Active Learning-Guided Seq2Seq Variational Autoencoder for Multi-target Inhibitor Generation

Júlia Vilalta-Mor [1 2] **Alexis Molina** [3] **Laura Ortega Varga** [4 5] **Isaac Filella-Merce** [1] **Victor Guallar** [1 3 6]

Abstract

Simultaneously optimizing molecules against multiple therapeutic targets remains a profound challenge in drug discovery, particularly due to sparse rewards and conflicting design constraints. We propose a structured active learning (AL) paradigm integrating a sequence-to-sequence (Seq2Seq) variational autoencoder (VAE) into iterative loops designed to balance chemical diversity, molecular quality, and multi-target affinity. Our method alternates between expanding chemically feasible regions of latent space and progressively constraining molecules based on increasingly stringent multi-target docking thresholds. In a proof-of-concept study targeting three related coronavirus main proteases (SARS-CoV-2, SARS-CoV, MERS-CoV), our approach efficiently generated a structurally diverse set of pan-inhibitor candidates. We demonstrate that careful timing and strategic placement of chemical filters within this active learning pipeline markedly enhance exploration of beneficial chemical space, transforming the sparse-reward, multi-objective drug design problem into an accessible computational task. Our framework thus provides a generalizable roadmap for efficiently navigating complex polypharmacological landscapes.

[1]Barcelona Supercomputing Center (BSC), Plaça d'Eusebi Güell, 1-3, 08034, Barcelona, Spain [2]PhD program in Biotechnology, Faculty of Pharmacy and Food Sciences, University of Barcelona (UB), 08028 Barcelona, Spain [3]Nostrum Biodiscovery S.L., Av. de Josep Tarradellas, 8-10, 3-2, 08029, Barcelona, Spain [4]Alzheimer's Research UK Oxford Drug Discovery Institute, Centre for Medicines Discovery, Nuffield Department of Medicine, Oxford, UK [5]Centre for Medicines Discovery, Nuffield Department of Medicine Research Building (NDMRB), University of Oxford, Oxford, UK [6]ICREA, Pg. Lluis Companys 23, 08010, Barcelona, Spain. Correspondence to: Victor Guallar <victor.guallar@bsc.es>, Isaac Filella-Merce <ifilell1@bsc.es>.

Proceedings of the Workshop on Generative AI for Biology at the 42nd International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025.

1. Introduction

Generative models (GMs), a machine learning approach widely applied in fields such as text generation, can be effectively used in drug discovery by treating molecular representations (e.g., SMILES strings) as a chemical language (Ahmad et al., 2024). These models enable the design of truly novel molecules by navigating unexplored regions of chemical space that are inaccessible to traditional screening methods reliant on predefined molecular libraries (Tang et al., 2024). This is particularly evident in ultra-large chemical libraries (Tingle et al., 2023; Enamine, 2024), which, despite containing tens of billions of compounds, often lack diversity due to their construction from repeated combinations of limited sets of building blocks and chemical reactions (Neumann & Klein, 2025). In contrast, GMs learn underlying chemical patterns from data and generate new molecular structures beyond the scope of existing libraries (Zhavoronkov et al., 2019; Grisoni et al., 2021; Korshunova et al., 2022; Swanson et al., 2024).

Among GMs, Variational Autoencoders (VAEs) have shown particular promise. Their minimal backbone consists of an encoder, which compresses data into a lower-dimensional latent space, and a decoder, which reconstructs the original data from this representation. Sequence-to-Sequence (Seq2Seq) VAEs are especially effective in learning latent representations of molecules, enabling a controlled sampling of new molecules with optimised physicochemical and pharmacological properties (Gómez-Bombarelli et al., 2018; Filella-Merce et al., 2023). Building on this directed generation, the integration of methods such as reinforcement learning (RL) (Sheikholeslami et al., 2025) and active learning (AL) into GMs workflows has further improved their ability to guide molecular generation toward specific objectives. For instance, by iteratively guiding the model with feedback from molecular modelling simulations such as docking, these workflows can prioritise molecules. Specifically, those with enhanced affinity towards a given target (Filella-Merce et al., 2023).

GM workflows have primarily focused on generating molecules with affinity for single targets. In recent years, increasing attention has been given to multi-target drug discovery, *i.e.* molecules with simultaneous affinity to multiple

targets. Molecular generation aimed at multi-target inhibition could pave the way for the development of polypharmacological drugs, offering a powerful new approach to treating complex diseases such as cancer (Cichońska et al., 2024; Isigkeit et al., 2024). Similarly, multi-target generation also aligns with the concept of pan-inhibitors, where a single molecule exerts therapeutic effects across different organisms by targeting homologous proteins (Shahhamzehei et al., 2022; Huang et al., 2023). Liu et al. (2021) and Munson et al. (2024) explored the use of RL to construct a multi-target GM. Their goal was to steer the generation toward predefined desirable properties, by rewarding generated molecules that successfully meet the specified criteria. However, these approaches, like many RL-based methods (Olivecrona et al., 2017; Blaschke et al., 2020; Sheikholeslami et al., 2025; Haddad et al., 2025), face challenges due to the sparse reward problem. This is particularly relevant in the case of target affinity, where the large imbalance between inactive and active molecules hinders the model's ability to learn effective strategies for obtaining the desired reward (Korshunova et al., 2022). As an alternative, AL can be used to iteratively select only active molecules to retrain the GM. In addition, AL can help overcome challenges commonly faced by standalone GMs, including poor target engagement under low-data regimes (Van Tilborg & Grisoni, 2024), limited synthesizability of the generated molecules, and the generalization outside the training set (Gangwal et al., 2024; Loeffler et al., 2024; Kyro et al., 2024).

In this paper, we propose a multi-target generative workflow aimed at generating molecules with predicted simultaneous affinity to multiple targets. As a test case, we evaluated our multi-target generative workflow on the design of a pan-inhibitor targeting the main protease (Mpro) of diverse coronaviruses: SARS-CoV-2, SARS-CoV, and MERS-CoV (Shahhamzehei et al., 2022; Huang et al., 2023).

2. Methods

2.1. Multi-target Generative Workflow

The multi-target generative workflow builds upon the concept of single-target molecular generation. This involves generating molecules with predicted affinity for a single target, and extends its application to multiple biological targets.

The workflow starts by training the generator, a Seq2Seq VAE (Gómez-Bombarelli et al., 2018; Gupta et al., 2018), with a general dataset of molecules represented in text-based format (SMILES). This first training teaches the VAE the underlying grammar, thereby enabling the generation of chemically feasible molecules. Subsequently, the Seq2Seq VAE is fine-tuned with a fixed specific dataset of molecules with known affinity towards multiple targets, without requiring affinity to all targets simultaneously (Section A). By doing so, the generation is biased towards molecules with affinity to the multiple targets.

Once the VAE is pretrained, the multi-target generative workflow initiates an iterative process of molecular generation and refinement through a two-level AL workflow. The first level AL cycle, or Chemical AL, promotes molecules based on physicochemical properties. The second level AL cycle, or the Affinity AL, promotes molecules based on simultaneous predicted affinity to multiple targets (hereafter referred to as multi-target affinity) (Figure 1).

The two-level AL workflow begins by performing the Chemical AL cycle n times. In each cycle, new molecules are generated and filtered based on the presence of undesired structural motifs and chemoinformatic predictor thresholds. The resulting molecules are then used to fine-tune the VAE, starting from the general training weights. This fine-tuning utilises all accumulated molecules from previous and current Chemical AL cycles (the accumulated specific dataset), as well as molecules in the fixed specific dataset. After completing the n Chemical AL cycles, an Affinity AL cycle is conducted, in which all molecules in the accumulated specific dataset are filtered based on their multi-target affinity. The VAE is then fine-tuned, starting from the general training weights, using the accumulated molecules from previous and current Affinity AL cycles (the updated specific dataset) along with those from the fixed specific dataset. After m Affinity AL cycles, this two-level AL approach enables the progressive refinement of molecule generation, first by enforcing favourable chemical properties, and then by optimising for multi-target affinity (Figure 1). Currently, the workflow runs on a single GPU, and each Affinity AL cycle can take approximately 18 hours to complete (Section G).

2.2. Seq2Seq Variational Autoencoder

We implemented a Seq2Seq VAE architecture that processes molecules as SMILES sequences, one-hot encoded with a vocabulary size of $D = 50$. The encoder consists of a single LSTM layer that processes the input sequence $x = (x_1, \ldots, x_T)$ producing a final hidden state h_T, which is then passed through a fully connected layer with 256 units and ReLU activation:

$$h' = ReLU(W^{enc} h_T + b^{enc}) \tag{1}$$

From h', the VAE models the latent space as a probabilistic Gaussian distribution with mean vector $\mu \in \mathbb{R}^{128}$ and log-variance vector $log\ \sigma^2 \in \mathbb{R}^{128}$ computed as:

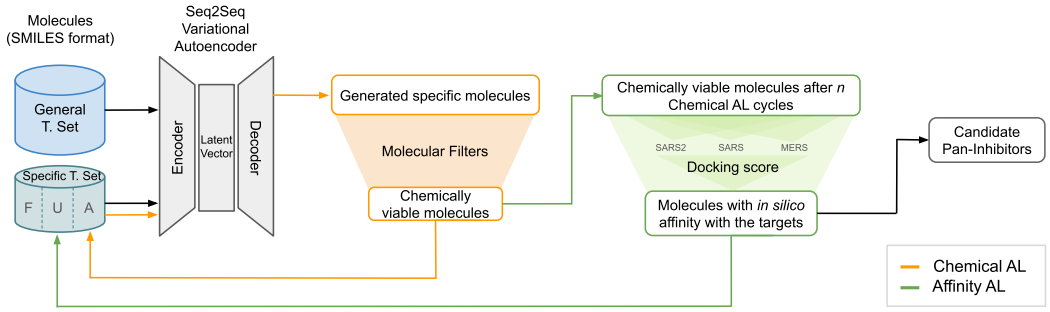

Figure 1. The multi-target generative workflow. The two-level AL workflow is illustrated with arrows of different colours: the first-level cycle, Chemical AL, is shown in orange, while the second-level cycle, Affinity AL, is shown in green. The specific training set stands for F: fixed, U: updated, and A: accumulated.

$$\mu = W_\mu h' + b_\mu, log\sigma^2 = W_\sigma h' + b_\sigma \tag{2}$$

where $W_\mu, W_\sigma \in \mathbb{R}^{128\times256}$ and $b_\mu, b_\sigma \in \mathbb{R}^{128}$. The latent vector z is sampled using the reparametrization trick:

$$z = \mu + \sigma \odot \epsilon, \epsilon \in \mathcal{N}(0, I), \sigma = e^{\frac{1}{2}log\sigma^2} \tag{3}$$

where ϵ is a random noise vector sampled from a standard normal distribution.

The decoder reconstructs the sequence by transforming z through a fully connected layer with 256 units and ReLU activation, which initializes the hidden state of a decoder LSTM. Finally, the decoder logits are passed to a softmax layer to obtain the output sequence.

During training, the VAE learns to reconstruct the input sequence x from a compressed latent representation z, while also shaping the latent space so that similar inputs map to nearby latent vectors. This is achieved by maximizing the Evidence Lower Bound (ELBO), a composite loss function that combines the Reconstruction loss and the Kullback-Leibler (KL) divergence (Odaibo, 2019).

The KL-divergence term in the VAE's loss regularizes the approximate posterior $q(z \mid x)$ toward the standard normal prior $p(z) = \mathcal{N}(0, I)$, yielding a continuous, densely populated latent space of valid molecular embeddings. To generate new molecules, one samples $z \sim N(0, I)$. The sampled vector z is then passed through a linear (fully connected) layer to initialize the decoder's LSTM hidden (and, if applicable, cell) state. From there, the decoder operates autoregressively: at each time step, it takes the embedding of the previously generated token (and optionally z again), updates its hidden state, and outputs a probability distribution over the next token. This process repeats until the end-of-sequence token is emitted or a predefined maximum length is reached, yielding a complete SMILES string.

2.3. Chemical Active Learning Cycle

The goal of the Chemical AL cycle is to guide molecular generation towards synthesisable, drug-like molecules that meet user-defined requirements for structural variability. A key concern in molecular generative AI is the limited synthetic accessibility and chemical quality of the generated molecules. To address this, the Chemical AL cycle starts by applying a substructure filter to remove generated molecules containing undesirable structural motifs that could hinder their synthesis and progression in a drug design campaign. Retention of such motifs could lead to their accumulation due to the iterative nature of the workflow (error propagation). Then, the filtered molecules are evaluated against three cheminformatic predictors designed to improve drug-likeness, further enhance synthetic accessibility, and guide molecular generation based on structural variability. These filtering predictors are: (1) Quantitative Estimate of Drug-likeness (QED) (Bickerton et al., 2012), ranging from 0 (least drug-like) to 1 (most drug-like), (2) Synthetic Accessibility (SA) score (Ertl & Schuffenhauer, 2009), ranging from 1 (easy to synthesise) to 10 (difficult to synthesise), and (3) Tanimoto similarity (TA) (Bajusz et al., 2015) against the specific dataset, computed using Morgan4 fingerprints, with values from 0 (no similarity) to 1 (identical molecules). Finally, medicinal chemistry filters are applied using 4 SMARTS-based catalogues to filter generated molecules containing undesirable structural motifs associated with toxicity, promiscuity, or poor pharmacokinetics. These catalogues (found in the Python package RDKit (Landrum et al., 2025)) include PAINS (Baell & Holloway, 2010), Brenk (Brenk et al., 2008), NIH (Doveston et al., 2014; Jadhav et al., 2010), and CHEMBL (Walters, 2025). Molecules that passed all filters are added to the accumulated specific dataset.

2.4. Affinity Active Learning Cycle

Affinity AL cycles are applied to guide molecular generation toward molecules with multi-target affinity. Ligand–protein affinities are predicted by docking molecules from the accumulated specific dataset to the multiple protein targets (Sections B and C). A key feature of this step is the use of two complementary docking score thresholds to filter molecules: (1) a global threshold computed as the mean docking score across the multiple targets, and (2) individual thresholds for each target, to prevent the global score from masking poor affinity to any single target. Molecules that satisfy all affinity thresholds are transferred to the updated specific dataset.

To progressively enhance affinity, a linear decay strategy is implemented, where docking score thresholds are decreased in each Affinity AL cycle (noting that lower docking scores correspond to higher predicted affinities). Let $T_g^{(i)}$ be the global threshold and $T_{ind}^{(i)}$ be the individual threshold at cycle i, with $\delta > 0$ being the decay rate. The thresholds are updated according to:

$$T_g^{(i+1)} = \begin{cases} T_g^{(i)} - \delta, & \text{if } N^{(i)} \geq N_{\min} \\ T_g^{(i)}, & \text{otherwise} \end{cases} \tag{4}$$

$$T_{\text{ind}}^{(i+1)} = \begin{cases} T_{\text{ind}}^{(i)} - \delta, & \text{if } N^{(i)} \geq N_{\min} \\ T_{\text{ind}}^{(i)}, & \text{otherwise} \end{cases} \tag{5}$$

where $N^{(i)}$ is the number of molecules passing the thresholds at cycle i, and N_{min} is the minimum number of molecules required to continue decreasing the thresholds. This decay is conditional upon the retention of at least N_{min} molecules meeting both threshold criteria; otherwise, the thresholds are maintained.

A stopping patience parameter of p is introduced, stopping the generative workflow if thresholds cannot be lowered in p consecutive Affinity AL cycles:

$$C = \begin{cases} 0, & \text{if } T_g^{(i+1)} < T_g^{(i)} \text{ or } T_{\text{ind}}^{(i+1)} < T_{\text{ind}}^{(i)} \\ C + 1, & \text{otherwise} \end{cases} \tag{6}$$

where C is a counter for consecutive cycles without threshold decay, stopping the generation process when $C \geq p$.

3. Results

The multi-target generative workflow was evaluated on the design of a coronavirus pan-inhibitor. Specifically, the goal was to generate a molecule capable of inhibiting the main protease (Mpro) of three distinct coronaviruses: SARS-CoV-2, SARS-CoV, and MERS-CoV. Mpro is a well-established antiviral target due to its essential role in viral replication. Its high conservation across coronavirus species (see Section B) makes it an ideal test case for our multi-target generative workflow.

We designed a case-specific generation pipeline with two main goals: to enhance molecular diversity and to progressively improve multi-target affinity toward the three Mpro targets. The pipeline starts with an initial Affinity AL cycle of 40 Chemical AL cycles, aimed at maximising molecular diversity through a TA threshold < 0.4. This is followed by shorter Affinity AL cycles, each with 10 Chemical AL cycles and a relaxed TA threshold < 0.6, intended to improve multi-target affinity. These shorter Affinity AL cycles enable more frequent filtering based on multi-target affinity, with effectiveness further increased through progressively decaying docking score thresholds across successive cycles ($N_{min} = 50$, $\delta = 0.1 kcal/mol$, and $p = 3$). All Affinity AL cycles use thresholds of QED ≥ 0.8 and SA ≤ 3 to favour drug-like, synthetically accessible molecules, along with starting docking score thresholds of -7.5 kcal/mol (global) and -7.0 kcal/mol (individual).

3.1. Generation performance under two configurations: Regular vs Ablated

Due to concerns that excessive filtering during the Chemical AL cycle could constrain molecular generation, we conducted two parallel generative cases. The first followed the full multi-target workflow as described in the Methods section (regular configuration). The second, serving as an ablation study, excluded SMARTS-based filtering catalogues to evaluate their impact within the Chemical AL cycles (ablated configuration).

Both generative configurations exhibited similar trends in terms of the number of generated molecules and progression through the Chemical and Affinity AL cycles (Figure 2A). Due to the reduced filtering of the ablated configuration, a total of 15 AL Affinity cycles were completed before reaching the stopping parameter p, with final docking thresholds of -8.5 kcal/mol (global) and -8 kcal/mol (individual). In contrast, the regular configuration, which included the SMARTS filters within, plateaued earlier, stopping after 13 AL Affinity cycles with slightly less stringent final thresholds of -8.3 kcal/mol (global) and -7.8 kcal/mol (individual).

Figure 3 (Section D) compares the global docking scores distributions of the generated molecules for the regular and the ablated configuration. In the ablated configuration, the distribution shifts leftward, displaying a higher molecular count in the high-affinity region (lower range of docking score), indicating better predicted multi-target affinity of the generated molecules. Importantly, the total number of

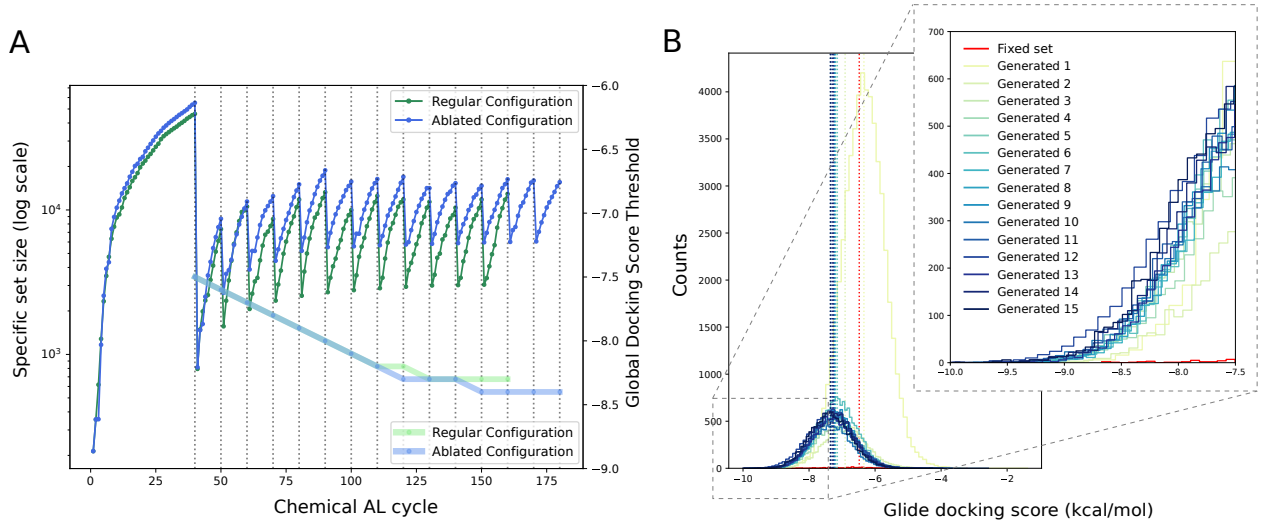

Figure 2. A) Evolution of specific set size (logarithmic scale) after each Chemical AL cycle for both regular and ablated configurations. Vertical dotted lines mark the end of each Affinity AL cycle. The secondary vertical axis represents the global docking score threshold applied to each Affinity AL cycle. The initial point represents the fixed specific set size. Points between two dotted vertical lines represent the evolution of the cumulative specific set size. The drop in size at the vertical lines corresponds to the Affinity AL filtering, highlighting the drop in size of the updated specific set. B) Global docking scores histograms of generated molecules in the ablated configuration across Affinity AL cycles compared to the histogram of the fixed specific dataset. The panel with a zoom-in shows the enrichment of generated molecules in the range of high multi-target affinity.

generated molecules passing all docking score thresholds and thus, accumulated in the updated specific dataset, was consistently higher in the ablated configuration, even after applying SMARTS filtering post-generation (Table 1). Specifically, we observed a 1.35-fold increase in molecular count under the -7.5 kcal/mol global and -7 kcal/mol individual thresholds, which increased to a 3-fold increase when applying the more stringent -9 kcal/mol global and -8 kcal/mol individual thresholds. This suggests that deferring such filters from the AL workflow and placing them as a post-generative filter offers a more permissive yet still chemically relevant approach to molecular generation.

3.2. Candidate pan-inhibitors

Given the enhanced capacity of the ablated configuration, followed by post-generation SMARTS filtering, to generate molecules with multi-target affinity, we focused the remainder of our analysis on this configuration. Figure 2B presents the distribution of the global docking scores for the generated molecules across the different Affinity AL cycles, as well as for those in the fixed specific dataset. Notably, the zoom-in panel reveals a consistent increase in the number of molecules in the lower docking score range with successive AL cycles. This effect is driven by the iterative nature of our AL workflow and the linear decay strategy applied to the docking score thresholds. As expected, we observe a marked enrichment of low-scoring molecules compared to the fixed set, which was not designed for multi-target affinity.

From this point onwards, we consider as a candidate pan-inhibitor any generated molecule with an individual docking score threshold of -8 kcal/mol for each target. Setting threshold values of -8 kcal/mol (global) and -8 kcal/mol (individual) yields a total of 310 molecules in the regular configuration, and a total of 650 molecules in the ablated configuration after applying SMARTS filtering (Table 1). Notably, only one molecule in the fixed specific set meets these docking score thresholds, highlighting the ability of our multi-target generative workflow to generalise toward the design of multi-target inhibitors.

3.3. Molecular diversity of the candidate pan-inhibitors

To assess the performance of our multi-target generative workflow and the diversity of generated molecules, we begin by examining the evolution of three key metrics across each Chemical AL cycle: validity, uniqueness, and novelty. Validity (the percentage of chemically valid molecules among those generated) shows considerable variation across cycles, generally fluctuating between 40% and 70% (on average 59.62% $\pm$ 11.96). This variability can be explained by the stochastic nature of the generation. In contrast, uniqueness (the proportion of non-duplicate valid molecules) and novelty (the percentage of unique molecules not found in the cumulative specific set) remain consistently high throughout the workflow, on average 98.51% $\pm$ 3.34 and 99.48% $\pm$

Table 1. Number of molecules under different global and individual docking score thresholds for the fixed specific set, and those generated from the regular and ablated configurations. It also includes molecules from the ablated configuration filtered post-generation based on the SMARTS-based catalogues. *Selected molecules as potential candidate pan-inhibitors.

Docking Thresholds		Number of Compounds			
Global	Individual	Fixed Specific Set	Generated Regular Config.	Generated Ablated Config.	Ablated Config. After Filtering
-7.5	-7	10	15,245	32,101	20,661
-8	-7.5	2	1,591	4,911	3,324
-8	-8	1	310*	932	650*
-8.5	-8	1	58	237	168
-9	-8	0	7	28	21
-9.5	-8	0	0	4	4

0.60, respectively (Figure 4, Section E). These results indicate that our multi-target generative model maintains high levels of molecular novelty and uniqueness over successive cycles.

To better understand the efficiency of the AL workflow, we analysed the percentage of molecules passing the two AL level filters across each Affinity AL cycle (Figure 5, Section E). On average, 51.54% $\pm$ 4.24 of generated molecules pass the Chemical AL filters. This indicates consistent performance throughout the generation process, excluding the first Affinity AL cycle, where only 31.6% passed due to more restrictive Chemical AL filters (a lower Tanimoto threshold of 0.4 versus 0.6 in subsequent cycles). In contrast, the percentage of molecules passing the Affinity AL filters is less consistent, with an average of 1.33% $\pm$ 1.64. Notably, there is a downward trend in molecules passing the Affinity filters, likely driven by the linear decay applied to the docking score thresholds.

To further evaluate the structural novelty of the generated molecules, we conducted an additional analysis focusing on the diversity of molecular scaffolds generated across Affinity AL cycles. Figure 6A (Section F) shows that our initial attempt to maximise variability by running a long and stringent Affinity AL cycle (with 40 Chemical AL cycles instead of 20 and a Tanimoto threshold of < 0.4 instead of < 0.6) was successfully achieved, with scaffold diversity increasing from 153 to 654 scaffold clusters. Although the plot appears to plateau around the sixth Affinity AL cycle, this does not imply a complete stop in scaffold generation. In fact, when we consider only the molecules that satisfied the final affinity thresholds of -8 kcal/mol (global and individual) (Figure 6B, Section F), we observe that the number of scaffold clusters continued to grow over time. This indicates that the model was still exploring new and diverse chemical regions within the high-affinity range.

Additionally, the expansion of chemical space throughout the generation was assessed using a UMAP representation of the generated molecules at each Affinity AL cycle (Figure 7, Section F). In the first large Affinity AL cycle, designed to promote variability, the model explored a broad and dispersed region of chemical space. Subsequent, shorter Affinity AL cycles showed more localised exploration, concentrating around previously visited regions, aligning with the design intent of the generation strategy. Notably, the candidate pan-inhibitors are evenly distributed across all chemical space rather than clustered in a single region, indicating that they originate from different and variable scaffolds.

Finally, the candidate pan-inhibitors were searched across the ultra-large chemical libraries Enamine REAL DB (≈6.7B molecules) (Enamine, 2024), ZINC22 (≈37B molecules) (Tingle et al., 2023), and WuXi (≈3.4B molecules) (WuXi, 2024). Among all generated candidates (969), only one molecule had an exact match in WuXi (generated in the regular configuration).

4. Discussion

We developed a multi-target generative workflow based on a VAE, iteratively refined through a two-level AL cycle. The first level, the Chemical AL cycle, promotes the generation of diverse, drug-like, and synthetically accessible molecules. The second, the Affinity AL cycle, guides generation toward molecules with multi-target affinity. To evaluate the workflow, we aimed to generate a diverse set of molecules with high simultaneous predicted affinity toward three homologous targets to obtain potential pan-inhibitor molecules. These targets were the main proteases (Mpro) of SARS-CoV-2, SARS-CoV, and MERS-CoV. This case study, involving structurally similar targets, provided an ideal test to validate the approach's multi-target applicability.

To assess the trade-off between restrictive filtering and flexibility in chemical space exploration, we compared two generative configurations: regular and ablated. For the Mpro targets, the ablated configuration, excluding SMARTS filtering during the Chemical AL cycle and instead applying them post-generation, proved more effective. This con-

figuration enabled broader exploration by not discarding molecules prematurely, leading to a higher yield of pan-inhibitor candidates. However, in other target contexts, the ablated configuration may lead to the accumulation of undesirable motifs that, while chemically unfavourable, could enhance affinity (a property prioritised by the workflow). In such cases, this issue will be exacerbated by the iterative nature of the multi-target generative process, which will propagate such motifs across generations. In situations where post-generation filtering significantly reduces candidates, it may be preferable to use the regular configuration with SMARTS filtering integrated into the chemical AL cycle.

Our diversity analysis demonstrated that the multi-target generative workflow successfully achieved its goal of enhancing molecular variability through the exploration of unseen regions of the chemical space. Broad exploration was prioritised during an initial, extended Affinity AL cycle with stricter similarity thresholds, enabling the generation of a diverse molecular pool. This pool was then refined for multi-target affinity in subsequent, shorter Affinity AL cycles. In these, similarity constraints were relaxed, and affinity thresholds became increasingly stringent. Notably, the candidate pan-inhibitors were evenly distributed across chemical space, highlighting the model's ability to generate structurally diverse molecules with strong multi-target affinity. Remarkably, nearly all of these (959 out of 960) were absent from several ultra-large chemical libraries, which collectively contain 37 billion molecules (without accounting for duplicates).

The application of our multi-target generative workflow on the discovery of a pan-inhibitor for coronaviruses led to the selection of 960 pan-inhibitor candidates. Nevertheless, while these results are encouraging, further *in silico* and experimental validation will be essential to narrow down our selection and identify candidates whose biological efficacy and safety can be confirmed. These molecules could serve as a starting point for addressing future coronavirus outbreaks.

These findings highlight the potential of our generative workflow in multi-target drug discovery. It is particularly well-suited for polypharmacological strategies aimed at treating complex diseases, where the simultaneous modulation of multiple targets could be transformative. In future applications, the workflow could also be extended to include not only multiple therapeutic targets but also antitargets, enabling simultaneous optimisation of efficacy and safety by actively avoiding undesirable interactions. Besides, accounting for potential antitarget interactions early in the drug discovery process will reduce the risk of costly downstream failures. However, as the number of targets and antitargets increases, so does the computational burden, particularly in the Affinity AL cycle, where docking calculations scale with each additional target. Therefore, computational efficiency through parallelization will be a critical consideration for the broader deployment of this approach.

Acknowledgements

This work has been supported by the predoctoral fellowship AGAUR-FI (2024 FI-1 00920) Joan Oró, funded by the Department of Research and Universities of the Government of Catalonia, with co-funding from the European Social Fund Plus. Additional funding was provided by the project CPP2022-009737, financed by MICIU/AEI /10.13039/501100011033 and by the European Union NextGenerationEU/PRTR.

Impact Statement

The impact of this work stems from the controllability and flexibility of the presented multi-target generative workflow. Control is achieved by sampling from the continuous latent space of the Seq2Seq VAE, followed by property refinement through a two-level AL workflow designed in close collaboration with medicinal chemistry experts. This allows the generation of chemically viable molecules with desired properties. Model flexibility, enables its application across a broad range of use cases. These include single-target inhibitor design, polypharmacology, pan-inhibitor development, and designing inhibitors for multiple conformations of a single target. Furthermore, structural variability can be reversed to favor the generation of molecules similar to the specific dataset, making the workflow well-suited for lead optimization. Additionally, the workflow can be used in low-data regimes using a virtual seed. This can be molecules obtained from virtual screening of ultra-large chemical libraries, which can guide generation when known inhibitors are unavailable. Ultimately, these capabilities contribute to accelerating the early-stage of the drug discovery process, and consequently, have a societal impact by reducing the cost and time of the overall drug development process.

References

Ahmad, S., Bano, N., Sharma, S., Sakina, S., Ahmad, N., and Raza, K. Generative AI in Drug Designing: Current State-of-the-Art and Perspectives. In Raza, K., Ahmad, N., and Singh, D. (eds.), *Generative AI: Current Trends and Applications*, pp. 427–463. Springer Nature, Singapore, 2024. ISBN 978-981-97-8460-8. doi: 10.1007/978-981-97-8460-8_20. URL https://doi.org/10.1007/978-981-97-8460-8_20.

Baell, J. B. and Holloway, G. A. New Substructure Filters for Removal of Pan Assay Interference Compounds (PAINS) from Screening Libraries and for Their Ex-

clusion in Bioassays. *Journal of Medicinal Chemistry*, 53(7):2719–2740, April 2010. ISSN 0022-2623. doi: 10.1021/jm901137j. URL https://doi.org/10.1021/jm901137j. Publisher: American Chemical Society.

Bajusz, D., Rácz, A., and Héberger, K. Why is Tanimoto index an appropriate choice for fingerprint-based similarity calculations? *Journal of Cheminformatics*, 7 (1):20, May 2015. ISSN 1758-2946. doi: 10.1186/s13321-015-0069-3. URL https://doi.org/10.1186/s13321-015-0069-3.

Bento, A. P., Gaulton, A., Hersey, A., Bellis, L. J., Chambers, J., Davies, M., Krüger, F. A., Light, Y., Mak, L., McGlinchey, S., Nowotka, M., Papadatos, G., Santos, R., and Overington, J. P. The ChEMBL bioactivity database: an update. *Nucleic Acids Research*, 42(Database issue): D1083–1090, January 2014. ISSN 1362-4962. doi: 10.1093/nar/gkt1031.

Bickerton, G. R., Paolini, G. V., Besnard, J., Muresan, S., and Hopkins, A. L. Quantifying the chemical beauty of drugs. *Nature Chemistry*, 4(2):90–98, February 2012. ISSN 1755-4349. doi: 10.1038/nchem.1243. URL https://www.nature.com/articles/nchem.1243. Publisher: Nature Publishing Group.

Blaschke, T., Engkvist, O., Bajorath, J., and Chen, H. Memory-assisted reinforcement learning for diverse molecular de novo design. *Journal of Cheminformatics*, 12(1):68, November 2020. ISSN 1758-2946. doi: 10.1186/s13321-020-00473-0. URL https://doi.org/10.1186/s13321-020-00473-0.

Brenk, R., Schipani, A., James, D., Krasowski, A., Gilbert, I. H., Frearson, J., and Wyatt, P. G. Lessons Learnt from Assembling Screening Libraries for Drug Discovery for Neglected Diseases. *ChemMedChem*, 3(3):435–444, 2008. ISSN 1860-7187. doi: 10.1002/cmdc.200700139. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/cmdc.200700139.

Cichońska, A., Ravikumar, B., and Rahman, R. AI for targeted polypharmacology: The next frontier in drug discovery. *Current Opinion in Structural Biology*, 84:102771, February 2024. ISSN 0959-440X. doi: 10.1016/j.sbi.2023.102771. URL https://www.sciencedirect.com/science/article/pii/S0959440X23002452.

Doveston, R. G., Tosatti, P., Dow, M., Foley, D. J., Li, H. Y., Campbell, A. J., House, D., Churcher, I., Marsden, S. P., and Nelson, A. A unified lead-oriented synthesis of over fifty molecular scaffolds. *Organic & Biomolecular Chemistry*, 13(3):859–865, December 2014. ISSN 1477-0539. doi: 10.1039/C4OB02287D. URL https://pubs.rsc.org/en/content/articlelanding/2015/ob/c4ob02287d. Publisher: The Royal Society of Chemistry.

Edgar, R. C. MUSCLE: multiple sequence alignment with high accuracy and high throughput. *Nucleic Acids Research*, 32(5):1792–1797, 2004. ISSN 0305-1048. doi: 10.1093/nar/gkh340. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC390337/.

Enamine. REAL Database - Enamine, 2024. URL https://enamine.net/compound-collections/real-compounds/real-database.

Ertl, P. and Schuffenhauer, A. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of Cheminformatics*, 1(1):8, June 2009. ISSN 1758-2946. doi: 10.1186/1758-2946-1-8. URL https://doi.org/10.1186/1758-2946-1-8.

Ester, M., Kriegel, H.-P., and Xu, X. A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. *KDD-96 Proceedings*, 1996.

Filella-Merce, I., Molina, A., Orzechowski, M., Díaz, L., Zhu, Y. M., Mor, J. V., Malo, L., Yekkirala, A. S., Ray, S., and Guallar, V. Optimizing Drug Design by Merging Generative AI With Active Learning Frameworks, May 2023. URL http://arxiv.org/abs/2305.06334. arXiv:2305.06334 [q-bio].

Gangwal, A., Ansari, A., Ahmad, I., Azad, A. K., and Wan Sulaiman, W. M. A. Current strategies to address data scarcity in artificial intelligence-based drug discovery: A comprehensive review. *Computers in Biology and Medicine*, 179:108734, September 2024. ISSN 0010-4825. doi: 10.1016/j.compbiomed.2024.108734. URL https://www.sciencedirect.com/science/article/pii/S0010482524008199.

Grisoni, F., Huisman, B. J. H., Button, A. L., Moret, M., Atz, K., Merk, D., and Schneider, G. Combining generative artificial intelligence and on-chip synthesis for de novo drug design. *Science Advances*, 7(24):eabg3338, June 2021. doi: 10.1126/sciadv.abg3338. URL https://www.science.org/doi/10.1126/sciadv.abg3338. Publisher: American Association for the Advancement of Science.

Gupta, A., Müller, A. T., Huisman, B. J. H., Fuchs, J. A., Schneider, P., and Schneider, G. Generative Recurrent Networks for De Novo Drug Design. *Molecular Informatics*, 37(1-2):1700111, January 2018. ISSN 1868-1743. doi: 10.1002/minf.201700111. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5836943/.

Gómez-Bombarelli, R., Wei, J. N., Duvenaud, D., Hernández-Lobato, J. M., Sánchez-Lengeling, B., Sheberla, D., Aguilera-Iparraguirre, J., Hirzel, T. D., Adams, R. P., and Aspuru-Guzik, A. Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. *ACS Central Science*, 4(2):268–276, February 2018. ISSN 2374-7943. doi: 10.1021/acscentsci.7b00572. URL https://doi.org/10.1021/acscentsci.7b00572. Publisher: American Chemical Society.

Haddad, R., Litsa, E. E., Liu, Z., Yu, X., Burkhardt, D., and Bhisetti, G. Targeted molecular generation with latent reinforcement learning. *Scientific Reports*, 15(1): 15202, April 2025. ISSN 2045-2322. doi: 10.1038/s41598-025-99785-0. URL https://www.nature.com/articles/s41598-025-99785-0. Publisher: Nature Publishing Group.

Halgren, T. A. Identifying and Characterizing Binding Sites and Assessing Druggability. *Journal of Chemical Information and Modeling*, 49(2):377–389, February 2009. ISSN 1549-9596. doi: 10.1021/ci800324m. URL https://doi.org/10.1021/ci800324m. Publisher: American Chemical Society.

Halgren, T. A., Murphy, R. B., Friesner, R. A., Beard, H. S., Frye, L. L., Pollard, W. T., and Banks, J. L. Glide: A New Approach for Rapid, Accurate Docking and Scoring. 2. Enrichment Factors in Database Screening. *Journal of Medicinal Chemistry*, 47(7):1750–1759, March 2004. ISSN 0022-2623. doi: 10.1021/jm030644s. URL https://doi.org/10.1021/jm030644s. Publisher: American Chemical Society.

Huang, C., Shuai, H., Qiao, J., Hou, Y., Zeng, R., Xia, A., Xie, L., Fang, Z., Li, Y., Yoon, C., Huang, Q., Hu, B., You, J., Quan, B., Zhao, X., Guo, N., Zhang, S., Ma, R., Zhang, J., Wang, Y., Yang, R., Zhang, S., Nan, J., Xu, H., Wang, F., Lei, J., Chu, H., and Yang, S. A new generation Mpro inhibitor with potent activity against SARS-CoV-2 Omicron variants. *Signal Transduction and Targeted Therapy*, 8(1):1–13, March 2023. ISSN 2059-3635. doi: 10.1038/s41392-023-01392-w. URL https://www.nature.com/articles/s41392-023-01392-w. Publisher: Nature Publishing Group.

Isigkeit, L., Hörmann, T., Schallmayer, E., Scholz, K., Lillich, F. F., Ehrler, J. H. M., Hufnagel, B., Büchner, J., Marschner, J. A., Pabel, J., Proschak, E., and Merk, D. Automated design of multi-target ligands by generative deep learning. *Nature Communications*, 15(1): 7946, September 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-52060-8. URL https://www.nature.com/articles/s41467-024-52060-8. Publisher: Nature Publishing Group.

Jadhav, A., Ferreira, R. S., Klumpp, C., Mott, B. T., Austin, C. P., Inglese, J., Thomas, C. J., Maloney, D. J., Shoichet, B. K., and Simeonov, A. Quantitative Analyses of Aggregation, Autofluorescence, and Reactivity Artifacts in a Screen for Inhibitors of a Thiol Protease. *Journal of Medicinal Chemistry*, 53(1):37–51, January 2010. ISSN 0022-2623. doi: 10.1021/jm901070c. URL https://doi.org/10.1021/jm901070c. Publisher: American Chemical Society.

Katoh, K. and Standley, D. M. MAFFT Multiple Sequence Alignment Software Version 7: Improvements in Performance and Usability. *Molecular Biology and Evolution*, 30(4):772–780, April 2013. ISSN 0737-4038. doi: 10.1093/molbev/mst010. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3603318/.

Kim, S., Chen, J., Cheng, T., Gindulyte, A., He, J., He, S., Li, Q., Shoemaker, B., Thiessen, P., Yu, B., Zaslavsky, L., Zhang, J., and Bolton, E. PubChem 2025 update. *Nucleic Acids Research*, 53(D1):D1516–D1525, January 2025. ISSN 1362-4962. doi: 10.1093/nar/gkae1059. URL https://doi.org/10.1093/nar/gkae1059.

Knox, C., Wilson, M., Klinger, C., Franklin, M., Oler, E., Wilson, A., Pon, A., Cox, J., Chin, N. E., Strawbridge, S., Garcia-Patino, M., Kruger, R., Sivakumaran, A., Sanford, S., Doshi, R., Khetarpal, N., Fatokun, O., Doucet, D., Zubkowski, A., Rayat, D., Jackson, H., Harford, K., Anjum, A., Zakir, M., Wang, F., Tian, S., Lee, B., Liigand, J., Peters, H., Wang, R. Q., Nguyen, T., So, D., Sharp, M., da Silva, R., Gabriel, C., Scantlebury, J., Jasinski, M., Ackerman, D., Jewison, T., Sajed, T., Gautam, V., and Wishart, D. DrugBank 6.0: the DrugBank Knowledgebase for 2024. *Nucleic Acids Research*, 52(D1):D1265–D1275, January 2024. ISSN 0305-1048. doi: 10.1093/nar/gkad976. URL https://doi.org/10.1093/nar/gkad976.

Korshunova, M., Huang, N., Capuzzi, S., Radchenko, D. S., Savych, O., Moroz, Y. S., Wells, C. I., Willson, T. M., Tropsha, A., and Isayev, O. Generative and reinforcement learning approaches for the automated de novo design of bioactive compounds. *Communications Chemistry*, 5 (1):1–11, October 2022. ISSN 2399-3669. doi: 10.1038/s42004-022-00733-0. URL https://www.nature.com/articles/s42004-022-00733-0. Publisher: Nature Publishing Group.

Kyro, G. W., Morgunov, A., Brent, R. I., and Batista, V. S. ChemSpaceAL: An Efficient Active Learning Methodology Applied to Protein-Specific Molecular Generation. *Journal of Chemical Information and Modeling*, 64(3):653–665, February 2024. ISSN 1549-9596. doi: 10.1021/acs.jcim.3c01456. URL https://doi.org/

10.1021/acs.jcim.3c01456. Publisher: American Chemical Society.

Landrum, G., Tosco, P., Kelley, B., Rodriguez, R., Cosgrove, D., Vianello, R., sriniker, Gedeck, P., Jones, G., NadineSchneider, Kawashima, E., Nealschneider, D., Dalke, A., Swain, M., Cole, B., Turk, S., Savelev, A., tadhurst cdd, Vaucher, A., Wójcikowski, M., Take, I., Walker, R., Scalfani, V. F., Faara, H., Ujihara, K., Probst, D., Lehtivarjo, J., godin, g., Pahl, A., and Monat, J. rdkit/rdkit: 2025_03_1 (Q1 2025) Release, March 2025. URL https://zenodo.org/records/15115844.

Liu, X., Ye, K., van Vlijmen, H. W. T., Emmerich, M. T. M., IJzerman, A. P., and van Westen, G. J. P. DrugEx v2: de novo design of drug molecules by Pareto-based multi-objective reinforcement learning in polypharmacology. *Journal of Cheminformatics*, 13(1): 85, November 2021. ISSN 1758-2946. doi: 10.1186/s13321-021-00561-9. URL https://doi.org/10.1186/s13321-021-00561-9.

Loeffler, H. H., Wan, S., Klähn, M., Bhati, A. P., and Coveney, P. V. Optimal Molecular Design: Generative Active Learning Combining REINVENT with Precise Binding Free Energy Ranking Simulations. *Journal of Chemical Theory and Computation*, 20(18):8308–8328, September 2024. ISSN 1549-9618. doi: 10.1021/acs.jctc.4c00576. URL https://doi.org/10.1021/acs.jctc.4c00576. Publisher: American Chemical Society.

Madhavi Sastry, G., Adzhigirey, M., Day, T., Annabhimoju, R., and Sherman, W. Protein and ligand preparation: parameters, protocols, and influence on virtual screening enrichments. *Journal of Computer-Aided Molecular Design*, 27(3):221–234, March 2013. ISSN 1573-4951. doi: 10.1007/s10822-013-9644-8. URL https://doi.org/10.1007/s10822-013-9644-8.

McInnes, L., Healy, J., and Melville, J. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction, September 2020. URL http://arxiv.org/abs/1802.03426. arXiv:1802.03426 [stat].

MedChemExpress. MedChemExpress: Master of Bioactive Molecules | Inhibitors, Screening Libraries & Proteins, 2024. URL https://www.medchemexpress.com/.

Munson, B. P., Chen, M., Bogosian, A., Kreisberg, J. F., Licon, K., Abagyan, R., Kuenzi, B. M., and Ideker, T. De novo generation of multi-target compounds using deep generative chemistry. *Nature Communications*, 15 (1):3636, May 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-47120-y. URL https://www.nature.com/articles/s41467-024-47120-y. Publisher: Nature Publishing Group.

Neumann, A. and Klein, R. A Benchmark Set of Bioactive Molecules for Diversity Analysis of Compound Libraries and Combinatorial Chemical Spaces, April 2025. URL https://chemrxiv.org/engage/chemrxiv/article-details/67ebaa8c81d2151a028aed21.

Odaibo, S. Tutorial: Deriving the Standard Variational Autoencoder (VAE) Loss Function, July 2019. URL http://arxiv.org/abs/1907.08956. arXiv:1907.08956 [cs].

Olivecrona, M., Blaschke, T., Engkvist, O., and Chen, H. Molecular de-novo design through deep reinforcement learning. *Journal of Cheminformatics*, 9(1):48, September 2017. ISSN 1758-2946. doi: 10.1186/s13321-017-0235-x. URL https://doi.org/10.1186/s13321-017-0235-x.

rdkit.Chem package. rdkit.Chem package — The RDKit 2025.03.1 documentation, 2025. URL https://www.rdkit.org/docs/source/rdkit.Chem.html.

Schrödinger, L. Schrödinger Release 2025-2: LigPrep, 2025. URL https://www.schrodinger.com/platform/products/ligprep/.

scikit learn. DBSCAN, 2025. URL https://scikit-learn/stable/modules/generated/sklearn.cluster.DBSCAN.html.

Shahhamzehei, N., Abdelfatah, S., and Efferth, T. In Silico and In Vitro Identification of Pan-Coronaviral Main Protease Inhibitors from a Large Natural Product Library. *Pharmaceuticals*, 15(3):308, March 2022. ISSN 1424-8247. doi: 10.3390/ph15030308. URL https://www.mdpi.com/1424-8247/15/3/308. Number: 3 Publisher: Multidisciplinary Digital Publishing Institute.

Sheikholeslami, M., Mazrouei, N., Gheisari, Y., Fasihi, A., Irajpour, M., and Motahharynia, A. DrugGen enhances drug discovery with large language models and reinforcement learning. *Scientific Reports*, 15(1): 13445, April 2025. ISSN 2045-2322. doi: 10.1038/s41598-025-98629-1. URL https://www.nature.com/articles/s41598-025-98629-1. Publisher: Nature Publishing Group.

Sievers, F. and Higgins, D. G. Clustal Omega for making accurate alignments of many protein sequences. *Protein Science*, 27(1):135–145, 2018. ISSN 1469-896X. doi: 10.1002/pro.3290. URL https://onlinelibrary.wiley.

com/doi/abs/10.1002/pro.3290. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/pro.3290.

Swanson, K., Liu, G., Catacutan, D. B., Arnold, A., Zou, J., and Stokes, J. M. Generative AI for designing and validating easily synthesizable and structurally novel antibiotics. *Nature Machine Intelligence*, 6(3): 338–353, March 2024. ISSN 2522-5839. doi: 10.1038/s42256-024-00809-7. URL https://www.nature.com/articles/s42256-024-00809-7.

Tang, X., Dai, H., Knight, E., Wu, F., Li, Y., Li, T., and Gerstein, M. A Survey of Generative AI for de novo Drug Design: New Frontiers in Molecule and Protein Generation, June 2024. URL http://arxiv.org/abs/2402.08703. arXiv:2402.08703 [q-bio].

Tingle, B. I., Tang, K. G., Castanon, M., Gutierrez, J. J., Khurelbaatar, M., Dandarchuluun, C., Moroz, Y. S., and Irwin, J. J. ZINC-22A Free Multi-Billion-Scale Database of Tangible Compounds for Ligand Discovery. *Journal of Chemical Information and Modeling*, 63(4):1166–1176, February 2023. ISSN 1549-9596. doi: 10.1021/acs.jcim.2c01253. URL https://doi.org/10.1021/acs.jcim.2c01253. Publisher: American Chemical Society.

tmtools. tmtools: Python bindings around the TM-align code for structural alignment of proteins, 2025. URL https://github.com/jvkersch/tmtools.

umap learn. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction — umap 0.5.8 documentation, 2025. URL https://umap-learn.readthedocs.io/en/latest/.

Van Tilborg, D. and Grisoni, F. Traversing Chemical Space with Active Deep Learning: A Computational Framework for Low-data Drug Discovery, February 2024. URL https://chemrxiv.org/engage/chemrxiv/article-details/65d8833ce9ebbb4db9098cb5.

Walters, P. PatWalters/rd_filters, March 2025. URL https://github.com/PatWalters/rd_filters. original-date: 2018-08-08T01:29:21Z.

Waskom, M., Botvinnik, O., O'Kane, D., Hobson, P., Lukauskas, S., Gemperline, D. C., Augspurger, T., Halchenko, Y., Cole, J. B., Warmenhoven, J., Ruiter, J. d., Pye, C., Hoyer, S., Vanderplas, J., Villalba, S., Kunter, G., Quintero, E., Bachant, P., Martin, M., Meyer, K., Miles, A., Ram, Y., Yarkoni, T., Williams, M. L., Evans, C., Fitzgerald, C., Brian, Fonnesbeck, C., Lee, A., and Qalieh, A. mwaskom/seaborn: v0.8.1 (September 2017), September 2017. URL https://zenodo.org/records/883859.

WuXi. Library-Chemistry, 2024. URL https://chemistry.wuxiapptec.com/library.

wwPDB consortium. Protein Data Bank: the single global archive for 3D macromolecular structure data. *Nucleic Acids Research*, 47(D1):D520–D528, January 2019. ISSN 0305-1048. doi: 10.1093/nar/gky949. URL https://doi.org/10.1093/nar/gky949.

Zhavoronkov, A., Ivanenkov, Y. A., Aliper, A., Veselov, M. S., Aladinskiy, V. A., Aladinskaya, A. V., Terentiev, V. A., Polykovskiy, D. A., Kuznetsov, M. D., Asadulaev, A., Volkov, Y., Zholus, A., Shayakhmetov, R. R., Zhebrak, A., Minaeva, L. I., Zagribelnyy, B. A., Lee, L. H., Soll, R., Madge, D., Xing, L., Guo, T., and Aspuru-Guzik, A. Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nature Biotechnology*, 37(9):1038–1040, September 2019. ISSN 1546-1696. doi: 10.1038/s41587-019-0224-x. URL https://www.nature.com/articles/s41587-019-0224-x. Publisher: Nature Publishing Group.

A. Datasets

The general training set was obtained from the ChEMBL 30 database (Bento et al., 2014) (2.7 million molecules) to which a drug-likeness filter was applied to ensure suitability for the generative model (molecular weight between 150 and 500 Da, free of salts, and compliant with Lipinski's Rule of Five), reducing its size to 247,199 molecules in SMILES format.

The fixed specific training set was assembled by collecting experimentally validated inhibitors (with known IC_{50} affinity values) targeting the Mpro of SARS-CoV-2, SARS-CoV, and MERS-CoV. These molecules were retrieved from multiple chemical databases, including PDB (wwPDB consortium, 2019), ChEMBL (Bento et al., 2014), PubChem (Kim et al., 2025), DrugBank (Knox et al., 2024), and MedChemExpress (MedChemExpress, 2024), resulting in a total of 477 molecules after removing peptides. Of this, 267 corresponded to SARS-CoV-2 inhibitors, 198 to SARS-CoV inhibitors, and 30 to MERS-CoV inhibitors. Given this imbalance between species, all retrieved molecules were subjected to a cross-docking protocol against the three targets to create a balanced specific training set. Molecules below a predicted affinity threshold of -5.9 kcal/mol (docking score from Glide) to all three proteases were retained. This docking score threshold was established by calculating the average docking score of all molecules across all targets. This process yielded a final specific training set comprising 214 unique molecules, each with a minimum signal of predicted multi-target affinity.

B. Target Selection and Preparation

All available crystallographic structures of the main protease (Mpro) from the Protein Data Bank (PDB) (wwPDB consortium, 2019) were collected and classified according to their viral origin: 479 from SARS-CoV-2, 32 from SARS-CoV, and 32 from MERS-CoV. Target structures were preprocessed by removing water molecules, ligands, and ions. In cases where structures contained multiple chains, reflecting both Mpro dimeric and monomeric states, individual chains were separated, and only monomeric chains within a defined residue range (minimum of 290 and maximum of 330) encompassing the catalytic site were retained, as the objective was to inhibit a single active monomer. For each virus, all structures were superimposed onto a designated reference structure using TMalign (tmtools, 2025), and further prepared using Schrödinger Protein Preparation Wizard (Madhavi Sastry et al., 2013), which completed the missing atoms and optimized the structural geometry.

To remove redundant structures (structures with identical catalytic site conformation), we cluster them based on the 3D volume of their catalytic site. To do so, we computed the 3D catalytic site volumes using Schrödinger SiteMap (Halgren, 2009) and constructed a pairwise catalytic site volume overlapping matrix. Then we hierarchically clustered this matrix using Seaborn (Waskom et al., 2017). This allowed us to identify clusters of very similar 3D catalytic site volumes from which we extract representatives, reducing the total number of Mpro structures from 543 to 195.

Cross-docking of the fixed specific set was subsequently conducted with these 195 structures, and receptor selection was based on the total number of inhibitors (the structure with the highest counts was selected) that were below the docking score threshold of -5.9 kcal/mol (same used in the construction of the specific set). Further evaluation based on resolution, model completeness, and electron density quality was done in the selected structures. Based on these criteria, the final selected target structures were 7RNW for SARS-CoV-2, 2GX4 for SARS-CoV, and 7ENE for MERS-CoV.

To assess the degree of similarity among homologous Mpro proteins across different coronavirus species, a sequence conservation analysis was performed. Pairwise sequence alignments were carried out using three independent tools, MAFFT (Katoh & Standley, 2013), MUSCLE (Edgar, 2004), and Clustal Omega (Sievers & Higgins, 2018), all of which produced consistent results. The resulting pairwise sequence percentage identities (%ID) are summarised in Table 2. In addition to sequence-based comparisons, we superimposed the target structures using TM-align (tmtools, 2025). Pairwise root-mean-square deviations (RMSD) reflecting high structural conservation are also reported in Table 2. From the previous structural superimpositions, we extracted their structure-based sequence alignments, from which we derive the %ID of the catalytic sites, the key regions targeted in the design of multi-target inhibitors (Table 3).

Table 2. Percentage identity and RMSD of the Mpro sequences for SARS-CoV-2 (SARS2), SARS-CoV (SARS), and MERS-CoV (MERS).

%ID AND RMSD	SARS2	SARS	MERS
SARS2	-	96.08% (RMSD=1.06)	50% (RMSD=1.32)
SARS	96.08% (RMSD=1.06)	-	50.98% (RMSD=1.55)
MERS	50% (RMSD=1.32)	50.98% (RMSD=1.55)	-

Table 3. Percentage identity derived from the structure-based sequence alignment just on the catalytic site of the Mpro for SARS-CoV-2 (SARS2), SARS-CoV (SARS), and MERS-CoV (MERS).

%ID	SARS2	SARS	MERS
SARS2	-	96.7%	63.33%
SARS	96.7%	-	66.67%
MERS	63.33%	66.67%	-

C. Docking Protocol

Docking simulations were computed using Schrödinger's Glide software (Halgren et al., 2004), operating in standard precision (SP) mode, with a grid centred on the catalytic active site (C145:SG) and fixed dimensions of 10Å for the inner box and 30Å for the outer box. For each molecule, up to five docking poses were evaluated, and no structural constraints were applied.

D. Cumulative Counts across Affinity AL Cycles

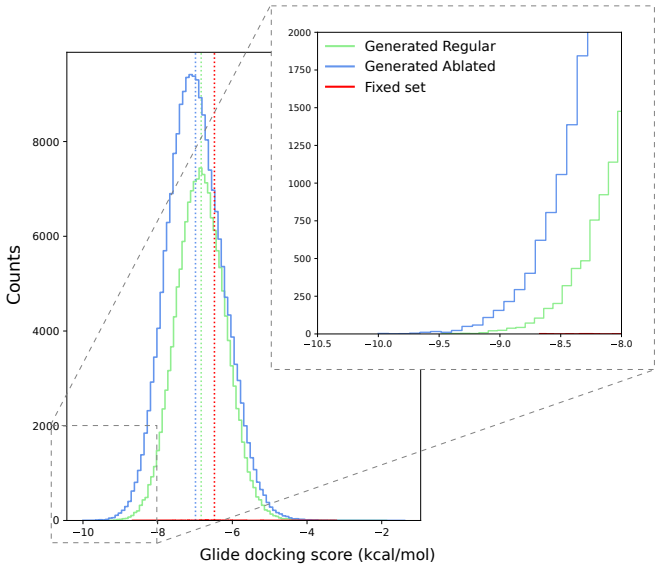

Figure 3. Global docking score cumulative histograms of generated molecules across all Affinity AL cycles for the regular and the ablated configurations, and also the fixed specific set. The zoomed-in section represents the molecules under the final defined threshold of -8 kcal/mol.

E. Generation Statistics

The validity, uniqueness, and novelty of the generated molecules at each Chemical AL cycle were calculated as follows:

$$\text{Validity} = \left(\frac{N_{\text{val}}}{N_{\text{gen}}}\right) \times 100 \tag{7}$$

$$\text{Uniqueness} = \left(\frac{N_{\text{uni}}}{N_{\text{val}}}\right) \times 100 \tag{8}$$

$$\text{Novelty} = \left(\frac{N_{\text{unk}}}{N_{\text{uni}}}\right) \times 100 \tag{9}$$

where N_{gen} stands for the total generated molecules, N_{val} for the total valid molecules from N_{gen}, N_{uni} for the total unique molecules (non-duplicated) from N_{val}, and N_{unk} for the total number of unique molecules not found in the cumulative specific set.

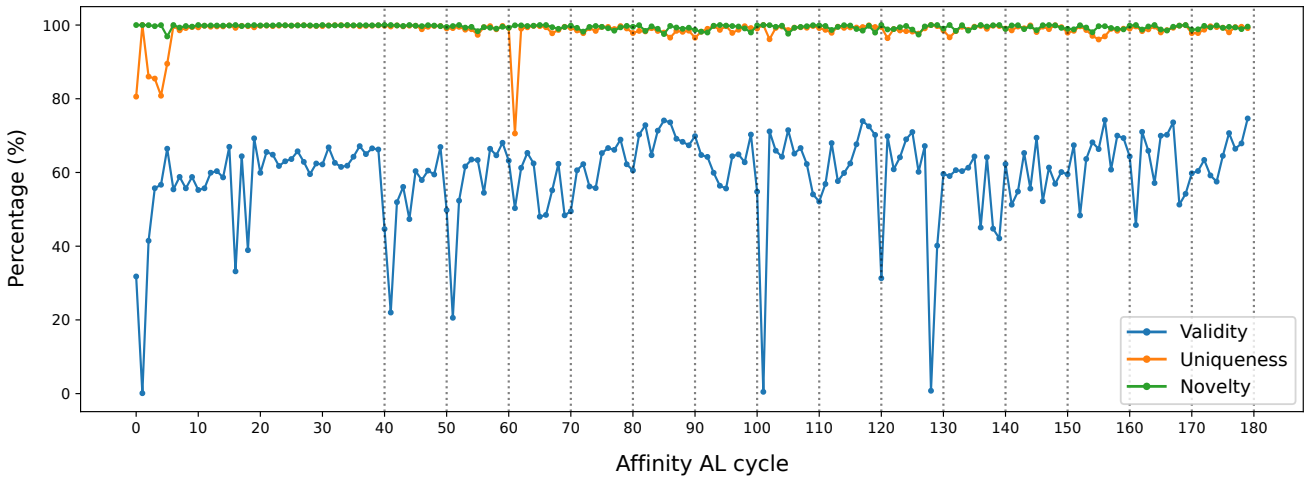

Figure 4. Validity, uniqueness, and novelty of the molecules generated at each Chemical AL cycle. The vertical dotted lines represent the division between Affinity AL cycles.

We tracked the number of molecules that fulfilled the criteria of the Chemical and Affinity AL cycles at each Affinity AL iteration.

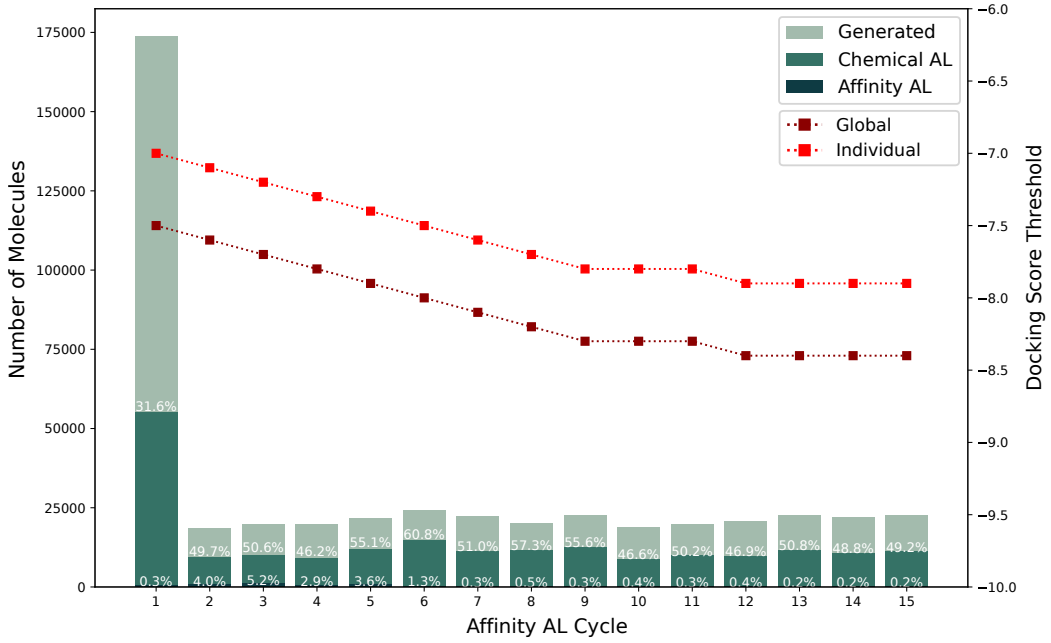

Figure 5. Number and percentage of generated molecules at each Affinity AL cycle fulfilling the Chemical AL filters and the Affinity AL filters. The docking score thresholds (global and individual) that were applied at each Affinity AL cycle are represented on the right-y axis.

F. Molecular Diversity

To extract scaffolds from molecules that met specified global and individual docking score thresholds, we used the Scaffolds.MurckoScaffold module from Chem package in the RDKit Python library (rdkit.Chem package, 2025). For clustering these molecular scaffolds, we applied the DBSCANS algorithm (Ester et al., 1996) from the scikit-learn Python library (scikit learn, 2025), varying the value of the epsilon parameter (minimum similarity between any pair of scaffolds within the same cluster) to ensure that the observed trends are not artifacts of the chosen epsilon value.

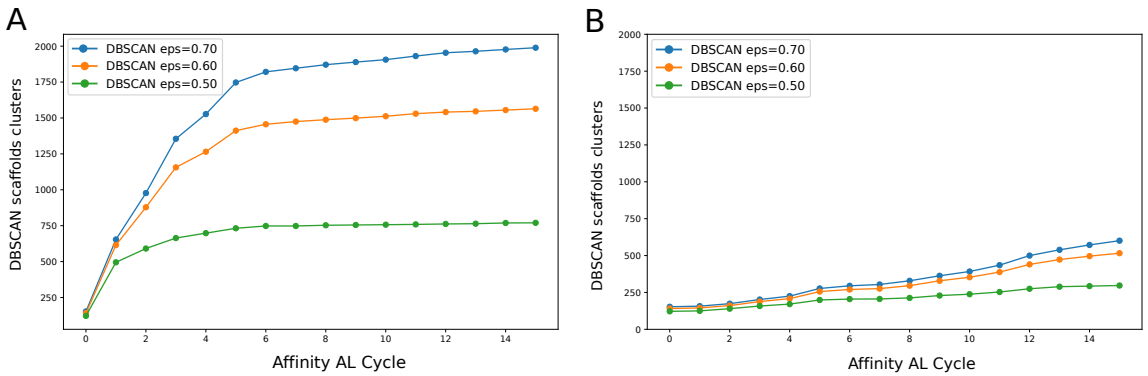

Figure 6. Evolution of scaffold clusters across Affinity AL cycles. Each connected line represents a DBSCAN clustering performed with a different epsilon parameter, as indicated in the legend. A) Evolution of scaffold clusters among generated molecules below the thresholds applied at each Affinity AL cycle, and B) Evolution of scaffold clusters among generated molecules below thresholds of -8 kcal/mol (global and individual).

The two-dimensional molecular representation provided by the Uniform Manifold Approximation and Projection (UMAP) algorithm (McInnes et al., 2020) enabled visualisation of the exploratory behaviour of our multi-target generative workflow. The UMAP in Figure 7 was obtained with the umap-learn Python library (umap learn, 2025), using Morgan4 fingerprints and Hamming distance.

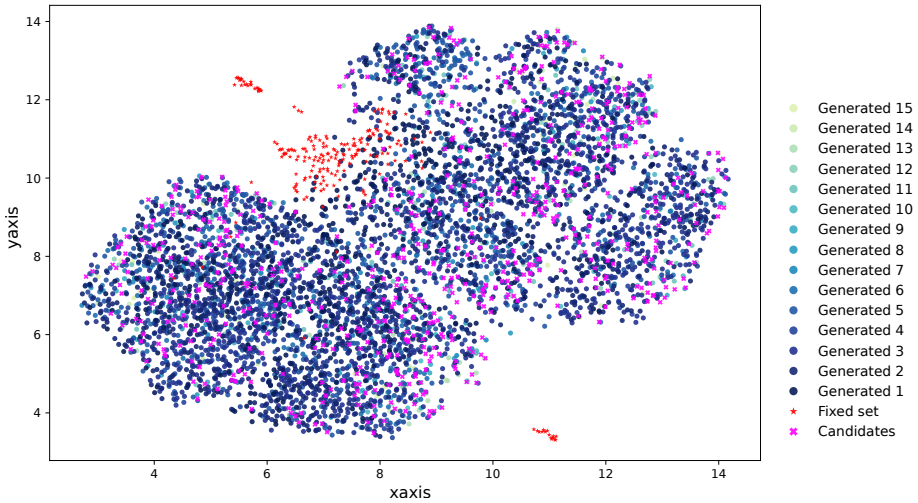

Figure 7. UMAP illustrating the generated molecules over the Affinity AL cycles, with marker styles and colours indicating different cycles as shown in the legend. Fucsin crosses represent the candidate pan-inhibitors, thus generated molecules that passed the docking score thresholds of -8 kcal (global) and -8 kcal/mol (individual).

G. Performance time

To evaluate the computational efficiency of the proposed multi-target generative workflow, execution times were calculated for each major component of the workflow. The presented times are real time (wall-clock time):

- The training time for the general dataset (≈200k molecules) required 38.26 $\pm$ 0.53 minutes on a single GPU H100.
- Fine-tuning on a specific dataset of 1,000 molecules, starting from the pretrained weights, required a total of 10.52 $\pm$ 0.73 minutes on a single GPU H1000. It is important to notice that the specific dataset has a cumulative nature, grows with each Chemical AL cycle, which eventually leads to a gradual rise in time over iterations.
- For generating one molecule in one GPU H100, the model takes 1.51 $\pm$ 0.15 seconds.
- Chemical AL filters took about 0.05 ± 8.5e-4 seconds per molecule, running on a single core of a 4th Generation Intel Xeon Scalable processor.
- A Chemical AL cycle, comprising a specific dataset fine-tuning, a generation of 3,500 molecules, and their consecutive chemical AL filtering, required a total of 1 hour and 42 minutes $\pm$ 8 minutes.
- An affinity AL cycle requires n Chemical AL cycles, followed by LigPrep (Schrödinger, 2025) for ligand preparation and Glide docking (Halgren et al., 2004). The ligprep and docking for 1.000 molecules on 3 targets took 48.58 $\pm$ 1.37 minutes, parallelised on 30 cores. Thus, an Affinity AL cycle with $n = 10$ Chemical AL cycles will take approximately 17 hours and 43 minutes.
- Assuming a regular execution of our multi-target generative workflow, including 10 Affinity AL cycles, the expected time would be of 177 hours and 15 minutes ($\approx$ 7.39 days).