
PaccMann^{RL} on SARS-CoV-2: Designing antiviral candidates with conditional generative models

Jannis Born^{1,2*} Matteo Manica^{1*} Joris Cadow^{1*} Greta Markert^{1,2} Nil Adell Mill^{1,2} Modestas Filipavicius^{1,2} María Rodríguez Martínez¹

Abstract

With the fast development of COVID-19 into a global pandemic, scientists around the globe are desperately searching for effective antiviral therapeutic agents. Bridging systems biology and drug discovery, we propose a deep learning framework for conditional *de novo* design of antiviral candidate drugs tailored against given protein targets. First, we train a multimodal ligand–protein binding affinity model on predicting affinities of antiviral compounds to target proteins and couple this model with pharmacological toxicity predictors. Exploiting this multi-objective as a reward function of a conditional molecular generator (consisting of two VAEs), we showcase a framework that navigates the chemical space toward regions with more antiviral molecules. Specifically, we explore a challenging setting of generating ligands against unseen protein targets by performing a leave-one-out-cross-validation on 41 SARS-CoV-2-related target proteins. Using deep RL, it is demonstrated that in 35 out of 41 cases, the generation is biased towards sampling more binding ligands, with an average increase of 83% comparing to an unbiased VAE. We present a case-study on a potential Envelope-protein inhibitor and perform a synthetic accessibility assessment of the best generated molecules is performed that resembles a viable roadmap towards a rapid in-vitro evaluation of potential SARS-CoV-2 inhibitors.

1 Introduction

The Severe Acute Respiratory Syndrome (SARS) Coronavirus disease (COVID 2019) is an acute respiratory disease caused by novel coronavirus SARS-CoV-2 that, to date, has

infected millions and killed hundreds of thousands. Despite longstanding efforts into understanding the pathogenicity of coronaviruses (CoVs) (Drosten et al., 2003), there are no approved drugs against CoV, and new systematic approaches to identify effective antiviral agents are urgently needed. Current efforts are predominantly focused on drug repurposing strategies, with a handful of promising candidates, including remdesivir and hydroxychloroquine. Initial hopes are currently balked, remdesivir does not significantly reduce time to clinical improvement (Wang et al., 2020) and hydroxychloroquine was not found effective in a meta-study of human clinical trials (Shamshirian et al., 2020). Gordon et al. (2020) recently identified 69 promising compounds by measuring binding affinities of 26 out of the 29 SARS-CoV-2 proteins against human proteins.

With high uncertainty in the outcome of drug repurposing strategies, it is worth exploiting *de novo* drug discovery approaches against SARS-CoV-2. Drug discovery is challenging, with costs of up to 3 billion US\$ per new FDA-approved drug, an attrition rate of 99.99%, more than 10 years until market release and a search space of 10^{60} compounds (Scannell et al., 2012). However, the availability of high-throughput screenings of compound–protein interactions (CPI) has enabled deep learning to set new benchmarks for large-scale QSAR prediction models for predicting protein–drug binding affinity (Karimi et al., 2019). Deep learning has further been proven feasible of *in silico* design of molecules with desired chemical properties and shown potential to accelerate discovery of DDR1 inhibitors (Zhavronkov et al., 2019). A few studies used deep generative models to release libraries of (unsynthesized) candidates to target 3C-like protease, a main therapeutic target of SARS-CoV-2 (Zhavronkov et al., 2020; Tang et al., 2020) but both studies manually curated datasets to target 3C-like protease inhibitors. Here, we aim to bridge systems biology and drug discovery, using deep learning to explore target-driven drug design with conditional generative models. Our framework (see Figure 1) for conditional molecular design is conceptually inspired by our previous work, PaccMann^{RL} (Born et al., 2020), however note that here we focus on protein-driven instead of omics-profile-driven drug generation. Our framework can be trained to design compounds against any

*Equal contribution ¹IBM Research Europe, Switzerland. ²ETH Zurich, Switzerland. Correspondence to: J.B. <jab@zurich.ibm.com>, M.M <tte@zurich.ibm.com>, J.C <dow@zurich.ibm.com>.

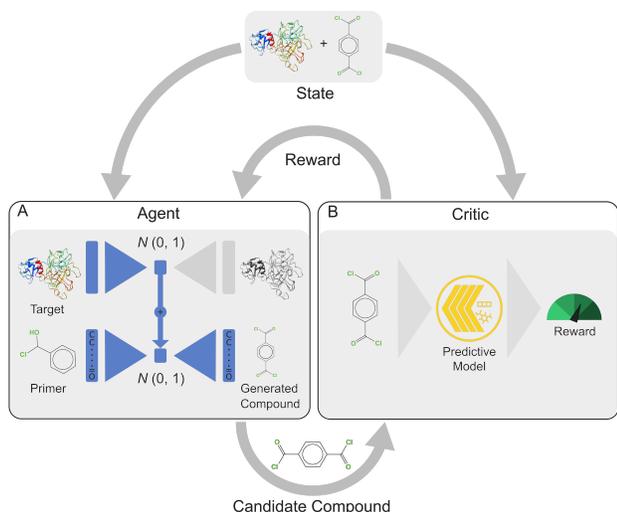


Figure 1. A drug discovery framework for antiviral small molecules against SARS-CoV-2. The conditional compound generator, called agent (see **A**), can produce novel structures specifically designed to target a protein of interest. The generative process starts with the encoding of the primary structure of the target protein into a latent space of protein sequences. The representation is fed into a molecular decoder of a separately pretrained molecule VAE to produce a candidate compound. Next, the proposed compound is evaluated by a critic (see **B**) composed by: a multimodal deep learning model that predicts protein-drug binding affinity using protein and compound sequences as input, and a QSAR-based score to punish toxicity. By means of the reward given by the critic, a closed-loop system is created and is trained with deep reinforcement learning to maximize a multi-objective reward.

primary protein structure(s). Deep learning for target-driven drug design was first formulated by Aumentado-Armstrong (2018) and similarly to Chenthamarakshan et al. (2020), our approach implements a conditional generator that can be applied to *unseen* protein targets. However, instead of using conditional sampling, we perform an RL-biased conditional generation (fusing the latent spaces of protein targets and small molecules) that is demonstrated to generalize to *unseen* targets. We further couple our model with IBM RXN¹, an AI-governed platform for automated chemical synthesis to promptly synthesize the best compounds (Schwaller et al., 2020).

2 Methods

SELFIES VAE. The molecular generator is a variational auto-encoder (VAE) that is pretrained on 1,576,904 bioactive compounds from ChEMBL (10% are held out as validation set). The VAE implementation mostly follows Born et al. (2020), i.e., it consists of two layers of

stack-augmented GRUs (Joulin & Mikolov, 2015) in both encoder and decoder. The latent space has a dimensionality of 256, molecules are represented as SELFIES (Krenn et al., 2019), a robust adaption of the molecular in-line notation SMILES (Simplified Molecular Input Line Entry Specification) devised for generative models, and one-hot encodings are used. During training, KL annealing, teacher forcing and token dropout are employed. During testing, the stochastic decoder is sampling from the softmax distribution over the output tokens.

Protein VAE. The protein VAE consists of 3 dense layers of sizes [768, 512, 256] in both encoder and decoder. The model is trained on ~400,000 proteins from UniProt (SwissProt). The proteins considered were selected by filtering out sequences longer than 8,190 amino acids. The maximum length of the sequences has been selected to accommodate the statistics for the SARS-CoV-2 relevant proteins compiled by UniProt². Note that the VAE is not trained on the raw sequences but on 768 dimensional latent representations obtained from TAPE (Rao et al., 2019). During training, KL annealing and dropout are employed.

Protein-ligand affinity prediction. To predict CPI, we utilize a bimodal neural network based on the multiscale convolutional attention model (Manica et al., 2019; Cadow et al., 2020) (MCA, for model architecture see appendix Figure 4). Drug-protein binding affinity data is obtained from BindingDB, a public database of 1,813,527 measured binding affinities between 7,044 proteins and 802,551 small drug-like compounds as of March 2020. As for the protein VAE data, the database was filtered from entries with target sequences longer than 8,190 amino acids. The remaining 1,361,076 entries with an average of 187 reported compounds per target protein were taken as binding examples. From compounds not reported as entry for a given target 187 compounds were randomly sampled as non-binding to the respective protein target to match the binding examples. Finally, the combined examples were filtered for invalid SMILES to a total of 2,723,726 binding/non-binding pairs of 771,839 compounds and split into random, stratified train (72%), validation (18%) and test (10%) folds.

Toxicity prediction. Using the Tox21 database available through DeepChem (Wu et al., 2018), we trained a MCA model on the augmented SMILES sequences (Bjerrum, 2017) to predict the 12 toxicity classes.

Conditional generation. The conditional generative model is obtained by encoding a protein target with the protein VAE and decoding it with the pretrained molecular decoder.

²covid-19.uniprot.org/as on 22 May 2020.

¹<https://rxn.res.ibm.com/>

The objective function of this hybrid VAE G_{Θ} is:

$$\Pi(\Theta|r) = \sum_{s_T \in S^*} P_{\Theta}(s_T|r) R(s_T, r) \quad (1)$$

where r is the protein target of interest, s_T is a SELFIES string at time T (terminated with the $\langle \text{END} \rangle$ token), S^* resembles the molecular space and $P_{\Theta}(s_T|r)$ is the probability to sample s_T given r . In detail, $P_{\Theta}(s_T|r) := \prod_{t=0:T} p(a_t|s_{t-1})$ where $s_0 = r$ and a_t is the action at time t sampled from the dictionary of SELFIES tokens. $R(s_T, r)$ is the output of the critic C , in our case a multi objective:

$$R(s_T, r) = A(s_T, r) + 0.5 T(s_T) \quad (2)$$

where $A(\cdot)$ is the affinity predictor and $T(\cdot)$ is the toxicity predictor that returns 1 iff s_T is inactive in all 12 assays. Since Equation 1 is intractable to compute, it is approximated using policy gradient and subject to maximization using REINFORCE (Williams, 1992).

Protein targets. We fetched the 41 protein targets that are labelled as relevant to SARS-CoV-2 by UniProt (as on 22 May 2020). A full list of targets is given in Table 2 and includes e.g. the 3C-like protease (M_{pro}) which was identified as most promising candidate for antiviral compound development (Wu et al., 2020) and was already investigated with generative models (Zavoronkov et al., 2020) and molecular docking studies (Khaerunnisa et al., 2020). Other proteins are the nucleocapsid (N-) protein or the spike glycoprotein which is the most important surface protein, the target of chloroquine and mediates entrance to human respiratory epithelial cells by interacting with the ACE2 receptor.

3 Results

Protein-ligand affinity prediction. Because the conditional generation focuses on antiviral drug design, it is important that the affinity predictor generalizes well for viral proteins. The results of the MCA model on validation and test data from BindingDB are displayed in Table 1 next to the performance on 10k viral proteins. The results shows

Table 1. Result of bimodal affinity predictor on BindingDB data.

| | Validation | Test | Viral |
|--------------------------|------------|-------|-------|
| ROC-AUC | 0.968 | 0.969 | 0.96 |
| Average precision | 0.963 | 0.965 | 0.92 |

that the model learned reasonably well to classify CPI as binding or non-binding.

Toxicity predictor. Because toxicity is a major cause of the high attrition rate in drug discovery, we decided to perform

a multi-objective optimization (see Equation 2) based on toxicity and binding affinity. Across 10 runs, this model achieved a ROC-AUC of 0.877 ± 0.04 , surpassing prior results on this benchmarked dataset. Both the affinity and toxicity predictor are not investigated further herein, but employed as reward function for the conditional generation.

Conditional generative model

In this study, we are not primarily interested in proposing the *best* possible compounds. We rather want to validate whether our framework can go beyond current approaches for target-driven compound design (Chenthamarakshan et al., 2020; Zavoronkov et al., 2019; 2020) in the sense that it does not require fine-tuning for specific targets. We therefore investigated the generalization capabilities of our framework by performing a leave-one-out-cross-validation (LooCV) on the 41 targets. The RL optimization was performed for 5 epochs and 500 molecules were sampled in each step. The results are depicted in Table 2 and demonstrate that in 35 out of 41 cases the model proposed more binding compounds against an unseen target, compared to the baseline SELFIES VAE.

From the baseline SELFIES VAE, a total of 3,000 molecules was sampled. The average ratio of compounds predicted to bind increased from 18% to 26% with the best epoch averaging 33%. For example density plots see the appendix (Figure 5). We additionally optimized the generator to propose less toxic compounds. This succeeded to a lesser extent, probably at least partially caused by the lower weight in the reward function. For a qualitative evaluation, Figure 2

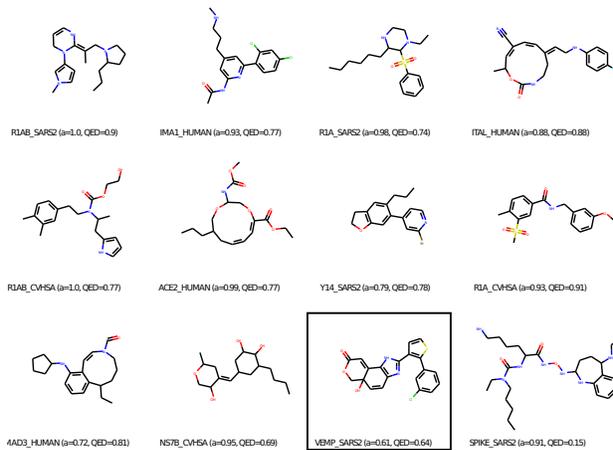


Figure 2. Molecules sampled against specific protein targets. For a selection of 12 targets, the generated compound with the highest reward is depicted. a stands for binding affinity. The encircled molecule is discussed further in the case study.

shows a selection of the sampled molecules alongside their QED score (Bickerton et al., 2012)).

Table 2. Generating antiviral compounds against unseen SARS-CoV-2 targets. For each of the 41 targets, Affinity_0 shows the fraction of binding molecules sampled *before* training. Aff_{best} shows the fraction at the best epoch of RL training, while $\text{Aff}_{\text{median}}$ shows the median across all 5 training epochs. The same applies to Tox_{best} and Tox_{med} , where Tox_0 was 8.7%.

| Target protein | Affinity_0 | $\text{Aff}_{\text{med}} \pm \text{SEM}$ | Aff_b | $\text{Tox}_{\text{med}} \pm \text{SED}$ | Tox_b |
|------------------|---------------------|--|----------------|--|----------------|
| VME1-CVHSA | 20% | 18% \pm 3% | 29% | 6% \pm 3% | 19% |
| IMAI-HUMAN | 88% | 97% \pm 1% | 100% | 5% \pm 3% | 18% |
| VEMP-SARS2 | 29% | 16% \pm 2% | 20% | 9% \pm 2% | 12% |
| NS7B-SARS2 | 25% | 30% \pm 5% | 33% | 7% \pm 5% | 25% |
| ITAL-HUMAN | 24% | 16% \pm 6% | 43% | 9% \pm 1% | 12% |
| NCAP-CVHSA | 17% | 11% \pm 1% | 15% | 12% \pm 2% | 14% |
| R1AB-CVHSA | 58% | 90% \pm 2% | 91% | 9% \pm 1% | 11% |
| NS8B-CVHSA | 9% | 12% \pm 2% | 20% | 7% \pm 4% | 25% |
| A0A663DJA2-SARS2 | 26% | 35% \pm 3% | 41% | 14% \pm 3% | 18% |
| NS8A-CVHSA | 21% | 47% \pm 4% | 55% | 10% \pm 1% | 10% |
| NS7A-SARS2 | 4% | 3% \pm 1% | 7% | 10% \pm 3% | 19% |
| Y14-SARS2 | 17% | 29% \pm 4% | 43% | 8% \pm 2% | 14% |
| NS6-SARS2 | 20% | 12% \pm 3% | 22% | 4% \pm 3% | 14% |
| SMAD3-HUMAN | 50% | 74% \pm 3% | 86% | 6% \pm 1% | 10% |
| SPIKE-CVHSA | 3% | 0% \pm 1% | 5% | 7% \pm 1% | 11% |
| DDX1-HUMAN | 9% | 14% \pm 2% | 20% | 9% \pm 1% | 10% |
| AP3A-SARS2 | 4% | 0% \pm 1% | 3% | 9% \pm 3% | 19% |
| R1A-CVHSA | 14% | 45% \pm 3% | 50% | 9% \pm 1% | 11% |
| NS8-SARS2 | 7% | 10% \pm 3% | 18% | 10% \pm 1% | 15% |
| PHB2-HUMAN | 4% | 3% \pm 0% | 4% | 11% \pm 3% | 23% |
| SGTA-HUMAN | 11% | 12% \pm 1% | 13% | 8% \pm 1% | 12% |
| NS7A-CVHSA | 18% | 35% \pm 5% | 59% | 11% \pm 2% | 15% |
| ORF9B-CVHSA | 9% | 11% \pm 2% | 17% | 6% \pm 1% | 11% |
| R1A-SARS2 | 62% | 82% \pm 3% | 89% | 8% \pm 2% | 14% |
| Y14-CVHSA | 14% | 15% \pm 2% | 23% | 11% \pm 2% | 15% |
| ORF9B-SARS2 | 18% | 12% \pm 1% | 15% | 12% \pm 2% | 16% |
| TMPS2-HUMAN | 6% | 5% \pm 1% | 6% | 6% \pm 1% | 10% |
| BST2-HUMAN | 10% | 5% \pm 3% | 16% | 10% \pm 2% | 14% |
| NS3B-CVHSA | 25% | 23% \pm 2% | 29% | 12% \pm 1% | 15% |
| SPIKE-SARS2 | 7% | 6% \pm 2% | 12% | 10% \pm 1% | 12% |
| FURIN-HUMAN | 28% | 27% \pm 4% | 36% | 9% \pm 3% | 20% |
| AP3A-CVHSA | 9% | 0% \pm 1% | 6% | 8% \pm 1% | 12% |
| VME1-SARS2 | 15% | 16% \pm 3% | 27% | 6% \pm 2% | 14% |
| NS7B-CVHSA | 21% | 26% \pm 1% | 27% | 7% \pm 1% | 11% |
| MPP5-HUMAN | 5% | 9% \pm 2% | 11% | 15% \pm 2% | 16% |
| ACE2-HUMAN | 51% | 77% \pm 4% | 85% | 5% \pm 2% | 12% |
| VEMP-CVHSA | 21% | 25% \pm 3% | 30% | 12% \pm 2% | 20% |
| NS6-CVHSA | 10% | 13% \pm 1% | 15% | 3% \pm 3% | 14% |
| PHB-HUMAN | 3% | 0% \pm 1% | 3% | 6% \pm 1% | 7% |
| R1AB-SARS2 | 83% | 100% \pm 0% | 100% | 5% \pm 1% | 7% |
| NCAP-SARS2 | 25% | 5% \pm 2% | 9% | 9% \pm 4% | 24% |
| Average | 18% | 26% \pm 4% | 33% | 9% \pm 0.5% | 15% |

To investigate the learned chemical space, we assembled a dataset of 10,000 random ChEMBL compounds, 3,000 molecules sampled from the unbiased VAE, 3,000 molecules sampled from the biased generator and 82 SARS-CoV-2 candidate drugs from the literature (top 15 matches on PubChem and 69 compounds identified via protein-interaction-maps (Gordon et al., 2020), excluding 2 duplicates). For all these molecules, binding affinities were computed alongside other pharmacological properties. Next, a UMAP (McInnes et al., 2018) was performed on the ECFP4 fingerprint (Rogers & Hahn, 2010) and visualized with Faerun/Tmap (Probst & Reymond, 2018; 2020)³. The interactive visualisation shows that the RL optimisation lead to over-sampling a manifold of the chemical space that is more densely populated with binding compounds. The 3D UMAP shows that the currently investigated candidate molecules (red) are structurally fairly dissimilar (i.e. scattered across

³The Faerun visualization with ECFP is available at: https://paccmann.github.io/assets/umap_fingerprints.html

the chemical space). But it gives evidence that our model successfully navigates the chemical space towards regions of high reward. While this shows that the generator succeeded in its objective, ultimately, the quality of the reward function remains the bottleneck of the framework.

Case study. For a more detailed assessment of the quality of the molecules, we ranked all $\sim 3,000$ conditionally generated molecules by their tanimoto similarity τ to the closest neighbour of the 81 literature candidates. Among the top 5 molecules, we found the molecule encircled in Figure 2, generated against VEMP_{SARS2} (UniProt ID: P0DTC4), the envelope small membrane protein (E-Protein), a key player for virion assembly and morphogenesis. Our candidate exhibits the highest tanimoto similarity to the compounds MZ1 and dBET6 ($\tau = 0.64$ based on RDKit fingerprint). These two pre-clinical SARS-CoV-2 drug candidates were identified by Gordon et al. (2020) to target the E-protein by degrading the human BRD2 and BRD4 proteins and thus preventing the virus from inducing changes in the host’s protein expression.

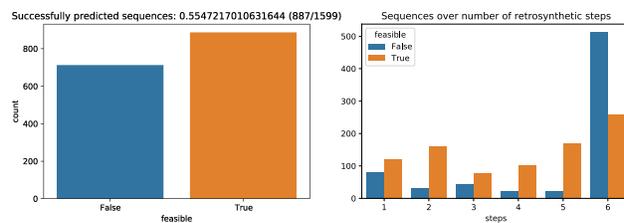


Figure 3. Results of retrosynthesis attempts. A retrosynthetic pathway is considered feasible if it leads to commercially available precursors within six reaction steps. Orange indicates feasibility while blue indicates non feasibility. On the left (A), overall feasibility of the predicted sequences. On the right (B), feasibility over the number of reaction steps.

The Top-5 candidate compounds for each protein target were further analyzed for synthetic feasibility using IBM RXN’s retrosynthesis engine (Schwaller et al., 2020). We performed the predictions using the Python package rxn4chemistry⁴. Figure 3 shows the predictions over the retrosynthetic sequences estimated for all molecules. Although the generated molecules are not optimized for synthetic accessibility, more than half of the synthesis routes predicted are feasible. It’s interesting to observe how more than 300 sequences requires only a single or two steps reactions, indicating, assuming a reasonable yield, an extremely efficient synthesis for part of the molecules generated.

4 Discussion

The dramatic effect of the COVID-19 pandemic is compounded by the lack of vaccines and therapeutic agents

⁴<https://github.com/rxn4chemistry/rxn4chemistry>

against SARS-CoV-2. Worse, traditional approaches to drug discovery are slow, inefficient, error-prone and costly.

Here, we proposed a novel framework for compound design that can be targeted towards *any viral target protein* with no retraining requirements. We showcased the potential of the framework by successfully tackling the problem of generating novel compounds with high binding affinity to unseen targets, while controlling toxicity of the generated molecules. Furthermore, for each target, we estimated retrosynthetic pathways of the most promising molecules, to assess the feasibility of the generated compounds. Large-scale screening data for 1,670 compounds tested against SARS-CoV-2 proteins that just became available (Heiser et al., 2020), will be used in the future to improve the affinity predictor, one of the bottlenecks of our approach.

References

- Aumentado-Armstrong, T. Latent molecular optimization for targeted therapeutic design. *arXiv preprint arXiv:1809.02032*, 2018.
- Bickerton, G. R., Paolini, G. V., Besnard, J., Muresan, S., and Hopkins, A. L. Quantifying the chemical beauty of drugs. *Nature chemistry*, 4(2):90, 2012.
- Bjerrum, E. J. Smiles enumeration as data augmentation for neural network modeling of molecules. *arXiv preprint arXiv:1703.07076*, 2017.
- Born, J., Manica, M., Oskoei, A., Cadow, J., and Martínez, M. R. Paccmann rl: Designing anticancer drugs from transcriptomic data via reinforcement learning. In *International Conference on Research in Computational Molecular Biology*, pp. 231–233. Springer, 2020.
- Cadow, J., Born, J., Manica, M., Oskoei, A., and Rodríguez Martínez, M. Paccmann: a web service for interpretable anticancer compound sensitivity prediction. *Nucleic Acids Research*, 2020.
- Chenthamarakshan, V., Das, P., Padhi, I., Strobelt, H., Lim, K. W., Hoover, B., Hoffman, S. C., and Mojsilovic, A. Target-specific and selective drug design for covid-19 using deep generative models. *arXiv preprint arXiv:2004.01215*, 2020.
- Drosten, C. et al. Identification of a novel coronavirus in patients with severe acute respiratory syndrome. *New England journal of medicine*, 348(20):1967–1976, 2003.
- Gordon, D. E., Jang, G. M., Bouhaddou, M., Xu, J., Obernier, K., White, K. M., O’Meara, M. J., Rezelj, V. V., Guo, J. Z., Swaney, D. L., et al. A sars-cov-2 protein interaction map reveals targets for drug repurposing. *Nature*, pp. 1–13, 2020.
- Heiser, K. et al. Identification of potential treatments for covid-19 through artificial intelligence-enabled phenomic analysis of human cells infected with sars-cov-2. *bioRxiv*, 2020.
- Joulin, A. and Mikolov, T. Inferring algorithmic patterns with stack-augmented recurrent nets. In *Advances in neural information processing systems*, pp. 190–198, 2015.
- Karimi, M., Wu, D., Wang, Z., and Shen, Y. Deepaffinity: interpretable deep learning of compound–protein affinity through unified recurrent and convolutional neural networks. *Bioinformatics*, 35(18):3329–3338, 2019.
- Khaerunnisa, S., Kurniawan, H., Awaluddin, R., Suhartati, S., and Soetjipto, S. Potential inhibitor of covid-19 main protease (mpro) from several medicinal plant compounds by molecular docking study. *Prepr. doi10.20944/preprints202003.0226.v1*, pp. 1–14, 2020.
- Krenn, M., Häse, F., Nigam, A., Friederich, P., and Aspuru-Guzik, A. Selfies: a robust representation of semantically constrained graphs with an example application in chemistry. *arXiv preprint arXiv:1905.13741*, 2019.
- Manica, M., Oskoei, A., Born, J., Subramanian, V., Sáez-Rodríguez, J., and Rodríguez Martínez, M. Toward explainable anticancer compound sensitivity prediction via multimodal attention-based convolutional encoders. *Molecular Pharmaceutics*, 2019.

- McInnes, L., Healy, J., and Melville, J. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- Probst, D. and Reymond, J.-L. Fun: a framework for interactive visualizations of large, high-dimensional datasets on the web. *Bioinformatics*, 34(8):1433–1435, 2018.
- Probst, D. and Reymond, J.-L. Visualization of very large high-dimensional data sets as minimum spanning trees. *Journal of Cheminformatics*, 12(1):1–13, 2020.
- Rao, R., Bhattacharya, N., Thomas, N., Duan, Y., Chen, X., Canny, J., Abbeel, P., and Song, Y. S. Evaluating protein transfer learning with tape. In *Advances in Neural Information Processing Systems*, pp. 9686–9698, 2019.
- Rogers, D. and Hahn, M. Extended-connectivity fingerprints. *Journal of chemical information and modeling*, 50(5):742–754, 2010.
- Scannell, J. W., Blanckley, A., Boldon, H., and Warrington, B. Diagnosing the decline in pharmaceutical r&d efficiency. *Nature reviews Drug discovery*, 11(3):191, 2012.
- Schwaller, P., Petraglia, R., Zullo, V., Nair, V. H., Haeuselmann, R. A., Pisoni, R., Bekas, C., Laino, T., et al. Predicting retrosynthetic pathways using transformer-based models and a hyper-graph exploration strategy. *Chemical Science*, 2020.
- Shamshirian, A., Hessami, A., Heydari, K., and Alizadeh-Navaei, R. e. a. Hydroxychloroquine versus covid-19: A periodic systematic review and meta-analysis. *medRxiv*, 2020. doi: 10.1101/2020.04.14.20065276.
- Tang, B., He, F., Liu, D., Fang, M., Wu, Z., and Xu, D. Ai-aided design of novel targeted covalent inhibitors against sars-cov-2. *bioRxiv*, 2020.
- Wang, Y., Zhang, D., Du, G., Du, R., Zhao, J., Jin, Y., Fu, S., Gao, L., Cheng, Z., Lu, Q., et al. Remdesivir in adults with severe covid-19: a randomised, double-blind, placebo-controlled, multicentre trial. *The Lancet*, 2020.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- Wu, C., Liu, Y., Yang, Y., Zhang, P., Zhong, W., Wang, Y., Wang, Q., Xu, Y., Li, M., Li, X., et al. Analysis of therapeutic targets for sars-cov-2 and discovery of potential drugs by computational methods. *Acta Pharmaceutica Sinica B*, 2020.
- Wu, Z., Ramsundar, B., Feinberg, E. N., Gomes, J., Geniesse, C., Pappu, A. S., Leswing, K., and Pande, V. Moleculenet: a benchmark for molecular machine learning. *Chemical science*, 9(2):513–530, 2018.
- Zhavoronkov, A. et al. Deep learning enables rapid identification of potent ddr1 kinase inhibitors. *Nature biotechnology*, 37(9): 1038–1040, 2019.
- Zhavoronkov, A. et al. Potential non-covalent sars-cov-2 3c-like protease inhibitors designed using generative deep learning approaches and reviewed by human medicinal chemist in virtual reality. *ChemRxiv*, 2020.

Appendix

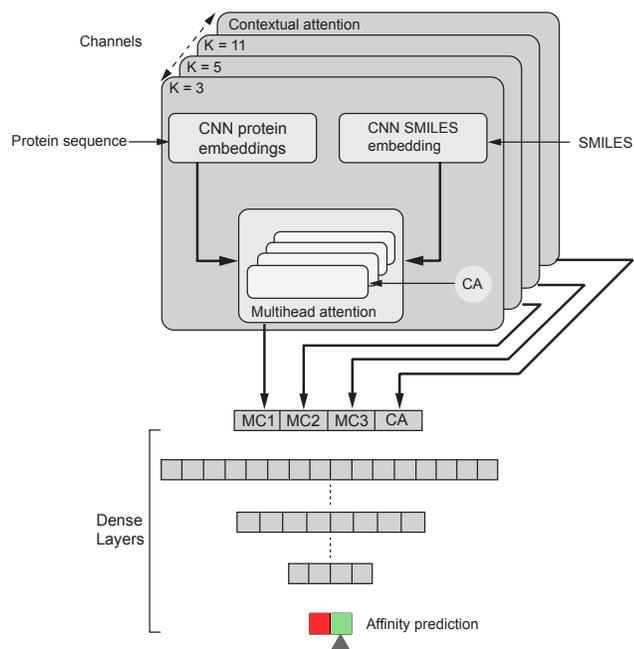


Figure 4. Multimodal protein-ligand affinity predictor for antiviral compounds. Inspired by the MCA architecture in Manica et al. (2019), this is a multimodal classification model that performs multiscale convolutions on SMILES embeddings (ligand) and amino acid embeddings (protein). The output is fed into multiple heads of contextual attention mechanism prior to a set of stacked dense layers.

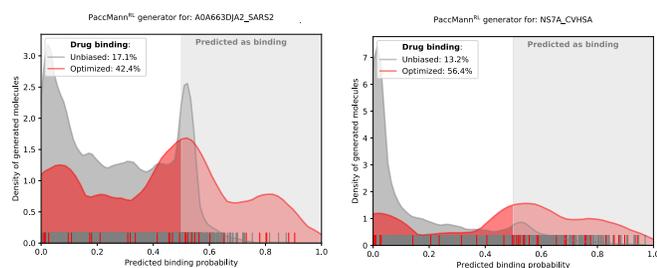


Figure 5. Exemplary density functions of conditional generation. Gray distributions show predicted binding affinities of $n=3,000$ molecules sampled from an unbiased SELFIES VAE. Depicted in red are the densities obtained by sampling from the RL optimized conditional generative model. It can be seen that the optimization biased the sampling process toward regions of the chemical space that are more densely populated with ligands that are predicted to bind.

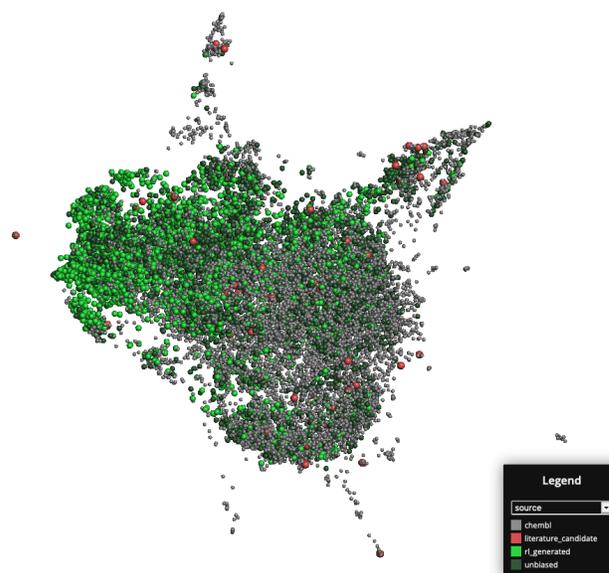


Figure 6. UMAP dimensionality reduction of the chemical space visualized with Faerun/TMap (Probst & Reymond, 2018; 2020). Snapshot of the Faerun visualization of a UMAP of 10,000 molecules randomly selected from ChEMBL (grey), alongside 3,000 molecules sampled from the unbiased generator (dark green), 3,500 molecules sampled against SARS-CoV-2 related target proteins and 82 drug candidates according to the literature.