

EVENING THE SCORE: TARGETING SARS-CoV-2 PROTEASE INHIBITION IN GRAPH GENERATIVE MODELS FOR THERAPEUTIC CANDIDATES

Jenna Bilbrey, Sutanay Choudhury, Neeraj Kumar

Pacific Northwest National Laboratory

Richland, WA, USA

{jenna.pope, sutanay.choudhury, neeraj.kumar}@pnnl.gov

Logan Ward, Ganesh Sivaraman

Argonne National Laboratory

Lemont, IL, USA

{lward, gsivaraman}@anl.gov

ABSTRACT

We examine a pair of graph generative models for the therapeutic design of novel drug candidates targeting SARS-CoV-2 viral proteins. Due to a sense of urgency, we chose well-validated models with unique strengths: an autoencoder that generates molecules with similar structures to a dataset of drugs with anti-SARS activity and a reinforcement learning algorithm that generates highly novel molecules. During generation, we explore optimization toward several design targets to balance druglikeness, synthetic accessibility, and anti-SARS activity based on IC_{50} . This generative framework¹ will accelerate drug discovery in future pandemics through the high-throughput generation of targeted therapeutic candidates.

1 INTRODUCTION

Discovering lead candidates for SARS-CoV-2 presents a challenge to the scientific community, and systematically developing AI-based automated workflows can accelerate the design of effective therapeutic compounds with a target set of properties [You et al. (2018); Li et al. (2018); Huang et al. (2020); Zhou et al. (2020); Horwood & Noutahi (2020); Khemchandani et al. (2020); Chenthamarashan et al. (2020); Hu et al. (2020); Sivaraman et al. (2020); Joshi & Kumar (2021)]. The potential chemical space is composed of over 10^{60} compounds that ideally need to be tested for the therapeutic design [Reymond (2015)] and discovery before identifying a lead for a given target protein of SARS-CoV-2. However, candidates with suitable activity against specific proteins only narrows the search space to $10^4 - 10^5$ compounds. There is a need for rapid development of machine-learning (ML) workflows that serve as a search-and-screen process of this reduced chemical space. Candidate molecules generated by ML models are passed to downstream verification via virtual high-throughput drug-protein binding techniques, chemical synthesis, biochemical and biophysical assay, and finally clinical trials [Batra et al. (2020); Smith & Smith (2020)]. In the current work, we develop, evaluate, and compare two leading graph-generative approaches to design candidate therapeutic compounds targeting SARS-CoV-2 protease inhibition (Figure 1). This systematic comparison of two leading approaches distinguishes our work in the context of the broad effort undertaken by the research community in its quest for novel therapeutics (see section 2 for a thorough review).

First, we train a junction-tree variational autoencoder (JT-VAE) [Jin et al. (2018)] to generate molecules with similar structures to a curated dataset of known drugs with anti-SARS activity. The trained JT-VAE is then used in Bayesian optimization to generate novel candidates with targeted properties. We next examine graph-based deep reinforcement learning (DQN) [Zhou et al. (2019)] to generate candidates that are not constrained by their structural proximity to known anti-SARS

¹<https://github.com/exalearn/covid-drug-design>

compounds. To benchmark the two graph generative models, we use the same set of property scoring functions as optimization targets.

Contributions. The goal of our study is to perform multi-objective optimization to generate molecular candidates by considering important bioactivity properties along with druglikeness and synthetic accessibility. We focus on pIC_{50} (the inverse log of the half maximal inhibitory concentration, IC_{50}), which is an experimentally measured property that captures the potency of a therapeutic candidate towards a protease target, where higher values indicate exponentially more potent inhibitors. Notably, pIC_{50} cannot be accurately modeled through *ab initio* methods [Bag & Ghorai (2016); Sebaugh (2011)]. Our ML workflow for the high-throughput generation of therapeutic candidates with anti-SARS activity contains the following key components: a surrogate model for pIC_{50} prediction, candidate generation via JT-VAE and DQN models, and validation of top-ranking candidates against a Drug Target Binding Affinity (DBTA) classifier [Huang et al. (2020)] to assess their potential activity against SARS-CoV-2. Alongside this paper, we release the ML workflow and a curated dataset of drug molecules with anti-SARS activity for use by the community.

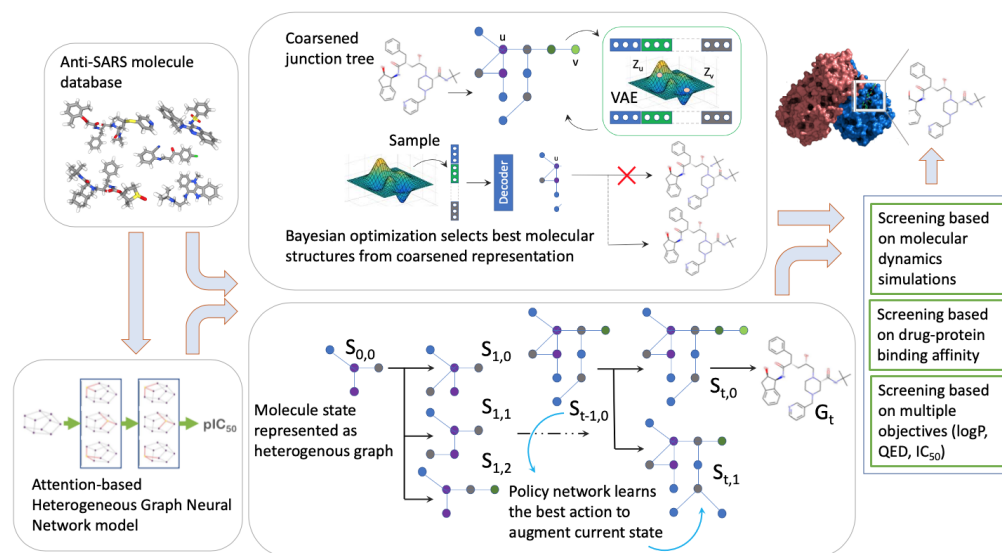


Figure 1: Depiction of workflow developed for this work. The anti-SARS database is used to train a MPNN to predict pIC_{50} and the JT-VAE model. The trained MPNN is used as the scoring function in both JT-VAE (top row) and DQN-based molecular generation (bottom). Candidate molecules are screened by pIC_{50} (>8) and validated by a Drug Target Binding Affinity classifier.

2 BACKGROUND AND RELATED WORK

Biological factors in drug design. A variety of properties can be considered for lead optimization [Copeland et al. (2006)]. Here we give an explanation of the various properties and their ideal values. Lipinski’s Rule of Five is the standard starting place when considering properties of potential oral drug compounds [Lipinski (2004)]. The Rule of Five identifies that molecules with no more than 5 hydrogen bond donors, no more than 10 hydrogen bond acceptors, a molecular mass under 500 Da, and an octanol–water partition coefficient ($\log P$) below 5 are more likely to show pharmacological activity. $\log P$ is a measure of lipophilicity, which provides an understanding of the behavior of a drug in the body. More recent analysis has shown that $\log P$ values between 1 and 3 may be more appropriate considering the effect of $\log P$ on absorption, distribution, metabolism, elimination, and toxicology (ADMET) properties [Waring (2010)]. Though oral bioavailability is an important factor, a sole focus on $\log P$ has the potential to screen out otherwise useful compounds [Zhang & Wilkinson (2007)]. QED has been proposed as a more holistic druglikeness metric [Bickerton et al. (2012)], which from 0 (low) to 1 (high). However, druglikeness does not indicate the activity or effectiveness of a drug towards a specific target. The half maximal inhibitory concentration (IC_{50}) provides a quantitative measure of the potency of a compound to inhibit a specific biological process. IC_{50}

is obtained by measurement, and no universal *ab initio* method of computing its value exists. A number of methods have been developed to approximate IC_{50} , many based on QSPR and recently some based on machine learning [Patankar & Jurs (2000); Bag & Ghorai (2016); Armutlu et al. (2008); Bjerrum (2017)]. Similarity to known drugs is also an important factor in drug discovery [Maggiore et al. (2014); Kumar & Zhang (2018)], as is the ability to synthesize the molecule, which can be estimated by the synthetic accessibility (SA) score – from 1 (easy) to 10 (difficult) [Ertl & Schuffenhauer (2009)].

COVID-19-focused efforts. In response to the COVID-19 pandemic, researchers have pushed to identify marketed drugs that can be repurposed for SARS-CoV-2 treatment [Zhou et al. (2020); Smith & Smith (2020); Gordon et al. (2020); Zhavoronkov et al. (2020); Skwark et al. (2020); Mittal et al. (2020); Chang (2020); Chen et al. (2020); Wu et al. (2020)]. Most of these the COVID-19-focused efforts relied on molecular simulations and docking studies to identify known drugs that interacted with 3-chymotrypsin-like main protease (3CLpro), also called the main protease (Mpro), of SARS-Cov-2. Machine learning (ML) approaches to drug repurposing were also reported. For instance, Batra et al. (2020) applied an ML-based screening approach to find known compounds with binding affinity to either the isolated SARS-CoV-2 S-protein at its host receptor region or to the S-protein-human ACE2 interface complex. Huang et al. (2020) developed DeepPurpose, a deep learning toolkit for drug repurposing, with the goal of recommending known candidates with high binding affinity to target amino acid sequences.

An array of *de novo* design approaches specifically targeting anti-SARS-CoV-2 drugs have also been reported. Hu et al. (2020) trained the GPT2 transformer-based model to generate SMILES strings of potential drug candidates follow the Lipinski’s ”Rule of Five”, optimizing towards 3CLpro–target interaction. Chenthamarakshan et al. (2020) developed the generative modeling framework CogMol which relies upon a variational autoencoder sampled by Controlled Latent attribute Space Sampling (CLaSS) to design candidates optimized toward binding affinity. Born et al. (2020) amended their PaccMann RL approach, which coupled reinforcement learning with variational autoencoders with reinforcement learning, to generate molecules optimized towards binding affinity and pharmacological toxicity predictors. Zhavoronkov et al. (2020) applied a similar approach of using reinforcement learning to optimize a generative model, but examined genetic algorithms, language models, generative adversarial networks as well as autoencoders. Tang et al. (2020) applied a deep Q learning approach with a structure-based optimization policy that targets QED, the presence of specific fragments, and the presence of pharmacophores. In a novel direction, Skwark et al. (2020) focused on the design of proteins that bind to SARS-CoV-2 in place of ACE2 using reinforcement learning.

Graph generative methods for molecule design. Heterogeneous graphs provide a natural representation for small-molecule organic compounds, with nodes representing atoms in the molecular structure and edges representing bonds between the atoms [Weininger (1990)]. This approach motivated the exploration of graph-generative models, such as graph convolutional policy networks [You et al. (2018)], variational autoencoders [Jin et al. (2018); Samanta et al. (2019); Jin et al. (2020)], and variants of deep reinforcement learning [Zhou et al. (2019); Ståhl et al. (2019)], for the target-driven optimization of drug molecules. In the current work, we develop, evaluate, and compare two leading graph-generative approaches for candidate therapeutic compounds, variational autoencoders and reinforcement learning, specifically targeting molecules able to inhibit the function of SARS-CoV-2 proteins. Our focus is to understand the distinction between the variational autoencoder and reinforcement learning approaches in terms of similarity versus novelty in the generated molecules, as well as the optimal approach to developing a wholistic scoring function that accounts for both druglikeness and SARS-Cov-2 protease inhibition.

3 METHOD

3.1 SURROGATE MODEL FOR IC_{50} PREDICTION

We trained a message-passing neural network (MPNN) [Scarselli et al. (2008); Gilmer et al. (2017)] to predict pIC_{50} (the inverse log of IC_{50}) for a given molecular structure. Following the formalism of Gilmer et al. (2017), our network is composed of message, update, and readout operations (eqns. 1-3) and our choices for these functions are based on networks developed by John et al. (2019) for polymer property prediction.

The original state of each atom (h_v) and bond (α_{vw}) in our molecule (G) is a 256-length vector with values defined by an embedding table based on the atomic number and bond type (e.g., single, double, aromatic). The states of these atoms are modified by eight successive “message” layers. Each message layer uses a two-layer multi-layer perceptron (MLP) with sigmoid activations to compute a message that uses the state of an atom (h_v), the state of the neighboring atom (h_w) and the bond which joins them (α_{vw}). The atom and bond states are updated according to the following equations:

$$m_v^{t+1} = \sum_{w \in \text{Neighbors}(v)} M_t(h_v^t, h_w^t, \alpha_{vw}^t) \quad (1)$$

$$h_v^{t+1} = h_v^t + m_v^{t+1} \quad (2)$$

$$\alpha_{vw}^{t+1} = \alpha_{vw}^t + M_t(h_v^t, h_w^t, \alpha_{vw}^t) \quad (3)$$

The atom states output from the last layer (h_v^T) are used to predict the pIC₅₀ of the molecule using a “readout” function (R):

$$\hat{y} = R(h_v^T | v \in G) \quad (4)$$

We examined several variants of the readout function in our study. We tested both “molecular fingerprints,” where the states of each node are combined *before* using a multi-layer perceptron (MLP) to reduce to compute pIC₅₀, and an “atomic contribution,” where we combine the nodes *after* MLP to compute a per-node contribution to pIC₅₀. We experimented with the use of five different functions to reduce the atomic state/contributions to a single value for each graph: summation, mean, maximum, softmax, and attention. The attention functions are created by learning an attention map by passing the node states through a MPNN. We tested all combinations of “molecular fingerprint” vs. “atomic contribution” and the five readout functions, for a total of 10 networks, training each on network the same 90% of our pIC₅₀ dataset and comparing its performance on the withheld 10% of the data. We used an MPNN that uses attention maps to reduce contributions from each atom to a single pIC₅₀ of a molecule in all subsequent experiments.

Junction-Tree Variational Autoencoder. We used a junction-tree variational autoencoder (JT-VAE) [Jin et al. (2018)] to generate molecules with high proximity to anti-SARS therapeutic molecules. The autoencoder generates novel molecular graphs by laying a tree-structured scaffold over substructures in the molecule, which are then combined into a valid molecule using a MPNN. JT-VAE allows for the expansion of a molecular graph through the addition of subgraphs, or “vocabulary” of valid components, derived from the training set (Fig. 1). The subgraphs are used to encode a molecule into its vector representation and decode latent vectors into valid molecular graphs. The use of subgraphs maintains chemical validity at each step, while also incorporating chemical units common to the training set. Chemical graphs generated from the vocabulary are structurally similar to those in the training set, which is a benefit when attempting to design molecules with similar properties to known drugs. We omit the details here for brevity and point the reader to Jin et al. (2018) for details on JT construction and Kusner et al. (2017) for Bayesian optimization.

Deep Reinforcement Learning. We follow the Q-learning approach of Zhou et al. for our deep reinforcement learning approach [Zhou et al. (2019)]. In this setting, we cast the molecule generation problem as a Markov Decision Process (MDP) [Mnih et al. (2013)] to learn a policy network π that determines the best sequence of actions to transform an initial molecular graph to a larger graph with desirable properties in a step-by-step fashion (Fig. 1). At each step, we enumerate all possible actions and select those which produce valid molecules (e.g., respect valency rules). Next, we train a MLP to predict the *value* of a certain action by passing the Morgan fingerprints [Rogers & Hahn (2010)] as input. The MLP approximates the value of an action computed using the Bellman equation, where the score of a state and the maximum score of the subsequent state is multiplied by a decay factor. As established with other Deep Q-Learning approaches, the addition of the value of the next state increases the value of moves which will lead to higher future rewards.

SARS-CoV-2-specific Scoring Functions. The scoring functions described in this section are used for both Bayesian optimization in the JT-VAE approach and reward computation for the deep reinforcement-learning approach.

$$\log P^P(m) = \log P(m) - SA(m) - cycle(m) \quad (5)$$

$$QED^P(m) = QED(m) - SA(m) - cycle(m) \quad (6)$$

$$pIC_{50}(m) = MPNN(m) \quad (7)$$

$$pIC_{50} + QED^P(m) = MPNN(m) + QED(m) - SA(m) - cycle(m) \quad (8)$$

Following Jin et al. (2018), we first compute a score that penalizes the partition coefficient between octanol and water ($\log P$) by the synthetic accessibility (SA) score (in which higher SA values are discouraged) and the number of cycles with more than 6 atoms (eqn. 5). Considering that the Quantitative Estimate of Druglikeness (QED) is a more comprehensive heuristic than $\log P$, we also use a similar scoring function composed of QED penalized by the SA score and number of long cycles (eqn. 6). We then examine the utility of a SARS-specific scoring function based on the pIC_{50} predicted by our MPNN (eqn. 7). Finally, we examine a multi-objective scoring function that includes both pIC_{50} and penalized QED (eqn. 8).

4 EXPERIMENTAL ANALYSIS

The focus of our study is two fold: **1)** Perform multi-objective optimization for generating therapeutic candidates with activity towards protease target of SARS-CoV-2, **2)** Build novel yet unique compounds by optimizing the scoring function including synthesizability score. To this end, we are poised to utilize best possible graph based ML models.

Dataset Preparation. We assembled and curated a protease dataset containing molecules active against various protease in enzymatic assays from experimental pharmacology databases, such as ChEMBL, BindingDB, and ToxCat [Bento et al. (2014)]. The dataset was filtered based on pIC_{50} activity and potency. Molecules larger than 1,000 Dalton were removed along with non-drug like molecules containing metals and polypeptides. In some instances, duplicate entries were generated from multiple experimental studies, in which case we used the mean pIC_{50} . The resulting dataset contains SMILES strings and experimental pIC_{50} values of 6,545 unique molecules.

Quantitative Characterization. We computed the metrics included in our scoring functions for each molecule in the dataset. The highest pIC_{50} is 10.89, while the lowest is 1.22. The most common pIC_{50} is 4.0, which is shared by 320 structures, and the vast majority of structures (91.5%) have pIC_{50} values greater than 4.0. $\log P$ values range from -10.36 to 16.65 in a near Gaussian distribution with a mean of 3.70; 77.0% of all structures meet the requirement of Lipinski’s Rule of 5 that $\log P$ be no greater than 5. QED values range from 0.01 to 0.94, while SA values range from 1.35 to 8.24, with 73.8% being below 4. We computed the Tanimoto similarity [Bajusz et al. (2015)] for all pairs of compounds to gain insight into the structural diversity of molecules in our dataset (Fig. S2). We observed that structures tend to become more similar to their neighbors as pIC_{50} increases, indicating that compounds with high pIC_{50} values tend to be structurally similar, supporting the consideration of molecular similarity during drug design and discovery.

Training pIC_{50} Surrogate Model. We found that limiting the pIC_{50} prediction to contributions from only a few specific atoms in the molecule improved performance. This improved performance can be explained by a physical mechanism. The presence or absence of a specific chemical pattern in a molecular structure (e.g., functional groups, substituents, aromatic rings) controls binding of the molecule to a certain portion of a target protein, while the other atoms in the molecule play a role in determining whether the molecule will stay affixed at the target site. We hypothesize that such contributions can have a critical effect in tuning the binding affinity, which influences pIC_{50} .

JT-VAE Setup. We trained the JT-VAE on our dataset for 8,300 iterations with the following hyperparameters: hidden state dimension of 450, latent code dimension of 56, and graph message passing depth of 3. To optimize towards the specified scoring functions, we trained a sparse Gaussian process (SGP) to predict a score given the latent representation learned by the JT-VAE and then performed 10 iterations of batched Bayesian optimization (sampling = 50) based on the expected improvement.

DQN Setup. Our DQN approach constructs molecules atom-by-atom and bond-by-bond in a step-wise manner. Each episode starts with either a single atom or a pre-seeded molecule and the RL

Table 1: Properties of the top-3 molecules generated using the specified scoring function.

Scoring Function	pIC ₅₀			QED			logP			SA Score		
	1st	2nd	3rd	1st	2nd	3rd	1st	2nd	3rd	1st	2nd	3rd
logP ^P (JT-VAE)	4.93	4.57	4.60	0.45	0.78	0.71	4.05	4.27	4.20	1.70	2.08	2.08
logP ^P (DQN)	6.10	8.17	4.98	0.04	0.07	0.11	13.86	12.64	12.52	3.59	2.97	2.93
QED ^P (JT-VAE)	4.15	4.23	4.71	0.91	0.84	0.91	3.72	2.19	2.20	1.80	1.71	2.25
QED ^P (DQN)	6.80	6.80	6.80	0.77	0.77	0.77	3.46	3.46	3.46	2.05	2.05	2.05
pIC ₅₀ (JT-VAE)	10.22	10.07	10.07	0.15	0.12	0.12	4.86	3.80	3.06	4.52	4.94	4.98
pIC ₅₀ (DQN)	10.57	10.56	10.39	0.09	0.11	0.41	1.51	2.03	-0.44	6.90	6.71	5.57
pIC ₅₀ +QED ^P (JT-VAE)	8.58	5.98	8.18	0.83	0.87	0.70	4.02	3.37	3.82	1.93	1.74	1.50
pIC ₅₀ +QED ^P (DQN)	10.27	10.27	10.27	0.80	0.80	0.80	3.02	3.02	3.02	2.90	2.90	2.90


agent is allowed up to 40 steps to construct an optimized molecule. A move is determined either randomly or by the Q-function model. We update the model to predict the move after each step in each episode and, as this model improves during training, we gradually reduce the probability that a random move is chosen over the model prediction. If we start with a single atom approach the DQN finds tens of thousands of candidate molecules that needed to be optimized. Starting from a pre-seeded molecule led to an increased number of molecules generated with high pIC₅₀, while starting from a single atom produced molecules with higher QED, as shown in Fig. S3.

Comparative Effect of the Scoring Function. Table 1 shows the top-3 molecules generated by the two generative models with each scoring function. DQN always outperforms JT-VAE in finding a molecule with a superior value of the scoring function being optimized. The performance disparity is particularly apparent when optimizing for logP: the maximum logP from DQN is 12.6 compared to only 4.1 for JT-VAE. We attribute the difference in performance to JT-VAE implicitly sampling from a distribution of drug-like molecules which typically have logP values between -0.4 and 5.6, while DQN has no such constraints. However, the JT-VAE-generated molecules have consistently better QED and SA scores even when those values are not explicitly optimized for. The molecules in our dataset were experimentally synthesized, leading the SA scores of these molecules to be quite low. Likewise, because they are drug molecules, their QED are high. This leads JT-VAE to generate similar molecules with high QED and low SA scores. On the other hand, the RL agent uses no information about the space of experimentally studied drug candidates during its training process and, accordingly, finds molecules far outside that region.

Overall, we find different purposes for JT-VAE and RL-based molecular optimization. JT-VAE implicitly uses the distribution of molecules in its training set to bias towards realistic molecules, albeit at the expense of highly novel candidates. The RL-based approach lacks such constraints and, for better or worse, optimizes without implicitly regarding synthesizability or other characteristics not explicitly encoded in the scoring function, leading to many more novel candidates which are otherwise impossible to optimize toward.

Qualitative Analysis. We observed an interesting structural trend in the JT-VAE-generated molecules optimized towards pIC₅₀. Fig. 2 shows the structures of generated molecules with pIC₅₀ > 8 and the anti-HIV drug Indinavir. A common backbone is shared between Indinavir and the top-6 predictions. The Tanimoto similarity of these six generated molecules against Indinavir range from 0.65 to 0.91. Indinavir has been proposed as a drug candidate to treat SARS-CoV-2 due to favorable binding affinity to the 3-chymotrypsin-like main protease (3CLpro) that is pivotal for the replication of SARS-CoV-2 [Yoshino et al. (2020); Chang (2020); Gyebi et al. (2020); Harrison (2020)]. Notably, three of the generated molecules have a higher predicted pIC₅₀ than Indinavir.

In silico Drug-Target Binding Affinity (DTBA) methods offer a way to evaluate drug-target interactions [He et al. (2017)]. We employ a ML-based DTBA model to validate the interaction of molecules generated by JT-VAE against 3CLpro [Chen et al. (2020)]. We trained a DBTA binary classification model using the extended connectivity fingerprint [Rogers & Hahn (2010)] encoding for the drug molecule and the target protease sequence encoding from the DeepPurpose toolkit [Huang et al. (2020)]. The DBTA model classified 4 of the top-11 molecules (including the top-2 in Fig. 2) with >50% probability of interaction with 3CLpro.

Figure 2: Top-6 molecules generated by the JT-VAE method optimized towards pIC₅₀.

5 CONCLUSIONS

To possibly explore the novelty and uniqueness of generated molecules, we compared two graph generative models, JT-VAE and DQN, for the task of generating small-molecule therapeutic candidates with activity against 3CLpro target of SARS-CoV-2. DQN always outperformed JT-VAE in finding molecules with a superior value of the scoring function being optimized. However, JT-VAE generated molecules that were more structurally similar to those in the dataset due to substructure representation, which produced a lower SA score and logP < 5.

REFERENCES

- P. Armutlu, M. E. Ozdemir, F. Uney-Yuksektepe, I. H. Kavakli, and M. Turkey. Classification of drug molecules considering their ic50 values using mixed-integer linear programming based hyper-boxes method. *BMC Bioinformatics*, 9:411, 2008. ISSN 1471-2105. doi: 10.1186/1471-2105-9-411.
- A. Bag and P. K. Ghorai. Development of quantum chemical method to calculate half maximal inhibitory concentration (ic50). *Mol Inform*, 35(5):199–206, 2016. ISSN 1868-1743. doi: 10.1002/minf.201501004.
- Dávid Bajusz, Anita Rácz, and Károly Héberger. Why is tanimoto index an appropriate choice for fingerprint-based similarity calculations? *Journal of cheminformatics*, 7(1):20, 2015.
- Rohit Batra, Henry Chan, Ganesh Kamath, Rampi Ramprasad, Mathew J. Cherukara, and Subramanian K.R.S. Sankaranarayanan. Screening of therapeutic agents for covid-19 using machine learning and ensemble docking studies. *The Journal of Physical Chemistry Letters*, 11(17):7058–7065, 2020. doi: 10.1021/acs.jpcelett.0c02278. PMID: 32787328.
- A Patrícia Bento, Anna Gaulton, Anne Hersey, Louisa J Bellis, Jon Chambers, Mark Davies, Felix A Krüger, Yvonne Light, Lora Mak, Shaun McGlinchey, et al. The chembl bioactivity database: an update. *Nucleic Acids Research*, 42(D1):D1083–D1090, 2014.
- G. Richard Bickerton, Gaia V. Paolini, Jérémy Besnard, Sorel Muresan, and Andrew L. Hopkins. Quantifying the chemical beauty of drugs. *Nature chemistry*, 4(2):90–98, 2012. ISSN 1755-4349 1755-4330. doi: 10.1038/nchem.1243. URL <https://pubmed.ncbi.nlm.nih.gov/22270643><https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3524573/>.
- Esben Jannik Bjerrum. SMILES enumeration as data augmentation for neural network modeling of molecules. *CoRR*, abs/1703.07076, 2017. URL <http://arxiv.org/abs/1703.07076>.
- Jannis Born, Matteo Manica, Joris Cadow, Greta Markert, Nil Adell Mill, Modestas Filipavicius, and María Rodríguez Martínez. Paccmannrl on sars-cov-2: Designing antiviral candidates with conditional generative models. *arXiv e-prints*, pp. arXiv–2005, 2020.

- Y. et al. Chang. Potential therapeutic agents for covid-19 based on the analysis of protease and rna polymerase docking. *Preprints*, pp. 2020020242, 2020. doi: 10.20944/preprints202002.0242.v1.
- Yu Wai Chen, Chin-Pang Benu Yiu, and Kwok-Yin Wong. Prediction of the sars-cov-2 (2019-ncov) 3c-like protease (3cl pro) structure: virtual screening reveals velpatasvir, ledipasvir, and other drug repurposing candidates. *F1000Research*, 9, 2020.
- Vijil Chenthamarakshan, Payel Das, Inkit Padhi, Hendrik Strobelt, Kar Wai Lim, Ben Hoover, Samuel C Hoffman, and Aleksandra Mojsilovic. Target-specific and selective drug design for covid-19 using deep generative models. *arXiv preprint arXiv:2004.01215*, 2020.
- Robert A Copeland, David L Pompliano, and Thomas D Meek. Drug–target residence time and its implications for lead optimization. *Nature Reviews Drug Discovery*, 5(9):730–739, 2006.
- Peter Ertl and Ansgar Schuffenhauer. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of Cheminformatics*, 1(1):8, 2009. ISSN 1758-2946. doi: 10.1186/1758-2946-1-8. URL <https://doi.org/10.1186/1758-2946-1-8>.
- Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *34th International Conference on Machine Learning–Volume 70*, pp. 1263–1272. JMLR. org, 2017.
- David E. Gordon, Gwendolyn M. Jang, Mehdi Bouhaddou, Jiewei Xu, Kirsten Obernier, Kris M. White, Matthew J. O’Meara, Veronica V. Rezelj, Jeffrey Z. Guo, Danielle L. Swaney, Tia A. Tummino, Ruth Huettenhain, Robyn M. Kaake, Alicia L. Richards, Beril Tutuncuoglu, Helene Foussard, Jyoti Batra, Kelsey Haas, Maya Modak, Minkyu Kim, Paige Haas, Benjamin J. Polacco, Hannes Braberg, Jacqueline M. Fabius, Manon Eckhardt, Margaret Soucheray, Melanie J. Bennett, Merve Cakir, Michael J. McGregor, Qiongyu Li, Bjoern Meyer, Ferdinand Roesch, Thomas Vallet, Alice Mac Kain, Lisa Miorin, Elena Moreno, Zun Zar Chi Naing, Yuan Zhou, Shiming Peng, Ying Shi, Ziyang Zhang, Wenqi Shen, Ilsa T. Kirby, James E. Melnyk, John S. Chorba, Kevin Lou, Shizhong A. Dai, Inigo Barrio-Hernandez, Danish Memon, Claudia Hernandez-Armenta, Jiankun Lyu, Christopher J. P. Mathy, Tina Perica, Kala B. Pilla, Sai J. Ganesan, Daniel J. Saltzberg, Ramachandran Rakesh, Xi Liu, Sara B. Rosenthal, Lorenzo Calviello, Sri-vats Venkataramanan, Jose Liboy-Lugo, Yizhu Lin, Xi-Ping Huang, YongFeng Liu, Stephanie A. Wankowicz, Markus Bohn, Maliheh Safari, Fatima S. Ugur, Cassandra Koh, Nastaran Sadat Savar, Quang Dinh Tran, Djoshkun Shengjuler, Sabrina J. Fletcher, Michael C. O’Neal, Yiming Cai, Jason C. J. Chang, David J. Broadhurst, Saker Klippsten, Phillip P. Sharp, Nicole A. Wenzell, Duygu Kuzuoglu, Hao-Yuan Wang, Raphael Trenker, Janet M. Young, Devin A. Cavero, Joseph Hiatt, Theodore L. Roth, Ujjwal Rathore, Advait Subramanian, Julia Noack, Mathieu Hubert, Robert M. Stroud, Alan D. Frankel, Oren S. Rosenberg, Kliment A. Verba, David A. Agard, Melanie Ott, Michael Emerman, Natalia Jura, et al. A sars-cov-2 protein interaction map reveals targets for drug repurposing. *Nature*, 2020. ISSN 1476-4687. doi: 10.1038/s41586-020-2286-9. URL <https://doi.org/10.1038/s41586-020-2286-9>.
- Gideon A. Gyebi, Olalekan B. Ogunro, Adegbenro P. Adegunloye, Oludare M. Ogunyemi, and Saheed O. Afolabi. Potential inhibitors of coronavirus 3-chymotrypsin-like protease (3cl(pro)): an in silico screening of alkaloids and terpenoids from african medicinal plants. *Journal of Biomolecular Structure and Dynamics*, pp. 1–13, 2020. ISSN 1538-0254 0739-1102. doi: 10.1080/07391102.2020.1764868. URL <https://pubmed.ncbi.nlm.nih.gov/32367767https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7256353/>.
- C Harrison. Coronavirus puts drug repurposing on the fast track. *Nature Biotechnology*, 38(4):379, 2020.
- Tong He, Marten Heidemeyer, Fuqiang Ban, Artem Cherkasov, and Martin Ester. Simboost: a read-across approach for predicting drug–target binding affinities using gradient boosting machines. *Journal of Cheminformatics*, 9(1):1–14, 2017.
- Julien Horwood and Emmanuel Noutahi. Molecular design in synthetically accessible chemical space via deep reinforcement learning. *arXiv preprint arXiv:2004.14308*, 2020.

- Fan Hu, Dongqi Wang, Yishen Hu, Jiaxin Jiang, and Peng Yin. Generating novel compounds targeting sars-cov-2 main protease based on imbalanced dataset. In *2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 432–436. IEEE, 2020.
- Kexin Huang, Tianfan Fu, Cao Xiao, Lucas Glass, and Jimeng Sun. Deeppurpose: a deep learning based drug repurposing toolkit. *arXiv preprint arXiv:2004.08919*, 2020.
- Wengong Jin, Regina Barzilay, and Tommi S. Jaakkola. Junction tree variational autoencoder for molecular graph generation. *CoRR*, abs/1802.04364, 2018. URL <http://arxiv.org/abs/1802.04364>.
- Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Multi-Objective Molecule Generation using Interpretable Substructures. *arXiv e-prints*, art. arXiv:2002.03244, February 2020.
- Peter C. St. John, Caleb Phillips, Travis W. Kemper, A. Nolan Wilson, Yanfei Guan, Michael F. Crowley, Mark R. Nimlos, and Ross E. Larsen. Message-passing neural networks for high-throughput polymer screening. *Journal of Chemical Physics*, 150(23):234111, June 2019. doi: 10.1063/1.5099132. URL <https://doi.org/10.1063/1.5099132>.
- Rajendra P Joshi and Neeraj Kumar. Artificial intelligence based autonomous molecular design for medical therapeutic: A perspective. *arXiv preprint arXiv:2102.06045*, 2021.
- Yash Khemchandani, Steve O’Hagan, Soumitra Samanta, Neil Swainston, Timothy J Roberts, Danushka Bollegala, and Douglas B Kell. Deepgraphmol, a multi-objective, computational strategy for generating molecules with desirable properties: a graph convolution and reinforcement learning approach. *bioRxiv*, 2020.
- Ashutosh Kumar and Kam Y. J. Zhang. Advances in the development of shape similarity methods and their application in drug discovery. *Frontiers in chemistry*, 6:315–315, 2018. ISSN 2296-2646. doi: 10.3389/fchem.2018.00315. URL <https://pubmed.ncbi.nlm.nih.gov/30090808https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6068280/>.
- Matt J. Kusner, Brooks Paige, and José Miguel Hernández-Lobato. Grammar variational autoencoder, 2017.
- Yibo Li, Liangren Zhang, and Zhenming Liu. Multi-objective de novo drug design with conditional graph generative model. *Journal of Cheminformatics*, 10(1):33, 2018. ISSN 1758-2946. doi: 10.1186/s13321-018-0287-6. URL <https://doi.org/10.1186/s13321-018-0287-6>.
- Christopher A Lipinski. Lead-and drug-like compounds: the rule-of-five revolution. *Drug Discovery Today: Technologies*, 1(4):337–341, 2004.
- Gerald Maggiora, Martin Vogt, Dagmar Stumpfe, and Jurgen Bajorath. Molecular similarity in medicinal chemistry: miniperspective. *Journal of medicinal chemistry*, 57(8):3186–3204, 2014.
- Lovika Mittal, Anita Kumari, Mitul Srivastava, Mrityunjay Singh, and Shailendra Asthana. Identification of potential molecules against covid-19 main protease through structure-guided virtual screening approach. *Journal of Biomolecular Structure and Dynamics*, pp. 1–19, 2020.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- S. J. Patankar and P. C. Jurs. Prediction of ic50 values for acat inhibitors from molecular structure. *Journal of Chemical Information and Computer Sciences*, 40(3):706–723, 2000. ISSN 0095-2338. doi: 10.1021/ci990125r. URL <https://doi.org/10.1021/ci990125r>.
- Jean-Louis Reymond. The chemical space project. *Accounts of Chemical Research*, 48(3):722–730, 2015.
- David Rogers and Mathew Hahn. Extended-connectivity fingerprints. *Journal of Chemical Information and Modeling*, 50(5):742–754, 2010.

- Bidisha Samanta, DE Abir, Gourhari Jana, Pratim Kumar Chattaraj, Niloy Ganguly, and Manuel Gomez Rodriguez. Nevae: A deep generative model for molecular graphs. In *AAAI Conference on Artificial Intelligence*, volume 33, pp. 1110–1117, 2019.
- Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008.
- JL Sebaugh. Guidelines for accurate ec50/ic50 estimation. *Pharmaceutical statistics*, 10(2):128–134, 2011.
- Ganesh Sivaraman, Nicholas Jackson, Benjamin Sanchez-Lengeling, Alvaro Vasquez-Mayagoitia, Alan Aspuru-Guzik, Venkatram Vishwanath, and Juan de Pablo. A machine learning workflow for molecular analysis: application to melting points. *Machine Learning: Science and Technology*, 2020.
- Marcin J Skwark, Nicolás López Carranza, Thomas Pierrot, Joe Phillips, Slim Said, Alexandre Laterre, Amine Kerkeni, Uğur Şahin, and Karim Beguir. Designing a prospective covid-19 therapeutic with reinforcement learning. *arXiv preprint arXiv:2012.01736*, 2020.
- Micholas Smith and Jeremy C. Smith. Repurposing therapeutics for covid-19: Supercomputer-based docking to the sars-cov-2 viral spike protein and viral spike protein-human ace2 interface, Feb 2020. URL https://chemrxiv.org/articles/Repurposing_Therapeutics_for_the_Wuhan_Coronavirus_nCov-2019_Supercomputer-Based_Docking_to_the_Viral_S_Protein_and_Human_ACE2_Interface/11871402/3.
- Niclas Ståhl, Göran Falkman, Alexander Karlsson, Gunnar Mathiason, and Jonas Bostrom. Deep reinforcement learning for multiparameter optimization in de novo drug design. *Journal of Chemical Information and Modeling*, 59(7):3166–3176, 2019.
- Bowen Tang, Fengming He, Dongpeng Liu, Meijuan Fang, Zhen Wu, and Dong Xu. Ai-aided design of novel targeted covalent inhibitors against sars-cov-2. *BioRxiv*, 2020.
- Michael J Waring. Lipophilicity in drug discovery. *Expert Opinion on Drug Discovery*, 5(3): 235–248, 2010. doi: 10.1517/17460441003605098. URL <https://doi.org/10.1517/17460441003605098>. PMID: 22823020.
- David Weininger. Smiles. 3. depict. graphical depiction of chemical structures. *Journal of Chemical Information and Computer Sciences*, 30(3):237–243, 1990.
- Canrong Wu, Yang Liu, Yueying Yang, Peng Zhang, Wu Zhong, Yali Wang, Qiqi Wang, Yang Xu, Mingxue Li, Xingzhou Li, et al. Analysis of therapeutic targets for sars-cov-2 and discovery of potential drugs by computational methods. *Acta Pharmaceutica Sinica B*, 10(5):766–788, 2020.
- Ryunosuke Yoshino, Nobuaki Yasuo, and Masakazu Sekijima. Identification of key interactions between sars-cov-2 main protease and inhibitor drug candidates. *Scientific reports*, 10(1):1–8, 2020.
- Jiaxuan You, Bowen Liu, Zhitao Ying, Vijay Pande, and Jure Leskovec. Graph convolutional policy network for goal-directed molecular graph generation. In *Advances in Neural Information Processing Systems*, pp. 6410–6421, 2018.
- Ming-Qiang Zhang and Barrie Wilkinson. Drug discovery beyond the ‘rule-of-five’. *Current Opinion in Biotechnology*, 18(6):478 – 488, 2007. ISSN 0958-1669. doi: <https://doi.org/10.1016/j.copbio.2007.10.005>. URL <http://www.sciencedirect.com/science/article/pii/S0958166907001279>. Chemical biotechnology / Pharmaceutical biotechnology.
- Alex Zhavoronkov, Vladimir Aladinskiy, Alexander Zhebrak, Bogdan Zagribelnyy, Victor Terentiev, Dmitry S Bezrukov, Daniil Polykovskiy, Rim Shayakhmetov, Andrey Filimonov, Philipp Orekhov, et al. Potential covid-2019 3c-like protease inhibitors designed using generative deep learning approaches. *Insilico Medicine Hong Kong Ltd A*, 307:E1, 2020.

Yadi Zhou, Yuan Hou, Jiayu Shen, Yin Huang, William Martin, and Feixiong Cheng. Network-based drug repurposing for novel coronavirus 2019-ncov/sars-cov-2. *Cell Discovery*, 6(1):14, 2020. ISSN 2056-5968. doi: 10.1038/s41421-020-0153-3. URL <https://doi.org/10.1038/s41421-020-0153-3>.

Zhenpeng Zhou, Steven Kearnes, Li Li, Richard N. Zare, and Patrick Riley. Optimization of molecules via deep reinforcement learning. *Scientific Reports*, 9(1):10752, 2019. ISSN 2045-2322. doi: 10.1038/s41598-019-47148-x. URL <https://doi.org/10.1038/s41598-019-47148-x>.

SUPPLEMENTARY MATERIAL

A DATASET ANALYSIS

We computed the metrics included in our scoring functions for each molecule in the database. The range of values can be seen in Figure S1. The highest pIC_{50} is 10.89, while the lowest pIC_{50} is 1.22. The most common pIC_{50} is 4.0, which is shared by 320 structures, and the vast majority of structures (91.5%) have pIC_{50} values greater than 4.0. LogP values range from -10.36 to 16.65 in a near Gaussian distribution with a mean of 3.70; 77.0% of all structures meet the requirement of Lipinski’s Rule of 5 that logP be no greater than 5. QED values range from 0.01 to 0.94, while SA values range from 1.35 to 8.24, with 96.3% being below 5.

Figure S1 also shows the range of properties of molecules generated by our trained JT-VAE without optimization. In general, JT-VAE reproduced the distribution of logP, QED, pIC_{50} , and SA scores present in the database. The pIC_{50} values of the JT-VAE-generated molecules were obtained from our trained MPNN, discussed in Section 3.1.




Figure S1: We generated 1,000 molecules using the trained JT-VAE, of which 560 were unique. The figure shows the comparison of the QED, logP, SA score, and pIC_{50} of compounds in the database and those generated by the JT-VAE. The JT-VAE reproduced the range of values present in the database, minus outliers. The similar values indicate that the JT-VAE is able to reproduce the wide range of structures present in the database. The pIC_{50} values for generated molecules were estimated by our MPNN.

We also computed the Tanimoto similarity for all pairs of compounds to gain insight into the structural diversity of molecules in our database (Figure S2). The entries in the matrix were ordered in increasing pIC_{50} values. The similarity is represented by the color bar, with yellow representing low similarity (0) and red high similarity (1). We observe that structures tend to become more similar to their neighbors as pIC_{50} increases, indicating that compounds with high pIC_{50} values tend to be structurally similar, supporting the consideration of molecular similarity during drug discovery.




Figure S2: Heat map of the Tanimoto similarity between all compounds in the database. Entries are ordered in increasing pIC₅₀. The Tanimoto similarity is generally higher among structures with high pIC₅₀ (located towards bottom right of the matrix), indicating the importance of considering structural similarity in drug discovery.

B EFFECTS OF PRE-SEEDING RL

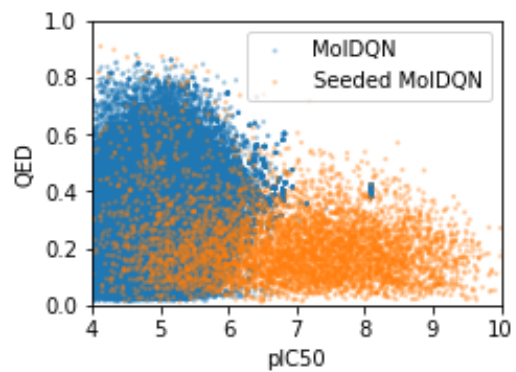


Figure S3: (QED vs pIC₅₀ for molecules generated by RL with starting from a single atom (MolDQN) or pre-seeded with a drug-like molecule (Seeded MolDQN) optimized towards the multi-objective scoring function in eqn. S8.