Drug Repurposing for COVID-19

Angela Kralevska[1], Ivana Vichentijevikj[2] and Monika Simjanoska Misheva[1]
[1]*Faculty of Computer Science and Engineering, Ss. Cyril and Methodius University, Rugjer Boshkovikj 16, P.O. 393, 1000, Skopje, North Macedonia*
[2]*iReason LLC,3rd Macedonian Brigade 37, 1000, Skopje, North Macedonia*

Keywords: Drug Repurposing, COVID-19, Deep Learning, Molecular Docking, Large Language Models, Virtual Screening.

Abstract: The COVID-19 pandemic has underscored the urgent need for rapid, adaptable drug discovery strategies capable of delivering viable therapeutic candidates in compressed time frames. Drug repurposing offers a cost-effective and time-efficient alternative to de novo drug development by identifying new indications for existing compounds. This study presents an integrated computational framework that combines deep learning–based drug–target interaction prediction, molecular docking, and large language model (LLM)–driven molecule generation for COVID-19 drug repurposing. Models from the DeepPurpose library were trained and fine-tuned on benchmark datasets (Davis and KIBA) and a COVID-19–specific high-throughput screening dataset for SARS-CoV-2 3CLPro, enabling large-scale virtual screening of candidate compounds. Molecular docking identified oleanolic acid and other compounds with strong predicted affinity to the SARS-CoV-2 main protease (3CL-Pro), while the DrugGen workflow generated sequence-conditioned SMILES and was able to recover known bioactive molecules such as Carfilzomib. These results show that combining deep learning–based screening, docking-based structural validation, and LLM-guided candidate generation can rapidly surface existing drugs with potential for repurposing. Rather than proposing new chemical matter, this work provides a reproducible in silico pipeline to prioritize repurposable compounds for follow-up in urgent infectious disease settings.

1 INTRODUCTION

Drug repurposing is a process of discovering new applications for authorized or investigational medications that are not covered by the original medical indication (Bakshi et al., 2025). This strategy is of growing importance to the pharmaceutical industry, as de novo drug discovery remains both costly and time-consuming. Recent large-scale economic analyses estimate that the median research and development (R&D) cost per newly approved drug is approximately $708 million, with mean costs exceeding $1.3 billion after accounting for the cost of capital and discontinued development programs (Mulcahy et al., 2025). Moreover, the full development timeline from discovery to approval typically spans over a decade. By contrast, drug repurposing offers a fast and cost-effective means for drug candidate discovery.

When an unexpected medical scenario arises, such as the coronavirus disease (COVID-19), the need for new drugs becomes inevitable, and drug repurposing is one of the most convenient approaches towards their discovery. In March 2020, the COVID-19 outbreak has been declared a global pandemic by the World Health Organization (WHO). The COVID-19 pandemic has caused more than 452 million infections worldwide, with more than 6 million deaths, as of March 2022 (Organization, 2022). Several drugs have been evaluated in clinical trials, with Remdesivir approved by the FDA for hospitalized patients requiring supplemental oxygen, and Paxlovid and molnupiravir authorized for emergency use.

Deep learning has become a central computational paradigm for drug repurposing, enabling scalable virtual screening and representation learning from heterogeneous molecular, structural, and biomedical data. Recent comprehensive reviews frame computational repurposing strategies into broader paradigms, including *target-based*, *phenotype-based*, and *predictive modeling* approaches (Tanoli et al., 2025). Within this framework, predictive modeling methods—largely driven by deep learning—operationalize target-based repurposing by learning drug–target interaction (DTI) patterns directly from molecular and protein representations. Frameworks such

Kralevska, A., Vichentijevikj, I. and Misheva, M. S.
Drug Repurposing for COVID-19.
DOI: 10.5220/0014372200004070

In *Proceedings of the 19th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2026) - Volume 2: BIODEVICES, BIOIMAGING, BIOINFORMATICS*, pages 633-640
ISBN: 978-989-758-802-0; ISSN: 2184-4305

as DeepChem[1] and the DeepPurpose[2] library provide standardized implementations of these predictive models. Molecular docking has complemented these approaches by serving as a structure-based validation and prioritization step, enabling the refinement of structure–activity relationships and the ranking of high-affinity candidates for further evaluation (Kralevska et al., 2021; Malar et al., 2024; Yasir et al., 2024).

More recently, large language models (LLMs) and multi-agent systems have introduced a new paradigm, combining structured molecular data with unstructured biomedical literature. One significant development is the **DrugAgent** architecure (Inoue et al., 2025), which introduces a multi-agent system for drug–target interaction (DTI) prediction. It includes an AI agent using machine learning (DeepPurpose), a knowledge graph (KG) agent leveraging databases such as DrugBank, STITCH, CTD, and DGIdb, and a literature-based search agent. Another notable advancement is **DrugReAlign** (Wei et al., 2024)system. This framework utilizes multi-source prompt engineering to feed LLMs with structural, interactional, and target-specific cues from PDB and other databases. The resulting predictions are validated via large-scale molecular docking using AutoDock Vina, showcasing a balance between reliability and exploration of novel DTIs. In parallel, the pathway-centric pipeline by (Zunzunegui Sanz et al., 2025) propose an LLM-based hypothesis validation pipeline that focuses on pathway-based repurposing. Their study demonstrates that carefully structured prompts and models like GPT-4o and DeepSeek can accurately validate or refute computationally predicted drug-disease associations. Building on these advances, this work presents a comparative and integrative computational framework for COVID-19 drug repurposing that combines deep learning–based DTI prediction, LLM-assisted molecular generation, and structure-based molecular docking for structure-based validation. Specifically, we evaluate established predictive modeling approaches implemented within the DeepPurpose framework alongside a streamlined LLM-driven pipeline, enabling a direct comparison between conventional encoder–decoder DTI models and generative language model–based candidate proposal.

[1]https://github.com/deepchem/deepchem

[2]https://github.com/kexinhuang12345/DeepPurpose

2 MATERIALS AND METHODS

2.1 Methods for Deep Learning Approach

In this study, deep learning models were used to predict drug-target interactions (DTIs) by learning representations directly from compound and protein inputs. Compounds were provided as SMILES strings or molecular fingerprints, while target proteins were represented by amino acid sequences. The predicted interaction scores were used to prioritize candidate compounds for downstream analysis.

2.1.1 Datasets

Four datasets were used in this study. Two benchmark drug–target affinity datasets, Davis (Davis et al., 2011) and KIBA (Pahikkala et al., 2015; He et al., 2017), were used for pretraining and model selection. A COVID-19–specific high-throughput screening bioassay for SARS-CoV-2 3CLPro inhibition (PubChem AID 1706)[3] was used for fine-tuning in a binary classification setting. Finally, the CAS COVID-19 Antiviral Candidate Compounds dataset[4] was used for large-scale virtual screening.

Davis contains kinase inhibitor interactions with measured dissociation constants (K_d), covering 442 targets and 72 compounds (Davis et al., 2011). KIBA integrates multiple bioactivity types (K_i, K_d, IC50) into a unified KIBA score; we used the filtered version containing 229 targets and 2111 compounds with at least 10 interactions per entity (Pahikkala et al., 2015; He et al., 2017).

2.1.2 Drug Representations

Ligands were represented using multiple molecular encodings. Sequence-based representations included SMILES strings, which are commonly used as direct input to deep learning models (Öztürk et al., 2018; Hirohara et al., 2018). Structure-based representations comprised molecular fingerprints, including key-based fingerprints (e.g., PubChem FP) and hashed fingerprints such as Morgan (1024 bits, radius 2) and Daylight (2048 bits) fingerprints, which encode circular and path-based substructures, respectively (Morgan, 1965; Rogers and Hahn, 2010).

Learned representations included word2vec-inspired embeddings (SMILES2vec, Mol2vec) and deep learning–based encoders such as autoencoders

[3]https://pubchem.ncbi.nlm.nih.gov/bioassay/1706

[4]https://www.cas.org/covid-19-antiviral-compounds-dataset

and transformer models (Goh et al., 2017; Jaeger et al., 2018; Lim et al., 2018; Xue et al., 2020). Finally, graph-based representations were used to model molecules as atom–bond graphs (David et al., 2020).

2.1.3 Protein Representations

Target proteins were represented primarily by their amino acid sequences. Depending on the selected model configuration, proteins were encoded either directly from sequence data or using derived features such as amino acid composition (AAC). When structural information was required for downstream analysis, predicted protein structures generated by AlphaFold were used (Jumper et al., 2021).

2.1.4 Models

DeepPurpose is a deep learning framework for drug–target interaction prediction based on encoder–decoder architectures. The framework takes as input a compound representation and a protein representation, encodes them separately, and combines the resulting embeddings to predict binding affinity or interaction scores. Figure 1 shows the DeepPurpose DTI prediction pipeline used in this study.

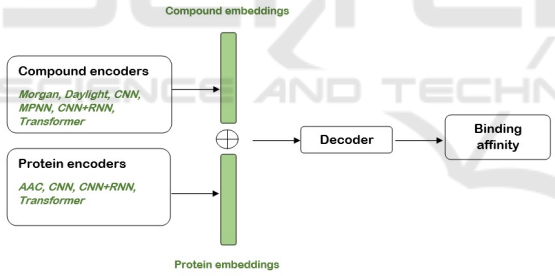

Figure 1: Drug-Target Interaction Prediction pipeline - DeepPurpose.

The DeepPurpose library has implementation of 8 compound encoders and 7 protein encoders ranging from fingerprints to different deep neural networks. In this study, six compound encoders were evaluated, including Multi-Layer Perceptrons (MLP) on Morgan and Daylight fingerprints, Convolutional Neural Networks (CNN) on SMILES strings, Message Passing Neural Networks (MPNN), and transformer-based encoders. Three protein encoders were used: MLP on amino acid composition (AAC), CNN on amino acid sequences, and transformer-based encoders.

2.1.5 Model Training, Fine-Tuning, and Virtual Screening

DeepPurpose models were trained in three stages. First, multiple compound–protein encoder combinations were evaluated on the Davis dataset (70%/10%/20% train/validation/test split) as a regression task optimized using mean squared error, and the best-performing configuration was selected. This configuration was then trained on the larger KIBA dataset using the same setup for 100 epochs.

For COVID-19 adaptation, selected pretrained models were fine-tuned on the 3CLPro QFRET bioassay (PubChem AID 1706) as a binary classification task using the recommended activity threshold of 15, with an 80%/10%/10% split, batch size 128, learning rate 0.001, and 10 training epochs. Performance was evaluated using AUROC, AUPRC, F1 score, and cross-entropy loss.

Finally, virtual screening was performed on the CAS COVID-19 Antiviral Candidate Compounds dataset. Compounds provided in `.sdf` format were converted to SMILES using Open Babel (O'Boyle et al., 2011). Both fine-tuned and non-fine-tuned models were applied, and predicted scores were aggregated using mean, max, and mean–max strategies. To reduce redundancy, screening was repeated on a reduced set of unique compounds generated with PubChemPy.[5]

2.2 Evaluation Metrics

2.2.1 Metrics for Regression Models

Regression performance was evaluated using mean squared error (MSE), Pearson's correlation coefficient, and concordance index (CI), which are standard metrics for continuous prediction tasks. MSE measures the average squared difference between predicted and observed values. The equation 1 is used to calculate the MSE.

$$MSE = \frac{1}{n}\sum_{i=1}^{n}\left(Y_i - \hat{Y}_i\right)^2 \quad (1)$$

Pearson's correlation coefficient quantifies the linear association between predicted and observed values. The formula for calculating the Pearson Correlation is shown on equation 2.

$$\rho_{X,Y} = \frac{cov(X,Y)}{\sigma_X \sigma_Y} \quad (2)$$

The concordance index (CI) evaluates the consistency between the predicted and true ranking of

[5]https://github.com/mcs07/PubChemPy

binding affinity values and is commonly used in drug–target interaction studies.

2.2.2 Metrics for Binary Classification Models

Evaluation metrics for binary classification models implemented in the DeepPurpose library are Area Under the Receiver Operating Characteristic curve (AUROC), Area Under the Precision-Recall Curve (AUPRC), F1 score, and Cross-entropy loss. AUROC and AUPRC evaluate model performance across different decision thresholds, while the F1 score captures the balance between precision and recall. Cross-entropy loss measures the discrepancy between predicted class probabilities and true labels and is commonly used for probabilistic classification models.

2.3 Methods for LLM-Based Molecule Generation Workflow

To complement conventional deep learning and molecular docking approaches, we incorporate a novel workflow based on large language models (LLMs) for structure-aware drug repurposing. The proposed pipeline utilizes DrugGen (Sheikholeslami et al., 2024), a generative model pretrained on molecular data, capable of producing novel SMILES strings conditioned on a target protein sequence.

The workflow starts with the amino acid sequence of SARS-CoV-2 3CLPro as input to DrugGen, which generates candidate inhibitor molecules in SMILES format. Generated compounds are screened against PubChem to identify previously reported molecules, which are retained for downstream analysis. These candidates are then evaluated using AutoDock Vina to assess binding affinity and spatial compatibility with 3CLPro, following the same docking protocol described in Section 2.4.

This approach enables the generation of structurally diverse, previously unseen compounds without relying on predefined drug libraries. It is particularly suitable for rapid early-stage screening in pandemic or low-data scenarios. A schematic representation of the LLM-based drug repurposing workflow is shown in Figure 2.

2.4 Molecular Docking for Drug Repurposing

Molecular docking was employed as a structure-based validation step to evaluate the binding modes and affinities of selected candidate compounds against the SARS-CoV-2 main protease (3CLPro). Docking experiments were performed using AutoDock Vina (Trott and Olson, 2010), which predicts ligand–protein binding poses and corresponding binding energies using a stochastic global optimization strategy combined with local refinement.

2.4.1 Protein and Ligand Preparation

The crystallographic structure of SARS-CoV-2 3CLPro (PDB ID: 6M2N) was retrieved from the RCSB Protein Data Bank.[6] All co-crystallized ligands, water molecules, and heteroatoms were removed prior to docking. Polar hydrogens were added, and the protein structure was converted from `.pdb` to `.pdbqt` format using AutoDock Tools.

Ligand structures were obtained from PubChem in `.sdf` format and converted to `.pdb` using PyMOL.[7] Subsequently, ligands were converted to `.pdbqt` format with AutoDock Tools. The docking set included compounds prioritized by deep learning and virtual screening, along with reference antiviral drugs Molnupiravir and Paxlovid for comparison.

2.4.2 Docking Setup and Execution

Binding sites were defined based on the co-crystallized ligands present in the 3CLPro structure. A single grid box encompassing all four catalytic cavities was used to ensure coverage of all relevant binding regions. The grid box dimensions were set to $(64, 126, 122)$ and centered at coordinates $(-48, -29, 27)$. The grid box position is shown in Figure 3

AutoDock Vina was executed with default scoring parameters. For each ligand, multiple binding poses were generated and ranked by predicted binding affinity. The pose with the lowest binding energy was selected for further analysis.

2.4.3 Post-Docking Analysis

Docking results were analyzed based on predicted binding affinities and interaction patterns. Selected ligand–protein complexes were visualized using PyMOL to inspect binding orientations and interactions with key catalytic residues of 3CLPro. Compounds exhibiting binding affinities comparable to or stronger than reference drugs were prioritized for further discussion.

[6] https://www.rcsb.org/
[7] https://pymol.org/

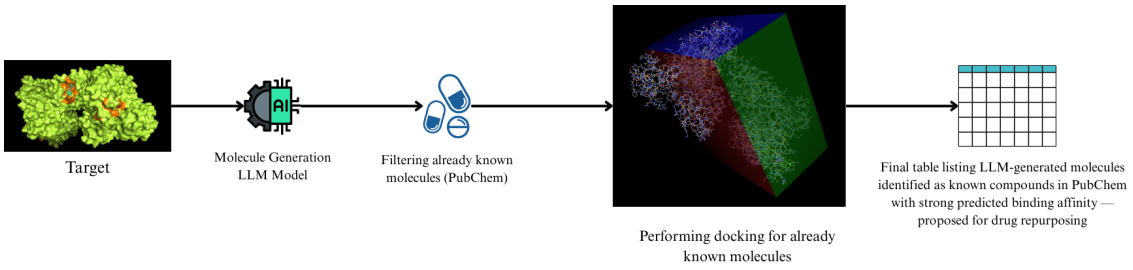

Figure 2: Overview of the proposed LLM-based workflow for COVID-19 drug repurposing.

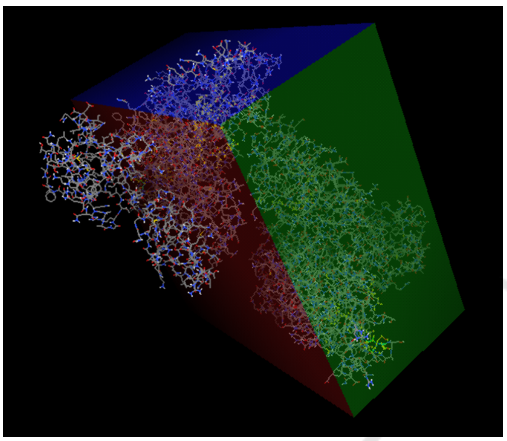

Figure 3: Grid box placement on the 3CLPro protein structure used for molecular docking.

3 RESULTS

Eight DeepPurpose model configurations were evaluated on the Davis dataset using different combinations of compound and protein encoders. The quantitative results are summarized in Table 1. Most models achieved comparable performance across all evaluation metrics. The lowest mean squared error (0.23) was obtained by the Morgan fingerprint–CNN encoder combination, while the Transformer–AAC configuration showed substantially lower performance. This weaker performance may be attributed to limitations of amino acid composition (AAC), which removes sequence order and motif information that could be important for binding, potentially leading to a mismatch between the input representation and the capacity of Transformer-based encoders.

The selected model was subsequently trained on the KIBA dataset. After training, the model achieved a mean squared error of 0.18, a Pearson correlation of 0.86, and a concordance index of 0.87.

Three models pretrained on the Davis dataset were further fine-tuned for SARS-CoV-2 3CLPro inhibition. Evaluation results for the fine-tuned models are presented in Table 2. AUROC values ranged from 0.71 to 0.77, with the Morgan+AAC configuration achieving the highest F1 score. The ROC curves corresponding to these models are shown in Figure 4.

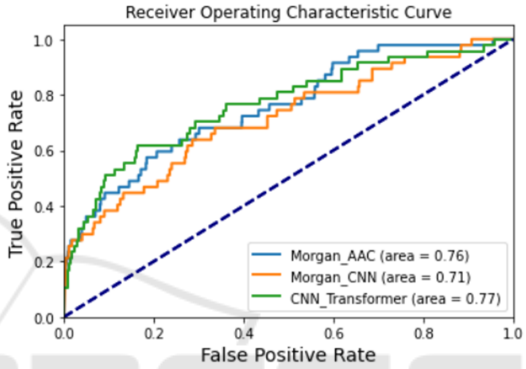

Figure 4: ROC curve from the evaluation of the fine-tuned models with the Davis dataset.

The model pretrained on the KIBA dataset was also fine-tuned on the same task. This model achieved an AUROC of 0.73, an AUPRC of 0.16, an F1 score of 0.07, and a cross-entropy loss of 1.3.

Virtual screening was performed on the CAS COVID-19 Antiviral Candidate Compounds dataset using both fine-tuned and non–fine-tuned models. Screening was additionally repeated on a reduced dataset consisting of 1052 unique compounds. Complete ranking results for all screening configurations are provided in Supplementary Material (Section 6).

With the fine-tuned version of the model pretrained on KIBA dataset only one drug was predicted to be suitable candidate for drug repurposing, namely Neodymium, with 0.58 probability to interact with the 3CLPro.

After the performed virtual screening on the CAS COVID-19 dataset and its reduced version, few compounds whose 3D conformer is available on PubChem have been selected for molecular docking. Docking scores are summarized in Table 3. Paxlovid and oleanolic acid showed the strongest binding affinities (-9.2 kcal/mol). Several other compounds, including aurintricarboxylic acid, sofosbuvir, amidogen, and C. I. Mordant Violet 16, also exhibited stronger pre-

Table 1: Results from evaluation of the different DeepPurpose models trained on the Davis dataset.

Compound encoder	Protein encoder	MSE	Pearson Correlation	Concordance index
CNN	***Transformer***	0.24	0.84	0.88
MPNN	***AAC***	0.31	0.79	0.86
MPNN	***CNN***	0.27	0.82	0.87
Morgan	***CNN***	0.23	0.84	0.89
Morgan	***AAC***	0.24	0.84	0.88
Daylight	***AAC***	0.25	0.84	0.89
CNN	***CNN***	0.24	0.84	0.89
Transformer	***AAC***	0.73	0.32	0.65

Table 2: Results from evaluation of the fine-tuned models pre-trained with the Davis dataset.

Compound encoder	Protein encoder	AUROC	AUPRC	F1	Cross-entropy loss
CNN	***Transformer***	0.77	0.24	0.26	1.43
Morgan	***CNN***	0.71	0.22	0.29	1.20
Morgan	***AAC***	0.76	0.26	0.32	1.28

Table 3: Results from the molecular docking with 3CLPro.

Ligand	**Affinity**
Oleanoic acid	-9.2
Paxlovid	-9.2
Aurintricarboxylic acid	-8.8
Sofosbuvir	-8.5
Amidogen	-8.3
C. I. Mordant Violet 16	-7.9
Molnupiravir	-7.3
Uridine pyruvate	-7.1
Tiracizine	-6.9
Thymidine-5'-phosphate	-6.8
[Hydroxy-(4-methylphenoxy) phosphoryl]formic acid	-5.9
13-Oxabicyclo[10.1.0]tridecane	-5.7
Piperazine-1-carbothioamide	-4.2

dicted binding than molnupiravir.

Oleanolic acid showed interactions with His41 and Cys44 residues of 3CLPro, consistent with binding within the active site. The corresponding protein–ligand complex is shown in Figure 5.

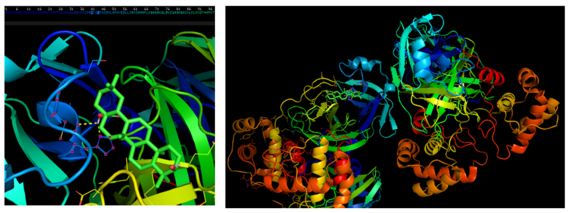

Figure 5: Complex of 3CLPro and Oleanoic acid with selected ligand sites in PyMOL.

Results from the generative DrugGen-based repurposing pipeline are reported in Table 4. After excluding non-therapeutic molecules like FAD, Carfilzomib exhibited the strongest predicted binding affinity among the generated candidates.

Table 4: Docking scores for top DrugGen-generated compounds targeting SARS-CoV-2 3CLPro.

Compound	Docking score (kcal/mol)
FAD	-10.940
Carfilzomib	-9.505
Flualprazolam	-7.897
Amlexanox	-7.366
2-Chloroadenosine	-7.358
CHEMBL4210738	-7.116
Midazolam	-7.102
2'-Deoxyisoguanosine	-7.053

4 DISCUSSION

This study integrates three complementary components for COVID-19 drug repurposing: deep learning–based drug–target interaction (DTI) prediction, structure-based docking, and LLM-driven molecule generation. On the Davis benchmark, most DeepPurpose encoder combinations achieved similar predictive performance (Table 1). These results fall within the range of values commonly reported for Davis in the literature, as documented by established DTI methods such as DeepDTA and Affinity2Vec (Öztürk et al., 2018; Thafar et al., 2022). Rather than serving as a direct comparison, this contextualization indicates that the selected DeepPurpose configurations provide a reasonable and representative baseline for downstream repurposing analyses.

Results from fine-tuning on the SARS-CoV-2 3CLPro bioassay (AUROC ≈0.71–0.77; Table 2), are

expected when converting biochemical screening data into a binary task, where thresholding and experimental variability can introduce label noise and class imbalance effects. Consequently, these scores are best interpreted as a prioritization signal rather than definitive evidence of inhibition.

Docking was used as a downstream structural filter for shortlisted candidates. Under a fixed AutoDock Vina protocol, several screened compounds achieved binding scores comparable to reference antivirals, with oleanolic acid among the strongest-ranked candidates (Table 3). While docking energies are approximate and depend on preparation and scoring assumptions, they provide a consistent structure-based sanity check and facilitate pose-level inspection (Trott and Olson, 2010).

The LLM-based workflow enables target-conditioned generation of candidate SMILES without reliance on large predefined libraries (Sheikholeslami et al., 2024). Unlike large-scale virtual screening, which requires enumerating and scoring extensive compound collections, this approach proposes candidates directly from the target sequence and thus operates as a fast, database-light screening route. Recovering known molecules such as Carfilzomib among the top docked hits suggests that LLM-guided generation can contribute meaningfully to repurposing-oriented discovery. When combined with lightweight database matching and docking, this strategy provides an efficient complement to predictive screening methods.

5 LIMITATIONS AND FUTURE WORK

This study is limited by the assumptions of the deep learning and LLM-based components. Fine-tuning relied on a binary SARS-CoV-2 3CLPro assay with a fixed activity threshold, which may introduce class imbalance and label noise and limit the ability to capture graded potency differences. In addition, CAS screening prioritized candidates solely by predicted interaction scores, without explicit ADMET, toxicity, or drug-likeness filtering. The DrugGen pipeline was applied in a repurposing-oriented manner by matching generated SMILES against PubChem to identify previously reported compounds; however, the workflow does not perform additional checks on molecular novelty or chemical validity beyond those reported in the original DrugGen study.

Future work will include the application of explicit ADMET and drug-likeness filtering during CAS screening to refine candidate prioritization. In addition, the LLM-driven filtering stage will be extended to annotate whether candidate compounds have documented FDA or EMA approval, providing better context for their repurposing potential. Ultimately, experimental validation will be required to confirm inhibitory activity and assess repurposing potential.

6 CONCLUSION

This work evaluates two complementary computational routes for COVID-19 drug repurposing. A deep learning–based drug–target interaction approach was first used to prioritize candidate compounds, followed by structure-based docking for validation, highlighting oleanolic acid as a top-ranked candidate. In parallel, an LLM-driven molecule generation strategy produced target-conditioned candidates that were similarly assessed by docking, identifying Carfilzomib among the strongest-scoring compounds. By applying a shared docking protocol as a common validation step, this study demonstrates how predictive screening and generative modeling can provide complementary insights for repurposing-oriented discovery. While further experimental validation is required, these results illustrate the potential of combining established computational methods with emerging LLM-based approaches to support rapid candidate prioritization in time-sensitive settings.

FUNDING

Funded by the European Union under Horizon Europe project (ChatMED Grant Agreement ID: 101159214). Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

SUPPLEMENTARY MATERIAL

Supplementary Material 1
https://docs.google.com/spreadsheets/d/1k_GThkAaIvucGeLHova4PqDhtIfffwKr/edit?usp=sharing&ouid=112078374933413143169&rtpof=true&sd=true.

REFERENCES

Bakshi, A., Gangopadhyay, K., Basak, S., De, R. K., Sengupta, S., and Dasgupta, A. (2025). Integrating state-space modeling, parameter estimation, deep learning, and docking techniques in drug repurposing: a case study on covid-19 cytokine storm. *Journal of the American Medical Informatics Association*, page ocaf035.

David, L., Thakkar, A., Mercado, R., and Engkvist, O. (2020). Molecular representations in ai-driven drug discovery: a review and practical guide. *Journal of Cheminformatics*, 12.

Davis, M. I., Hunt, J. P., Herrgard, S., Ciceri, P., Wodicka, L. M., Pallares, G., Hocker, M., Treiber, D. K., and Zarrinkar, P. P. (2011). Comprehensive analysis of kinase inhibitor selectivity. *Nature biotechnology*, 29(11):1046–1051.

Goh, G., Hodas, N., Siegel, C., and Vishnu, A. (2017). Smiles2vec: An interpretable general-purpose deep neural network for predicting chemical properties. *arXiv*.

He, T., Heidemeyer, M., Ban, F., Cherkasov, A., and Ester, M. (2017). Simboost: a read-across approach for predicting drug–target binding affinities using gradient boosting machines. *Journal of cheminformatics*, 9(1):1–14.

Hirohara, M., Saito, Y., Koda, Y., Sato, K., and Sakakibara, Y. (2018). Convolutional neural network based on smiles representation of compounds for detecting chemical motif. *BMC Bioinformatics*, 19.

Inoue, Y., Song, T., Wang, X., Luna, A., and Fu, T. (2025). Drugagent: Multi-agent large language model-based reasoning for drug–target interaction prediction. *arXiv*.

Jaeger, S., Fulle, S., and Turk, S. (2018). Mol2vec: Unsupervised machine learning approach with chemical intuition. *Journal of chemical information and modeling*, 58 1:27–35.

Jumper, J. M., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Zídek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Ballard, A., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., Back, T., Petersen, S., Reiman, D. A., Clancy, E., Zielinski, M., Steinegger, M., Pacholska, M., Berghammer, T., Bodenstein, S., Silver, D., Vinyals, O., Senior, A. W., Kavukcuoglu, K., Kohli, P., and Hassabis, D. (2021). Highly accurate protein structure prediction with alphafold. *Nature*, 596:583 – 589.

Kralevska, A., Velichkovska, M., Cicimov, V., Eftimov, T., and Simjanoska, M. (2021). Finding potential inhibitors of covid-19. In *BIOINFORMATICS*, pages 110–117.

Lim, J., Ryu, S., Kim, J., and Kim, W. (2018). Molecular generative model based on conditional variational autoencoder for de novo molecular design. *Journal of Cheminformatics*, 10.

Malar, D. S., Verma, K., Prasanth, M. I., et al. (2024). Network analysis-guided drug repurposing strategies targeting lpar receptor in the interplay of covid, alzheimer's, and diabetes. *Scientific Reports*, 14:4328.

Morgan, H. L. (1965). The generation of a unique machine description for chemical structures-a technique developed at chemical abstracts service. *Journal of Chemical Documentation*, 5:107–113.

Mulcahy, A., Rennane, S., Schwam, D., Dickerson, R., Baker, L., and Shetty, K. (2025). Use of clinical trial characteristics to estimate costs of new drug development. *JAMA Network Open*, 8(1):e2453275–e2453275.

O'Boyle, N. M., Banck, M., James, C. A., Morley, C., Vandermeersch, T., and Hutchison, G. R. (2011). Open babel: An open chemical toolbox. *Journal of cheminformatics*, 3(1):1–14.

Organization, W. H. (2022). Who coronavirus (covid-19) dashboard. https://covid19.who.int/. Accessed: 2022-03-15.

Pahikkala, T., Airola, A., Pietilä, S., Shakyawar, S., Szwajda, A., Tang, J., and Aittokallio, T. (2015). Toward more realistic drug–target interaction predictions. *Briefings in bioinformatics*, 16(2):325–337.

Rogers, D. and Hahn, M. (2010). Extended-connectivity fingerprints. *Journal of Chemical Information and Modeling*, 50(5):742–754. PMID: 20426451.

Sheikholeslami, M., Mazrouei, N., Gheisari, Y., Fasihi, A., Irajpour, M., and Motahharynia, A. (2024). Druggen: Advancing drug discovery with large language models and reinforcement learning feedback. *arXiv*.

Tanoli, Z., Fernández-Torras, A., Özcan, U. O., et al. (2025). Computational drug repurposing: approaches, evaluation of in silico resources and case studies. *Nature Reviews Drug Discovery*, 24(7):521–542.

Thafar, M. A., Alshahrani, M., Albaradei, S., Essack, M., and Bajic, V. B. (2022). Affinity2vec: Drug–target binding affinity prediction through representation learning, graph mining, and machine learning. *Scientific Reports*, 12(1):4751.

Trott, O. and Olson, A. J. (2010). Autodock vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *Journal of computational chemistry*, 31(2):455–461.

Wei, J., Fu, X., Zeng, X.-X., Wang, L., and Cao, D. (2024). Drugrealign: a multisource prompt framework for drug repurposing based on large language models. *BMC Biology*.

Xue, D., Zhang, H., Xiao, D., Gong, Y., Chuai, G., Sun, Y., Tian, H., Wu, H., Li, Y.-K., and Liu, q. (2020). X-mol: large-scale pre-training for molecular understanding and diverse molecular analysis.

Yasir, M., Park, J., Han, E.-T., Park, W. S., Han, J.-H., and Chun, W. (2024). Drug repositioning via graph neural networks: Identifying novel jak2 inhibitors from fda-approved drugs through molecular docking and biological validation. *Molecules*, 29(6).

Zunzunegui Sanz, I., Otero-Carrasco, B., and Rodríguez-González, A. (2025). Accelerating drug repurposing with ai: The role of large language models in hypothesis validation. *bioRxiv*.

Öztürk, H., Ozkirimli, E., and Ozgur, A. (2018). Deepdta: Deep drug-target binding affinity prediction. *Bioinformatics*, 34.