

PAPER • OPEN ACCESS

Data-driven molecular design for discovery and synthesis of novel ligands: a case study on SARS-CoV-2

To cite this article: Jannis Born *et al* 2021 *Mach. Learn.: Sci. Technol.* **2** 025024

View the [article online](#) for updates and enhancements.

You may also like

- [Emerging nanotechnology role in the development of innovative solutions against COVID-19 pandemic](#)
Zeeshan Ahmad Bhutta, Ayesha Kanwal, Moazam Ali et al.
- [Inactivation of SARS-CoV-2 in clinical exhaled breath condensate samples for metabolomic analysis](#)
Shuang Hu, Mitchell M McCartney, Juan Arredondo et al.
- [Binding of SARS-CoV-2/SARS-CoV spike protein with human ACE2 receptor](#)
Rajendra P Koirala, Bidhya Thapa, Shyam P Khanal et al.



PAPER

OPEN ACCESS

RECEIVED

27 November 2020

REVISED

9 February 2021

ACCEPTED FOR PUBLICATION

19 February 2021

PUBLISHED

25 March 2021

Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Data-driven molecular design for discovery and synthesis of novel ligands: a case study on SARS-CoV-2

Jannis Born^{1,2,4} , Matteo Manica^{1,4} , Joris Cadow^{1,4} , Greta Markert¹ , Nil Adell Mill¹ , Modestas Filipavicius¹, Nikita Janakarajan¹ , Antonio Cardinale^{1,3}, Teodoro Laino¹ and María Rodríguez Martínez¹

¹ IBM Research Europe, Zurich, Switzerland

² Department of Biosystems Science and Engineering, ETH Zurich, Switzerland

³ Università di Pisa, Pisa, Italy

⁴ Shared first-authorship

E-mail: jab@zurich.ibm.com, tte@zurich.ibm.com, dow@zurich.ibm.com and mrm@zurich.ibm.com


Keywords: drug discovery, deep learning, COVID-19, generative chemistry, SARS-CoV-2, machine learning, compound protein interaction

Supplementary material for this article is available [online](#)

Abstract

Bridging systems biology and drug design, we propose a deep learning framework for *de novo* discovery of molecules tailored to bind with given protein targets. Our methodology is exemplified by the task of designing antiviral candidates to target SARS-CoV-2 related proteins. Crucially, our framework does not require fine-tuning for specific proteins but is demonstrated to generalize in proposing ligands with high predicted binding affinities against unseen targets. Coupling our framework with the automatic retrosynthesis prediction of *IBM RXN for Chemistry*, we demonstrate the feasibility of swift chemical synthesis of molecules with potential antiviral properties that were designed against a specific protein target. In particular, we synthesize an antiviral candidate designed against the host protein angiotensin converting enzyme 2 (ACE2); a surface receptor on human respiratory epithelial cells that facilitates SARS-CoV-2 cell entry through its spike glycoprotein.

This is achieved as follows. First, we train a multimodal ligand–protein binding affinity model on predicting affinities of bioactive 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 that consists of two variational autoencoders (VAE), our framework steers the generation toward regions of the chemical space with high-reward 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 reinforcement learning, it is demonstrated that in 35 out of 41 cases, the generation is biased towards sampling binding ligands, with an average increase of 83% comparing to an unbiased VAE. The generated molecules exhibit favorable properties in terms of target binding affinity, selectivity and drug-likeness. We use molecular retrosynthetic models to provide a synthetic accessibility assessment of the best generated hit molecules. Finally, with this end-to-end framework, we synthesize 3-Bromobenzylamine, a potential inhibitor of the host ACE2 protein, solely based on the recommendations of a molecular retrosynthesis model and a synthesis protocol prediction model. We hope that our framework can contribute towards swift discovery of *de novo* molecules with desired pharmacological properties.



1. Introduction

1.1. COVID-19 antivirals

The severe acute respiratory syndrome (SARS) coronavirus disease (COVID-19) is an acute respiratory disease caused by SARS-CoV-2 that, to date, has infected more than a hundred million humans and killed more than two million. Despite longstanding efforts to understand the pathogenicity of coronaviruses [1], no drugs were approved before the outbreak of the COVID-19 pandemic, and thus, new systematic approaches to identify effective antiviral agents are urgently needed. Current efforts are predominantly focused on drug repurposing strategies, and lots of effort was initially devoted to the clinical investigation of a handful of promising candidates, including remdesivir and hydroxychloroquine. Initial hopes regarding hydroxychloroquine have been diminished as it was not found effective in a meta-study of human clinical trials [2]. While evidence for remdesivir has been conflicting [3, 4], it has been granted an emergency approval in North America, Asia and Europe [5]. Recently, protein-protein-interaction studies identified 69 promising compounds by measuring binding affinities of 26 out of the 29 SARS-CoV-2 proteins against human proteins [6].

With limited success in the drug repurposing strategies, it is worth exploiting *de novo* drug discovery approaches against SARS-CoV-2. Drug discovery is a daunting challenge, with a search space of 10^{60} compounds [7], ~10 years from design to market [8] and costs of up to \$3 billion per new FDA-approved drug [9]. Given that so far there are only ~1500 FDA-approved drugs on the market [10], while the number of already synthesized molecules is at least 60 million [11], the total attrition rate of drug discovery is above 99.99%. Additionally, it takes around 10 years until a typical compound reaches the market [8]. 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 [12]. Deep learning has further been proven capable of *in silico* design of molecules with desired chemical properties [13–16] and shown potential to accelerate the discovery of DDR1 inhibitors [17]. 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 [18–20] but these studies manually curated datasets to design 3C-like protease inhibitors.

1.2. Our contribution

Here, we aim to bridge systems biology and drug discovery, and use deep learning to explore target-driven drug design with conditional generative models. Our framework (see (A) and (B)) for conditional molecular design builds upon our previous work, *PaccMann^{RL}* [21]; 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 primary protein structure. Deep learning for target-driven drug design was first formulated in

2018 by Aumentado-Armstrong [22] and has since been investigated by others [23–25]. Similarly to Chenthamarakshan *et al* [25], our approach implements a conditional generator that can be applied to *unseen* protein targets. However, instead of using a conditional rejection sampling approach, we use deep reinforcement learning and conditionally generate molecules by fusing the latent spaces of protein targets and small molecules. Importantly, our approach is demonstrated to generalize to *unseen* targets — it can generate high fractions of potentially binding ligands for protein sequences it has never been exposed to. We further couple our model with IBM RXN, an AI-governed platform for retrosynthesis prediction [26, 27], and demonstrate the entire cycle from design to synthesis of one of the most promising compounds.

We emphasize that our framework is generic in its applicability for protein-target driven molecular design and synthesis. To showcase the methodology, we focus on designing novel antiviral candidates against SARS-CoV-2-related virus and host proteins.

Notably, the main contribution of our work is not claiming to find the most promising SARS-CoV-2 antiviral candidates, but rather to present new technology, namely how a generic, machine-learning-driven, fully automatic pipeline for design and chemical synthesis may potentially accelerate the discovery of hit-like molecules.

2. Materials and methods

2.1. Problem formulation

The goal is to develop a conditional generative model G_{Θ} that can be queried with a protein and returns novel ligands that have high binding affinities to the target and, as a secondary objective, low toxicity. Let \mathcal{P} denote the space of proteins and \mathcal{M} the molecular space. We are interested to learn a mapping $G_{\Theta} : \mathcal{P} \rightarrow \mathcal{M}$ s.t. the reward $R(p, m)$ with $p \in \mathcal{P}, m \in \mathcal{M}$ is maximal.

2.1.1. Predictive models

$R(\cdot)$ is a multimodal reward function depending on p and m . Specifically, let

$$R(p, m) := \text{Aff}(p, m) - \gamma \cdot \text{Tox}(m) \quad (1)$$

with $\text{Aff}(\cdot)$ as the target-compound binding affinity of p and m , $\text{Tox}(\cdot)$ as the toxicity of m and $\gamma \in \mathbb{R}^+$ as a user-defined parameter to control the importance of toxicity (we used $\gamma = 0.5$). Since the exact computation of Aff and Tox are intractable to compute *in-silico* (as they require *in-vitro* experiments) they are approximated using Φ_{Tox} and Φ_{Aff} . Precisely, $\Phi_{\text{Tox}} : \mathcal{M} \rightarrow \hat{\mathbf{y}}_{\text{Tox}}$ maps a molecule to a toxicity vector of 12 toxicity classes where 1 means toxic and 0 non-toxic. A positive reward is given if and only if the molecule is not predicted active in any of the 12 toxicity classes. Furthermore $\Phi_{\text{Aff}} : \mathcal{P} \times \mathcal{M} \rightarrow [0, 1]$ maps a protein-ligand tuple to a probability that the ligand binds to the protein. For details on this model and the data (e.g. the threshold for binding strength) see subsection 2.2.2 and S1.2 (available online at stacks.iop.org/MLST/2/025024/mmedia).

2.1.2. Generative models

The two predictive models for binding affinity and toxicity are employed as reward functions for the conditional generator, which is modelled as follows.

2.1.2.1. Pretraining

First, two separate models $\Theta_{\text{Prot}} : \mathcal{P} \rightarrow \mathcal{P}$ and $\Theta_{\text{Mol}} : \mathcal{M} \rightarrow \mathcal{M}$ are devised. The parameters of both models can be optimized without supervision using training data of unlabelled proteins and molecules, respectively. These models implicitly define a mapping from discrete structures (proteins and molecules) to continuous representations. Θ_{Prot} and Θ_{Mol} are independent generative models that are parameterized with VAEs [28] so that both can be represented as a composition of encoder/decoder modules: $\Theta_{\text{Prot}} = [\Theta_{\text{Prot}}^{\text{Dec}} \circ \Theta_{\text{Prot}}^{\text{Enc}}]$ and $\Theta_{\text{Mol}} = [\Theta_{\text{Mol}}^{\text{Dec}} \circ \Theta_{\text{Mol}}^{\text{Enc}}]$. The two models are tuned to optimize the ELBO: $\mathcal{L}_{\text{VAE}} := \mathbb{E}[\log p(\mathbf{z}|\mathbf{x})] - \mathcal{D}_{\text{KL}}[q(\mathbf{z}|\mathbf{x}), p(\mathbf{z})]$ where $q(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\vec{0}, \mathbf{I})$, i.e. the latent code is modelled using a multivariate unit Gaussian following standard VAE formulation [28]. After training, the generative model can be utilized by sampling from $p(\mathbf{z})$ and applying $\Theta_{\text{Mol}}^{\text{Dec}} : \mathcal{Z} \rightarrow \mathcal{M}$ which constitutes our baseline for molecule generation.

2.1.2.2. Conditional generation

After the above models are trained, the conditional generative model is defined by $G_{\Theta} : [\Theta_{\text{Mol}}^{\text{Dec}} \circ \Theta_{\text{Prot}}^{\text{Enc}}] : \mathcal{P} \rightarrow \mathcal{Z} \rightarrow \mathcal{M}$. In other words, the conditional generator is obtained by encoding a protein with the protein VAE and decoding the latent code with the molecular decoder. This means, a molecule m is obtained from a

protein p by: $m = G_{\Theta}(p) = \Theta_{\text{Mol}}^{\text{Dec}}(\Theta_{\text{Prot}}^{\text{Enc}}(p))$. The final training objective function of this hybrid VAE (G_{Θ}) is

$$\mathcal{R}_{G_{\Theta}} = \sum_{m \in \mathcal{M}} P_{\Theta}(m|p)R(p, m), \quad (2)$$

where $P_{\Theta}(m|p)$ indicates the conditional probability approximated by G_{Θ} . Since equation (2) is intractable to compute, it is approximated using policy gradient and subject to maximization using REINFORCE [29], as proposed in ReLeaSE [15]. Notably, the conditional generation is not limited to binary classifiers, as $R(\cdot)$ can be instantiated with any function. Figures 1(A) and (B) summarize the structure of the proposed system.

2.2. Data and implementation

2.2.1. Toxicity prediction

2.2.1.1. Data

For toxicity prediction, we utilize the Tox21 database [30]. This dataset consists of $N_{\text{Tox}} = 11\,765$ training molecules that are screened against 12 toxicity assays of nuclear receptor and stress response pathways, and labelled with binary values for each class (toxic vs non-toxic). The test dataset consists of 648 molecules (data split is provided).

2.2.1.2. Model

The toxicity prediction model $\Phi_{\text{Tox}} : \mathcal{M} \rightarrow \hat{y}_{\text{Tox}}$ was implemented through a multiscale convolutional attention model that ingested augmented SMILES sequences [31, 32] to predict the 12 toxicity endpoints. For details see S1.1. The purpose of this *in silico* screening was to assess the toxicity of parent molecules and their metabolites. Our work on toxicity prediction is further detailed in [33].

2.2.2. Protein-ligand affinity prediction

2.2.2.1. Data

Drug-protein binding affinity data is obtained from BindingDB [34], a public database of measured binding affinities between proteins and small drug-like compounds. After data processing (details in S1.2) 1361 076 entries (7302 protein targets, 772 634 compounds) were taken as positive binding examples. Negative samples, on the other hand, were obtained by randomly assigning 187 compounds to each target from the list of compounds not reported to bind to that particular target yielding a total of 2273 726 samples. For reasons on the choice for classification instead of regression see S1.2.

2.2.2.2. Model

To predict CPI, we model $\Phi_{\text{Aff}} : \mathcal{P} \times \mathcal{M} \rightarrow [0, 1]$ with a bimodal neural network based on the multiscale convolutional attention model (MCA [35, 36]). This model ingests a molecule $m \in \mathcal{M}$ represented as a SMILES sequence and a protein $p \in \mathcal{P}$ represented as an amino acid sequence, applies convolutions (to aggregate local information) and a contextual attention mechanism (to focus on relevant substructures), before a continuous prediction $y \in [0, 1]$ is obtained that is interpreted as probability that a binding occurs. For model details see S1.2.

2.2.3. Protein VAE

2.2.3.1. Data

For learning the protein space \mathcal{P} we utilize 404 552 proteins from UniProt [37]. The proteins considered were selected by filtering out sequences longer than 8190 amino acids. We choose our training data to be vectorial representations of 768 dimensions rather than primary structures), specifically embeddings generated by a BERT [38] language model released in TAPE [39].

TAPE adapted self-supervised natural language models like BERT to learn continuous representations of proteins on large scale. The amino acid embeddings were pretrained on 32.6 M Pfam sequences [40] following the original BERT architecture and learned via masked-token prediction [39]. The TAPE model contains 110 M parameters and we exclusively used it for feature generation to obtain the training data for the protein VAE. To accommodate the occasional unusually long protein sequences, maximum sequence length was increased from the original 512 to 8192 tokens. In this work, we used the pretrained amino acid features from TAPE 'as is' by averaging them, i.e. without fine-tuning them for a downstream classification task. For conditional generation, however, we used the pooled instead of the averaged embeddings, without observing a notable difference in performance. The approach to use TAPE was pursued to circumvent the need to train a large-scale protein language model from scratch.

2.2.3.2. Model

As stated above, we model $\Theta_{\text{Prot}} : \mathcal{P} \rightarrow \mathcal{P}$ using a VAE [28] with three dense layers of sizes [768, 512, 256] with ReLU activation and batch normalization in both encoder and decoder. During training, we use KL annealing [41], dropout ($p = 0.2$), a learning rate of 3×10^{-3} and optimize RMSE.

2.2.4. SELFIES VAE

2.2.4.1. Data

For learning the molecular space \mathcal{M} we utilize 1576 904 bioactive compounds from ChEMBL. Ten percent are held out as validation set. We choose our training data $m \in \mathcal{M}$ not to be SMILES sequences (like for the toxicity and affinity predictor), but rather SELFIES strings [42], a robust adaption of SMILES that was devised for generative models and solves the validity problem in SMILES generation.

2.2.4.2. Model

The model for the chemical space $\Theta_{\text{Mol}} : \mathcal{M} \rightarrow \mathcal{M}$ is implemented using a VAE, following Born *et al* [21], i.e. it consists of two layers of stack-augmented GRUs [15] in both encoder and decoder. For details see S1.3. Notably, after Θ_{Mol} is pretrained, we sample from $p(\mathbf{z})$ and apply the decoder $\Theta_{\text{Mol}}^{\text{Dec}} : \mathcal{Z} \rightarrow \mathcal{M}$ as described above to constitute our baseline for molecular generation.

2.2.5. Conditional generation

2.2.5.1. Data

We retrieved 41 SARS-CoV-2 related protein targets as labelled in UniProt (as on 22 May 2020)⁵. A full list of targets is given in table 2 and includes among others the 3C-like protease (M_{pro}), a promising candidate for antiviral compound development [43] that has already been investigated with generative models [18] and molecular docking studies [44]. Other included proteins are the nucleocapsid (N-) protein and the spike glycoprotein. The latter is a surface protein, which mediates entrance to human respiratory epithelial cells by interacting with the ACE2 receptor and is the target of chloroquine. Notably, the targets used for conditional generation do not need to be present in the training data of the affinity predictor or the protein VAE. Hence, we can generate potential ligands for proteins without the need for any binding data, a key feature to combat new diseases. Here, 9/41 SARS-CoV-2 related proteins are present in the training data of the affinity predictor and 27/41 are present in the training data of the protein VAE.

2.2.5.2. Model

The conditional generator G_{Θ} consists of the pretrained protein encoder and the pretrained molecular decoder, s.t. $G_{\Theta} : [\Theta_{\text{Mol}}^{\text{Dec}} \circ \Theta_{\text{Prot}}^{\text{Enc}}]$ and learns a mapping $\mathcal{P} \rightarrow \mathcal{Z} \rightarrow \mathcal{M}$. In practice, we set $\gamma = 0.5$ in our multi-objective reward function since optimizing toxicity was a secondary objective of this work compared to optimizing binding affinity for a target protein. Moreover, we perform a leave-one-out cross-validation (LooCV) during conditional generation (i.e. $N_{\text{CG}} = 40$) and evaluate the model by the fraction of generated ligands with high predicted affinity and low toxicity ($\Theta_{\text{Aff}}(p, m) > 0.5$ and $\Theta_{\text{Tox}}(m) = 0$).

2.3. AI-assisted synthesis of a candidate compound

2.3.1. Retrosynthesis prediction

First, we predicted the best synthetic routes using the interface of IBM RXN⁶, specifically, a transformer-based retrosynthesis prediction engine [26]. The prediction is built upon sequence-based models that operate on SMILES strings, specifically two molecular transformers [45]; one for forward reaction prediction (reactants to product) and one for backward reaction prediction (product to reactants). The two models are combined to explore the retrosynthesis hypertree, using a beam search based on the forward model confidence scores, to find the most likely routes to synthesize a molecule given a set of commercially available chemicals. Given a SMILES string, this model predicts a reaction tree composed of starting materials and intermediates (all represented as SMILES) as nodes and reaction types as edges.

Next, we used a transformer model [46] to predict the optimal synthesis protocol using a text representation of the predicted synthesis steps. The architecture is trained to predict a sequence of actions given a reaction encoded as SMARTS. The synthesis action generation model is trained with a large number of chemical recipes extracted from a corpus of organic chemistry procedures [27]. For details on the methodology see [26, 27, 45, 46].

⁵ covid-19.uniprot.org/.

⁶ <https://rxn.res.ibm.com/>.

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

2.3.2. Chemical synthesis

In the following, the predicted procedure is described in detail. In a glass reactor of 100 ml, 9 ml of anhydrous THF were added at room temperature. Nine milliliter of a solution of 3-bromo-benzonitrile in THF (0.11 M, 1 mmol) were added under gentle (100 rpm) stirring. While maintaining the temperature of the reactor at 25 °C, 1 ml of LiAlH₄ in THF (2 M concentration) was added dropwise across 180 s. The reaction mixture was stirred for 5 min at 25 °C and then the excess of LiAlH₄ quenched with 2 ml saturated NaCl aqueous solution and stirred for 60 s. Although the brine was not provided in the predicted synthetic route, it was favored over water or an alcohol since it prevents the formation of a colloidal dispersion of aluminum hydroxide. The organic layer was collected and further analysed. 0.3 ml of the organic layer was diluted 50 times and then analyzed with an LC/MS (Agilent TOF6230). The spectrum is reported in supplementary figure S7.

3. Results

3.1. Protein-ligand affinity prediction

The results of the binding affinity prediction model on validation and test data from BindingDB are displayed in table 1. The results show that the model learned reasonably well to classify compound-protein-interaction samples as binding (positive class) or non-binding (negative class). Because the conditional generation focuses on antiviral drug design, we had to ensure that the affinity predictor generalizes well for viral proteins. We therefore additionally measured model performance on ~10k held-back samples from viral proteins and find reasonable generalization to viral proteins.

3.2. Toxicity predictor

Because toxicity is a major cause of the high attrition rate in drug discovery, we decided to perform a multi-objective optimization based on toxicity and binding affinity. Across ten runs on the Tox21 dataset, this model achieved a ROC-AUC of 0.877 ± 0.04 , in predicting toxic vs non-toxic in the 12 Tox21 assays (ROC-AUC is obtained by concatenating the predictions for each of the 12 classes) 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. For details see [33].

3.3. Conditional generative model

3.3.1. Workflow of RL optimization

The following procedure of conditional generative design is similar to the one described in [21] with the difference that here, the biomolecular context is a protein target instead of an omic profile. During the deep RL optimization, the agent G_{Θ} , a hybrid-VAE receives as input a SARS-CoV-2-related protein, such as 3C-like proteases. First, the protein $p \in \mathcal{P}$ is encoded into the latent space of proteins using $\Theta_{\text{Prot}}^{\text{Enc}} : \mathcal{P} \rightarrow \mathcal{Z}$. Next the molecular decoder generates a compound using $\Theta_{\text{Mol}}^{\text{Dec}} : \mathcal{Z} \rightarrow \mathcal{M}$. This is a valid operation due to the variational constraint in the bottleneck layer of both VAEs, as imposed by the Kullback-Leibler (KL) divergence in the ELBO of Θ_{Mol} and Θ_{Prot} . In other words, the independent pre-training of the protein VAE as well as the SELFIES VAE steered both unimodal models to encode their data points (i.e. proteins and molecules) as samples of a multivariate Gaussian distribution. This observation is key for the construction of our hybrid-VAE consisting of a protein-encoder and a molecular decoder. Therefore, during the RL optimization, initial decodings will thus still be valid molecules. Finally, the compound-protein pair is evaluated using the reward function $R : \mathcal{M} \times \mathcal{P} \rightarrow \mathbb{R}$. Over time, the stochastic optimization led to generating more compounds with high predicted binding affinities, as shown in table 2.

3.3.2. Results of LooCV

In this study, we aim to validate whether our conditional generative framework can go beyond current approaches for target-driven compound design [17, 18, 25] in the sense that it does not require explicit optimization for a specific target. We therefore investigated the generalization capabilities of our framework by performing a leave-one-out-cross-validation (LooCV) on the 41 targets. Prior to starting the RL optimization, we sampled 3000 molecules from the pre-trained SELFIES VAE and predicted binding affinities (Affinity_0) and toxicity scores (Tox_0) for all those molecules. This constitutes our baseline for later

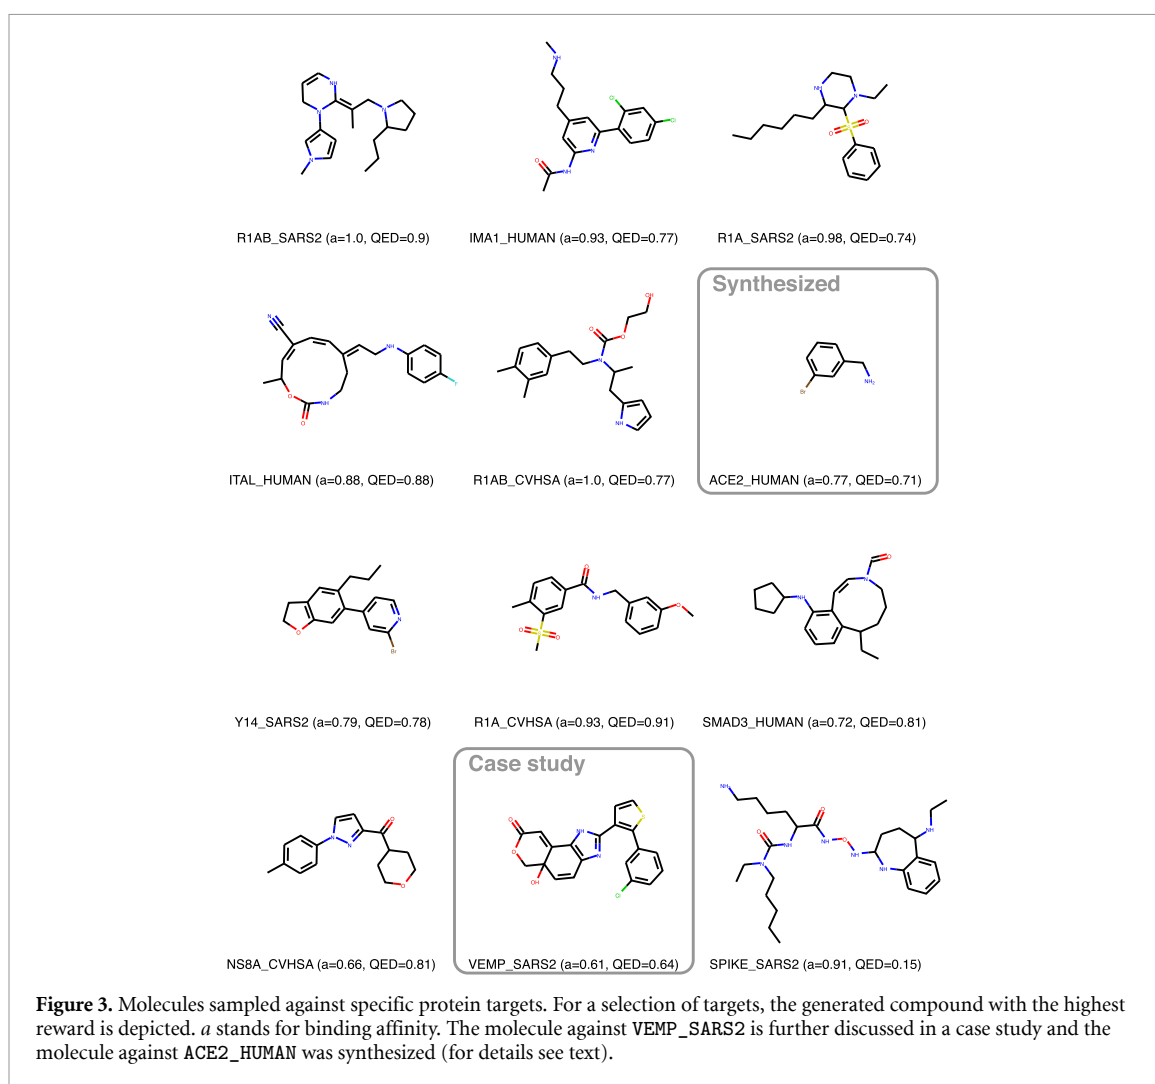
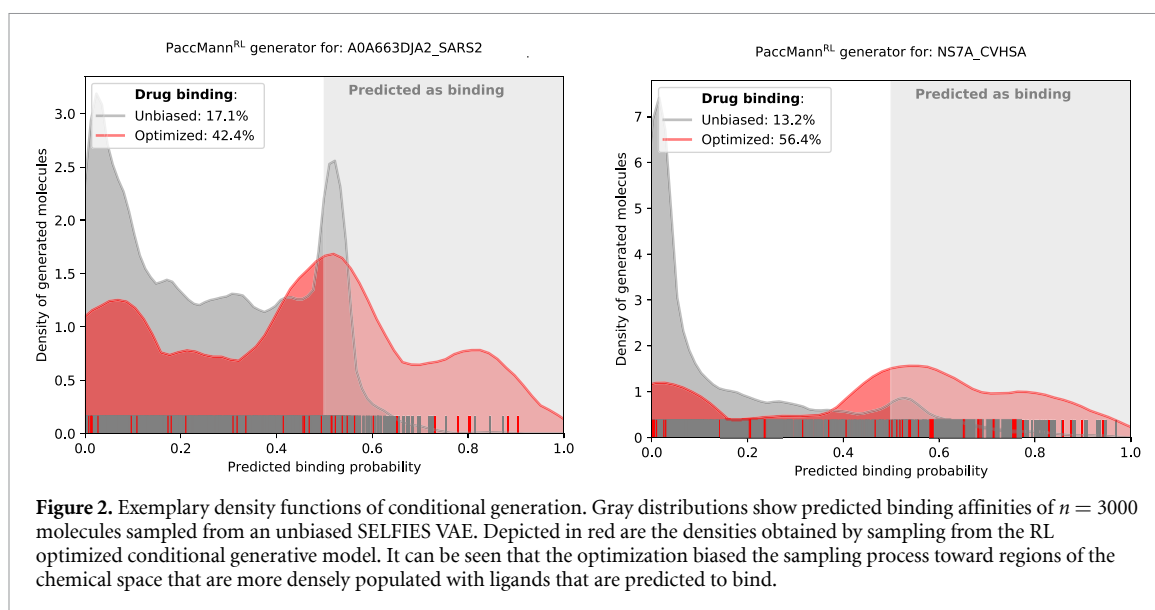
Table 2. Generating antiviral compounds against unseen SARS-CoV-2 targets. For each of the 41 targets, Aff_{0} shows the fraction of binding molecules sampled *before* training, Aff_{best} at the best epoch of RL training, and Aff_{median} the median across all five training epochs. The same applies to Tox_{best} and Tox_{med} . Note that Tox_0 is independent of the protein and was 8.7%. The best model is highlighted in bold for each target protein SEM abbreviates standard error of the mean.

Target protein	Aff_{0}	$Aff_{med} \pm SEM$	Aff_b	$Tox_{med} \pm SED$	Tox_b
VME1-CVHSA	20%	18% \pm 3%	29%	6% \pm 3%	19%
IMA1-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%

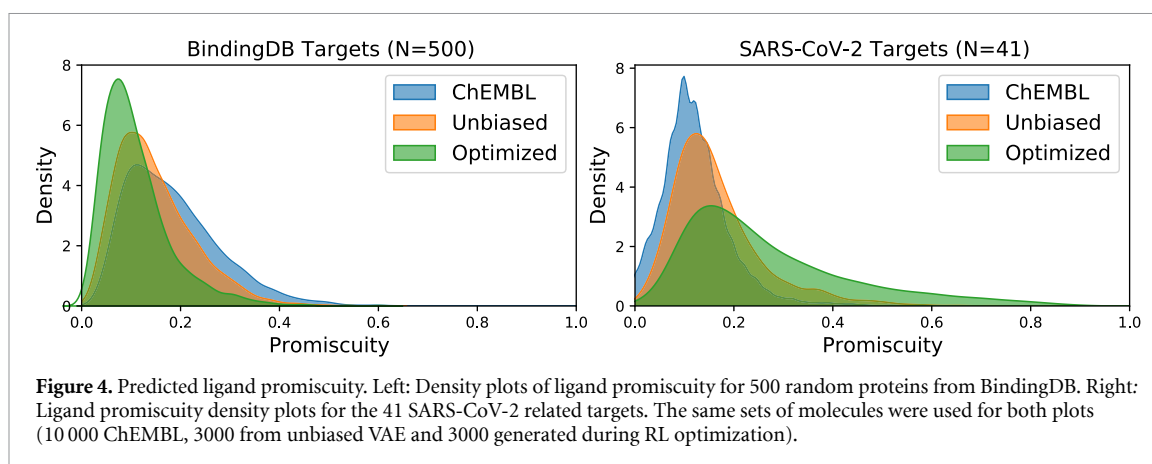
comparison. Next, the RL optimization was performed for 5 epochs and 500 molecules were sampled in each epoch. In table 2 we report the percentage of molecules predicted to bind and to be toxic for each of the 41 targets obtained by 41 runs in a LooCV. The results demonstrate that in 35 out of 41 cases the model proposed more binding compounds against an unseen target, compared to the baseline SELFIES VAE. The average ratio of compounds predicted to bind increased from 18% to 26% with the best epoch averaging 33% across all targets. Example density plots for 2 out of the 41 individual optimizations are shown in figure 2. 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 3 shows a selection of the sampled molecules alongside their QED score [47].

3.3.3. Learned chemical space

Knowing that text-based deep molecular generative models can memorize large fractions of the chemical space [48], we sought to investigate the learned chemical space, assembled by a dataset of 10 000 random ChEMBL compounds, 3000 molecules sampled from the unbiased VAE, 3000 molecules sampled during the



RL optimisation and 82 SARS-CoV-2 candidate drugs from the literature (top 15 matches on PubChem and 69 compounds identified via protein-interaction-maps [6], excluding two duplicates). For all these molecules, binding affinities were computed alongside other pharmacological properties. Next, a UMAP [49] was performed on the ECFP4 fingerprints, a well-established index for molecular featurization, of all



molecules [50] and visualized with Faerun/Tmap [51, 52]. The interactive visualisation (available online⁷) shows that the RL optimisation concentrates the compound sampling on a manifold of the chemical space that is more densely populated with binding compounds (for a snapshot see supplementary figure S4). The 3D UMAP shows that the currently investigated candidate molecules (red) are structurally fairly dissimilar, i.e. widely scattered across 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 of generating more ligands with high affinities we are aware that the quality of the reward function remains a bottleneck of the framework.

3.3.4. Target selectivity

Small molecules designed to bind to a specific target can be expected to bind to at least a dozen different targets [53]. This promiscuity of ligands is an important challenge in targeted drug design as it can induce side effects such as off-target cytotoxicity or lowered efficacy. Accounting for target selectivity has been proposed as a means to reduce attrition rates in downstream clinical trials [54]. To assess the selectivity of the generated molecules we computed a promiscuity score $P_{m,T}$ for each molecule m and a set of targets T by measuring the fraction of targets to which the molecule is predicted to bind. The same set of molecules as described above were used and the results are shown in figure 4. Interestingly, the promiscuity of the molecules generated during RL optimization was significantly lower than the promiscuity of ChEMBL molecules as well as molecules from the unbiased VAE (Tukey's HSD test, $p < 0.001$ in all pairwise differences, mean for ChEMBL: 0.19, mean for Optimized: 0.11, see figure 4(A)). Note that this was achieved without explicitly penalizing high promiscuity to unrelated targets. Additionally, investigating promiscuity within the set of 41 examined SARS-CoV-2 related targets reveals that promiscuity is significantly *higher* for the optimized molecules compared to the other two sets (Tukey's HSD test, $p < 0.001$ in all pairwise differences, mean for ChEMBL: 0.12, mean for Optimized: 0.27, see figure 4(B)). While further work to improve selectivity can certainly be done, these results indicate that our conditionally generated molecules are not only less prone to off-target binding effects but also more likely to bind to related SARS-CoV-2 targets than the other sets of molecules.

3.4. Case study

For a more detailed assessment of the quality of the molecules, we ranked all ~3000 conditionally generated molecules by their Tanimoto similarity $\tau(F_1, F_2) = \frac{|F_1 \cap F_2|}{|F_1 \cup F_2|}$ (where F_1 and F_2 are binary fingerprints of molecules $m_1, m_2 \in \mathcal{M}$) to the closest neighbour of the 82 literature candidates. Among the top five molecules, we found the molecule encircled in figure 3 generated against VEMP_{SARS2} (UniProt ID: PODTC4), the envelope small membrane protein (E-protein), a key player for virion assembly and morphogenesis. From all 82 literature candidates, our molecule exhibits the highest Tanimoto similarity to the compounds MZ1 and dBET6 ($\tau = 0.64$ based on RDKit fingerprint). Notably, these two pre-clinical SARS-CoV-2 drug candidates were identified by Gordon *et al* [6] as targeting the E-protein, exactly the protein which was used to condition the generation. MZ1 and dBET6 target E-Protein by degrading the human BRD2 and BRD4 proteins and thus preventing the virus from inducing changes in the host's protein expression.

⁷ The Faerun visualization of the ECFPs is available at: https://pacmann.github.io/assets/umap_fingerprints.html.

3.5. Retrosynthesis

The top-5 candidate compounds for each protein target were further analyzed for synthetic feasibility using IBM RXN's retrosynthesis engine [26]. We performed the predictions using the Python package `rxn4chemistry`⁸. We used the default settings for the hypertree exploration: limiting the maximum number of synthesis steps to 6 (to obtain a reasonable yield), applying a confidence threshold on the forward model score of 0.6 (FAP, forward acceptance probability) and setting the number of beams for the search to 10 (more details can be found in the documentation⁹).

Although the generated molecules are not optimized for synthetic accessibility, more than half of the >2000 predicted synthesis routes are feasible. Overall, for 29% of the top-5 candidates per target, a synthetic route could be successfully predicted with at most six synthesis steps. Interestingly, almost half of the successfully predicted molecules only require a single or two step reactions, indicating that many of the generated molecules can be synthesized from commercially available materials in a few steps. Moreover, a correlation analysis between chemical and pharmacological properties indicates that some properties like QED and synthetic feasibility are highly correlated (for details see supplementary figure S5).

3.6. Selection of synthesis candidates

3.6.1. Selection of ACE2 target

The selected target for the first synthesis was ACE2, a host protein that is widely regarded a promising target for SARS-CoV-2 antiviral drug design [55–57] and was even argued a priority role [58]. Despite this encouraging evidence in favour of ACE2, studies using generative models against SARS-CoV-2 host targets are almost absent; with one exception [59]. For details on docking and drug repurposing studies on ACE2 for COVID-19 see S1.6. We aim to fill this gap and here exemplify the process of generating and synthesizing ligands predicted to bind to host targets.

3.6.1.1. Role of ACE2

ACE2 is a type 1 membrane protein that regulates the renin-angiotensin-aldosterone system (RAAS) and is predominantly expressed in lung alveolar epithelial and endothelial cells [60–62]. It plays an important role in regulating cardiovascular homeostasis [63, 64], inhibition of cell growth, and protection from alveolar epithelial cell injury [65–67]. ACE2 has previously been identified as a functional receptor for SARS-CoV to mediate cell entry by its spike protein [68, 69]. SARS-CoV-2 also utilizes ACE2 as a receptor [70, 71], specifically by a fusion of ACE2 with the densely glycosylated spike (S) proteins [70], but with an increased binding affinity due to modifications around the centre of the binding domain [72, 73]. Owing to the importance of the S-protein in viral cell entry and fusion, developing ACE2 inhibiting drugs seems like a logical step in tackling SARS-CoV-2. The plausibility of this approach was highlighted using a recombinant RBD protein to prevent SARS-CoV-2's RBD from binding to ACE2's peptidase domain [74, 75]. With strong evidence of the role played by ACE2 in viral infection, much of the focus is on finding drugs that can prevent the S-protein from interacting with this enzyme.

3.6.2. Selection of molecule

As a demonstration of the fully autonomous pipeline for swift discovery and synthesis of molecules generated against unseen SARS-CoV-2-related protein targets, we synthesized 3-Bromobenzylamine, a molecule proposed by our generative model against ACE2. Note that on rare occasions, our generative model proposes molecules that already exist in chemical databases and as such are not *de-novo*. This was the case for 3-Bromobenzylamine, a bioactive compound which has been crystalized with the histidyl-RNA synthetase of *T. cruzi* in PDB 4YRI and is known to inhibit neurological enzymes like PNMT [76].

3-Bromobenzylamine was selected anyway for several other reasons. First, a maximum common subgraph similarity search [77] within the set of 82 COVID-19 literature candidates revealed that 3-Bromobenzylamine is a full substructure of Arbidol (Umifenovir), a broad-spectrum antiviral drug used in Asia against influenza and hepatitis [78, 79]. Notably, Arbidol was proposed as an antiviral drug for COVID-19 specifically for its interaction with the ACE2 receptor [80], exactly the target against which 3-Bromobenzylamine was generated. It was first hypothesized that Arbidol may act as a virus host fusion inhibitor for SARS-CoV-2 [80, 81] (thus preventing viral entry to the target cell, just like in influenza and hepatitis viruses [82, 83]) and later experimentally confirmed in docking studies [84, 85] as well as *in-vitro* [86]. An antiviral effect of Arbidol on SARS-CoV has been known for many years [87] and while COVID-19 studies are not yet fully conclusive [88], Arbidol was found effective in decreasing mortality [89]

⁸ <https://github.com/rxn4chemistry/rxn4chemistry>.

⁹ https://rxn4chemistry.github.io/rxn4chemistry/_modules/rxn4chemistry.html#rxn4chemistry.core.RXN4ChemistryWrapper.predict_automatic_retrosynthesis.

and in increasing negative rate of PCR [90–92]. We therefore hypothesized that 3-Bromobenzylamine as a smaller and broadly available compound, could operate in a similar mechanism of action by interacting with ACE2, especially since the presence of bromine was found important for the efficacy of Arbidol [93].

Secondly, the biochemical properties predicted *in-silico* were desirable, with a predicted ACE2 affinity of 0.77, a high drug-likeness (QED = 0.71), a penalized LogP of 0.25, the presence of an aromatic ring, an estimated solubility (ESOL [94]) of -2.66 , a relatively low promiscuity to the remaining protein targets (0.13) and a relatively high promiscuity for the other SARS-CoV-2 related targets (0.27). The molecule has a molecular weight of 186 Dalton, passes the Lipinski rule of five and was predicted to be non-toxic in all but one Tox21 assays (NR-AR-LBD, the androgen receptor ligand-binding domain). Thirdly, the retrosynthetic route (shown in supplementary figure S6) was comparably simple (one reaction, eight reactants) and predicted with high confidence (98.5%) by the retrosynthesis model [26]. Although the reaction itself, a nitrile reduction, reducing 3-Bromobenzonitrile with lithium aluminium hydride, is challenging, it is a known, well-understood reaction that minimizes the risk of complication thus making the target compound the ideal candidate to demonstrate the validity of the end-to-end concept.

3.7. Chemical synthesis

The LC/MS shows a clear signal related to the presence of the 3-Bromobenzylamine with a score of more than 99%. No signals connected to the precursors were identified by the target screening analysis. Although it is not possible to use the qualitative analysis for quantitative arguments, the lack of evidences pointing to the presence of the precursor in the LC/MS analysis corroborates the indications that the automatically suggested synthesis route successfully completed in a quantitative way (report in supplementary figure S7).

4. Discussion

4.1. Summary

Here, we proposed a novel framework for compound design that can be targeted towards *any target protein* without retraining requirements. We showcased the potential of our generative framework by tackling the problem of designing novel antiviral candidates with high binding affinity to unseen SARS-CoV-2 related targets, while controlling toxicity of the generated molecules. Without explicitly accounting for target selectivity, the generated molecules exhibited comparably low promiscuity within a random set of protein targets but comparably high promiscuity within the set of SARS-CoV-2 related targets. A future endeavour could be to optimize scaffolds of existing drugs against specific targets by coupling our approach with the recently proposed deep scaffold generator [95] or to directly optimize binding scores. Furthermore, to assess the feasibility of synthesizing the generated compounds, we estimated the retrosynthetic pathways of a subset of candidates for each target. We exemplified how our framework for target-driven *de novo* discovery of molecules with potential antiviral properties can be coupled with an AI-assisted synthesis platform for swift chemical synthesis. One molecule was selected for further testing and successfully synthesized using an automatically derived synthesis route. This molecule, 3-Bromobenzylamine, is a substructure of Arbidol, a broad-spectrum antiviral drug [79] with known efficacy in combatting COVID-19 [89]. We are aware that our synthesized molecule is commercially available, relatively small, and that the true biochemical activity remains unclear at this point. However, the ultimate goal of our contribution is to showcase a generic pipeline for rapid generation and chemical synthesis of molecules with desired properties, in this case, potential binding to SARS-CoV-2 related target proteins.

4.2. Limitations

We are aware that the true bioactivity of the proposed molecules can only be consolidated by *in vitro* and *in vivo* experiments. Another bottleneck in our pipeline is the accuracy of the predictive models such as the affinity predictor, which could be further improved using recently available data from a large-scale screening consisting of 1670 compounds tested against SARS-CoV-2 proteins [96]. On the one hand, it is advantageous that our model operates only on protein primary structures and thus does not necessitate the availability of 3D protein structure, a key limitation especially in situations of new infectious diseases such as COVID-19 where it takes months before tertiary structure of key proteins becomes available [97]. On the other hand, docking simulations could be beneficial to collect further evidence that the molecules indeed bind to the target structure prior to synthesis, particularly for target-site specific modeling which has not been tackled directly in here. It is therefore conceivable that some of the generated ligands bind to unintended drug binding pockets that do not or only mildly allow the drug to exert its mechanism of action. Predicting binding affinities for specific sites and conditioning the generative model on specific pockets can in principle be achieved given the pocket is well defined in terms of primary structure. However, related studies with deep generative models on SARS-CoV-2 have shown that primary structure can suffice to generate molecules with

favorable binding free energies that reliably identify the druggable binding pocket within the 3D protein structure [25].

4.3. Conclusion

A key advantage of our proposed method is that it can generate ligands with high predicted binding affinities for entirely novel protein sequences, as we found in the leave-one-out-cross-validation on 41 SARS-CoV-2 related proteins. To the best of our knowledge, this is the first report of a conditional molecular generator that generalizes to unseen protein sequences. Another key novelty of the proposed system is the integration of molecular generative models with retrosynthesis and synthesis protocol prediction models. This enables the generation of *de novo* molecules and the compilation of the synthesis's experimental procedures without human intervention. Conclusively, we hope that our framework can be a building block towards the swift discovery of *de novo* molecules with desired pharmacological properties.

Code & data availability

The data that support the findings of this study are openly available at: <https://ibm.ent.box.com/v/paccmann-sarscov2>. Under the link, we provide processed versions of all datasets used in the study as well as pretrained models. The following datasets were used: ChEMBL [98]; SELFIES VAE; UniProt [37]; Protein VAE; BindingDB [34]; Protein-ligand affinity prediction; Tox21 [30]; toxicity prediction; SARS-CoV-2 targets [6]; conditional generation. All models are implemented in PyTorch 1.3.1 and training was done on POWER8 processors and an NVIDIA Tesla P100. The source code for training the models presented in this submission is publicly available at: https://github.com/PaccMann/paccmann_sarscov2.



Acknowledgment

The authors would like to thank Dr Maria Gabrani for her continuous support and acknowledge funding from the European Union's Horizon 2020 research and innovation programme under Grant No. 826121. The authors acknowledge the useful feedback from anonymous reviewers that helped to improve this article.

Author Contributions

M M, J B and M R M designed the study. M M supervised the model development by J B, J C and G M. J B developed the conditional generator and J B and J C analyzed the results. M M performed the retrosynthesis analysis and M M, A C and T L selected the molecule and oversaw the synthesis. J B drafted the initial version of the manuscript to which all authors contributed substantially.

ORCID iDs

Jannis Born  <https://orcid.org/0000-0001-8307-5670>
Matteo Manica  <https://orcid.org/0000-0002-8872-0269>
Joris Cadow  <https://orcid.org/0000-0002-4410-2805>
Greta Markert  <https://orcid.org/0000-0001-5254-5596>
Nil Adell Mill  <https://orcid.org/0000-0003-0676-7547>
Nikita Janakarajan  <https://orcid.org/0000-0001-7886-8385>
Teodoro Laino  <https://orcid.org/0000-0001-8717-0456>
María Rodríguez Martínez  <https://orcid.org/0000-0003-3766-4233>

References

- [1] Drosten C *et al* 2003 Identification of a novel coronavirus in patients with severe acute respiratory syndrome *New Engl. J. Med.* **348** 1967–76
- [2] Shamshirian A, Hessami A, Heydari K, and Alizadeh-Navaei R *et al* 2020 Hydroxychloroquine versus COVID-19: a periodic systematic review and meta-analysis *medRxiv*
- [3] Beigel J H *et al* 2020 Remdesivir for the treatment of COVID-19—Final report *New Engl. J. Med.* (<https://doi.org/10.1056/NEJMoa2007764>)
- [4] Wang Y *et al* 2020 Remdesivir in adults with severe COVID-19: a randomised, double-blind, placebo-controlled, multicentre trial *Lancet* **395** 1569–78
- [5] Lamb Y N 2020 Remdesivir: first approval *Drugs* 1–9
- [6] Gordon D E *et al* 2020 A SARS-CoV-2 protein interaction map reveals targets for drug repurposing *Nature* 1–13
- [7] Polishchuk P G, Madzhidov T I and Varnek A 2013 Estimation of the size of drug-like chemical space based on GDB-17 data *J. Comput.-Aided Mol. Des.* **27** 675–9

- [8] Scannell J W, Blanckley A, Boldon H and Warrington B 2012 Diagnosing the decline in pharmaceutical R&D efficiency *Nat. Rev. Drug Discovery* **11** 191
- [9] Schneider G 2019 Mind and machine in drug design *Nat. Mach. Intell.* 1
- [10] Kinch M S, Haynesworth A, Kinch S L and Hoyer D 2014 An overview of FDA-approved new molecular entities: 1827–2013 *Drug Discovery Today* **19** 1033–9
- [11] Kim S et al 2016 PubChem substance and compound databases *Nucleic Acids Res.* **44** D1202–13
- [12] Karimi M, Wu D, Wang Z and Shen Y 2019 DeepAffinity: interpretable deep learning of compound–protein affinity through unified recurrent and convolutional neural networks *Bioinformatics* **35** 3329–38
- [13] Gomez-Bombarelli R et al 2018 Automatic chemical design using a data-driven continuous representation of molecules *ACS Cent. Sci.* **4** 268–76
- [14] Blaschke T, Olivecrona M, Engkvist O, Bajorath J and Chen H 2018 Application of generative autoencoder in *de novo* molecular design *Molecular Inform.* **37** 1700123
- [15] Popova M, Isayev O and Tropsha A 2018 Deep reinforcement learning for *de novo* drug design *Sci. Adv.* **4** eaa7885
- [16] Popova M, Shvets M, Oliva J, and Isayev O 2019 MolecularRNN: generating realistic molecular graphs with optimized properties (arXiv:1905.13372)
- [17] Zhavoronkov A et al 2019 Deep learning enables rapid identification of potent DDR1 kinase inhibitors *Nat. Biotechnol.* **37** 1038–40
- [18] Zhavoronkov A et al 2020 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*
- [19] Tang B, He F, Liu D, Fang M, Wu Z, and Xu D 2020 AI-aided design of novel targeted covalent inhibitors against SARS-CoV-2 *bioRxiv*
- [20] Bung N, Krishnan S R, Bulusu G and Roy A 2021 *De novo* design of new chemical entities for SARS-CoV-2 using artificial intelligence *Future Medicinal Chemistry* (<https://doi.org/10.4155/fmc-2020-0262>)
- [21] Born J, Manica M, Oskooei A, Cadow J and Martínez M R 2020 PaccMann RL: designing anticancer drugs from transcriptomic data via reinforcement learning *Int. Conf. on Research in Computational Molecular Biology* (Springer) pp 231–3
- [22] Aumentado-Armstrong T 2018 Latent molecular optimization for targeted therapeutic design (arXiv:1809.02032)
- [23] Krishnan S R, Bung N, Bulusu G and Roy A 2021 Accelerating *de novo* drug design against novel proteins using deep learning *J. Chem. Inf. Model.* **61** 621–30
- [24] Skalic M, Sabbadin D, Sattarov B, Sciabola S and De Fabritiis G 2019 From target to drug: generative modeling for the multimodal structure-based ligand design *Mol. Pharmaceutics* **16** 4282–91
- [25] Chenthamarakshan V et al 2020 Cogmol: target-specific and selective drug design for COVID-19 using deep generative models *Advances in Neural Information Processing Systems* 33
- [26] Schwaller P et al 2020 Predicting retrosynthetic pathways using transformer-based models and a hyper-graph exploration strategy *Chem. Sci.* (<https://doi.org/10.1039/C9SC05704H>)
- [27] Vaucher A C, Zipoli F, Gelykens J, Nair V H, Schwaller P and Laino T 2020 Automated extraction of chemical synthesis actions from experimental procedures *Nat. Commun.* (<https://doi.org/10.1038/s41467-020-17266-6>)
- [28] Kingma D P, and Welling M 2013 Auto-encoding variational bayes (arXiv:1312.6114)
- [29] Williams R J 1992 Simple statistical gradient-following algorithms for connectionist reinforcement learning *Mach. Learn.* **8** 229–56
- [30] Huang R, Xia M, Nguyen D-T, Zhao T, Sakamuru S, Zhao J, Shahane S A, Rossoshek A and Simeonov A 2016 Tox21Challenge to build predictive models of nuclear receptor and stress response pathways as mediated by exposure to environmental chemicals and drugs *Front. Environ. Sci.* **3** 85
- [31] Bjerrum E J 2017 Smiles enumeration as data augmentation for neural network modeling of molecules (arXiv:1703.07076)
- [32] Weininger D 1988 SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules *J. Chem. Inf. Comput. Sci.* **28** 31–6
- [33] Markert G, Born J, Manica M, Schneider G and Rodríguez Martínez M 2020 Chemical representation learning for toxicity prediction *PharML Workshop at ECML-PKDD (Conf. on Machine Learning and Principles and Practice of Knowledge Discovery in Databases)*
- [34] Gilson M K, Liu T, Baitaluk M, Nicola G, Hwang L and Chong J 2016 BindingDB in 2015: a public database for medicinal chemistry, computational chemistry and systems pharmacology *Nucleic Acids Res.* **44** D1045–53
- [35] Manica M, Oskooei A, Born J, Subramanian V, Sáez-Rodríguez J and Rodríguez Martínez M 2019 Toward explainable anticancer compound sensitivity prediction via multimodal attention-based convolutional encoders *Mol. Pharm.* **16** 4797–806
- [36] Cadow J, Born J, Manica M, Oskooei A and Rodríguez Martínez M 2020 PaccMann: a web service for interpretable anticancer compound sensitivity prediction *Nucleic Acids Res.* **48** W502–8
- [37] Consortium U 2019 UniProt: a worldwide hub of protein knowledge *Nucleic Acids Res.* **47** D506–15
- [38] Devlin J, Chang M-W, Lee K, and Toutanova K 2018 Bert: pre-training of deep bidirectional transformers for language understanding (arXiv:1810.04805)
- [39] Rao R, Bhattacharya N, Thomas N, Duan Y, Chen X, Canny J, Abbeel P and Song Y S 2019 Evaluating protein transfer learning with TAPE *Advances in Neural Information Processing Systems* pp 9686–98
- [40] El-Gebali S et al 2019 The Pfam protein families database in 2019 *Nucleic Acids Res.* **47** D427–32
- [41] Bowman S R, Vilnis L, Vinyals O, Dai A M, Jozefowicz R, and Bengio S 2015 Generating sentences from a continuous space (arXiv:1511.06349)
- [42] Krenn M, Häse F, Nigam A, Friederich P and Aspuru-Guzik A 2020 Self-referencing embedded strings (SELFIES): a 100% robust molecular string representation *Mach. Learn.: Sci. Technol.* **1** 045024
- [43] Wu C et al 2020 Analysis of therapeutic targets for SARS-CoV-2 and discovery of potential drugs by computational methods *Acta Pharm. Sin. B* (<https://doi.org/10.1016/j.apsb.2020.02.008>)
- [44] Khaerunnisa S, Kurniawan H, Awaluddin R, Suhartati S, and Soetjipto S 2020 Potential inhibitor of COVID-19 main protease (Mpro) from several medicinal plant compounds by molecular docking study *Preprint* pp 1–14
- [45] Schwaller P, Laino T, Gaudin T, Bolgar P, Hunter C A, Bekas C and Lee A A 2019 Molecular transformer: a model for uncertainty-calibrated chemical reaction prediction *ACS Cent. Sci.* **5** 1572–83
- [46] Vaucher A C, Schwaller P, Gelykens J, Nair V H, Iuliano A and Laino T 2020 Inferring Experimental Procedures from Text-Based Representations of Chemical Reactions
- [47] Bickerton G R, Paolini G V, Besnard J, Muresan S and Hopkins A L 2012 Quantifying the chemical beauty of drugs *Nat. Chem.* **4** 90
- [48] Arús-Pous J, Blaschke T, Ulander S, Reymond J-L, Chen H and Engkvist O 2019 Exploring the GDB-13 chemical space using deep generative models *J. Chem. Inf.* **11** 1–14

- [49] McInnes L, Healy J, and Melville J 2018 Umap: uniform manifold approximation and projection for dimension reduction (arXiv:1802.03426)
- [50] Rogers D and Hahn M 2010 Extended-connectivity fingerprints *J. Chem. Inf. Model.* **50** 742–54
- [51] Probst D and Reymond J-L 2018 FUn: a framework for interactive visualizations of large, high-dimensional datasets on the web *Bioinformatics* **34** 1433–5
- [52] Probst D and Reymond J-L 2020 Visualization of very large high-dimensional data sets as minimum spanning trees *J. Chem. Inf.* **12** 1–13
- [53] Peón A, Naulaerts S and Ballester P J 2017 Predicting the reliability of drug-target interaction predictions with maximum coverage of target space *Sci. Rep.* **7** 1–11
- [54] Miljković F and Bajorath J 2018 Data-driven exploration of selectivity and off-target activities of designated chemical probes *Molecules* **23** 2434
- [55] Li S-r, Tang Z-j, Li Z-h and Liu X 2020 Searching therapeutic strategy of new coronavirus pneumonia from angiotensin-converting enzyme 2: the target of COVID-19 and SARS-CoV *Eur. J. Clin. Microbiol. Inf. Dis.* **39** 1021
- [56] Zhang H, Penninger J M, Li Y, Zhong N and Slutsky A S 2020 Angiotensin-converting enzyme 2 (ACE2) as a SARS-CoV-2 receptor: molecular mechanisms and potential therapeutic target *Intensive Care Med.* **46** 586–90
- [57] McKee D L, Sternberg A, Stange U, Laufer S and Naujokat C 2020 Candidate drugs against SARS-CoV-2 and COVID-19 *Pharm. Res.* **104859**
- [58] Terah K, Baddal B and Gülcan H O 2020 Prioritizing potential ACE2 inhibitors in the COVID-19 pandemic: insights from a molecular mechanics-assisted structure-based virtual screening experiment *J. Mol. Graph. Model.* **107697**
- [59] Ray S, Lall S, Mukhopadhyay A, Bandyopadhyay S, and Schönhuth A 2020 Predicting potential drug targets and repurposable drugs for COVID-19 via a deep generative model for graphs (arXiv:2007.02338)
- [60] Hamming I, Timens W, Bulthuis M, Lely A, Navis G v and van Goor H 2004 Tissue distribution of ACE2 protein, the functional receptor for SARS coronavirus. A first step in understanding SARS pathogenesis *J. Pathol.* **203** 631–7
- [61] Donoghue M et al 2000 A novel angiotensin-converting enzyme-related carboxypeptidase (ACE2) converts angiotensin I to angiotensin 1-9 *Circ. Res.* **87** e1–9
- [62] Tipnis S R, Hooper N M, Hyde R, Karran E, Christie G and Turner A J 2000 A human homolog of angiotensin-converting enzyme cloning and functional expression as a captopril-insensitive carboxypeptidase *J. Biol. Chem.* **275** 33238–43
- [63] Oudit G Y, Crackower M A, Backx P H and Penninger J M 2003 The role of ACE2 in cardiovascular physiology *Trends Cardiovasc. Med.* **13** 93–101
- [64] Crackower M A et al 2002 Angiotensin-converting enzyme 2 is an essential regulator of heart function *Nature* **417** 822–8
- [65] le Tran Y and Forster C 1997 Angiotensin-(1-7) and the rat aorta: modulation by the endothelium *J. Cardiovasc. Pharmacol.* **30** 676–82
- [66] Schindler C, Bramlage P, Kirch W and Ferrario C M 2007 Role of the vasodilator peptide angiotensin-(1-7) in cardiovascular drug therapy *Vasc. Health Risk Manag.* **3** 125
- [67] Li X, Molina-Molina M, Abdul-Hafez A, Uhal V, Xaubet A and Uhal B D 2008 Angiotensin converting enzyme-2 is protective but downregulated in human and experimental lung fibrosis *Am. J. Physiol. Lung Cell Mol. Physiol.* **295** L178–85
- [68] Li W et al 2003 Angiotensin-converting enzyme 2 is a functional receptor for the SARS coronavirus *Nature* **426** 450–4
- [69] Li F, Li W, Farzan M and Harrison S C 2005 Structure of SARS coronavirus spike receptor-binding domain complexed with receptor *Science* **309** 1864–8
- [70] Hoffmann M et al 2020 SARS-CoV-2 cell entry depends on ACE2 and TMPRSS2 and is blocked by a clinically proven protease inhibitor *Cell* (<https://doi.org/10.1016/j.cell.2020.02.052>)
- [71] Zhou P et al 2020 A pneumonia outbreak associated with a new coronavirus of probable bat origin *Nature* 1–4
- [72] Wang Q et al 2020 Structural and functional basis of SARS-CoV-2 entry by using human ACE2 *Cell*
- [73] Yan R, Zhang Y, Li Y, Xia L, Guo Y and Zhou Q 2020 Structural basis for the recognition of SARS-CoV-2 by full-length human ACE2 *Science* **367** 1444–8
- [74] Tai W, He L, Zhang X, Pu J, Voronin D, Jiang S, Zhou Y and Du L 2020 Characterization of the receptor-binding domain (RBD) of 2019 novel coronavirus: implication for development of RBD protein as a viral attachment inhibitor and vaccine *Cell. Mol. Immunol.* 1–8
- [75] Chen W-H, Hotez P J and Bottazzi M E 2020 Potential for developing a SARS-CoV receptor-binding domain (RBD) recombinant protein as a heterologous human vaccine against coronavirus infectious disease (COVID)-19 *Hum. Vaccines Immunother.* **1–4**
- [76] Grunewald G, Sall D and Monn J 1988 Conformational and steric aspects of the inhibition of phenylethanolamine N-methyltransferase by benzylamines *J. Med. Chem.* **31** 433–44
- [77] Cao Y, Jiang T and Girke T 2008 A maximum common substructure-based algorithm for searching and predicting drug-like compounds *Bioinformatics* **24** i366–74
- [78] Pshenichnaya N Y, Bulgakova V, Lvov N, Poromov A, Selkova E, Grekova A, Shestakova I, Maleev V and Leneva I 2019 Clinical efficacy of umifenovir in influenza and ARVI (study ARBITR) *Ther. Arch.* **91** 56–63
- [79] Boriskin Y S, Pêcheur E-I and Polyak S J 2006 Arbidol: a broad-spectrum antiviral that inhibits acute and chronic HCV infection *Viol. J.* **3** 56
- [80] Liu C et al 2020 Research and development on therapeutic agents and vaccines for COVID-19 and related human coronavirus diseases *ACS Cent. Sci.* **6** 315–31
- [81] Mascolo A, Scavone C, Rafaniello C, Ferrajolo C, Racagni G, Berrino L, Paolisso G, Rossi F and Capuano A 2020 Renin-angiotensin system and coronavirus disease 2019: a narrative review *Front. Cardiovascular Med.* **7**
- [82] Boriskin Y, Leneva I, Pêcheur E-I and Polyak S 2008 Arbidol: a broad-spectrum antiviral compound that blocks viral fusion *Curr. Med. Chem.* **15** 997–1005
- [83] Kadam R U and Wilson I A 2017 Structural basis of influenza virus fusion inhibition by the antiviral drug Arbidol *Proc. Natl Acad. Sci.* **114** 206–14
- [84] Choudhary S and Silakari O 2020 Scaffold morphing of arbidol (umifenovir) in search of multi-targeting therapy halting the interaction of SARS-CoV-2 with ACE2 and other proteases involved in COVID-19 *Virus Res.* **289** 198146
- [85] Padhi A, Seal A and Tripathi T 2020 How does arbidol inhibit the novel coronavirus SARS-CoV-2? Atomistic insights from molecular dynamics simulations
- [86] Zhao H et al 2020 Cross-linking peptide and repurposed drugs inhibit both entry pathways of SARS-CoV-2 *Nat. Res.* submitted
- [87] Khamitov R, Loginova S, Shchukina V, Borisevich S, Maksimov V and Shuster A 2008 Antiviral activity of arbidol and its derivatives against the pathogen of severe acute respiratory syndrome in the cell cultures *Vopr. Virusol.* **53** 9–13

- [88] Jun C *et al* 2020 Efficacies of lopinavir/ritonavir and abidol in the treatment of novel coronavirus pneumonia *Chin. J. Inf. Dis.* E008
- [89] Wang Z, Yang B, Li Q, Wen L and Zhang R 2020 Clinical features of 69 cases with coronavirus disease 2019 in Wuhan, China *Clin. Inf. Dis.* (<https://doi.org/10.1093/cid/ciaa538>)
- [90] Leneva I A and Pshenichnaya B V N.Y.2 2020 Umifenovir and coronavirus infections: a review of research results and clinical practice *Ther. Arch.* **11** 5
- [91] Deng L, Li C, Zeng Q, Liu X, Li X, Zhang H, Hong Z and Xia J 2020 Arbidol combined with LPV/r versus LPV/r alone against corona virus disease 2019: a retrospective cohort study *J. Inf.* (<https://doi.org/10.1016/j.jinf.2020.03.002>)
- [92] Huang D, Yu H, Wang T, Yang H, Yao R and Liang Z 2020 Efficacy and safety of umifenovir for coronavirus disease 2019 (COVID-19): a systematic review and meta-analysis *J. Med. Virol.*
- [93] Di Mola A, Peduto A, La Gatta A, Delang L, Pastorino B, Neyts J, Leyssen P, de Rosa M and Filosa R 2014 Structure–activity relationship study of arbidol derivatives as inhibitors of chikungunya virus replication *Bioorg. Med. Chem.* **22** 6014–25
- [94] Delaney J S 2004 ESOL: estimating aqueous solubility directly from molecular structure *J. Chem. Inf. Comput. Sci.* **44** 1000–5
- [95] Arús-Pous J, Patronov A, Bjerrum E J, Tyrchan C, Reymond J-L, Chen H and Engkvist O 2020 SMILES-based deep generative scaffold decorator for *de-novo* drug design *J. Chem. Inf.* **12** 1–18
- [96] Heiser K *et al* 2020 Identification of potential treatments for COVID-19 through artificial intelligence-enabled phenomic analysis of human cells infected with SARS-CoV-2 *bioRxiv*
- [97] Lan J *et al* 2020 Structure of the SARS-CoV-2 spike receptor-binding domain bound to the ACE2 receptor *Nature* **581** 215–20
- [98] Bento A P *et al* 2013 The ChEMBL bioactivity database: an update *Nucleic Acids Res.* **42** D1083–90