Active Learning to Select the Most Suitable Reagents and One-Step Organic Chemistry Reactions for Prioritizing Target-Specific Hits from Ultralarge Chemical Spaces

Vladimir Kozyrev, François Sindt, Didier Rognan

Active learning to select the most suitable reagents and one-step organic chemistry reactions for prioritizing target-specific hits from ultra-large chemical spaces.

Vladimir Kozyrev,[†] François Sindt[†] and Didier Rognan[†,*]

[†]Laboratoire d'innovation thérapeutique, UMR7200 CNRS-Université de Strasbourg, F-67400 Illkirch, France.

* To whom correspondence should be addressed (phone: +33 3 68 85 42 35, fax: +33 3 68 85 43 10, email: rognan@unistra.fr)

ABSTRACT

Designing chemically novel and synthesizable ligands from the largest possible chemical space is a major issue in modern drug discovery to identify early hits easily amenable to medicinal chemistry optimization. Starting from the sole three-dimensional structure of a protein binding site, we herewith describe a fully automated active learning protocol to propose the commercial chemical reagents and one-step organic chemistry reactions necessary to enumerate target-specific primary hits from ultra-large chemical spaces. When applied in different scenarios (single transform, multiple transforms) addressing chemical spaces of various sizes (from 670 million to 4.5 billion compounds), the method was able to recover up to 98% of virtual hits discovered by an exhaustive docking-based approach while scanning only 5% of the full chemical space. It is therefore applicable to the structure-based screening of trillion-sized chemical spaces at a very high throughput with minimal computational resources.

INTRODUCTION

Virtual screening (VS) is currently experiencing a revolution in early hit identification with the rise of ultra-large chemical spaces in which potent and chemically novel ligands can be searched. [1-3] Instead of relying on the stock of ca. 15 million on-the-shelf commercially available drug-like molecules,[4] there is a clear trend to virtually screen much larger combinatorial spaces of ligands not yet available but easily synthesizable from cheap chemical reagents and hundreds of reproducible and parallelizable organic chemistry reactions.[5-11] Interestingly, several groups using either ligand-based or structure-based VS methods independently reported that virtual screening of ultra-large chemical spaces (> 100 million compounds) tends to yield both higher hit rates and more potent hits[12-27] than that achieved by conventional screening of on-the shelf-libraries.[28] These promising advances push chemical suppliers to offer on-demand chemical spaces of increasing sizes[29] with remarkable synthesis success rates (ca. 80% for Enamine REAL space)[5] to deliver intended compounds at high speed (a few weeks) and high purity (>90%). At the time of writing this manuscript, over trillions compounds are potentially accessible either on-demand (Enamine's REAL: 70 billion, Enamine's xREAL: 2.4 trillion, WuXi's GalaXi: 12 billion, Otava's CHEMryia: 12 billion) or on a "do-it-yourself" approach (Ambinter's Ambrosia: 110 billion, e-molecules eXplore: 5 trillion, ChemSpace's Freedom space: 5.1 billion).[29] Such gigantic chemical spaces cannot be enumerated anymore and require *ad-hoc* algorithms and virtual screening tools[3] for fast and cost-effective screening. With the advent of trillion-sized virtual spaces,[30] even the fastest workflows reach their limit and compromises must be made with respect to a representative fraction of the full chemical space to be scanned while still retaining most of promising hits. Supervised machine learning (ML) approaches,[31-32] particularly within the active learning (AL) framework, hold promises at this task by training numerical models to predict experimental affinities[33] or proxies thereof (e.g. docking scores,[34-43]protein-ligand interactions[44]) from simple ligand descriptors or

protein-ligand docking poses. Typically, a small training set (e.g. 0.5% of the full chemical space) is first randomly chosen to train a ML model to predict the expected biological property, the model being then applied to all remaining compounds. Updating the initial model with a novel acquisition batch composed of user-selected samples (e.g. best scores, most uncertain predictions) and predicting the rest of the compounds, is iteratively performed by AL until a sufficiently large subset (1-10%) of the full chemical space has been scanned. Existing retrospective studies[33-43] agree to conclude that active learning is able to prioritize a large fraction (75-90%) of virtual hits obtained by exhaustive virtual screening after a just few iterations, permitting to scan less to 10% of the full chemical space and to save a lot of computational resources. Interestingly, simple models (e.g. linear regression)[41] seem to perform almost on-par with much more sophisticated methods (e.g. deep neural networks)[35] thereby offering the possibility to screen trillion-sized chemical spaces at low cost. However, true prospective VS screens benefiting from AL acceleration are still very rare[45-46] and remain disappointing with respect to hit rate and hit potency, when compared to exhaustive VS of comparable ultra-large spaces.[1] A possible explanation lies in the fact that docking scores, often used as affinity proxies, are well known to poorly correlate with absolute binding free energies.[47-49]

We therefore recently developed a docking score-independent structure-based virtual screening tool (SpaceDock) exploiting the combinatorial nature of reaction-driven chemical spaces.[16] SpaceDock orients commercial chemical reagents in any binding site of interest and then recombine them into full ligands, directly in the three-dimensional (3D) target space, according to both geometrical and organic chemistry cross-compatibilities. When applied to receptors of pharmaceutical interest, the workflow was able to recover a very close analog of a ground-truth drug (eticlopride), as well as chemically novel compounds with high hit rates and potencies.[16] Although it is yet applicable to billions of on-demand compounds easily synthesizable by 41 one-step organic chemistry reactions, it remains intellectually

unsatisfactory to spend 99.9% of computing resources to evaluate reagents pairs (two-component one-step reactions) or triads (three-component one-step reactions) that do not respect user-defined geometrical and chemical recombination rules. The aim of the present study is therefore to investigate whether AL on simple reagent and reaction descriptors may efficiently prioritize the best primary hits before SpaceDocking. We here relied on two previous virtual screening campaigns of ultra-large spaces (670 million compounds targeting human dopamine D3 receptor,[16] 4.5 billion compounds targeting the CBLB ubiquitin-protein ligase) [50] for which SpaceDock has exhaustively been applied to all chemical space members (sets 1 and 3, **Figure 1**). For both chemical spaces, a pilot study was first applied to 1 million randomly chosen reagent pairs (sets 2 and 4, **Figure 1**) to find the best AL parameters that were next applied to the prediction of hits (pairs of reagents yielding to a SpaceDock success) from the full chemical spaces.

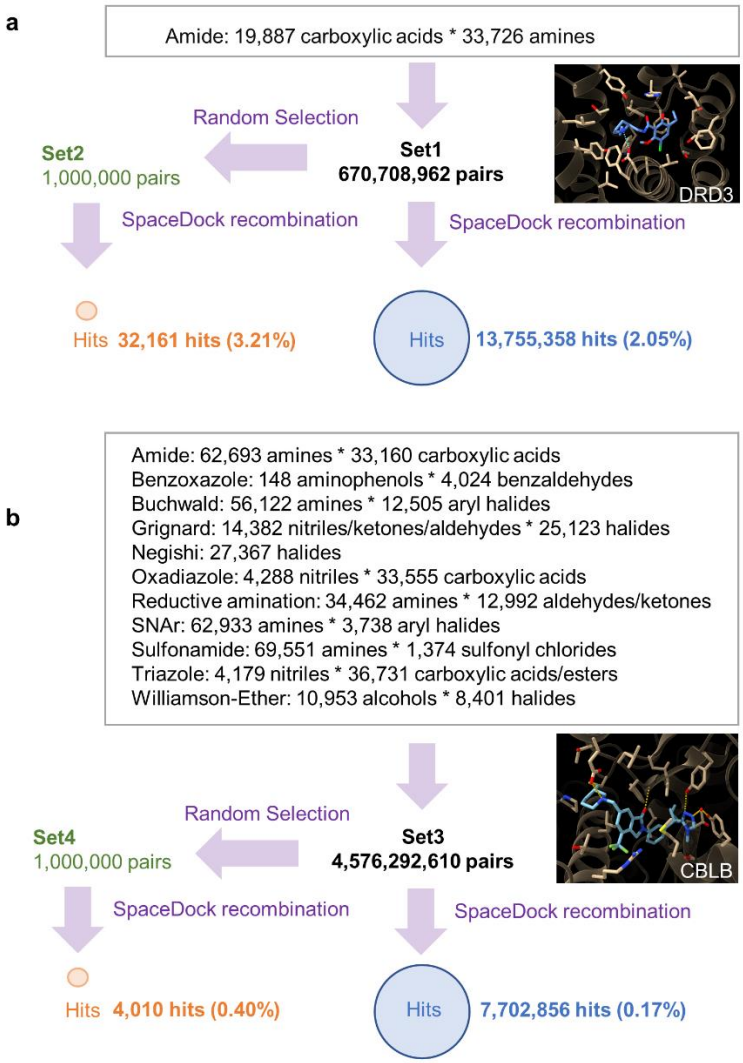

Figure 1. Setting-up datasets for predicting target-specific assembly of reagent pairs. **a)** DRD3 dataset: chemical space of 670 million carboxamides (set 1) obtained by amide bond formation between 19,887 carboxylic acids and 33,726 amines. 13,755,358 pairs of reagents could be successfully assembled by SpaceDock to enumerate full ligands.[16] A random selection of 1,000 carboxylic acids and 1,000 amines affords a set of 1,000,000 pairs (set 2). **b)** CBLB dataset: chemical space of 4.5 billion ligands (set 3) obtained by 11 organic chemistry reactions. 7,702,8566 pairs of reagents could be successfully assembled by SpaceDock to enumerate full ligands. A random selection of 1,000,000 pairs of reagents defines set 4.

MATERIAL AND METHODS

Datasets

DRD3 chemical space. 19,887 commercially available carboxylic acids and 33,726 amines were already docked to the human dopamine D3 receptor (PDB 3PBL),[16] defining a full chemical space of 670,708,962 possible carboxamides (set 1, **Figure 1a**). 13,755,358 recombinations were successfully generated with SpaceDock, in other words pairs of reagents (2D) for which at least one of their docking poses fulfilled the two following conditions: (i) 3D geometrical and chemical compatibility with an amide bond formation, (ii) protein-ligand interaction fingerprint (IFP) similarity of the merged pose to the ground-truth ligand (eticlopride, **Figure 1a**) above well-defined thresholds (full similarity ≥ 0.60, polar interaction similarity ≥ 0.50). Each corresponding pair (carboxylic acid, amine) was assigned a positive label ("1"), other unsuccessful pairs being annotated with a negative label ("0"). A random selection of 1,000 carboxylic acids and 1,000 amines defined a pilot subset of 1 million pairs (set 2, **Figure 1a**), out of which 32,161 pairs (3.21%) had a positive label.

CBLB chemical space. A chemical space of 4.5 billion compounds made of 11 different two-component reactions (set 3, **Figure 1b**, **Supporting Table S1**) was investigated by docking the corresponding reagents to the X-ray structure of the TKB domain of the CBLB E3 ubiquitin-protein ligase (PDB 8GCY).[51] Atomic coordinates of the CBLB protein and chemical reagents were prepared as previously described for the dopamine D3 receptor.[16] 7,702,856 recombinations were successfully generated with SpaceDock, considering IFP minimal similarity to the 8GCY ligand of 0.6 (all interactions) and 0.5 (polar interactions), respectively. Each corresponding triad (reagent 1, reagent 2, reaction) was assigned a positive label ("1"), other unsuccessful triads being labelled with a negative label ("0"). A random selection of pairs

of reagents defined a pilot subset of 1 million pairs (set 4, **Figure 1b**), out of which 4,010 pairs (0.40%) had a positive label according to a slightly modified polar IFP Tanimoto similarity threshold set to 0.25.

SpaceDock methodology. Although described in details elsewhere,[16] we here just briefly recapitulate the basic SpaceDock operations. First, commercial reactants, carefully chosen for each reaction from the stock of Enamine's reagents collection are converted from 2D (SD file) to 3D coordinates (MOL2 file), ionized at pH 7.4 and docked to the target of interest using CCDC GOLD v.2022 docking algorithm[52] with standard parameters. If the commercial reactant has a well-defined stereochemistry, it will be kept upon 2D to 3D conversion for docking. In case of a racemic building block, all stereoisomers are generated and docked to the target. For each reactant, 20 poses were prioritized according to the PLP score. Pairs of reactant poses (400 per enumerated ligand) are then checked for compatibility to be recombined into a final ligand by either a unique reaction (DRD3 sets) or multiple reactions (CBLB sets) according to both geometrical (critical distances, angles and dihedrals) and chemical reactivity criteria.[16] During the recombination step that operates in the protein 3D coordinates frame, exit atoms are removed and atom types/bond orders are directly updated. From here on, we define a SpaceDock hit as any compound successfully enumerated by at least of the 400 possible solutions.

Reagents and reaction descriptors. Each chemical reagent was represented by a Morgan fingerprint (radius 2; size 2,048) and generated using the RDKit library (release 2023.09.6). To avoid a fingerprint that is independent on the order by which reactants are read, bitwise AND and XOR operations were applied each pair of reactant fingerprints and then concatenated. Thus, the first part of the pair

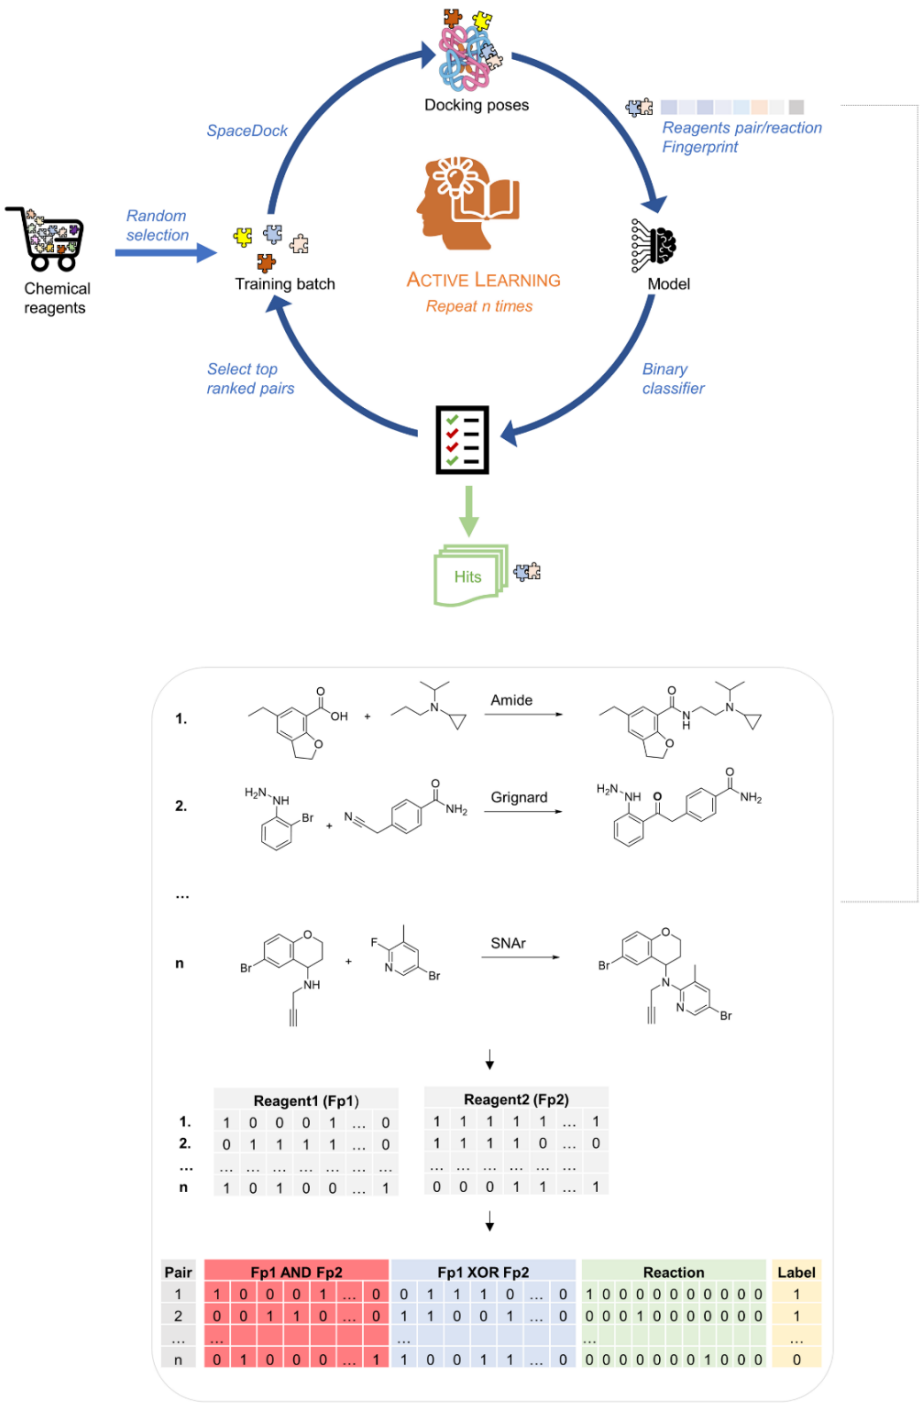

Figure 2. Overall active learning cycle to select the most likely chemical reagents and reactions to enumerate target-specific ligands. The chemical space comprises n pairs of reagents undergoing a one organic chemistry reaction. The two chemical reagents are described by radius-2 Morgan fingerprints of size 2,048. Reagent fingerprints are last converted two consecutive blocks describing common (first block) and different (second block) activated bits. A label ("0", "1") is added to each descriptor to qualify the previously reported SpaceDock success to recombine the corresponding pair in the intended target binding site. A one-hot encoding of organic chemistry reactions (11 bits) is finally added. The active learning cycle is repeated until a user-defined percentage of the full chemical space has been explored.

descriptor consists of common activated bits in the two fingerprints, while the second part consists of the activated bits that are different between the fingerprints (**Figure 2**). For the CBLB dataset, a reaction descriptor vector was also added to the described pair descriptors. This vector is a one-hot encoding of the reaction that leads to the formation of the pair. The length of the reaction descriptor is 11 bits, corresponding to the number of 11 reactions that form the CBLB dataset (**Supporting Table S1**).

Active learning. The protocol utilized for optimizing AL parameters on sets 2 and 4 of one million pairs, can be described in the following steps:

(i) A set of *n* building block pairs is randomly selected from the library. The number *n* is defined by the initial batch size parameter set as a percentage of the full chemical space size (typically ranging from 0.1% to 2.5%);

(ii) A surrogate ML model (Logistic Regression, Random Forest, Multilayer Perceptron) is trained on acquired data;

(iii) The probability of belonging to the positive class is predicted for every reagent pair in the entire library;

(iv) The next subset is selected from the library as *m* reagent pairs with the highest predicted probability values, excluding the pairs selected at previous iterations. The number *m* is defined by the exploration batch size parameter set as a percentage of the full chemical space size (typically ranging from 0.25% to 2.5%).

Steps ii to iv are iteratively repeated until the selected batch size reaches the desired portion of the full chemical space (ranging from 5% to 10%), with the last exploration batch potentially smaller than the others, to reach the desired chemical space coverage.

After optimizing parameters on sets 2 and 4, an initial batch size of 0.25% and an exploration batch size of 1% of the full chemical space size were selected for the application to the full DRD3 (set 1) and CBLB (set 3) chemical spaces

Machine learning models. We utilized three different classification models to predict the probability of a data point being a positive sample: logistic regression (LogReg), random forest (RF), and multilayer perceptron (MLP). All models were implemented using the scikit-learn library version 1.4.0,[53] using standard values of models' hyperparameters.

Logistic Regression (LogReg). Logistic regression is a linear model employed for binary classification tasks. It estimates the target variable as a linear combination of input features, subsequently applying a sigmoid activation function, which maps it to a value between 0 and 1, representing the probability of the data point belonging to the positive class. The output of the sigmoid function is utilized directly as the probability score.

Random Forest (RF). Random forest is an ensemble learning method that constructs a multitude of decision trees during training and outputs the class that is determined by the majority vote of the individual trees. For probability estimation, the portion of trees in the forest that predict a certain class is interpreted as the probability of being assigned that class.

Multilayer Perceptron (MLP). The multilayer perceptron is a feedforward artificial neural network comprising multiple layers of nodes. Each node (neuron) in one layer connects to every node in the subsequent layer. The MLP employs a non-linear activation function to model complex relationships in the data. We used an MLP model with one hidden layer containing

100 neurons and a ReLU activation function. The output layer uses a sigmoid activation function, the output of which is used directly as a probability estimate.

Cross validation. During training, the batch is divided into *n* equally-sized folds, and *n* models are trained, each using *n-1* folds. After training, each model is used to make predictions for the entire dataset. For each data point, the mean predicted probability $\hat{\mu}(p)$ and the standard deviation $\sigma(p)$ are calculated.

$$\hat{\mu}(p) = \frac{\sum_{i=1}^{n} p_i}{n}$$

$$\sigma(p) = \sqrt{\frac{\sum_{i=1}^{n} (\hat{\mu}(p) - p_i)^2}{n}}$$

Where p_i is the prediction of individual model and *n* the number of folds.

Acquisition functions. An acquisition function in active learning is a heuristic used to determine which data points should be selected next for labeling or evaluation. The acquisition function is applied to the surrogate model predictions, and the argmax function is applied to the calculated values to obtain the indexes of pairs that should be retrieved on the current iteration.

$$S = \underset{\{x_i\}_{i=i}^{b} \subset X}{\operatorname{argmax}} \sum_{i=1}^{b} \alpha(x, \hat{f})$$

where *S* is a new subset of pairs to acquire, *X* is the chemical space, α the acquisition function, *x* a data point from the library, $\hat{f}$ are predictions of the model for the point *x*, *b* is the size of the chemical space.

The goal of the acquisition function is to balance exploration and exploitation. In this work, we employed four different acquisition functions: greedy, Thompson, standard deviation (SD), and entropy. Please note that entropy was only used without cross validation.

Greedy: $\alpha(p) = \hat{\mu}(p)$ Data points with the highest mean predicted probabilities $\hat{\mu}(p)$are selected for the next batch. This function prioritizes data points that the model currently believes are most likely to be hits.

Thompson: $\alpha(p) \sim N(\hat{\mu}(p), \sigma(p))$ For each data point, a random number is drawn from a normal distribution with the pre-calculated mean $\hat{\mu}(p)$and standard deviation of the predicted probabilities $\sigma(p)$. This approach allows for both the exploration of uncertain regions and the exploitation of high-probability regions in the chemical space.

SD: $\alpha(p) = \sigma(p)$ Data points with the highest standard deviation *σ(p)* are selected for the next batch. This function targets data points where the model is most uncertain, aiming to reduce uncertainty in subsequent iterations.

Entropy: $\alpha(p) = -(p * ln(p) + (1-p)ln(1-p))$ When entropy is used as an acquisition function, a single model is trained on the full batch. Predictions are then made for the entire dataset, where the probability of a positive outcome *(p)* is predicted. Then the points with the highest entropy are selected for the next batch. In a binary classification task, selecting points with the highest entropy is equivalent to selecting points with probabilities closest to the decision boundary (0.5). This approach aims to retrieve points where the model is least certain about its predictions.

Evaluation metrics. The performance of each classification model was evaluated by computing the percentage of SpaceDock hits (pairs of reagents), identified at each active learning iteration.

Synthesizability of FDA-approved drugs. 2D structures of all FDA-approved drugs (n=2,619) were downloaded from the DrugBank knowledgebase (https://go.drugbank.com/) and filtered for compounds verifying two properties: (i) a molecular weight between 90 and 1,000 Da (ii) a known target with a Uniprot identifier. The 280 remaining compounds (SMILES strings) were used as queries to estimate their synthetic route by the SynFormer algorithm.[11] For each query, the synthetic route yielding the closest possible analog in less than 5 steps was saved and post-processed to keep only compounds with an objective score above 0.7.

RESULTS AND DISCUSSION

The general workflow used herein is depicted in **Figure 2**. A random selection of a few commercial chemical reagents (from 0.1% to 2.5% of the full set) is first docked to the 3D structure of the protein target with the SpaceDock method,[16] which directly recombines pairs of reagents' poses into full ligands according to their topological and chemical cross-compatibility. Pairs of reagents and the corresponding organic chemistry reaction(s) to recombine them are encoded in a fixed-sized 2D fingerprint (see Methods) and labeled positive (hit) or negative according to as whether SpaceDock was successful or not to link them in the protein binding site, according to strict topological, protein-ligand interactions and chemical reactivity criteria.[16] The reactants-reaction fingerprint is indeed 3D agnostic and is therefore identical for two enantiomeric ligands. This is however not crucial since

reactants are later converted into 3D coordinates before docking, explicitly considering stereochemistry if necessary. If the commercial reactant has a well-defined stereochemistry, it will be kept upon 2D to 3D conversion for docking. In case of a racemic building block, all stereoisomers are generated and docked to the target. The spacedocking procedure is therefore able to pick the correct enantiomer after the docking step. Likewise, the conformational freedom of the ligand is not addressed by the fingerprint but explicitly considered later during the docking-recombination steps.

A machine learning model (logistic regressor, random forest, multilayer perceptron) is then trained to distinguish positive from negative pairs and then used to enrich the initial training batch with a new exploration batch (from 0.25 to 2.5% of the full set), defining an active learning cycle that terminates after a user-defined number of iterations covering up to 5-10% of the entire chemical space. The three selected ML models correspond to increasing complexity and computing demand (LogReg < RF < MLP) and are easy to implement in scikit-learn with a low parametrization level. Indeed, many other architectures could have been chosen, e.g. decision trees or deep neural networks (DNNs). A previous study[41] evidenced that DNNs do not provide any substantial advantage to the ones selected in the current study, for predicting docking scores. Moreover, training DNN models on 10-100 million samples, as requested for set 2 would require multiple graphic processing units and a quite significant computing demand, thereby limiting the interest of the corresponding AL-acceleration protocol.

Preliminary trials on a baseline model. We first performed early trials on the 1 million pairs sets (DRD3 set 2, CBLB set 4) using logistic regression as baseline model, and setting initial and exploration batch sizes to 0.5% of the full chemical space. Obtained results clearly indicated that cross-validation, whatever the number of folds and acquisition metric did not yield to significantly higher hit recall values than a single model trained on all data using either entropy or the greedy metric to acquire new

samples at each iteration (**Supporting Tables 2, 3**). From here on, we have thus run single models (in triplicates) with the greedy acquisition method, and further studied the importance of key AL parameters (surrogate ML model, initial batch size, exploration batch size), starting with the above representative chemical spaces of one million pairs of reagents. Whereas set 2 addresses a single reaction in a focused chemical space (670 million carboxamides), set 4 is representative of a much larger chemical space (11 reactions) defined by 11 diverse organic chemistry reactions (**Supporting Table S1**). Since AL is supposed to retrieve the highest possible number of true hits within the smallest possible exploration of the full chemical space,[35, 40, 43] the current study will primarily focus on the hit recall at each AL iteration as a performance indicator. The best compromise between speed and performance will then be applied to the exploration of the entire chemical spaces of 670 million (set 1) and 4.5 billion (set 3) possible compounds.

Influence of the initial batch size. We first investigated the size of the initial batch used for training ML models to predict hits from small-sized sets 2 and 4, keeping the exploration batch size fixed to 5,000 reagent pairs (0.5% of the chemical space). Reassuringly, AL is shown to be profitable to all scenarios whatever the dataset (set 2, set 4), the size of the training set (ranging from 0.1% to 2.5% of the chemical space), and the ML model (**Figure 3**). From 60 to 80% of SpaceDock hits could indeed be recovered while exploring only 10% of the entire chemical space. In agreement with a previous study aimed at predicting docking scores,[41] hit recall tends to decrease when increasing the initial training batch size from 0.1 to 2.5% of the full space. For both datasets, smaller initial batch sizes generally result in better performance during the early stages of exploration for all models. For instance, with a 0.1% initial batch size, the models exhibit higher percentages of hits retrieved at 1-3% exploration compared to larger initial batch sizes (e.g. 1-2.5%). As the initial batch size increases, performance at

the early stages (1% to 3%) tends to be lower. However, performance converges at higher exploration percentages (9% to 10%), and the initial batch size had minimal effect on the total number of hits retrieved after exploring 10% of the library.

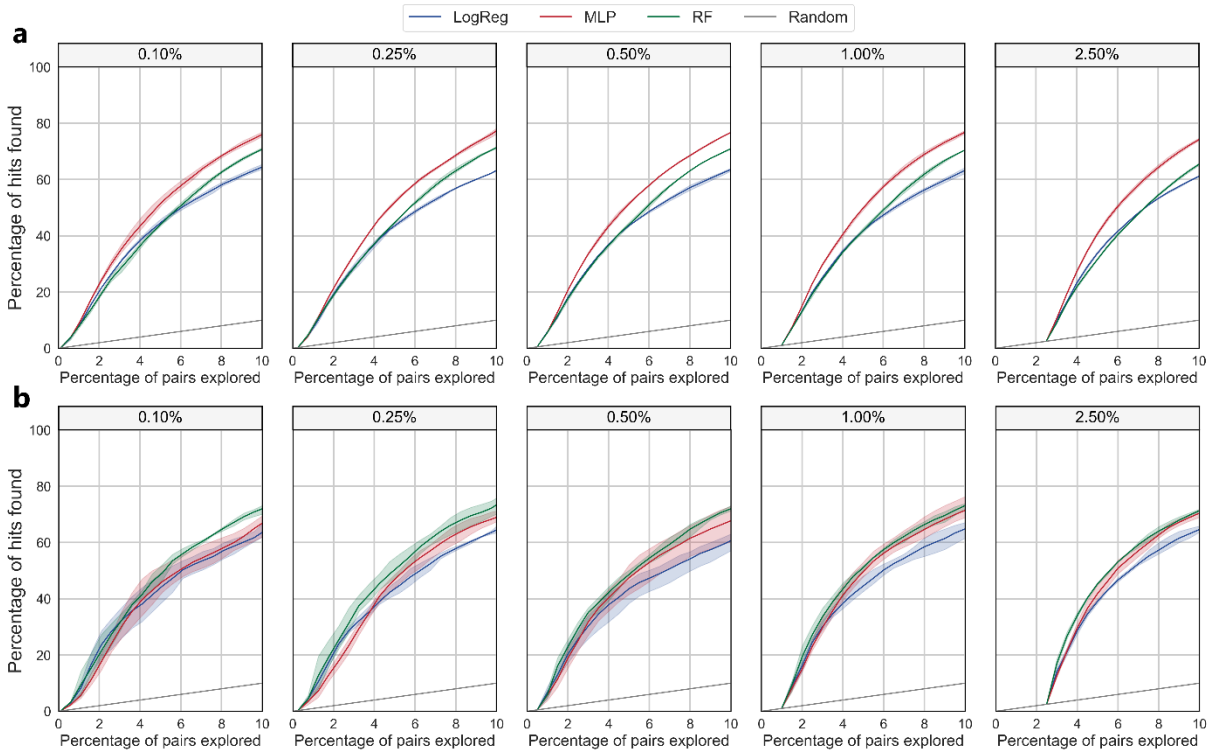

Figure 3. Influence of the size of initial batches on hit recall. Initial batches of various sizes (from 1,000 to 25,000 reagent pairs) with exploration batches of fixed size (5,000 pairs) were investigated using various machine learning models (logistic regression, random forest, multilayer perceptron) until 100,000 pairs were sampled. The size of the last acquisition batch has been reduced to achieve 10% of chemical space exploration. The results are averaged across three independent runs, with the line representing the mean value, and the shaded area showing the range of maximum and minimum values across the three runs. **a)** DRD3 set 2 of one million pairs; **b)** CBLB set 4 of one million pairs.

At 10% library exploration, the best results were obtained with an initial batch size of 0.25% for both datasets (**Supporting Tables S4, S5**). The observed trends are dataset-independent and equally applicable to single-reaction and multi-reaction spaces. Last, we do observe variations in the performance depending on the chosen ML model. Whereas the MLP model was clearly the most suitable for set 2, RF tends to be the best choice for set 4 (**Figure 3**). The lightweight regression model,

although clearly less accurate than more sophisticated models, still exhibits a remarkable performance, given the significantly reduced computing times required for training and prediction. The best performance for DRD3 set 2 was achieved with an initial batch size of 0.25% using the MLP model, retrieving 51.78% and 77.26% of hits at 5% and 10% of exploration, respectively (**Supporting Table S4**).

For the CBLB set 4, the best results were achieved with an initial batch size of 0.25% using the RF model. It retrieved 50.31% and 73.39% of the hits at the 5% and 10% exploration thresholds, respectively (**Supporting Table S5**). The explicit encoding of the reaction used to couple reagents marginally improves the performance of all models probably because ca. 80% of all selected reagent pairs can only be coupled by a single organic chemistry reaction (**Supporting Figure S1**). The overall accuracy of the binary classifiers, exemplified here by the best models (MLP for set 2, RF for set 4) trained on initial and acquisition batches of 5,000 pairs is excellent for both datasets (**Supporting Figure S2**). Interestingly, the optimal probability threshold set to distinguish positive from negative samples is lower than 0.50 for both sets (0.40 for set 2, 0.20 for set 4) and leads to excellent balanced accuracies of 0.883 and 0.882 for DRD3 and CBLB sets 2 and 4, respectively. The models perfectly assign a negative label to almost all true negatives and efficiently rank most of positive samples before the first negative samples, as shown by very high values of the area under the ROC curves (0.93 and 0.89 for sets 2 and 4, respectively; **Supporting Figure S2**).

Influence of the exploration batch size. We next investigated effect of the exploration batch size, used for extending the training set during the AL cycle, on the AL performance (**Figure 4**). To speed up calculations, the initial batch size was set to 0.5% of the full space (5,000 pairs) and not to the optimal value of 0.25 % (see above), since the former setting require one AL cycle less to explore a fixed subset of the entire chemical space while only marginally decreasing the hit recall (less than 1.5%; **Supporting**

Tables S4, S5). Exploration batches ranging from 0.25% (2,500 reagent pairs) to 2.5% (25,000 pairs) were investigated, for the same ML models. For all ML models and both chemical spaces, we observed a clear trend where decreasing the batch size led to improved performance (**Figure 4**).

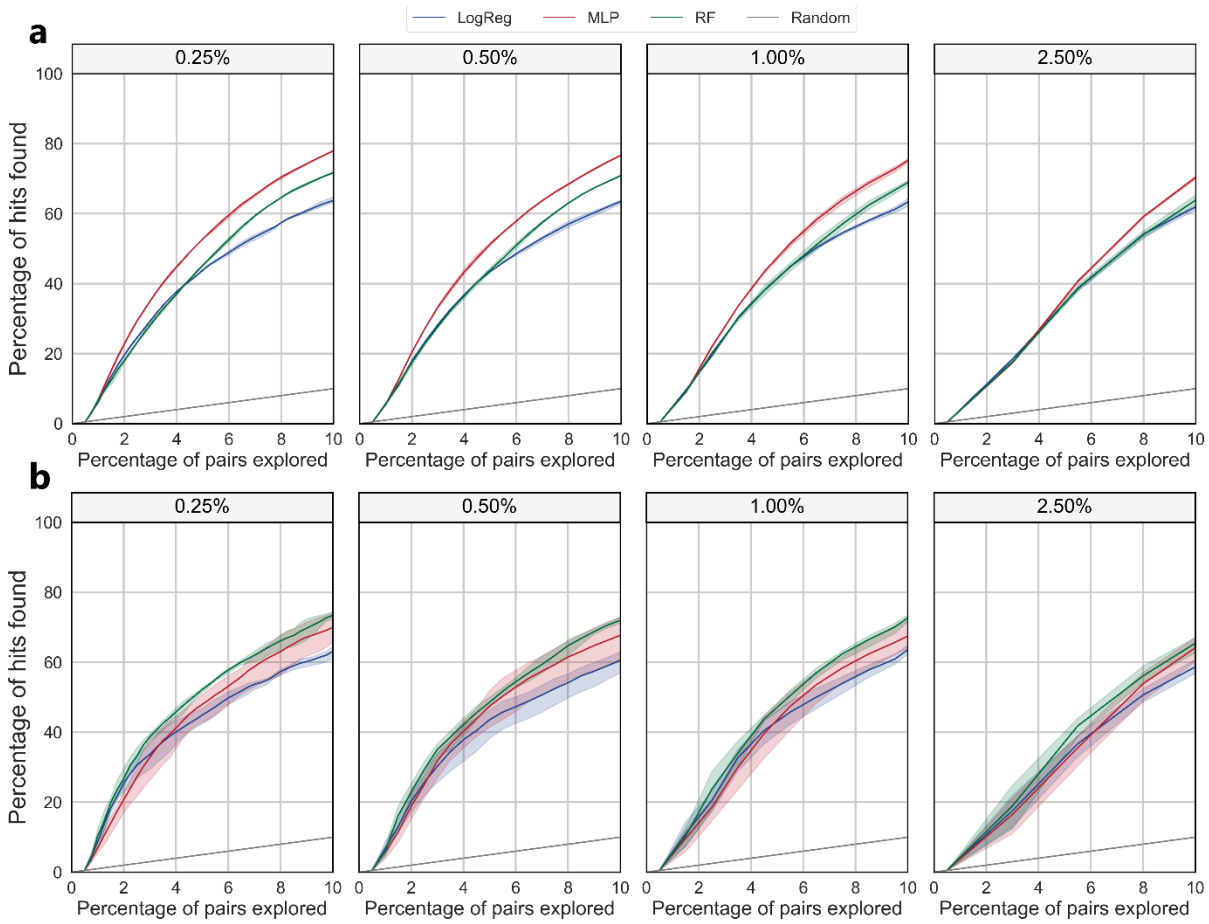

Figure 4. Influence of the size of the exploration batch on hit recall. Exploration batches of various sizes (from 2,500 to 25,000 reagent pairs) with initial batches of fixed size (5,000 pairs) were investigated using various machine learning model (logistic regression, random forest, multilayer perceptron) and acquisition methods (greedy, standard deviation, Thompson sampling, random) until 100,000 pairs of the full chemical space have been sampled. The size of the last acquisition batch has been reduced to achieve 10% of chemical space exploration. The results are averaged across three independent runs, with the line representing the mean value, and the shaded area showing the range of maximum and minimum values across the three runs **a)** DRD3 set 2 of one million pairs; **b)** CBLB set 4 of one million pairs.

For DRD3 set 2 the best performance was achieved with the MLP model and exploration batch size of 0.25%, which resulted in hit recalls of 52.78% and 77.99 at 5% and 10% of exploration of the DRD3 set 2, respectively (**Supporting Table S6**). For the CBLB set 4, the best performance was obtained by a RF model and the smallest possible exploration batch of 0.25%, with hit recall values of 52.17% and 73.39% upon exploring 5 and 10% of the full chemical space (**Supporting Table S7**). As to be expected, maximal hit recalls are slightly higher for set 2 than for set 4, due to the increased chemical diversity and number of organic chemistry reactions required to enumerate CBLB set 4 compounds. However, explicitly encoding the nature of the chemical reaction along with the reagents, still enables to achieve high hit recall values, and presents the advantage of avoiding reaction-specific ML models.

Influence of the ML model. For the DRD3 set 2, the MLP model consistently retrieves the highest percentage of hits with varying initial and exploration batch sizes (**Figures 3, 4**). It outperforms both LogReg and RF, particularly as exploration advances. For instance, with an initial batch size of 0.1%, MLP achieves 75.97% hit recovery at 10% exploration, significantly surpassing RF (70.75%) and LogReg (64.42%). Generally, the performance of the MLP model is followed by RF, with LogReg showing the lowest performance.

For the CBLB set 4, the RF model exhibited superior performance across all initial and exploration batch sizes (**Figures 3, 4**). The MLP model demonstrated competitive performance, especially in the later stages of exploration, while the LogReg model consistently showed the least performance. For instance, with an initial batch size of 0.5%, RF retrieved 71.98% of hits after exploring 10% of the dataset, while MLP and LogReg could only recover 67.71% and 60.60% of hits, respectively.

Analysis of retrieved hits. The analysis of the hits retrieved in the experiments described above shows that AL effectively focuses on retrieving the best reagent pairs and reactions from the corresponding chemical spaces. Since hit definition is based on SpaceDock ability to properly orient a pair of reagents for a well-defined reaction within the target binding site, we next checked how hit recall values depend on the hit definition rules. Keeping the geometrical and chemical cross- compatibility rules of reagent pairs unchanged, we defined stricter hit lists by just increasing the threshold of protein-ligand interaction fingerprint (IFP) similarity of the docked reagent pairs to the ground truth ligand. Interestingly, the stricter the SpaceDock hit definition, the higher the hit recall values for the DRD3 set 2 (**Supporting Figure S3**). Interestingly, the same trend could be observed for the multi-reaction CBLB set 4, demonstrating that the AL protocol can select not only the best possible reagents but also the proper organic chemistry reaction to link them (**Supporting Figure S4**). Using small initial batch sizes (0.1 to 0.5%), hit recall values increase from ca. 70 to 80% when the full IFP similarity threshold to qualify a SpaceDock hit is increased from 0.60 to 0.70 (compare **Figure 3B** and **Figure S4**).

Effect of initial and exploration batch sizes on computing times. AL requires multiple training and prediction steps requiring computer resources depending on the chosen ML model, the size of the initial training batch, the size of the exploration batches, and the desired number of AL iterations. If the computing time is modest for small-sized chemical spaces, it can rapidly increase for ultra-large spaces, notably for the demanding training steps. We therefore checked the cpu time necessary to explore 10% of the pilot DRD3 set 2 and CBLB set 4 of one million samples (**Supporting Table S8**). The time required for the active learning procedure across different exploration batch sizes for the two chemical spaces shows a consistent trend: increasing the exploration batch size significantly reduces the time needed for the process. This very strong effect can be easily explained by the reduction of

iterations needed to explore 10% of the full chemical space. The effect of initial batch size on the time required for the AL procedure is here more complex. The Logistic Regression model's time varied without a clear trend across different initial batch sizes, while the MLP and RF models showed some reductions in time with increasing initial batch sizes, but the relationship was not straightforward. The results for the LogReg model show that it is significantly faster (seconds to minutes scale) than both the MLP and RF models that, across all tested conditions, requires from 30 to 90 minutes. This speed advantage makes the LogReg model a very quick solution for active learning in large datasets, providing a balance between efficiency and performance. Its rapid processing capability makes it particularly suitable for scenarios where computational resources are limited, or fast iterations are required.

Application of active learning to the structure-based screening of ultra-large chemical spaces. Having made a first proof of concept that chemical reagents and organic chemistry reactions may be predicted by an AL protocol to enumerate target-specific ligands from small-sized DRD3 and CBLB chemical spaces, we next applied the AL workflow to the corresponding entire spaces (set 1 and set 3, **Figure 1**) comprising 670 million and 4.5 billion possible compounds, respectively. To screen the two chemical spaces, a unique set of AL parameters was chosen: a MLP model was trained on 0.25% of the full chemical spaces and updated by exploration batches of 1 % until 5% of the chemical space was explored. The selection of the exploration batch size of 1% was influenced by computational performance considerations. Smaller batch sizes required more iterations, and given the large sizes of the two spaces, we chose an exploration batch size of 1% to limit the number of iterations to six. This decision balanced computational efficiency and performance, as our previous study on small subspaces showed that an exploration batch size of 1% did not drastically reduce performance compared to a smaller batch size of 0.25%.

Inspection of hit recall curves (**Figure 5**) clearly shows that most conclusions drawn from the previous preliminary studies on one-million pairs sets 2 and 4 also apply to ultra-large chemical space screening: (i) AL was profitable to select the most suitable reagents for single (**Figure 5a**) and multiple-reaction (**Figure 5b**) ultra-large spaces; (ii) hit recovery ranging from 70 to 98%, depending on the hit definition, could be achieved by just scanning 5% of the full chemical spaces; (iii) hit recovery was even higher for stricter hit definitions using higher IFP similarities.

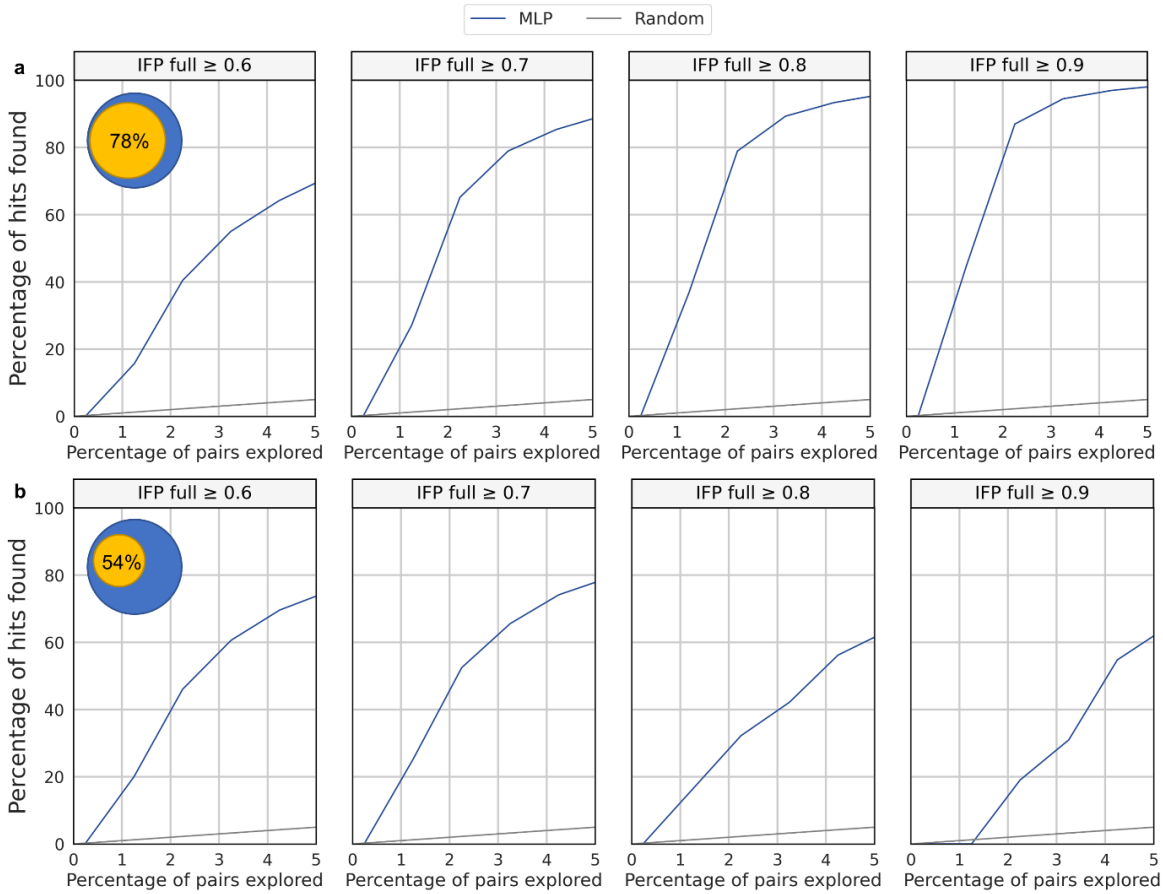

Figure 5. Hit recall by active learning of the best reagents and reactions on full-sized chemical spaces. A surrogate MLP model is trained on 0.25% of the chemical space and updated by acquisition batch sizes of 1%. The size of the last acquisition batch has been reduced to achieve 5% of chemical space exploration. Hits are defined as pairs of reagents successfully recombined with SpaceDock and exhibiting a full IFP similarity to the ground truth ligand greater or equal to a user-defined threshold (0.6, 0.7, 0.8, 0.9). Random selection (grey line) is given for comparison. **a)** DRD3 set 1 of 670,708,962 pairs of reagents and a single reaction. IFP similarity thresholds correspond to the top 2.05% (13,755,358 pairs: IFP ≥ 0.6), 0.52% (3,475,458 pairs: IFP ≥ 0.7), 0.01% (693,735 pairs: IFP ≥ 0.8), and 0.0004% scorers (26,279 pairs: IFP ≥ 0.9) of the entire chemical space, respectively; **b)** CBLB set 3 of 4.5 billion pairs of reagents and 11 possible reactions. IFP similarity thresholds correspond to the top 0.12% (7,702,856 pairs: IFP ≥ 0.6), 0.02% (1,411,712 pairs: IFP ≥ 0.7), 0.0004% (25,829 pairs: IFP ≥ 0.8), and 6.10^{-}

[7]% top scorers (42 pairs: IFP ≥ 0.9) of the entire chemical space, respectively. The inserts in the left plots display the percentage of Bemis-Murcko scaffolds recovered by AL-accelerated SpaceDock (yellow) with respect to exhaustive SpaceDock (blue).

Interestingly, most chemotypes described by unique Bemis-Murcko (BM) scaffolds,[54] were recovered upon AL acceleration (78% for set 1; 54% for set 3) demonstrating that AL was not trapped in a few regions of the full chemical space. Importantly, the chemotype describing the ground-truth DRD3 antagonist (eticlopride) was successfully retrieved, the closest AL-SpaceDock hit being even more similar to eticlopride than the molecule selected by the exhaustive SpaceDock protocol (**Supporting Figure S6**).

Due to much lower SpaceDock hit rates in the CBLB space, hit and BM scaffold recovery rates were lower than those obtained for DRD3 space at stricter hit definitions. However, it is remarkable that 26 out of 42 possible hits can still be retrieved, considering the strictest hit definition (IFP similarity ≥ 0.9; **Figure 5b**), from a chemical space of 4.5 billion possible compounds Hit recall curves were very steep at the first steps of the AL cycle, illustrating good early enrichment values, and begin to level off after exploring 2-3% of the full chemical spaces. Interestingly, a plateau was obtained sooner for stricter hit definitions.

Comparing the computing time required by exhaustive SpaceDock with the AL-accelerated SpaceDock protocol, we noticed a significant acceleration by a factor 9.5 for the 670 million DRD3 space, and of 15.7 for the bigger 4.7 billion CBLB space (**Supporting Figure S5**). To even speed-up the AL cycle, MLP models were implemented in PyTorch v2.3.1,[55] thereby enabling much faster AL iterations running on a single NVIDIA® RTX A5000 graphics processing unit (GPU). Obtained hit recall values for the CBLB set 3 were almost identical to that described above using the Scikit-learn framework,[53] at a very modest computing time (40 gpu hours for MLP training/prediction, 670 cpu hours for SpaceDocking).

AL-SpaceDock applicability domain. It is important to recall that the presented computational method is not aimed at finding drugs but just easily synthesizable hit compounds to start a medicinal chemistry optimization. Drugs are usually far more complex than hits since many properties (especially selectivity, pharmacokinetics and tox/safety) need to be optimized. In most cases, this optimization results in structural modifications (e.g. blocking metabolism sites, adding solubilizing groups) requiring several synthetic steps, making final drugs more difficult to synthesize (at least in terms of number of synthetic steps) than primary hits. Despite being simplistic, generating ultra-large chemical spaces by simple organic chemistry reactions, even from a single step, generates a huge amount of chemically diverse chemotypes that are suitable enough for primary hit identification.[56]

To ascertain that close analogs of approved drugs might be found by our structure-based approach, the synthetic accessibility of 280 small molecular weight and target-annotated FDA-approved drugs[57] was estimated by the SynFormer algorithm,[11] a generative artificial Intelligence software for navigating the Enamine REAL synthesizable chemical space. For 152 out of 280 queries, the drug or a very close analog thereof are indeed predicted to be synthesizable in at least 2 steps (out of which 25% are obtained in a single step) from a list of 115 reactions involving 223,244 commercial building blocks (**Figure S7**). We can therefore conclude that AL-accelerated SpaceDock is suitable to prioritize high-value primary hits in many drug discovery scenarios.

CONCLUSIONS

We herein demonstrate that active learning may be efficiently used to prioritize reagents and organic chemistry reactions to enumerate target-specific ligands under 3D constraints of a given binding site. Contrary to many AL-accelerated docking procotols,[34-43] the presented workflow does not aim at predicting docking scores but just the topological and chemical complementarity of two chemical

reagents to enumerate a full ligand exhibiting shared protein-ligand interactions with at least a ground-truth ligand. Given the notorious absence of relationships between docking scores and absolute binding free energies, we believe that our approach is much safer to prioritize the most relevant hits than ML-accelerating docking protocols predicting raw docking scores. It is important to recall that the current method is aimed at predicting target-specific primary hits, amenable to medicinal chemistry optimization, and easily synthesizable by simple one-step organic chemistry reactions. By no way can it be applied to predict the multi-step retrosynthesis of much more complex final drug candidates.

Several reaction-based chemical space exploration methods starting from a known compound to propose synthetically viable and diverse analogs have been recently proposed for that purpose.[7-11] They significantly differ from the SpaceDock approach by three points: (i) they are driven from simple SMILES arbitrary target specification (SMARTS) language, (ii) they operate in a retrosynthesis mode (lead to retrosynthetic routes to constrained enumeration of analogs), (iii) they do not explicitly take into account a protein binding site topology during the library enumeration step. The latter methods are in fact quite complementary to ours in enumerating scaffold-focused chemical spaces from primary hits identified by our 3D-driven and pocket-oriented approach.

When applied to both single and multiple one-step reactions encoding ultra-large chemical spaces of various sizes (from 670 million to 4.5 billion tangible compounds), AL-SpaceDock was able to recover from 75 to 98% of the good reagents and corresponding organic chemistry reactions to enumerate the expected true hits. Each ML model is indeed target-specific and has to be generated beforehand before each AL-accelerated virtual screening. In the current two examples, although the peak performance is not observed for the same ML model-acquisition strategy, hit recovery rates remain within a small range whatever the strategy and ML model. The best compromise has to be chosen on a case-by-case basis with respect to the size of the chemical space and the available computing power. For chemical

sizes up to a few billion compounds, any ML model will be beneficial. For trillion-sized spaces, a simpler logistic regression model appears to be a good compromise between speed and accuracy.

Once selected, the reagent pairs may be either submitted to a SpaceDock full ligand triaging,[16] or recombined in 3D before a conventional all-atom docking to the binding site of interest. Our AL-accelerated docking procedure simplifies the structure-based screening of billion-sized combinatorial chemical spaces to an easily manageable list of a few million compounds that are worth investigating further. It should however be reminded that AL-accelerated docking does not scale linearly with the size of the chemical space to screen since the training step becomes more computationally expensive along the AL iterations. Therefore, extending the present protocol to a combinatorial space of one trillion compounds although doable requires optimizing the time needed to train ML models on a few billion compounds (e.g. by switching to gpu computing), the following prediction steps being achievable on standard cpu architectures with a dozen of cores. To fully delineate its applicability domain, the AL-accelerated workflow still needs to be challenged with the corresponding exhaustive approach, on additional true prospective screening campaigns, with respect to hit rate and hit overlap. Interestingly, it offers the possibility to cope in the future with combinatorial chemical spaces of ever-increasing sizes[58] to select the most promising ligands under precise constraints of any druggable protein cavity.

ACKNOWLEDGMENTS

The Calculation Centre of the IN2P3 (CNRS, Villeurbanne, France) is acknowledged for the allocation of computing time and excellent support. We deeply thank Michael Bossert and Yurii Moroz (Enamine

Ltd.) for sharing the list of the chemical reagents and organic chemistry reactions contributing to REAL space.

FUNDING

The authors are thankful to the doctoral school of chemical sciences (EDSC, University of Strasbourg) for a Ph.D. grant to F.S.

Author Contribution

Vladimir Kozyrev: development of reagent-reactions descriptors and of all active learning models, methodology, validation and calculations, data analysis, writing the manuscript; Francois Sindt: enumeration of exhaustive SpaceDock hits, list of reagent pairs and reactions; D. Rognan: conceptualization of the study, methodology, data analysis, writing the manuscript.

Conflict of interests

D.R. is cofounder and shareholder of BIODOL Therapeutics. The other authors declare no competing interests.

Supporting Information.

Influence of cross-validation and acquisition metric on hit retrieval for the DRD3 chemical space of 1 million pairs of reagents (set 2); Influence of cross-validation and acquisition metric on hit retrieval for

the DRD3 chemical space of 1 million pairs of reagents (set 4); Organic chemistry reactions and the number of unique chemical reagents that define the CBLB chemical space of 4,576,492,610 compounds; Influence of the initial batch size on hit retrieval (in percentage) for the DRD3 chemical space of 1 million pairs of reagents (set 2); Influence of the initial batch size on hit retrieval (in percentage) for the CBLB chemical space of 1 million pairs of reagents (set 4); Influence of the exploration batch size on hit retrieval (in percentage) for the DRD3 chemical space of 1 million pairs of reagents (set 2); Influence of the exploration batch size on hit retrieval (in percentage) for the CBLB chemical space of 1 million pairs of reagents (set 4); Cpu time required to explore 10% of the DRD3 set 2 and CBLB set 4 of 1 million reagent pairs; Coupling of two chemical reagents according to well-defined organic chemistry reactions; Overall accuracy of a MLP model in the binary classification of reactant pairs for DRDR3 set 2 and CBLB set 4; Influence of the size of the initial batch size on hit recall in DRD3 set 2, considering different hit definitions; Influence of the size of the initial batch size on hit recall from CBLB set 4, considering different hit definitions; Ground-truth DRD3 ligand (Eticlopride, DrugBankID DB15492, left panel) and closest analogs identified by exhaustive SpaceDock (middle panel) and AL-accelerated SpaceDock (right panel); Number of predicted synthetic steps to synthesize 280 FDA-approved drugs from Enamine commercial building blocks

Data and Software Availability

The one-million pairs datasets (set 2, set 4) used in this study (SMILES strings of Enamine reagents, hit lists) are freely available at https://github.com/VAKozyrev/ML-SpaceDock.

SpaceDock processing scripts are available at https://github.com/litfsindt/LIT-SpaceDock (accessed 08-29-2024). Active learning protocols (python scripts) for training ML models and predicting the

probability of recombining reagents pairs can be accessed from https://github.com/VAKozyrev/ML-SpaceDock.

REFERENCES

1. Sadybekov, A. V.; Katritch, V., Computational Approaches Streamlining Drug Discovery. *Nature,* **2023,** *616*, 673-685.
2. Lyu, J.; Irwin, J. J.; Shoichet, B. K., Modeling the Expansion of Virtual Screening Libraries. *Nat. Chem. Biol,* **2023,** *19*, 712-718.
3. Warr, W. A.; Nicklaus, M. C.; Nicolaou, C. A.; Rarey, M., Exploration of Ultralarge Compound Collections for Drug Discovery. *J. Chem. Inf. Model.,* **2022,** *62*, 2021-2034.
4. Lucas, X.; Gruning, B. A.; Bleher, S.; Gunther, S., The Purchasable Chemical Space: A Detailed Picture. *J. Chem. Inf. Model.,* **2015,** *55*, 915-924.
5. 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.
6. Wahl, J.; Sander, T., Fully Automated Creation of Virtual Chemical Fragment Spaces Using the Open-Source Library Openchemlib. *J. Chem. Inf. Model.,* **2022,** *62*, 2202-2211.
7. 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.
8. Dolfus, U.; Briem, H.; Rarey, M., Synthesis-Aware Generation of Structural Analogues. *J. Chem. Inf. Model.,* **2022,** *62*, 3565-3576.
9. Levin, I.; Fortunato, M. E.; Tan, K. L.; Coley, C. W., Computer-Aided Evaluation and Exploration of Chemical Spaces Constrained by Reaction Pathways. *AiChE J.,* **2023,** *69*.

10. Sankaranarayanan, K.; Jensen, K. F., Similarity Based Functionalization for Enumeration of Synthetically Plausible Chemical Libraries Surrounding a Target. *Chem. Sci.,* **2024,** *15*, 10221-10231.
11. Gao, W.; Luo, S.; Coley, C. W., Generative Artificial Intelligence for Navigating Synthesizable Chemical Space. *ArXiv: 2410.03494v1,* **2024**.
12. Sadybekov, A. A.; Sadybekov, A. V.; Liu, Y.; Iliopoulos-Tsoutsouvas, C.; Huang, X. P.; Pickett, J.; Houser, B.; Patel, N.; Tran, N. K.; Tong, F.; Zvonok, N.; Jain, M. K.; Savych, O.; Radchenko, D. S.; Nikas, S. P.; Petasis, N. A.; Moroz, Y. S.; Roth, B. L.; Makriyannis, A.; Katritch, V., Synthon-Based Ligand Discovery in Virtual Libraries of over 11 Billion Compounds. *Nature,* **2022,** *601*, 452-459.
13. Beroza, P.; Crawford, J. J.; Ganichkin, O.; Gendelev, L.; Harris, S. F.; Klein, R.; Miu, A.; Steinbacher, S.; Klingler, F. M.; Lemmen, C., Chemical Space Docking Enables Large-Scale Structure-Based Virtual Screening to Discover ROCK1 Kinase Inhibitors. *Nat Commun,* **2022,** *13*, 6447.
14. Lyu, J.; Wang, S.; Balius, T. E.; Singh, I.; Levit, A.; Moroz, Y. S.; O'Meara, M. J.; Che, T.; Algaa, E.; Tolmachova, K.; Tolmachev, A. A.; Shoichet, B. K.; Roth, B. L.; Irwin, J. J., Ultra-Large Library Docking for Discovering New Chemotypes. *Nature,* **2019,** *566*, 224-229.
15. Muller, J.; Klein, R.; Tarkhanova, O.; Gryniukova, A.; Borysko, P.; Merkl, S.; Ruf, M.; Neumann, A.; Gastreich, M.; Moroz, Y. S.; Klebe, G.; Glinca, S., Magnet for the Needle in Haystack: "Crystal Structure First" Fragment Hits Unlock Active Chemical Matter Using Targeted Exploration of Vast Chemical Spaces. *J. Med. Chem.,* **2022,** *65*, 15663-15678.
16. Sindt, F.; Seyller, A.; Eguida, M.; Rognan, D., Protein Structure-Based Organic Chemistry-Driven Ligand Design from Ultralarge Chemical Spaces. *ACS Cent. Sci.,* **2024,** *10*, 615-627.
17. Petrovic, D.; Scott, J. S.; Bodnarchuk, M. S.; Lorthioir, O.; Boyd, S.; Hughes, G. M.; Lane, J.; Wu, A.; Hargreaves, D.; Robinson, J.; Sadowski, J., Virtual Screening in the Cloud Identifies Potent and Selective ROS1 Kinase Inhibitors. *J. Chem. Inf. Model.,* **2022,** *62*, 3832-3843.

18. Gahbauer, S.; DeLeon, C.; Braz, J. M.; Craik, V.; Kang, H. J.; Wan, X.; Huang, X. P.; Billesbolle, C. B.; Liu, Y.; Che, T.; Deshpande, I.; Jewell, M.; Fink, E. A.; Kondratov, I. S.; Moroz, Y. S.; Irwin, J. J.; Basbaum, A. I.; Roth, B. L.; Shoichet, B. K., Docking for EP4R Antagonists Active against Inflammatory Pain. *Nat. Commun.,* **2023,** *14*, 8067.

19. Singh, I.; Li, F.; Fink, E. A.; Chau, I.; Li, A.; Rodriguez-Hernandez, A.; Glenn, I.; Zapatero-Belinchon, F. J.; Rodriguez, M. L.; Devkota, K.; Deng, Z.; White, K.; Wan, X.; Tolmachova, N. A.; Moroz, Y. S.; Kaniskan, H. U.; Ott, M.; Garcia-Sastre, A.; Jin, J.; Fujimori, D. G.; Irwin, J. J.; Vedadi, M.; Shoichet, B. K., Structure-Based Discovery of Inhibitors of the SARS-CoV-2 NSP14 N7-Methyltransferase. *J. Med. Chem.,* **2023,** *66*, 7785-7803.

20. Singh, I.; Seth, A.; Billesbolle, C. B.; Braz, J.; Rodriguiz, R. M.; Roy, K.; Bekele, B.; Craik, V.; Huang, X. P.; Boytsov, D.; Pogorelov, V. M.; Lak, P.; O'Donnell, H.; Sandtner, W.; Irwin, J. J.; Roth, B. L.; Basbaum, A. I.; Wetsel, W. C.; Manglik, A.; Shoichet, B. K.; Rudnick, G., Structure-Based Discovery of Conformationally Selective Inhibitors of the Serotonin Transporter. *Cell,* **2023,** *186*, 2160-2175 e2117.

21. Chisholm, T. S.; Mackey, M.; Hunter, C. A., Discovery of High-Affinity Amyloid Ligands Using a Ligand-Based Virtual Screening Pipeline. *J. Am. Chem. Soc.,* **2023,** *145*, 15936-15950.

22. Fink, E. A.; Xu, J.; Hubner, H.; Braz, J. M.; Seemann, P.; Avet, C.; Craik, V.; Weikert, D.; Schmidt, M. F.; Webb, C. M.; Tolmachova, N. A.; Moroz, Y. S.; Huang, X. P.; Kalyanaraman, C.; Gahbauer, S.; Chen, G.; Liu, Z.; Jacobson, M. P.; Irwin, J. J.; Bouvier, M.; Du, Y.; Shoichet, B. K.; Basbaum, A. I.; Gmeiner, P., Structure-Based Discovery of Nonopioid Analgesics Acting through the Alpha(2a)-Adrenergic Receptor. *Science,* **2022,** *377*, eabn7065.

23. Kaplan, A. L.; Confair, D. N.; Kim, K.; Barros-Alvarez, X.; Rodriguiz, R. M.; Yang, Y.; Kweon, O. S.; Che, T.; McCorvy, J. D.; Kamber, D. N.; Phelan, J. P.; Martins, L. C.; Pogorelov, V. M.; DiBerto, J. F.; Slocum, S. T.; Huang, X. P.; Kumar, J. M.; Robertson, M. J.; Panova, O.; Seven, A. B.; Wetsel, A. Q.;

Wetsel, W. C.; Irwin, J. J.; Skiniotis, G.; Shoichet, B. K.; Roth, B. L.; Ellman, J. A., Bespoke Library Docking for 5-HT(2a) Receptor Agonists with Antidepressant Activity. *Nature,* **2022,** *610*, 582-591.

24. Luttens, A.; Gullberg, H.; Abdurakhmanov, E.; Vo, D. D.; Akaberi, D.; Talibov, V. O.; Nekhotiaeva, N.; Vangeel, L.; De Jonghe, S.; Jochmans, D.; Krambrich, J.; Tas, A.; Lundgren, B.; Gravenfors, Y.; Craig, A. J.; Atilaw, Y.; Sandstrom, A.; Moodie, L. W. K.; Lundkvist, A.; van Hemert, M. J.; Neyts, J.; Lennerstrand, J.; Kihlberg, J.; Sandberg, K.; Danielson, U. H.; Carlsson, J., Ultralarge Virtual Screening Identifies SARS-CoV-2 Main Protease Inhibitors with Broad-Spectrum Activity against Coronaviruses. *J. Am. Chem. Soc.,* **2022,** *144*, 2905-2920.
25. Alon, A.; Lyu, J.; Braz, J. M.; Tummino, T. A.; Craik, V.; O'Meara, M. J.; Webb, C. M.; Radchenko, D. S.; Moroz, Y. S.; Huang, X. P.; Liu, Y. F.; Roth, B. L.; Irwin, J. J.; Basbaum, A. I.; Shoichet, B. K.; Kruse, A. C., Structures of the σ2 Receptor Enable Docking for Bioactive Ligand Discovery. *Nature,* **2021,** *600*, 759-764.
26. Gorgulla, C.; Boeszoermenyi, A.; Wang, Z. F.; Fischer, P. D.; Coote, P. W.; Padmanabha Das, K. M.; Malets, Y. S.; Radchenko, D. S.; Moroz, Y. S.; Scott, D. A.; Fackeldey, K.; Hoffmann, M.; Iavniuk, I.; Wagner, G.; Arthanari, H., An Open-Source Drug Discovery Platform Enables Ultra-Large Virtual Screens. *Nature,* **2020,** *580*, 663-668.
27. Stein, R. M.; Kang, H. J.; McCorvy, J. D.; Glatfelter, G. C.; Jones, A. J.; Che, T.; Slocum, S.; Huang, X. P.; Savych, O.; Moroz, Y. S.; Stauch, B.; Johansson, L. C.; Cherezov, V.; Kenakin, T.; Irwin, J. J.; Shoichet, B. K.; Roth, B. L.; Dubocovich, M. L., Virtual Discovery of Melatonin Receptor Ligands to Modulate Circadian Rhythms. *Nature,* **2020,** *579*, 609-614.
28. Liu, F.; Wu, C. G.; Tu, C. L.; Glenn, I.; Meyerowitz, J.; Kaplan, A. L.; Lyu, J.; Cheng, Z.; Tarkhanova, O. O.; Moroz, Y. S.; Irwin, J. J.; Chang, W.; Shoichet, B. K.; Skiniotis, G., Large Library Docking Identifies Positive Allosteric Modulators of the Calcium-Sensing Receptor. *Science,* **2024,** *385*, eado1868.

29. https://www.biosolveit.de/chemical-spaces/ (accessed 10-31-2024).
30. Neumann, A.; Marrison, L.; Klein, R., Relevance of the Trillion-Sized Chemical Space "Explore" as a Source for Drug Discovery. *ACS Med. Chem. Lett.,* **2023,** *14*, 466-472.
31. Bedart, C.; Simoben, C. V.; Schapira, M., Emerging Structure-Based Computational Methods to Screen the Exploding Accessible Chemical Space. *Curr. Opin. Struct. Biol.,* **2024,** *86*, 102812.
32. Cavasotto, C. N.; Di Filippo, J. I., The Impact of Supervised Learning Methods in Ultralarge High-Throughput Docking. *J. Chem. Inf. Model.,* **2023,** *63*, 2267-2280.
33. Gorantla, R.; Kubincova, A.; Suutari, B.; Cossins, B. P.; Mey, A., Benchmarking Active Learning Protocols for Ligand-Binding Affinity Prediction. *J. Chem. Inf. Model.,* **2024,** *64*, 1955-1965.
34. Berenger, F.; Kumar, A.; Zhang, K. Y. J.; Yamanishi, Y., Lean-Docking: Exploiting Ligands' Predicted Docking Scores to Accelerate Molecular Docking. *J. Chem. Inf. Model.,* **2021,** *61*, 2341-2352.
35. Gentile, F.; Agrawal, V.; Hsing, M.; Ton, A. T.; Ban, F.; Norinder, U.; Gleave, M. E.; Cherkasov, A., Deep Docking: A Deep Learning Platform for Augmentation of Structure Based Drug Discovery. *ACS Cent. Sci.,* **2020,** *6*, 939-949.
36. Roggia, M.; Natale, B.; Amendola, G.; Di Maro, S.; Cosconati, S., Streamlining Large Chemical Library Docking with Artificial Intelligence: The PyRMD2Dock Approach. *J. Chem. Inf. Model.,* **2023,** *64*, 2143-2149.
37. Cao, Z.; Sciabola, S.; Wang, Y., Large-Scale Pretraining Improves Sample Efficiency of Active Learning-Based Virtual Screening. *J. Chem. Inf. Model.,* **2024,** *64*, 1882-1891.
38. Klarich, K.; Goldman, B.; Kramer, T.; Riley, P.; Walters, W. P., Thompson Sampling-an Efficient Method for Searching Ultralarge Synthesis on Demand Databases. *J. Chem. Inf. Model.,* **2024,** *64*, 1158-1171.

39. Popov, K. I.; Wellnitz, J.; Maxfield, T.; Tropsha, A., Hit Discovery Using Docking Enriched by Generative Modeling (HIDDEN GEM): A Novel Computational Workflow for Accelerated Virtual Screening of Ultra-Large Chemical Libraries. *Mol. Inf.,* **2024,** *42, e202300207*.

40. Yang, Y.; Yao, K.; Repasky, M. P.; Leswing, K.; Abel, R.; Shoichet, B. K.; Jerome, S. V., Efficient Exploration of Chemical Space with Docking and Deep Learning. *J. Chem. Theory Comput.,* **2021,** *17*, 7106-7119.

41. Marin, E.; Kovaleva, M.; Kadukova, M.; Mustafin, K.; Khorn, P.; Rogachev, A.; Mishin, A.; Guskov, A.; Borshchevskiy, V., Regression-Based Active Learning for Accessible Acceleration of Ultra-Large Library Docking. *J. Chem. Inf. Model.,* **2024,** *64*, 2612-2623.

42. Mehta, S.; Laghuvarapu, S.; Pathak, Y.; Sethi, A.; Alvala, M.; Priyakumar, U. D., MEMES: Machine Learning Framework for Enhanced Molecular Screening. *Chem. Sci.,* **2021,** *12*, 11710-11721.

43. Graff, D. E.; Shakhnovich, E. I.; Coley, C. W., Accelerating High-Throughput Virtual Screening through Molecular Pool-Based Active Learning. *Chem. Sci.,* **2021,** *12*, 7866-7881.

44. Jastrzebski, S.; Szymczak, M.; Pocha, A.; Mordalski, S.; Tabor, J.; Bojarski, A. J.; Podlewska, S., Emulating Docking Results Using a Deep Neural Network: A New Perspective for Virtual Screening. *J. Chem. Inf. Model.,* **2020,** *60*, 4246-4262.

45. Gentile, F.; Fernandez, M.; Ban, F.; Ton, A. T.; Mslati, H.; Perez, C. F.; Leblanc, E.; Yaacoub, J. C.; Gleave, J.; Stern, A.; Wong, B.; Jean, F.; Strynadka, N.; Cherkasov, A., Automated Discovery of Noncovalent Inhibitors of SARS-CoV-2 Main Protease by Consensus Deep Docking of 40 Billion Small Molecules. *Chem. Sci.,* **2021,** *12*, 15960-15974.

46. Gutkin, E.; Gusev, F.; Gentile, F.; Ban, F.; Koby, S. B.; Narangoda, C.; Isayev, O.; Cherkasov, A.; Kurnikova, M. G., In Silico Screening of LRRK2 WDR Domain Inhibitors Using Deep Docking and Free Energy Simulations. *Chem. Sci.,* **2024,** *15*, 8800-8812.

47. Li, Y.; Su, M.; Liu, Z.; Li, J.; Liu, J.; Han, L.; Wang, R., Assessing Protein-Ligand Interaction Scoring Functions with the CASF-2013 Benchmark. *Nat. Protoc.,* **2018,** *13*, 666-680.
48. Warren, G. L.; Andrews, C. W.; Capelli, A. M.; Clarke, B.; LaLonde, J.; Lambert, M. H.; Lindvall, M.; Nevins, N.; Semus, S. F.; Senger, S.; Tedesco, G.; Wall, I. D.; Woolven, J. M.; Peishoff, C. E.; Head, M. S., A Critical Assessment of Docking Programs and Scoring Functions. *J. Med. Chem.,* **2006,** *49*, 5912-5931.
49. Pantsar, T.; Poso, A., Binding Affinity Via Docking: Fact and Fiction. *Molecules,* **2018,** *23*.
50. Unpublished Data from Cache Challenge #4, https://cache-challenge.org/challenges/finding-ligands-targeting-the-tkb-domain-of-cblb (accessed 10-31-2024).
51. Kimani, S. W.; Perveen, S.; Szewezyk, M.; Zeng, H.; Dong, A.; Li, F.; Ghiabi, P.; Li, Y.; Chau, I.; Arrowsmith, C. H.; Barsyte-Lovejoy, D.; Santhakumar, V.; Vedadi, M.; Halabelian, L., The Co-Crystal Structure of Cbl-B and a Small-Molecule Inhibitor Reveals the Mechanism of CBL-B Inhibition. *Commun. Biol.,* **2023,** *6*, 1272.
52. Jones, G.; Willett, P.; Glen, R. C.; Leach, A. R.; Taylor, R., Development and Validation of a Genetic Algorithm for Flexible Docking. *J. Mol. Biol,* **1997,** *267*, 727-748.
53. Scikit-Learn: Machine Learning in Python, https://scikit-learn.org/ (accessed 10-31-2024).
54. Bemis, G. W.; Murcko, M. A., The Properties of Known Drugs. 1. Molecular Frameworks. *J. Med. Chem.,* **1996,** *39*, 2887-2893.
55. Pytorch, https://pytorch.org/ (accessed 10-31-2024).
56. Tomberg, A.; Boström, J., Can Easy Chemistry Produce Complex, Diverse, and Novel Molecules? . *Drug Discov. Today,* **2020,** *25*, 2174-2181.
57. Knox, C.; Wilson, M.; Klinger, C. M.; Franklin, M.; Oler, E.; Wilson, A.; Pon, A.; Cox, J.; Chin, N. E. L.; Strawbridge, S. A.; Garcia-Patino, M.; Kruger, R.; Sivakumaran, A.; Sanford, S.; Doshi, R.; Khetarpal, N.; Fatokun, O.; Doucet, D.; Zubkowski, A.; Rayat, D. Y.; Jackson, H.; Harford, K.; Anjum, A.; Zakir,

M.; Wang, F.; Tian, S.; Lee, B.; Liigand, J.; Peters, H.; Wang, R. Q. R.; Nguyen, T.; So, D.; Sharp, M.; da Silva, R.; Gabriel, C.; Scantlebury, J.; Jasinski, M.; Ackerman, D.; Jewison, T.; Sajed, T.; Gautam, V.; Wishart, D. S., Drugbank 6.0: The Drugbank Knowledgebase for 2024. *Nucleic Acids Res.,* **2024,** *52*, D1265-D1275.

58. Korn, M.; Ehrt, C.; Ruggiu, F.; Gastreich, M.; Rarey, M., Navigating Large Chemical Spaces in Early-Phase Drug Discovery. *Curr. Opin. Struct. Biol.,* **2023,** *80*, 102578.

Table of content graphics