

Target-Specific Novel Molecules with their Recipe: Incorporating Synthesizability in the Design Process

Sowmya Ramaswamy Krishnan, Navneet Bung, Rajgopal Srinivasan, Arijit Roy*
TCS Research (Life Sciences Division), Tata Consultancy Services Limited, Hyderabad
500081, India

*Corresponding author – roy.arijit3@tcs.com

Abstract

Application of Artificial intelligence (AI) in drug discovery has led to several success stories in recent times. While traditional methods mostly relied upon screening large chemical libraries for early-stage drug-design, the AI-based approaches can help identify novel target-specific molecules by sampling from a much larger chemical space. Although this has increased the possibility of finding diverse and novel molecules from previously unexplored chemical space, this has also posed a great challenge for medicinal chemists to synthesize at least some of the AI-designed novel molecules for experimental validation. To address this challenge, in this work, we propose a novel forward synthesis-based generative AI method, which is used to explore the synthesizable chemical space. The method uses a structure-based drug design framework, where the target protein structure and a target-specific seed fragment from co-crystal structures can be the initial inputs. A random fragment from a purchasable fragment library can also be the input if a target-specific fragment is unavailable. Then a template-based forward synthesis route prediction and molecule generation is performed in parallel using the Monte Carlo Tree Search (MCTS) method where, the subsequent fragments for molecule growth can again be obtained from a purchasable fragment library. To the best of our knowledge, this is the first model to utilize MCTS for forward synthesis route prediction. The rewards for each iteration of MCTS are computed using a drug-target affinity (DTA) model based on the docking pose of the generated reaction intermediates at the binding site of the target protein of interest. With the help of the proposed method, it is now possible to overcome one of the major obstacles posed to the AI-based drug design approaches through the ability of the method to design novel target-specific synthesizable molecules.

Keywords: *Forward synthesis, Monte Carlo Tree Search, Generative model, Drug design, Structure-based framework*

Introduction

Deep learning models have transformed the space of drug discovery by enabling molecule generation, property prediction and on-the-fly molecular optimization in the recent years (Olivecrona et al., 2017; Segler et al., 2018; Popova et al., 2018; Bung et al., 2020; Born et al., 2021; He et al., 2021; Krishnan et al., 2021; Krishnan et al., 2022; Das et al., 2023; Vangala et al., 2023). Several testaments to the success of deep neural networks in identifying clinically potent and efficacious small molecules have also emerged (Zhavarov et al., 2019; Stokes et al., 2020). Although the extent of chemical space currently explored by deep generative models remains unquantified (Bender and Cortes-Ciriano, 2021; Vogt, 2022), it is safer to say that they can explore a sufficiently diverse and relevant space of potentially novel small molecules with the aid of techniques such as transfer learning (TL) and reinforcement learning (RL). With the models currently available in literature, it is technically possible to design drug-like small molecules for any given target protein of interest using either the protein sequence (Grechishnikova, 2021; Born et al., 2022a; Born et al., 2022b) or structure (Aumentado-Armstrong, 2018; Skalic et al., 2019; Xu et al., 2021; Krishnan et al., 2022; Isert et al., 2022) or a set of known inhibitors (Segler et al., 2018; Bung et al., 2020; Krishnan et al., 2021a) identified *a priori*. However, further experimental validation of the designed small molecules requires that the molecule be easily synthesizable.

Several studies have tried to address the problem of synthetic accessibility of generated small molecules essentially through one of the two methods: 1) quantification of synthesizability with metrics and 2) *in silico* prediction of the probable synthesis route(s) for the molecule. Traditionally, the synthesizability of a small molecule is quantified in terms of the contribution of fragments toward the overall complexity of the molecular structure (Ertl and Schuffenhauer, 2009). Few examples of such metrics include synthetic accessibility score, SAS (Ertl and Schuffenhauer, 2009), and SCScore (Coley et al., 2018). Recently, metrics such as RAscore (Thakkar et al., 2021), SYBA (Voršilák et al., 2020), GASA (Yu et al., 2022), and CMPNN (Li and Chen, 2022) utilize machine learning and deep learning models, with a varied set of molecular features underlying each model, to predict synthesizability. Based on a comparative analysis (Skoraczynski et al., 2022) on the performance of these metrics alongside existing retrosynthesis prediction models such as AiZynthFinder (Genheden et al., 2020), it is evident that these metrics need to be finetuned according to the target model for synthetic route prediction. This suggests that the generalizability of these metrics is to be carefully analyzed to discriminate molecules according to their degree of synthesizability. Several examples of molecules with conflicting synthesizability predictions have also been reported earlier (Gao and Coley, 2020; Skoraczynski et al., 2023), suggesting a need to improve existing metrics and reach consensus.

On the other hand, multiple groups have approached the problem of synthetic route prediction for small molecules, resulting in multiple approaches for both forward and retrosynthetic route prediction. The existing approaches can be categorized into either reaction template-based or template-free methods. The template-dependent approaches (Gottipati et al., 2020; Li et al., 2022; Button et al., 2019; Horwood and Noutahi, 2020; Noh et al., 2022) require a set of expertly curated chemical reaction rules, which generalize transformations from reactants to products. The template-independent approaches (Bradshaw et al., 2019; Schwaller et al., 2019) try to learn these rules from a large corpus of reactions in organic chemistry, therefore limiting their generalization capability to novel reactant datasets. Based on literature, it has been observed that the existing template-dependent methods explicitly limit the reactions to single-step or two-step synthesis

problems (Li et al., 2020) to avoid traversing the highly diverse space of available fragments. Two of the existing methods (Button et al., 2019; Noh et al., 2022) have enforced generation of molecules which are similar to existing approved drugs or known synthesizable small molecules, thereby significantly restricting the exploration of chemical space by the model. While other approaches have tried to overcome this by allowing user-defined starting reactants (Li et al., 2022), the molecules are again generated atom-by-atom through policy learning, which has been shown before to reduce chemical space exploration in previous studies (Podda et al., 2020; Chen et al., 2021). Only one of the existing approaches (Li et al., 2022) has tried to provide a structure-level perspective of the evolution of the molecule and its interactions with the binding site residues through 3D conformation generation. Although the coordinate space of the generation process was anchored to the binding site, the lack of minimization and analysis of short contacts post-generation leads to incomplete validation of the generated conformations.

In this work, we introduce a synthesis-aware generative model which can perform template-based forward synthesis route prediction and molecule generation simultaneously using the Monte Carlo Tree Search (MCTS) method. To the best of our knowledge, this is the first model to utilize MCTS for forward synthesis route prediction. The rewards for each iteration of MCTS are computed using a drug-target affinity (DTA) model dependent on the docking pose of the generated reaction intermediates at the binding site of the target protein of interest. The model utilizes a tethered docking approach wherein, crystallographic fragments can be provided as the starting point for molecule generation, and the corresponding fragment coordinates in the crystal structure can be used as the seed pose for docking. In this way, the model is designed to consider any available information regarding the protein structure while designing molecules, and predict their forward synthesis route using libraries of globally purchasable reactants. The application of the method is showcased through generation of molecules specific to the human cAMP-dependent protein kinase A (PKA).

Materials and Methods

Target-specific fragments were considered as initial starting points

Fragment-based drug discovery can identify very small molecules also known as fragments that can bind to the target protein, using biophysical and biochemical methods. Fragment screening hits from existing crystal structures were considered as initial reactants to grow the molecule based on chemical reaction rules, till the target-specific molecules with high bioactivity are obtained. To make the molecule target specific, bioactivity of the intermediate and final molecules were calculated (as described in the section titled “***Incorporation of structural information from the target protein during MCTS reward calculation***”). In this article, fragment hits are obtained from Protein Data Bank (PDB)-deposited structures for the target protein of interest. In real world scenario, such fragment hits can also be obtained from fragment-based screening assay. To generalize the method for cases where no initial fragment hit is available, generic, random fragments obtained from a purchasable fragment library, such as Enamine, were also used as starting fragments to initialize the generation process.

Datasets utilized for molecule generation

Reaction templates: A dataset of 121 reaction templates was obtained in reaction SMIRKS format from an earlier study (Horwood and Noutahi, 2020). The validity of these reaction SMIRKS was first checked using the rdchemReactions module in RDKit. The RDKit Rxn objects of valid reaction templates (Table S1) were saved for mapping the reactants to their probable template reactions.

Subsequent reactants: The dataset of subsequent reactants was curated from the Enamine comprehensive building blocks catalog (Grygorenko et al., 2020). 1,169,054 reactants part of the catalog were pre-processed using RDKit to map each reactant SMILES to their corresponding Enamine product identifier. The Enamine database was chosen to ensure that the reactants corresponding to the intermediates identified as part of the forward synthesis route of any molecule are easily purchasable globally.

Unimolecular and bimolecular reactions: The USPTO MIT dataset (Jin et al., 2017) with organic reactions extracted from patent literature was considered for reaction template mapping. Since previous studies have already shown that majority of all existing reactions can be categorized into either unimolecular or bimolecular reactions (Hartenfeller et al., 2011; Konze et al., 2019; Gottipati et al., 2020; Horwood and Noutahi, 2020; Ucak et al., 2022), only these reaction categories were isolated from the dataset using custom Python scripts to parse the reaction SMILES format. The resultant set of 116,633 unimolecular reactions and 350,893 bimolecular reactions were considered for reaction template mapping. Unimolecular reactions are those where a single reactant undergoes transformation to obtain the final product. Bimolecular reactions are those where two reactants are required for the reaction to happen.

Preparing the Enamine reactants dataset for model training

The reactants from Enamine globally purchasable dataset were mapped to the probable reaction templates which can be applied on them, using the RDKit Rxn objects and the RunReactants function. Such an approach can reduce the search space and can allow focusing on only the subsequent reactants that can be mapped with previous/starting reactant and the corresponding reaction template. With the SMILES of Enamine building blocks dataset as input, the template mapping step could map 99% (1,157,468 reactants) of data to their corresponding probable reaction templates. However, upon analyzing the number of reactants mapped per reaction template it was observed that the dataset exhibited severe class imbalance, with the least populated class (reaction template) containing only 27 reactants mapped to it (Table S2). To handle this, Butina clustering (Butina, 1999) was employed to cluster reactants mapped to each reaction template, using Tanimoto coefficient (TC) as the distance metric. With a within-cluster similarity of 60% (TC of at least 0.6) as the criteria, only the representative reactants and singleton reactants from the clustering results were chosen for each template to create the dataset (Table S2). Singleton reactants without any clustering were also included in the dataset to retain the reactant diversity after downsampling. This downsampled dataset included 281,765 (24.34%) of the template-mapped reactants obtained.

Preparing the USPTO reactants dataset for model training

The procedure explained in the previous subsection, makes sure to map reactants from Enamine library with their respective reaction templates. Similarly, the reactants from 13.78% (64,439) of the USPTO reactions, both unimolecular and bimolecular, could be mapped to the reaction templates using RDKit. The similarity of products obtained with the templates identified were compared to the actual products from the USPTO reactions using ECFP4 fingerprints (Rogers and Hahn, 2010) and Tanimoto coefficient (TC) as the similarity metric, to ensure that the products obtained are highly similar, if not identical, to the actual products reported. From the mapping, more than 90% of the products obtained with RDKit were identical to the actual products (supplementary information - Fig. S1). Only the Enamine reactants and USPTO reactions with templates available were used for building the generative models.

Models involved in the MCTS formulation

(a) Reaction template selection policy: Given a starting reactant of interest or a reaction intermediate as the input, the reaction template selection policy is used to decide a suitable reaction template for the starting reactant or reaction intermediate, to grow the MCTS search tree. Hence, a template classification model is pre-trained with the downsampled Enamine reactants dataset as input. For any reactant, this pre-trained policy can provide a probability distribution over the possible reaction templates that can be used at that step of the search. With the atomistic molecular graph representation of the reactant molecule, a graph convolution network (GCN) (Kipf and Welling, 2017) was used to predict the reaction template. The architecture included a single GCN layer followed by a fully connected layer and a final dense layer with softmax activation. The output hidden state from the GCN layer was subject to global mean pooling to get graph-level readout from the model. The downsampled Enamine reactants dataset of 281,765 molecules was split in a 9:1:1 ratio using the stratified split method in scikit-learn. Categorical cross-entropy loss was used to train the model and the parameters were optimized using the Adam optimizer. The initial learning rate was set to 0.005 and the model was trained for 100 epochs with a dropout rate of 0.5 to prevent overfitting. The model was implemented with the Deep Graph Library (DGL) framework (Zheng et al., 2021) using PyTorch as the backend. The architecture of the model described above is depicted in supplementary information (Fig. S2).

(b) Second reactant selection policy: After the prediction of most probable reaction template with the reaction template selection policy, it is further categorized into unimolecular or bimolecular reactions. In case of bimolecular reaction templates, a second reactant is necessary to obtain the final product by the application of the chosen template. Hence, a second reactant selection policy was pre-trained with the template-mapped bimolecular reactions from the USPTO dataset (Jin et al., 2017). For this model, the first reactant is represented as a molecular graph and the predicted bimolecular reaction template is represented using a binary one-hot encoded vector. The graph-level readout for the molecular graph is obtained from a graph convolutional (GCN) layer and the reaction template vector is concatenated to this graph embedding vector (Fig. 1). The concatenated vector is passed through two fully connected layers with ReLU activation and a final dense layer with tanH activation, to get a vector of 35 properties normalized between -1 and 1 (Gottipati et al., 2020). This property vector is considered as the output representation of the probable second reactant for the bimolecular reaction. The second reactant selection policy is trained with mean squared error (MSE) loss function and the parameters are optimized with the Adam optimizer. The initial learning rate was set to 0.005 and the model was trained for 100 epochs with a dropout rate of 0.5 to prevent overfitting. The model was implemented with the Deep Graph Library (DGL) framework (Zheng et al., 2021) using PyTorch as the backend. The architecture of the second reactant selection policy model is provided in supplementary information (Fig. S3).

(c) kNN search: The property vector predicted by the second reactant selection policy is subject to k-nearest neighbors (kNN) search against the pre-computed and normalized property vectors of all Enamine reactants from the downsampled dataset, to obtain a set of k possible second reactants for the bimolecular reaction (Gottipati et al., 2020). Euclidean distance between the predicted property vector and the Enamine property vectors is considered as the distance metric for the kNN step. The k reactants obtained are used to expand the search tree further. Since distance calculations against the entire downsampled dataset of Enamine reactants poses a linear time complexity ($O(n)$), searching against only the Enamine reactants mapped to the predicted reaction template was used as a heuristic during the kNN search.

Incorporation of structural information from the target protein during MCTS reward calculation

The drug-target affinity (DTA) prediction model based on our previous studies (Krishnan et al., 2022; Krishnan et al., 2023), is used to compute rewards during MCTS. The model takes as input an extended connectivity interaction fingerprint (ECIF) of the protein-ligand complex, and predicts the corresponding binding affinity in log-scale (Sánchez-Cruz et al., 2021). To obtain the ECIF fingerprints during MCTS, the predicted reaction intermediates and products were docked on-the-fly to the target protein using the rDock program (Ruiz-Carmona et al., 2014). This program was chosen since it supports significantly faster docking in both tethered and untethered modes, enabling consideration of available crystallographic fragment positions as seed coordinates for docking pose identification. In this way, the MCTS process is made more specific to a target protein of interest by pre-training the DTA model to be specific to the protein of interest. The DTA model so obtained was used to compute the rewards of reaction intermediates and products obtained during the expansion and simulation steps of MCTS, which are detailed below.

Monte Carlo Tree Search (MCTS) formulation for forward synthesis route prediction

MCTS is a reinforcement learning-based process to utilize information from simulations of the future steps and select current actions by learning from the simulated episodes (Browne et al., 2012; Segler et al., 2017a). MCTS involves four major steps namely, selection, expansion, rollout or simulation, and update (Segler et al., 2017a). To generate molecules along with their forward synthesis route using the template-mapped reactants as the starting point, two policies (reaction template selection policy and second reactant selection policy) were pre-trained with the datasets prepared. In the MCTS search tree, every node is a reactant or intermediate or product molecule, and every edge is a reaction template applied to arrive at the child nodes or leaf nodes.

Selection (Tree policy): In every step of MCTS, the Upper Confidence Bounds (UCB) condition (Auer and Ortner, 2010), also known as the tree policy, is used to select the action (reaction template) at the current state (node or reactant) based on the number of times the node was visited during simulation. A node or reactant that was visited multiple times, will have a larger visit count and hence, more probability to be generated in the path of maximum reward during MCTS. The action or reaction template which led to the most visited node will be chosen for the current state based on the value computed as follows (1):

$$a_t = \arg \max_{a \in A(s_t)} \left(\frac{Q(s_t, a)}{N(s_t, a)} + c \frac{\sqrt{\log N(s_{t-1}, a)}}{N(s_t, a)} \right) \quad (1)$$

Here, $Q(s_t, a)$ is the reward obtained for the node from the reinforcement learning critic, $N(s_t, a)$ is the visit count for the current state or node, $N(s_{t-1}, a)$ is the visit count for the parent of the current node, and c is a parameter controlling the exploration-exploitation trade-off during the MCTS process. In this study, a c value of 2 is used for all results reported. The rewards for the child nodes obtained from each step are computed in terms of their predicted log-scale bioactivity, using a pre-trained drug-target affinity (DTA) prediction model from our previous studies (Sánchez-Cruz et al., 2021; Krishnan et al., 2022).

Expansion: During expansion from a selected node (starting reactant or reaction intermediate), the probable child nodes are predicted in a two step process. First the top

10 probable reaction templates which can be applied on the selected node are predicted using the pre-trained reaction template selection policy. From the predicted templates, the unimolecular and bimolecular templates are separated. For the bimolecular templates, the most probable second reactants to complete the reaction are predicted using the pre-trained second reactant selection policy and kNN search over the Enamine library. Finally, the probable child nodes for every predicted reaction template are obtained using the RDKit RunReactants function.

Rollout or Simulation: Each predicted child node from the expansion step is subject to 10 rounds of simulation, to gauge their reward potential (estimated future bioactivity) before expanding the MCTS search tree further. The average reward of the simulations is thresholded to a cut-off bioactivity value to select the child candidate for further expansion. It is to be noted that for every reward calculation, the reaction intermediate is docked on-the-fly to the target protein to obtain the ECIF representation, which is used as input to the DTA model. Hence, information on the structure of the target protein and the exact binding site location are used during every step of the MCTS search, to guide the forward synthesis process.

Update: During update, the rewards ($Q(s_t, a)$) and visit counts ($N(s_t, a)$) of the edges and nodes of the search tree, respectively, are updated based on the simulation results from the current state (s_t) to the root node. The average reward from the simulation is assigned to the node from where the simulation was performed. The state-action pair which produced the maximum average reward post-simulation is selected for the next expansion-simulation-update iteration of MCTS. After several iterations of MCTS expansion, simulation and update, the path containing reactants and products with the most number of visits and highest average reward is reported, with the final product being the generated molecule of interest. The complete formulation of the forward synthesis prediction problem using MCTS is provided as a schematic below (Fig. 1).

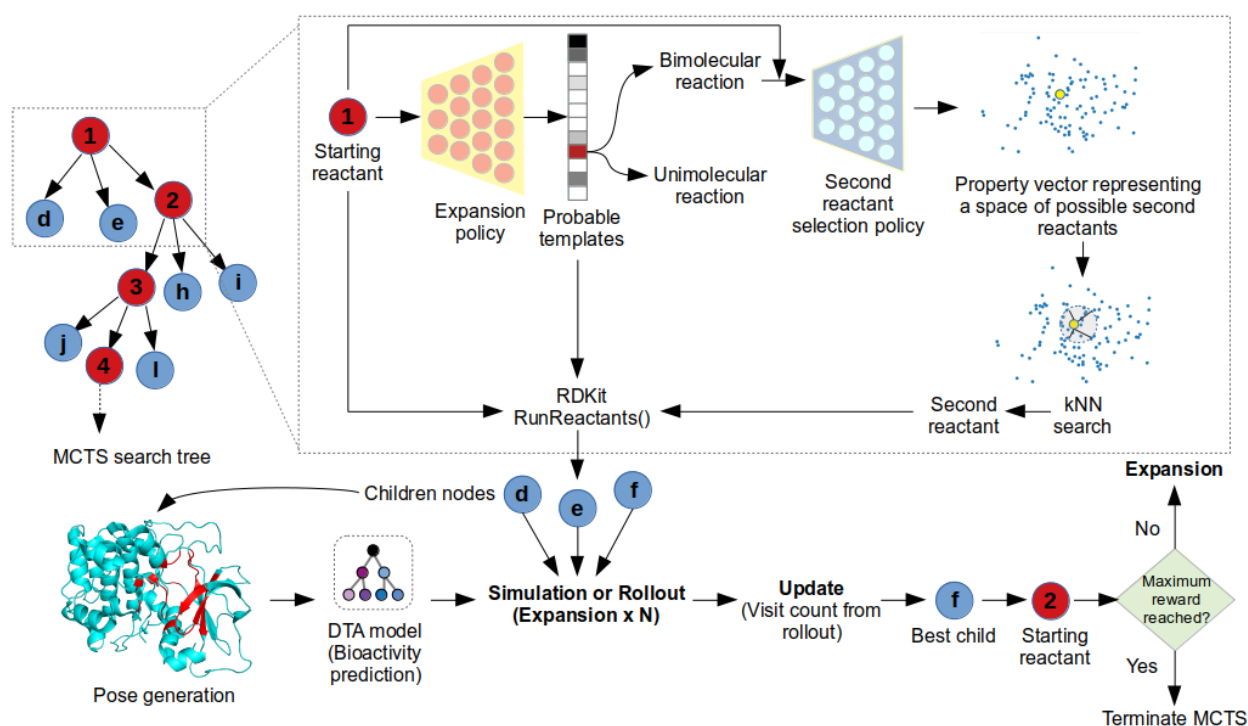


Figure 1: A detailed schematic of the MCTS formulation of forward synthesis prediction cum small molecule generation problem. The three pre-trained models involved are the expansion policy, second reactant selection policy and DTA model. Information regarding

the structure of the target protein is utilized for pose generation through standard docking programs, for MCTS reward calculation (predicted bioactivity from the DTA model).

Application of MCTS for generation of molecules targeting human cAMP-dependent Protein Kinase A (PKA)

To generate molecules specific to the human cAMP-dependent Protein Kinase A (PKA) along with their forward synthesis routes using MCTS, the set of starting reactants from Enamine downsampled dataset, with identity to the substructures present in PKA inhibitors were identified using RDKit substructure matching functions. Further, with the dataset of 537 PKA inhibitors and their bioactivity collected from ChEMBL database, the DTA model was pre-trained to be specific to PKA, with the PKA-inhibitor complex structures predicted using the GNINA docking program. To perform the untethered docking and for DTA model development, the structure of PKA with PDB ID: 3OWP was considered, and the binding site coordinates were defined based on the ligand bound to the protein in the crystal structure.




Figure 2: The crystallographic fragments and their corresponding scaffolds chosen for the tethered docking-based MCTS experiments from (a) 3NX8; (b) 3OOG; (c) 5BX6; (d) 5BX7; (e) 5N3Q.

For the tethered docking experiments, five experimental structures of the human cAMP-dependent protein kinase A (PKA) with a co-crystallized fragment were finalized (PDB IDs: 3NX8, 3OOG, 5BX6, 5BX7, 5N3Q). These five structures were chosen due to the identical binding sites of the fragments involved. For each structure, rDock parameter files were prepared for tethered docking in an automated fashion and integrated with the DTA model for bioactivity prediction. The fragments present in each crystal structure chosen, along with the scaffold from the fragment used for tethered docking are shown above (Fig. 2). The human cAMP-dependent protein kinase A (PKA)-specific DTA model was found to have an r value of 0.72 and an RMSE of 0.66 on the test set. The regression plot obtained with the PKA-specific DTA model is provided below (Fig. 3) along with the line of best fit obtained.




Figure 3: Regression plot showing the performance of the PKA-specific DTA model for the test dataset.

Results and Discussion

Case Study against human cAMP-dependent protein kinase A: The human cAMP-dependent protein kinase A (PKA) is a heterotetramer composed of two regulatory subunits (R) and two catalytic subunits (C). Binding of cAMP to the regulatory subunit leads to dissociation of the catalytic subunits from the regulatory subunits. The free catalytic subunits of PKA (PRKACA) are further responsible for phosphorylation of a wide variety of cellular substrates involved in metabolism, gene expression and cellular proliferation (Cheung et al., 2015). Mutations in the catalytic subunit have been identified to lead to Cushing's syndrome, a kidney disorder leading to excessive cortisol production, and tumors such as fibrolamellar hepatocellular carcinoma (FL-HCC) (Cheung et al., 2015). Increase in cAMP levels and PKA substrate levels has also been shown to be linked to incidence of hepatitis C virus infections in humans (Farquhar et al., 2008). Due to the unusual mode of regulation of PRKACA in comparison with other kinases, it has been considered as a vital target for design of kinase inhibitors in cancers (Herberg and Taylor; 1993; Wen and Taylor, 1994; Viht et al., 2007).

Table 1: Performance of the three models developed in this study – reaction template selection policy, second reactant selection policy and the DTA model is tabulated below. Further details on model performance are provided in supplementary information (Section S1 and Fig. S4).

Model	Performance metrics
Reaction template selection policy	ROC = 0.88
Second reactant selection policy	$R_p = 0.90$ $r^2 = 0.81$ MAE = 0.115
DTA model	$r^2 = 0.72$

In the following sections, the results obtained from the application of the proposed forward synthesis prediction method using Monte Carlo Tree Search to design potential small molecules against the catalytic subunit of human PKA are discussed in detail. The performance of the pre-trained models to perform MCTS are summarized above (Table 1). Also, a detailed comparison of the proposed method with existing forward synthesis prediction models in literature is provided in the supplementary information (Table S3).




Figure 4: Examples of four tethered docking-based experiments with different crystallographic starting reactants obtained from structures with PDB IDs: (a) 3NX8, (b) 5BX7, and (c) 5BX6. Each reaction step contains the name of the reaction template

chosen and the enamine identifiers of the second reactants and products, wherever possible.

Structural perspective of the forward synthesis route generated based on crystallographic fragments

Crystallographic fragments can be considered as valuable starting points for exploration of the interacting groups within the binding site of a drug target. To retain the information available from such screening studies, a tethered docking-based approach was used. This approach would anchor the scaffolds of reaction intermediates and products to the scaffold position of the fragment hits selected. Based on comparative analysis, the rDock program (Ruiz-Carmona et al., 2014) was found to provide better speed with both tethered and untethered docking calculations than GNINA (McNutt et al., 2021) (Supplementary information – Section S2). A few examples of the molecules generated from the tethered docking based MCTS runs, along with their forward synthesis routes are shown below (Fig. 4).




Figure 5: Examples of scaffolds of generated reaction intermediates and products anchored to the scaffold from the crystallographic fragment hit (green sticks). The scaffold alignments were obtained based on the tethered docking results collected using the rDock program during the MCTS experiments.

On comparison of the docking poses obtained for the starting reactant, reaction intermediates and the final product with the initial seed coordinates, it was observed that the centre of mass of the scaffold of interest in all pairs of molecules was aligned within 0.5 Å to that of the crystallographic fragment. A few examples of the aligned molecules are shown below (Fig. 5). The results indicate that the tethered docking approach could successfully constrain the scaffold of interest to the desired region of the kinase binding site. Analysis of the intermolecular interactions observed over the course of the forward synthesis route indicated a steady increase in bioactivity with the growth of the molecule in the binding site. Due to the abundance of non-polar amino acids in the human PKA binding site, most of the interactions observed between the generated molecules and the binding

site residues were hydrophobic, followed by hydrogen bonds and stacking interactions (Fig. 6). This was found to be in accordance with the predominant interactions of human PKA inhibitors reported in literature (Toyota et al., 2022) (Fig. S5). The following binding site residues were found to consistently interact with both existing inhibitors and generated small molecules from this study: Arg18B, Arg19B, Asn20B, Ala21B, His23B, Leu49A, Gly50A, Thr51A, Gly52A, Ser53A, Phe54A, Gly55A, Arg56A, Val57A, Ala70A, Lys72A, Ile73A, Leu74A, Gln84A, Thr88A, Glu91A, Val104A, Met120A, Glu121A, Tyr122A, Val123A, Glu127A, Asp166A, Lys168A, Glu170A, Asn171A, Leu173A, Thr183A, Asp184A, Gly186A, Phe187A, Phe327A, Tyr330A.


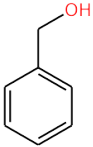
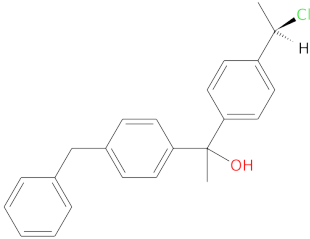
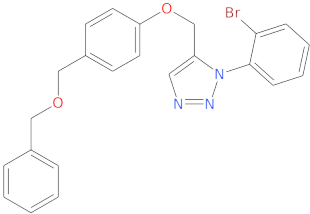
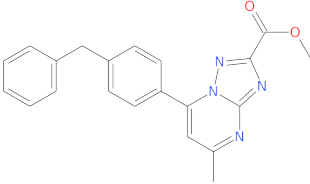


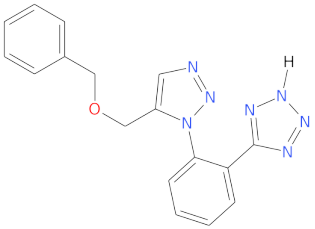
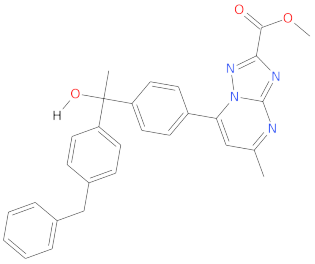
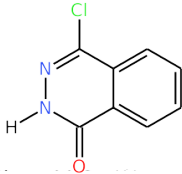
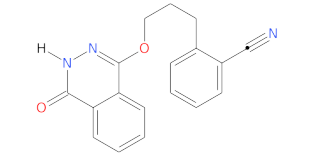
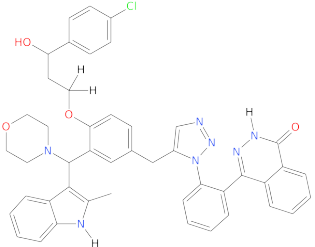
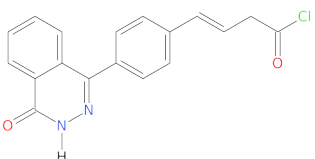
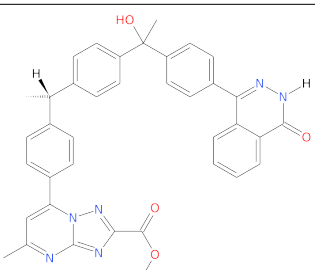
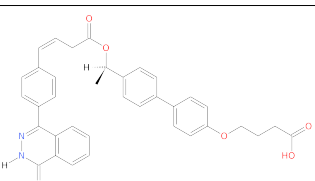
Figure 6: Examples of interactions between the generated small molecules (cyan sticks) and the PKA binding site residues (green lines). Hydrogen bonds and ionic interactions are highlighted using yellow dotted lines, while residues involved in hydrophobic interactions are also shown for each molecule. The predicted bioactivity of the generated small molecule is also provided in red figures within parantheses. The figures were generated using PyMOL.

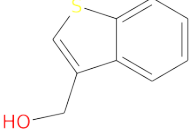
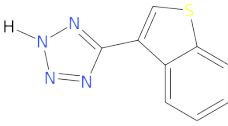
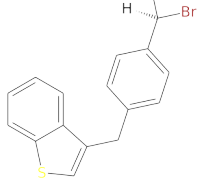
Molecular diversity observed among multiple runs of MCTS using same starting reactant

By performing multiple runs of MCTS using same starting reactant, different final products were obtained for each run. The average similarity of different small molecules obtained with the same starting reactant was found to be 0.65, in terms of Tanimoto coefficient. This indicates that even across different runs starting from the same reactant, the generated molecules show significant molecular diversity. Upon comparison of the generated molecules with existing PKA inhibitors curated from ChEMBL, several molecules with better bioactivity and interaction profiles were identified. The results of top 5 MCTS runs for three of the chosen crystallographic starting reactants are tabulated below (Table 2). Interestingly, most of the molecules with better bioactivity also had very low similarity (less than 0.5) with existing PKA inhibitors, further highlighting the diversity of small molecules generated by the MCTS workflow proposed.

Table 2: Results from the MCTS calculations for starting reactant from PDB IDs: 3NX8, 5BX6 and 5BX7. Tethered docking-based MCTS calculations were performed 5 times and the final product molecules obtained are shown along with their predicted bioactivity and maximum similarity with existing PKA actives. The ChEMBL ID of the similar PKA inhibitor and its experimental bioactivity value are also provided.

Starting reactant	No. of reactions in the predicted route	Final product molecule	Predicted bioactivity against PKA	Similarity (TC) to existing PKA actives	Similar ChEMBL molecule (bioactivity)
	3		6.338	0.582	CHEMBL35 44960 (7.13)
	2		6.609	0.461	CHEMBL19 72568 (5.50)
	1		6.137	0.506	CHEMBL20 04872 (6.36)

	3		6.350	0.475	CHEMBL19 77135 (6.99)
	2		6.061	0.525	CHEMBL45 76489 (7.37)
	2		6.161	0.488	CHEMBL19 77135 (6.99)
	4		6.628	0.748	CHEMBL45 76489 (7.37)
	2		6.109	0.552	CHEMBL20 00481 (5.30)
	4		6.075	0.521	CHEMBL20 00481 (5.30)
	3		5.587	0.601	CHEMBL45 76489 (7.37)

	2		5.295	0.427	CHEMBL19 65170 (6.09)
	3		6.006	0.439	CHEMBL19 77135 (6.99)

Performance of generative model with generic and target-specific starting reactants

In addition to utilizing the fragment screening hits from existing crystal structures for molecule generation, random and target-specific starting reactants from a purchasable fragment library, such as Enamine, were also used to initialize the generation process. To obtain the target-specific starting reactants, exact substructure matching was performed to identify fragments from the Enamine library which are enriched in the existing PKA inhibitors. A generic (random) set of starting reactants was also identified and confirmed to be non-overlapping with that of the target-specific reactant dataset. The MCTS workflow was used to design inhibitors with both these reactant datasets to investigate the generalization capability of the approach.




Figure 7: Comparison of generated small molecules from the proposed method using random starting reactants and target-specific starting reactants.

The complete forward synthesis route predicted by MCTS for each molecule shown in first 5 rows of Table 2 is provided in the supplementary information (Figs. S7-S11). Upon comparison of the small molecules generated using random starting reactants and target-specific reactants, only two molecules had above 0.75 Tanimoto similarity between the two datasets (Fig. 7). However, in terms of the distribution of predicted bioactivity values, both datasets showed significant overlap, indicating that the MCTS process samples across a highly diverse region of the chemical space, without compromising on the bioactivity against the target protein of interest (Fig. S6). By embedding the generated small

molecules and the existing PKA inhibitors using t-distributed stochastic neighbor embedding (tSNE), the different chemical spaces occupied by generated small molecules and existing PKA inhibitors could also be clearly visualized (Fig. 8). To further understand the diversity and synthesizability of generated molecules, all final products obtained were compared with ~13.8 million readily purchasable small molecules available in the ZINC20 database. Based on the comparison, only 60 generated molecules had significant similarity (above 0.75 TC) to the ZINC20 compounds (Fig. 9). These results also substantiate the diversity of the generated molecules and their synthesis routes, using only reactants from a purchasable fragment library such as Enamine.

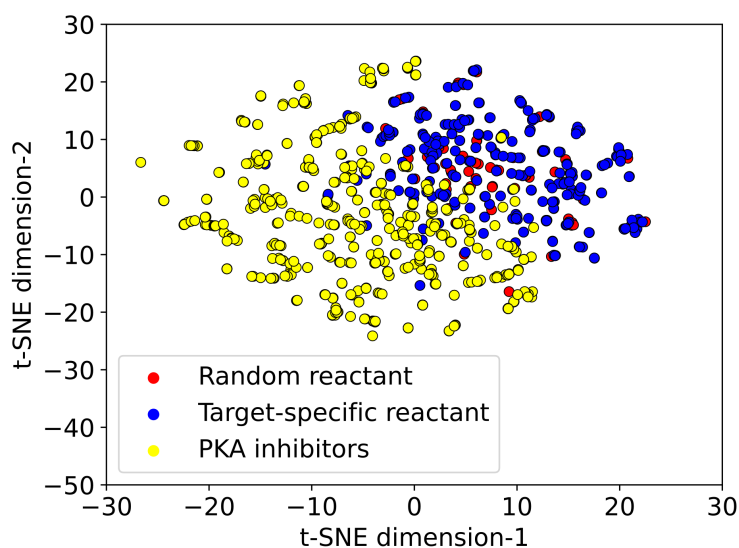


Figure 8: tSNE plot showing the different chemical spaces occupied by existing PKA inhibitors and generated small molecules.




Figure 9: Similarity of the generated small molecules to readily purchasable compounds from the ZINC20 database. The similarity of the small molecules is quantified as Tanimoto coefficients (TC). The ZINC IDs of the highly similar compounds are also provided.

Conclusions

In this study, a novel forward synthesis prediction method for synthesis-aware small molecule generation was proposed using Monte Carlo Tree Search (MCTS). The proposed approach provides a structure-level perspective of the evolution of potential small molecules at the binding site of the target protein. The method can be utilized to design molecules with large, diverse and readily purchasable fragment libraries such as Enamine. The method can also consider crystallographic fragment screening hits as starting points to grow the molecule within the binding site. The proposed approach was validated by designing potential inhibitors against the human cAMP-dependent protein kinase (PKA), which was considered as the target protein of interest. The results indicate that the MCTS-based approach could design highly diverse small molecules with better bioactivity compared to existing inhibitors, along with their predicted forward synthesis route and interaction profile. The generated small molecules sample a completely novel region of the chemical space upon comparison with synthesizable large compound databases and embedding approaches, highlighting the applicability of this method to discover potential novel small molecules with better bioactivity. The ability of the proposed method to design target-specific and synthesizable novel drug molecules can help to overcome of the major challenge posed to the AI-models, which is synthesizability of novel molecules proposed by AI-models.

Authorship contribution statement

Sowmya Ramaswamy Krishnan: Data curation, Validation, Formal analysis, Software, Writing – original draft, Writing – review & editing, Visualization.

Navneet Bung: Conceptualization, Validation, Formal analysis, Investigation, Resources, Writing – review & editing, Supervision.

Rajgopal Srinivasan: Validation, Resources, Writing – review & editing, Supervision.

Arijit Roy: Conceptualization, Validation, Investigation, Resources, Writing – review & editing, Supervision.

Conflicts of Interest

All authors involved are employed by Tata Consultancy Services Limited.

Acknowledgements

We thank our colleagues Mr. Dibyajyoti Das, Dr. Broto Chakraborty, Mr. Sarveswara Rao Vangala, Mrs. Akriti Jain and Dr. Gopalakrishnan Bulusu for helpful discussions.

References

- Auer, P.; Ortner, R. UCB revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*. **2010**, *61*, 1-2.
- Aumentado-Armstrong, T. Latent Molecular Optimization for Targeted Therapeutic Design. *ArXiv*. **2018**. <https://arxiv.org/abs/1809.02032>
- Bender, A.; Cortes-Ciriano, I. Artificial intelligence in drug discovery: what is realistic, what are illusions? Part 2: a discussion of chemical and biological data. *Drug Discov. Today*. **2021**, *26*, 1040-1052.
- Born, J.; Manica, M.; Oskooei, A.; Cadow, J.; Markert, G.; Martínez, M. R. PacMann^{RL}: De novo generation of hit-like anticancer molecules from transcriptomic data via reinforcement learning. *IScience*. **2021**, *24*, 102269.
- Born, J.; Shoshan, Y.; Huynh, T.; Cornell, W. D.; Martin, E. J.; Manica, M. On the Choice of Active Site Sequences for Kinase-Ligand Affinity Prediction. *J. Chem. Inf. Model*. **2022**, *62*, 4295-4299.
- Born, J.; Huynh, T.; Stroobants, A.; Cornell, W. D.; Manica, M. Active Site Sequence Representations of Human Kinases Outperform Full Sequence Representations for Affinity Prediction and Inhibitor Generation: 3D Effects in a 1D Model. *J. Chem. Inf. Model*. **2022**, *62*, 240-257.
- Bradshaw, J.; Paige, B.; Kusner, M. J.; Segler, M. H. S.; Hernández-Lobato, J. M. A model to search for synthesizable molecules. *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. **2019**, *713*, 7937-7949.
- Browne, C.; Powley, E.; Whitehouse, D.; Lucas, S.; Cowling, P. I.; Rohlfshagen, P.; Tavener, S.; Perez, D.; Samothrakis, S.; Colton, S. A Survey of Monte Carlo Tree Search Methods. *IEEE Trans. Comput. Intell. AI Games*. **2010**, *4*, 1-43.
- Bung, N.; Krishnan, S. R.; Bulusu, G.; Roy, A. De novo design of new chemical entities for SARS-CoV-2 using artificial intelligence. *Future Med. Chem*. **2021**, *13*, 575-585.
- Butina, D. Unsupervised Data Base Clustering Based on Daylight's Fingerprint and Tanimoto Similarity: A Fast and Automated Way To Cluster Small and Large Data Sets. *J. Chem. Inf. Comput. Sci*. **1999**, *39*, 747-750.
- Button, A.; Mark, D.; Hiss, J. A.; Schneider, G. Automated de novo molecular design by hybrid machine intelligence and rule-driven chemical synthesis. *Nat. Mach. Intell*. **2019**, *1*, 307-315.
- Chen, B.; Fu, X.; Barzilay, R.; Jaakkola, T. Fragment-based Sequential Translation for Molecular Optimization. *arXiv*. **2021**.
- Cheung, J.; Ginter, C.; Cassidy, M.; Franklin, M. C.; Rudolph, M. J.; Robine, N.; Darnell, R. B.; Hendrickson, W. A. Structural insights into mis-regulation of protein kinase A in human tumors. *Proc. Natl. Acad. Sci. U. S. A*. **2015**, *112*, 1374-9.
- Coley, C. W.; Rogers, L.; Green, W. H.; Jensen, K. F. SCScore: Synthetic Complexity Learned from a Reaction Corpus. *J. Chem. Inf. Model*. **2018**, *58*, 252-261.
- Das, D.; Chakrabarty, B.; Srinivasan, R.; Roy, A. Gex2SGen: Designing Drug-like Molecules from Desired Gene Expression Signatures. *J. Chem. Inf. Model*. **2023**, *63*, 1882-1893.
- Ertl, P.; Schuffenhauer, A. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *J. Cheminform*. **2009**, *1*, 8.
- Farquhar, M. J.; Harris, H. J.; Diskar, M.; Jones, S.; Mee, C. J.; Nielsen, S. U.; Brimacombe, C. L.; Molina, S.; Toms, G. L.; Maurel, P.; Howl, J.; Herberg, F. W.; van Ijzendoorn, S. C. D.; Balfe, P.; McKeating, J. A. Protein kinase A-dependent step(s) in hepatitis C virus entry and infectivity. *J. Virol*. **2008**, *82*, 8797-811.
- Gao, W.; Coley, C. W. The Synthesizability of Molecules Proposed by Generative Models. *J. Chem. Inf. Model*. **2020**, *60*, 5714-5723.
- Genheden, S.; Thakkar, A.; Chadimová, V.; Reymond, J.; Engkvist, O.; Bjerrum, E. AiZynthFinder: a fast, robust and flexible open-source software for retrosynthetic planning. *J. Cheminform*. **2020**, *12*, 70.
- Gottipati, S. K.; Sattarov, B.; Niu, S.; Pathak, Y.; Wei, H.; Liu, S.; Blackburn, S.; Thomas, K.; Coley, C.; Tang, J.; Chandar, S.; Bengio, Y. Learning to Navigate The Synthetically Accessible Chemical Space Using Reinforcement Learning. *Proceedings of the 37th International Conference on Machine Learning Research*. **2020**, *119*, 3668-3679.
- Grechishnikova, D. Transformer neural network for protein-specific de novo drug generation as a machine translation problem. *Sci. Rep*. **2021**, *11*, 321.
- Grygorenko, O. O.; Radchenko, D. S.; Dziuba, I.; Chuprina, A.; Gubina, K. E.; Moroz, Y. S. Generating Multibillion Chemical Space of Readily Accessible Screening Compounds. *IScience*. **2020**, *23*, 101681.
- Hartenfeller, M.; Eberle, M.; Meier, P.; Nieto-Oberhuber, C.; Altmann, K.; Schneider, G.; Jacoby, E.; Renner, S. A collection of robust organic synthesis reactions for in silico molecule design. *J. Chem. Inf. Model*. **2011**, *51*, 3093-8.
- He, J.; You, H.; Sandström, E.; Nittinger, E.; Bjerrum, E. J.; Tyrchan, C.; Czechtizky, W.; Engkvist, O. Molecular optimization by capturing chemist's intuition using deep neural networks. *J. Cheminform*. **2021**, *13*, 26.

Herberg, F. W.; Taylor, S. S. Physiological inhibitors of the catalytic subunit of cAMP-dependent protein kinase: effect of magnesium-ATP on protein-protein interactions. *Biochemistry*. **1993**, *32*, 14015-22.

Horwood, J.; Noutahi, E. Molecular Design in Synthetically Accessible Chemical Space via Deep Reinforcement Learning. *ACS Omega*. **2020**, *5*, 32984-32994.

Isert, C.; Atz, K.; Schneider, G. Structure-based drug design with geometric deep learning. *ArXiv*. **2022**. <https://arxiv.org/abs/2210.11250>

Jin, W.; Coley, C. W.; Barzilay, R.; Jaakkola, T. Predicting Organic Reaction Outcomes with Weisfeiler-Lehman Network. *Proceedings of the 31st Conference on Neural Information Processing Systems*. **2017**, Long Beach, CA, USA.

Kipf, T. N.; Welling, M. Semi-Supervised Classification With Graph Convolutional Networks. *Proceedings of the 5th International Conference on Learning Representations*. **2017**, Toulon, France.

Konze, K. D.; Bos, P. H.; Dahlgren, M. K.; Leswing, K.; Tubert-Brohman, I.; Bortolato, A.; Robbason, B.; Abel, R.; Bhat, S. Reaction-Based Enumeration, Active Learning, and Free Energy Calculations To Rapidly Explore Synthetically Tractable Chemical Space and Optimize Potency of Cyclin-Dependent Kinase 2 Inhibitors. *J. Chem. Inf. Model*. **2019**, *59*, 3782-3793.

Krishnan, S. R.; Bung, N.; Bulusu, G.; Roy, A. Accelerating *De Novo* Drug Design against Novel Proteins Using Deep Learning. *J. Chem. Inf. Model*. **2021**, *61*, 621-630.

Krishnan, S. R.; Bung, N.; Vangala, S. R.; Srinivasan, R.; Bulusu, G.; Roy, A. *De Novo* Structure-Based Drug Design Using Deep Learning. *J. Chem. Inf. Model*. **2022**, *62*, 5100-5109.

Li, B.; Chen, H. Prediction of Compound Synthesis Accessibility Based on Reaction Knowledge Graph. *Molecules*. **2022**, *27*, 1039.

Li, Y.; Pei, J.; Lai, L. Synthesis-driven design of 3D molecules for structure-based drug discovery using geometric transformers. *arXiv*. **2022**. <https://arxiv.org/abs/2301.00167>

McNutt, A. T.; Francoeur, P.; Aggarwal, R.; Masuda, T.; Meli, R.; Ragoza, M.; Sunseri, J.; Koes, D. R. GNINA 1.0: molecular docking with deep learning. *J. Cheminform*. **2021**, *13*, 43.

Noh, J.; Jeong, D.; Kim, K.; Han, S.; Lee, M.; Lee, H.; Jung, Y. Path-Aware and Structure-Preserving Generation of Synthetically Accessible Molecules. *Proceedings of the 39th International Conference on Machine Learning*. **2022**, 162.

Pflug, A.; de Oliveira, T. M.; Bossemeyer, D.; Engh, R. A. Mutants of protein kinase A that mimic the ATP-binding site of Aurora kinase. *Biochem. J*. **2011**, *440*, 85-93.

Podda, M.; Bacciu, D.; Micheli, A. A Deep Generative Model for Fragment-Based Molecule Generation. *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*. **2020**, 108.

Popova, M.; Isayev, O.; Tropsha, A. Deep reinforcement learning for de novo drug design. *Sci. Adv*. **2018**, *4*, eaap7885.

Rogers, D.; Hahn, M. Extended connectivity fingerprints. *J. Chem. Inf. Model*. **2010**, *50*, 742-54.

Ruiz-Carmona, S.; Alvarez-Garcia, D.; Foloppe, N.; Garmendia-Doval, A. B.; Juhos, S.; Schmidtke, P.; Barril, X.; Hubbard, R. E.; Morley, S. D. rDock: a fast, versatile and open source program for docking ligands to proteins and nucleic acids. *PLoS Comput. Biol*. **2014**, *10*, e1003571.

Sánchez-Cruz, N.; Medina-Franco, J. L.; Mestres, J.; Barril, X. Extended connectivity interaction features: improving binding affinity prediction through chemical description. *Bioinformatics*. **2021**, *37*, 1376-1382.

Schwaller, P.; Laino, T.; Gaudin, T.; Bolgar, P.; Hunter, C. A.; Bekas, C.; Lee, A. A. Molecular Transformer: A Model for Uncertainty-Calibrated Chemical Reaction Prediction. *ACS Cent. Sci*. **2019**, *5*, 1572-1583.

Segler, M. H. S.; Preuss, M.; Waller, M. P. Learning to Plan Chemical Syntheses. *ArXiv*. **2017**. <https://arxiv.org/abs/1708.04202>

Segler, M. H. S.; Waller, M. P. Neural-Symbolic Machine Learning for Retrosynthesis and Reaction Prediction. *Chemistry*. **2017**, *23*, 5966-5971.

Segler, M. H. S.; Kogej, T.; Tyrchan, C.; Waller, M. P. Generating Focused Molecule Libraries for Drug Discovery with Recurrent Neural Networks. *ACS Cent. Sci*. **2018**, *4*, 120-131.

Skalic, M.; Sabbadin, D.; Sattarov, B.; Sciabola, S.; De Fabritiis, G. From Target to Drug: Generative Modeling for the Multimodal Structure-Based Ligand Design. *Mol. Pharm*. **2019**, *16*, 4282-4291.

Skoraczynski, G.; Kitlas, M.; Miasojedow, B.; Gambin, A. Critical assessment of synthetic accessibility scores in computer-assisted synthesis planning. *J. Cheminform*. **2023**, *15*, 6.

Stokes, J. M.; Yang, K.; Swanson, K.; Jin, W.; Cubillos-Ruiz, A.; Donghia, N. M.; MacNair, C. R.; French, S.; Carfrae, L. A.; Bloom-Ackermann, Z.; Tran, V. M.; Chiappino-Pepe, A.; Badran, A. H.; Andrews, I. W.; Chory, E. J.; Church, G. M.; Brown, E. D.; Jaakkola, T. S.; Barzilay, R.; Collins, J. J. A Deep Learning Approach to Antibiotic Discovery. *Cell*. **2020**, *180*, 688-702.

Thakkar, A.; Chadimová, V.; Bjerrum, E. J.; Engkvist, O.; Reymond, J. Retrosynthetic accessibility score (RAscore) - rapid machine learned synthesizability classification from AI driven retrosynthetic planning. *Chem. Sci*. **2021**, *12*, 3339-3349.

Toyota, A.; Goto, M.; Miyamoto, M.; Nagashima, Y.; Iwasaki, S.; Komatsu, T.; Momose, T.; Yoshida, K.; Tsukada, T.; Matsufuji, T.; Ohno, A.; Suzuki, M.; Ubukata, O.; Kaneta, Y. Novel protein kinase cAMP-Activated Catalytic Subunit Alpha (PRKACA) inhibitor shows anti-tumor activity in a fibrolamellar hepatocellular carcinoma model. *Biochem. Biophys. Res. Commun.* **2022**, *621*, 157-161.

Ucak, U. V.; Ashyrmamatov, I.; Ko, J.; Lee, J. Retrosynthetic reaction pathway prediction through neural machine translation of atomic environments. *Nat. Commun.* **2022**, *13*, 1186.

Vangala, S. R.; Krishnan, S. R.; Bung, N.; Srinivasan, R.; Roy, A. pBRICS: A Novel Fragmentation Method for Explainable Property Prediction of Drug-Like Small Molecules. *J. Chem. Inf. Model.* **2023**, doi: 10.1021/acs.jcim.3c00689

Viht, K.; Schweinsberg, S.; Lust, M.; Vaasa, A.; Raidaru, G.; Lavogina, D.; Uri, A.; Herberg, F. W. Surface-plasmon-resonance-based biosensor with immobilized bisubstrate analog inhibitor for the determination of affinities of ATP- and protein-competitive ligands of cAMP-dependent protein kinase. *Anal. Biochem.* **2007**, *362*, 268-277.

Vogt, M. Using deep neural networks to explore chemical space. *Expert Opin. Drug Discov.* **2022**, *17*, 297-304.

Voršilák, M.; Kolář, M.; Čmelo, I.; Svozil, D. SYBA: Bayesian estimation of synthetic accessibility of organic compounds. *J. Cheminform.* **2020**, *12*, 35.

Wen, W.; Taylor, S. S. High affinity binding of the heat-stable protein kinase inhibitor to the catalytic subunit of cAMP-dependent protein kinase is selectively abolished by mutation of Arg133. *J. Biol. Chem.* **1994**, *269*, 8423-30.

Xu, M.; Ran, T.; Chen, H. De Novo Molecule Design Through the Molecular Generative Model Conditioned by 3D Information of Protein Binding Sites. *J. Chem. Inf. Model.* **2021**, *61*, 3240-3254.

Yu, J.; Wang, J.; Zhao, H.; Gao, J.; Kang, Y.; Cao, D.; Wang, Z.; Hou, T. Organic Compound Synthetic Accessibility Prediction Based on the Graph Attention Mechanism. *J. Chem. Inf. Model.* **2022**, *62*, 2973-2986.

Zheng, D.; Wang, M.; Gan, Q.; Song, X.; Zhang, Z.; Karypis, G. Scalable Graph Neural Networks with Deep Graph Library. *Proceedings of the 14th ACM International Conference on Web Search and Data Mining.* **2021**, 1141-1142.

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.; Aspuru-Guzik, A. Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nat. Biotechnol.* **2019**, *37*, 1038-1040.