

Accelerating Prototype-Based Drug Discovery using Conditional Diversity Networks

Shahar Harel

Technion - Israel Institute of Technology
Haifa, Israel
sshahar@cs.technion.ac.il

Kira Radinsky

Technion - Israel Institute of Technology
Haifa, Israel
kirar@cs.technion.ac.il

ABSTRACT

Designing a new drug is a lengthy and expensive process. As the space of potential molecules is very large ($10^{23} - 10^{60}$), a common technique during drug discovery is to start from a molecule which already has some of the desired properties. An interdisciplinary team of scientists generates hypothesis about the required changes to the prototype. In this work, we develop an algorithmic unsupervised-approach that automatically generates potential drug molecules given a prototype drug. We show that the molecules generated by the system are valid molecules and significantly different from the prototype drug. Out of the compounds generated by the system, we identified 35 FDA-approved drugs. As an example, our system generated Isoniazid – one of the main drugs for Tuberculosis. The system is currently being deployed for use in collaboration with pharmaceutical companies to further analyze the additional generated molecules.

KEYWORDS

Prototype-Based Drug Discovery, Drug Design, Deep Learning for Medicine

ACM Reference Format:

Shahar Harel and Kira Radinsky. 2018. Accelerating Prototype-Based Drug Discovery using Conditional Diversity Networks. In *KDD '18: The 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, August 19–23, 2018, London, United Kingdom*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3219819.3219882>

1 INTRODUCTION

Producing a new drug is an expensive and lengthy process that might take over 500 million dollars and over 10–15 years. The first stage is drug discovery, in which potential drugs are identified before selecting a candidate drug to progress to clinical trials. Although historically, some drugs have been discovered by accident (e.g., Minoxidil and Penicillin), today more systematic approaches are common. The most common method involves screening large libraries of chemicals in high-throughput screening assays (HTS) to identify an effect on potential targets (usually proteins). The goal

of such a process is to identify compounds that might modify the target activity, which might often result in a therapeutic effect.

While HTS is a commonly used method for novel drug discovery, it is common to start from a molecule which already has some of the desired properties. Such a molecule, usually called a “prototype”, might be extracted from a natural product or a drug on the market which could be improved upon. Intuitively, producing a chemically and structurally related substance to an existing active pharmaceutical compound usually improves on the efficacy of the prototype drug – reduces adverse effects, works on patients that are resistant to the prototype, and might be less expensive [7].

During this process of prototype-based drug discovery, an interdisciplinary team of scientists generates hypothesis about the required changes to the prototype. One might consider this process as a pattern recognition process – chemists, through their work, gain experience in identifying correlations between chemical structure retrosynthetic routes and pharmacological properties [30]. They rely on their expertise and medicinal chemistry intuition to create chemical hypotheses, which have been shown to be biased [29]. However, the chemical space is virtually infinite – the amount of synthetically valid chemicals which are potentially drug-like molecules is estimated to be between $10^{23} - 10^{60}$ [27]. In this work, we develop an algorithmic unsupervised approach to automatically generate potential drug molecules given a prototype drug.

It is common to encode molecular structures into SMILES notations (simplified molecular-input line-entry system) that preserve the chemical structural information. For example, Methyl isocyanate can be encoded using the following string: CN=C=O. We learn embeddings of drug-like molecules in molecule space represented by SMILES. To identify drug-like molecules, which are used to train our algorithm, we use the Lipinski criteria – a common chemical drug-design qualitative measure that estimates the structure bioavailability, solubility in water, potency, etc. [22]

Variational Auto Encoders (VAE) [18] are encoder-decoder architecture that attempts to learn the data distribution in a way that can later be sampled from to generate new examples. State-of-the-art results have been shown for generating images that resemble natural images, yet not identical to the train data [18, 21]. Training a vanilla VAE on drug-like molecules provides an ability to sample new molecules which intuitively should be drug-like [10]. In this work, we extend VAE to allow a conditional sampling – sampling an example from the data distribution (drug-like molecules) which is closer to a given input. This allows sampling molecules closer to a prototype drug, and thus increase probability of generating a valid drug with similar characteristics. Additionally, we add a diversity component that allows the sampling to be different from the prototype drug as well. We present a deep-learning approach

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '18, August 19–23, 2018, London, United Kingdom

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5552-0/18/08...\$15.00

<https://doi.org/10.1145/3219819.3219882>

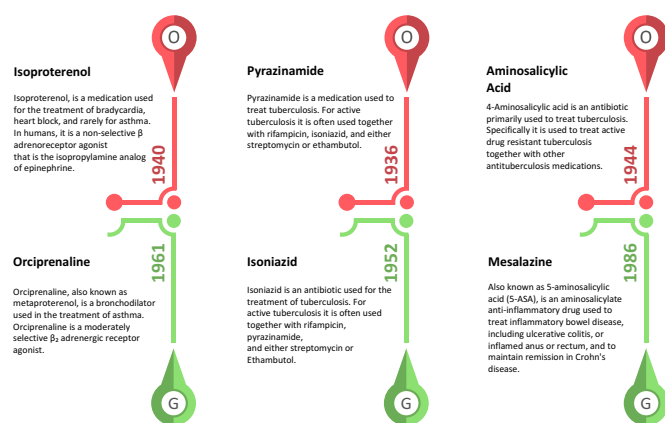


Figure 1: Drug development timeline, with example of drugs generated by CDN (bottom), using FDA approved drugs as prototypes (top).

which we call Conditional Diversity Networks (CDN), which allows the diverse conditioned sampling. The results show that the molecules CDN generates are similar to the prototype drugs yet significantly diverse. We show empirical results that the system generates high percentage of valid molecules. Additionally, we perform retrospective experiments and use drugs developed in the 1930's and 1940's as prototypes. The system was then able to generate new drugs, some of which discovered dozens of years after the prototype discovery (Figure 1).

One such example is the system discovery of the main drug for Tuberculosis – Isoniazid. Discovered in 1952, it is on the World Health Organization's List of "Essential Medicines, the most effective and safe medicines needed in a health system" [26]. In the retrospective experiment, we used as prototypes only drugs discovered until 1940. For the drug Pyrazinamide, first discovered in 1936, the system generated the SMILES notation of what today is known as Isoniazid. Pyrazinamide, although discovered in 1936, was not used until 1972 for treating Tuberculosis. Tuberculosis can become resistant to treatment if Pyrazinamide is used alone and therefore is always used in combination with Isoniazid and others. The combination reduces treatment time from 9 months to less than 3 months. This example shows promise on how substances that could not be used at the time of discovery can serve as a prototype for discovering new drugs. In collaboration with pharmaceutical companies additional generated molecules are being tested today. We believe our system lays the foundations to build algorithmically-directed HTS based on prototype drugs.

2 RELATED WORK

Over the past decade, deep neural networks (DNN) has been a game changer in various areas of machine learning research, such as computer vision [19, 34], natural language processing [24] and speech recognition [12]. Deep neural network most prominent success stories are observed in domains with access to large, raw (unprocessed) datasets. In such scenarios deep learning was able to achieve above human level performance. Compared with those domains, DNN in

chemistry relies heavily on engineered features representing a molecule [4, 23, 28]. Such approaches are semi optimal as they restrict the space of potential representations through the assumptions made by limiting to the chosen features [9].

More recent methods overcome this issue by leveraging advanced deep neural network models to learn chemical continuous representations (i.e., embeddings) based on a large datasets of molecular raw data. Molecular raw data can be represented in few ways, and processed with different deep architectures. Among those we can find 2D/3D images served as input to a convolutional neural network (CNN) [9, 35], molecular graph representation paired with neural graph embedding methods [5, 39], and SMILES strings – modeled as a language model with recurrent neural network (RNN) [2, 10, 31]. Coley et al., Jin et al., Schwaller et al. leverage the embeddings for numerous supervised prediction task, e.g. predicting outcomes of complex organic chemistry reactions.

Recently, deep generative models have opened up new avenues for leveraging molecular embeddings for unsupervised tasks such as molecule generation, and drug design. Most methods aim at generating valid molecules. For example Segler et al. train RNN as a language model to predict the next character in a SMILES string. After training, the model can be used to generate new sequences corresponding to new molecules. Gómez-Bombarelli et al. leverage on the VAE [18] generative model, to learn dense molecular embedding space. At test time, the model is able to generate new molecules from samples of the prior distribution enforced on the latent representation during training. In this general form of generation, we can only hope to achieve the task of generating molecule libraries with no specific chemical/biological characteristics, but the characteristics of the training data. Others extend this approach by tuning the model on a dataset of molecules with specific characteristics [32], or by applying post processing methods, such as Bayesian Optimization [10, 13] and Reinforcement Learning [25].

In this work, we target the problem of generating drug-like molecules and show that training vanilla generative models on this family shows limited results (Section 5), both for generating diverse novel molecules and for generating drugs. Following the common chemical approach, we focus the generative approach on a given prototype. This helps "guide" the search process around the prototype in the chemical space. Given prototypes can be drug-like molecules or known drugs. We introduce parametrized diversity and design an end-to-end neural network solution to train the model to represent the chemical space, and to allow for diversity driven prototype based exploration and novel molecule generation.

3 METHODS

We define the problem of prototype-driven hypothesis generation as a conditional data generation process. The model operates on a given molecule prototype and generates various molecules as candidates. The generated molecules should be novel and share desired properties with the prototype. We start by reviewing how molecules are represented as text (Section 3.1), and then present the generative model (Section 3.2). Our generative model builds upon recent methods for deep representation learning. We train a stochastic neural network to learn internal molecule representation (embedding). After obtaining the molecule embedding we further

utilize the stochastic component of the neural architectures to introduce parametrized diversity into the generation process. The architecture of our proposed solution is presented in Section 3.3.

3.1 Molecule Representation

The choice of representation of molecules is at the heart of any computer-based chemical analysis. For molecule generation, it is of crucial importance, as the task is to both analyze and generate objects of the same representation. Cadeddu et al. showed that organic molecules contain fragments whose rank distribution is essentially identical to that of sentence fragments. The consequence of this discovery is that the vocabulary of organic chemistry and human language follow very similar laws. Intuitively, there is an analogy between a chemist understanding of a compound and a language speaker understanding of a word. This introduces a potential to leverage recent advances in linguistics-based analysis, and deep sequence models in particular.

A SMILES string is a commonly-used text encoding for organic molecules. SMILES represents a molecule as a sequence of characters corresponding to atoms as well as special characters denoting opening and closure of rings and branches. For example c and C represent aromatic and aliphatic carbon atoms, O represents oxygen, -, = and # represent single, double and triple bonds [36]. Then a molecules, such as Benzene, is represented in SMILES notation as c1ccccc1. It has already been shown that SMILES representation of molecules has been effective in chemoinformatics [2, 10, 31, 32]. This has strengthened our belief that recent advances in the field of deep computational linguistics and generative models might have an immense impact on prototype based drug development.

3.2 Molecule Driven Hypothesis Generation

Generative models have been applied for many tasks, e.g., image generation. The models synthesized new images which resembled the database the models were trained on [18, 21]. One of the most popular generative frameworks are Variational Autoencoders (VAE) [18]. VAE are encoder-decoder models that use a variational approach for latent representation learning. Intuitively, the encoder is constrained to generate latent representations that follow a prior. During generation, latent vectors are sampled from the priors and passed to the decoder that generates the new representation. We leverage VAE for the task of molecule generation. The stochasticity allows integrating chemical diversity into the generation process. However, application of generative models for molecule generation have shown limited results [10]. Unlike image generation, where each image is valid, when we aim at molecule generation, not each representation is a valid molecule representation. Intuitively, when we sample from the prior for image generation – the space of images is much more dense than that of valid molecules. Therefore, many image samples are valid compared to randomly generated molecules representations. We hypothesize that a constrained generation next to a known prototype, rather than a non-constrained sampling, will yield better molecule generation. We extend VAE generation process to condition on a prototype, i.e., generate molecules closer to a given drug. Intuitively, directing the sampling process closer to existing prototype drugs might yield valid molecules that carry similar characteristics to the prototype yet provide diversity. Our

results provide evidence that a conditioned sample along side a diversity component yields more valid and novel results. If conditioned on known drugs, the system is able to generate drugs discovered years after the prototype (Section 3.2).

More formally, we assume a molecule M has a latent representation z that captures the main factors of variation in the chemical space. We model the covariates $z_i | M$ as Gaussians ($z_i \sim \mathcal{N}(\mu_i, \sigma_i)$). With the latent representation z at hand, we want to generate a candidate molecule in a SMILES discrete form, therefore we define the generative model $\hat{y} | z \sim \text{Multi}(\theta)$, where y is the generated candidate, Formally:

$$q(z | M) = \prod_{i=1}^{D_z} \mathcal{N}(\mu_i = \hat{\mu}_i, \sigma_i = \hat{\sigma}_i) \quad (1)$$

$$p(\hat{Y} | z) = \text{Multi}(\theta = \hat{\theta}) \quad (2)$$

Where q is approximated via encoder neural network function, applied on molecule M as input, and outputs the latent feature parameterization $(\hat{\mu}, \hat{\sigma})$ of the molecule. We then sample instance from this parametrization to obtain the final encoded output z ; p is represented via decoder neural network, applied on the molecule sampled feature instance z as input, and generates the output molecule as described below.

Generating a molecule as a SMILES string reflects multinomial distribution over the atoms space. Each atom is represented via a character. We form the character generation process as an iterative process, each character y_i is generated based on the hidden encoded representation z and the formerly generated y_i 's. In total, the output of this step is a string $\hat{y}_i = \{y_1, \dots, y_N\}$, where N is a pre-defined maximal generation length. Formally, for a single character y_i

$$P(\hat{y}_i | \hat{y}_1, \dots, y_{i-1}, z) = f(s_i, y_{i-1}) \quad (3)$$

Where y_{i-1} is the character embedding correspond to the last character generated, and s_i is a state at step i representing the current processed information of both the molecule latent representation z , and the formerly generated y_i 's up to $i-2$.

Our goal is to create a molecule that is different from the original molecule M . Intuitively, we wish to explore the chemical space around the molecule M . Therefore, during generation process we introduce a diversity component noising the multidimensional Gaussian parameters used for sampling the hidden vector z . More formally, to introduce *diversity* to our generation process, we instantiate our encoder output parameters with a diversity layer. Intuitively, the diversity layer outputs a noisy sample from a distribution centered as the encoder suggested, but with larger variance. This allows us to explore the molecule space around an origin molecule, with tune-able amount of diversity, corresponds to variability in chemical space. The diversity layer samples noisy instance according to the encoded Gaussian parameters and a diversity parameter D .

The output of the diversity layer is a sample from a conditional diverse distribution, described as follows:

Given the encoder outputs: vector of means $\hat{\mu}_i$ and standard deviations $\hat{\sigma}_i$, and random noise sample n - from Gaussian distribution with diversity parameter D - [$n \sim \mathcal{N}(0, D)$].

$$\text{Diverse } z = (n \times \hat{\sigma}_i) + \hat{\mu}_i \sim \mathcal{N}(\hat{\mu}_i, \hat{\sigma}_i^2 \times D) \quad (4)$$

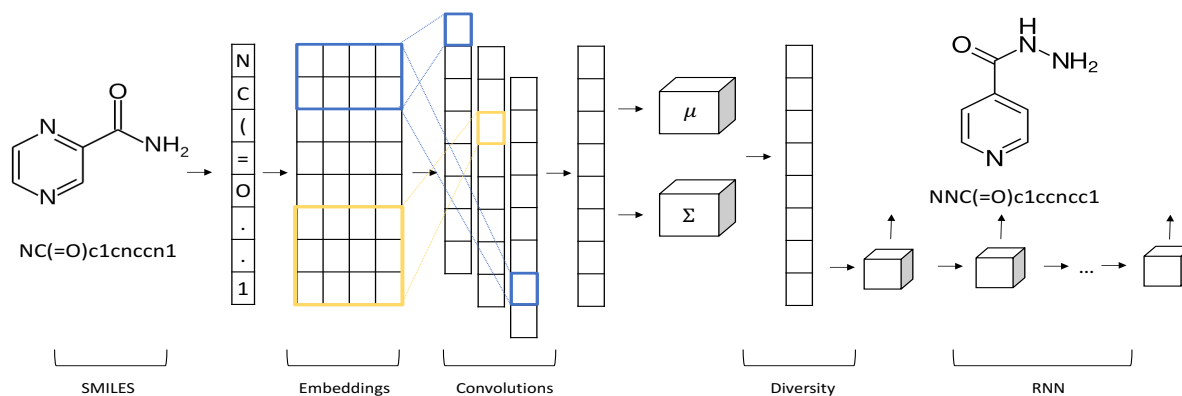


Figure 2: CDN end-to-end neural net architecture

We obtain instance from the diverse distribution as our final noisy encoded representation (z) for the compound M , used as the base for the decoder diversity-driven molecule generation.

We note that during training our diversity parameter D is set to 1. Thus z instance is sampled from the non-diverse distribution suggested by the computed parameters. Tuning this parameter at generation time allow us to explore the space around the prototype.

3.3 CDN Architecture

We leverage recent advances in generative models and deep learning for natural language process (NLP) to form the prototype hypothesis generation process as an end-to-end deep neural network solution. Figure 2 presents *CDN* (Conditional Diversity Network) architecture. *CDN* starts by encoding the molecule (in SMILES notation) using the encoder function. First, encoding each character in the SMILES representation into its d dimensional embedding, then applying convolutions over various substring (filter) sizes (e.g. correspond to chemical substructures). A similar encoder architecture was suggested for NLP tasks, such as sentence classification [16]. The extracted features are then concatenated and fully connected layers are applied. The outputs of the encoder are considered as a vector of means and a vector of standard deviations, representing the distributions of features for the prototype. In VAE, samples based on those parameter vectors are then fed into a decoder. The goal is to optimize reconstruction of the original input and constraint the representation to a known prior. During generation, the vectors of features are sampled from the prior distribution and their output is passed to the decoder that generates a new representation.

We extend the VAE generation process by adding a *diversity layer*. During generation, instead of sampling from the prior means and standard deviations, we first feed a prototype. We sample from the prototype feature distribution with parametrized diversity (Section 3.2) to form the prototype latent representation – served as input for the decoder.

As described in section 3.2, our decoder is a sequential generator. By generating sequentially, we form another parameter of variability in the generated data, by introducing minor variations into the molecule generated during generation. This is the main component

of many other works on molecule generation [6, 11, 32] to introduce diversity into the generation process. We later show that our diversity layer can introduce diversity beyond this component.

We represent our decoder as a recurrent neural network (LSTM). The decoder receives the encoder output as its input. The encoded representation forms the first state of the decoder. The decoder then generates the compound sequentially (character by character) by operating on the distribution over characters at each time step, based on its updated state and the input character from former step.

During training, we feed the decoder with the correct next symbol, even if it was predicted wrongly [37]. During generation, we experiment with two options for generating the next symbol: one by selecting the best scored character from distribution over symbols (argmax), and the second is by sampling from the same distribution. By introducing sampling into the generation, we are able to increase the amount of variability we generate during generation. The model is trained to reconstruct the input prototype from a low dimensional continuous representation, by minimizing a discrete reconstruction loss. Formally, to minimize the reconstruction error on a discrete molecule representation, we use the cross-entropy loss, defined as:

$$H(y, \hat{y}) = \sum y_i \log \hat{y}_i \quad (5)$$

We note that we minimize the variational lower bound [18], which is essentially optimizing the reconstruction error while constraining the latent distribution with a prior. To reconstruct syntactically valid SMILES, the generative model would have to learn the SMILES, which includes keeping track of atoms, rings and brackets to eventually close them. In this case, the lower dimension representation that can be reconstructed into a valid prototype, is a highly informative representation. In other words, by minimizing the reconstruction error, we want learn a prototype continuous representation that captures the coordinates along the main factors of variation in the chemical space. This representation is the base for further diversifying the molecule generation process.

4 EXPERIMENTAL SETTINGS

In this section we provide details on the datasets, hyperparameter setting, and the training in general. Then, we mention the methods compared and used in our experiments.

4.1 Model Details

CDN was trained using a Tensorflow API [1]. We use the Adam algorithm [17] to optimize all the parameters of the network jointly, regarding weights initialization - the atoms embedding were initialized using a random uniform distribution ranging from -0.1 to 0.1, convolution weights used truncated normal with std 0.1, all other weights used the Xavier initialization [8], biases were initialized with constant. To reduce overfitting, we included an early stopping criteria based on the validation set reconstruction error. We use exponential decay factor on learning rate, and the teacher forcing method [37] during training. In total, table 1 presents CDN hyperparameter configuration.

The code for our system is available over github¹ for further research in the community

Parameter	Value
max molecule length	50
char embedding size	128
filter sizes	3, 4, 5, 6
number of filters	128
latent z dimension	300
batch size	64
initial learning rate	0.001
LSTM cell units	150

Table 1: CDN hyperparameter configuration

4.2 Datasets

4.2.1 Drug-like molecules database. The ZINC database [14] contains commercially available compounds for structure based virtual screening. In addition, the database have subsets of ZINC filtered by physical properties. One such filtering is based on Lipinski’s rule of five [22] – a heuristic method to evaluate if a molecule can be a drug. The subset contains over 10 million unique drug-like compounds. CDN was trained on a subset with approximately 200k drug-like compounds extracted at random from the ZINC drug-like database. The subset was further divided to train/validation/test sets, with 5k compounds for validation/test sets, and the rest for training set. The subsets are used for training the model (train), evaluating hyperparameters and stopping criteria (validation), and for method evaluation and experiments (test).

4.2.2 Drug database. The DrugBank database [38] is a bioinformatics and cheminformatics resource that combines detailed drug data with comprehensive drug target information. For retrospective experiments (Section 5.2), we extracted a test set of 869 FDA approved drugs from the DrugBank database. Note, our system is not trained on drugs, but rather presented with drug prototypes only during generation.

4.3 Compared Methods

As discussed in Section 2, not much work has been done in the area of deep drug generation and specifically not on the diversity aspect of the generation. To the best of our knowledge, the works that do consider the task, do not aim at prototyping specific compound, but training for unconditional molecule generation and later apply post-processing to achieve general molecular characteristics. We compare our methods to the state of the art models for molecule

generation on the reconstruction criteria, and further show that our model is able to build on top of those models to apply diversity. Specifically, we compare all following methods

- (1) *Seq2Seq* [33] - An autoencoder architecture applied on sequence data for prediction of sequences. Both encoder and decoder are recurrent neural networks (RNN). Although the model is in general deterministic, it is able to bring stochasticity (and thus novelty) into the molecules generation process by setting the RNN decoder to sample from the distribution over characters in every time step instead of predicting the topmost next character. We therefore consider two baselines – one using the Argmax method and the other utilizing the Sampling method to reach diversity.
- (2) *Conv2Seq* - To conform better with CDN parameter setting that utilize on CNNs, we implement a second auto encoder same as the previous method but with convolution encoder.
- (3) VAE [18] – a vanilla implementation of VAE. This model generates new molecules from unit Gaussian random samples, regardless of prototypes.
- (4) *CDN-VAE* - Our diversity model on top of variational auto encoder. D is the diversity parameter of Equation 4 The higher the D , the higher the diversity induced. We note, that for $D = 1$, the model extends VAE for a conditional setting but without diversity.

5 EXPERIMENTS

Designing an evaluation framework for novel generated molecules is a challenging task. Evaluating a new chemical compound might take years. We start by evaluating the most basic characteristics of the generated SMILE, namely, its chemical structure validity, its syntactical similarity with respect to the prototype, and how novel the generated molecules with respect to the prototype drug (Section 5.1). Once we confirm CDN ability to generate valid molecules that share similarity with a prototype, we conduct several drug related experiments to show CDN capabilities in the real world for generating new drugs. We design a retrospective experiment and measure how many of the generated molecules are indeed known drugs (Section 5.2). Next, we present evaluation of the trade off between the molecules reconstruction accuracy and novelty as a function of CDN diversity component (Sections 5.4, 5.5). Finally, we take a deeper look at the molecule feature representation constructed with CDN in the context of drug families and show that it is able to embed close drug families via similar embeddings (Section 5.6).

5.1 Novel Molecules Generation

Our main goal is to create novel molecules that carry similarities to the prototype. Thus, the metric of reconstruction is an important metric. We examine the methods on the task of prototype reconstruction on a test set of 5k ZINC drug-like compounds. To explicitly address the reconstruction accuracy and validity vs. the generated molecules diversity, we measure the following metrics:

- (1) *Reconstruction Accuracy (Acc)* - Character-level accuracy with the input prototype served also as target.
- (2) *Valid Molecule Percentage (Valid)* - Percentage of valid molecules. There are several numeric validations performed on molecules

¹https://github.com/shaharharel/CDN_Molecule

Model	Acc	Valid	Novel	Acc @ 1k	Valid @ 1k	Novel @ 1k
Seq2Seq - Argmax	0.94	0.93	0.13	-	-	-
Seq2Seq - Sampling	0.91	0.88	0.19	0.92	0.89	32.5
Conv2Seq - Argmax	0.92	0.85	0.14	-	-	-
Conv2Seq - Sampling	0.89	0.77	0.18	0.88	0.76	35.2
VAE	- ²	0.58	-	-	-	-
CDN - D=1	0.91	0.89	0.19	0.9	0.89	8
CDN - D=2	0.82	0.81	0.26	0.81	0.8	66.6
CDN - D=3	0.64	0.63	0.37	0.65	0.65	227

Table 2: Evaluation of CDN and baselines for diversity and validity of generated molecules

Input Drug	Input SMILES	Generated Drug	Generated SMILES
Aminosalicylic	<chem>Nc1ccc(C(=O)O)c(O)c1</chem>	Mesalazine	<chem>Nc1ccc(O)c(C(=O)O)c1</chem>
Pyrazinamide	<chem>NC(=O)c1cnccn1</chem>	Isoniazid	<chem>NNC(=O)c1cncc1</chem>
Protriptyline	<chem>CNCCCC1c2ccccc2C=Cc2ccccc21</chem>	Desipramine	<chem>CNCCCC1c2ccccc2CCc2ccccc21</chem>
Phenelzine	<chem>NNCCc1ccccc1</chem>	Isoniazid	<chem>NNC(=O)c1cncc1</chem>
Isoprotorenol	<chem>CC(C)NCC(O)c1ccc(O)c(O)c1</chem>	Orciprenaline	<chem>CC(C)NCC(O)c1cc(O)cc(O)c1</chem>
Pheniramine	<chem>CN(C)CCC(c1ccccc1)c1cccn1</chem>	Tripelennamine	<chem>CN(C)CCN(Cc1ccccc1)c1cccn1</chem>

Table 3: Sample of automatically generated drugs and the drug served as prototype to the generation process

representation to validate its correctness. We used Rdkit [20] library to measure validity of the generated compounds.

- (3) *Novel Molecule Percentage (Novel)* - A novel molecule is both a valid molecule, and different from the prototype.

To be able to measure the molecule generation capabilities over various Gaussian samples for the same prototype compound (we want to be able to generate *several* compounds related to the origin compound), we also measure all the above metrics with $@k$ notation. In our context, $@k$ represents that for each prototype compound, we run *CDN* generation process with k instances of random noises parametrized with diversity D . We note that to measure *Novel@k*, we count how many unique molecules generated – a novel molecule is counted only once, even if it was generated with various Gaussian samples for the prototype. The metric of *Novel@k* is not normalized, thus, the semantics of this metric should be, intuitively, interpreted as how many unique molecules were generated for a prototype and 1000 Gaussian samples.

Table 2 presents the results of *CDN* and the baselines on the metrics above. Analyzing *Acc*, *Valid* and *Novel* metrics, we observe that with diversity level of 1 (non-diverse sampling), *CDN* generates similar diversity to the baselines. Increasing D , significantly increases the amount of novel generated molecules, while reducing the level of accuracy and the valid molecules rate. This result stems from the intuition that as the representation becomes noisier, it is harder for the model to reconstruct the original prototype. Addressing the $@k$ metrics, we observe *CDN* is able to maintain the accuracy and validity levels with many random samples used for generation, while generating various unique molecules for the same prototype input, with the number of unique molecule significantly increasing with the diversity parameter D .

5.2 Drug Generation

The main aim of this work is to generate novel molecules with desired properties (characterized by the prototype molecule), by

²As no prototype is given, there is no reconstruction to measure

searching the chemical space around the prototype. To check the immediate benefits (i.e., without further screening the generated compound) of our approach to a real world task, we conduct a retrospective experiment in the drug domain. We apply our method on a test set of FDA approved drugs as prototypes. We note that none of the drugs was observed in the training data, which was composed of only drug-like molecules.

Evaluation on this task is harder since our goal is to generate drugs, and we cannot a-priori know if the generated molecule has the desired characteristics of a drug without further experimenting with the compound. We therefore consider as gold standard a test set of 869 approved known drugs. Although this test set is very small in compare with the enormous molecule space, some approved drugs are chemically similar and share similar therapeutic characteristics, thus we hypothesize that by applying *CDN* on FDA approved drugs as prototypes, we might be able to generate other known compounds / drugs with similar characteristics.

Interestingly, targeting some existing drugs as prototypes, our model was able to generate molecules that also appear in the FDA approved drugs list and are closely related to the prototype, both in the chemical aspects and by their medical use (i.e. targeting the same biological mechanism of action). Table 3 presents a sample of the drugs generated. In total, we run the baselines and *CDN* variants over all 869 approved drugs dataset as prototypes, with 1000 Gaussian samples in each run. Table 4 presents the number of FDA approved generated drugs with each method. We also present the percentage those drugs constitute from the valid molecules generated. We draw the reader attention to the negligible chance of generating a drug using exhaustive search without constraints (e.g., using HTS). We observe that the VAE could not produce any known drug. We hypothesize that this stems from the fact that VAE randomly generates a molecule and not based on a prototype. *CDN* with no diversity and the other baselines generated 9–12 drugs. This result emphasize how the variability that the decoders present during sampling contributes to the generation of known drugs.

Model	#Drugs	% from Generated Valid Molecules
VAE	0	0%
Seq2Seq	12	0.002%
Conv2Seq	9	0.0018%
CDN-VAE D=1	12	0.0023%
CDN-VAE D=2	22	0.005%
CDN-VAE D=3	35	0.01%

Table 4: Automatically generated FDA approved drugs. We present the percentage of the FDA-approved drug from the total valid molecules generated by each method

One would expect that higher D values, that increase the distance from the prototype drug (see analysis in Section 5.4), might create molecules with low probability of being drugs. However, we observe that for higher D values of our diversity layer (CDN-VAE D=2 and CDN-VAE D=3), the amount of known drugs increases significantly. One should remember that the model doesn't have any "drug" understanding – the model was only trained given drug-like molecules, and *all known drugs were eliminated from the training*. The key here is the chemical similarity drugs share. Thus, by targeting a drug molecule as prototype to the generation process, our model is able to chemically diversify the prototype drug in a way that generated another known drugs. We are encouraged by the results that *CDN* was able to generate a significant number of already known drugs. We are currently testing with pharmaceuticals companies the additional generated molecules.

5.3 Qualitative Examples

We present a few qualitative examples of the drugs generated. We would like to explore whether the application of the system on drugs developed up until a certain year might find drugs that will be discovered years later. During training we eliminate all known drugs from the ZINC database and we present as prototypes a single drug. Figure 1 presents a timeline with example pairs of origin (top row) and generated (bottom row) molecules, with the year of the drug first use. By using *CDN* we could have generated the bottom molecules directly when we knew the origin molecules, possibly sparing years of research. The system was able to identify the main drug for Tuberculosis – Isoniazid using an initial prototype of the disease that was never used due to its side effects (Pyrazinamide). Additional intriguing example is the generation of Orciprenaline which is used to treat Asthma from a prototype drug that was mainly used for heart block, and very rarely for asthma. These pairs are closely related in their therapeutic effect, but a few changes for the molecule were needed to reposition it for Asthma treatment. Another interesting discovery was Mesalazine, used to treat inflammatory bowel disease based on an antibiotic primarily used to treat tuberculosis. discovered about 40 years before.

5.4 Diversity Mechanisms

A common method to employ diversity in encoder-decoder models is to employ a sampling decoder into the architecture [6, 11, 32]. The diversity is introduced by sampling from distribution over characters in each time step of generation, rather than choosing the

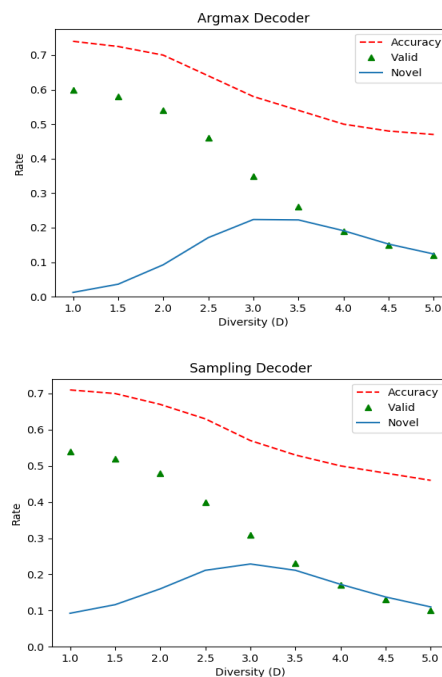


Figure 3: Diversity parameter effect on performance.

Model	Acc	Valid	Novel	A@1k	V@1k	N@1k
CDND=1	0.88	0.78	0.19	0.88	0.79	39.5
CDND=2	0.8	0.69	0.25	0.79	0.68	94
CDND=3	0.57	0.39	0.27	0.56	0.38	179

Table 5: CDN performance using a sampling decoder.

topmost (argmax) character at test time. We analyze the contribution of the diversity layer D for *CDN* presented in this work along side a sampling decoder as well. Table 5 presents *CDN* performance on the previous metrics (Section 5.1), but with a sampling decoder. We compare *CDN* with sampling but with no diversity component ($D = 1$) to *CDN* with sampling with higher values of D and observe that the diversity parameter is able to introduce additional diversity beyond the sampling decoder component.

To analyze the behavior of the diversity parameter D on the accuracy/validity and novelty trade-off in the drug domain, we generated samples for the FDA approved test set (Section 4.2), with various configurations of the diversity parameter D . Figure 3 presents the results for the two types of decoder functions. As we hypothesized, with both decoders, increasing the value of the diversity parameter D , significantly increases the amount of novel molecules generated. As we expected, the novelty is not free, we observe lower accuracy and lower valid rates for increased diversity. Comparing argmax with sampling decoders, we observe that in general, sampling has lower accuracy and valid rate, but for low diversity value the sampling method generates significantly more novelty than the argmax. This behavior reduces for higher values of the diversity parameter, were both methods generates similar rates of novelty. We also observe the novelty rate reduces at some point of increased diversity value. This is quite expected because for large values of

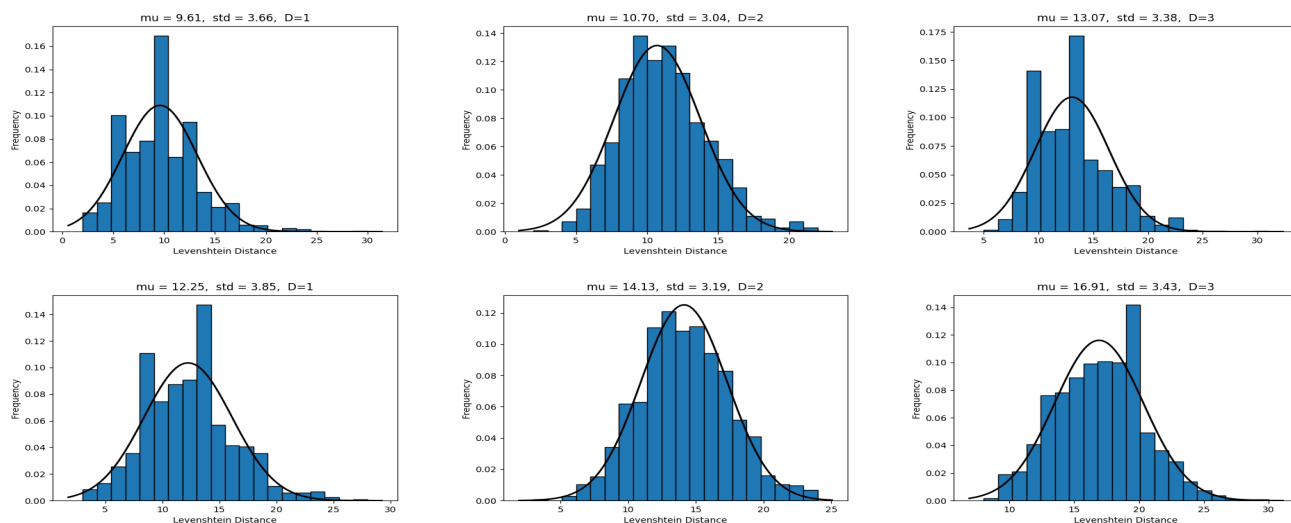


Figure 4: Levenshtein distance histograms for analyzing the diversity generated by *CDN*. Top - Prototype molecule vs Generated molecule distances. Bottom - within generated molecule population distances.

diversity, the latent molecule representation sampled with larger noise, thus at some point the generator is not able to recover much valid molecules in general, and novel ones in particular.

5.5 Molecular Variations

We would like to analyze not only whether a molecule is different from the prototype molecule but also quantify the diversity of the molecules with respect to the prototype molecule. Additionally, we would like to validate that the generated molecules originating from a prototype are also diverse with respect to one another. We compare the Levenshtein distances of the generated SMILES within the generated population and with respect to the prototype, used as input for a specific generation instance. We apply *CDN* on the drug-like test set as prototypes. We note that we count only valid molecules generated in all evaluations. Figure 4 presents histograms of the Levenshtein distances for the generated molecules, with approximated Gaussian parameters and curve on top of the histograms. The top row represents the input prototype compared to the generated molecules Levenshtein distance distribution for different configurations of the diversity parameter D (increasing D from left to right). The bottom row represents the inner-generated population Levenshtein distance distribution for various values of D . On both type of distance evaluations (rows), we observe significantly larger Levenshtein distances for larger values of D , thus indicating positive effect of the diversity parameter D on both the distance from the prototype molecule, and the average inner distance between molecules that were generated with different random samples to the same prototype. Additionally, we observe *CDN* diversity in generation is not limited to generating diversity with respect to the origin molecule, but also generates diversity within the generated population for a specific prototype, with higher amount of diversity tuned with the diversity parameter D .

Class	Cosine	L2	L1
Thiazide Diuretics	0.872	.95	.908
Benzodiazepines	0.923	.883	.859
β -Blocker	0.866	.849	.822
NSAIDs	0.955	.853	.833
Across Drugs	1.0	1.0	1.0

Table 6: In class and across drugs normalized distances computed on various drug classes.

5.6 Molecule Representation in Latent Space

Encoder-decoder settings produce intermediate representations of their input. In this section, we analyze the quality of those representations. During *CDN* generation process, we first encode molecule into a low dimensional vector space with the encoder function. We refer to the output vector as the molecule embedding. To evaluate the embeddings, we leverage them for the task of drug classification. Intuitively, if the embeddings captures enough information for drug classification, we might rely on this representation for molecule generation. We note that for the task of encoding the molecule feature representation, we set the diversity parameter $D = 1$, but one should remember that the representation is still instantiated from unit Gaussian, and thus is not deterministic.

A drug class is a set of medications that have similar chemical structures, or the same mechanism of action (i.e., bind to the same biological target). In Table 6 we report embedding vector normalized distances in-class and across various drug classes. Thiazide and Benzodiazepines are chemical classes while β -Blocker and NSAIDs³ are classes representing mechanism of action. We observe all in-classes distances are significantly lower than across class. We conclude that although our molecule representation is noisy by the stochastic nature of *CDN*, similarities in the embedding space are able to reflect significant similarities among various drug classes.

³NSAIDs - Nonsteroidal anti-inflammatory agents

6 CONCLUSIONS

Drug discovery is the process of identifying potential molecules that can be targeted for drugs. Common methods include systematic generation and testing of molecules via HTS. However, the molecular space is very large. Additional approaches require chemist to identify potential drugs based on their knowledge. Usually, they would start from a known compound in nature or known drug and identify potential changes. Approaches in machine learning today mainly focused on non-controlled molecule generation using generative mechanisms, such as VAE. The approaches were limited in their ability to generate both valid and novel molecules. In this work, we presented a prototype-based approach for generating drug-like molecules. We adopt the chemist approach of "borrowing" from nature or focusing on known drugs. We hypothesize that biasing the molecule generation towards known drugs might yield valid molecules. We train our model on drug-like molecules, and during generation extend VAE to, intuitively, search closer to the prototype (which can be a drug). We add additional component to diversify the molecules generated. We present results that show that many of the molecules generated are both valid and novel. When conditioning on drugs, we observe our system was able to generate known drugs that it never encountered before. The system is currently being deployed for use in collaboration with pharmaceutical companies to further analyze the additional generated molecules.

REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. TensorFlow: A System for Large-Scale Machine Learning. In *OSDI*, Vol. 16. 265–283.
- [2] Esben Jannik Bjerrum. 2017. Smiles enumeration as data augmentation for neural network modeling of molecules. *arXiv preprint arXiv:1703.07076* (2017).
- [3] Andrea Cadeddu, Elizabeth K Wylie, Janusz Jurczak, Matthew Wampler-Doty, and Bartosz A Grzybowski. 2014. Organic chemistry as a language and the implications of chemical linguistics for structural and retrosynthetic analyses. *Angewandte Chemie* 126, 31 (2014), 8246–8250.
- [4] Connor W Coley, Regina Barzilay, William H Green, Tommi S Jaakkola, and Klavs F Jensen. 2017. Convolutional embedding of attributed molecular graphs for physical property prediction. *Journal of chemical information and modeling* 57, 8 (2017), 1757–1772.
- [5] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. 2015. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in neural information processing systems*. 2224–2232.
- [6] Peter Ertl, Richard Lewis, Eric Martin, and Valery Polyakov. 2017. In silico generation of novel, drug-like chemical matter using the LSTM neural network. *arXiv preprint arXiv:1712.07449* (2017).
- [7] Silvio Garattini. 1997. Are me-too drugs justified? *Journal of Nephrology* 10, 6 (1997), 283–294.
- [8] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. 249–256.
- [9] Garrett B Goh, Charles Siegel, Abhinav Vishnu, and Nathan O Hodas. 2017. ChemNet: A Transferable and Generalizable Deep Neural Network for Small-Molecule Property Prediction. *arXiv preprint arXiv:1712.02734* (2017).
- [10] Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamin Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. 2016. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science* (2016).
- [11] Anvita Gupta, Alex T Müller, Berend JH Huisman, Jens A Fuchs, Petra Schneider, and Gisbert Schneider. 2017. Generative Recurrent Networks for De Novo Drug Design. *Molecular informatics* (2017).
- [12] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al. 2012. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine* 29, 6 (2012), 82–97.
- [13] Hisaki Ikebata, Kenta Hongo, Tetsu Isomura, Ryo Maezono, and Ryo Yoshida. 2017. Bayesian molecular design with a chemical language model. *Journal of computer-aided molecular design* 31, 4 (2017), 379–391.
- [14] John J Irwin, Teague Sterling, Michael M Mysinger, Erin S Bolstad, and Ryan G Coleman. 2012. ZINC: a free tool to discover chemistry for biology. *Journal of chemical information and modeling* 52, 7 (2012), 1757–1768.
- [15] Wengong Jin, Connor Coley, Regina Barzilay, and Tommi Jaakkola. 2017. Predicting Organic Reaction Outcomes with Weisfeiler-Lehman Network. In *Advances in Neural Information Processing Systems*. 2604–2613.
- [16] Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882* (2014).
- [17] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [18] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.
- [20] Greg Landrum. 2006. RDKit: Open-source cheminformatics. *Online*. <http://www.rdkit.org>. Accessed 3, 04 (2006), 2012.
- [21] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. 2015. Autoencoding beyond pixels using a learned similarity metric. *arXiv preprint arXiv:1512.09300* (2015).
- [22] Christopher A Lipinski. 2000. Drug-like properties and the causes of poor solubility and poor permeability. *Journal of pharmacological and toxicological methods* 44, 1 (2000), 235–249.
- [23] Andreas Mayr, Günter Klambauer, Thomas Unterthiner, and Sepp Hochreiter. 2016. DeepTox: toxicity prediction using deep learning. *Frontiers in Environmental Science* 3 (2016), 80.
- [24] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.
- [25] Marcus Olivcrona, Thomas Blaschke, Ola Engkvist, and Hongming Chen. 2017. Molecular de-novo design through deep reinforcement learning. *Journal of cheminformatics* 9, 1 (2017), 48.
- [26] World Health Organization et al. 2003. The selection and use of essential medicines: report of the WHO Expert Committee, 2002:(including the 12th model list of essential medicines). (2003).
- [27] Pavel G Polishchuk, Timur I Madzhidov, and Alexandre Varnek. 2013. Estimation of the size of drug-like chemical space based on GDB-17 data. *Journal of computer-aided molecular design* 27, 8 (2013), 675–679.
- [28] Bharath Ramsundar, Steven Kearnes, Patrick Riley, Dale Webster, David Konerding, and Vijay Pande. 2015. Massively multitask networks for drug discovery. *arXiv preprint arXiv:1502.02072* (2015).
- [29] Volker Schneck and Jonas Boström. 2006. Computational chemistry-driven decision making in lead generation. *Drug discovery today* 11, 1-2 (2006), 43–50.
- [30] Gisbert Schneider. 2017. Automating drug discovery. *Nature Reviews Drug Discovery* (2017).
- [31] Philippe Schwaller, Theophile Gaudin, David Lanyi, Costas Bekas, and Teodoro Laino. 2017. "Found in Translation": Predicting Outcome of Complex Organic Chemistry Reactions using Neural Sequence-to-Sequence Models. *arXiv preprint arXiv:1711.04810* (2017).
- [32] Marwin HS Segler, Thierry Kogej, Christian Tyrchan, and Mark P Waller. 2017. Generating focussed molecule libraries for drug discovery with recurrent neural networks. *arXiv preprint arXiv:1701.01329* (2017).
- [33] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*. 3104–3112.
- [34] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1–9.
- [35] Izhar Wallach, Michael Dzamba, and Abraham Heifets. 2015. Atomnet: A deep convolutional neural network for bioactivity prediction in structure-based drug discovery. *arXiv preprint arXiv:1510.02855* (2015).
- [36] David Weininger. 1988. SMILES, 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.
- [37] Ronald J Williams and David Zipser. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural computation* 1, 2 (1989), 270–280.
- [38] David S Wishart, Craig Knox, An Chi Guo, Savita Shrivastava, Murtaza Hassanali, Paul Stothard, Zhan Chang, and Jennifer Woolsey. 2006. DrugBank: a comprehensive resource for in silico drug discovery and exploration. *Nucleic acids research* 34, suppl_1 (2006), D668–D672.
- [39] Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. 2018. MoleculeNet: a benchmark for molecular machine learning. *Chemical Science* 9, 2 (2018), 513–530.