDK4 and CDK6 of the generated molecules and their similarity scores to the reference drug Ribociclib: (a, b) Distributions of the averages of the LV-BP and 2DFP-BP predicted ΔG s to the two targets, respectively. (c) Similarity score distribution to the reference drug Ribociclib.

predictions are 1.7%, 1.2%, 3.7%, 3.4%, 4.8%, and 1.7%, respectively. Most of them have better binding affinities than the reference drug Ribociclib. For example, to CDK4, the affinities of the first and fourth compounds are predicted to be higher than that of Ribociclib, and the other three have similar ones with Ribociclib; to CDK6, all the top six candidates have better binding affinities. Moreover, their similarity scores to the reference drug are between 0.65 and 0.75. They are not in the SciFinder database, which means these six candidates are novel. The toxicity (LD50), log P, log S values, and synthetic accessibility scores (SASs) calculated by our model, RdKit,



Figure 13. Correlations between the LV-BP and 2DFP-BP predicted ΔGs of the generated molecules to the targets CDK4 and CDK6; the ones having high relative errors between the two predictions (>5%) are already eliminated.

Article



Figure 15. Drug Ribociclib and its top six generated molecules. The predicted ΔG s to the targets CDK4 and CDK6, similarity scores (Sim) to the drug, calculated toxicity (Tox), log P, log S values, and synthetic accessibility scores (SASs) are also reported.



Figure 16. (a) Illustration of the interactions between the drug Ribociclib and the target CDK6 extracted from the experimental structure (PDB ID 5L2T). (b) 3D alignment of the common pharmacophores obtained from all the active compounds to the target CDK6 with the 3D experimental structure of the drug Ribociclib; the top six generated molecules are also aligned to it.



Figure 17. Correlation plots between the LV-BP and 2DFP-BP predicted ΔGs of the generated molecules to the six drug targets; the ones with high relative errors between the two predictions are already eliminated.

and Alog PS are also shown in the figure; their toxicity, log P, log S values, and SASs are comparable to those of the reference drug.

Interaction and Pharmacophore Analysis. Figure 16a shows the interaction details between the drug Ribociclib and the target CDK6 in the 3D crystal structure of their complex (PDB ID 5L2T⁵³). It indicates the main interactions are hydrogen bonds and hydrophobic interactions. The pharmacophore analysis in Figure 16b reveals that the critical pharmacophores are the pyrrolopyrimidine, pyridine, and cyclohexane, which can form hydrogen bonds and hydrophobic interactions with the binding site.

Among the top six candidates, except the first one, the others contain all these critical pharmacophores; the first one has most of them. This suggests all the six compounds are potential inhibitors to the targets.

III.B.3. Tests on the Other Single-Target Drugs. We also apply our GNC to the rest six drugs listed in Table 1 and design novel drug candidates. The similarity score restraints are from 0.30 to 0.90, with an increment of 0.025. The ΔG objective ranges are chosen to contain the ΔGs of the drugs as well as lots of other training samples. We also only collect the generated molecules with the relative errors between the LV-BP and 2DFP-BP predicted ΔGs below 5%.

Figure 17 indicates, since the ones with high discrepancies between the two predictions are eliminated, our generated compounds to all these six drug targets have promising correlations.

Now, we focus on whether highly active drug candidates are created or not. This relies on the binding affinity and similarity score distributions of the training sets. Figure 18 provides the sizes and the ΔG distributions of these six target-specified training sets. The red bars indicate the intervals containing the ΔG s of the reference drugs. Figure 19 plots their similarity score distributions.

Figure 18a shows that the training set of the drug Acalabrutinib contains 1459 compounds. Among them, more than 400 compounds have higher binding affinities than that of the drug, which is beneficial to generate compounds more potent than the drug. Also, Figure 19a reveals that, in the training set, over 500 molecules have the similarity scores to the drug over 0.3, which is quite helpful to generate novel compounds with similarity scores also in this range. These promising statistical information helps to explain why our GNC can create as many as 879 new compounds for this drug target with high confidence, and 17 of them have higher or similar binding affinities with that of Acalabrutinib; the best ΔG is -11.90 kcal/mol (see Figures S1 and S3).

As revealed by 18b, the training set of the drug Idelalisib is larger than that of Acalabrutinib, and over 500 molecules in this set have higher binding affinities than that of Idelalisib. As a result, 73 of the 794 generated compounds are more potent than Idelalisib, with the best ΔG being -11.27 kcal/mol (see Figures S1 and S4).

Figure 18c exhibits that the data set of Macimorelin is quite small. However, the correlation of the generated molecules is still satisfactory, and 13 generated molecules have similar ΔGs with that of the drug (see Figure S1). This is because its training set contains 205 compounds with higher binding affinities than that of the drug, and also 182 compounds with the similarity scores to the drug over 0.3.

Figure 18d and e depict the data sets of the drugs Dabrafenib and Enzalutamide. The affinity values of these two drugs are close to the upper boundaries of their training sets' affinity



Figure 18. Sizes and experimental ΔG distributions of the training sets to the six drug targets. The red bars indicate the intervals containing the ΔG s of the six reference drugs.

domains. In other words, few molecules in their training sets are more active than the drugs. Since the interpolation nature of machine learning models tends to provide reliable predictions inside the populated range, it is tough to generate compounds more potent than these drugs. Although their training set sizes and similarity scores are favorable, the numbers of generated compounds with better binding affinities than that of the drugs are only 2 and 4, respectively (see Figures S1, S6, and S7).

We perform our last experiment on the drug Panobinostat. Panobinostat is an extreme example: as illustrated in Figure 18f, its binding affinity is the highest in its training set. Therefore, although our GNC model generates 319 molecules, the top active one among them only reaches a ΔG of -11.56 kcal/mol, which is far less potent than the drug itself ($\Delta G = -12.46$ kcal/mol).

In the Supporting Information, Figures S3 to S8 provide the top six generated alternative compounds for the drugs Acalabrutinib, Idelalisib, Macimorelin, Dabrafenib, Enzalutamide, and Panobinostat, as well as their predicted ΔGs .

IV. DISCUSSIONS

With the availability of deep learning technologies, an increasing number of *in silico* molecule generation models have been proposed. These models can be classified into three categories: randomized output, controlled output, and optimized output.⁵⁴



Figure 19. Similarity score distributions of the other molecules to the six reference drugs in their training sets.

One of the challenges is how to generate new molecules with desired chemical properties, especially drug-like molecules. Another challenge is how to improve the utility of *in silico* molecule generation without direct experimental validations. To address these challenges, we propose a new GNC to generate drug-like molecules based on multiproperty optimizations via gradient descent.

IV.A. Essential Circumstances for Reliable and Desired Molecule Generation. Based on our experiments, there are two essential circumstances to generate molecules with reliable and desired predicted chemical properties:

1. An objective property value should always be in a region with many training samples. Based on the nature of machine learning methods, the predictor can be built with high accuracy at an objective value with many training samples around it. It can be seen from Section III.A.3 that when a binding affinity objective is in the middle of the training set's distribution, the latent space generator can always generate novel molecules with reliable predictions confirmed by the 2DFP-BP. However, when an objective is close to or even at the edge of the training-set distribution, one may still generate many novel compounds, but the predictions to them are quite risky, such as high discrepancies between the LV-BP and 2DFP-BP predictions. pubs.acs.org/jcim

2. Generated compounds should have some high similarity scores to some existing molecules in the training set. If some molecules (not necessarily reference molecules) in the training set are similar to generated compounds, then the predictions are reliable and can be verified by the 2DFP-BP.

IV.B. Necessity of Both Similarity and Property Value Restraints. The two points above also explain why both similarity score and property value restraints should be included in our generator.

The goal of property restraint is twofold. First, it restricts property values to desired ones. Additionally, it can be used to achieve high reliability. As discussed above, if the property value is limited to a populated region of the training-set distribution, the resulting generated molecules will most likely have reliable predictions.

Similarity restraint is also to ensure prediction accuracy. High similarity scores to existing molecules make predictions more accurate and reliable. Moreover, in a regular drug discovery procedure for a given target, one typically starts from some lead compounds or even current drugs and then carries out lead optimizations to make candidates more "drug-like", e.g., higher activity and lower side effects.² Therefore, similarly, in a drugspecified generative model, the similarity to a reference compound or drug must be controlled to guarantee that new drug-like compounds still bind to the target. For example, with similarity restraint, generated molecules share pharmacophores with the reference drug.

IV.C. Multiple Property Restraints. Drug design is sophisticated. To develop a drug, plenty of properties must be carefully studied, such as binding affinity, toxicity, partition coefficient (log P), aqueous solubility (log S), and off-target effect. The failure in any one of these properties can prevent drug candidates from the market. In other words, drug design is a multiproperty optimization process.

Technically, our generator can handle this multiproperty optimization. In our framework, the restraint to each property is realized by one term in the loss function. Therefore, multiproperty optimization can be satisfied simultaneously in our GNC. In this work, multiple property restraints are tested on one drug with two targets (Ribociclib). In fact, the generated new candidates have ideal binding affinities to the two targets simultaneously; this means our model can work on the multiple property restraints. Multiple properties can also be specified as toxicity, log P, log S, etc. To avoid side effects, one can also control drug candidates to have a high binding affinity to one target but low binding affinities to other targets.

V. CONCLUSION

The search for alternative drugs is important for improving the quality of existing drugs and making new drugs cheaply available to low-income people. In this work, we develop a new generative network complex (GNC) for automated generation of drug-like molecules based on the multiproperty optimization via gradient descent in the latent space. In this new GNC, multiple chemical properties, particularly binding affinity and similarity score, are optimized to generate new molecules with desired chemical and drug properties. To ensure the prediction reliability of these new compounds, we reevaluate them by independent 2D-fingerprint based predictors. Molecules without consistent predictions from the latent-vector model and the 2D-fingerprint model are not accepted. After the consistent check, hundreds of potential drug

candidates are reported. Performed on a supercomputer, the generation process from one seed to a large number of new molecules takes less than 10 min. Therefore, our GNC is an efficient new paradigm for discovering new drug candidates. To demonstrate the utility of the present GNC, we first test its reliability on the BACE1 target and then further apply this model to generate thousands of new alternative drug candidates for a few market existing drugs, namely, Ceritinib, Ribociclib, Acalabrutinib, Idelalisib, Dabrafenib, Macimorelin, Enzalutamide, and Panobinostat.

We also discuss the keys to generating drug-like candidates with reliable predictions. First, an objective property value should be in a populated region of the training-set distribution. Second, generated molecules need to have good similarity scores to some existing compounds in the training set.

ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at https://pubs.acs.org/doi/10.1021/acs.jcim.0c00599.

Figures S1 to S8 are the ΔG and similarity score distributions of the generated drug-like molecules for the drugs Acalabrutinib, Idelalisib, Dabrafenib, Macimorelin, Enzalutamide, and Panobinostat, as well as the 2D plots and predicted ΔG s of the top six among them. (PDF)

AUTHOR INFORMATION

Corresponding Author

Guo-Wei Wei – Department of Mathematics, Department of Biochemistry and Molecular Biology, and Department of Electrical and Computer Engineering, Michigan State University, East Lansing, Michigan 48824, United States; ◎ orcid.org/ 0000-0002-5781-2937; Email: wei@math.msu.edu

Authors

- Kaifu Gao Department of Mathematics, Michigan State University, East Lansing, Michigan 48824, United States
- **Duc Duy Nguyen** Department of Mathematics, Michigan State University, East Lansing, Michigan 48824, United States
- **Meihua Tu** Pfizer Medicine Design, Cambridge, Massachusetts 02139, United States

Complete contact information is available at: https://pubs.acs.org/10.1021/acs.jcim.0c00599

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

This work was supported in part by NSF Grants DMS-1721024, DMS-1761320, and IIS1900473, NIH Grant GM126189, and Michigan Economic Development Corporation. D.D.N. and G.W.W. are also funded by Bristol-Myers Squibb and Pfizer.

REFERENCES

(1) DiMasi, J. A.; Grabowski, H. G.; Hansen, R. W. Innovation in the pharmaceutical industry: new estimates of r&d costs. *Journal of health economics* **2016**, *47*, 20–33.

(2) Hughes, J. P.; Rees, S.; Kalindjian, S. B.; Philpott, K. L. Principles of early drug discovery. *British journal of pharmacology* **2011**, *162* (6), 1239–1249.

(3) Macarron, R.; Banks, M. N.; Bojanic, D.; Burns, D. J.; Cirovic, D. A.; Garyantes, T.; Green, D. V. S.; Hertzberg, R. P.; Janzen, W. P.;

Paslay, J. W.; et al. Impact of high-throughput screening in biomedical research. *Nat. Rev. Drug Discovery* **2011**, *10* (3), 188.

(4) Graul, A. I.; Revel, L.; Rosa, E.; Cruces, E. Overcoming the obstacles in the pharma/biotech industry: 2008 update. *Drug News Perspect.* 2009, 22 (1), 39.

(5) Khan, M. T. H. Predictions of the admet properties of candidate drug molecules utilizing different qsar/qspr modelling approaches. *Curr. Drug Metab.* **2010**, *11* (4), 285–295.

(6) Waring, M. J.; Arrowsmith, J.; Leach, A. R.; Leeson, P. D.; Mandrell, S.; Owen, R. M.; Pairaudeau, G.; Pennie, W. D.; Pickett, S. D.; Wang, J.; et al. An analysis of the attrition of drug candidates from four major pharmaceutical companies. *Nat. Rev. Drug Discovery* **2015**, *14* (7), 475.

(7) Strømgaard, K.; Krogsgaard-Larsen, P.; Madsen, U.Textbook of drug design and discovery; CRC Press: 2017.

(8) Alqahtani, S. In silico adme-tox modeling: progress and prospects. *Expert Opin. Drug Metab. Toxicol.* **2017**, *13* (11), 1147–1158.

(9) Balakin, K. V.; Ivanenkov, Y. A.; Savchuk, N. P. Compound library design for target families. In *Chemogenomics*; Springer: 2009; pp 21–46.
(10) Kapetanovic, I. M. Computer-aided drug discovery and

development (caddd): in silico-chemico-biological approach. *Chem.-Biol. Interact.* **2008**, 171 (2), 165–176.

(11) Huang, R.; Leung, I. Protein-directed dynamic combinatorial chemistry: a guide to protein ligand and inhibitor discovery. *Molecules* **2016**, *21* (7), 910.

(12) Szymański, P.; Markowicz, M.; Mikiciuk-Olasik, E. Adaptation of high-throughput screening in drug discovery-toxicological screening tests. *Int. J. Mol. Sci.* **2012**, *13* (1), 427–452.

(13) Schneider, G.; Fechner, U. Computer-based de novo design of drug-like molecules. *Nat. Rev. Drug Discovery* **2005**, *4* (8), 649.

(14) Min, S.; Lee, B.; Yoon, S. Deep learning in bioinformatics. *Briefings Bioinf.* **2016**, *18* (5), 851–869.

(15) Mamoshina, P.; Vieira, A.; Putin, E.; Zhavoronkov, A. Applications of deep learning in biomedicine. *Mol. Pharmaceutics* **2016**, *13* (5), 1445–1454.

(16) Sutskever, I.; Vinyals, O.; Le Quoc, V. Sequence to sequence learning with neural networks. *Advances in neural information processing systems* **2014**, 3104–3112.

(17) Kingma, D. P.; Welling, M. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114, **2013**.

(18) Makhzani, A.; Shlens, J.; Jaitly, N.; Goodfellow, I.; Frey, B. Adversarial autoencoders. arXiv preprint arXiv:1511.05644, **2015**.

(19) Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Sherjil, O.; Courville, A.; Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems* **2014**, 2672–2680.

(20) Kaelbling, Leslie Pack; Littman, Michael L; Moore, Andrew W Reinforcement learning: A survey. *Journal of artificial intelligence research* **1996**, *4*, 237–285.

(21) Winter, Robin; Montanari, Floriane; Steffen, Andreas; Briem, Hans; Noé, Frank; Clevert, Djork-Arné Efficient multi-objective molecular optimization in a continuous latent space. *Chemical science* **2019**, *10* (34), 8016–8024.

(22) Winter, Robin; Montanari, Floriane; Noé, Frank; Clevert, Djork-Arné Learning continuous and data-driven molecular descriptors by translating equivalent chemical representations. *Chemical science* **2019**, *10* (6), 1692–1701.

(23) Gómez-Bombarelli, Rafael; Wei, Jennifer N; Duvenaud, David; Hernández-Lobato, José Miguel; Sánchez-Lengeling, Benjamín; Sheberla, Dennis; Aguilera-Iparraguirre, Jorge; Hirzel, Timothy D; Adams, Ryan P; Aspuru-Guzik, Alán Automatic chemical design using a data-driven continuous representation of molecules. *ACS Cent. Sci.* **2018**, 4 (2), 268–276.

(24) Skalic, Miha; Jiménez, José; Sabbadin, Davide; De Fabritiis, Gianni Shape-based generative modeling for de novo drug design. J. Chem. Inf. Model. 2019, 59 (3), 1205–1214.

(25) Kadurin, Artur; Nikolenko, Sergey; Khrabrov, Kuzma; Aliper, Alex; Zhavoronkov, Alex druGAN: an advanced generative adversarial autoencoder model for de novo generation of new molecules with

pubs.acs.org/jcim

desired molecular properties in silico. *Mol. Pharmaceutics* **2017**, *14* (9), 3098–3104.

(26) Sattarov, Boris; Baskin, Igor I; Horvath, Dragos; Marcou, Gilles; Bjerrum, Esben Jannik; Varnek, Alexandre De novo molecular design by combining deep autoencoder recurrent neural networks with generative topographic mapping. *J. Chem. Inf. Model.* **2019**, *59* (3), 1182–1196.

(27) Olivecrona, Marcus; Blaschke, Thomas; Engkvist, Ola; Chen, Hongming Molecular de-novo design through deep reinforcement learning. J. Cheminf. **2017**, 9 (1), 48.

(28) Popova, Mariya; Isayev, Olexandr; Tropsha, Alexander Deep reinforcement learning for de novo drug design. *Science advances* **2018**, *4* (7), No. eaap7885.

(29) Kang, Seokho; Cho, Kyunghyun Conditional molecular design with deep generative models. *J. Chem. Inf. Model.* **2019**, *59* (1), 43–52.

(30) Zhou, Zhenpeng; Kearnes, Steven; Li, Li; Zare, Richard N; Riley, Patrick Optimization of molecules via deep reinforcement learning. *Sci. Rep.* **2019**, *9* (1), 10752.

(31) Kerns, Edward H; Di, Li; Carter, Guy T In vitro solubility assays in drug discovery. *Curr. Drug Metab.* **2008**, *9* (9), 879–885.

(32) Leo, Albert; Hansch, Corwin; Elkins, David Partition coefficients and their uses. *Chem. Rev.* **1971**, *71* (6), 525–616.

(33) Cho, K.; van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; Bengio, Y. Learning phrase representations using rnn encoder-decoder for statistical machine translation. arXiv preprint arXiv:1406.1078, **2014**.

(34) Hochreiter, Sepp; Schmidhuber, Jürgen Long short-term memory. *Neural computation* **1997**, *9* (8), 1735–1780.

(35) Weininger, D. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *J. Chem. Inf. Model.* **1988**, *28* (1), 31–36.

(36) Lipinski, Christopher A; Lombardo, Franco; Dominy, Beryl W; Feeney, Paul J Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Adv. Drug Delivery Rev.* **1997**, 23 (1–3), 3–25.

(37) Gaulton, Anna; Bellis, Louisa J; Bento, A Patricia; Chambers, Jon; Davies, Mark; Hersey, Anne; Light, Yvonne; McGlinchey, Shaun; Michalovich, David; Al-Lazikani, Bissan; et al. Chembl: a large-scale bioactivity database for drug discovery. *Nucleic Acids Res.* **2012**, *40* (D1), D1100–D1107.

(38) Paszke, A.; Gross, S.; Chintala, S.; Chanan, G. Pytorch: Tensors and dynamic neural networks in python with strong gpu acceleration. *PyTorch: Tensors and dynamic neural networks in Python with strong GPU acceleration* **2017**, 6.

(39) Gao, Kaifu; Nguyen, Duc Duy; Sresht, Vishnu; Mathiowetz, Alan M.; Tu, Meihua; Wei, Guo-Wei Are 2d fingerprints still valuable for drug discovery? *Phys. Chem. Chem. Phys.* **2020**, *22*, 8373–8390.

(40) Rogers, David; Hahn, Mathew Extended-connectivity fingerprints. J. Chem. Inf. Model. 2010, 50 (5), 742–754.

(41) Hall, Lowell H; Kier, Lemont B Electrotopological state indices for atom types: a novel combination of electronic, topological, and valence state information. *J. Chem. Inf. Model.* **1995**, 35 (6), 1039–1045.

(42) Landrum, G. et al. Rdkit: Open-source cheminformatics, 2006.

(43) Schapire, R. E. The boosting approach to machine learning: An overview. In *Nonlinear estimation and classification;* Springer: 2003; pp 149–171.

(44) Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine learning in python. *Journal of machine learning research* **2011**, *12* (Oct), 2825–2830.

(45) Caruana, R. Multitask learning. *Machine learning* **1997**, 28 (1), 41–75.

(46) Salentin, S.; Schreiber, S.; Haupt, V. J.; Adasme, M. F.; Schroeder, M. Plip: fully automated protein–ligand interaction profiler. *Nucleic Acids Res.* **2015**, *43* (W1), W443–W447.

(47) Dixon, Steven L; Smondyrev, Alexander M; Knoll, Eric H; Rao, Shashidhar N; Shaw, David E; Friesner, Richard A Phase: a new engine for pharmacophore perception, 3d qsar model development, and 3d database screening: 1. methodology and preliminary results. J. Comput.-Aided Mol. Des. 2006, 20 (10–11), 647–671.

(48) Vassar, Robert; Kovacs, Dora M; Yan, Riqiang; Wong, Philip C The β -secretase enzyme bace in health and alzheimer's disease: regulation, cell biology, function, and therapeutic potential. *J. Neurosci.* **2009**, *29* (41), 12787–12794.

(49) Prati, Federica; Bottegoni, Giovanni; Bolognesi, Maria Laura; Cavalli, Andrea Bace-1 inhibitors: From recent single-target molecules to multitarget compounds for alzheimer's disease: Miniperspective. *J. Med. Chem.* **2018**, *61* (3), *61*9–637.

(50) Tetko, Igor V; Gasteiger, Johann; Todeschini, Roberto; Mauri, Andrea; Livingstone, David; Ertl, Peter; Palyulin, Vladimir A; Radchenko, Eugene V; Zefirov, Nikolay S; Makarenko, Alexander S; et al. Virtual computational chemistry laboratory–design and description. J. Comput.-Aided Mol. Des. **2005**, 19 (6), 453–463.

(51) Friboulet, Luc; Li, Nanxin; Katayama, Ryohei; Lee, Christian C; Gainor, Justin F; Crystal, Adam S; Michellys, Pierre-Yves; Awad, Mark M; Yanagitani, Noriko; Kim, Sungjoon; et al. The alk inhibitor ceritinib overcomes crizotinib resistance in non–small cell lung cancer. *Cancer Discovery* **2014**, *4* (6), 662–673.

(52) Schwede, T.; Kopp, J.; Guex, N.; Peitsch, M. C. Swiss-model: an automated protein homology-modeling server. *Nucleic acids research* **2003**, *31* (13), 3381–3385.

(53) Chen, Ping; Lee, Nathan V; Hu, Wenyue; Xu, Meirong; Ferre, Rose Ann; Lam, Hieu; Bergqvist, Simon; Solowiej, James; Diehl, Wade; He, You-Ai; et al. Spectrum and degree of cdk drug interactions predicts clinical performance. *Mol. Cancer Ther.* **2016**, *15* (10), 2273–2281.

(54) Grow, C.; Gao, K.; Nguyen, D. D.; Wei, G.-W. Generative network complex (gnc) for drug discovery. *Communications in Information and Systems* **2019**, *19*, 241–277.